


# Conference on Computing in High Energy and Nuclear Physics

Saturday, 19 October 2024 - Friday, 25 October 2024



## Book of Abstracts



# Contents

HPCNeuroNet: A Neuromorphic Approach Merging SNN Temporal Dynamics with Transformer Attention for FPGA-based Particle Physics . . . . .	1
Jet Discrimination with Quantum Complete Graph Neural Network . . . . .	1
Library for Functional analysis in CERN's ROOT . . . . .	1
Remote3: Public Engagement over 1 km underground –and beyond . . . . .	2
The CERN Open Source Program Office: changing the (HENP) software world, pragmatically . . . . .	2
Time-Of-Flight Correction for Hadrons at BESIII experiment . . . . .	3
Development of Auto-Validation System of BOSS . . . . .	4
Charged Particle Track Reconstruction in CLAS12 using Artificial Intelligence . . . . .	4
Online Electron Reconstruction at CLAS12 . . . . .	5
The automated Bandwidth Division for the LHCb first-level trigger . . . . .	5
Data Preservation in High Energy Physics: a 10 years perspective . . . . .	5
HPSS Batch Application . . . . .	6
Management of the data processing & Run-by-Run simulations in KM3NeT . . . . .	6
Hyperparameter Optimization for Deep Learning Models Using High Performance Computing . . . . .	7
Improving multi-track reconstruction algorithms in the Mu2e Experiment . . . . .	7
The Italian Summer Students Program at Fermilab and other US Laboratories: 40 years of education in particle physics and technology . . . . .	8
Streamlining ATLAS Monte-Carlo Generator Validation with PAVER . . . . .	8
Normalizing Flows for Physics Data Analyses . . . . .	9
dCache project status & update . . . . .	9
dE/dx Software in the BESIII Experiment . . . . .	10
dCache CI/CD migration to Kubernetes . . . . .	10

The evolution of INFN's Cloud Platform: improvements in Orchestration and User Experience . . . . .	11
Use of Intelligent Compression in Waveform like Data of Neutrino Oscillation Experiments . . . . .	12
Common Analysis Tools in CMS . . . . .	12
Transfer learning for Smart Background Simulation at Belle II . . . . .	12
ATLAS ITk Production Database use and tools . . . . .	13
Anomaly Detection on BESIII EMC using Machine Learning . . . . .	13
Versal ACAP processing for ATLAS-TileCal signal reconstruction . . . . .	14
Porting MADGRAPH to FPGA using High-Level Synthesis (HLS) . . . . .	14
WLCG transition from X.509 to Tokens: Progress and Outlook . . . . .	15
Data Challenge 2024 - CMS activities . . . . .	15
Model Building with Non-Parametric and Parametric Components for Partial Wave Analysis . . . . .	16
Advanced monitoring capabilities of the CMS Experiment for LHC Run3 and beyond . . . . .	16
Cluster counting reconstruction with classical supervised learning and transfer learning . . . . .	17
Docu-Bot: AI assisted user support . . . . .	18
Hardware acceleration for next-to-leading order event generation within MadGraph5_aMC@NLO . . . . .	18
Optimal use of timing measurement in vertex reconstruction at CMS . . . . .	18
A new SymPy backend for vector: uniting experimental and theoretical physicists . . . . .	19
Investigation of SVM-based 2D Heterojunctions for Data Filtering in Tracking System . . . . .	19
Benchmarking OSDF services to develop XrootD best practices . . . . .	20
Graph-based Full Event Interpretation: a graph neural network for event reconstruction in Belle II . . . . .	20
On Demand Column Joining for End User Analysis . . . . .	21
Running version control systems and continuous integration to fulfill CERN and LHC experiments needs . . . . .	21
Global Networking Challenges for the Coming Decade . . . . .	22
Monte Carlo challenges for Non Perturbative QED . . . . .	22
Monitoring particle accelerators with wireless IoT . . . . .	23
Ending the Analogue Telephony Era . . . . .	23

Upcoming database developments at CERN . . . . .	24
Accounting of (HPC) computing resources with AUDITOR . . . . .	24
Scalable and multi-tenant Kubernetes ingress infrastructure . . . . .	25
FTS3 Token Support for a Proxy-less WLCG world . . . . .	25
Ceph at CERN in the multi-datacentre era . . . . .	26
FTS as a part of the SKA data movement pipeline . . . . .	26
S3 Compatibility: Enabling Seamless Integration with EOS, CERN's Large-Scale Disk Storage System . . . . .	26
WLCG Operations: evolution and challenges on the way to High Luminosity LHC . . . . .	27
The neXt Dirac incarnation . . . . .	27
Advancements and Operations for LHC Run-3 and beyond . . . . .	28
Next-Gen Storage Infrastructure for ALICE: Paving the Road Toward Hi-Luminosity LHC . . . . .	28
Physics Data Forge: Unveiling the Power of IO Systems in CERN's Massive IO Scale Test Infrastructure . . . . .	29
Tracking efficiency studies for LHCb in Run 3 . . . . .	29
Exploring Data Caching Policy with Data Access Patterns from dCache Logs . . . . .	29
Comparing Cache Utilization Trends for Regional Data Caches . . . . .	31
Computing and machine learning for the SHiP experiment . . . . .	32
Carbon, Power, and Sustainability in ATLAS Computing . . . . .	32
The First Release of ATLAS Open Data for Research . . . . .	33
Integration of the ACTS track reconstruction toolkit in the ATLAS software for HL-LHC operations . . . . .	33
Containerization in the ATLAS Offline Code Management System . . . . .	34
ATLAS analysis workflows using the EventIndex and the Event Picking Server for massive event picking and enhanced processing . . . . .	34
Recent Advances in the GAN-based Fast Calorimeter Simulation of the ATLAS Experiment . . . . .	35
Advancements in the ATLAS Fast Chain for HL-LHC: Towards Efficient MC Production . . . . .	35
Towards Machine-Learning Particle Flow with the ATLAS Detector at the LHC . . . . .	36
New approaches for fast and efficient graph construction on CPU, GPU and heterogeneous architectures for the ATLAS event reconstruction . . . . .	36
ATLAS High-Luminosity LHC demonstrators with Data Carousel: Data-on-Demand and Tape Smart Writing . . . . .	36

Optimizing the ATLAS Geant4 detector simulation . . . . .	37
Using and Visualizing Graphs and Graph Algorithms . . . . .	37
Future scheduling in Athena on Accelerators and Multiple Nodes . . . . .	38
Towards an experiment-independent toolkit for fast calorimeter simulation . . . . .	38
AtlFast3: Fast Simulation in ATLAS for LHC Run 3 and beyond . . . . .	39
Exploring Client-Server Scalability with RNTuple & EOS: Comparative Analysis of Physics Data Formats for Analysis . . . . .	40
Quantum error mitigation for Fourier moments computation . . . . .	40
Diving into large-scale congestion with NOTED as a network controller and machine learning- based traffic forecasting . . . . .	41
Scitags: A Standardized Framework for Traffic Identification and Network Visibility in Data-Intensive Research Infrastructures . . . . .	41
Machine Learning Inference in Athena with ONNXRuntime . . . . .	42
AthenaTriton: A Tool for running Machine Learning Inference as a Service in Athena . . . . .	42
The Electronics Simulation Software in JUNO . . . . .	43
Building a Columnar Analysis Demonstrator for ATLAS PHYSLITE Open Data using the Python Ecosystem . . . . .	43
A Pilot Analysis Facility at CERN, Architecture, Implementation and First Evaluation . . . . .	44
interTwin - an interdisciplinary Digital Twin Engine for Science . . . . .	44
Jet reconstruction with quantum algorithms . . . . .	45
Improvements of the GPU Processing Framework for ALICE . . . . .	46
Enhancing software-hardware co-design for HEP by low-overhead profiling of single- and multi-threaded programs on diverse architectures with AdaptivePerf . . . . .	46
Evolution of the CERN Tape Archive scheduling system . . . . .	47
Challenges of repack in the era of the high-capacity tape cartridge . . . . .	47
Navigating the Multilingual Landscape of Scientific Computing: Python, Julia, and Awk- ward Array . . . . .	48
Scaling TraceWin beam dynamics simulations on Kubernetes for Reinforcement Learning training . . . . .	48
Thoroughly testing and integrating hundreds of Pull Requests per month: ROOT's new Cost-efficient and Feature Rich GitHub-based CI . . . . .	49
Distributed analysis in production with RDataFrame . . . . .	49
The ROOT Project at the end of Run 3 and towards HL-LHC . . . . .	50

Integration of the Goettingen HPC cluster Emmy to the WLCG Tier-2 centre GoeGrid and performance tests . . . . .	50
Parameter Estimation in ATLAS with Neural Simulation-Based Inference . . . . .	51
ML-based Adaptive Prefetching and Data Placement for US HEP systems . . . . .	51
The Multi-Threaded Detector Simulation in JUNO . . . . .	53
EDM4hep.jl: Analysing EDM4hep files with Julia . . . . .	53
On-the-fly data set joins and concatenations with ROOT RNTuple . . . . .	54
CMS FlashSim: end-to-end simulation with ML . . . . .	54
A Multi-Objective Optimization Tool for Track Reconstruction in CMS . . . . .	54
Adoption of the alpaka performance portability library in the CMS software . . . . .	55
Challenges of the CMS Level-1 Trigger for the HL-LHC Upgrade . . . . .	56
Simulating the CMS High Granularity Calorimeter with ML . . . . .	56
Pile-up estimation with Graph Neural Networks for LHC Experiment . . . . .	56
Neural network clusterization for the ALICE TPC online computing . . . . .	57
Data discovery, analysis and reproducibility in Virtual Research Environments . . . . .	57
Extending ALICE's GPU tracking capabilities: Towards a comprehensive accelerated barrel reconstruction . . . . .	58
Leveraging public cloud resources for the processing of CMS open data . . . . .	58
Madgraph on GPUs and vector CPUs: towards production . . . . .	59
TGeoArbN - A ROOT compatible triangle mesh geometry implementation . . . . .	59
Whole-node scheduling in the ALICE Grid: Initial experiences and evolution opportunities . . . . .	60
Web-based graphics in ROOT . . . . .	60
Optimized Graph Convolution for calorimetry event classification . . . . .	61
Ranking-based machine learning for track seed selection in ACTS . . . . .	61
Low-latency AI for triggering on electrons at High Luminosity LHC with the CMS Level-1 hardware Trigger. . . . .	62
Migration of CADI to Fence . . . . .	62
Enabling Alternative Architectures in the ALICE Grid . . . . .	63
AllocMonitor: An Extensible Memory Monitor . . . . .	63
Unprivileged subdivision of job resources within the ALICE Grid . . . . .	64

RNTuple: A CMS Perspective . . . . .	64
The CMS Interactive Job Tracer . . . . .	65
CMS L1 Data Scouting for HL-LHC . . . . .	65
ROOT RNTuple: Next Generation Event Data I/O for HENP . . . . .	66
Development of machine-learning based app for anomaly detection in CMSWEB . . . . .	66
Real-time Anomaly Detection at the L1 Trigger of CMS Experiment . . . . .	66
A CXL-based Memory Lake Architecture for 40 MHz L1 Data Scouting at CMS . . . . .	67
Particle Flow Reconstruction with Alpaka Portability Library . . . . .	67
Direct I/O for RNTuple Columnar Data . . . . .	68
Leveraging distributed resources through high throughput analysis platforms for enhancing HEP data analyses . . . . .	68
Comparative efficiency of HEP codes across languages and architectures . . . . .	69
Towards a GPU-enabled electron seeding algorithm in the CMS experiment . . . . .	69
Continuous integration of analysis workflows on a distributed analysis facility . . . . .	70
The git based ATLAS data acquisition configuration service in LHC Run 3 . . . . .	70
Leveraging the Run 3 experience for the evolution of the ATLAS software-based readout towards HL-LHC . . . . .	71
GPU Acceleration and EDM Developments for the ATLAS 3D Calorimeter Clustering in the Software Trigger . . . . .	71
Development of an FPGA based track reconstruction pipeline for the ATLAS Event Filter . . . . .	72
Turning CephFS into a collaborative cloud storage with CERNBox . . . . .	72
Online track reconstruction with graph neural networks on FPGAs for the ATLAS experiment . . . . .	73
Performance of the ATLAS GNN4ITk Particle Track Reconstruction GPU pipeline . . . . .	73
Implementation and Performance Analysis of the ALICE grid middleware JAliEn's Job Optimizer . . . . .	73
QDIPS: Deep Sets Network for FPGA investigated for high speed inference on ATLAS . . . . .	74
Evolution of the ATLAS TDAQ online software framework towards Phase-II upgrade: use of Kubernetes as an orchestrator of the ATLAS Event Filter computing farm . . . . .	75
Updates of the ATLAS High-Level Trigger in Run 3 . . . . .	75
The upgrade of the ATLAS Trigger and Data Acquisition system for the High Luminosity LHC . . . . .	76
Integrating FPGA Accelerators in Athena for use at the ATLAS Event Filter . . . . .	76



The ATLAS Trigger System . . . . .	77
ML-based classification in an open-source framework for the ALICE heavy-flavour analysis . . . . .	77
Unified Experiment Monitoring . . . . .	78
Deployment of inference as a service at the US CMS Tier-2 data centers . . . . .	78
An automation framework for condition data processing using data-streaming techniques and ML models . . . . .	78
Cold data support for the CERN Open Data Portal . . . . .	79
The Global Network Advancement: A Next Generation System for Data Intensive Sciences . . . . .	79
QiboML: a full-stack quantum machine learning framework . . . . .	80
Canadian Belle-II Tier1 Raw Data Centre . . . . .	81
Using S3 storage with xrootd vs Dynafed . . . . .	81
The BaBar Long Term Data Preservation and Computing Infrastructure . . . . .	81
Data Challenges in JUNO distributed computing infrastructure towards JUNO data-taking	82
Hard scatter vertex identification in ATLAS using graph neural networks . . . . .	82
Adoption of ROOT RNTuple for the next main event data storage technology in the ATLAS production framework Athena . . . . .	83
Improving Computational Performance of ATLAS GNN Track Reconstruction Pipeline .	83
GNN based 3D calorimeter cell clustering pipeline for the ATLAS detector . . . . .	84
ATLAS WLCG Data Challenge 2024 planning and implementation . . . . .	85
ATLAS software tools to handle ROOT RNTuple . . . . .	85
Impact of RNTuple on storage resources for ATLAS production . . . . .	86
High Performance Graph Segmentation for ATLAS GNN Track Reconstruction . . . . .	86
Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS . . . . .	87
Building Scalable Analysis Infrastructure for ATLAS . . . . .	87
Measurement and Analysis of ATLAS applications (CPU and Storage) energy consumption and data center optimization . . . . .	88
Generative machine learning for fast silicon detector simulation . . . . .	88
An implementation of cloud-based grid CE and SE for ATLAS and Belle II . . . . .	89
Leveraging Language Models for Enhanced Code Review in Particle Physics Software De- velopment . . . . .	90

Application of quantum graph neural network on detector data . . . . .	90
Fast Jet Reconstruction in Julia . . . . .	90
Advancing Large-Scale Scientific Collaborations with Rucio: A Data Management Story	91
R&D towards heterogenous frameworks for Future Experiments . . . . .	92
Particle Identification Algorithms Based on Machine Learning for STCF . . . . .	92
TechWeekStorage24: Open Source Data Technologies for the Exabyte Era . . . . .	92
dCache-CTA Integration: One Year in Production at DESY . . . . .	93
Evaluating FPGA Acceleration with Intel® oneAPI Toolkit for High-Speed Data Processing . . . . .	93
Generalization and self supervision for physics simulation . . . . .	94
Prometheus-Powered Insight: Monitoring Koji's performance . . . . .	94
Managing accelerator control systems with Kubernetes at CERN . . . . .	95
Extending the Gaudi Software Framework outside of C++ . . . . .	95
Archive Metadata for efficient data colocation on tape . . . . .	96
An online data processing system for the CMS Level-1 Trigger data scouting demonstrator . . . . .	96
PATOF: From the Past To the Future: Legacy Data in Small and Medium-Scale "PUNCH" Experiments - a Blueprint for PUNCH and Other Disciplines . . . . .	97
Facilitating Scientific Reproducibility and Interoperability through CWL Integration in the Dirac Grid Middleware . . . . .	98
Architecting software applications in containerized environment for CMS data acquisition . . . . .	99
A Lightweight Analysis & Grid Facility for the DARWIN Experiment . . . . .	100
Thread-safe N-tuple Writing in Gaudi with TTree and Migration to RNTuple . . . . .	100
Next generation geometry model for Tracking in ACTS . . . . .	101
Improving overall GPU sharing and usage efficiency with Kubernetes . . . . .	101
Preliminary findings and recommendations from the Token Trust and Traceability Working Group. . . . .	102
Portable HCAL reconstruction in the CMS detector using the Alpaka library . . . . .	103
AdePT - Enabling GPU electromagnetic transport with Geant4 . . . . .	103
Towards an Introspective Dynamic Model of Globally Distributed Computing Infrastruc- tures . . . . .	103
Exploiting GPU Resources at VEGA for CMS Software Validation . . . . .	104

First Deployment of XCache for Workflow and Efficiency Optimizations on Opportunistic HPC Resources in Germany . . . . .	105
The roles of the SCAB Nominations and Activities systems in the ATLAS-CERN speakers selection . . . . .	105
SYCL-based online data processing framework concept for PANDA . . . . .	106
Enhancing product management in the ATLAS Management Glance team: challenges and solutions . . . . .	106
The Glance project common infrastructure dependencies upgrade from the ATLAS Glance perspective . . . . .	107
ATLAS Qualification interface refactoring strategy . . . . .	108
HEP-Help: a first-stop helpline for particle physics software . . . . .	108
The evolution of the CMS@Home project . . . . .	109
Application of TRACCC seeding to the CEPC vertex detector . . . . .	109
Tracking for long-lived particles at STCF . . . . .	110
Opticks : GPU ray trace accelerated optical photon simulation . . . . .	111
ATLAS usage of the Czech national HPC center: HyperQueue, cvmfsexec, and other news	111
The new configuration system for ATLAS event processing applications . . . . .	111
pip install ROOT: experiences making a complex multi-language package accessible for Python users . . . . .	112
Zero-overhead training of machine learning models with ROOT data . . . . .	112
Computing Activities at the Spanish Tier-1 and Tier-2s for the ATLAS experiment in the LHC Run3 period and towards the High-Luminosity Phase (HL-LHC) . . . . .	113
Improving the ATLAS Production System for Run 3: Achievements, Lessons, and the Future . . . . .	114
Enhancements and resource optimisations for ATLAS use of HammerCloud . . . . .	114
Looking back on 23 years of the ATLAS Metadata Interface (AMI): Reflections on database architecture and ongoing developments in preparation for Run 4 . . . . .	115
Delivering large scale public engagement open weeks at STFC . . . . .	115
Data Placement Optimization for ATLAS in a Multi-Tiered Storage System within a Data Center . . . . .	116
Using the ATLAS experiment software on heterogeneous resources . . . . .	116
Optimisation of ATLAS computing resource usage through a modern HEP Benchmark Suite via HammerCloud and PanDA . . . . .	117
CREST Demonstrator for HL-LHC in ATLAS . . . . .	117

Performance studies for ATLAS workloads in many-core Grid and HPC environments . . . . .	118
Integrating the Perlmutter HPC system in the ALICE Grid . . . . .	118
Modernizing ATLAS PanDA for a sustainable multi-experiment future . . . . .	119
Advancements in the in-file metadata system for the ATLAS experiment . . . . .	120
Advancing ATLAS DCS Data Analysis with a Modern Data Platform . . . . .	120
Microservices framework and configuration database for ATLAS ITk . . . . .	121
Updating the software description of the ATLAS Detector . . . . .	121
A simplified configuration for common algorithms for ATLAS analysis . . . . .	122
Total Cost of Ownership and Evaluation of Google Cloud Resources for the ATLAS Experiment at the LHC . . . . .	122
Enhancing Network Analytics through Machine Learning . . . . .	122
Infrastructure for deployment and evaluation of LHCb trigger configurations . . . . .	123
Investigating Data Access Models for ATLAS: A Case Study with FABRIC Across Borders and ServiceX . . . . .	124
Evolution of the ATLAS event data model for the HL-LHC . . . . .	124
NeuroMCT: Fast Monte Carlo Tuning with Generative Machine Learning in the JUNO Experiment . . . . .	125
Evolving MonALISA: Modernizing the Communication Layer of a Subagent-Based Distributed Monitoring System . . . . .	125
Decode the Workload: Training Deep Learning Models for Efficient Compute Cluster Representation . . . . .	126
CERN OSPO's Catalogue Initiative: Amplifying Open Science for the CERN Community	127
Porting and optimizing the performance of LArTPC detector simulations on Intel GPU with various programming models and compilers . . . . .	127
ML-Assisted Charged Particle Tracking at GlueX . . . . .	128
How to make kubernetes easy to use for thousands users . . . . .	128
EggNet: An Evolving Graph-based Graph Attention Network for End-to-end Particle Track Reconstruction . . . . .	128
Xiwu: A basic flexible and learnable LLM for High Energy Physics . . . . .	129
Boost physics study at HEP experiments with Dr. Sai . . . . .	129
Unsupervised Learning Techniques for Identification of Anomalous LZ Data . . . . .	130
Storage of nuclear waste suitable for non-invasive monitoring using muon scattering tomography . . . . .	130

End-to-end event simulation with Flow Matching and generator Oversampling . . . . .	131
End-to-end ML-based reconstruction . . . . .	132
ATLAS Tile Calorimeter Software Tools for Data Quality Monitoring . . . . .	132
Optimizing cut-based algorithms to specific physics acceptance regions . . . . .	133
Data Movement Model for the Vera C. Rubin Observatory . . . . .	133
Surrogate Modeling for Scalable Evaluation of Distributed Computing Systems for HEP applications . . . . .	134
Towards end-to-end particle flow in high granularity detectors . . . . .	134
EDM4hep - The common event data model for the Key4hep project . . . . .	135
OMS Data Aggregation and Management in the CMS Experiment . . . . .	135
Lessons Learned from Writing My Dissertation With Quarto . . . . .	136
Validation of Shoveler XRootD monitoring . . . . .	137
Heterogeneous computing at INFN-T1 . . . . .	137
The Big Data Processing Infrastructure for monitoring and analysing the ATLAS experiment processing activities at INFN-CNAF Tier-1 . . . . .	138
Real-time compression of high-speed Streaming Readout signals with Artificial Intelligence algorithms. . . . .	138
Evaluating a File-Based Event Builder to enhance the Data Acquisition in the CMS Experiment . . . . .	139
ML-based classification of photons and neutral mesons for direct photon measurement in ALICE . . . . .	140
A Tape RSE for Extremely Large Data Collection Backups . . . . .	141
Reconstruction of Full Decays using Transformers and Hyperbolic Embedding at Belle II	141
Integrated Research Infrastructure with JIRIAF . . . . .	142
So FAIR, so good: the INFN strategy for Data Stewardship . . . . .	142
Design and developing status of the data acquisition and the slow control systems in the Mu2e experiment at Fermilab . . . . .	143
Adoption of the alpaka performance portability library in the CMS software . . . . .	143
Adapting to Change: A look at the evolution of ALICE's Quality Control framework . .	144
A graph neural network based cosmic muon trigger for the Mu3e experiment . . . . .	144
Reconstruction in Key4hep using Gaudi . . . . .	145
Building the Key4hep Software Stack with Spack . . . . .	145

Particle Identification at LHCb combining Neural Networks and evolutionary computation . . . . .	146
Experience with ARM WNs at the WLCG Tier1 GridKa . . . . .	146
Prévessin Data Centre Powers Up . . . . .	147
The End of an Era: Termination of the Meyrin Data Centre Console Service . . . . .	147
Automated Template Testing and Management with GitLab Continuous Integration for AMS Offline Computing . . . . .	148
Recent Experience with the CMS Data Management System . . . . .	148
The Fermilab experience transitioning to CERN Tape Archive (CTA) . . . . .	149
ALICE Event Display - from the legacy ROOT-based visualization to the web-based application . . . . .	149
ROOT RNTuple implementation in Julia programming language . . . . .	150
GIL-free scaling of Uproot with Python 3.13 subinterpreters . . . . .	150
Offline data processing software for the High Energy cosmic-Radiation Detection facility	151
Implementation of an Open Drift Chamber in ACTS . . . . .	152
Implementation and development of a DAQ system DELILA at ELI-NP . . . . .	152
Carbon costs of storage: a UK perspective. . . . .	153
Implementing a Domain-Specific Language (DSL) for Physics Data Analysis in RootInteractive . . . . .	153
Evolution of the ATLAS TRT Gas Gain Stabilization System Software . . . . .	154
Advancements in Computing and Simulation Techniques for the HIBEAM-NNBAR Experiment . . . . .	154
AI-based approach for provider selection in the INDIGO PaaS Orchestration System of INFN Cloud . . . . .	155
Tracking and vertexing downstream the LHCb magnet at the first stage of the trigger . .	156
A Streamlined Neural Model for Real-Time Analysis at the First Level of the LHCb Trigger . . . . .	156
Six months of Single Sign On on Kubernetes . . . . .	156
BuSca: a Buffer Scanner at 30 MHz for New Long-Lived Particle Searches at LHCb . . .	157
Heterogeneous Computing and Power Efficiency in HEP . . . . .	157
Taking on RISC for Energy-Efficient Computing in HEP . . . . .	158
Parsing TTreeFormula in Python . . . . .	158
CppInterOp: Advancing Interactive C++ for High Energy Physics . . . . .	159

On-Grid GPU development via interactive HTCondor jobs and Analysis Facility style workflows . . . . .	160
User sharing of computational workflows in the REANA reproducible analysis platform	160
Benchmarking massively-parallel Analysis Grand Challenge workflows using Snakemake and REANA . . . . .	161
Reading Tea Leaves - Understanding internal events and addressing performance issues within a CephFS/XRootD Storage Element. . . . .	161
EvtGen –on its first steps towards thread safety . . . . .	162
CMS Token Transition . . . . .	162
Heterogeneous reconstruction of hadronic Particle Flow clusters with the Alpaka Portability Library . . . . .	163
Efficiency, Reproducibility, and Portability in HEP Machine Learning Training - ML Training Facility at Vanderbilt University . . . . .	163
A Multi-objective Graph Neural Network for enhanced HGAL hadronic reconstruction	164
Object storage model for CMS data . . . . .	164
Anomaly detection for data quality monitoring of the Muon system at CMS . . . . .	164
Enhancing GRB Detection: Machine Learning Optimization of Triggerless Data Analysis Algorithms for LHAASO-WCDA . . . . .	165
A RoCE-based network framework for science workloads in HEPS data center . . . . .	165
Keep-up Production in JUNO’s Offline Data Processing . . . . .	166
Enhancing XRootD Load Balancing for High-Throughput transfers . . . . .	166
Primary vertex timing reconstruction with the LHCb Ring Imaging Cherenkov detectors	167
Research on Wide Area Network Performance Anomaly Detection Technology Based on Machine Learning . . . . .	167
Taking derivatives of Geant4 - closer than you might think? . . . . .	168
TrackHHL: A Quantum Computing Algorithm for Track Reconstruction at the LHCb . .	168
TrackNET: Deep Learning-Based Track Recognition in Pixel and Strip-Based Particle Detectors . . . . .	169
Guided Quantum Compression for Higgs Identification . . . . .	169
Evolving StoRM WebDAV: delegation of file transfers to NGINX and support for SciTags	170
Benchmark Studies of ML Inference with TMVA SOFIE . . . . .	170
Front-End RDMA Over Converged Ethernet, lightweight RoCE endpoint . . . . .	171
Use of topological correlations in ML-based conditions for the CMS Level-1 Global Trigger upgrade for the HL-LHC . . . . .	171

The CMS Phase-2 High Granularity Calorimeter Endcap Event Reconstruction with the TICL Framework . . . . .	172
Fermilab's Transition to Token Authentication . . . . .	173
Automation and Job Management for LZ Simulations at NERSC . . . . .	173
tracc: GPU track reconstruction library for HEP experiments . . . . .	174
Packaging HEP heterogeneous mini-apps for portable benchmarking and facility evaluation on modern HPCs . . . . .	174
HEP-CCE Phase 2 Overview and Outlook . . . . .	175
Reconstruction Framework Advancements for Streaming Readout for the ePIC Experiment at the EIC . . . . .	175
Particle Identification at STCF DTOF Detector Based on Classical/Quantum Convolutional Neural Network . . . . .	176
Anomaly Detection using Autoencoders on Fundamental LZ Signals . . . . .	176
Enhancing CMS XCache efficiency: A comparative study of Machine Learning techniques and LRU mechanisms . . . . .	177
R&D Adoption and Progress in Full Simulation of the CMS experiment . . . . .	177
Time Based Clustering for Muon Chamber(MuCh) in the CBM experiment at FAIR, Germany . . . . .	178
Di-Muon cocktail reconstruction using Machine Learning technique in CBM experiment at FAIR SIS100 energies . . . . .	179
DUNE Rucio development and monitoring . . . . .	179
GPU-Accelerated Point Clouds Library for GNN-based Reconstruction and Beyond . . . . .	180
Performance of the parallelized General Triplet Track Fit implemented on the GPU . . . . .	180
Efficient and fast container execution using image snapshotters . . . . .	181
FPGA implementation of the General Triplet Track Fit . . . . .	181
XKIT for GridPP: (XRootD Kubentes Integration Testing for GridPP) . . . . .	182
Achieving 100Gb/s data rates with XRootD - Preparing for HL-HLC and SKA . . . . .	182
Fair Universe HiggsML Uncertainty Challenge . . . . .	183
The path to exabyte astronomy: SRCNet v0.1 for the Square Kilometre Array . . . . .	184
High-throughput data distribution for CBM online computing . . . . .	184
Scientific NREC Cluster - a fast and flexible solution for virtual computing clusters on OpenStack . . . . .	185
The NextGen Triggers project: overview, plans and first actions . . . . .	185



Allocating Carbon Costs to Computing Payloads across Heterogeneous Infrastructures . . .	186
GlitchFlow, a Digital Twin for transient noise in Gravitational Wave Interferometers . . .	186
Design and construction of High Energy Photon Source (HEPS) scientific data storage system . . . . .	187
Software defined network control for LHC Experiments . . . . .	187
Distributed Data Management with Rucio for the Einstein Telescope . . . . .	188
GPU-friendly surface model for Monte-Carlo detector simulations . . . . .	189
Prospects (and Hurdles) of the Computing Facility for FAIR Research . . . . .	189
Data Movement Manager (DMM) for the SENSE-Rucio Interoperation Prototype . . . . .	189
Declarative paradigms for data analysis description and implementation . . . . .	190
Optimization of distributed compute resources utilization in the CMS Global Pool . . . . .	191
Operational experience from the Spanish CMS Analysis Facility at CIEMAT . . . . .	191
Benchmarking XRootD-HTTPS on 400Gbps Links with Variable Latencies . . . . .	192
Simulation Comparison for the mSTS Geometry based on Primitive ROOT/TGeo Solids and Tessellated Solids . . . . .	192
The Cherenkov Telescope Array Observatory Production System Status and Development	193
Commissioning and exploitation of the MareNostrum5 cluster at the Barcelona Supercomputing Center for CMS computing . . . . .	194
Cache Rules Everything Around Me: Building ePIC Containers With Spack . . . . .	194
Towards the Einstein Telescope Computing Model . . . . .	195
Just-in-time workflow management for DUNE . . . . .	195
The 200Gbps Challenge: Imagining HL-LHC analysis facilities . . . . .	196
Orchestrated columnar-based analysis with columnflow . . . . .	196
Novel Fitting Approach Based on a Neural Network for JUNO . . . . .	197
A high-throughput input interface for the CBM FLES . . . . .	197
Towards an open-source hybrid quantum operating system . . . . .	198
Data-flow parallelism for high-energy and nuclear physics frameworks . . . . .	199
Multi-package development at Fermilab with Spack . . . . .	199
Checkpoint-Restart for HPC . . . . .	200
Machine Learning for Optimized Polarization at Jefferson Lab . . . . .	200
Efficient metadata management with the AMI ecosystem . . . . .	200

Near-Data Computing Model for Accelerating LHC Data Filtering . . . . .	201
Automatic and effective discovery of quantum kernels . . . . .	201
Overview of the data acquisition system architecture for the DarkSide-20k experiment . . . . .	202
Hydra: An AI-Based Framework for Interpretable and Portable Data Quality Monitoring . . . . .	203
Event Workflow Management System - A SaaS Solution for Massively Divisible and Distributed Workflows . . . . .	203
Development and Testing of DUNE's Production System . . . . .	204
Application of linear and non-linear constraints in a brute-force based alignment approach for CBM . . . . .	204
LHCb Open Data Ntupling Service: On-demand production and publishing of custom LHCb Open Data . . . . .	205
Sprucing and Analysis Productions: Offline data processing in LHCb without the pain . . . . .	206
IceSONIC - Network AI Inference on Coprocessors for IceCube Offline Processing . . . . .	206
Online Data Processing and monitoring for DarkSide-20k . . . . .	206
The Real-Time Data Workflow of LZ Dark Matter Experiment at NERSC . . . . .	207
Quantum-Assisted Generative AI for Simulation of the Calorimeter Response . . . . .	207
Comprehensive Description and Orchestration of Complex Data Processing Pipelines . . . . .	208
Development of a WebRTC-Based Remote Desktop for Neutron Scattering Data Analysis on Openstack . . . . .	208
Likelihood analysis methods for the PID system of STCF experiment . . . . .	209
Preparation for Multi-Site Processing at the Vera C. Rubin Observatory . . . . .	209
ProtoDUNE Run Conditions Database . . . . .	210
LHCb Stripping Project: Continuing to Fully and Efficiently Utilize Legacy Data . . . . .	210
Advancements in Offline Analysis Software for Run III at LHCb . . . . .	211
Leveraging workflow engines and computing frameworks for physics analysis scalability and reproducibility . . . . .	212
Provisioning of Grid computing resources in the Norwegian Research and Education Cloud . . . . .	212
Track classification in Active Target Time Projection Chamber using supervised machine learning methods . . . . .	213
Tuning the CMS Coffea-casa facility for 200 Gbps Challenge . . . . .	214
Integration of A Supercomputing Center in LHAASO Distributed Computing System . . . . .	214
Designing Operational Security systems: People, Processes and Technology . . . . .	215

Simulation of a MPGD-based hadronic calorimeter for Muon Collider experiments . . . . .	216
Belle II Network Performance Analysis in the Context of the WLCG Data Challenge 2024	216
The Belle II Raw Data Transfer System . . . . .	217
Kalman filter for muon reconstruction in the CMS Phase-2 endcap calorimeter . . . . .	217
Enhancing Data Management in Nuclear Physics and Cross-Domain Metadata Integration . . . . .	218
Exploiting Kubernetes to Simplify the Deployment and Management of the Multi-purpose CMS Pilot Job Factory . . . . .	218
Geometry Simplification Methods for Virtual Reality Applications . . . . .	219
Quantum Machine Learning for track reconstruction . . . . .	220
HEPS scientific computing system design for interactive data analysis scenarios . . . . .	220
RICH ring reconstruction using machine learning for CBM . . . . .	221
Distributed Computing Infrastructure for HERD experiment . . . . .	221
The CloudVeneto's Container-as-a-Service ecosystem . . . . .	222
BAT.jl, the Bayesian Analysis Toolkit in Julia . . . . .	222
Evolving INDIGO IAM towards the next challenges . . . . .	223
The new hardware trigger processor at NA62 experiment: Status of the System and First Results . . . . .	223
Distributed management and processing of ALICE monitoring data with Onedata . . . . .	224
Revolutionising Radio Astronomy: The UK's Role in SKA's SRCNet Deployment . . . . .	225
Designing the Identity Infrastructure for the Square Kilometre Array Regional Centre Net- work . . . . .	226
A Successful Data Centre Refurbishment Project . . . . .	226
CVMFS: Pushing performance on highly parallel, many-core clients . . . . .	227
Machine learning based event reconstruction for the MUonE experiment . . . . .	227
Implementing XRootD/SciToken-Based Access to Lustre Storage at GSI: A First Step To- ward Data Federation for FAIR . . . . .	228
LbMCSUBMIT: Streamlined production and submission of LHCb MC requests . . . . .	228
Versioning of the LHCb DD4Hep-based Detector Description . . . . .	229
Towards an IPv6-only WLCG: more successes in reducing IPv4 . . . . .	229
A Comprehensive Bandwidth Testing Framework for the LHCb Upgrade Trigger System	230

Prospects for novel track reconstruction algorithms based on Graph Neural Network models using telescope detector testbed . . . . .	231
SPECTRUM: towards a Computing Strategy for Data-intensive Science Infrastructures in Europe . . . . .	232
The ePIC Simulation Campaign Workflow on the Open Science Grid . . . . .	232
Latest developments of the PUNCH4NFDI compute and storage infrastructures . . . . .	233
Performance of the LHCb heterogeneous software trigger . . . . .	234
Track reconstruction in the STCF detector . . . . .	234
KKMCee: Multiphoton MC for lepton and quark pair production at lepton colliders . . . . .	235
Recent developments in the Gaussino simulation software . . . . .	235
Julia in HEP . . . . .	235
Generative AI for fast simulations in LHCb . . . . .	236
A point library for the fast simulation of the LHCb Calorimeter . . . . .	237
A Geometry Agnostic Heterogenous Framework for Track Reconstruction for HEP Experiments . . . . .	237
MEASUREMENT OF PROPAGATION TIME WITH THE USE OF FPGA DEVICES IN ALICE EXPERIMENT AT CERN . . . . .	238
Up-scaling for measuring the spatial distribution of radiation dose for applications in the preparation of individual patient treatment plans . . . . .	238
Experiences from the CBM collaboration: CAD to ROOT conversion for detector geometries . . . . .	239
Data-driven efficiencies of the LHCb High Level Trigger in Run 3 . . . . .	240
HEPCloud Facility Operations at Fermilab—The First Six Years . . . . .	240
Development of the platform for simulation of spatial distribution for therapeutic dose in Dose-3D phantom . . . . .	241
Evolution and Broadening of the National Analysis Facility at DESY . . . . .	242
HIGH PERFORMANCE ALGORITHMS FOR LOW POWER SUSTAINABLE HARDWARE IN HEP AT VALENCIA . . . . .	242
Operations Developments at the National Analysis Facility at DESY . . . . .	243
Modular Experiment Control System packages for the CBM experiment . . . . .	243
label-based virtual directories in the dCache storage system . . . . .	244
Ultra-Fast Geometry-Independent Highly-Granular Calorimeter Simulation with Diffusion Point Clouds . . . . .	244
A Managed Tokens Service for Securely Keeping and Distributing Grid Tokens . . . . .	245

zfit: general likelihood model fitting in Python . . . . .	245
Moving a data center keeping availability at the top . . . . .	246
Adaptive Hough Transform for Charged Particles Tracking at the LHC . . . . .	246
AccGPT: A CERN Knowledge Retrieval Chatbot . . . . .	247
Collaborative software and maintainability for ePIC experiment at EIC . . . . .	248
Monitoring large-scale dCache installations with storage events using Kafka streams . . . . .	248
FPGA-RICH: A low-latency, high-throughput online particle identification system for the NA62 experiment . . . . .	249
Collaborative Tools for the ePIC Experiment . . . . .	249
From Hope to Heuristic: Realistic Runtime Estimates for Quantum Optimisation in NHEP . . . . .	250
Reshaping Analysis for Fast Turnaround . . . . .	251
Hardware and software design of APenetX: a custom FPGA-based NIC for scientific computing . . . . .	251
Supporting the development of Machine Learning for fundamental science in a federated Cloud with the AI_INFN platform . . . . .	252
An optimized C++ software for the management of Timepix4 data acquisition and analysis . . . . .	253
Supporting medium/small-sized experiments in the transition from X.509 to JWTs . . . . .	253
Efficient Tracking Algorithm Evaluations through Multi-Level Reduced Simulations . . . . .	254
Identification of muon-electron elastic events using Graph Neural Networks for precision measurements. . . . .	255
Free-streaming online tracking in CBM . . . . .	255
New GridKa Tape Storage System –from design to production deployment . . . . .	256
Accelerating detector simulations with Celeritas: performance improvements and new capabilities . . . . .	256
Open Data at ATLAS: Bringing TeV collisions to the World . . . . .	257
Recent Updates to the Popular ATLAS Virtual Visit Programme . . . . .	257
Unlocking the compute continuum: scaling out from cloud to HPC and HTC resources . . . . .	258
Evolution of Regional, Age and Gender Demographics in the ATLAS Collaboration . . . . .	259
A Cloud oriented platform to exploit ANSYS application . . . . .	259
Efficient ML-Assisted Particle Track Reconstruction Designs . . . . .	260
An online GPU hit finder for the STS detector in the CBM experiment . . . . .	260

Energy-efficient graph-based algorithm for tracking at the HL-LHC . . . . .	261
Event generation with quantum computers through particle-oriented simulation . . . . .	261
Hadronization effects using String and Cluster Models in Herwig 7 . . . . .	262
GlideinBenchmark: collecting resource information to optimize provisioning . . . . .	262
Open Science and Compressed Baryonic Matter experiment . . . . .	263
Parallel Photon Simulation for IceCube In C++ . . . . .	263
Geant4 models for nuclear de-excitation . . . . .	264
Zero Degree Calorimeter fast simulation with normalizing flows . . . . .	264
Status of DUNE Offline Computing . . . . .	265
Numerical studies of space charge effect on particle tracking in a small TPC . . . . .	265
Towards more efficient job scheduling in ALICE: predicting job execution time using machine learning . . . . .	266
A data Quality-Assurance framework for online and offline applications for the CBM experiment . . . . .	266
A technical overview of industry-science R&D projects for the High Luminosity LHC under CERN openlab . . . . .	267
Real-time Level-1 Trigger Data Scouting at CMS using CXL Memory Lake . . . . .	267
Harnessing the power of neural networks to identify and locate primary vertices in proton-proton collision data . . . . .	268
Exploring new directions in enhancing the ACTS parameter optimization suite . . . . .	268
Addressing tokens dynamic generation, propagation, storage and renewal to secure the GlideinWMS pilot based jobs and system. . . . .	269
Real-time pattern recognition with FPGA at LHCb, an O(n) complexity architecture . . . . .	269
Real-time monitoring of LHCb interaction region with a fast trackless methodology . . . . .	270
Refining FastSim with Machine Learning . . . . .	271
Simulating the Carbon Cost of Grid Sites . . . . .	271
Quantifying the computational speedup with MG4GPU for CMS workflows . . . . .	272
Bamboo: A high-level HEP analysis library for ROOT::RDataFrame . . . . .	272
Using containers to speed up development, to run integration tests and to teach about distributed systems . . . . .	272
ProtoDUNE HD Offline Computing Infrastructure . . . . .	273
FORM, a Fine-grained Object Reading/Writing Model for DUNE . . . . .	273

Navigating Phase Space for Event Generation –interfacing Sherpa with BAT.jl . . . . .	274
Bridging Reproducibility Gaps in Partial Wave Analysis . . . . .	274
New RooFit PyROOT interfaces for connections with Machine Learning . . . . .	275
Monte Carlo efficiency via negative weight reduction in Herwig . . . . .	275
Multi-Experiment dCache Storage Service at BNL . . . . .	276
Leveraging Language Models to Navigate Conference Abstracts: An Open-Source Approach . . . . .	276
Saving through Sharing - Reflections on Sustainability in HEP . . . . .	277
An automation framework for the calibration of the CMS Precision Proton Spectrometer	277
Testing abstract . . . . .	278
The State of Event Generators for the HL-LHC . . . . .	278
Large Language Models in Physics . . . . .	278





**Parallel (Track 5) / 3**

## **HPCNeuroNet: A Neuromorphic Approach Merging SNN Temporal Dynamics with Transformer Attention for FPGA-based Particle Physics**

**Author:** Murat Isik<sup>1</sup>**Co-authors:** Hiruna Vishwamith<sup>2</sup>; I. Can Dikmen<sup>3</sup>; Jonathan Naoukin<sup>4</sup><sup>1</sup> *Drexel University*<sup>2</sup> *University Of Moratuwa*<sup>3</sup> *Temsa R&D Center*<sup>4</sup> *University of Texas at Austin***Corresponding Authors:** jnaoukin@utexas.edu, vishwamithpgh.20@uom.lk, can.dikmen@temsa.com, mci38@drexel.edu

This paper presents the innovative HPCNeuroNet model, a pioneering fusion of Spiking Neural Networks (SNNs), Transformers, and high-performance computing tailored for particle physics, particularly in particle identification from detector responses. Drawing from the intrinsic temporal dynamics of SNNs and the robust attention mechanisms of Transformers, our approach capitalizes on these synergies to achieve heightened performance in discerning intricate particle interactions. At the heart of HPCNeuroNet lies the integration of the sequential dynamism inherent in SNNs with the context-aware attention capabilities of Transformers, enabling the model to precisely decode and interpret complex detector data. HPCNeuroNet is realized through the HLS4ML framework, optimized for deployment on FPGA environments. This architectural choice not only enhances computing speed but also enhances the accuracy and scalability of the models. Benchmarked against traditional particle physics models, HPCNeuroNet showcases superior performance metrics, underlining its transformative potential in high-energy physics. Our findings illuminate the groundbreaking potential of conjoining SNNs, Transformers, and FPGA-based high-performance computing in the domain of particle physics, marking a significant stride forward and establishing a robust foundation for future endeavors in the field.

**Parallel (Track 5) / 4**

## **Jet Discrimination with Quantum Complete Graph Neural Network**

**Author:** Yi-An Chen<sup>1</sup><sup>1</sup> *National Taiwan University (TW)***Corresponding Author:** maplexworkx0302@gmail.com

Machine learning, particularly deep neural networks, has been widely utilized in high energy physics and has shown remarkable results in various applications. Moreover, the concept of machine learning has been extended to quantum computers, giving rise to a new research area known as quantum machine learning. In this paper, we propose a novel variational quantum circuit model, **Quantum Complete Graph Neural Network (QCGNN)**, designed for learning complete graphs. We argue that QCGNN has a polynomial speedup against its classical counterpart, due to the property of quantum parallelism. In this paper, we study the application of QCGNN through the challenging jet discrimination, where the jets are represented with complete graphs. Subsequently, we conduct a comparative analysis with classical graph neural networks to establish a benchmark.

**Poster session / 5**

## Library for Functional analysis in CERN's ROOT

**Author:** Tomasz Bold<sup>1</sup>

<sup>1</sup> *AGH University of Krakow (PL)*

**Corresponding Author:** tomasz.bold@cern.ch

The poster will present FunRootAna library.

This is a simple framework allowing to do ROOT analysis in a more functional way. In comparison to RDataFrame it offers more functional feel for the data analysis and can be used in any circumstances, not only with ROOT trees. Collections processing is inspired by Scala Apache Spark and the histograms creation and filling is much simplified. As consequence, a single line containing selection, data extraction & histogram definition is sufficient to obtain one unit of result. Basically, with FunRootAna the number of lines of analysis code per histogram is converging to 1. More here: <https://tboldagh.github.io/FunRootAna/>

**Poster session / 6**

## Remote3: Public Engagement over 1 km underground –and beyond

**Author:** Lauren Mowberry<sup>None</sup>

**Co-authors:** Alexander Murphy ; Emma Meehan ; Ian Collier <sup>1</sup>; Laura Murgatroyd ; Megan O'Flynn ; Sally Shaw ; Sophy Palmer ; Tom Dack ; Will Furnell

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

**Corresponding Authors:** tom.dack@cern.ch, will.furnell@stfc.ac.uk, emma.meehan@stfc.ac.uk, a.s.murphy@ed.ac.uk, lauren.mowberry@stfc.ac.uk, megan.o'flynn@stfc.ac.uk, laura.murgatroyd@stfc.ac.uk, ian.peter.collier@cern.ch, sophy.palmer@stfc.ac.uk, sally.shaw@ed.ac.uk

The Remote<sup>3</sup> (Remote Cubed) project is an STFC Public Engagement Leadership Fellowship funded activity, organised in collaboration between the University of Edinburgh (UoE), and STFC's Public Engagement Team, Scientific Computing Department, and Boulby Underground Laboratory –part of STFC Particle Physics.

Remote<sup>3</sup> works with school audiences to challenge teams of young people to design, build, and program their own LEGO Mindstorms "Mars Rover", which will be tested at the Boulby Underground Laboratory's Mars Yard, 1.1 km underground. Teams, with the assistance of mentors from UoE and STFC, will design their rover to complete various space-exploration themed challenges – ranging from taking a panoramic environment scan to navigating the Mars Yard landscape looking for LEGO brick samples. The project aims to engage with audiences who do not usually interact with STFC Public Engagement, such as more remote locations or areas of higher deprivation and give them the opportunity to work hands on with engineering and computing, whilst learning from and interacting with real scientists and engineers.

Since its inception in 2019, Remote<sup>3</sup> has flourished in a wide variety of different environments and through multiple mediums, from entirely virtual during the lockdowns of 2020-21, deep underground, in schools and storytelling at libraries, and in tents in fields at festivals.

This year Remote<sup>3</sup> is building on the lessons learnt through this varied programme to deliver a series of engagement activities in conjunction with STFC's Rutherford Appleton Laboratory Public Open Week, which has an expected audience of 20,000 people.

**Parallel (Track 8) / 7**

## The CERN Open Source Program Office: changing the (HENP) software world, pragmatically

**Authors:** Alexander Yohei Huss<sup>1</sup>; Andriy Boychenko<sup>1</sup>; Antonia Winkler<sup>2</sup>; Axel Naumann<sup>1</sup>; Dane Tacchini<sup>1</sup>; Giacomo Tenaglia<sup>1</sup>; Hamza Boukabache<sup>1</sup>; Han Hubert Dols<sup>1</sup>; Iris Kyranou<sup>None</sup>; Javier Serrano<sup>1</sup>; Joao Antunes Pequenao<sup>1</sup>; Matthias Bonora<sup>1</sup>; Philip Elson<sup>1</sup>; Sunje Dallmeier-Tiessen<sup>1</sup>; Victoria Stephany Huisman Sigcha<sup>3</sup>

<sup>1</sup> CERN

<sup>2</sup> Humboldt University of Berlin (DE)

<sup>3</sup> Saxion University of Applied Scienc (NL)

**Corresponding Authors:** andriy.boychenko@cern.ch, antonia.winkler@cern.ch, dane.tacchini@cern.ch, alexander.yohei.huss@cern.ch, matthias.bonora@cern.ch, axel.naumann@cern.ch, javier.serrano@cern.ch, iris.kyranou@cern.ch, hamza.boukabache@cern.ch, joao.pequenao@cern.ch, v.h.s@cern.ch, philip.elson@cern.ch, han.dols@cern.ch, sunje.dallmeier-tiessen@cern.ch, giacomo.tenaglia@cern.ch

High Energy (Nuclear) Physics and Open Source are a perfect match with a long history. CERN has created an Open Source Program Office (CERN OSPO [1]) to help open-source hardware and software in the CERN community - for CERN staff and the experiments' users. In the wider context, open source and CERN's OSPO have key roles in CERN's Open Science Policy [2]. With the OSPO, open-source projects should have more visibility inside and outside the organization, as contributions to society; the OSPO's team of practitioners want to make open source at CERN an easier, obvious task.

This presentation will provide you with an overview of the mission and objectives of the CERN Open Source Program Office (OSPO). This contribution exposes how the OSPO can and needs to help, what the OSPO wants to achieve, and what an OSPO's role might be in the HE(N)P software ecosystem. After more than a year in active engagement, we will share insights encountered so far, including the different challenges of open source in different parts of CERN. The presentation will share some behind-the-scenes stories: what the challenges were in creating it, what makes it special compared to other OSPOs, and why the OSPO won't do some things you might expect it to do. We will present the initial set of technical recommendations ("best practices") as proposed by the CERN OSPO; some alignment across institutions might be beneficial for the global HE(N)P community.

By sharing the CERN OSPO's journey, challenges, and lessons learned, we hope to provide valuable insights relevant to other HE(N)P centers, open-source projects, and the wider open source community.

[1] <https://opensource.cern/mandate>

[2] <https://openscience.cern/policies>

## Poster session / 8

### Time-Of-Flight Correction for Hadrons at BESIII experiment

**Authors:** Xinnan Wang<sup>1</sup>; Shengsen Sun<sup>1</sup>; Huaimin Liu<sup>1</sup>

<sup>1</sup> Institute of High Energy Physics Chinese Academy of Sciences

**Corresponding Authors:** xnwang@ihep.ac.cn, liuhm@ihep.ac.cn, sunss@mail.ihep.ac.cn

Beijing Spectrometer (BESIII) detector is used for high-precision studies of hadron physics and tau-charm physics. Accurate and reliable particle identification (PID) is crucial to improve the signal-to-noise ratio, especially for K/ $\pi$  separation. The time-of-flight (TOF) system, which is based on plastic scintillators, is a powerful tool for particle identification at BESIII experiment. The measured time is obtained using an empirical formula, which is used for time walk and hit position corrections, with Bhabha events used as calibration samples. Time difference is defined as the difference between the measured time and the expected time. Systematic time deviations of charged hadrons have been observed in the time differences for different particle species. This kind of systematic time deviation, which depends on the momentum and particle species, has been reported in several experiments using TOF based on plastic scintillation counters. Similar behaviors have also been observed in simulations with different deviations. In this study, the dependence of time deviations on pulse heights and hit positions is systematically investigated using different species of hadron control

samples. By applying corrections to the measured time, the time deviations are substantially reduced to nearly zero. The PID efficiencies of hadrons are enhanced both for real data and MC samples, and the systematic uncertainties of PID efficiencies are also optimized with further tuning. This study offers a new perspective on investigating time deviation in scintillation TOF detectors and provides a reference for improving detection accuracy.

Poster session / 9

## Development of Auto-Validation System of BOSS

**Author:** Di Jiang<sup>1</sup>

**Co-authors:** Qiumei MA maqm ; Xiaobin Ji <sup>2</sup>; Yao Zhang ; Ye Yuan <sup>3</sup>

<sup>1</sup> *Institute of High Energy Physics, Chinese Academy of Sciences*

<sup>2</sup> *IHEP, CAS*

<sup>3</sup> *Institute of High Energy Physics, Beijing*

**Corresponding Authors:** yuany@ihep.ac.cn, jiangdi@ihep.ac.cn, maqm@ihep.ac.cn, zhangyao@ihep.ac.cn, jixb@ihep.ac.cn

A modern version control system is capable of performing Continuous Integration (CI) and Continuous Deployment (CD) in a safe and reliable manner. Many experiments and software projects of High Energy Physics are now developing based on such modern development tools, GitHub for example. However, refactoring a large-scale running system can be challenging and difficult to execute. This is the reason why the BES Offline Software System (BOSS) continues to be developed using an outdated version control system, specifically, Concurrent Versions System (CVS). CVS does not automatically check the committed code during the commit process. To address this issue, a new auto-validation system has been developed, which overrides parts of the 'cvs' subcommand, enabling automatic code checks immediately after committing. Besides, with the integration of Gitlab, it includes functions designed for the convenience of developers and system managers, allowing them to work on multiple tasks simultaneously and automatically collects validated code. This approach strikes a balance between stability and innovation, allowing developers and system managers to enjoy the benefits of a modern-like version control system without having to much alter their work habits. The system is currently in use for the development and maintenance of BOSS.

Parallel (Track 2) / 10

## Charged Particle Track Reconstruction in CLAS12 using Artificial Intelligence

**Author:** Gagik Gavalian<sup>1</sup>

<sup>1</sup> *Jefferson Lab*

**Corresponding Author:** gagik.gavalian@gmail.com

The increasing complexity and data volume of Nuclear Physics experiments require significant computing resources to process data from experimental setups. The entire experimental data set has to be processed to extract sub-samples for physics analysis. The advancements in Artificial Intelligence and Machine Learning fields provide tools and procedures that can significantly enhance the throughput of data processing and significantly reduce the computational resources needed to process and categorize the experimental data in the raw data stream. In CLAS12 machine learning methods are developed to perform track reconstruction in real-time, allowing the identification of physics reactions from the raw data stream with the rates exceeding the data acquisition rates. In this paper, we present the Neural Network-driven track reconstruction that allows event classification and physics analysis in real time. We present a complete physics analysis of the data processed in the online.

**Poster session / 11****Online Electron Reconstruction at CLAS12****Author:** Richard Tyson<sup>1</sup><sup>1</sup> *Thomas Jefferson National Accelerator Facility (JLab)***Corresponding Author:** tyson@jlab.org

Online reconstruction is key for monitoring purposes and real time analysis in High Energy and Nuclear Physics (HEP) experiments. A necessary component of reconstruction algorithms is particle identification (PID) that combines information left by a particle passing through several detector components to identify the particle's type. Of particular interest to electro-production Nuclear Physics experiments such as CLAS12 is electron identification which is used to trigger data recording. A machine-learning approach was developed for CLAS12 to reconstruct and identify electrons by combining raw signals at the data acquisition level from several detector components. This approach achieves a high electron identification purity whilst retaining a 99.95% efficiency. The machine learning tools are capable of running at high rates exceeding the data acquisition rates and will allow electron reconstruction in real-time. This framework can then be expanded to other particle types. This work enhances online analyses and monitoring at CLAS12. Improved electron identification in the trigger also contributes to the reduction in recorded data volumes and improves data processing times. This approach to triggering will be employed when transitioning to higher luminosity experiments at CLAS12 where the data volume will increase significantly.

**Parallel (Track 2) / 12****The automated Bandwidth Division for the LHCb first-level trigger****Author:** Joshua Ethan Horswill<sup>1</sup>**Co-authors:** Conor Fitzpatrick<sup>1</sup>; Timothy David Evans<sup>1</sup><sup>1</sup> *University of Manchester (GB)***Corresponding Authors:** timothy.david.evans@cern.ch, conor.fitzpatrick@cern.ch, joshua.ethan.horswill@cern.ch

The first level of the trigger system of the LHCb experiment (HLT1) reconstructs and selects events in real-time at the LHC bunch crossing rate in software using GPUs. It must carefully balance a broad physics programme that extends from kaon physics up to the electroweak scale. An automated procedure to determine selection criteria is adopted that maximises the physics output of the entirety of this programme while satisfying constraints from the higher-level components of the trigger system, which cap the output rate of HLT1 to around 1MHz. In this talk, the method by which this optimisation is achieved will be described in detail, which uses a variant of the ADAM algorithm popular in machine learning tools, customised in order to solve discrete minimisation problems. The impact of this optimisation on the first data taken by the LHCb experiment in its nominal Run 3 configuration will also be shown.

**Parallel (Track 8) / 14****Data Preservation in High Energy Physics: a 10 years perspective****Author:** Cristinel Diaconu<sup>1</sup>

<sup>1</sup> CPPM, Aix-Marseille Université, CNRS/IN2P3 (FR)

**Corresponding Author:** diaconu@cppm.in2p3.fr

Data Preservation (DP) is a mandatory specification for any present and future experimental facility and it is a cost-effective way of doing fundamental research by exploiting unique data sets in the light of the ever increasing theoretical understanding. When properly taken into account, DP leads to a significant increase in the scientific output (10% typically) for a minimal investment overhead (0.1%). DP relies on and stimulates cutting-edge technology developments and is strongly linked to Open Science and FAIR data paradigms. A recently released report (Eur.Phys.J.C 83 (2023) 9, 795 | 2302.03583 [hep-ex] ) summarizes the status of data preservation in high energy physics from a perspective of more than ten years of experience with a structured effort at international level (DPHEP).

**Poster session / 15**

## HPSS Batch Application

**Author:** Justin Spradley<sup>None</sup>

**Corresponding Author:** jspradley@bnl.gov

How to effectively and efficiently stage a large number of requests from an IBM HPSS environment using a MariaDB database to keep track of requests and use Python for all business logic and to consume the HPSS API. The goal is to be able to scale to handle a large number of requests and to meet different needs of different experiments, and to make the program adaptable enough to allow for each experiment to have its own unique business logic. This update will take advantage of features of the newest versions of HPSS, as well as MariaDB, Python, and Linux. Furthermore, the hope is that the application will be able to log and handle a wider array of errors and exceptions, and allow for more in depth monitoring as the status of each request will be stored in a database which allows for easy querying. Furthermore this may allow for additional enhancements such as staging requests by priority.

**Poster session / 16**

## Management of the data processing & Run-by-Run simulations in KM3NeT

**Author:** Anna Sinopoulou<sup>1</sup>

**Co-authors:** Benjamin Trocme ; Carla Distefano <sup>2</sup>; Carla Distefano <sup>3</sup>; Carla Distefano ; Chiara F Lastoria <sup>4</sup>; Luigi Antonio Fusco <sup>5</sup>; Luigi Fusco <sup>6</sup>; Mieke Bouwhuis <sup>7</sup>; Valentin Pestel <sup>8</sup>

<sup>1</sup> INFN - Sezione di Catania

<sup>2</sup> Laboratori Nazionali del Sud (IT)

<sup>3</sup> LNS-INFN

<sup>4</sup> LPC Caen, CNRS/IN2P3, Caen (France)

<sup>5</sup> University of Salerno

<sup>6</sup> CERN

<sup>7</sup> NIKHEF

<sup>8</sup> LPC Caen - CNRS - IN2P3

**Corresponding Authors:** carla.distefano@cern.ch, distefano\_c@lns.infn.it, sinopoulou@inp.demokritos.gr, bouwhuim@nikhef.nl, trocme@apc.in2p3.fr, vpestel@lpccaen.in2p3.fr, chiara.filomena.lastoria@cern.ch, luigi.fusco@cern.ch, lfusco@bo.infn.it, carla.distefano@lns.infn.it

The KM3NeT collaboration is constructing two underwater neutrino detectors in the Mediterranean Sea sharing the same technology: the ARCA and ORCA detectors. ARCA is optimized for the observation of astrophysical neutrinos, while ORCA is designed to determine the neutrino mass hierarchy by detecting atmospheric neutrinos. Data from the first deployed detection units are being analyzed and several physics analyses have already been presented. As the detector configurations are growing and therefore, the amount of the recorded data, efficient data quality and processing management are essential.

Data reconstruction and Monte Carlo simulations are handled separately for each data taking period (run), to achieve complete processing output and optimal computing performance. A Run-by-Run simulation procedure is followed, to reproduce the conditions, possible seawater environment variations as well as the acquisition setup for each run. To handle computing requirements such as portability, reproducibility and scalability, the collaboration implemented this approach using Snakemake, a trending workflow management system.

Poster session / 17

## Hyperparameter Optimization for Deep Learning Models Using High Performance Computing

**Authors:** Domenico Diacono<sup>1</sup>; Francesco Grancagnolo<sup>2</sup>; Guang Zhao<sup>3</sup>; Linghui Wu<sup>3</sup>; Marcello Abbrescia<sup>1</sup>; Mingyi Dong<sup>3</sup>; Muhammad Numan Anwar<sup>1</sup>; Nicola De Filippis<sup>1</sup>; Shengsen Sun<sup>3</sup>

<sup>1</sup> *Politecnico di Bari and Istituto Nazionale di Fisica Nucleare Bari, Italy*

<sup>2</sup> *Istituto Nazionale di Fisica Nucleare Lecce, Italy*

<sup>3</sup> *Institute of High Energy Physics, 19B Yuquan Road, Beijing, 100049, Beijing, China*

**Corresponding Authors:** muhammad.anwar@ba.infn.it, domenico.diacono@ba.infn.it, zhaog@ihep.ac.cn, nicola.de.filippis@cern.ch, franco.grancagnolo@le.infn.it, marcello.abbrescia@ba.infn.it

Clusters counting in a drift chamber represents the most promising breakthrough in particle identification (PID) techniques in particle physics experiments. In this paper, neural network models, such as the Long Short-Term Memory (LSTM) Model and Convolutional Neural Network (CNN) Model, are trained using various hyperparameters like loss functions, activation functions, different numbers of neurons, batch sizes, and varying numbers of epochs etc. These models are trained for a two-step reconstruction algorithm, which involves peak finding and clusterization. For the peak finding algorithm, a trained Long Short-Term Memory (LSTM) model is used to discriminate between ionization signals (primary and secondary peaks) and noise in the waveform, addressing a classification problem. Concurrently, a Convolutional Neural Network model is utilized to determine the number of primary ionization clusters based on the detected peaks, dealing with a regression problem. The trained models (LSTM and CNN) are applied to the simulation, based on Garfield++, of particles traversing a gas mixture made of 90% Helium (He) and 10% Isobutane (C<sub>4</sub>H<sub>10</sub>) filling drift tubes with the same geometry as the ones used for the beam test at CERN in 2023 of the prototype of the IDEA detector for FCC. The simulation parameters included a cell size of 1.5 cm, a sampling rate of 1.2 GHz, a time window of 2000 ns, 5000 events, a number of ionization clusters with a mean value of 25.25, a number of electrons per cluster with a mean value of 1.468, and momentum of pion particles ranging from 4 to 180 GeV/c.

Poster session / 18

## Improving multi-track reconstruction algorithms in the Mu2e Experiment

**Author:** Alessandro Maria Ricci<sup>1</sup>

<sup>1</sup> *INFN Pisa*

**Corresponding Author:** [alessandro.ricci@df.unipi.it](mailto:alessandro.ricci@df.unipi.it)

Mu2e will search for the neutrinoless coherent  $\mu^- \rightarrow e^-$  conversion in the field of an Al nucleus, a Charged Lepton Flavor Violation (CLFV) process. The experiment is expected to start in 2026 and will improve the current limit by 4 orders of magnitude.

Mu2e consists of a straw-tube tracker and crystal calorimeter in a 1T B field complemented by a plastic scintillation counter veto to suppress cosmic ray backgrounds. The tracker geometry, composed of 36 tracking planes equally spaced with straws transverse to the beamline, makes track reconstruction a quite unique problem.

The first step of track reconstruction is hit clustering, in space and time. Pattern recognition is performed for each time cluster to identify hit combinations compatible with a 3D helix and remove background hits. Track fitting acts on the hit combinations to determine precision track parameters. The existing algorithms are robust and efficient for the topologies of interest for the principal physics analyses. However, we developed pattern recognition algorithms to improve track reconstruction of multi-particle events which are important for background estimates through data-driven procedures.

**Parallel (Track 8) / 19**

## **The Italian Summer Students Program at Fermilab and other US Laboratories: 40 years of education in particle physics and technology**

**Authors:** Emanuela Barzi<sup>None</sup>; Giorgio Bellettini<sup>1</sup>; Marco Mambelli<sup>2</sup>; Simone Donati<sup>3</sup>

<sup>1</sup> *Dipartimento di Fisica*

<sup>2</sup> *Fermilab (US)*

<sup>3</sup> *Istituto Nazionale di Fisica Nucleare - Pisa*

**Corresponding Authors:** [giorgio.bellettini@pi.infn.it](mailto:giorgio.bellettini@pi.infn.it), [simone.donati@pi.infn.it](mailto:simone.donati@pi.infn.it), [marcom@fnal.gov](mailto:marcom@fnal.gov), [barzi@fnal.gov](mailto:barzi@fnal.gov)

Since 1983 the Italian groups collaborating with Fermilab (US) have been running a 2-month summer training program for Master students. While in the first year the program involved only 4 physics students, in the following years it was extended to engineering students. Many students have extended their collaboration with Fermilab with their Master Thesis and PhD.

The program has involved almost 600 Italian students from more than 20 Italian universities. Each intern is supervised by a Fermilab Mentor responsible for the training program. Training programs spanned from Tevatron, CMS, Muon (g-2), Mu2e and SBN and DUNE design and data analysis, development of particle detectors, design of electronic and accelerator components, development of infrastructures and software for tera-data handling, quantum computing and research on superconductive elements and accelerating cavities.

In 2015 the University of Pisa included the program within its own educational programs. Summer Students are enrolled at the University of Pisa for the duration of the internship and at the end of the internship they write summary reports on their achievements. After positive evaluation by a University of Pisa Examining Board, interns are acknowledged 6 ECTS credits for their Diploma Supplement. In the years 2020 and 2021 the program was canceled due to the sanitary emergency but in 2022 it was restarted and allowed a cohort of 21 students in 2022, and a cohort of 27 students in 2023 to be trained for nine weeks at Fermilab. We are now organizing the 2024 program.

**Poster session / 20**

## **Streamlining ATLAS Monte-Carlo Generator Validation with PAVER**

**Author:** Mustafa Andre  Schmidt<sup>1</sup>



<sup>1</sup> *Bergische Universitaet Wuppertal (DE)*

**Corresponding Author:** mustafa.schmidt@cern.ch

In ATLAS and other high-energy physics experiments, the integrity of Monte-Carlo (MC) simulations is crucial for reliable physics analysis. The continuous evolution of MC generators necessitates regular validation to ensure the accuracy of simulations. We introduce an enhanced validation framework incorporating the Job Execution Monitor (JEM) resulting in the established Physics Modeling Group (PMG) Architecture for Validating Evgen with Rivet (PAVER). This setup automates the validation process, facilitating systematic evaluation of MC generator updates and their compliance with experimental data.

This approach allows for early detection of discrepancies in simulation outputs, ensuring that potential issues and bugs are addressed before the production of large-scale samples for the ATLAS collaboration. MC generator Validation is specially important to save energy and money and to reduce the carbon footprint in future simulation campaigns significantly which aligns very well with the importance of reaching sustainability within ATLAS. The result is a streamlined, robust, and accessible validation system that supports sustainable MC production in ATLAS.

This presentation will summarize the implementation of PAVER by highlighting its impact on enhancing simulation reliability and efficiency. It will furthermore include an overview of the massive validation program throughout the past years resulting in many successfully validated generator and software updates. In addition, this talk will present insights into the challenges and solutions in MC generator validation, with implications for future developments in high-energy physics simulations.

**Poster session / 21**

## Normalizing Flows for Physics Data Analyses

**Author:** Jan Gavranovic<sup>1</sup>

**Co-author:** Borut Paul Kersevan<sup>1</sup>

<sup>1</sup> *Jozef Stefan Institute (SI)*

**Corresponding Authors:** jan.gavranovic@cern.ch, borut.kersevan@cern.ch

Monte Carlo (MC) simulations are a crucial component when analysing the Standard Model and New physics processes at the Large Hadron Collider. The goal of this work is to explore the performance of generative models for complementing the statistics of classical MC simulations in the final stage of data analysis by generating additional synthetic data that follows the same kinematic distributions for a limited set of analysis-specific observables to a high precision. A normalizing flow architecture was adapted for this task and its performance was systematically evaluated using a well-known benchmark sample containing the Higgs boson production beyond the Standard Model and the corresponding irreducible background. The applicability of normalizing flows under different model parameters and a restricted number of initial events used in training was investigated. The best performing model was then chosen for further evaluation with a set of statistical procedures and a simplified physics analysis. We demonstrate that the the number of events used in training coupled with the flow architecture are crucial for the physics performance of the generative model. By implementing and performing a series of statistical tests and evaluations we show that a machine-learning-based generative procedure can be used to generate synthetic data that matches the original samples closely enough and that it can therefore be incorporated in the final stage of a physics analysis with some given systematic uncertainty.

**Parallel (Track 1) / 22**

## dCache project status & update

**Authors:** Christopher Green<sup>1</sup>; Dmitry Litvintsev<sup>1</sup>; Krishnaveni Chitrapu<sup>2</sup>; Lea Morschel<sup>None</sup>; Marina Sahakyan<sup>None</sup>; Paul Millar<sup>None</sup>; Svenja Meyer<sup>None</sup>; Tigran Mkrtchyan<sup>3</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

<sup>2</sup> *National Supercomputer Centre in Sweden*

<sup>3</sup> *DESY*

**Corresponding Authors:** krishnaveni@nsc.liu.se, greenc@fnal.gov, paul.millar@desy.de, lea.morschel@desy.de, marina.sahakyan@desy.de, litvinse@fnal.gov, svenja.meyer@desy.de, tigran.mkrtchyan@desy.de

The dCache project provides open-source software deployed internationally to satisfy ever-more demanding storage requirements. Its multifaceted approach provides an integrated way of supporting different use-cases with the same storage, from high throughput data ingest, data sharing over wide area networks, efficient access from HPC clusters, and long term data persistence on tertiary storage. Although dCache was originally developed for HEP experiments, today it is used by various scientific communities, including astrophysics, biomed, and life science, each with its specific requirements. To match the requirements of these new communities and keep up with the scaling demands of existing experiments, dCache evolution is a permanent ongoing process. With this contribution, we would like to highlight the recent developments in dCache regarding integration with CERN Tape Archive (CTA), advanced metadata handling, token-based authorization support, bulk API for QoS transitions, RESTAPI to control interaction with the tape system, and future development directions.

**Poster session / 23**

## dE/dx Software in the BESIII Experiment

**Author:** Wenxing Fang<sup>None</sup>

**Co-authors:** Shengsen Sun<sup>1</sup>; Weidong Li<sup>2</sup>

<sup>1</sup> *Institute of High Energy Physics Chinese Academy of Sciences*

<sup>2</sup> *IHEP, Beijing*

**Corresponding Authors:** liwd@ihep.ac.cn, fangwx@ihep.ac.cn, sunss@mail.ihep.ac.cn

The BESIII experiment operates as an electron-positron collider in the tau-charm energy region, pursuing a range of physics goals related to charm, charmonium, light hadron decays, and so on. Among these objectives, achieving accurate particle identification (PID) plays a crucial role, ensuring both high efficiency and low systematic uncertainty. In the BESIII experiment, PID performance heavily relies on two key measurements: the energy deposit per unit length (dE/dx) obtained from the main drift chamber (MDC) sub-detector, and the time of flight (TOF) measurement from the TOF sub-detector.

This contribution focuses specifically on the dE/dx aspect and provides a comprehensive overview of the dE/dx software employed in the BESIII experiment. The presentation encompasses simulation, calibration, and reconstruction techniques implemented in the analysis pipeline. Last but not least, with the help of machine learning (ML) technique, a study of ML-based dE/dx simulation will also be presented.

**Parallel (Track6) / 24**

## dCache CI/CD migration to Kubernetes

**Authors:** Christopher Green<sup>1</sup>; Dmitry Litvintsev<sup>1</sup>; Elena Gapon<sup>2</sup>; Lea Morschel<sup>None</sup>; Marina Sahakyan<sup>None</sup>; Markus Wengert<sup>3</sup>; Michael Schuh<sup>None</sup>; Stefan Bujack<sup>None</sup>; Svenja Meyer<sup>None</sup>; Tigran Mkrtchyan<sup>3</sup>; Tobias Klann<sup>None</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

<sup>2</sup> *Deutsches Elektronen-Synchrotron DESY*

<sup>3</sup> *DESY*

**Corresponding Authors:** tobias.klann@desy.de, greenc@fnal.gov, michael.schuh@desy.de, elena.gapon@desy.de, litvinse@fnal.gov, svenja.meyer@desy.de, tigran.mkrtchyan@desy.de, marina.sahakyan@desy.de, stefan.bujack@desy.de, lea.morschel@desy.de, markus.wengert@desy.de

For over two decades, the dCache project has provided open-source to satisfy ever-more demanding storage requirements. More than 80 sites around the world, rely on dCache to provide services for LHC experiments, Belle-II, EuXFEL and many others. This can be achieved only with a well-established process from a whiteboard, where ideas are created, through development, packaging and testing. The project's build and test infrastructure is based on Jenkins CI and a set of virtual machines. This infrastructure is maintained by dCache developers. With the introduction of the DESY-central Gitlab server, the developers have started migrating from VM-based testing to container-based deployments in the onsite Kubernetes cluster. As a result, we have packaged dCache containers and Helm charts that can be used by other sites to reproduce our test and build steps quickly or to evaluate new releases on their pre-production systems, and, eventually, become a standard model of dCache deployment at the sites.

This presentation will show challenges that we have faced, the techniques how they were solved and issues that still need to be addressed.

Poster session / 25

## The evolution of INFN's Cloud Platform: improvements in Orchestration and User Experience

**Author:** Marica Antonacci<sup>1</sup>

**Co-authors:** Barbara Martelli ; Claudio Grandi <sup>2</sup>; Daniele Spiga ; Enrico Vianello ; Ettore Serra ; Giacinto Donvito <sup>3</sup>; Giovanni Savarese ; Luca Giommi <sup>4</sup>

<sup>1</sup> *INFN*

<sup>2</sup> *INFN - Bologna*

<sup>3</sup> *INFN-Bari*

<sup>4</sup> *Universita e INFN, Bologna (IT)*

**Corresponding Authors:** barbara.martelli@cnaif.infn.it, luca.giommi@cern.ch, giovanni.savarese@ba.infn.it, marica.antonacci@ba.infn.it, enrico.vianello@cnaif.infn.it, etttore.serra@ba.infn.it, danielle.spiga@pg.infn.it, giacinto.donvito@ba.infn.it, claudio.grandi@cern.ch

The Italian National Institute for Nuclear Physics (INFN) has recently developed a national cloud platform to enhance access to distributed computing and storage resources for scientific researchers. A critical aspect of this initiative is the INFN Cloud Dashboard, a user-friendly web portal that allows users to request high-level services on demand, such as Jupyter Hub, Kubernetes, and Spark clusters. The platform is based on INDIGO-PaaS middleware, which integrates a TOSCA-based orchestration system. This system supports a lightweight federation of cloud sites and automates resource scheduling for optimal resource allocation.

Through the internal INFN DataCloud project and European initiatives like interTwin, INFN is undertaking a comprehensive overhaul of its PaaS system to adapt to evolving technologies and replace outdated software components. To further improve the orchestration system, INFN is exploring the use of artificial intelligence to enhance deployment scheduling.

Additionally, the dashboard, serving as a user interface for orchestrating and deploying services, has recently undergone significant renovations to boost usability and security. This contribution aims to

highlight key advancements in the PaaS orchestration system designed to offer a reliable, scalable, and user-friendly environment for the computational needs of the scientific community.

Parallel (Track 1) / 26

## Use of Intelligent Compression in Waveform like Data of Neutrino Oscillation Experiments

**Author:** Amit Bashyal<sup>1</sup>

**Co-authors:** Barnali Chowdhury<sup>1</sup>; Peter van Gemmeren<sup>1</sup>; Sheng Di<sup>1</sup>; Zelimir Djurcic<sup>1</sup>

<sup>1</sup> Argonne National Laboratory

**Corresponding Authors:** zdjurcic@anl.gov, sdi1@anl.gov, bchowdhury@anl.gov, abashyal@anl.gov, peter.van.gemmeren@cern.ch

Many High Energy Physics experiments store their data in compressed format using lossless algorithms. The use of lossy compression algorithms is less common in the HEP. However, as the storage requirements of the HEP experiments will grow during the HL-LHC and DUNE era, intelligent compression can enable significant storage optimization. The ability to highly compress data while preserving enough fidelity for subsequent physics analysis and interpretation could allow HEP experiments to limit storage needs so compressed objects can be part of down-stream/derived data products. This could allow experiments to use these objects when processing the derived data without reading the original raw data which is often much larger and harder to access, specially for experiments like DUNE.

In this talk, we will present our exploratory work of a framework that uses intelligent lossy compression algorithms like SZ3, MGARD and IDEAL-EM to compress and test the fidelity of the DUNE-like waveform data which are simple but very large due to the large far detector. Our work builds on the studies done by HEP-CCE Storage Optimization and will identify the potential of compression algorithms which are developed and widely used outside of the HEP field and their possible integration within the HEP community.

Poster session / 27

## Common Analysis Tools in CMS

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

The CMS experiment has recently established a new Common Analysis Tools (CAT) group. The CAT group implements a forum for the discussion, dissemination, organization and development of analysis tools, broadly bridging the gap between the CMS data and simulation datasets and the publication-grade plots and results. In this talk we discuss some of the recent developments carried out in the group, including the structure of the group, the facilities and services provided, the communication channels, the ongoing developments in the context of frameworks for data processing, strategies for the management of analysis workflows and their preservation and tools for the statistical interpretation of analysis results.

Parallel (Track 5) / 29

## Transfer learning for Smart Background Simulation at Belle II

**Author:** Nikolai Hartmann<sup>1</sup>

**Co-authors:** Boyang Yu ; Thomas Kuhr<sup>1</sup>

<sup>1</sup> *Ludwig Maximilians Universitat (DE)*

**Corresponding Authors:** thomas.kuhr@cern.ch, boyang.yu@physik.uni-muenchen.de, nikolai.marcel.hartmann@cern.ch

Accurate modeling of backgrounds for the development of analyses requires large enough simulated samples of background data. When searching for rare processes, a large fraction of these expensively produced samples is discarded by the analysis criteria that try to isolate the rare events. At the Belle II experiment, the event generation stage takes only a small fraction of the computational cost of the whole simulation chain, motivating filters for the simulation at this stage. Deep neural network architectures based on graph neural networks have been proven useful to predict approximately which events will be kept after the filter, even in cases where there is no simple correlation between generator and reconstruction level quantities. However, training these models requires large training data sets, which are hard to obtain for filters with very low efficiencies. In this presentation we show how a generic model, pre-trained on filters with high efficiencies can be fine-tuned to also predict filters where only little training data is available. This also opens opportunities for online learning during the simulation process where no separate training step is required.

**Parallel (Track6) / 30**

## ATLAS ITk Production Database use and tools

**Author:** Kenneth Gibb Wraight<sup>1</sup>

**Co-author:** Monika Wielers<sup>2</sup>

<sup>1</sup> *University of Glasgow (GB)*

<sup>2</sup> *RAL (UK)*

**Corresponding Authors:** monika.wielers@cern.ch, kenneth.gibb.wraight@cern.ch

The ATLAS experiment will undergo major upgrades for operation at the high luminosity LHC. The high pile-up interaction environment (up to 200 interactions per 40MHz bunch crossing) requires a new radiation-hard tracking detector with a fast readout.

The scale of the proposed Inner Tracker (ITk) upgrade is much larger than the current ATLAS tracker. The current tracker consists of ~4000 modules while ITk will be made of ~28,000 modules. To ensure a good production quality, all the items to build modules as well as bigger structures on which they will be placed need to be tracked along with the relevant quality control and quality assurance information. Hence, the ITk production database (PDB) is vital to follow the complex production flow for each item and institutes around the globe. The database also allows close monitoring of the production quality and production speed. After production the information will be stored for 10 years of data-taking to trace potential operational issues to specific production items.

A PDB API allows development of tools for database interaction by different user types: technicians, academics, engineers and vendors. Several options have been pursued to meet the needs by the collaboration: pythonic API wrapper, data-acquisition GUIs with integrated scripts, commandline scripts distributed via git repositories, containerised applications, and CERN hosted resources.

This presentation promotes information exchange and collaboration for tools which supports detector construction in a large-scale experiment. Examples of front-end development and reporting will be shown. Through these examples, the general themes of large-scale data management and multi-user global accessibility will be discussed. These concepts are relevant not only for modern high-energy particle physics (HEP) but also for large experiments beyond HEP.

**Poster session / 31****Anomaly Detection on BESIII EMC using Machine Learning****Author:** Mingrun Li<sup>1</sup>**Co-authors:** Chunxiu Liu<sup>1</sup>; Xiaobin Ji<sup>1</sup><sup>1</sup> *IHEP, CAS***Corresponding Authors:** liucx@ihep.ac.cn, jixb@ihep.ac.cn, mrli@ihep.ac.cn

The BESIII at the BEPCII electron-positron accelerator, which is located at IHEP, Beijing, China, is an experiment for the studies of hadron physics and  $\tau$ -charm physics with the highest accuracy achieved until now. It has collected several world's largest  $e^+e^-$  samples in  $\tau$ -charm region. Anomaly detection on BESIII detectors is an important segment of improving data quality, enhancing data acquisition efficiency and monitoring detectors' status. An offline unsupervised autoencoder-based anomaly detection method is applied on CsI(Tl) electromagnetic calorimeter (EMC). This method checks over histograms generated by each crystal using Jensen-Shannon Distance as loss function. Comparing to traditional method, this method is able to provide more accurate anomaly information with less manpower consuming.

**Parallel (Track 2) / 32****Versal ACAP processing for ATLAS-TileCal signal reconstruction****Author:** Francisco Hervas Alvarez<sup>1</sup>**Co-authors:** Alberto Valero Biot<sup>1</sup>; Luca Fiorini<sup>1</sup><sup>1</sup> *Univ. of Valencia and CSIC (ES)***Corresponding Authors:** alberto.valero@cern.ch, fran.hervas@ific.uv.es, luca.fiorini@cern.ch

Particle detectors at accelerators generate large amount of data, requiring analysis to derive insights. Collisions lead to signal pile up, where multiple particles produce signals in the same detector sensors, complicating individual signal identification. This contribution describes the implementation of a deep learning algorithm on a Versal ACAP device for improved processing via parallelization and concurrency. Connected to a host computer via PCIe, this system aims for enhanced speed and energy efficiency over CPUs and GPUs. In the contribution, we will describe in detail the data processing and the hardware, firmware and software components of the signal reconstruction of the Tile Calorimeter (TileCal) of the ATLAS detector which will be running in real time in the HL-LHC era. The contribution presents the implementation of the deep learning algorithm on Versal ACAP device, as well as the system for transferring data in an efficient way. In addition, the system integration tests and results from the tests with beam performed at CERN will be presented.

**Poster session / 33****Porting MADGRAPH to FPGA using High-Level Synthesis (HLS)****Author:** Hector Gutierrez Arance<sup>1</sup>**Co-authors:** Alberto Valero Biot<sup>1</sup>; Luca Fiorini<sup>1</sup><sup>1</sup> *Univ. of Valencia and CSIC (ES)*

**Corresponding Authors:** hector.gutierrez@ific.uv.es, alberto.valero@cern.ch, luca.fiorini@cern.ch

The escalating demand for data processing in particle physics research has spurred the exploration of novel technologies to enhance efficiency and speed of calculations. This study presents the development of a porting of MADGRAPH, a widely used tool in particle collision simulations, to FPGA using High-Level Synthesis (HLS).

Experimental evaluation is ongoing, but preliminary assessments suggest a promising enhancement in calculation speed compared to traditional CPU implementations. This potential improvement could enable the execution of more complex simulations within shorter time frames.

This study describes the complex process of adapting MADGRAPH to FPGA using HLS, focusing on optimizing algorithms for parallel processing. A key aspect of the FPGA implementation of the MADGRAPH software is reduction of the power consumption, which important implications for the scalability of computer centers and for the environment. These advancements could enable faster execution of complex simulations, highlighting FPGA's crucial role in advancing particle physics research and its environmental impact.

Parallel (Track 4) / 34

## WLCG transition from X.509 to Tokens: Progress and Outlook

**Authors:** Maarten Litmaath<sup>1</sup>; Tom Dack<sup>None</sup>

**Co-author:** WLCG Authorization Working Group <sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> WLCG

**Corresponding Authors:** tom.dack@cern.ch, maarten.litmaath@cern.ch

Since 2017, the Worldwide LHC Computing Grid (WLCG) has been working towards enabling token-based authentication and authorization throughout its entire middleware stack.

Taking guidance from the WLCG Token Transition Timeline, published in 2022, substantial progress has been achieved not only in making middleware compatible with the use of tokens, but also in understanding the limitations of the WLCG Common JWT Profiles, first published in 2019. Significant scalability experience has been gained from Data Challenge 2024, during which millions of files were transferred with only tokens used as credentials.

Besides describing the state of affairs in the transition to tokens, revisions to the WLCG token profile, and the evolving roadmaps, this contribution also covers the corresponding transition from VOMS-Admin to INDIGO-IAM services, with continuing improvements in terms of functionality as well as deployment.

Parallel (Track 1) / 35

## Data Challenge 2024 - CMS activities

**Authors:** Christoph Wissing<sup>1</sup>; Katy Ellis<sup>2</sup>

<sup>1</sup> Deutsches Elektronen-Synchrotron (DE)

<sup>2</sup> Science and Technology Facilities Council STFC (GB)

**Corresponding Authors:** katy.ellis@stfc.ac.uk, christoph.wissing@desy.de

To verify the readiness of the data distribution infrastructure for the HL-LHC, which is planned to start in 2029, WLCG is organizing a series of data challenges with increasing throughput and complexity. This presentation addresses the contribution of CMS to Data Challenge 2024, which

aims to reach 25% of the expected network throughput of the HL-LHC. During the challenge CMS tested various network flows, from the RAW data distribution to the “flexible” model, which adds network traffic resulting from data reprocessing and MC production between most CMS sites.

The overall throughput targets were met on the global scale utilizing several hundred links. Valuable information was gathered regarding scaling capabilities of key central services such as Rucio and FTS. During the challenge about half of the transferred volume was carried out via token based authentication. In general sufficient performance of individual links was observed and sites coped with the target throughput. For links that did not reach the target, attempts were made to identify the bottleneck, whether in the transfer tools, the network link, the involved storage systems or any other component.

Poster session / 36

## Model Building with Non-Parametric and Parametric Components for Partial Wave Analysis

Author: Lawrence Ng<sup>None</sup>

Corresponding Author: lng@jlab.org

NIFTy[1], a probabilistic programming framework developed for astrophysics, has recently been adapted to be used in partial wave analyses (PWA) at the COMPASS [2] experiment located in CERN. A non-parametric model, described as a correlated field, is used to characterize kinematically-smooth complex-binned amplitudes. Parametric models, like a Breit-Wigner distribution, can also be mixed in. This method is being investigated for use in the GlueX experiment located in Jefferson Lab.

I will introduce iftpwa[3], a flexible model-building framework that can construct and interfere both parametric and non-parametric amplitudes. A single configuration file is used to build a model and describe the optimization procedure resulting in a variationally approximated posterior distribution. This framework is designed to be modular which provides an avenue for inter-collaboration use and development.

### References

- [1] G. Edenhofer, P. Frank, J. Roth, R. H. Leike, M. Guerdi, L. I. Scheel-Platz, M. Guardiani, V. Eberle, M. Westerkamp, and T. A. Enßlin. Re-Envisioning Numerical Information Field Theory (NIFTy.re): A Library for Gaussian Processes and Variational Inference, 2024.
- [2] F. M. Kaspar, J. Beckers, and J. Knollmüller. Progress in the Partial-Wave Analysis Methods at COMPASS. EPJ Web Conf., 291:02014, 2024.
- [3] F. M. Kasper and L. Ng. iftpwa, 2024. Github Repository.

Parallel (Track 4) / 38

## Advanced monitoring capabilities of the CMS Experiment for LHC Run3 and beyond

Authors: Brij Kishor Jashal<sup>1</sup>; CMS Collaboration<sup>None</sup>; Federica Legger<sup>2</sup>

<sup>1</sup> Rutherford Appleton Laboratory

<sup>2</sup> Università e INFN Torino (IT)

Corresponding Author: borislav.pavlov@cern.ch



The CMS computing infrastructure spread globally over 150 WLCG sites forms a intricate ecosystem of computing resources, software and services. In 2024, the production computing cores breached half a million mark and storage capacity is at 250 PetaBytes on disk and 1.20 ExaBytes on Tape. To monitor these resources in real time, CMS working closely with CERN IT has developed a multi-faceted monitoring system providing real time insights using about 100 production dashboards.

In preparation of Run3, the CMS monitoring infrastructure underwent significant evolution to broaden the scope of monitored applications and services while enhancing sustainability and ease of operation. Leveraging open-source solutions, provided either by the CERN IT department or managed internally, monitoring applications have transitioned from bespoke solutions to standardized data flow and visualization services. Notably, monitoring applications for distributed workload management and data handling have migrated to utilize technologies like OpenSearch, VictoriaMetrics, InfluxDB, and HDFS, with access facilitated through programmatic APIs, Apache Spark, or Sqoop jobs, and visualization primarily via Grafana.

The majority of CMS monitoring applications are now deployed on Kubernetes clusters based microservices architecture. This contribution unveils the comprehensive stack of CMS monitoring services, showcasing how the integration of common technologies enables versatile monitoring applications and addresses the computation demands of LHC Run 3. Additionally, it explores the incorporation of analytics into the monitoring framework, demonstrating how these insights contribute to the operational efficiency and scientific output of the CMS experiment.

Parallel (Track 3) / 39

## Cluster counting reconstruction with classical supervised learning and transfer learning

**Authors:** Guang Zhao<sup>1</sup>; Zhefei Tian<sup>2</sup>

**Co-authors:** Francesco Grancagnolo<sup>3</sup>; Linghui Wu<sup>4</sup>; Mingyi Dong<sup>4</sup>; Nicola De Filippis<sup>5</sup>; Shengsen Sun<sup>4</sup>; Zhenyu Zhang<sup>2</sup>

<sup>1</sup> *Institute of High Energy Physics (CAS)*

<sup>2</sup> *Wuhan University*

<sup>3</sup> *INFN - Lecce*

<sup>4</sup> *Institute of High Energy Physics, 19B Yuquan Road, Beijing, 100049, Beijing, China*

<sup>5</sup> *Politecnico/INFN Bari (IT)*

**Corresponding Authors:** nicola.de.filippis@cern.ch, zhaog@ihep.ac.cn, tianzhefei@whu.edu.cn, zhenyuzhang@whu.edu.cn, franco.grancagnolo@le.infn.it

Particle identification (PID) is crucial in particle physics experiments. A promising breakthrough in PID involves cluster counting, which quantifies primary ionizations along a particle's trajectory in a drift chamber (DC), rather than relying on traditional dE/dx measurements. However, a significant challenge in cluster counting lies in developing an efficient reconstruction algorithm to recover cluster signals from DC cell waveforms.

In PID, machine learning algorithms have emerged as the state-of-the-art. For simulated samples, an updated supervised model based on LSTM and DGCNN achieves a remarkable 10% improvement in separating K from  $\pi$  compared to traditional methods. For test beam data samples collected at CERN, due to label scarcity and data/MC discrepancy, a semi-supervised domain adaptation model, which exploits Optimal Transport to transfer information between simulation and real data domains, is developed. The model is validated using pseudo data and further applied to real data. The performance is superior to the traditional methods and maintains consistent across varying track lengths.

There are two related papers that have been submitted to journals: 2402.16270 and 2402.16493. The previous one about the transfer learning has been accepted by the Computer Physics Communications (<https://doi.org/10.1016/j.cpc.2024.109208>).

**Poster session / 40**

## Docu-Bot: AI assisted user support

**Authors:** Jiri Chudoba<sup>1</sup>; Michal Chudoba<sup>2</sup>**Co-author:** Lukáš Hejtmanek<sup>3</sup><sup>1</sup> *Czech Academy of Sciences (CZ)*<sup>2</sup> *Faculty of Mathematics and Physics, Charles University*<sup>3</sup> *Institute of Computer Science, Masaryk University***Corresponding Authors:** xhejtman@ics.muni.cz, michal.chudoba.praha@gmail.com, jiri.chudoba@cern.ch

Users may have difficulties to find the needed information in the documentation for products, when many pages of documentation are available on multiple web pages or in email forums. We have developed and tested an AI based tool, which can help users to find answers to their questions. The Docu-bot uses Retrieval Augmentation Generation solution to generate answers to various questions. It uses github or open gitlab repositories with documentation as a source of information. Zip files with documentation in a plain text or markdown format can also be used for input. Sentence transformer model and Large Language Model generate answers.

Different LLM models can be used. For performance reasons, in most tests we use the model Mistral-7B-Instruct-v0.2, which fits into the memory of the Nvidia T4 GPU. We have also tested a larger model Mixtral-8x7B-Instruct-v0.1, which requires more GPU memory, available for example on Nvidia A100, A40 or H100 GPU cards. Another possibility is to use the API of OpenAI models like gpt-3.5-turbo, but the user has to provide his/her own API access key to cover expenses.

**Parallel (Track 5) / 41**

## Hardware acceleration for next-to-leading order event generation within MadGraph5\_aMC@NLO

**Authors:** Andrea Valassi<sup>1</sup>; Marco Zaro<sup>2</sup>; Olivier Mattelaer<sup>3</sup>; Stefan Roiser<sup>1</sup>; Zenny Jovi Joestar Wettersten<sup>1</sup><sup>1</sup> *CERN*<sup>2</sup> *Università degli Studi e INFN Milano (IT)*<sup>3</sup> *UCLouvain***Corresponding Authors:** andrea.valassi@cern.ch, stefan.roiser@cern.ch, zenny.wettersten@cern.ch, olivier.mattelaer@uclouvain.be, marco.zaro@cern.ch

As the quality of experimental measurements increases, so does the need for Monte Carlo-generated simulated events —both with respect to total amount, and to their precision. In perturbative methods this involves the evaluation of higher order corrections to the leading order (LO) scattering amplitudes, including real emissions and loop corrections. Although experimental uncertainties today are larger than those of simulations, at the High Luminosity LHC experimental precision is expected to increase above the theoretical one for events generated below next-to-leading order (NLO) precision. As forecasted hardware resources do not meet CPU requirements for these simulation needs, speeding up NLO event generation is a necessity for particle physics research.

In recent years, collaborators across Europe and the United States have been working on CPU vectorisation of LO event generation within the MadGraph5\_aMC@NLO framework, as well as porting it to GPUs, to major success. Recently, development has also started on vectorising NLO event generation. Due to the more complicated nature of NLO amplitudes this development faces several difficulties not accounted for in LO development. Nevertheless, this development seems promising, and a status report as well as the latest results will be presented in this contribution.

Parallel (Track 3) / 42

## Optimal use of timing measurement in vertex reconstruction at CMS

**Author:** Ksenia de Leo<sup>1</sup>

<sup>1</sup> *INFN Trieste (IT)*

**Corresponding Author:** ksenia.de.leo@cern.ch

The upgrade of the CMS apparatus for the HL-LHC will provide unprecedented timing measurement capabilities, in particular for charged particles through the Mip Timing Detector (MTD). One of the main goals of this upgrade is to compensate the deterioration of primary vertex reconstruction induced by the increased pileup of proton-proton collisions by separating clusters of tracks not only in space but also in time.

This contribution discusses the ongoing algorithmic developments to optimally exploit such new information, going beyond the initial studies at the time of the detector proposal, both from the physics and computational performance point of view. Different possible approaches are evaluated, comparing improvements of traditional methods and innovative techniques.

Poster session / 43

## A new SymPy backend for vector: uniting experimental and theoretical physicists

**Authors:** Jim Pivarski<sup>1</sup>; Saransh Chopra<sup>2</sup>

<sup>1</sup> *Princeton University*

<sup>2</sup> *Princeton University (US)*

**Corresponding Authors:** saransh.chopra@cern.ch, pivarski@princeton.edu

Vector is a Python library for 2D, 3D, and Lorentz vectors, especially arrays of vectors, to solve common physics problems in a NumPy-like way. Vector currently supports creating pure Python Object, NumPy arrays, and Awkward arrays of vectors. The Object and Awkward backends are implemented in Numba to leverage JIT-compiled vector calculations. Furthermore, vector also supports JAX and Dask operations on Awkward arrays of vectors.

We introduce a new SymPy backend in vector to allow symbolic computations on high energy physics vectors. Along with experimental physicists using vector for numerical computations, the SymPy backend will enable theoretical physicists to utilize the library for symbolic computations. Since the SymPy vector classes and their momentum equivalents operate on SymPy expressions, all of the standard SymPy methods and functions work on the vectors, vector coordinates, and the results of operations carried out on vectors. Moreover, vector's SymPy backend will create a stronger connection between software used by experimentalists and software used by theorists.

This talk will introduce vector and its backends to the users and funnel down to the SymPy backend. Finally, vector's SymPy backend is relatively new; hence, we aim to collect suggestions and recommendations from both theoretical and experimental physicists.

Parallel (Track 2) / 44

## Investigation of SVM-based 2D Heterojunctions for Data Filtering in Tracking System

**Author:** Tupendra Oli<sup>1</sup>

**Co-authors:** Jinlong Zhang<sup>1</sup>; Salman Habib<sup>1</sup>

<sup>1</sup> *Argonne National Laboratory*

**Corresponding Authors:** habib@anl.gov, zhangjl@anl.gov, toli@anl.gov

As particle physics experiments evolve to achieve higher energies and resolutions, handling the massive data volumes produced by silicon pixel detectors, which are used for charged particle tracking, poses a significant challenge. To address the challenge of data transport from high-resolution tracking systems, we investigate a support vector machine (SVM)-based data classification system designed to reject low-momentum particles in real-time. Such a classification can allow for the rejection of less important low-momentum charge clusters, thereby reducing the data volume before the data down streaming. This SVM system achieves high accuracy through the use of a custom kernel function, mixed-kernel, which is specifically adapted to the data recorded by a silicon tracker. Moreover, this custom kernel can be implemented using highly efficient, novel 2D heterojunction devices. In this talk, we will discuss about the SVM model, its performance, and the co-design approach that may be adapted to meet future device and processing needs in high-energy physics (HEP) collider experiments.

**Poster session / 45**

## Benchmarking OSDF services to develop XrootD best practices

**Authors:** Fabio Andrijauskas<sup>1</sup>; Igor Sfiligoi<sup>2</sup>; Frank Wurthwein<sup>3</sup>

<sup>1</sup> *Univ. of California San Diego (US)*

<sup>2</sup> *University of California San Diego*

<sup>3</sup> *UCSD*

**Corresponding Authors:** fkw@ucsd.edu, isfiligoi@ucsd.edu, fabio.andrijauskas@cern.ch

Research has become dependent on processing power and storage, with one crucial aspect being data sharing. The Open Science Data Federation (OSDF) project aims to create a scientific global data distribution network, expanding on the StashCache project to add new data origins and caches, access methods, monitoring, and accounting mechanisms. OSDF does not develop any new software, relying on the XrootD and Pelican projects, instead. Nevertheless, it is vital for OSDF to understand the XrootD limits under various configuration options, including transfer rate limits, proper buffer configuration and storage type effect. We have thus executed a set of benchmarks with the goal of creating a set of recommendations to share with the XrootD and Pelican teams. In this work we describe the tests and the results performed using hosts on the National Research Platform (NRP). The tests cover a wide range of file sizes and parallel streams, and use clients located at various distances from the server host. We also used several standalone clients (wget, curl, pelican) and the native HTCondor file transfer mechanisms.

**Parallel (Track 5) / 46**

## Graph-based Full Event Interpretation: a graph neural network for event reconstruction in Belle II

**Authors:** Corentin Santos<sup>1</sup>; Giulio Dujany<sup>2</sup>; Jacopo Cerasoli<sup>3</sup>; Merna AbuMusabh<sup>None</sup>

<sup>1</sup> *University of Strasbourg*

<sup>2</sup> *IPHC - CNRS*

<sup>3</sup> *Université de Strasbourg, CNRS, IPHC, UMR 7178, 67037 Strasbourg, France*

**Corresponding Authors:** corentin.santos@etu.unistra.fr, merna.abumusabh2@etu.unistra.fr, giulio.dujany@iphc.cnrs.fr, jacopo.cerasoli@cern.ch

In this work we present the Graph-based Full Event Interpretation (GraFEL), a machine learning model based on graph neural networks to inclusively reconstruct events in the Belle II experiment. Belle II is well suited to perform measurements of  $B$  meson decays involving invisible particles (e.g. neutrinos) in the final state. The kinematical properties of such particles can be deduced from the energy-momentum imbalance obtained after reconstructing the companion  $B$  meson produced in the event. This task is performed by reconstructing it either from all the particles in an event but the signal tracks, or using the Full Event Interpretation, an algorithm based on Boosted Decision Trees and limited to specific, hard-coded decay processes. A recent example involving the use of the aforementioned techniques is the search for the  $B^+ \rightarrow K^+ \nu \bar{\nu}$  decay, that provided an evidence for this process at about 3 standard deviations.

The GraFEL model is trained to predict the structure of the decay chain by exploiting the information from the detected final state particles only, without making use of any prior assumptions about the underlying event. By retaining only signal-like decay topologies, the model considerably reduces the amount of background while keeping a relatively high signal efficiency. The performances of the model when applied to the search for  $B^+ \rightarrow K^+ \nu \bar{\nu}$  are presented. The implementation of the model in the Belle II Analysis Software Framework is discussed.

**Parallel (Track 5) / 47**

## On Demand Column Joining for End User Analysis

**Authors:** Benjamin Galewsky<sup>1</sup>; Burt Holzman<sup>2</sup>; Nick Manganelli<sup>3</sup>

**Co-authors:** Gordon Watts<sup>4</sup>; Keith Ulmer<sup>5</sup>; Lindsey Gray<sup>2</sup>

<sup>1</sup> *Univ. Illinois at Urbana Champaign (US)*

<sup>2</sup> *Fermi National Accelerator Lab. (US)*

<sup>3</sup> *University of Colorado Boulder (US)*

<sup>4</sup> *University of Washington (US)*

<sup>5</sup> *University of Colorado, Boulder (US)*

**Corresponding Authors:** burt@fnal.gov, lindsey.gray@cern.ch, gwatts@uw.edu, nicholas.james.manganelli@cern.ch, keith.ulmer@cern.ch, bengal1@illinois.edu

The high luminosity LHC (HL-LHC) era will deliver unprecedented luminosity and new detector capabilities for LHC experiments, leading to significant computing challenges with storing, processing, and analyzing the data. The development of small, analysis-ready storage formats like CMS NanoAOD (4kB/event), suitable for up to half of physics searches and measurements, helps achieve necessary reductions in data processing and storage. However, a large fraction of analyses frequently require very computationally expensive machine learning output or data only stored in larger and less accessible formats, such as CMS MiniAOD (45kB/event) or AOD (450kB/event). This necessitates the non-volatile storage of derived data in custom formats. In this work, we present research on the development of workflows and integration of tools with ServiceX to efficiently fetch, cache, and join together data for use with columnar analysis tools.

We leverage scaleable, distributed SQL query engines like Trino to join disparate columns sourced from multiple files and without a restriction on relative row ordering. By replacing many customized datasets, containing largely overlapping contents, with smaller and unique sets of information that can be joined on demand with common central data, duplication can be reduced. Caching of these results keeps the cost of subsequent retrieval low, fitting well with modern physics analysis paradigms.

**Poster session / 48**

## Running version control systems and continuous integration to fulfill CERN and LHC experiments needs

**Author:** Ismael Posada Trobo<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** ismael.posada.trobo@cern.ch

GitLab Runners have been deployed at CERN since 2015. A GitLab runner is an application that works with GitLab Continuous Integration and Continuous Delivery (CI/CD) to run jobs in a pipeline. CERN provides runners that are available to the whole GitLab instance and can be used by all eligible users. Until 2023, CERN was providing a fixed amount of Docker runners executing in OpenStack virtual machines, following an in-house, customized solution that utilized the Docker+machine executor. This solution served its purpose for several years, however needed to be reviewed due to the deprecation by Docker Inc., and only a fork maintained by GitLab Inc.

During the last few years, the demand and the number of running pipelines have substantially increased, as the adoption of Continuous Integration and Delivery has been rapidly growing.

In view of the above, CERN needed to provide a supported, scalable infrastructure that would accommodate our users' demand.

This paper describes the process of how CERN migrated from the legacy in-house solution to a new scalable, reliable and easy-to-maintain solution of runners based on Kubernetes, including the challenges and lessons learned that have been faced during this complex migration process.

Parallel (Track 7) / 49

## Global Networking Challenges for the Coming Decade

**Authors:** Edoardo Martelli<sup>1</sup>; Tony Cass<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** edoardo.martelli@cern.ch, tony.cass@cern.ch

Back in the late 1990's when planning for LHC computing started in earnest, arranging network connections to transfer the huge LHC data volumes between participating sites was seen as a problem. Today, 30 years later, the LHC data volumes are even larger, WLCG traffic has switched from a hierarchical to a mesh model and yet almost nobody worries about the network.

Some people still do worry, however. Even if LHC data transfers still account for over 50% of NREN traffic, other data-intensive experiments are coming on stream and network engineers worry about managing the overall traffic efficiently.

We present here the challenges likely to be keeping network engineers busy in the coming decade: how to monitor traffic from different communities, how to avoid congestion over transoceanic links; how to smooth traffic flows to maximise throughput, hand-over of large flows at interconnection points; cyber security and more.

Parallel (Track 5) / 50

## Monte Carlo challenges for Non Perturbative QED

**Author:** Anthony Hartin<sup>1</sup>

<sup>1</sup> LMU

**Corresponding Author:** anthony.hartin@cern.ch

Non perturbative QED is used to predict beam backgrounds at the interaction point of colliders, in calculations of Schwinger pair creation and in precision QED tests with ultra-intense lasers. In order to predict these phenomena, custom built monte carlo event generators based on a suitable non perturbative theory have to be developed. One such suitable theory uses the Furry Interaction Picture, in which a background field is taken into account non perturbatively at Lagrangian level. This theory is precise, but the transition probabilities are in general, complicated. This poses a challenge for the monte carlo which struggles to implement the theory computatively. The monte carlo must in addition taken into account the behaviour of the background field at every space-time point at which an event is generated. We introduce here just such a monte carlo package, called IPstrong, and the techniques implemented to deal with the specific challenges outlined above.

**Parallel (Track 7) / 51**

## Monitoring particle accelerators with wireless IoT

**Author:** Rodrigo Sierra<sup>1</sup>

**Co-author:** Xoan Carlos Cosmed Peralejo <sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** xoan.carlos.cosmed.peralejo@cern.ch, rodrigo.sierra@cern.ch

Although wireless IoT devices are omnipresent in our homes and workplaces, their use in particle accelerators is still uncommon. Although the advantages of movable sensors communicating over wireless networks are obvious, the harsh radiation environment of a particle accelerator has been an obstacle to the use of such sensitive devices. Recently, though, CERN has developed a radiation-hard LoRaWAN based platform that can be adapted to support multiple sensors.

We report here on this platform, the deployment of an LPWAN network based on LoRaWAN technology in the underground areas at CERN, the infrastructure and tools developed to support device integration and data collection, and, finally, on some of the positive benefits that have been delivered through the use of these sensors in CERN's accelerator complex.

**Poster session / 52**

## Ending the Analogue Telephony Era

**Author:** Rodrigo Sierra<sup>1</sup>

**Co-authors:** Ihor Olkhovskiy <sup>1</sup>; Xoan Carlos Cosmed Peralejo <sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** ihor.olkhovskiy@cern.ch, xoan.carlos.cosmed.peralejo@cern.ch, rodrigo.sierra@cern.ch

2024 marks not just CERN's 70th birthday but also the end of analogue telephony at the laboratory. Traditional phone exchanges and the associated copper cabling cannot deliver 21st-century communication services and a decade-long project to modernize CERN's telephony infrastructure was completed earlier this year.

We report here on CERN's modern fixed telephony infrastructure, firstly our in-house development of an exchange which, based on open-source components and standard VoIP protocols, supports softphones, call centers, safety communications, interconnections with other voice services and an automatic switchboard, and secondly the two CERNphone applications that have replaced fixed phones, and which are used by more than 6000 users each week.

Parallel (Track 1) / 53

## Upcoming database developments at CERN

**Author:** Andrzej Nowicki<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** andrzej.nowicki@cern.ch

In this presentation, I will outline the upcoming transformations set to take place within CERN's database infrastructure. Among the challenges facing our database team during the Long Shutdown 3 (LS3) will be the upgrade of Oracle databases.

The forthcoming version of Oracle database is introducing a significant internal change as the databases will be converted to a container architecture. This presentation will introduce the concept of a container database and show the challenges and new opportunities it brings.

This major upgrade will also introduce numerous new features. Some of the expected functionalities are vector data types and vector indexes, as well as JSON Relational duality. Additionally, we'll explore the potential benefits of adopting new schema-level privileges and Oracle True Cache technology.

Seamless and efficient transition necessitates close collaboration between the database team and its users. This presentation aims to inform the community of the upcoming changes and the efforts required to ensure a smooth transition.

Poster session / 54

## Accounting of (HPC) computing resources with AUDITOR

**Author:** Michael Boehler<sup>1</sup>

**Co-authors:** Benjamin Rottler <sup>1</sup>; Dirk Sammel <sup>1</sup>; Florian Von Cube <sup>2</sup>; Manuel Giffels <sup>2</sup>; Matthias Schnepf ; Max Fischer <sup>3</sup>; Michael Huebner <sup>4</sup>; Oliver Freyermuth <sup>4</sup>; Raghuvar Vijayakumar <sup>1</sup>; Raphael Kleinemuehle <sup>5</sup>

<sup>1</sup> Albert Ludwigs Universitaet Freiburg (DE)

<sup>2</sup> KIT - Karlsruhe Institute of Technology (DE)

<sup>3</sup> Karlsruhe Institute of Technology

<sup>4</sup> University of Bonn (DE)

<sup>5</sup> Uni Wuppertal

**Corresponding Authors:** kleinemuehl@uni-wuppertal.de, ralf.florian.von.cube@cern.ch, matthias.schnepf@kit.edu, benjamin.rottler@cern.ch, raghuvar.vijayakumar@cern.ch, max.fischer@kit.edu, manuel.giffels@cern.ch, dirk.sammel@cern.ch, michael.huebner@uni-bonn.de, oliver.freyermuth@cern.ch, michael.boehler@cern.ch

New strategies for the provisioning of compute resources, e.g. in the form of dynamically integrated resources enabled by the COBaLD/TARDIS software toolkit, require a new approach of collecting accounting data. AUDITOR (AccoUnting DatahandlIng Toolbox for Opportunistic Resources), a flexible and expandable accounting ecosystem that can cover a wide range of use cases and infrastructures, was developed specifically for this purpose. Accounting data is collected via so-called collectors and stored in a database. So-called plugins can access the data and act based on the accounting information. Access to the data is handled by the core component of AUDITOR, which provides a REST API together with a Rust and a Python client library.

An HTCondor collector, a Slurm collector and a TARDIS collector are currently available, and a Kubernetes collector is already in the works.

The APEL plugin enables, for example, the creation of APEL accounting summaries and their transmission to the APEL accounting server. Although the original aim of the development of AUDITOR was to enable the accounting of opportunistic resources managed by COBaLD/TARDIS, it can also be used for standard accounting of a WLCG computing resource. As AUDITOR uses a highly flexible



data structure to store accounting data, extensions such as GPU resource accounting can be added with minimal effort.

This contribution provides insights into the design of AUDITOR and shows how it can be used to enable a number of different use cases.

**Poster session / 55**

## Scalable and multi-tenant Kubernetes ingress infrastructure

**Author:** Jack Henschel<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** jack.henschel@cern.ch

Since 2016, CERN has been using the OpenShift Kubernetes Distribution to host a platform-as-a-service (PaaS). This service is optimized for hosting web applications and has grown to tens of thousands of individual websites. By now, we have established a reliable framework that deals with varied use cases: thousands of websites per ingress controller (8K+ hostnames), handling with long-lived connections (30K+ concurrent sessions) and high traffic applications (25TB+ per day).

This session will discuss:

- CERN's web hosting infrastructure based on OpenShift Kubernetes clusters;
- usage of open source and in-house developed software for providing a seamless user experience;
- integrations for registering hostnames (local DNS, LanDB, external providers)
- provisioning of certificates (automatic with external-dns / ACME HTTP-01, manual provisioning)
- access control policies and “connecting” different components with OpenPolicyAgent
  - enforcing unique hostnames across multiple Kubernetes clusters
- strategies for setting up Kubernetes Ingress Controllers for multi-tenant clusters;
- methods for scaling and sharding ingress controllers according to application requirements (specifically HAProxy ingress controllers);

**Parallel (Track 1) / 57**

## FTS3 Token Support for a Proxy-less WLCG world

**Author:** Mihai Patrascoiu<sup>1</sup>

**Co-authors:** Joao Pedro Lopes ; Luca Mascetti <sup>1</sup>; Shubhangi Misra ; Steven Murray <sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** steven.murray@cern.ch, joao.pedro.batista.lopes@cern.ch, shubhangi.misra@cern.ch, luca.mascetti@cern.ch, mihai.patrascoiu@cern.ch

The WLCG community, with the main LHC experiments at the forefront, is moving away from x509 certificates, replacing the Authentication and Authorization layer with OAuth2 tokens. FTS, as a middleware and core component of the WLCG, plays a crucial role in the transition from x509 proxy certificates to tokens. The paper will present in-detail the FTS token design and how this will serve the needs of the community, WLCG and non-WLCG alike. Finally, a chapter will also report on performance measurements and lessons learned during the DataChallenge 2024

**Parallel (Track 1) / 58**

## Ceph at CERN in the multi-datacentre era

**Authors:** Abhishek Lekshmanan<sup>1</sup>; Enrico Bocchi<sup>1</sup>; Roberto Valverde Cameselle<sup>1</sup>; Zachary Goggin<sup>None</sup><sup>1</sup> CERN**Corresponding Authors:** abhishek.lekshmanan@cern.ch, enrico.bocchi@cern.ch, zachary.goggin@cern.ch, roberto.valverde.cameselle@cern.ch

The recent commissioning of CERN's Preveessin Data Centre (PDC) brings the opportunity for multi-datacentre Ceph deployments, bringing advantages for business continuity and disaster recovery. However, the simple extension of a single cluster across data centres is impractical due to the impact of latency on Ceph's strong consistency requirements. This paper reports on our research towards building a multi datacentre Ceph deployment in production. Due to the nature of different transaction semantics for blocks, objects and files, geo-distributing a ceph cluster needs a different approach for each protocol in use. This paper will detail the challenges with Ceph across data centres, the various solutions we evaluated and a roadmap for the future at CERN.

**Parallel (Track 1) / 59**

## FTS as a part of the SKA data movement pipeline

**Author:** Rose Cooper<sup>None</sup>**Co-author:** Tom Dack**Corresponding Authors:** tom.dack@cern.ch, rose.cooper@stfc.ac.uk

The File Transfer Service (FTS) is a bulk data mover responsible for queuing, scheduling, dispatching and retrying file transfer requests, making it a critical infrastructure component for many experiments. FTS is primarily used by the LHC experiments, namely ATLAS, CMS and LHCb, but is also used by some non-LHC experiments, including both AMS and DUNE. FTS is an essential part in the data movement pipeline for these experiments and is responsible for moving their data across the world via the worldwide LHC computing Grid (WLCG).

The Square Kilometre Array (SKA) is a multi-purpose radio telescope that will play a major role in answering key questions in modern astrophysics and cosmology. The SKA will have a survey speed a hundred times that of current radio telescopes and its capabilities will allow transformational experiments to be conducted in a wide variety of science areas. Whilst the headquarters for this project is located at Jodrell Bank in the UK, the main telescope sites are located in South Africa and Australia. The two telescope sites will produce approximately 700 PB of data per year, which will need to be moved to one of the SKA regional centres located in member countries around the world to be stored, before being accessed by scientists. It is evident that there will be several similarities between the computing requirements for the LHC and SKA experiments, in particular the challenges posed by moving large quantities of data around a global network.

In this talk, we will discuss the usage of FTS by SKA and its ability to enable long-range data transfer across the developing SKA regional centre network of sites. We will also discuss some alterations to the FTS service ran at STFC to better support SKA, most notably the migration to token based authentication away from X509 certificates.

**Poster session / 60**

## S3 Compatibility: Enabling Seamless Integration with EOS, CERN's Large-Scale Disk Storage System

**Authors:** Andreas Joachim Peters<sup>1</sup>; Elvin Alin Sindrilaru<sup>1</sup>; Luca Mascetti<sup>1</sup>; Mano Segransan<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> 42 Lausanne (CH)

**Corresponding Authors:** elvin.alin.sindrilaru@cern.ch, mano.segransan@cern.ch, luca.mascetti@cern.ch, andreas.joachim.peters@cern.ch

Amazon S3 is a leading object storage service known for its scalability, data reliability, security and performance. It is used as a storage solution for data lakes, websites, mobile applications, backup, archiving and more. With its management features, users can optimise data access to meet specific requirements and compliance standards. Given its popularity, many tools utilise the S3 interfaces. To enhance CERN's EOS Big Data storage, we are integrating an S3 interface into XRootD that is customised for EOS. This article describes the design, progress and future plans for the integration of the S3 API.

Parallel (Track 4) / 61

## WLCG Operations: evolution and challenges on the way to High Luminosity LHC

**Authors:** Julia Andreeva<sup>1</sup>; Maarten Litmaath<sup>1</sup>; Panos Paparrigopoulos<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** maarten.litmaath@cern.ch, panos.paparrigopoulos@cern.ch, julia.andreeva@cern.ch

The WLCG infrastructure is quickly evolving thanks to technology evolution in all areas of LHC computing: storage, network, alternative processor architectures, new authentication & authorization mechanisms, etc. This evolution also has to address challenges like the seamless integration of HPC and cloud resources, the significant rise of energy costs, licensing issues and support changes. WLCG Operations Coordination serves to organize parts of this evolution in close collaboration with the stakeholders: LHC experiments, sites, EGI and OSG infrastructure providers and middleware projects. This contribution describes how WLCG Operations Coordination helps ensure the smooth functioning of the WLCG infrastructure during Run 3 addressing the challenges described above and its preparation for the High-Luminosity LHC phase.

Plenary session / 62

## The neXt Dirac incarnation

**Author:** Federico Stagni<sup>1</sup>

**Co-authors:** Alexandre Franck Boyer<sup>1</sup>; Andrei Tsaregorodtsev<sup>2</sup>; Chris Burr<sup>1</sup>; Christophe Haen<sup>1</sup>

<sup>1</sup> CERN

<sup>2</sup> Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France

**Corresponding Authors:** alexandre.franck.boyer@cern.ch, federico.stagni@cern.ch, atsareg@in2p3.fr, christophe.denis.haen@cern.ch, christopher.burr@cern.ch

The Dirac interware has long served as a vital resource for user communities seeking access to distributed computing resources. Originating within the LHCb collaboration around 2000, Dirac has undergone significant evolution. A pivotal moment occurred in 2008 with a major refactoring, resulting in the development of the experiment-agnostic core Dirac, which paved the way for customizable extensions like LHCbDirac and BelleDirac, among others.

Despite its efficacy in meeting experiment-specific requirements, Dirac has accrued technical debt over its 15-year history. Installation management remains intricate, with significant entry barriers and a reliance on bespoke infrastructure. Additionally, the software development process lacks alignment with contemporary standards, impeding the onboarding process for new developers. Notably,

integral components such as the network protocol and authentication mechanisms are proprietary and pose challenges for seamless integration with external applications.

In response to these challenges, the Dirac consortium has embarked on the development of DiracX. Drawing upon two decades of experience and battle-tested technological frameworks, DiracX heralds a new era in distributed computing solutions. This contribution describes technical decisions, roadmap and timelines for the development of DiracX.

This article presents an overview of the architecture underpinning DiracX, shedding light on the technological decisions guiding its development. Recognizing the criticality of maintaining a continuously operational Dirac system for numerous user communities, we delve into the intricacies of the migration process from Dirac to DiracX.

Poster session / 63

## Advancements and Operations for LHC Run-3 and beyond

**Authors:** Abhishek Lekshmanan<sup>1</sup>; Andreas Joachim Peters<sup>1</sup>; Cedric Caffy<sup>1</sup>; David Smith<sup>1</sup>; Elvin Alin Sindrilaru<sup>1</sup>; Gianmaria Del Monte<sup>1</sup>; Guilherme Amadio<sup>1</sup>; Ioanna Vrachnaki<sup>None</sup>; Luca Mascetti<sup>1</sup>; Maria Arsuaga Rios<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** david.smith@cern.ch, andreas.joachim.peters@cern.ch, amadio@cern.ch, ioanna.vrachnaki@cern.ch, luca.mascetti@cern.ch, elvin.alin.sindrilaru@cern.ch, cedric.caffy@cern.ch, abhishek.lekshmanan@cern.ch, maria.arsuaga.rios@cern.ch, gianmaria.del.monte@cern.ch

The aim of this paper is to give an overview of the progress made in the EOS project - the large scale data storage system developed at CERN - during the preparation and during LHC Run-3. Developments consist of further simplification of the service architecture, metadata performance improvements, new memory inventory and cost & value interfaces, a new scheduler implementation, a generated REST API derived from the GRPC protocol, and new or better integration of features such as SciTags and SciTokens. We will report on operational experiences and the massive migration process to ALMA9, improvements in the quality assurance process and results achieved. Looking to the future, we will describe the development and evolution of EOS for Run-4 and highlight various software R&D and technology evaluation activities (e.g. SMR support) that have the potential to help realize the Run-4 requirements for physics storage at CERN and elsewhere.

Parallel (Track 1) / 64

## Next-Gen Storage Infrastructure for ALICE: Paving the Road Toward Hi-Luminosity LHC

**Authors:** Abhishek Lekshmanan<sup>1</sup>; Andreas Joachim Peters<sup>1</sup>; Cedric Caffy<sup>1</sup>; David Smith<sup>1</sup>; Elvin Alin Sindrilaru<sup>1</sup>; Gianmaria Del Monte<sup>1</sup>; Guilherme Amadio<sup>1</sup>; Ioanna Vrachnaki<sup>None</sup>; Luca Mascetti<sup>1</sup>; Maria Arsuaga Rios<sup>1</sup>; Volodymyr Yurchenko<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> National Academy of Sciences of Ukraine (UA)

**Corresponding Authors:** luca.mascetti@cern.ch, elvin.alin.sindrilaru@cern.ch, gianmaria.del.monte@cern.ch, ioanna.vrachnaki@cern.ch, maria.arsuaga.rios@cern.ch, andreas.joachim.peters@cern.ch, cedric.caffy@cern.ch, volodymyr.yurchenko@cern.ch, amadio@cern.ch, abhishek.lekshmanan@cern.ch, david.smith@cern.ch

ALICE introduced ground-breaking advances in data processing and storage requirements and presented the CERN IT data centre with new challenges with the highest data recording requirement of all experiments. For these reasons, the EOS O2 storage system was designed to be cost-efficient, highly redundant and maximise data resilience to keep data accessible even in the event of unexpected disruptions or hardware failures. With 150 PB of usable storage space, EOS O2 is now the largest disk storage system in use at CERN. We will report on our experience and the effectiveness of

operating this full production system in Run-3 and during the LHC heavy-ions run and on how this will help in paving the road towards the data deluge coming with Hi-Luminosity LHC. In particular, we will report on our experience with RS(10+2) erasure coding in production, the achievable performance of EOS O2, reliability figures, life cycle management, capacity extension and rebalancing operations.

Parallel (Track 1) / 65

## Physics Data Forge: Unveiling the Power of IO Systems in CERN's Massive IO Scale Test Infrastructure

**Authors:** Andrea Sciabà<sup>1</sup>; Andreas Joachim Peters<sup>1</sup>; David Smith<sup>1</sup>; Guilherme Amadio<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** david.smith@cern.ch, andrea.sciaba@cern.ch, andreas.joachim.peters@cern.ch, amadio@cern.ch

In this paper, IO scale tests at CERN are investigated in preparation for the further development of IO systems. Remote file access is critical in High Energy Physics (HEP) and is currently facilitated by XRootD and HTTP(S) protocols. With a tenfold increase in data volume expected for Run-4, higher throughput is critical. We compare some client-server implementations on 100- and 200-GE LANs connected to high-throughput storage devices. A joint project between IT and EP departments aims to evaluate RNTuple as a replacement for ROOT's TTree format, with a focus on verifying the scalability of EOS storage using a new data format. Based on this project we run a large-scale experiment to stress CERN's CPU, network and memory by using massively parallel analysis workflows and replicated datasets in EOS physics instances.

Poster session / 66

## Tracking efficiency studies for LHCb in Run 3

**Authors:** Peilian Li<sup>1</sup>; Rowina Caspary<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> Heidelberg University (DE)

**Corresponding Authors:** rowina.caspary@cern.ch, peilian.li@cern.ch

The LHCb detector, a multi-purpose detector with a main focus on the study of hadrons containing b- and c-quarks, has been upgraded to enable precision measurements at an instantaneous luminosity of  $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$  at  $\sqrt{s} = 14$  TeV, five times higher than the previous detector capacity. With the almost completely new detector, a software-only trigger system has been developed and all track reconstruction algorithms have been redesigned.

The knowledge of the track reconstruction efficiency at different momenta and regions of the detector is essential for many analyses including cross-section and asymmetry measurements. A tag-and-probe method is developed to estimate the tracking efficiency using muonic tracks from  $J/\psi \rightarrow \mu^+ \mu^-$  decays, where the probe tracks are reconstructed excluding hits from the tracking subdetectors under scrutiny.

A complementary method is exploited to address tracking efficiency corrections due to the hadronic interactions with the detector material using pions from  $D^0 \rightarrow K\pi$  and  $D^0 \rightarrow K\pi\pi\pi$  decays. In this talk, these data-driven methods and their applications to the data taken in 2023 and 2024 are presented.

Poster session / 67

## Exploring Data Caching Policy with Data Access Patterns from dCache Logs

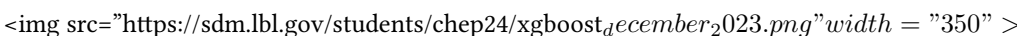
**Author:** Jacob Aldrich<sup>1</sup>**Co-authors:** Alex Sim ; Eric Christian Lancon <sup>2</sup>; Hironori Ito <sup>2</sup>; John Wu <sup>3</sup>; Shinjae Yoo ; Vincent Garonne <sup>2</sup><sup>1</sup> *University of California, Berkeley*<sup>2</sup> *Brookhaven National Laboratory (US)*<sup>3</sup> *LAWRENCE BERKELEY NATIONAL LABORATORY***Corresponding Authors:** elancon@bnl.gov, vincent.garonne@cern.ch, hito@bnl.gov, jacobaldrich11@berkeley.edu, sjyoo@bnl.gov, kwu@lbl.gov, asim@lbl.gov

The dCache storage management system at Brookhaven National Lab plays a vital role as a disk cache, storing extensive datasets from high-energy physics experiments, mainly the ATLAS experiment. Given that dCache's storage is significantly smaller than the total ATLAS data, it's crucial to have an efficient cache management policy. A common approach is to keep files that are accessed often, ready for future use. In our research, we analyze both recent and past patterns of file usage to predict the chances of them being needed again. Although dCache considers each file separately, we've observed that files within a dataset tend to be used together. Therefore, the system manager often gets requests to retain entire datasets in the cache, especially if they're expected to be in high demand soon. Our main focus is to determine if we could accurately forecast a dataset's future demand to automate the process of deciding which datasets to prioritize in the cache.

Our approach's cornerstone is a dynamic learning mechanism that regularly analyzes recent access logs. This process updates our machine learning models, enabling them to forecast the popularity of various datasets shortly. Specifically, our predictive model estimates the expected number of accesses for each dataset in the upcoming days. We then synchronize these predictions with the cache space allocated for monitoring sought-after datasets. This allows us to proactively load the most in-demand datasets into the disk cache. This strategic reservation method operates in conjunction with the current file removal policy, collectively improving the overall efficiency of the system.

To develop a predictive model for our caching system, we assessed several techniques and metrics to distinguish popular datasets from less popular ones effectively. Employing k-means clustering, we categorized datasets based on their popularity and explored diverse methods to precisely measure dataset usage. Given our constrained disk space, our aim was to optimize the selection of retained datasets, thereby improving cache efficiency.

Prior study [1] has demonstrated the feasibility of detecting popular datasets using a machine learning approach. In this study, we compare the predictive efficacy of two distinct models: a neural network model and a gradient-boosted trees regression model (XGBoost). The models, configured with 17 input variables, are trained on 127 million data points, collected over a span of three years from our data processing pipeline. Additionally, both models underwent hyperparameter tuning via Optuna, conducted on Perlmutter at NERSC.



*Fig.1. December 2023 comparison of predicted and actual dataset accesses using XGBoost. Axes represent next-day actual (x) vs. predicted (y) accesses. Points are colored based on the recency of last access, with lighter points indicating more recent accesses.*

Despite the inherent difficulty in forecasting future dataset accesses, our models showed promising performance. Notably, the XGBoost model displayed a lower root mean squared error (RMSE) for testing datasets compared to the neural network. Specifically, the relative ratios of testing RMSE to standard deviation were 0.28 for XGBoost and 0.84 for the neural network models.

Our research confirms that predicting dataset popularity is feasible through careful analysis of data features and the application of well-designed models. While the real-world application of these models in live caching policies requires further testing, our study underscores the potential of machine learning in improving dCache systems. Future endeavors will concentrate on implementing, benchmarking, and validating the efficacy of these proposed methods.

## REFERENCES

[1] J. Bellavita, C. Sim, K. Wu, A. Sim, S. Yoo, H. Ito, V. Garonne, and E. Lancon, "Understanding data access patterns for dcache system," in 26th International Conference on Computing in High Energy & Nuclear Physics (CHEP2023), 2023.

Poster session / 68

## Comparing Cache Utilization Trends for Regional Data Caches

**Author:** Alex Sim<sup>None</sup>

**Co-authors:** Chin Guok ; Diego Davila Foyo <sup>1</sup>; Frank Wurthwein <sup>2</sup>; Harvey Newman <sup>3</sup>; INDER MONGA <sup>4</sup>; John Wu <sup>5</sup>; Justas Balcas <sup>3</sup>; Ronak Monga <sup>6</sup>

<sup>1</sup> Univ. of California San Diego (US)

<sup>2</sup> UCSD

<sup>3</sup> California Institute of Technology (US)

<sup>4</sup> LAWRENCE BERKELEY NATIONAL L

<sup>5</sup> LAWRENCE BERKELEY NATIONAL LABORATORY

<sup>6</sup> Indiana University

**Corresponding Authors:** kwu@lbl.gov, ronmonga@iu.edu, fkw@ucsd.edu, asim@lbl.gov, imonga@es.net, justas.balcas@cern.ch, harvey.newman@cern.ch, diego.davila@cern.ch, chin@es.net

The surge in data volumes from large scientific collaborations, like the Large Hadron Collider (LHC), poses challenges and opportunities for High Energy Physics (HEP). With annual data projected to grow thirty-fold by 2028, efficient data management is paramount. The HEP community heavily relies on wide-area networks for global data distribution, often resulting in redundant long-distance transfers. This work studies how regional data caches [1] mitigate network congestion and enhance application performance, using millions of access records from regional caches in Southern California, Chicago, and Boston, serving the LHC's CMS experiment [2]. Our analysis demonstrates the potential of in-network caching to revolutionize large-scale scientific data dissemination.

We analyzed the cache utilization trends, see examples in Figure 1, and identified distinct patterns from the three deployments. Our exploration further employed machine learning models for predicting data access patterns within regional data caches. These findings offer insights crucial for resource usage estimation, including storage cache and network requirements, and can inform improvements in application workflow performance within regional data cache systems. Figure 2 shows the sample prediction result of the hourly cache utilization for SoCal Cache. While the predictions capture volume spikes, they may not precisely match the heights.

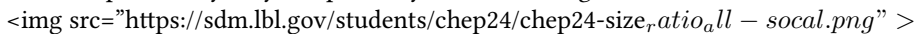


Fig.1(a).SoCalCache(capacity2PB)fromJune2020toDec.2023,with29.48milliondataaccessrequestswhere23.8PBw

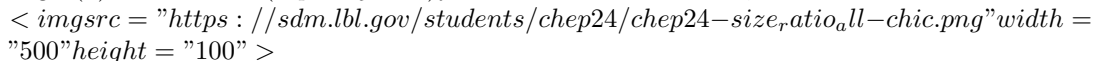
 <imgsrc="https://sdm.lbl.gov/students/chep24/chep24-size\_ratio\_all-chic.png"width="500"height="100" >

Fig.1(b).ChicagoCache(capacity340TB)fromOct.2022toDec.2023,with5.7milliondataaccessrequestswhere11.0PBw

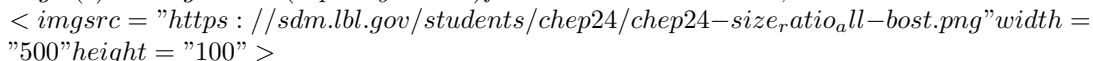
 <imgsrc="https://sdm.lbl.gov/students/chep24/chep24-size\_ratio\_all-bost.png"width="500"height="100" >

Fig.1(c).BostonCache(capacity170TB)fromAug.2023toDec.2023,with27.7milliondataaccessrequestswhere5.0PBw

Fig. 1. Fraction (by data volume) of cache hits (blue) and cache misses (orange) every day for three different in-network cache installations

<br>

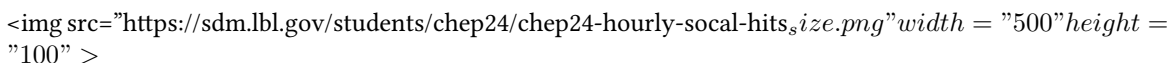
 <imgsrc="https://sdm.lbl.gov/students/chep24/chep24-hourly-social-hits\_size.png"width="500"height="100" >

Fig.2.HourlycachehitvolumefromSoCalCache,showingactualmeasurements(bluebars),trainingpredictions(orange)

In this study, we present a detailed analysis of cache utilization trends across the SoCal, Chicago, and Boston regional caches serving the LHC CMS experiment. Our investigation reveals distinct

patterns in cache usage unique to each region, highlighting the need for tailored approaches in cache design. Our exploration of neural network models were demonstrated to provide accurate predictions of cache usage trends, these models could be used anticipate future needs for storage and network resources. By understanding and leveraging these insights, we could significantly enhance the efficiency of resource allocation and optimize application workflow performance.

#### REFERENCES

- [1] E. Fajardo, D. Weitzel, M. Rynge, M. Zvada, J. Hicks, M. Selmecci, B. Lin, P. Paschos, B. Bockelman, A. Hanushevsky, F. Wurthwein, and I. Sfiligoi, "Creating a content delivery network for general science on the internet backbone using XCaches," EPJ Web of Conferences, vol. 245, p. 04041, 2020. [Online]. Available: <https://doi.org/10.1051/epjconf/202024504041>
- [2] C. Sim, K. Wu, A. Sim, I. Monga, C. Guok, D. Hazen, F. Wurthwein, D. Davila, H. Newman, and J. Balcas, "Predicting resource utilization trends with southern california petabyte scale cache," in 26th International Conference on Computing in High Energy & Nuclear Physics (CHEP2023), 2023.

Poster session / 69

## Computing and machine learning for the SHiP experiment

**Author:** Oliver Lantwin<sup>1</sup>

<sup>1</sup> *INFN Napoli*

**Corresponding Author:** [oliver.lantwin@cern.ch](mailto:oliver.lantwin@cern.ch)

The recently approved SHiP experiment aims to search for new physics at the intensity frontier, including feebly interacting particles and light dark matter, and perform precision measurements of tau neutrinos.

To fulfill its full discovery potential, the SHiP software framework is crucial, and faces some unique challenges due to the broad range of models under study, and the extreme statistics necessary for the background studies. The SHiP environment also offers unique opportunities for machine learning for detector design and anomaly detection.

This talk will give an overview of the general software framework and of past, ongoing and planned simulation and machine learning studies.

Parallel (Track 7) / 70

## Carbon, Power, and Sustainability in ATLAS Computing

**Authors:** Daniel Schien<sup>None</sup>; Emma Sian Kuwertz<sup>1</sup>; Paul Shabajee<sup>1</sup>; Rodney Walker<sup>2</sup>; Zach Marshall<sup>3</sup>

<sup>1</sup> *University of Bristol*

<sup>2</sup> *Ludwig Maximilians Universitat (DE)*

<sup>3</sup> *Lawrence Berkeley National Lab. (US)*

**Corresponding Authors:** [rodney.walker@physik.uni-muenchen.de](mailto:rodney.walker@physik.uni-muenchen.de), [zach.marshall@cern.ch](mailto:zach.marshall@cern.ch), [paul.shabajee@bristol.ac.uk](mailto:paul.shabajee@bristol.ac.uk), [emma.sian.kuwertz@cern.ch](mailto:emma.sian.kuwertz@cern.ch), [daniel.schien@bristol.ac.uk](mailto:daniel.schien@bristol.ac.uk)

The ATLAS Collaboration operates a large, distributed computing infrastructure: almost 1M cores of computing and almost 1 EB of data are distributed over about 100 computing sites worldwide. These resources contribute significantly to the total carbon footprint of the experiment, and they are expected to grow by a large factor as a part of the experimental upgrades for the HL-LHC at the end of the decade. This contribution describes various efforts to understand, monitor, and reduce the carbon footprint of the distributed computing of the experiment. This includes efforts to construct



a full life-cycle assessment (LCA) model for the carbon impact of ATLAS distributed computing, all with the goal of making recommendations for sites to reduce their carbon footprint for the HL-LHC.

**Plenary session / 71**

## The First Release of ATLAS Open Data for Research

**Authors:** Lukas Alexander Heinrich<sup>1</sup>; Mariana Isabel Vivas Albornoz<sup>2</sup>; Stephane Willocq<sup>2</sup>; Zach Marshall<sup>3</sup>

<sup>1</sup> *Technische Universitat Munchen (DE)*

<sup>2</sup> *University of Massachusetts (US)*

<sup>3</sup> *Lawrence Berkeley National Lab. (US)*

**Corresponding Authors:** stephane.willocq@cern.ch, mariana.isabel.vivas.albornoz@cern.ch, lukas.heinrich@cern.ch, zach.marshall@cern.ch

The ATLAS Collaboration has released an extensive volume of data for research use for the first time. The full datasets of proton collisions from 2015 and 2016, alongside a wide array of matching simulated data, are all offered in the PHYSLITE format. This lightweight format is chosen for its efficiency and is the preferred standard for ATLAS internal analyses. Additionally, the inclusion of Heavy Ion collision data considerably widens the scope for research within the particle physics community. To ensure accessibility and usability, the release includes a comprehensive suite of software tools and detailed documentation, catering to a varied audience. Code examples, from basic Jupyter notebooks to more complex C++ analysis packages, aim to facilitate engagement with the data. This contribution details the available data, corresponding metadata, software, and documentation, and initial interactions with researchers outside the ATLAS collaboration, underscoring the project's potential to foster new research and collaborations.

**Parallel (Track 3) / 72**

## Integration of the ACTS track reconstruction toolkit in the ATLAS software for HL-LHC operations

**Author:** Carlo Varni<sup>1</sup>

<sup>1</sup> *University of California Berkeley (US)*

**Corresponding Author:** carlo.varni@cern.ch

In view of the High-Luminosity LHC era the ATLAS experiment is carrying out an upgrade campaign which foresees the installation of a new all-silicon Inner Tracker (ITk) and the modernization of the reconstruction software.

Track reconstruction will be pushed to its limits by the increased number of proton-proton collisions per bunch-crossing and the granularity of the ITk detector. In order to remain within CPU budgets while retaining high physics performance, the ATLAS Collaboration plans to use ACTS, an experiment-independent toolkit for track reconstruction. The migration to ACTS involves the redesign of the track reconstruction components as well as the ATLAS Event Data Model (EDM), resulting in a thread-safe and maintainable software.

In this contribution, the current status of the ACTS integration for the ATLAS ITk track reconstruction is presented, with emphasis on the improvements of the track reconstruction software and the implementation of the ATLAS EDM.

**Parallel (Track6) / 73**

## Containerization in the ATLAS Offline Code Management System

**Authors:** Alexander Undrus<sup>1</sup>; Asoka De Silva<sup>2</sup>; Shuwei Ye<sup>1</sup><sup>1</sup> *Brookhaven National Laboratory (US)*<sup>2</sup> *TRIUMF (CA)***Corresponding Authors:** yesw@bnl.gov, asoka.de.silva@cern.ch, undrus@bnl.gov

The ATLAS offline code management system serves as a collaborative framework for developing a code base totaling more than 5 million lines. Supporting up to 50 nightly release branches, the ATLAS Nightly System offers abundant opportunities for updating existing software and developing new tools for forthcoming experimental stages within a multi-platform environment. This paper describes the utilization of container technology for the ATLAS nightly jobs. By conducting builds and tests of offline releases within containers, we ensure portability across various build nodes. The controlled container environment enhances stability by removing dependencies on operating system updates. Furthermore, it sets the base and facilitates the production of containerized software across different user activity areas and pipelines. The ATLAS experiment has accumulated data since 2009. It is important to maintain access to software for processing and analyzing historical data developed on outdated operating systems. Container technology plays an indispensable role in providing secure and operationally sound environments for building and testing on such operating systems. This document provides details on the organizational support for OS containers used in software building, including methods for setting up runtime environments.

**Poster session / 74**

## ATLAS analysis workflows using the EventIndex and the Event Picking Server for massive event picking and enhanced processing

**Authors:** Aleksandr Iakovlev<sup>1</sup>; Aleksandra Dimitrievska<sup>2</sup>; Dario Barberis<sup>3</sup>; Evgeny Alexandrov<sup>1</sup>; Igor Alexandrov<sup>1</sup>; Semen Turchikhin<sup>4</sup>; Simone Pagan Griso<sup>5</sup><sup>1</sup> *Joint Institute for Nuclear Research (RU)*<sup>2</sup> *University of Birmingham (GB)*<sup>3</sup> *Università e INFN Genova (IT)*<sup>4</sup> *INFN e Università Genova (IT)*<sup>5</sup> *Lawrence Berkeley National Lab. (US)***Corresponding Authors:** alexand@jinr.ru, alexander.iakovlev@cern.ch, simone.pagan.griso@cern.ch, evgeny.alexandrov@cern.ch, aleksandra.dimitrievska@cern.ch, dario.barberis@cern.ch, semen.turchikhin@cern.ch

The ATLAS detector produces a wealth of information for each recorded event. Standard calibration and reconstruction procedures reduce this information to physics objects that can be used as input to most analyses; nevertheless, there are very specific analyses that need full information from some of the ATLAS subdetectors, or enhanced calibration and/or reconstruction algorithms. For these use cases, a novel workflow has been developed that involves the selection of events satisfying some basic criteria, their extraction in RAW data format using the EventIndex data catalogue and the Event Picking Server, and their specialised processing. This workflow allows us in addition to commission and use new calibration and reconstruction techniques before launching the next full reprocessing (important given the longer and longer expected time between full reprocessing campaigns), to use algorithms and tools that are too CPU or disk intensive if run over all recorded events, and in the future to apply AI/ML methods that start from low-level information and could profit from rapid development/use cycles. This presentation describes the tools involved, the procedures followed and the current operational performance.

**Poster session / 75**

## Recent Advances in the GAN-based Fast Calorimeter Simulation of the ATLAS Experiment

**Authors:** Federico Andrea Corchia<sup>1</sup>; Lorenzo Rinaldi<sup>1</sup><sup>1</sup> *Universita e INFN, Bologna (IT)***Corresponding Authors:** lorenzo.rinaldi@cern.ch, corchia@bo.infn.it

Simulation of the detector response is a major computational challenge in modern High-Energy Physics experiments, accounting for about 40% of the total computational resources used in ATLAS. The simulation of the calorimeter response is particularly demanding, consuming about 80% of the total simulation time.

In order to make the best use of the available computational resources, fast simulation tools based on Machine Learning techniques have been developed to simulate the calorimeter response faster than Geant4 while maintaining a high level of accuracy. One such tool, developed by the ATLAS Collaboration and currently in production for LHC Run 3, is FastCaloGAN, which uses Generative Adversarial Networks (GANs) to generate electromagnetic and hadronic showers.

To facilitate the training and optimisation of the GANs, and to enable a more efficient use of computational resources, a container-based system, FastCaloGANtainer, facilitates the deployment of the FastCaloGAN training on complementary high-performance resources such as High Performance Computing (HPC) farms and ensures its operational independence from the underlying system.

This talk presents the latest developments in FastCaloGAN and FastCaloGANtainer, discussing their technical details and recent improvements in terms of Physics and computational performance. For FastCaloGAN, these improvements include an improved voxelisation and extension to further use cases (e.g. particle types not yet covered), while for FastCaloGANtainer they concern its deployment on a wider variety of resources with multi-CPU/GPU nodes and different architectures (including cutting-edge HPC clusters such as Leonardo at CINECA in Bologna, Italy).

**Parallel (Track 5) / 76**

## Advancements in the ATLAS Fast Chain for HL-LHC: Towards Efficient MC Production

**Author:** Fang-Ying Tsai<sup>1</sup><sup>1</sup> *Stony Brook University (US)***Corresponding Author:** fang-ying.tsai@cern.ch

The ATLAS Fast Chain represents a significant advancement in streamlining Monte Carlo (MC) production efficiency, specifically for the High-Luminosity Large Hadron Collider (HL-LHC). This project aims to simplify the production of Analysis Object Data (AODs) and potentially Derived Analysis Object Data (DAODs) from generated events with a single transform, facilitating rapid reproduction of the entire MC dataset multiple times per year. By eliminating intermediate formats and optimizing CPU utilization, the Fast Chain offers substantial savings in disk space while staying within the CPU budget by employing fast simulation methodologies instead of full MC campaigns. Central to the success of the Fast Chain is the seamless integration of fast simulation and reconstruction techniques. Leveraging AtlFast3 methodologies for efficient calorimeter shower simulation and employing Fast Track Simulation (FATRAS) for charged particles in the Inner Detector, the project aims at accelerated processing without compromising accuracy. Notably, muon simulations rely on Geant4 due to minimal CPU overhead. Pileup effects are incorporated through MC overlay, with potential future integration of data overlay. Reconstruction speed optimization focuses on Inner Detector track reconstruction. Strategies such as dedicated reconstruction configurations and track overlay from pre-mixed pileup datasets are being explored. In summary, the ATLAS Fast Chain project demonstrates a paradigm shift in MC production methodologies, offering a scalable and efficient solution tailored to the demands of the HL-LHC era. This abstract provides an overview

of the project's objectives, methodologies, and ongoing developments, showcasing its potential to revolutionize MC production within the ATLAS experiment.

Parallel (Track 5) / 78

## **Towards Machine-Learning Particle Flow with the ATLAS Detector at the LHC**

**Author:** Luca Clissa<sup>1</sup>

<sup>1</sup> *Universita e INFN, Bologna (IT)*

**Corresponding Author:** luca.clissa@cern.ch

Particle flow reconstruction at colliders combines various detector subsystems (typically the calorimeter and tracker) to provide a combined event interpretation that utilizes the strength of each detector. The accurate association of redundant measurements of the same particle between detectors is the key challenge in this technique. This contribution describes recent progress in the ATLAS experiment towards utilizing machine-learning to improve particle flow in the ATLAS detector. In particular, point-cloud techniques are utilized to associate measurements from the same particle, leading to reduced confusion compared to baseline techniques. Next steps towards further testing and implementation will be discussed.

Poster session / 79

## **New approaches for fast and efficient graph construction on CPU, GPU and heterogeneous architectures for the ATLAS event reconstruction**

**Author:** Christophe COLLARD<sup>1</sup>

<sup>1</sup> *Laboratoire des 2 Infinis - Toulouse, CNRS / Univ. Paul Sabatier*

**Corresponding Author:** christophe.collard@l2it.in2p3.fr

Graph neural networks (GNN) have emerged as a cornerstone of ML-based reconstruction and analysis algorithms in particle physics. Many of the proposed algorithms are intended to be deployed close to the beginning of the data processing chain, e.g. in event reconstruction software of running and future collider-based experiments. For GNN to operate, the input data are represented as graphs. The creation of the graphs and the associated cost are often limiting factors in high-throughput production environments. We discuss the specific example of charged-particle track reconstruction in the ATLAS detector. The HL-LHC upgrade of the ATLAS detector brings an unprecedented track reconstruction challenge, both in terms of the large number of silicon hit cluster readouts, and the throughput required. The GNN4ITk project has designed GNN-based algorithms for tracking with a similar level of physics performance to traditional techniques, that scale sub-quadratically, provided that the large input graphs can be created efficiently. In this contribution, we present novel methods that are able to produce these graphs quickly and efficiently, and describe their computing performance.

Parallel (Track 1) / 80

## **ATLAS High-Luminosity LHC demonstrators with Data Carousel: Data-on-Demand and Tape Smart Writing**

**Authors:** Alexei Klimentov<sup>1</sup>; Doris Ressmann<sup>2</sup>; Haykuhi Musheghyan<sup>2</sup>; Julien Leduc<sup>3</sup>; Mario Lassnig<sup>3</sup>; Misha Borodin<sup>4</sup>; Tadashi Maeno<sup>1</sup>; Tatiana Korchuganova<sup>5</sup>; Xin Zhao<sup>1</sup>

<sup>1</sup> *Brookhaven National Laboratory (US)*

<sup>2</sup> *Karlsruhe Institute of Technology (KIT)*

<sup>3</sup> *CERN*

<sup>4</sup> *University of Iowa (US)*

<sup>5</sup> *University of Pittsburgh (US)*

**Corresponding Authors:** mario.lassnig@cern.ch, tatiana.korchuganova@cern.ch, doris.ressmann@kit.edu, tmaeno@bnl.gov, haykuhi.musheghyan@cern.ch, julien.leduc@cern.ch, alexei.klimentov@cern.ch, xzhao@bnl.gov, mikhail.borodin@cern.ch

The High Luminosity upgrade to the LHC (HL-LHC) is expected to generate scientific data on the scale of multiple exabytes. To tackle this unprecedented data storage challenge, the ATLAS experiment initiated the Data Carousel project in 2018. Data Carousel is a tape-driven workflow in which bulk production campaigns with input data resident on tape are executed by staging and promptly processing a sliding window to disk buffer such that only a small fraction of the input files are pinned on disk at any one time. Put in ATLAS production before Run3, Data Carousel continues to be our focus for seeking new opportunities in disk space savings, and enhancing tape usage throughout the ATLAS Distributed Computing (ADC) environment. These efforts are highlighted by two recent ATLAS HL-LHC demonstrator projects: data-on-demand and tape smart writing. We will discuss the recent studies and outcomes from these projects, along with various related improvements across the ATLAS distributed computing software. The research was conducted together with site experts at CERN and Tier-1 centers.

Parallel (Track 5) / 81

## Optimizing the ATLAS Geant4 detector simulation

**Author:** Benjamin Morgan<sup>1</sup>

<sup>1</sup> *University of Warwick (GB)*

**Corresponding Author:** ben.morgan@warwick.ac.uk

The ATLAS experiment at the LHC heavily depends on simulated event samples produced by a full Geant4 detector simulation. This Monte Carlo (MC) simulation based on Geant4 is a major consumer of computing resources and is anticipated to remain one of the dominant resource users in the HL-LHC era. ATLAS has continuously been working to improve the computational performance of this simulation for the Run 3 MC campaign. This update highlights the implementation of recent and upcoming optimizations. These improvements include enhancements to the core Geant4 software, strategic choices in simulation configuration, simplifications in geometry and magnetic field descriptions, and technical refinements in the interface between ATLAS simulation code and Geant4. Overall, these improvements have resulted in a more than 100% increase in throughput compared to the baseline simulation configuration utilized during Run 2.

Parallel (Track 5) / 82

## Using and Visualizing Graphs and Graph Algorithms

**Author:** Julius Hrivnac<sup>1</sup>

<sup>1</sup> *Université Paris-Saclay (FR)*

**Corresponding Author:** julius.hrivnac@cern.ch

Representing HEP and astrophysics data as graphs (i.e. networks of related entities) is becoming increasingly popular. These graphs are not only useful for structuring data storage but are also increasingly utilized within various machine learning frameworks.

However, despite their rising popularity, numerous unused opportunities exist, particularly concerning the utilization of graph algorithms and intuitive visualization techniques.

This presentation will introduce a comprehensive graph framework designed for handling the HEP and astronomical data. The framework supports the storage, manipulation and analyses of the graph data, facilitating the use of elementary graph algorithms. Additionally, it enables the export of graph data to specialized external toolkit for more sophisticated processing and analysis.

An integral feature of the presented framework is its highly interactive, web-based graphical front-end. This interface provides users with deep insights into the graph structures of their data, enabling interactive analysis and multi-faceted visualization of graph properties. It also offers integration capabilities with other related frameworks.

The practical application of this framework will be demonstrated through its use in analyzing relationships between astronomical alerts, specifically from the Zwicky Transient Facility (ZTF) and the Rubin Observatory. By leveraging the collective properties and relationships within these data, the framework facilitates comprehensive analyses and provides recommendations based on object similarities and neighborhood characteristics. This approach paves the way for novel insights and methodologies in approach to data.

**Parallel (Track 3) / 83**

## **Future scheduling in Athena on Accelerators and Multiple Nodes**

**Authors:** Beojan Stanislaus<sup>1</sup>; Charles Leggett<sup>2</sup>; Julien Esseiva<sup>1</sup>; Paolo Calafiura<sup>1</sup>; Vakho Tsulaia<sup>1</sup>; Xiangyang Ju<sup>1</sup>

<sup>1</sup> Lawrence Berkeley National Lab. (US)

<sup>2</sup> Lawrence Berkeley National Lab (US)

**Corresponding Authors:** charles.g.leggett@gmail.com, beojan.stanislaus@cern.ch, xiangyang.ju@cern.ch, vakhtang.tsulaia@cern.ch, julien.esseiva@cern.ch, pcalafiura@lbl.gov

The large increase in luminosity expected from Run 4 of the LHC presents the ATLAS experiment with a new scale of computing challenge, and we can no longer restrict our computing to CPUs in a High Throughput Computing paradigm. We must make full use of the High Performance Computing resources available to us, exploiting accelerators and making efficient use of large jobs over many nodes.

Here we show our current developments in introducing these capabilities to Athena, ATLAS's general software framework. We will show how we have used MPI to distribute processing over multiple nodes, and how this can be used to run real ATLAS jobs from the Grid on an HPC. We will also show how we have integrated a first-class capability to offload work to an accelerator without blocking the CPU, by making use of suspendable lightweight threads, and an example of how this capability can be used in a real workload.

**Parallel (Track 5) / 84**

## **Towards an experiment-independent toolkit for fast calorimeter simulation**

**Authors:** Jana Schaarschmidt<sup>1</sup>; Joshua Falco Beirer<sup>2</sup>; Marilena Bandieramonte<sup>3</sup>; Rui Zhang<sup>4</sup>

<sup>1</sup> *University of Washington (US)*

<sup>2</sup> *CERN*

<sup>3</sup> *University of Pittsburgh (US)*

<sup>4</sup> *University of Wisconsin Madison (US)*

**Corresponding Authors:** jana.schaarschmidt@cern.ch, joshua.beirer@cern.ch, marilena.bandieramonte@cern.ch, rui.zhang@cern.ch

For high-energy physics experiments, the generation of Monte Carlo events, and in particular the simulation of the detector response, is a very computationally intensive process. In many cases, the primary bottleneck in detector simulation is the detailed simulation of the electromagnetic and hadronic showers in the calorimeter system. For the ATLAS experiment, about 80% of the total CPU usage for detector simulation is devoted to the simulation of the secondary particles produced in the calorimeter system.

To make best use of the available resources, ATLAS is currently using its state-of-the-art fast simulation tool AtlFast3, which uses classical parametric and ML-based solutions for shower generation. AtlFast3 is deployed in a heterogeneous simulation infrastructure known as the Integrated Simulation Framework (ISF), which was originally developed over a decade ago, and is becoming increasingly difficult to maintain by the collaboration. In an effort to greatly simplify its simulation infrastructure for Run 4 and beyond, the collaboration is in the process of phasing out ISF and implementing its fast simulation library directly as a Geant4 fast simulation model.

In addition, efforts have started to develop a fully experiment-independent library for fast calorimeter simulation, providing a universal interface for the lateral and longitudinal parameterisation of calorimeter shower development, as well as for ML-based approaches to shower generation.

This talk will give an overview of the current and future (fast) simulation infrastructure in ATLAS. The new experiment-independent library for fast calorimeter simulation will be presented and its use for other experiments will be motivated and discussed.

**Parallel (Track 5) / 85**

## **AtlFast3: Fast Simulation in ATLAS for LHC Run 3 and beyond**

**Authors:** Jana Schaarschmidt<sup>1</sup>; Joshua Falco Beirer<sup>2</sup>; Marilena Bandieramonte<sup>3</sup>; Rui Zhang<sup>4</sup>

<sup>1</sup> *University of Washington (US)*

<sup>2</sup> *CERN*

<sup>3</sup> *University of Pittsburgh (US)*

<sup>4</sup> *University of Wisconsin Madison (US)*

**Corresponding Authors:** marilena.bandieramonte@cern.ch, jana.schaarschmidt@cern.ch, joshua.beirer@cern.ch, rui.zhang@cern.ch

As we are approaching the high-luminosity era of the LHC, the computational requirements of the ATLAS experiment are expected to increase significantly in the coming years. In particular, the simulation of MC events is immensely computationally demanding, and their limited availability is one of the major sources of systematic uncertainties in many physics analyses. The main bottleneck in the detector simulation is the detailed simulation of electromagnetic and hadronic showers in the ATLAS calorimeter system using Geant4.

In order to increase the MC statistics and to leverage the available CPU resources for LHC Run 3, the ATLAS collaboration has recently put into production a refined and significantly improved version of its state-of-the-art fast simulation tool AtlFast3. AtlFast3 uses classical parametric and machine

learning based approaches such as Generative Adversarial Networks (GANs) for the fast simulation of LHC events in the ATLAS detector.

This talk will present the newly improved version of AtlFast3 that is currently in production for the simulation of Run 3 samples. In addition, ideas and plans for the future of fast simulation in ATLAS will also be discussed.

#### Plenary session / 86

## Exploring Client-Server Scalability with RNTuple & EOS: Comparative Analysis of Physics Data Formats for Analysis

**Authors:** Andrea Sciabà<sup>1</sup>; Andreas Joachim Peters<sup>1</sup>; Guilherme Amadio<sup>1</sup>; Jakob Blomer<sup>1</sup>; Jonas Hahnfeld<sup>2</sup>; Markus Schulz<sup>1</sup>; Vincenzo Eduardo Padulano<sup>1</sup>

**Co-authors:** Alaettin Serhan Mete<sup>3</sup>; Danilo Piparo<sup>1</sup>; Florine de Geus<sup>4</sup>; Matti Kortelainen<sup>5</sup>; Philippe Canal<sup>5</sup>

<sup>1</sup> CERN

<sup>2</sup> CERN & Goethe University Frankfurt

<sup>3</sup> Argonne National Laboratory (US)

<sup>4</sup> CERN/University of Twente (NL)

<sup>5</sup> Fermi National Accelerator Lab. (US)

**Corresponding Authors:** vincenzo.eduardo.padulano@cern.ch, amadio@cern.ch, markus.schulz@cern.ch, jakob.blomer@cern.ch, florine.de.geus@cern.ch, andrea.sciaba@cern.ch, andreas.joachim.peters@cern.ch, alaettin.serhan.mete@cern.ch, philippe.canal@cern.ch, jonas.hahnfeld@cern.ch, matti.kortelainen@cern.ch, danilo.piparo@cern.ch

Collaboratively, the IT and EP departments have launched a formal project within the Research and Computing sector to evaluate a novel data format for physics analysis data utilized in LHC experiments and other fields. The objective of this initiative is to substitute the current TTree data format of ROOT with a more efficient format known as RNTuple, which provides superior support for multi-threading and enhanced compression capabilities. This aspect of the project focuses on verifying the scalability of the storage back-end EOS during the migration from the old to the new format, utilizing replicated and erasure coded profiles.

#### Parallel (Track 5) / 87

## Quantum error mitigation for Fourier moments computation

**Author:** Oriel Orphee Moira Kiss<sup>1</sup>

**Co-authors:** Alessandro Roggero<sup>2</sup>; Michele Grossi<sup>3</sup>

<sup>1</sup> Universite de Geneve (CH)

<sup>2</sup> University of Trento

<sup>3</sup> CERN

**Corresponding Authors:** oriel.kiss@cern.ch, a.roggero@unitn.it, michele.grossi@cern.ch

Hamiltonian moments in Fourier space - expectation values of the unitary evolution operator under a Hamiltonian at different times - provide a convenient framework to understand quantum systems. They offer insights into the energy distribution, higher-order dynamics, response functions, correlation information and physical properties. This paper focuses on the computation of Fourier moments within the context of a nuclear effective field theory on superconducting quantum hardware.



The study integrates echo verification and noise renormalization into Hadamard tests using control reversal gates. These techniques, combined with purification and error suppression methods, effectively address quantum hardware decoherence. The analysis, conducted using noise models, reveals a significant reduction in noise strength by two orders of magnitude. Moreover, quantum circuits involving up to 266 CNOT gates over five qubits demonstrate high accuracy under these methodologies when run on IBM superconducting quantum devices.

Parallel (Track 7) / 88

## Diving into large-scale congestion with NOTED as a network controller and machine learning-based traffic forecasting

**Author:** Maria Del Carmen Misa Moreira<sup>1</sup>

**Co-author:** Edoardo Martelli<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** edoardo.martelli@cern.ch, carmen.misa.moreira@cern.ch

The Network Optimised Experimental Data Transfer (NOTED) has undergone successful testing at several international conferences, including the International Conference for High Performance Computing, Networking, Storage and Analysis (also known as SuperComputing). It has also been tested at scale during the WLCG Data Challenge 2024, in which NREN's and WLCG sites conducted testing at 25% of the rates foreseen for the HL-LHC. During these events, NOTED has demonstrated its ability to detect network congestion and dynamically reconfigure the network by executing actions, thereby enhancing network capacity. Recently, the integration of NOTED with the CERN's Network Monitoring System has increased its ability to detect and respond to congestion in the LHCOPN (Tier 0 to Tier 1's links) and LHCONE (Tier 1's to Tier 2's links) networks. We report here on NOTED's enhanced ability to identify congested WLCG sites and DC24 experiences with network reconfiguration to alleviate the detected congestion. Previous work has demonstrated the feasibility of improving NOTED's ability to predict network traffic using machine learning with LSTM (Long Short-Term Memory) networks, given its capacity to learn from historical data. We present here new findings on the beneficial impact of various machine learning approaches, including encoders, transformers, and other algorithms, on the NOTED's performance in relation to traffic forecasting.

Parallel (Track 1) / 89

## Scitags: A Standardized Framework for Traffic Identification and Network Visibility in Data-Intensive Research Infrastructures

**Authors:** Andrew Hanushevsky<sup>1</sup>; Dale Carder<sup>None</sup>; Garhan Attebury<sup>2</sup>; Marian Babik<sup>3</sup>; Shawn Mc Kee<sup>4</sup>; Tim Chown<sup>None</sup>; Tristan Sullivan<sup>5</sup>

**Co-authors:** Andrew Lake ; Bruno Heinrich Hoefft<sup>6</sup>; James Letts<sup>7</sup>; Joe Mambretti<sup>8</sup>; Karl Newell ; Michael Lambert<sup>9</sup>

<sup>1</sup> Stanford University/SLAC

<sup>2</sup> University of Nebraska-Lincoln

<sup>3</sup> CERN

<sup>4</sup> University of Michigan (US)

<sup>5</sup> University of Victoria

<sup>6</sup> KIT - Karlsruhe Institute of Technology (DE)

<sup>7</sup> Univ. of California San Diego (US)

<sup>8</sup> *International Center for Advanced Internet Research Northwestern University*

<sup>9</sup> *Pittsburgh Supercomputing Center*

**Corresponding Authors:** lambert@psc.edu, andy@es.net, attebury@cse.unl.edu, knewell@internet2.edu, shawn.mckee@cern.ch, marian.babik@cern.ch, bruno.hoeft@kit.edu, letts@physics.ucsd.edu, abh@slac.stanford.edu, j-mambretti@northwestern.edu, tim.chown@jisc.ac.uk, tssulliv@uvic.ca, dwarder@es.net

High-Energy Physics (HEP) experiments rely on complex, global networks to interconnect collaborating sites, data centers, and scientific instruments. Managing these networks for data-intensive scientific projects presents significant challenges because of the ever-increasing volume of data transferred, diverse project requirements with varying quality of service needs, multi-domain infrastructure, WAN distances, and limited visibility into network traffic flows. This lack of visibility hinders network operators' ability to understand actual user behavior across different network segments, optimize performance, undertake effective traffic engineering and shaping, and effectively debug and troubleshoot issues.

This project addresses these challenges by focusing on improving network visibility through standardized packet marking and flow labeling techniques. We present the Scitags initiative, a collaborative effort formed within the Research Networking Technical Working Group (RNTWG) in 2020. Scitags aims to develop a generic framework and standards for identifying the owner and associated scientific activity of network traffic. This framework extends beyond HEP/WLCG experiments, and it has a potential to benefit all global communities using Research and Education (R&E) networks.

The presentation will detail the current state of the Scitags initiative, including the evolving framework, the underlying technologies being explored (e.g., eBPF, IPv6, HbH, etc.), and the roadmap for production deployment within R&E networks. By enabling improved network visibility, Scitags will empower network operators to optimize performance, troubleshoot issues more effectively, and ultimately support the growing needs of data-intensive scientific collaborations.

**Poster session / 90**

## Machine Learning Inference in Athena with ONNXRuntime

**Authors:** Attila Krasznahorkay<sup>1</sup>; Beojan Stanislaus<sup>2</sup>; Charles Leggett<sup>3</sup>; Johannes Elmsheuser<sup>4</sup>; Julien Esseiva<sup>2</sup>; Paolo Calafiura<sup>2</sup>; Vakho Tsulaia<sup>2</sup>; Xiangyang Ju<sup>2</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *Lawrence Berkeley National Lab. (US)*

<sup>3</sup> *Lawrence Berkeley National Lab (US)*

<sup>4</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** johannes.elmsheuser@cern.ch, xiangyang.ju@cern.ch, beojan.stanislaus@cern.ch, charles.g.leggett@gmail.com, pcalafiura@lbl.gov, attila.krasznahorkay@cern.ch, julien.esseiva@cern.ch, vakhtang.tsulaia@cern.ch

Machine Learning (ML)-based algorithms play increasingly important roles in almost all aspects of data processing in the ATLAS experiment at CERN. Diverse ML models are used in detector simulation, event reconstruction, and data analysis. They are being deployed in the ATLAS software framework, Athena. Our primary approach to perform ML inference in Athena is to use ONNXRuntime. ONNXRuntime is a cross-platform ML model acceleration library, with a flexible interface to integrate hardware-specific libraries. In this talk, we will describe the ONNXRuntime interface in Athena and the impact of advanced ONNXRuntime settings on various ML models and workflows at ATLAS.

**Parallel (Track 3) / 91**

## AthenaTriton: A Tool for running Machine Learning Inference as a Service in Athena

**Authors:** Beojan Stanislaus<sup>1</sup>; Charles Leggett<sup>2</sup>; Haoran Zhao<sup>3</sup>; Julien Esseiva<sup>1</sup>; Paolo Calafiura<sup>1</sup>; Shih-Chieh Hsu<sup>4</sup>; Vakho Tsulaia<sup>1</sup>; Xiangyang Ju<sup>1</sup>; Yuan-Tang Chou<sup>3</sup>

<sup>1</sup> Lawrence Berkeley National Lab. (US)

<sup>2</sup> Lawrence Berkeley National Lab (US)

<sup>3</sup> University of Washington (US)

<sup>4</sup> University of Washington Seattle (US)

**Corresponding Authors:** vakhtang.tsulaia@cern.ch, schsu@uw.edu, yuan-tang.chou@cern.ch, haoran.zhao@cern.ch, charles.g.leggett@gmail.com, xiangyang.ju@cern.ch, pcalafiura@lbl.gov, julien.esseiva@cern.ch, beojan.stanislaus@cern.ch

Machine Learning (ML)-based algorithms play increasingly important roles in almost all aspects of the data analyses in ATLAS. Diverse ML models are used in detector simulations, event reconstructions, and data analyses. They are being deployed in the ATLAS software framework, Athena. The primary approach to perform ML inference in Athena is to use the ONNXRuntime. However, some ML models could not be converted to ONNXRuntime because certain ML operations, such as the MultiAggregation in pyG as of writing, are not supported. Furthermore, a scalable inference strategy that maximises the event processing throughput is needed to cope with the ever-increasing simulation and collision data. A key element in that strategy should be enabling these ML algorithms to run on coprocessors like GPUs because not all computing sites have coprocessors. To that end, we introduce AthenaTriton, a tool that runs ML inference as a service based on the NVIDIA Triton Inference Server. With AthenaTriton, we give Athena the capability to act as a Triton client that sends requests to a remote or local server that performs the model inference. We will present the AthenaTriton design and its scalability in running ML-based algorithms. We emphasise that AthenaTriton can be used in both online and offline computing.

Poster session / 93

## The Electronics Simulation Software in JUNO

**Author:** 张豪森 zhanghaosen<sup>None</sup>

**Corresponding Author:** zhanghaosen@ihep.ac.cn

The Jiangmen Underground Neutrino Observatory (JUNO), located in Southern China, is a multi-purpose neutrino experiment that consists of a central detector, a water Cherenkov detector and a top tracker. The primary goal of the experiment is to determine the neutrino mass ordering (NMO) and precisely measure neutrino oscillation parameters. The central detector contains 20,000 ton liquid scintillator and is instrumented with 17,612 20-inch PMTs and 25,600 3-inch PMTs for anti-neutrino detection with an energy resolution of 3% at 1MeV. The electronics simulation is the crucial module of JUNO offline software (JUNOSW). It takes the photoelectron information from Geant4 based detector simulation as input to simulate the PMT response, trigger logic and electronics response of sub-detectors using an implementation based on SNIPEr managed dynamically-loadable elements(DLE). Electronics simulation incorporates a “hit-level” event mixing implementation which combines different event types with different rates that mimic the data streaming of real experimental data. The event mixing uses a “pull” based workflow using SNIPEr incident schema. The electronics simulation outputs become inputs to the online event classification algorithms (OEC) used for event tagging and then saved to file using ROOT I/O services. In this talk, a detailed introduction of the electronics simulation software will be presented.

Parallel (Track 5) / 94

## Building a Columnar Analysis Demonstrator for ATLAS PHYS-LITE Open Data using the Python Ecosystem

**Authors:** Alexander Held<sup>1</sup>; Gordon Watts<sup>2</sup>

<sup>1</sup> *University of Wisconsin Madison (US)*

<sup>2</sup> *University of Washington (US)*

**Corresponding Authors:** gwatts@uw.edu, alexander.held@cern.ch

The ATLAS experiment is in the process of developing a columnar analysis demonstrator, which takes advantage of the Python ecosystem of data science tools. This project is inspired by the analysis demonstrator from IRIS-HEP.

The demonstrator employs PHYSLITE OpenData from the ATLAS collaboration, the new Run 3 compact ATLAS analysis data format. The tight integration of ROOT features within PHYSLITE presents unique challenges when integrating with the Python analysis ecosystem. The demonstrator is constructed from ATLAS PHYSLITE OpenData, ensuring the accessibility and reproducibility of the analysis.

The analysis pipeline of the demonstrator incorporates a comprehensive suite of tools and libraries. These include uproot for data reading, awkward-array for data manipulation, Dask for parallel computing, and hist for histogram processing. For the purpose of statistical analysis, the pipeline integrates cabinetry and pyhf, providing a robust toolkit for analysis. A significant component of this project is the custom application of corrections, scale factors, and systematic errors using ATLAS software. The infrastructure and methodology for these applications will be discussed in detail during the presentation, underscoring the adaptability of the Python ecosystem for high-energy physics analysis.

**Parallel (Track 9) / 95**

## **A Pilot Analysis Facility at CERN, Architecture, Implementation and First Evaluation**

**Authors:** Andrea Sciabà<sup>1</sup>; Ben Jones<sup>1</sup>; Diogo Castro<sup>1</sup>; Enric Tejedor Saavedra<sup>1</sup>; Enrique Garcia Garcia<sup>1</sup>; Giovanni Guerrieri<sup>1</sup>; Markus Schulz<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** diogo.castro@cern.ch, markus.schulz@cern.ch, andrea.sciaba@cern.ch, enric.tejedor.saavedra@cern.ch, giovanni.guerrieri@cern.ch, ben.dylan.jones@cern.ch, enrique.garcia.garcia@cern.ch

Experiment analysis frameworks, physics data formats and expectations of scientists at the LHC have been evolving towards interactive analysis with short turnaround times. Several sites in the community have reacted by setting up dedicated Analysis Facilities, providing tools and interfaces to computing and storage resources suitable for interactive analysis. It is expected that this demand will increase towards the HL-LHC era and will be only met by scaling out to allow interactive processing of large datasets.

CERN IT launched a Pilot of an Analysis Facility based on established, proven services such as SWAN, HTCondor and EOS. This facilitates the access to massive resources by enabling the use of HTCondor managed resources from SWAN, offering parallel execution via frameworks such as ROOT RDataFrame and Coffea and their Dask backends.

In this contribution we will discuss the architecture of the Pilot Analysis Facility at CERN, giving the rationale for the decisions. For deciding on the next steps the evaluation of the impact of different resource allocation strategies at the CERN HTCondor pool is critical. One especially interesting strategy consists in combining a set of dedicated resources for interactive analysis with the use of the general resources that are subject to experiment quotas. Concrete performance and scalability measurements will be presented. We will put a special focus on the feedback we received from the early testers from the experiments.

Poster session / 97

## interTwin - an interdisciplinary Digital Twin Engine for Science

**Authors:** Matteo Bunino<sup>1</sup>; Kalliopi Tsolaki<sup>1</sup>; Alexander Zoechbauer<sup>1</sup>; Maria Girone<sup>1</sup>; Sofia Vallecorsa<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** alexander.zoechbauer@cern.ch, matteo.bunino@cern.ch, maria.girone@cern.ch, sofia.vallecorsa@cern.ch, kalliopi.tsolaki@cern.ch

The interTwin project, funded by the European Commission, is at the forefront of leveraging ‘Digital Twins’ across various scientific domains, with a particular emphasis on physics and earth observation. One of the most advanced use-cases of interTwin is event generation for particle detector simulation at CERN. interTwin enables particle detector simulations to leverage AI methodologies on cloud to high-performance computing (HPC) resources by using itwinai - the AI workflow and method lifecycle module of interTwin.

The itwinai module, a comprehensive solution for AI workflow and method lifecycle developed collaboratively by CERN and the Jülich Supercomputing Center (JSC), serves as the cornerstone for researchers, data scientists, and software engineers engaged in developing, training, and maintaining AI-based methods for scientific applications, such as the particle event generation. Its role is advancing interdisciplinary scientific research through the synthesis of learning and computing paradigms. This framework stands as a testament to the commitment of the interTwin project towards co-designing and implementing an interdisciplinary Digital Twin Engine. Its main functionalities and contributions are:

**Distributed Training:** itwinai offers a streamlined approach to distributing existing code across multiple GPUs and nodes, automating the training workflow. Leveraging industry-standard backends, including PyTorch Distributed Data Parallel (DDP), TensorFlow distributed strategies, and Horovod, it provides researchers with a robust foundation for efficient and scalable distributed training. The successful deployment and testing of itwinai on JSC’s HDFML cluster underscore its practical applicability in real-world scenarios.

**Hyperparameter Optimization:** One of the core functionalities of itwinai is its hyperparameter optimization, which plays a crucial role in enhancing model accuracy. By intelligently exploring hyperparameter spaces, itwinai eliminates the need for manual parameter tuning. The functionality, empowered by RayTune, contributes significantly to the development of more robust and accurate scientific models.

**Model Registry:** A key aspect of itwinai is its provision of a robust model registry. This feature allows researchers to log and store models along with associated performance metrics, thereby enabling comprehensive analyses in a convenient manner. The backend, leveraging MLFlow, ensures seamless model management, enhancing collaboration and reproducibility.

In line with the “Computing infrastructure” track of CHEP 2024, interTwin and its use-cases empowered by itwinai are positioned at the convergence of computation and physics and showcase the significant potential of AI research supported by HPC resources. Together, they contribute to a narrative of interconnected scientific frontiers, where the integration of digital twins, AI frameworks, and physics research broadens possibilities for exploration and discovery through itwinai’s user-friendly interface and powerful functionalities.

In conclusion, itwinai is a valuable and versatile resource, empowering researchers and scientists to embark on collaborative and innovative scientific research endeavors across diverse domains. The integration of physics-based digital twins and AI frameworks broadens possibilities for exploration and discovery through itwinai’s user-friendly interface and powerful functionalities.

Parallel (Track 5) / 98

## Jet reconstruction with quantum algorithms

**Author:** Hideki Okawa<sup>1</sup>

<sup>1</sup> *Chinese Academy of Sciences (CN)*

**Corresponding Author:** hideki.okawa@cern.ch

Jets are key observables to measure the hadronic activities at high energy colliders such as the Large Hadron Collider (LHC) and future colliders such as the High Luminosity LHC (HL-LHC) and the Circular Electron Positron Collider (CEPC). Yet jet reconstruction is a computationally expensive task especially when the number of final-state particles is large. Such a clustering task can be regarded as an optimization problem, which can be formulated in terms of an Ising Hamiltonian and searching for its ground state would provide the answer. Quantum algorithms such as the quantum approximate optimization algorithm (QAOA) provide promising solutions to tackle the problem. Performance of jet reconstruction is presented for both the quantum hardware and simulator.

**Parallel (Track 2) / 99**

## Improvements of the GPU Processing Framework for ALICE

**Author:** David Rohr<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Author:** drohr@jwtdt.org

ALICE is the dedicated heavy ion experiment at the LHC at CERN and records lead-lead collisions at a rate of up to 50 kHz.

The detector with the highest data rate of up to 3.4 TB/s is the TPC.

ALICE performs the full online TPC processing corresponding to more than 95% of the total workload on GPUs, and when there is no beam in the LHC, the online computing farm's GPUs are used to speed up the offline processing.

After the deployment of the first version of the online TPC processing needed for data taking, ALICE has implemented many improvements to its GPU processing framework.

These include a run time compilation mode applying on the fly optimizations, improvements to parallelize / speed up the GPU compilation, debugging modes to guarantee reproducible and deterministic results in concurrent reconstruction, and framework features to leverage common components in the code of different detectors.

The talk will give an overview of the ALICE experience with GPUs in online and offline processing and present the latest GPU processing framework features.

**Parallel (Track6) / 100**

## Enhancing software-hardware co-design for HEP by low-overhead profiling of single- and multi-threaded programs on diverse architectures with AdaptivePerf

**Author:** Maksymilian Graczyk<sup>1</sup>

**Co-author:** Stefan Roiser<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** maksymilian.graczyk@cern.ch, stefan.roiser@cern.ch

Given the recent slowdown of the Moore's Law and increasing awareness of the need for sustainable and edge computing, physicists and software developers can no longer just rely on computer hardware becoming faster and faster or moving processing to the cloud to meet the ever-increasing computing demands of their research (e.g. the data rate increase in HL-LHC). However, algorithmic optimisations alone are also starting to be insufficient, so novel computing paradigms spanning both software and hardware appear. Adapting existing and new software to them may be difficult though, especially for large and complex applications. This is where profiling can help bridge the gap, but finding a suitable profiler is challenging when a low overhead, wide architectural support, and reliability are important.

As a response to the above problem, AdaptivePerf was developed. It is an open-source, architecture-portable, and low-overhead profiling tool with custom-patched Linux perf as its main foundation, currently available on GitHub. Thanks to the extensive research and modifications, AdaptivePerf improves the main shortcomings of perf such as incomplete stack traces. It profiles how threads and processes are created within a program and what code segments within each thread/process should be considered on- or off-CPU bottlenecks, in terms of both consumed time and other hardware metrics like cache misses. If a user-friendly visualisation is needed, AdaptivePerf can present results as a timeline with the process tree, where corresponding non-time-ordered and time-ordered flame graphs can be browsed along with functions spawning new threads/processes.

The tool has already been shown to work on x86-64 and RISC-V and is designed in the context of the SYCLOPS EU project, which CERN is part of and where solutions for heterogeneous architectures are developed, e.g. custom RISC-V cores tailored to a specific problem, RISC-V support for SYCL, and SYCL-accelerated algorithms in ROOT. In this presentation, we will talk about the profiler, its place within the project, and how it can be used for software-hardware co-design for HEP.

Parallel (Track 1) / 101

## Evolution of the CERN Tape Archive scheduling system

**Author:** Jaroslav Guenther<sup>1</sup>

**Co-authors:** David Smith<sup>1</sup>; Eric Cano<sup>1</sup>; Joao Afonso<sup>1</sup>; Julien Leduc<sup>1</sup>; Konstantina Skovola; Michael Davis<sup>1</sup>; Niels Alexander Bugel; Pablo Oliver Cortes<sup>1</sup>; Richard Bachmann<sup>1</sup>; Sergio Alfageme Perez; Vladimir Bahyl<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** richard.bachmann@cern.ch, jaroslav.guenther@cern.ch, julien.leduc@cern.ch, david.smith@cern.ch, sergio.alfageme@cern.ch, michael.davis@cern.ch, vladimir.bahyl@cern.ch, pablo.oliver.cortes@cern.ch, joao.afonso@cern.ch, eric.cano@cern.ch, konstantina.skovola@cern.ch, niels.alexander.bugel@cern.ch

The CERN Tape Archive (CTA) scheduling system implements the workflow and lifecycle of Archive, Retrieve and Repack requests. The transient metadata for queued requests is stored in the Scheduler backend store (Scheduler DB). In our previous work, we presented the CTA Scheduler together with an objectstore-based implementation of the Scheduler DB. Now with four years of experience in production, the strengths and limitations of this implementation are better understood. While the objectstore-based implementation is highly efficient for FIFO queueing operations (archive/retrieve), non-FIFO operations (delete, priority queues) require some workarounds. The objectstore backend implementation imposes constraints on how the CTA Scheduler code can be modified and is an additional software dependency and technology for developers to learn. This paper discusses an alternate Scheduler DB implementation, based on relational database technology. We include a status report and roadmap.

Parallel (Track 1) / 102

## Challenges of repack in the era of the high-capacity tape cartridge

**Authors:** Jaroslav Guenther<sup>1</sup>; Joao Afonso<sup>1</sup>; Jorge Camarero Vera<sup>1</sup>; Julien Leduc<sup>1</sup>; Konstantina Skovola<sup>1</sup>; Michael Davis<sup>1</sup>; Niels Buegel<sup>1</sup>; Pablo Oliver Cortes<sup>1</sup>; Richard Bachmann<sup>1</sup>; Vladimir Bahyl<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** niels.alexander.bugel@cern.ch, richard.bachmann@cern.ch, joao.afonso@cern.ch, michael.davis@cern.ch, jorge.camarero@cern.ch, vladimir.bahyl@cern.ch, konstantina.skovola@cern.ch, jaroslav.guenther@cern.ch, pablo.oliver.cortes@cern.ch, julien.leduc@cern.ch

The latest tape hardware technologies (LTO-9, IBM TS1170) impose new constraints on the management of data archived to tape. In the past, new drives could read the previous one or even two generations of media, but this is no longer the case. This means that repacking older media to new media must be carried out on a more aggressive schedule than in the past. An additional challenge is the large capacity of the newer media. A 50 TB tape can contain a vast number of files, whose metadata must be tracked during repacking. Repacking an entire tape also requires a significant amount of disk storage. At CERN Tier-0, these challenges have created new operational problems to solve, in particular contention for resources between physics archival and repack operations. This contribution details these problems and describes the various approaches we have taken to mitigate and solve them. We include a roadmap for future repack developments.

Parallel (Track 9) / 103

## Navigating the Multilingual Landscape of Scientific Computing: Python, Julia, and Awkward Array

**Authors:** Ianna Osborne<sup>1</sup>; Jim Pivarski<sup>1</sup>

<sup>1</sup> Princeton University

**Corresponding Authors:** ianna.osborne@cern.ch, pivarski@princeton.edu

Scientific computing relies heavily on powerful tools like Julia and Python. While Python has long been the preferred choice in High Energy Physics (HEP) data analysis, there's a growing interest in migrating legacy software to Julia. We explore language interoperability, focusing on how Awkward Array data structures can connect Julia and Python. We discuss memory management, data buffer copies, and dependency handling, highlighting performance gains from invoking Julia from Python and vice versa. Particularly, we look into distributed array-oriented calculations involving large-scale HEP data and a unique role of Awkward Array in these workflows. We examine the advantages and challenges of achieving interoperability between Julia and Python in scientific computing.

Poster session / 104

## Scaling TraceWin beam dynamics simulations on Kubernetes for Reinforcement Learning training

**Authors:** Davide Marcato<sup>1</sup>; Daniel Lupu<sup>2</sup>

**Co-authors:** Luca Bellan<sup>2</sup>; Massimo Biasotto<sup>3</sup>; Michele Comunian; Sergio Fantinel<sup>4</sup>; Michele Gulmini<sup>5</sup>; Lisa Zangrando<sup>4</sup>

<sup>1</sup> Istituto Nazionale di Fisica Nucleare

<sup>2</sup> INFN-LNL

<sup>3</sup> INFN Legnaro

<sup>4</sup> Universita e INFN, Padova (IT)



<sup>5</sup> *Universita e INFN, Legnaro (IT)*

**Corresponding Authors:** daniel.lupu@lnl.infn.it, luca.bellan@lnl.infn.it, sergio.fantinel@lnl.infn.it, lisa.zangrando@pd.infn.it, massimo.biasotto@lnl.infn.it, michele.comunian@lnl.infn.it, michele.gulmini@cern.ch, davide.marcato@lnl.infn.it

Reinforcement Learning is emerging as a viable technology to implement autonomous beam dynamics setup and optimization in particle accelerators. A Deep Learning agent can be trained to efficiently explore the parameter space of an accelerator control system and converge to the optimal beam setup much faster than traditional methods. Training these models requires programmatic execution of a high volume of simulations. This contribution introduces pytracewin, a Python wrapper of the TraceWin beam dynamics simulator, which exposes simple methods to run simulations and retrieve results. It can be easily combined with the large Python ecosystem of Machine Learning and Reinforcement Learning libraries to develop optimization models. Still, the training process is computationally constrained by the number of simulations that can be run in a reasonable time. It is thus crucial to scale such workload on a dedicated computing infrastructure while retaining a simple high-level user interface.

We exploit Dask, an open-source library, to enable embarrassingly parallel execution of TraceWin simulations on Kubernetes, using a dynamically scalable number of workers and requiring minimal user code modifications. Workers are instantiated with a custom docker image combining Dask and pytracewin. The approach is validated using two Kubernetes clusters on INFN Cloud and Cloud-Veneto to simulate the ADIGE beam line at Legnaro National Laboratories.

**Parallel (Track6) / 105**

## Thoroughly testing and integrating hundreds of Pull Requests per month: ROOT's new Cost-efficient and Feature Rich GitHub-based CI

**Author:** Danilo Piparo<sup>1</sup>

**Co-authors:** Axel Naumann<sup>1</sup>; Malik Shahzad Muzaffar<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** axel.naumann@cern.ch, danilo.piparo@cern.ch, shahzad.malik.muzaffar@cern.ch

ROOT is an open source framework, freely available on GitHub, at the heart of data acquisition, processing and analysis of HE(N)P experiments, and beyond.

It is developed collaboratively: contributions are not authored only by ROOT team members, but also by a veritable nebula of developers and scientists from universities, labs as well as the private sector. More than 1500 GitHub Pull Requests are merged on average per year. It is in this context that code integration acquires a primary role: not only code contributions need to be reviewed, but they need to be thoroughly tested through a powerful CI infrastructure on several different platforms to comply with the high code quality standards of the project. Since the end of 2023, ROOT moved its continuous integration system from a Jenkins one to a GitHub Actions based one.

In this contribution, we characterise the transition to the GitHub CI, focussing our strategy, its implementation and the lesson learned, as well as the advantages the new system offers with respect to the previous one. Particular emphasis will be given to the evaluation of the cost-benefit ratio for Jenkins and GitHub Actions for the ROOT project. We'll also describe how we manage to run in less than one hour thousands of unit, integration, functional and end-to-end tests on different flavours of Windows, four versions of macOS, as well as about ten of the most used Linux distributions, taking advantage of the CERN computing infrastructure.

**Parallel (Track 9) / 106**

## Distributed analysis in production with RDataFrame

**Author:** Marta Czurylo<sup>1</sup>

**Co-authors:** Danilo Piparo<sup>1</sup>; Vincenzo Eduardo Padulano<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** vincenzo.eduardo.padulano@cern.ch, marta.maja.czurylo@cern.ch, danilo.piparo@cern.ch

The ROOT software package provides the data format used in High Energy Physics by the LHC experiments. It offers a data analysis interface called RDataFrame, which has proven to adapt well to the requirements of modern physics analyses. However, with increasing data collected by the LHC experiments, the challenge to perform an efficient analysis expands. One of the solutions to ease this challenge, is the leverage of modern high performing distributed computing environments for which RDataFrame provides an easy-to-use interface layer - the distributed RDataFrame.

In this talk, we show that the Distributed RDataFrame is out of the experimental testing phase, and it is now ready for production thanks to a stabilized user interface. We delve into recent improvements of the distributed RDataFrame, including Pythonizations of the interface that allow running the workflows seamlessly (for example, with the XGBoost library). As the variety and geographical locations of distributed environments are available, we show the reproducibility and compare the performance across several of them.

**Plenary session / 107**

## The ROOT Project at the end of Run 3 and towards HL-LHC

**Author:** Danilo Piparo<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Author:** danilo.piparo@cern.ch

In this contribution, we'll review the status of the ROOT project towards the end of LHC Run 3. We'll review its structure, available effort and management strategy, allowing to push innovation while guaranteeing long term support. In particular, we'll describe how ROOT became a veritable community effort attracting contributions not only from the ROOT team, but from collaborators at labs, universities and the private sector thanks to its open source philosophy. We'll review the main features of the 2025 data taking release 6.32 as well as the forthcoming development release. We'll also discuss the evolution of ROOT towards the HL-LHC era, in its 7th release cycle.

**Parallel (Track 7) / 108**

## Integration of the Goettingen HPC cluster Emmy to the WLCG Tier-2 centre GoeGrid and performance tests

**Authors:** Arnulf Quadt<sup>1</sup>; Daniel Schindler<sup>1</sup>; Sebastian Wozniowski<sup>1</sup>; Uday Saidev Polisetty<sup>1</sup>

<sup>1</sup> *Georg August Universitaet Goettingen (DE)*

**Corresponding Authors:** daniel.schindler@uni-goettingen.de, uday.saidev.polisetty@cern.ch, sebastian.wozniowski@cern.ch, arnulf.quadt@cern.ch

The German university-based Tier-2 centres successfully contributed a significant fraction of the computing power required for Runs 1-3 of the LHC. But for the upcoming Run 4, with its increased

need for both storage and computing power for the various HEP computing tasks, a transition to a new model becomes a necessity. In this context, the German community under the FIDIUM project is making interdisciplinary resources of the National High Performance Computing (NHR) usable within the WLCG and centralising mass storage at the Helmholtz centres.

The Goettingen campus hosts both a WLCG Tier-2 site, GoeGrid, and the HPC cluster Emmy that is part of the National High-Performance Computing (NHR) center NHR-Nord@Göttingen. The integration is done by virtually extending the GoeGrid batch system with containers, turning the HPC nodes into virtual worker nodes with their own partitionable job scheduling in order to run GoeGrid HEP jobs for the ATLAS collaboration. Submission and management of these containers are automated using COBalD (the Opportunistic Balancing Daemon) and TARDIS (The Transparent Adaptive Resource Dynamic Integration System). Data are provided via the GoeGrid mass storage for which a dedicated network connection has been established. Continuous production of ATLAS jobs is currently being tested in a one-year pilot phase. The setup, experience, performance tests and outlook are presented.

Parallel (Track 5) / 109

## Parameter Estimation in ATLAS with Neural Simulation-Based Inference

**Author:** Aishik Ghosh<sup>1</sup>

<sup>1</sup> *University of California Irvine (US)*

**Corresponding Author:** aishik.ghosh@cern.ch

Neural Simulation-Based Inference (NSBI) is a powerful class of machine learning (ML)-based methods for statistical inference that naturally handle high dimensional parameter estimation without the need to bin data into low-dimensional summary histograms. Such methods are promising for a range of measurements at the Large Hadron Collider, where no single observable may be optimal to scan over the entire theoretical phase space under consideration, or where binning data into histograms could result in a loss of sensitivity. This work develops an NSBI framework that, for the first time, allows NSBI to be applied to a full-scale LHC analysis, by successfully incorporating a large number of systematic uncertainties, quantifying the uncertainty coming from finite training statistics, developing a method to construct confidence intervals, and demonstrating a series of intermediate diagnostic checks that can be performed to validate the robustness of the method. As an example, the power and feasibility of the method are demonstrated for an off-shell Higgs boson couplings measurement in the four lepton decay channel, using simulated samples. The proposed method is a generalisation of the standard statistical framework at the LHC, and can benefit a large number of physics analyses. This work serves as a blueprint for measurements at the LHC using NSBI.

Parallel (Track 1) / 110

## ML-based Adaptive Prefetching and Data Placement for US HEP systems

**Author:** Venkat Sai Suman Lamba Karanam<sup>1</sup>

**Co-authors:** Byrav Ramamurthy<sup>1</sup>; Derek Weitzel<sup>2</sup>; Sarat Sasank Barla<sup>1</sup>

<sup>1</sup> *University of Nebraska-Lincoln*

<sup>2</sup> *University of Nebraska Lincoln (US)*

**Corresponding Authors:** ramamurthy@unl.edu, saisuman@huskers.unl.edu, sbarla2@huskers.unl.edu, derek.weitzel@cern.ch

Although caching-based efforts [1] have been in place in the LHC infrastructure in the US, we show that integrating intelligent prefetching and targeted dataset placement into the underlying caching strategy can improve job efficiency further. Newer experiments and experiment upgrades such as HL-LHC and DUNE are expected to produce 10x the amount of data than currently being produced. This means that the available network, storage and compute resources must be utilized efficiently. Additionally, newer generations of DAQs are moving towards streaming readout systems, navigating away from the traditional triggered systems. These newer DAQ systems offer continuous/real-time data calibration, reconstruction and storage by offloading to remote sites results [2, 3]. These observations imply that the available network, storage and compute resources must be used efficiently. The benefits in CPU efficiency and job turnaround times from colocating the datasets near computation are well-known, especially using dedicated or opportunistic cache storages using XCache or dCache systems [4]. Our prior work using the data transfer logs collected from the OSG dashboard revealed two major observations. First, there is a correlation between transfer time and the choice of storage site. The choice of source site, in case of storage redundancy, was found to be more important in transfer time than the actual file size (the dataset consists of many files). Second, there is not only a popularity skew in the remote files accessed by the jobs, but also files that are part of the same dataset are read more than others in that dataset.

<https://lh3.googleusercontent.com/drive-viewer/AKGpihZWfXLIPkKNZzPjOgiAymVeXWtwcXtXoO33uDYMsbrR5PJvJhL-8mAVoHFuu3SkMk-6JTpLRBDlDCWLXuF8Zyg=s2560?source=screenshot.guru> >< <https://lh3.googleusercontent.com/drive-viewer/AKGpihZWfXLIPkKNZzPjOgiAymVeXWtwcXtXoO33uDYMsbrR5PJvJhL-8mAVoHFuu3SkMk-6JTpLRBDlDCWLXuF8Zyg=s2560> / >< /a >

File size vs. Transfer Time for data chunks with markers colored by the Data Source. The data in the figures are collected over 24-hour period in 2019. There are two clearly demarcated groups of transfer time values. Mid-Top left in figures show Group 1, a group of transfer with low transfer times. For this group, the files are served by a single data source (identified by dark yellow color). The second group, Group 2 (see the bottom group of data points) consists of smaller file sizes, and they show a wide distribution of transfer times.

<https://lh3.googleusercontent.com/drive-viewer/AKGpihbPia71bJWtxq3TIuqafRALmtXmyrK64kAzCELSzvubEFZTRtkU=s2560?source=screenshot.guru> >< <https://lh3.googleusercontent.com/drive-viewer/AKGpihbPia71bJWtxq3TIuqafRALmtXmyrK64kAzCJzmbOcwU77nbWn98UFCWgxnfJk69zvGELSVzvbEFZTRtkU=s2560> / >< /a >

File size vs. Time of the Day plot of a single day from March of 2020 on the OSG. Left-side plot spans a single day in the March of 2020 and the right-side is a 2.5-hour snapshot of the same day. Each star represents a single file transfer. The color of the star represents a unique source from where the file was transferred from.

Using our analysis of the experiment and data pipelines, and the data access patterns in the US HEP environment, we present intelligent Machine Learning (ML)-based approaches to reduce the latency and improve job efficiency. We incorporate our prefetching and data placement techniques into a simulator by extending the WRENCH [5] simulator to reflect the experiment and data workflows typical in the US HEP infrastructure. The simulator is designed to reflect the US HEP environment (data, storage and computation workflows, and infrastructure). Simulators have long had an established history in planning and development of HEP infrastructure like WLCG (MONARC simulator) [6]. Simulations make it feasible to compare different network, storage and compute settings without building real testbeds for each setting. We present the proof-of-concept of our simulator with tight agreements to the real-world behavior of the experiments in US, specifically those belonging to LHC (CMS etc.) and DUNE.

## REFERENCES

1. Fajardo, Edgar, Derek Weitzel, Mats Rynge, Marian Zvada, John Hicks, Mat Selmecci, Brian Lin et al. "Creating a content delivery network for general science on the internet backbone using XCaches." In EPJ Web of Conferences, vol. 245, p. 04041. EDP Sciences, 2020.

2. Lawrence D. Streaming Readout and Remote Compute. Thomas Jefferson National Accelerator Facility (TJNAF), Newport News, VA (United States); 2023 May 1.
3. Suiu, Alice-Florența. “EPN2EOS Data Transfer System.” PhD diss., University POLITEHNICA of Bucharest, 2023.
4. Acosta-Silva C, Casals J, Peris AD, Molina JF, Hernández JM, Pérez CM, Dengra CP, Yzquierdo AP, Rodríguez FJ. A case study of content delivery networks for the CMS experiment.
5. Casanova, Henri, Suraj Pandey, James Oeth, Ryan Tanaka, Frédéric Suter, and Rafael Ferreira Da Silva. “Wrench: A framework for simulating workflow management systems.” In 2018 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS), pp. 74-85. IEEE, 2018.
6. Iosif C. Legrand and Harvey B. Newman. “The MONARC Toolset for Simulating Large Network-Distributed Processing Systems”. Proceedings of the 32nd Conference on Winter Simulation. WSC '00. Orlando, Florida: Society for Computer Simulation International, 2000, pp. 1794–1801. ISBN: 0780365828.

Parallel (Track 5) / 111

## The Multi-Threaded Detector Simulation in JUNO

**Author:** Peidong Yu<sup>1</sup>

<sup>1</sup> IHEP

**Corresponding Author:** yupd@ihep.ac.cn

The Jiangmen Underground Neutrino Observatory (JUNO) is a multi-purpose experiment under construction in southern China. JUNO is designed to determine the mass ordering of neutrinos and precisely measure neutrino oscillation parameters by detecting reactor neutrinos from the Yangjiang and Taishan Nuclear Power Plants. Atmospheric neutrinos, solar neutrinos, geo-neutrinos, supernova burst neutrinos and DSNB (Diffuse Supernova Neutrino Background), nucleon decay can also be studied with JUNO. The main detector of JUNO is a 20,000-ton liquid scintillator detector. The JUNO detector simulation software is a key component of the JUNO offline software (JUNOSW), developed based on the SNIPEr framework. Due to the large size of the detector and broad range of energies of interest, detector simulation poses challenges in terms of CPU time and memory consumption. With computing nodes gradually incorporating multiple integrated CPU cores, traditional single-threaded computing models may result in significant memory usage and inefficient system resource utilization. Implementing multithreading processing on many-core architectures can significantly improve system resource utilization efficiency. This report will introduce the design and implementation status of multi-threaded detector simulation in JUNOSW.

Poster session / 112

## EDM4hep.jl: Analysing EDM4hep files with Julia

**Author:** Pere Mato<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** pere.mato@cern.ch

EDM4hep aims to establish a standard event data model for the store and exchange of event data in HEP experiments, thereby fostering collaboration across various experiments and analysis frameworks. The Julia package EDM4hep.jl is capable of generating Julia-friendly structures for the EDM4hep data model and reading event data files in ROOT format (either TTree or RNTuple) that are written by C++ programs, utilising the UnROOT.jl package. This paper explores the motivations behind the primary design choices of this package, such as the exclusive use of structure of arrays (SoA) to access the stored collections, which then empower users to develop ergonomic data analyses using Julia's high-level concepts and functionality, while maintaining performance comparable to C++

programs. Several examples are given to illustrate how efficient data analysis can be achieved using high-level objects, eliminating the need to resort to flat n-tuples.

**Parallel (Track 5) / 113**

## **On-the-fly data set joins and concatenations with ROOT RNTuple**

**Authors:** Florine de Geus<sup>1</sup>; Vincenzo Eduardo Padulano<sup>2</sup>; Jakob Blomer<sup>2</sup>; Philippe Canal<sup>3</sup>; Ana-Lucia Varbanescu<sup>4</sup>

<sup>1</sup> *CERN/University of Twente (NL)*

<sup>2</sup> *CERN*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

<sup>4</sup> *University of Twente*

**Corresponding Authors:** vincenzo.eduardo.padulano@cern.ch, philippe.canal@cern.ch, a.l.varbanescu@utwente.nl, jakob.blomer@cern.ch, florine.de.geus@cern.ch

With the large data volume increase expected for HL-LHC and the even more complex computing challenges set by future colliders, the need for more elaborate data access patterns will become more pressing. ROOT's next-generation data format and I/O subsystem, RNTuple, is designed to address those challenges, currently already showing a clear improvement in storage and I/O efficiency with respect to its predecessor, TTree. These improvements provide a solid baseline to introduce extensions that directly target common HENP workflow features not easily achievable before. Notably, many workflows benefit from the ability to join and concatenate data sets during application runtime, with the aim to reduce overall storage requirements and improve application ergonomics. The successful implementation of such compositions requires taking several factors into careful consideration, especially for large data sets that do not fit in memory. These factors include the transparent handling of (in)compatibility between different data sets, the rules that determine how data set compositions are processed, and their effects on runtime performance. In this contribution, we will present the ongoing work to support advanced composition of RNTuple data sets. We will discuss the main design considerations through a selection of concrete workflow use cases, the interfaces and internal machinery that enable the compositions, and an initial set of performance evaluation results.

**Plenary session / 114**

## **CMS FlashSim: end-to-end simulation with ML**

Detailed event simulation at the LHC is taking a large fraction of computing budget. CMS developed an end-to-end ML based simulation that can speed up the time for production of analysis samples of several orders of magnitude with a limited loss of accuracy. As the CMS experiment is adopting a common analysis level format, the NANO AOD, for a larger number of analyses, such an event representation is used as the target of this ultra fast simulation that we call FlashSim. Generator level events, from PYTHIA or other generators, are directly translated into NANO AOD events at several hundred Hz rate with FlashSim. We show how training FlashSim on a limited number of full simulation events is sufficient to achieve very good accuracy on larger datasets for processes not seen at training time. Comparisons with full simulation samples in some simplified benchmark analysis are also shown. With this work, we aim at establishing a new paradigm for LHC collision simulation workflows in view of HL-LHC.

**Parallel (Track 3) / 115**

## A Multi-Objective Optimization Tool for Track Reconstruction in CMS

**Author:** Simone Rossi Tisbeni<sup>1</sup><sup>1</sup> *Universita e INFN, Bologna (IT)*

Efficient and precise track reconstruction is critical for the results of the Compact Muon Solenoid (CMS) experiment. The current CMS track reconstruction algorithm is a multi-step procedure based on the combinatorial Kalman filter as well as a Cellular Automaton technique to create track seeds. Multiple parameters regulate the reconstruction steps, populating a large phase space of possible solutions. The fine-tuning of these parameters is necessary to ensure an optimal reconstruction. The CMS tracker featured robust performance and efficient tracking in Run 3 condition, however, the High Lumi environment is expected to be much more demanding. The upgrade will lead to higher rates and pile-up that require further improvement in all the reconstruction processes, with more complex algorithms featuring additional parameters. Alternative techniques to help the experts in properly tuning these environments are thus being investigated.

This report presents an original tool based on the established Particle Swarm heuristic optimization algorithm (PSO) to perform parameter tuning of the pixel track reconstruction software currently employed in the CMS experiment. The software enables multi-objective optimization against tracking efficiency and fake rate, resulting in the individuation of a Pareto front of valid parameters' sets for reconstruction.

The algorithm has been tested at the end of the data-taking period of 2023 on the pixel track reconstruction algorithm with excellent results. The parameters obtained with the optimization resulted in comparable reconstruction efficiency for both phase 1 data and phase 2 simulations, with a 50% improvement in fake rates, especially for low transverse momentum of the particles.

Further research and development can explore the application of this tool to other aspects of the CMS reconstruction process. Additionally, investigating the integration of this tool within the existing CMS framework can streamline the optimization workflow for future data-taking periods.

**Parallel (Track 2) / 117**

## Adoption of the alpaka performance portability library in the CMS software

**Author:** Andrea Bocci<sup>1</sup><sup>1</sup> *CERN*

To achieve better computational efficiency and exploit a wider range of computing resources, the CMS software framework (CMSSW) has been extended to offload part of the physics reconstruction to NVIDIA GPUs. To support additional back-ends, as well to avoid the need to write, validate and maintain a separate implementation of the reconstruction algorithms for each back-end, CMS has adopted the Alpaka performance portability library.

Alpaka (Abstraction Library for Parallel Kernel Acceleration) is a header-only C++ library that provides performance portability across different back-ends, abstracting the underlying levels of parallelism. It supports serial and parallel execution on CPUs, and extremely parallel execution on NVIDIA, AMD and Intel GPUs.

This contribution will show how Alpaka is used in the CMS software to develop and maintain a single code base; to use different toolchains to build the code for each supported back-end, and link them into a single application; to seamlessly select the best back-end at runtime, and implement

portable reconstruction algorithms that run efficiently on CPUs and GPUs from different vendors. It will describe the validation and deployment of the Alpaka-based implementation in the CMS High Level Trigger, and highlight how it achieves near-native performance.

**Plenary session / 118**

## Challenges of the CMS Level-1 Trigger for the HL-LHC Upgrade

**Authors:** CMS Collaboration<sup>None</sup>; Santiago Folgueras<sup>1</sup>

<sup>1</sup> *Universidad de Oviedo (ES)*

**Corresponding Author:** borislav.pavlov@cern.ch

For the High-Luminosity Large Hadron Collider era, the trigger and data acquisition system of the Compact Muon Solenoid experiment will be entirely replaced. Novel design choices have been explored, including ATCA prototyping platforms with SoC controllers and newly available interconnect technologies with serial optical links with data rates up to 28 Gb/s. Trigger data analysis will be performed through sophisticated algorithms, including widespread use of Machine Learning, in large FPGAs, such as the Xilinx Ultrascale family. The system will process over 50 Tb/s of detector data with an event rate of 750 kHz. The talk will discuss the technological and algorithmic aspects of the upgrade of the CMS trigger system, emphasizing the use of low-latency Machine Learning and AI algorithms with several examples.

**Parallel (Track 5) / 119**

## Simulating the CMS High Granularity Calorimeter with ML

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

Detector simulation is a key component of physics analysis and related activities in CMS. In the upcoming High Luminosity LHC era, simulation will be required to use a smaller fraction of computing in order to satisfy resource constraints. At the same time, CMS will be upgraded with the new High Granularity Calorimeter (HGCAL), which requires significantly more resources to simulate than the existing CMS calorimeters. This computing challenge motivates the use of generative machine learning models as surrogates to replace full physics-based simulation. We study the application of state-of-the-art diffusion models to simulate particle showers in the CMS HGCAL. We will discuss methods to overcome the challenges posed by the high-dimensional, irregular geometry of the HGCAL. The quality of the showers produced by the diffusion model will be assessed by comparison to the full GEANT4-based simulation. The increase in simulation throughput will be quantified and methods to accelerate the diffusion model inference will also be discussed.

**Parallel (Track 3) / 122**

## Pile-up estimation with Graph Neural Networks for LHC Experiment

**Author:** Mahtab Jalal Vandi<sup>1</sup>

**Co-authors:** Hamed Bakhshiansohi<sup>1</sup>; Saleh z



<sup>1</sup> *Isfahan University of Technology (IR)*

**Corresponding Authors:** hamed.bakhshian@cern.ch, saleh.zrh@gmail.com, mahtab.jalal.vandi@cern.ch

High-energy physics experiments, like the CMS experiment at the Large Hadron Collider (LHC), face growing challenges in simulating and analyzing the pile-up of simultaneous proton-proton collisions as luminosity levels increase with LHC upgrades. This paper presents a novel approach for pile-up estimation that leverages the power of Graph Neural Networks (GNNs). GNNs have emerged as a powerful tool for representing complex event data, as shown in recent applications for particle reconstruction and pile-up mitigation. We propose a GNN-based method to estimate the number of pile-up interactions within each event. Preliminary results are promising, indicating that the method can be used to study the linearity response of the luminometers, a critical step for luminosity measurement. This research explores the transformative potential of GNNs for pile-up estimation and luminosity measurement in high-energy physics experiments. We expect this method to become a valuable addition to existing techniques, contributing to advancements in the field.

Parallel (Track 2) / 123

## Neural network clusterization for the ALICE TPC online computing

**Author:** Christian Sonnabend<sup>1</sup>

<sup>1</sup> *CERN, Heidelberg University (DE)*

**Corresponding Author:** christian.sonnabend@cern.ch

The ALICE Time Projection Chamber (TPC) is the detector with the highest data rate of the ALICE experiment at CERN and is the central detector for tracking and particle identification. Efficient online computing such as clusterization and tracking are mainly performed on GPU's with throughputs of approximately 900 GB/s. Clusterization itself has a well known background with a variety of algorithms in the field of machine learning. This work investigates a neural network approach to cluster rejection and regression on a topological basis. Central to its task are the center-of-gravity, sigma and total charge estimation as well as rejection of clusters in the TPC readout. Additionally, a momentum vector estimate is made from the 3D input across readout rows in combination with reconstructed tracks which can benefit track seeding. Performance studies on inference speed as well as model architectures and physics performance on Monte-Carlo data can be presented, showing that tracking performance can be maintained while rejecting 5-10% of raw clusters with a O(30%) reduced fake-rate for clusterization itself compared to the current GPU clusterizer.

Parallel (Track 9) / 124

## Data discovery, analysis and reproducibility in Virtual Research Environments

**Author:** Enrique Garcia Garcia<sup>1</sup>

**Co-authors:** Giovanni Guerrieri<sup>1</sup>; Hugo Gonzalez Labrador<sup>1</sup>; Xavier Espinal<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** enrique.garcia.garcia@cern.ch, giovanni.guerrieri@cern.ch, hugo.gonzalez.labrador@cern.ch, xavier.espinal@cern.ch

During the ESCAPE project, the pillars of a pilot analysis facility were built following a bottom-up approach, in collaboration with all the partners of the project. As a result, the CERN Virtual Research Environment (VRE) initiative proposed a workspace that facilitates the access to the data

in the ESCAPE Data Lake, a large scale data management system defined by Rucio, along with the interactive analysis that jupyter notebooks enable. The VRE also provisions a variety of scientific software stacks via CVMFS, and can be connected to local data processing resources through REANA. The latter is an open source software, developed within the CERN IT department, that provides a framework focussed on the reanalysis and reproducibility of scientific results. The CERN VRE has deployed an instance of REANA, allowing users to make use of the platform functionalities together with the rest of the services in the analysis facility.

Having a single interface that integrates different services with the underlying infrastructure certainly eases the user experience. Furthermore, in line with the ESCAPE Open Collaboration, the development of open source tools that can be reused in different physics communities with similar analysis strategies would lay the foundation of common lifecycle analysis practices. Therefore, in order to foster accessibility, as well as interactivity and reproducibility to more complex infrastructure services, the development of user-friendly middleware should be prioritized.

This contribution focuses on the connection of REANA to the CERN VRE's interface through a Jupyter extension. The development of this extension makes it possible to use the VRE as a single workspace to enhance the lifecycle of a research analysis: from discovery and data access, through interactive analysis and offload to computing resources, to reproducibility of results.

**Parallel (Track 3) / 125**

## **Extending ALICE's GPU tracking capabilities: Towards a comprehensive accelerated barrel reconstruction**

**Author:** Matteo Concas<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** [matteo.concas@cern.ch](mailto:matteo.concas@cern.ch)

During Run 3, ALICE has enhanced its data processing and reconstruction chain by integrating GPUs, a leap forward in utilising high-performance computing at the LHC.

The initial 'synchronous' phase engages GPUs to reconstruct and compress data from the TPC detector. Subsequently, the 'asynchronous' phase partially frees GPU resources, allowing further offloading of additional reconstruction tasks to enhance efficiency. Notably, ITS tracking has been ported as an independent module for two major GPU platforms.

This presentation will detail the integration of ITS GPU tracking within the existing framework, aiming to develop a unified GPU-based reconstruction pipeline.

This pipeline minimises memory transfer latency by coordinating various simultaneous processing stages.

Performance metrics of the integrated system will be discussed, highlighting the technical strategies and outcomes of this implementation.

**Parallel (Track 8) / 126**

## **Leveraging public cloud resources for the processing of CMS open data**

**Author:** Kati Lassila-Perini<sup>1</sup>

**Co-author:** CMS Collaboration

<sup>1</sup> *Helsinki Institute of Physics (FI)*

**Corresponding Authors:** kati.lassila-perini@cern.ch, borislav.pavlov@cern.ch

The CMS experiment at the Large Hadron Collider (LHC) regularly releases open data and simulations, enabling a wide range of physics analyses and studies by the global scientific community. The recent introduction of the NanoAOD data format has provided a more streamlined and efficient approach to data processing, allowing for faster analysis turnaround. However, the larger MiniAOD format retains richer information that may be crucial for certain research endeavors.

To ensure the long-term usability of CMS open data to their full extent, this work explores the potential of leveraging public cloud resources for the computationally intensive processing of the MiniAOD format. Many open data users may not have access to the necessary computing resources for handling the large MiniAOD datasets. By offloading the heavy lifting to scalable cloud infrastructure, researchers can benefit from increased processing power and improved overall efficiency in their data analysis workflows, with a moderate short-term cost.

The study investigates best practices and challenges for effectively utilizing public cloud platforms to handle the processing of CMS MiniAOD data, with a focus on quantifying the overall time and cost of using these resources. The ultimate aim is to empower the CMS open data community to maximize the scientific impact of this valuable resource.

**Parallel (Track 5) / 128**

## Madgraph on GPUs and vector CPUs: towards production

**Authors:** Andrea Valassi<sup>1</sup>; Taylor Childers<sup>None</sup>; Stephan Hageboeck<sup>1</sup>; Olivier Mattelaer<sup>2</sup>; Nathan Nichols<sup>3</sup>; Filip Zdzislaw Optolowicz<sup>4</sup>; Stefan Roiser<sup>1</sup>; Jorgen Teig<sup>5</sup>; Zenny Jovi Joestar Wettersten<sup>1</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *UCLouvain*

<sup>3</sup> *Argonne National Laboratory*

<sup>4</sup> *University of Wrocław*

<sup>5</sup> *Norwegian University of Science and Technology (NTNU) (NO)*

**Corresponding Authors:** filip.optolowicz@cern.ch, nnichols@anl.gov, jchilders@anl.gov, stephan.hageboeck@cern.ch, stefan.roiser@cern.ch, olivier.mattelaer@uclouvain.be, jorgen.teig@cern.ch, zenny.wettersten@cern.ch, andrea.valassi@cern.ch

The effort to speed up the Madgraph5\_aMC@NLO generator by exploiting CPU vectorization and GPUs, which started at the beginning of 2020, is expected to deliver the first production release of the code for QCD leading-order (LO) processes in 2024. To achieve this goal, many additional tests, fixes and improvements have been carried out by the development team in recent months, both to carry out its internal workplan and to respond to the feedback from the LHC experiments about the current and required functionalities of the software. Several new physics processes, including both Standard Model and Beyond Standard Model calculations, have been tested and extensively debugged. Support for AMD GPUs via native HIP has been added to the CUDA/C++ baseline implementation of the code; work is in progress to also add support for Intel GPUs to this CUDA/C++ plugin, based on the parallel SYCL implementation developed in the past. The user interface and packaging of the software, and the usability challenges coming from the large number of events that must be generated in parallel on a GPU, have also been an active area of development. In this contribution, we will report on these activities and on the status of the LO software at the time of the CHEP2024 conference. The status and outlook for one of the main further directions of our development effort, notably the support of next-to-leading-order (NLO) processes, is described in a separate contribution to this conference.

**Parallel (Track 3) / 132**

## TGeoArbN - A ROOT compatible triangle mesh geometry implementation

**Authors:** Ben Salisbury<sup>1</sup>; Christoph Schmidt<sup>1</sup>; Ulrike Thoma<sup>1</sup>

<sup>1</sup> *HISKP Bonn*

**Corresponding Authors:** s6besali@uni-bonn.de, schmidt@hiskp.uni-bonn.de, thoma@hiskp.uni-bonn.de

To increase the automation to convert Computer-Aided-Design detector components as well as entire detector systems into simulatable ROOT geometries, TGeoArbN, a ROOT compatible geometry class, was implemented allowing the use of triangle meshes in VMC-based simulation. To improve simulation speed a partitioning structure in form of an Octree can be utilized. TGeoArbN in combination with a CADToROOT-Converter (based on [1]) allowed e.g. for a high level of automation for the conversion of the forward endcap geometry of the PANDA electromagnetic calorimeter. The aim of the talk is to give an overview on TGeoArbN and the modified CADToROOT-Converter version.

[1] T. Stockmanns, "STEP-to-ROOT –from CAD to Monte Carlo Simulation", Journal of Physics: Conference Series 396 (2012) 022050, url: <https://doi.org/10.1088/1742-6596/396/2/022050>

**Parallel (Track 4) / 133**

## Whole-node scheduling in the ALICE Grid: Initial experiences and evolution opportunities

**Author:** Marta Bertran Ferrer<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Author:** marta.bertran.ferrer@cern.ch

JAliEn, the ALICE experiment's Grid middleware, utilizes whole-node scheduling to maximize resource utilization from participating sites. This approach offers flexibility in resource allocation and partitioning, allowing for customized configurations that adapt to the evolving needs of the experiment. This scheduling model is gaining traction among Grid sites due to its initial performance benefits. Additionally, understanding common execution patterns for different workloads allows for more efficient scheduling and resource allocation strategies.

However, managing the entire set of resources on a node requires careful orchestration. JAliEn employs custom mechanisms to dynamically allocate idle resources to running workloads, ensuring overall resource usage stays within the node's capacity.

This paper evaluates the experiences of the first sites using whole-node scheduling. It highlights its suitability for accommodating jobs with varying resource demands, particularly those with high memory requirements.

**Parallel (Track 9) / 134**

## Web-based graphics in ROOT

**Authors:** Bertrand Bellenot<sup>1</sup>; Olivier Couet<sup>1</sup>; Serguei Linev<sup>2</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

**Corresponding Authors:** olivier.couet@cern.ch, sergei.linev@cern.ch, bertrand.bellenot@cern.ch

The ROOT framework provides various implementations of graphics engines tailored for different platforms, along with specialized support of batch mode. Over time, as technology evolves and new versions of X11 or Cocoa are released, maintaining the functionality of correspondent ROOT components becomes increasingly challenging. The TWebCanvas class in ROOT represents an attempt to unify all these flavors together and provide a single implementation for all supported platforms –based on web technologies. The presentation will explore the utilization of the latest JavaScript ROOT v7, which offers comprehensive support for all ROOT classes, ensuring seamless integration and enhanced functionality across various platforms. Additionally, the presentation will discuss the implementation of HMAC-based authentication to bolster security for server-client communication, ensuring data integrity and confidentiality in web-based interactions. Furthermore, the presentation will showcase efforts in automating the testing of web graphics functionality, streamlining the development process and ensuring the reliability and robustness of the ROOT Web GUI.

Parallel (Track 3) / 135

## Optimized Graph Convolution for calorimetry event classification

**Authors:** Frederic Magniette<sup>1</sup>; Matthieu Martin Melennec<sup>1</sup>; Shamik Ghosh<sup>1</sup>

<sup>1</sup> Centre National de la Recherche Scientifique (FR)

**Corresponding Authors:** mmelennec@llr.in2p3.fr, frederic.bruno.magniette@cern.ch, shamik.ghosh@cern.ch

In the recent years, high energy physics discoveries have been driven by the increasing of detector volume and/or granularity. This evolution gives access to bigger statistics and data samples, but can make it hard to process results with current methods and algorithms. Graph neural networks, particularly graph convolution networks, have been shown to be powerful tools to address these challenges. These methods however raise some difficulties with their computing resource needs. In particular, representing physics events as graphs is a tricky problem that demands a good balance between resource consumption and graph quality, which can greatly affect the accuracy of the model. We propose a graph convolution network pipeline architecture to perform classification and regression tasks on calorimeter events and discuss its performances. It is designed for resource constrained environments, and in particular to efficiently represent calorimeter events as graphs, allowing up to a quadratic improvement in complexity with satisfying accuracy. Finally, we discuss possible applications to other high energy physics detectors.

Parallel (Track 3) / 136

## Ranking-based machine learning for track seed selection in ACTS

**Authors:** Corentin Allaire<sup>1</sup>; David Rousseau<sup>2</sup>; Françoise BOUVET<sup>None</sup>; Hadrien Benjamin Grasland<sup>3</sup>

<sup>1</sup> IJCLab, Université Paris-Saclay, CNRS/IN2P3

<sup>2</sup> IJCLab-Orsay

<sup>3</sup> IJCLab - CNRS

**Corresponding Authors:** corentin.allaire@cern.ch, rousseau@lal.in2p3.fr, hadrien.grasland@cern.ch, francoise.bouvet@ijclab.in2p3.fr

The reconstruction of particle trajectories is a key challenge of particle physics experiments, as it directly impacts particle identification and physics performances while also representing one of the primary CPU consumers of many high-energy physics experiments. As the luminosity of particle colliders increases, this reconstruction will become more challenging and resource-intensive. New

algorithms are thus needed to address these challenges efficiently. During track reconstruction, many more tracks are reconstructed than truth particles. This is partially due to fake tracks resulting from an arbitrary combination of detector hits and redundant duplicates of truth particle tracks. Reducing their amount could thus directly improve the speed of our tracking chain.

Those extra tracks are usually removed at the end of the tracking chain in a step called ambiguity resolution. We previously demonstrated (<https://indico.jlab.org/event/459/contributions/11453/>) that machine learning could speed up this step. Unfortunately, when a track is removed at the end of the tracking chain, all the time spent reconstructing it is lost. Eliminating fake and duplicated tracks before they are reconstructed would thus significantly speed up the reconstruction chain.

Cutting is usually applied to the seeds to mitigate this effect based on seed quality computed by the seeding algorithm. But those algorithms might not always be the most effective, requiring a lot of hand tuning and might not always keep the seed leading to the best possible track.

We propose using a ranking-based machine learning algorithm to select the track seeds before the track finding reconstructs them. The problem is fundamentally the same as with ambiguity resolution but with much less information available on the seeds than the tracks. With a clustering algorithm (such as DBSCAN), we can bundle together nearby seeds that appear to come from the same truth particle. Afterwards, we can apply a Neural Network (NN) using a novel Margin Ranking Loss Function to compare the seeds in each group and only keep one, which should lead to the closest reconstructed track to the truth. In order to fully evaluate this approach's potential, we implemented it within the A Common Tracking Software (ACTS) framework and tested it on the Open Data Detector (ODD), a realistic virtual detector similar to a future ATLAS one. This evaluation showed an up to ten times speedup of the track finding and an improvement in the quality of the reconstructed tracks at the cost of a slight decrease in efficiency.

**Poster session / 137**

## **Low-latency AI for triggering on electrons at High Luminosity LHC with the CMS Level-1 hardware Trigger.**

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

In preparation for the High Luminosity LHC (HL-LHC) run, the CMS collaboration is working on an ambitious upgrade project for the first stage of its online selection system: the Level-1 Trigger. The upgraded system will use powerful field-programmable gate arrays (FPGA) processors connected by a high-bandwidth network of optical fibers. The new system will access highly granular calorimeter information and online tracking: their combination for identifying physics objects is a key asset to cope with the harsh HL-LHC environment without compromising physics acceptance. The track matching is particularly relevant for identifying calorimeter deposits originating from electron particles. Traditional identification techniques rely on several independent selection stages applied to the calorimeter and track primitives, followed by an angular matching procedure. A new machine learning approach is presented, combining track and calorimeter information into a single identification and matching step. The new algorithm leverages new technologies for running fast inference on FPGA.

The talk will report on the system design, the implementation in firmware, and the performance obtained on simulated events.

**Poster session / 138**

## **Migration of CADI to Fence**

**Authors:** CMS Collaboration<sup>None</sup>; Muhammad Imran<sup>1</sup>

<sup>1</sup> National Centre for Physics (PK)

**Corresponding Author:** borislav.pavlov@cern.ch

CMS Analysis Database Interface (CADI) is a management tool for physics publications in the CMS experiment. It acts as a central database for the CMS collaboration, keeping track of the various analysis projects being conducted by researchers. Each analysis paper written by the authors goes through an extensive journey from early analysis to publication. There are various stakeholders involved in that process who can provide their comments/feedback and may be involved in the approval/disapproval process of the analysis. Front End Engine for Glance (FENCE) is a technology developed by the UFRJ team that emerged to unify and facilitate the development of UFRJ-CERN collaboration systems. It allows system interfaces to be created by simply editing a configuration file in JSON, without the need for deep programming knowledge of users and changing the system's internal source code. Thus, the current system of ATLAS experiment, which uses the Glance technology in its foundation and FENCE as an abstraction layer above, is developed, allowing users to access the heterogeneous data sources related to the experiments in a simple and efficient way. Originally developed by ATLAS, it was recently redesigned by LHCb following a more modular architecture – splitting the code base in a PHP based RestAPI backend and a VueJS based frontend service – and this version was also adopted for use in the LHCb and Alice experiments. CMS decided to migrate CADI to the new version of the FENCE system. For CMS, two subsystems of the FENCE system are initially considered: the “membership” and “analysis life cycle management”(ALCM). The membership subsystem is a prerequisite of ALCM. It contains information on members, institutes, authorships, and various reports. In contrast, the ALCM subsystem is primarily used for the management of publication workflows like CADI. In this talk, we'll describe the procedure that we followed to migrate CADI to FENCE. We encountered various issues during this process and will report the lessons learned while doing this migration so that other experiments in future will not have to undergo these issues if they migrate their system to FENCE.

**Parallel (Track 4) / 139**

## Enabling Alternative Architectures in the ALICE Grid

**Author:** Maksim Melnik Storetvedt<sup>1</sup>

**Co-authors:** Latchezar Betev<sup>1</sup>; Kalana Wijethunga<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> University of Moratuwa (LK)

**Corresponding Authors:** latchezar.betev@cern.ch, kalana.wijethunga@cern.ch, mstoretv@cern.ch

The ALICE Collaboration has begun exploring the use of ARM resources for the execution of Grid payloads. This was prompted by both their recent availability in the WLCG, as well as their increased competitiveness with traditional x86-based hosts in terms of both cost and performance. With the number of OEMs providing ARM offerings aimed towards servers and HPC growing, the presence of these resources in the Grid is anticipated to rise further. Consequently, it becomes a priority to ensure the underlying middleware is capable of running across architectures, ensuring available resources in the Grid are fully utilised.

This contribution outlines a reworked middleware stack, now used in production within ALICE, capable of running jobs across both Amd64 and Aarch64 ISAs, and the initial findings when used to execute Grid jobs compiled for the latter. Furthermore, it will examine how the middleware stack is able to dynamically match packages and binaries depending on the host. In turn, making both the selection process and executing architecture transparent from the end-user. At the same time, an overview is provided on how the middleware remains agnostic to the underlying architecture, allowing it to scale across various other types of CPUs - enabling support for additional architectures beyond ARM if needed, such as RISC-V.

**Parallel (Track6) / 140**

## AllocMonitor: An Extensible Memory Monitor

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

A major constraint on CMS production jobs is the amount of memory they require. CMS needs the ability to monitor and investigate memory usage in its software. General-purpose memory profilers often significantly slow down the monitored application, and require lots of additional memory. For many cases, the detailed information about memory allocations and deallocations recorded by the general-purpose profiler are not necessary.

The CMS AllocMonitor package provides a generalized facility for adding callbacks for when C++ allocates and deallocates memory using the C++ standard approved APIs. The callback system is fast and has low memory overhead. The work done by the callbacks primarily determines the full extent of the time and memory overheads of the profiling, and therefore can be tuned for the different use cases. For example, sometimes one only needs a very fast thread-safe overview of the amount of memory used by a job. At other times, one might want to know for a specific algorithm the per-thread memory usage.

This contribution will discuss three main items about AllocMonitor. First, an overview of the mechanisms used to acquire the data from the standard C++ API are described. Second, details about how one can use and customize the gathered data for different use cases are presented. Third, real world examples of useful customizations are shown.

**Poster session / 141**

## Unprivileged subdivision of job resources within the ALICE Grid

**Author:** Maksim Melnik Storetvedt<sup>1</sup>

**Co-authors:** Marta Bertran Ferrer<sup>1</sup>; Latchezar Betev<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** marta.bertran.ferrer@cern.ch, mstoretv@cern.ch, latchezar.betev@cern.ch

Job pilots in the ALICE Grid have become increasingly tasked with how to best manage the resources given to each job slot. With the emergence of more complex and multicore oriented workflows, this has since become an increasingly challenging process, as users often request arbitrary resources, in particular CPU and memory. This is further exacerbated by often having several user payloads running in parallel in the same slot, and with useful management utilities generally needing elevated privileges to function.

To alleviate resource management within each given job slot, the ALICE Grid has begun utilising novel features introduced in later Linux kernels, such as Cgroups v2, to provide means for fine-grained resource controls. By allowing specific controllers to be delegated down a Cgroup hierarchy, it enables users to access and tune these resource controls as needed - unprivileged. When further used in conjunction with the ALICE job pilot, it enables each job slot to be subpartitioned. In turn, allowing the pilot to act as its own local resource management system in its given slot - with a full "box-in" of each subjob to its own subset of the given resources.

This contribution describes the updated ALICE job pilot and its management and delegation process. Specifically, how it utilises kernel features to create individual resource groups for its jobs, while accommodating for the variety of configurations and computing elements used across participating sites - enabling these features to be used across the ALICE Grid.



**Parallel (Track 1) / 142**

## RNTuple: A CMS Perspective

**Author:** CMS Collaboration<sup>None</sup>**Corresponding Author:** borislav.pavlov@cern.ch

ROOT is planning to move from TTree to RNTuple as the data storage format for HL-LHC in order to, for example, speed up the IO, make the files smaller, and have a modern C++ API. Initially, RNTuple was not planned to support the same set of C++ data structures as TTree supports. CMS has explored the necessary transformations in its standard persistent data types to switch to RNTuple. Many challenges were encountered as alternative data structures were explored. This contribution will discuss the challenges uncovered and how collaboration with the ROOT team allowed them to be overcome. The solution to the challenges allows progressive changes to the CMS data types rather than requiring a sudden change to all data types to be stored in RNTuple. Once the solution was achieved, storage performance comparisons using the CMS data types were possible between RNTuple and TTree. This contribution will also present the results of those comparisons.

**Parallel (Track6) / 143**

## The CMS Interactive Job Tracer

**Author:** CMS Collaboration<sup>None</sup>**Corresponding Author:** borislav.pavlov@cern.ch

CMS applications are generally complex: they can have up to many thousands of components that the CMSSW framework schedules that are run on tens of threads. Understanding the timing characteristics of such complex applications is difficult, especially if one is looking for correlations between the components. To aid in understanding the runtime behavior of the applications CMS has developed a web browser based interactive data visualization tool. The tool is composed of three parts: one for gathering the data from the framework in a compact format, a second for reorganizing the data in the compact format for a specific use case, and a third for visualizing the reorganized data. This contribution will describe each part and how they work with each other. For the first part we will describe the generic hooks the CMSSW framework provides that enable real time tracing of the activities of the framework. For the second part, we will outline the data reorganization needed to support different use cases. And for the third part, we will show how the generic web viewer supports the different use cases without use-case-specific code.

**Poster session / 145**

## CMS L1 Data Scouting for HL-LHC

**Author:** CMS Collaboration<sup>None</sup>**Corresponding Author:** borislav.pavlov@cern.ch

The CMS Experiment at the CERN Large Hadron Collider (LHC) relies on a Level-1 Trigger system (L1T) to process in real time all potential collisions, happening at a rate of 40 MHz, and select the most promising ones for data acquisition and further processing. The CMS upgrades for the upcoming high-luminosity LHC run will vastly improve the quality of the L1T event reconstruction, providing opportunities for a complementary Data Scouting approach where physics analysis is performed on a data stream containing all collisions but limited to L1T reconstruction. This poster describes the future Data Scouting system, some first estimates of its physics capabilities, and the demonstration setups used to assess its technical feasibility.

**Plenary session / 146**

## ROOT RNTuple: Next Generation Event Data I/O for HENP

**Authors:** Florine de Geus<sup>1</sup>; Jakob Blomer<sup>2</sup>; Jonas Hahnfeld<sup>3</sup>; Philippe Canal<sup>4</sup>; Vincenzo Eduardo Padulano<sup>2</sup><sup>1</sup> *CERN/University of Twente (NL)*<sup>2</sup> *CERN*<sup>3</sup> *CERN & Goethe University Frankfurt*<sup>4</sup> *Fermi National Accelerator Lab. (US)***Corresponding Authors:** vincenzo.eduardo.padulano@cern.ch, philippe.canal@cern.ch, jonas.hahnfeld@cern.ch, florine.de.geus@cern.ch, jakob.blomer@cern.ch

For several years, the ROOT team is developing the new RNTuple I/O subsystem in preparation of the next generation of collider experiments. Both HL-LHC and DUNE are expected to start data taking by the end of this decade. They pose unprecedented challenges to event data I/O in terms of data rates, event sizes and event complexity. At the same time, the I/O landscape is getting more diverse. HPC cluster file systems and object stores, NVMe disk cache layers in analysis facilities, and S3 storage on cloud resources are mixing with traditional XRootD managed spinning disk pools.

The ROOT team will finalize a first production version of the RNTuple binary format by the end of the year. After this point, ROOT will provide backwards compatibility for RNTuple data. This contribution provides an overview of the RNTuple feature set, the related R&D activities, and the long-term vision for RNTuple. We report on performance, interface design, tooling, robustness, integration with experiment frameworks, and validation results as well as recent R&D on parallel reading and writing and exploitation of modern hardware and storage systems. We will give an outlook on possible future features after a first production release.

**Parallel (Track 7) / 147**

## Development of machine-learning based app for anomaly detection in CMSWEB

**Corresponding Author:** borislav.pavlov@cern.ch

The CMS experiment's operational infrastructure hinges significantly on the CMSWEB cluster, which serves as the cornerstone for hosting a multitude of services critical to the data taking and analysis. Operating on Kubernetes ("k8s") technology, this cluster powers over two dozen distinct web services, including but not limited to DBS, DAS, CRAB, WMarchive, and WMCORE.

In this talk, we propose and develop an application which is specifically tailored to the task of anomaly detection within this ecosystem of services. The core approach involves harnessing the capabilities of machine/deep learning methods, alongside a comprehensive exploration of various service parameters, to identify irregularities and potential threats effectively. The application is designed with the goal that continually monitors these services for any deviations from their expected behavior. Leveraging diverse machine/deep learning techniques and scrutinizing service-specific parameters, the application will be equipped to discern anomalies and aberrations that might signify security breaches or performance issues. Once an anomaly is detected, the system will not only record this event but will also promptly generate alerts. These alerts will be intelligently routed to the relevant service developers or administrators responsible for maintaining the affected components. This proactive alerting mechanism ensures that any emerging issues are swiftly addressed, minimizing potential disruptions and fortifying the overall reliability of the CMSWEB cluster and its critical services.

Parallel (Track 2) / 149

## Real-time Anomaly Detection at the L1 Trigger of CMS Experiment

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

We present the preparation, deployment, and testing of an autoencoder trained for unbiased detection of new physics signatures in the CMS experiment Global Trigger (GT) test crate FPGAs during LHC Run 3. The GT makes the final decision whether to readout or discard the data from each LHC collision, which occur at a rate of 40 MHz, within a 50 ns latency. The Neural Network makes a prediction for each event within these constraints, which can be used to select anomalous events for further analysis. The GT test crate is a copy of the main GT system, receiving the same input data, but whose output is not used to trigger the readout of CMS, providing a platform for thorough testing of new trigger algorithms on live data, but without interrupting data taking. We describe the methodology to achieve ultra low latency anomaly detection, and present the integration of the DNN into the GT test crate, as well as the monitoring, testing, and validation of the algorithm during proton collisions.

Poster session / 151

## A CXL-based Memory Lake Architecture for 40 MHz L1 Data Scouting at CMS

**Authors:** CMS Collaboration<sup>None</sup>; Giovanna Lazzari Miotto<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** borislav.pavlov@cern.ch

The Level-1 Data Scouting (L1DS) is a novel data acquisition subsystem at the CMS Level-1 Trigger (L1T) that exposes the L1T primitives involved in event selection for online processing at the LHC's full 40 MHz bunch-crossing rate, enabling zero-bias and unconventional analyses. Since Run 3, a L1DS demonstrator has relied on a shared ramdisk for its incoming and intermediate data. While the HL-LHC and CMS'Phase 2 upgrade are projected to augment trigger resolutions, processing and storage concerns have prompted the development of a new L1DS processing pipeline with a performant shared memory lake. For this, we leverage the emerging Compute Express Link (CXL) standard, whose protocols provide uniform, cache-coherent memory access to heterogeneous processing units, e.g., smart NICs and GPUs. In this contribution, we present the integration of CXL-compliant shared memory into the L1DS pipeline at CMS, expounding on the observed benefits and limitations of our approach. Furthermore, we perform a comprehensive evaluation of the demonstrator system's performance in realistic analyses and discuss use cases for the CMS community.

Poster session / 153

## Particle Flow Reconstruction with Alpaka Portability Library

**Authors:** CMS Collaboration<sup>None</sup>; Jonathan Samudio<sup>1</sup>

<sup>1</sup> Baylor University (US)

**Corresponding Author:** borislav.pavlov@cern.ch

In response to increasing data challenges, CMS has adopted the use of GPU offloading at the High-Level Trigger (HLT). However, GPU acceleration is often hardware specific, and increases the maintenance burden on software development. The Alpaka (Abstraction Library for Parallel Kernel Acceleration) portability library offers a solution to this issue, and has been implemented into the CMS software (CMSSW) for use online at HLT.

A portion of the final-state particle candidate reconstruction algorithm, Particle Flow, has been ported to Alpaka and deployed at HLT for 2024 data taking. The formation of hadronic Particle Flow clusters represented a target for increased performance through parallel operation. We will discuss the port of hadronic Particle Flow clustering to Alpaka, and the validation of physics and performance at HLT.

**Parallel (Track 3) / 154**

## Direct I/O for RNTuple Columnar Data

**Authors:** Jonas Hahnfeld<sup>1</sup>; Jakob Blomer<sup>2</sup>; Philippe Canal<sup>3</sup>; Thorsten Kollegger<sup>4</sup>

<sup>1</sup> *CERN & Goethe University Frankfurt*

<sup>2</sup> *CERN*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

<sup>4</sup> *Goethe University Frankfurt*

**Corresponding Authors:** jonas.hahnfeld@cern.ch, jakob.blomer@cern.ch, philippe.canal@cern.ch, kollegger@em.uni-frankfurt.de

RNTuple is the new columnar data format designed as the successor to ROOT's TTree format. It allows to make use of modern hardware capabilities and is expected to be used in production by the LHC experiments during the HL-LHC. In this contribution, we will discuss the usage of Direct I/O to fully exploit modern SSDs, especially in the context of the recent addition of parallel RNTuple writing. In contrast to buffered I/O where files are accessed via the operating system's page cache, Direct I/O circumvents all caching by the kernel and thereby enables higher bandwidths. However, to achieve this advantage, Direct I/O imposes strict alignment requirements on the I/O requests sent to the operating system: In particular, file offsets, byte counts and userspace buffer addresses must be aligned appropriately. This is challenging for columnar data formats and RNTuple pages that have variable size after compression. We will discuss possible strategies and performance results for both synthetic benchmarks as well as real-world applications.

**Parallel (Track 9) / 155**

## Leveraging distributed resources through high throughput analysis platforms for enhancing HEP data analyses

**Authors:** ATLAS Collaboration<sup>None</sup>; Adelina D'Onofrio<sup>1</sup>; Bernardino Spisso<sup>2</sup>; CMS Collaboration<sup>None</sup>; Elvira Rossi<sup>2</sup>; Federica Maria Simone<sup>3</sup>; Francesco Giuseppe Gravili<sup>4</sup>; Salvatore Loffredo<sup>5</sup>; Tommaso Diotallevi<sup>6</sup>

<sup>1</sup> *INFN Napoli (IT)*

<sup>2</sup> *Universita Federico II e INFN Sezione di Napoli (IT)*

<sup>3</sup> *Universita e INFN, Bari (IT)*

<sup>4</sup> *INFN Lecce e Universita del Salento (IT)*

<sup>5</sup> *INFN - National Institute for Nuclear Physics*

<sup>6</sup> *Universita e INFN, Bologna (IT)*

**Corresponding Authors:** borislav.pavlov@cern.ch, adelina.d'onofrio@cern.ch, bernardino.spisso@cern.ch, francesco.giuseppe.gravili@cern.ch, elvira.rossi@cern.ch

The analysis of data collected by the CMS and ATLAS experiments at CERN, ahead of the next phase of high-luminosity at the LHC, requires a flexible and dynamic access to big amounts of data, as well as an environment capable of dynamically accessing distributed resources. An interactive high throughput platform, based on a parallel and geographically distributed back-end, has been developed in the framework of the “HPC, Big Data e Quantum Computing Research Centre” Italian National Center (ICSC), providing experiment-agnostic resources. Starting from container technology and orchestrated via Kubernetes, the platform provides analysis tools via Jupyter interface and Dask scheduling system, masking complexity for frontend users and rendering cloud resources flexibly.

An overview of the technologies involved and the results on benchmark use cases will be provided, with suitable metrics to evaluate preliminary performance of the workflow. The comparison between the legacy analysis workflows and the interactive and distributed approach will be provided based on several metrics from event throughput to resource consumption. The use cases include the search for direct pair production of supersymmetric particles and for dark matter in events with two opposite-charge leptons, jets and missing transverse momentum using data collected by the ATLAS detector in Run-2 (JHEP 04 (2021) 165), and searches for rare flavor decays at the CMS experiment in Run-3 using large datasets collected by high-rate dimuon triggers.

**Parallel (Track6) / 156**

## Comparative efficiency of HEP codes across languages and architectures

**Author:** Samuel Cadellin Skipsey<sup>None</sup>

**Corresponding Author:** samuel.cadellin.skipsey@cern.ch

Recently, interest in measuring and improving the energy (and carbon) efficiency of computation in HEP, and elsewhere, has grown significantly. Measurements have been, and continue to be, made of the efficiency of various computational architectures in standardised benchmarks... but those benchmarks tend to compare only implementations in single programming languages. Similarly, comparisons of the efficiency of various languages tend to focus on a single architecture, although it is the case that some abstractions in a given language can match specific architectural choices (in, say, memory ordering strictness) better than others.

The existence of the JetReconstruction.jl project, implementing a subset of the FastJet C++ code's functionality in performant Julia, allows us to usefully compare how the relative efficiencies of implementations in the two languages are influenced by the architecture they are executed on.

We report on the results of comparing benchmarks on these codes, and others, on x86 and various aarch64 implementations, amongst others.

**Parallel (Track 3) / 157**

## Towards a GPU-enabled electron seeding algorithm in the CMS experiment

**Author:** Charis Kleio Koraka<sup>1</sup>

<sup>1</sup> University of Wisconsin Madison (US)

Electrons are one of the key particles that are detected by the CMS experiment and are reconstructed using the CMS software (CMSSW). Reconstructing electrons in CMSSW is a computational intensive task that is split into several steps, seeding being the most time consuming one. During the electron seeding process, the collection of tracker hits (seeds) is significantly reduced by selecting only seeds that are compatible with a hypothesized electron trajectory. This contribution will describe the

process of redesigning the electron seeding algorithm in a parallelizable way that will exploit the massive parallelism that GPUs can offer. The new algorithm code base is implemented using the Alpaka library, a performance portability library that allows having a single code base for execution on different types of hardware.

Poster session / 160

## Continuous integration of analysis workflows on a distributed analysis facility

**Authors:** CMS Collaboration<sup>None</sup>; Matteo Bartolini<sup>1</sup>

<sup>1</sup> *Universita e INFN, Firenze (IT)*

Data analysis in the field of High Energy Physics presents typical big data requirements, such as the vast amount of data to be processed efficiently and quickly. The Large Hadron Collider in its high luminosity phase will produce about 100 PB/year of data, ushering in the era of high precision physics. Currently, analysts are building and sharing their software on git-based platforms which improve reproducibility and offer a high level of workflow automatization. On the other hand, it's becoming more and more critical to complement this aspect with an easy and user-friendly access to distributed resources for CPU-intensive calculations. In this talk, it will be shown how it is possible to enable Continuous Integration (CI) with CMS datasets by using the XRootD IO protocol and dynamic proxy generation and, in combination with the GitLab CI/CD functionalities, how to trigger an analysis execution with a simple commit. By using dynamic auth access tokens it's possible to offload all the CPU-heavy work from the gitlab workers to on-demand computing resources: from regional CMS Tier-2 resources to the national-wide datalake model currently under deployment within the ICSC (the italian national center for research in HPC, big data and quantum computing) project. Thanks to this alternative approach, in particular, integrating the submission of jobs to HTCondor into the gitlab CI will become easier, automatising the handling of big datasets. In this way analysts will be able to quickly run different tests on their data, perform different analyses in parallel and, at the same time, keep tracks of all the changes made.

Poster session / 161

## The git based ATLAS data acquisition configuration service in LHC Run 3

**Author:** ATLAS TDAQ<sup>None</sup>

The ATLAS experiment at the LHC at CERN uses a large, distributed trigger and data acquisition system composed of many computing nodes, networks, and hardware modules. Its configuration service is used to provide descriptions of control, monitoring, diagnostic, recovery, dataflow and data quality configurations, connectivity, and parameters for modules, chips, and channels of various online systems, detectors, and the whole ATLAS experiment. Those descriptions have historically been stored in more than one thousand interconnected XML files, which are updated by various experts many times per day. Maintaining error-free and consistent sets of such files and providing reliable and fast access to current and historical configurations is a major challenge. This paper gives details of the configuration service upgrade on the modern git version control system backend for LHC Run 3 and its exploitation experience. It may be interesting for developers using human-readable file formats, where consistency of the files, performance, access control, traceability of modifications, and effective archiving are key requirements.

Parallel (Track 2) / 162

## Leveraging the Run 3 experience for the evolution of the ATLAS software-based readout towards HL-LHC

Author: ATLAS TDAQ<sup>None</sup>

The High-Luminosity Large Hadron Collider (HL-LHC), scheduled to start operating in 2029, aims to increase the instantaneous luminosity by a factor of 10 compared to the LHC. To match this increase, the ATLAS experiment has been implementing a major upgrade program divided into two phases. The first phase (Phase-I), completed in 2022, introduced new trigger and detector systems that have been used during the Run 3 data taking period which began in July 2022. These systems have been used in conjunction with the new Data Acquisition (DAQ) Readout system, based on a software application called Software Readout Driver (SW ROD). SW ROD receives and aggregates data from the front-end electronics via the Front-End Link eXchange (FELIX) system and passes aggregated data fragments to the High-Level Trigger (HLT) system. During Run 3, SW ROD operates in parallel with the legacy Readout System (ROS) at an input rate of 100 kHz. For the Phase-II, the legacy ROS will be completely replaced with a new system based on the next generation of FELIX and an evolution of the SW ROD application called Data Handler. Data Handler has the same functional requirements as SW ROD but must be able to operate at an input rate of 1 MHz. To facilitate this evolution the SW ROD has been implemented using plugin architecture.

This contribution presents the design and implementation of the SW ROD application for Run 3, along with the strategy for its evolution to the Phase-II Readout system. It discusses the lessons learned during Run 3 and describes the challenges that have been addressed to accomplish the demanding performance requirements of HL-LHC.

Parallel (Track 2) / 163

## GPU Acceleration and EDM Developments for the ATLAS 3D Calorimeter Clustering in the Software Trigger

Author: ATLAS TDAQ<sup>None</sup>

ATLAS is one of the two general-purpose experiments at the Large Hadron Collider (LHC), aiming to detect a wide variety of physics processes. Its trigger system plays a key role in selecting the events that are detected, filtering them down from the 40 MHz bunch crossing rate to the 1 kHz rate at which they are committed to storage. The ATLAS trigger works in two stages, Level-1 and the High-Level Trigger (HLT), with the first being a hardware-based coarse filtering applied using custom electronics and FPGAs, and the second relying on offline-like algorithms implemented fully in software, running on a farm of commodity CPUs. The LHC will undergo the High-Luminosity Upgrade soon (scheduled to be finished by 2029), which represents an additional challenge for the ATLAS trigger. The increased pile-up leads to events that are typically more complex and thus more computationally demanding to reconstruct, and a broad-ranging suite of upgrades to the ATLAS detector itself also encompasses increasing the input and output rates of the High Level Trigger by a factor of 10. As such, both the processing power required to handle a single event and the overall number of events that will need to be processed will increase, placing greater pressure on the trigger farm. One possibility of answering these increased computational demands in a cost- and energy-effective

way is the use of hardware accelerators, in particular leveraging the massive parallelism and general computational capabilities offered by GPUs for problems that are suited to their mode of operation.

Among the algorithms being assessed for GPU acceleration, Topological Clustering, the main and most computationally demanding stage of calorimeter reconstruction, has reached the significant milestone of 100% agreement with the CPU algorithm and maximum speed-ups in excess of a factor of 10. This is achieved through a more GPU-friendly variant of the algorithm, dubbed Topo-Automaton Clustering. A significant bottleneck remains in the time taken to convert between the data representation used within the GPU and the equivalent CPU data structures, which can be up to two thirds of the total execution time of the algorithm. This contribution will describe the development, optimization and integration of Topo- Automaton Clustering with the ATLAS trigger, including the latest benchmarks and ongoing efforts to develop an EDM framework that could allow for a general description of GPU-friendly data structures in order to alleviate the main bottleneck.

Parallel (Track 2) / 164

## Development of an FPGA based track reconstruction pipeline for the ATLAS Event Filter

Author: ATLAS TDAQ<sup>None</sup>

For the upcoming HL-LHC upgrade of the ATLAS experiment, the deployment of GPU or FPGA co-processors within the online Event Filter system is being studied as a measure to increase throughput and save power. End-to-end track reconstruction pipelines are currently being developed using commercially available FPGA accelerator cards. These utilize FPGA base partitions, drivers and runtime tools supplied by the manufacturer to reduce design effort. Algorithms are implemented both in hardware description language (HDL) and high-level synthesis (HLS), and integrated as kernels into an OpenCL host software interfacing with the ATLAS main software framework Athena. This contribution summarizes the algorithmic developments, the integration workflow and status.

Poster session / 165

## Turning CephFS into a collaborative cloud storage with CERN-Box

Authors: Giuseppe Lo Presti<sup>1</sup>; Sebastian Bukowiec<sup>1</sup>; Diogo Castro<sup>1</sup>; Javier Ferrer<sup>None</sup>; Hugo Gonzalez Labrador<sup>1</sup>

<sup>1</sup> CERN

Corresponding Authors: javier.ferrer@cern.ch, giuseppe.lopresti@cern.ch, hugo.gonzalez.labrador@cern.ch, sebastian.bukowiec@cern.ch, diogo.castro@cern.ch

CERNBox is an innovative scientific collaboration platform, built using solely open-source components to meet the unique requirements of scientific workflows. Used at CERN for the last decade, the service satisfies the 35K users at CERN and seamlessly integrates with batch farms and Jupyter-based services. Powered by Reva, an open-source HTTP and gRPC server written in Go, CERNBox has demonstrated the provision of Sync&Share capabilities on top of multiple storage systems such as EOS and CephFS, as well as enabling federated sharing with other institutions.



In this contribution, we present the evolution of CERNBox in supporting CephFS, which has been chosen as the storage system to address the Windows applications use-cases at CERN. As we are migrating out of DFS, the legacy Windows storage provided by Microsoft, and commissioning Windows Workspaces powered by CephFS, we show how CERNBox provides a flexible software stack to seamlessly integrate the Windows-based community, which includes the Engineering sector of the Organization.

We conclude by emphasizing the multiple synergies enabled by this approach. On one hand, Windows-based data-centric workflows can leverage the multi-protocol accesses (sync, web, SMB) provided by CERNBox. On the other hand, the widespread adoption of CephFS within the scientific community positions CERNBox as an out-of-the-box solution for implementing a scalable collaborative cloud storage service.

Parallel (Track 2) / 166

## Online track reconstruction with graph neural networks on FPGAs for the ATLAS experiment

Author: ATLAS TDAQ<sup>None</sup>

For the HL-LHC upgrade of the ATLAS TDAQ system, a heterogeneous computing farm deploying GPUs and/or FPGAs is under study, together with the use of modern machine learning algorithms such as Graph Neural Networks (GNNs). We present a study on the reconstruction of tracks in the ATLAS Inner Tracker using GNNs on FPGAs for the Event Filter system. We explore each of the steps in a GNN-based tracking pipeline: graph construction, edge classification using an interaction network, and segmentation of the graph into track candidates. We investigate optimizations of the GNN approach that aim to minimize FPGA resources utilization and maximize throughput while retaining high track reconstruction efficiency and low fake rates required for the ATLAS Event Filter tracking system. These studies include model hyperparameter tuning, model pruning and quantization-aware training, and sequential processing of sub-graphs across the detector.

Poster session / 167

## Performance of the ATLAS GNN4ITk Particle Track Reconstruction GPU pipeline

Author: ATLAS TDAQ<sup>None</sup>

With the upcoming upgrade of High Luminosity LHC, the need for computation power will increase in the ATLAS trigger system by more than an order of magnitude. Therefore, new particle track reconstruction techniques are explored by the ATLAS collaboration, including the usage of Graph Neural Networks (GNN). The project focusing on that research, GNN4ITk, considers several heterogeneous computing options, including the usage of Graphics Processing Units (GPU). The framework can reconstruct tracks with high efficiency, however, the computing requirements of the pipeline are high. We will report on the efforts to reduce the memory consumption and inference time enough to enable the usage of commercially available and affordable GPUs for the future ATLAS trigger system while maintaining high tracking performance.

Parallel (Track 4) / 168

## Implementation and Performance Analysis of the ALICE grid middleware JAliEn's Job Optimizer

**Authors:** Bjarte Kileng<sup>1</sup>; Costin Grigoras<sup>2</sup>; Haakon Andre Reme-Ness<sup>1</sup>; Haavard Helstrup<sup>1</sup>; Latchezar Betev<sup>2</sup>; Maksim Melnik Storetvedt<sup>2</sup>

<sup>1</sup> *Western Norway University of Applied Sciences (NO)*

<sup>2</sup> *CERN*

**Corresponding Authors:** bki@hvl.no, haakon.andre.reme-ness@cern.ch, mstoretv@cern.ch, haavard.helstrup@cern.ch, costin.grigoras@cern.ch, latchezar.betev@cern.ch

This paper presents a comprehensive analysis of the implementation and performance enhancements of the new job optimizer service within the JAliEn (Java ALICE environment) middleware framework developed for the ALICE grid. The job optimizer service aims to efficiently split large-scale computational tasks into smaller grid jobs, thereby optimizing resource utilization and throughput of the grid by ensuring more grid resources are able to match with grid jobs. New functionalities for users of the grid are described, while also delving into back-end changes that have improved the job optimizer service.

Through testing and evaluation in a production environment, significant improvements in database performance, faster job splitting, and better scalability have been observed when doing comparative analysis against the legacy job optimization service. Further potential improvements in the future will also be explored.

This paper will also provide a look into the technical intricacies of the new job optimizer service, highlighting functionalities, implementation strategies, and integration within the existing JAliEn framework. Furthermore, insights into the lessons learned and challenges encountered during the implementation phase, deployment, and operationalization of the job optimizer service will be discussed.

Poster session / 169

## QDIPS: Deep Sets Network for FPGA investigated for high speed inference on ATLAS

**Author:** ATLAS TDAQ<sup>None</sup>

Deep sets network architectures have useful applications in finding correlations in unordered and variable length data input, thus having the interesting feature of being permutation invariant. Its use on FPGA would open up accelerated machine learning in areas where the input has no fixed length or order, such as inner detector hits for clustering or associated particle tracks for jet tagging. We adapted DIPS (Deep Impact Parameter Sets), a deep sets neural network flavour tagging algorithm previously used in ATLAS offline low-level flavour tagging and online b-jet trigger preselections, for use on FPGA with the aim to assess its performance and resource costs. QKeras and HLS4ML are used for quantisation-aware training and translation for FPGA implementation, respectively. Some challenges are addressed, such as finding replacements for functionality not available in HLS4ML (e.g. Time Distributed layers) and implementations of custom HLS4ML layers. Satisfactory implementations are tested on an actual FPGA board for the assessment of true resource consumption and latency. We show the optimal FPGA-based algorithm performance relative to CPU-based full precision performance previously achieved in the ATLAS trigger, as well as performance trade-offs when reducing FPGA resource usage as much as possible. The project aims to demonstrate a viable solution for performing sophisticated Machine Learning-based tasks for accelerated reconstruction or particle identification for early event rejection while running in parallel to other more intensive tasks on FPGA.

**Parallel (Track 2) / 170**

## **Evolution of the ATLAS TDAQ online software framework towards Phase-II upgrade: use of Kubernetes as an orchestrator of the ATLAS Event Filter computing farm**

**Author:** ATLAS TDAQ<sup>None</sup>

The ATLAS experiment at the Large Hadron Collider (LHC) at CERN continuously evolves its Trigger and Data Acquisition (TDAQ) system to meet the challenges of new physics goals and technological advancements. As ATLAS prepares for the Phase-II Run 4 of the LHC, significant enhancements in the TDAQ Controls and Configuration tools have been designed to ensure efficient data collection, processing, and management. This abstract presents the evolution of ATLAS TDAQ Controls and Configuration system leading up to Phase-II Run4. As part of the evolution towards Phase-II, Kubernetes has been chosen to orchestrate the Event Filter farm. By leveraging Kubernetes, ATLAS can dynamically allocate computing resources, scale processing capacity in response to changing data taking conditions, and ensure high availability of data processing services. The integration of the Kubernetes with the TDAQ Run Control framework enables perfect synchronisation between the experiment's data acquisition components and the computing infrastructure. We will discuss the architectural considerations and implementation challenges involved in Kubernetes integration with the ATLAS TDAQ controls and configuration system. We will highlight the benefits of using Kubernetes as an event filter farm orchestrator, including improved resource utilization, enhanced fault tolerance, and simplified deployment and management of data processing workflows. In addition, we will report on the extensive testing of Kubernetes that was conducted using a farm of 2500 servers within the experiment data taking environment, demonstrating its scalability and robustness in handling the demands of the ATLAS TDAQ system for Phase-II. The adoption of Kubernetes represents a significant step forward in the evolution of ATLAS TDAQ controls and configuration system, aligning with industry best practices in container orchestration and cloud-native computing.

**Poster session / 171**

## **Updates of the ATLAS High-Level Trigger in Run 3**

**Author:** ATLAS TDAQ<sup>None</sup>

The main reconstruction and simulation software framework of the ATLAS experiment, Athena, underwent a major change during the LHC Run 3 in the way the configuration step of its applications is performed. The new configuration system, called ComponentAccumulator, emphasises modularity and provides a way for standalone execution of parts of a job, as long as the inputs are available, which allows unit-testing of individual components or groups of components, as well as easier debugging.

The switch to the new configuration system of the High-Level Trigger (HLT) software, which utilises Athena algorithms for object reconstruction and hypothesis testing, required designing a special approach to prevent disruption of data taking during the code migration to ComponentAccumulator. An additional challenge is brought by a large amount of HLT chains, where in many cases copies of the same algorithm with varying configurations are used, which significantly increases the number of configured parameters compared to offline reconstruction jobs.

This report describes migration of the HLT software to ComponentAccumulator along with further improvements in the data acquisition introduced for Run 3

data taking.

**Parallel (Track 2) / 172**

## **The upgrade of the ATLAS Trigger and Data Acquisition system for the High Luminosity LHC**

**Author:** ATLAS TDAQ<sup>None</sup>

The ATLAS experiment at CERN is constructing upgraded system for the “High Luminosity LHC”, with collisions due to start in 2029. In order to deliver an order of magnitude more data than previous LHC runs, 14 TeV protons will collide with an instantaneous luminosity of up to  $7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ , resulting in much higher pileup and data rates than the current experiment was designed to handle. While this is essential to realise the physics programme, it presents a huge challenge for the detector, trigger, data acquisition and computing. The detector upgrades themselves also present new requirements and opportunities for the trigger and data acquisition system.

The design of the TDAQ upgrade comprises: a hardware-based low-latency real-time Trigger operating at 40 MHz, Data Acquisition which combines custom readout with commodity hardware and networking to deal with 4.6 TB/s input, and an Event Filter running at 1 MHz which combines offline-like algorithms on a large commodity compute service with the potential to be augmented by commercial accelerators. Commodity servers and networks are used as far as possible, with custom ATCA boards, high speed links and powerful FPGAs deployed in the low-latency parts of the system. Offline-style clustering and jet-finding in FPGAs, and accelerated track reconstruction are designed to combat pileup in the Trigger and Event Filter respectively.

This contribution will report recent progress on the design, technology and construction of the system. The physics motivation and expected performance will be shown for key physics processes.

**Poster session / 173**

## **Integrating FPGA Accelerators in Athena for use at the ATLAS Event Filter**

**Author:** ATLAS TDAQ<sup>None</sup>

In the realm of high-energy physics research, the demand for computational power continues to increase, particularly in online applications such as Event Filter. Innovations in performance enhancement are sought after, leading to exploration in integrating FPGA accelerators within existing software frameworks like Athena, extensively employed in the ATLAS experiment at CERN. This presentation delves into the intricacies of this integration, focusing on the system-level challenges posed by the simultaneous utilization of FPGA resources by multiple Athena algorithms in the heterogeneous computing environment explored for the TDAQ Phase II upgrade.

Central to this discussion is the notion of shared state management,

particularly concerning the loading of FPGA bitstreams. As multiple algorithms contend for access to the same FPGA, efficient management of the FPGA's state becomes crucial to ensure optimal performance and resource utilization. This work addresses this challenge, presenting insights and strategies for orchestrating FPGA resource sharing within the Athena framework.

While still a work in progress, this contribution provides valuable insights into the ongoing efforts to seamlessly integrate FPGA accelerators into complex research environments, paving the way for enhanced computational capabilities.

**Parallel (Track 2) / 174**

## The ATLAS Trigger System

**Author:** ATLAS TDAQ<sup>None</sup>

The ATLAS experiment in the LHC Run 3 uses a two-level trigger system to select events of interest to reduce the 40 MHz bunch crossing rate to a recorded rate of up to 3 kHz of fully-built physics events. The trigger system is composed of a hardware based Level-1 trigger and a software based High Level Trigger. The selection of events by the High Level Trigger is based on a wide variety of reconstructed objects, including leptons, photons, jets, b-jets, missing transverse energy, and B-hadrons in order to cover the full range of the ATLAS physics programme. We will present an overview of improvements in the reconstruction, calibration, and performance of the different trigger objects, as well as computational performance of the High Level Trigger system.

**Poster session / 176**

## ML-based classification in an open-source framework for the ALICE heavy-flavour analysis

**Author:** Maria Teresa Camerlingo<sup>1</sup>

<sup>1</sup> *Universita e INFN, Bari (IT)*

**Corresponding Author:** maria.teresa.camerlingo@cern.ch

The ALICE Collaboration aims to precisely measure heavy-flavour (HF) hadron production in high-energy proton-proton and heavy-ion collisions since it can provide valuable tests of perturbative quantum chromodynamics models and insights into hadronization mechanisms. Measurements of the  $\Xi_c^+$  and  $\Lambda_c^+$  production decaying in a proton (p) and charged  $\pi$  and K mesons are remarkable examples of investigation in the HF sector. Like in other ALICE analyses, a quite novel approach based on Boosted Decision Tree (BDT) classifiers has been adopted to discriminate the signal yields from the background processes. Especially for  $\Xi_c^+ \rightarrow p\pi K$  process, the Machine Learning (ML)-based approach is required and particularly challenging due to its large combinatorial background, small branching ratio, and short  $O(100 \mu\text{m})$  decay length of  $\Xi_c^+$  baryon. FAIR, a European project synergic to the ALICE experiment, aims to set up an open-source, user-friendly and interactive pytorch-based environment external to the official ALICE framework to perform BDT-based multivariate analyses. The FAIR benchmark imports different ML packages (XGBoost, Sklearn and Ray) to prepare the data and configure the BDT models in Jupyter Notebooks. Currently, the training is performed on a preliminary dataset with limited statistics using a partitioned shared GPU available through an Apache Mesos cluster at the ReCaS-Bari datacenter. In the future, when a larger dataset will be available, we intend to leverage a GPU-powered Kubernetes cluster for processing large-scale applications, including ML tool training. This contribution will present a performance comparison

of the investigated BDT architectures trained with simulated signal events and background Run 3 data provided by ALICE.

**Parallel (Track 4) / 177**

## Unified Experiment Monitoring

**Author:** Ewoud Ketele<sup>1</sup>

**Co-author:** Domenico Giordano<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** domenico.giordano@cern.ch, ewoud.rob.l.ketele@cern.ch

The Unified Experiment Monitoring (UEM) is the project in WLCG with the objective to harmonise the WLCG job accounting reports across the LHC experiments, in order to provide aggregated reports of the compute capacity used by WLCG along time. This accounting overview of all LHC experiments is vital for the strategy planning of WLCG and therefore it finds the strong support of the LHC Committee (LHCC). However, creating common overviews is challenging, due to the different internals of each experiment monitoring system and also due to the long time scale of the reports to cover at least a decade of data. These monitoring systems evolved largely independently over time, implying that the UEM project has to design and implement different approaches to couple the multiple data sources within the CERN IT monitoring tools which will be used. Last but not least, the different terminologies have to be aligned into a useful and coherent set. This contribution will drive the audience through the motivations of the project, the challenges faced, the design adopted to overcome them, and the presentation of the state of the art.

**Poster session / 178**

## Deployment of inference as a service at the US CMS Tier-2 data centers

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

Coprocessors, especially GPUs, will be a vital ingredient of data production workflows at the HL-LHC. At CMS, the GPU-as-a-service approach for production workflows is implemented by the SONIC project (Services for Optimized Network Inference on Coprocessors). SONIC provides a mechanism for outsourcing computationally demanding algorithms, such as neural network inference, to remote servers, where requests from multiple clients are intelligently distributed across multiple GPUs by a load-balancing service. This talk highlights the recent progress in deploying SONIC at selected U.S. CMS Tier-2 data centers. Using realistic CMS Run3 data processing workflows, such as those containing transformer-based algorithms, we demonstrate how SONIC is integrated into the production-like environment to enable accelerated inference offloading. We will present developments from both the client and server sides, including production job and data center configurations for NVIDIA and AMD GPUs. We will also present performance scaling benchmarks and discuss the challenges of operating SONIC in CMS production, such as server discovery, GPU saturation, fallback server logic, etc.

**Parallel (Track 1) / 180**

## An automation framework for condition data processing using data-streaming techniques and ML models

**Author:** Elton Shumka<sup>1</sup>

<sup>1</sup> *University of Sofia - St. Kliment Ohridski (BG)*

A novel approach for the periodic aggregation, synchronization, and automated analysis of non-physics data relevant to the operation of modern high-energy physics experiments will be introduced during the talk. We present an implementation of this approach tailored to the Resistive Plate Chambers (RPC) subsystem of the CMS experiment at CERN LHC accelerator, known as “RPC Automation”. The RPC system is a part of the CMS Muon spectrometer and consists of 1056 detectors with a total number of 123432 digital data channels, read out synchronously with a 40 MHz clock. Main RPC parameters such as currents and rates are continuously correlated with running conditions such as LHC luminosity, environmental humidity, pressure and temperature. All raw data is stored in Oracle databases by the WinCC OA framework. The RPC Automation leverages periodic data streaming within discrete time intervals to optimize the execution of database queries. Our methodology encompasses robust data streaming, real-time analysis of non-physics data, synchronization of data points with different timestamps, and the subsequent storage of processed data in new tables. Furthermore, the automation framework includes modules for preprocessing data intended for training machine learning-based models used in detector operation monitoring via anomaly detection.

**Parallel (Track 8) / 181**

## Cold data support for the CERN Open Data Portal

**Authors:** Jose Benito Gonzalez Lopez<sup>1</sup>; Pablo Saiz<sup>1</sup>; Zacharias Zacharodimos<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** pablo.saiz@cern.ch, jose.benito.gonzalez@cern.ch, zacharias.zacharodimos@cern.ch

The CERN Open Data Portal holds over 5 petabytes of high-energy physics experiment data, serving as a hub for global scientific collaboration. Committed to Open Science principles, the portal aims to democratize access to these datasets for outreach, training, education, and independent research. Recognizing the limitations of current disk-based storage, we are starting a project to expand our data storage methodologies. Our approach involves integrating hot storage (such as spinning disks) for immediate data access and cold storage (such as tape, or even interfaces to the experiment frameworks) for cost-effective long-term preservation. This innovative strategy will significantly expand the portal’s capacity to accommodate more experiment data. However, we anticipate challenges in navigating technical complexities and logistical hurdles. These challenges include the latency to access cold data, monitoring and automatizing the transitions between hot and cold and ensuring the long-term preservation of data in the experiment frameworks. The strategy is to integrate existing solutions like EOS, FTS, CTA and Rucio.

In our presentation, we will discuss these challenges, present our prototype solution, and outline future developments aimed at enhancing the accessibility, efficiency, and resilience of the CERN Open Data Portal’s data ecosystem.

**Parallel (Track 7) / 182**

## The Global Network Advancement: A Next Generation System for Data Intensive Sciences

**Author:** Harvey Newman<sup>1</sup>

<sup>1</sup> *California Institute of Technology (US)*

**Corresponding Author:** harvey.newman@cern.ch

The Global Network Advancement Group (GNA-G) and its Data Intensive Sciences (DIS) and SENSE/AutoGOLE Working Groups, a worldwide collaboration bringing together major science programs, research and education networks, and advanced network R&D projects spanning the U.S, Europe, Asia, Latin America and Oceania, are developing a next generation network-integrated system that will meet the challenges of data intensive sciences, showing the way towards the next generation of intelligent operations of R&E networks, as well as new methods and modes of network operation and capacity management that will benefit both network research and production teams. The continued use of the Waveserver Ais discussed in this proposal will be a key enabler of these developments.

These advances, being developed within the GNA-G framework and its SENSE/AutoGOLE and Data Intensive Sciences (DIS) working groups, the Global Research Platform, and its many science, computer science and R&E network teams are working to address the key challenges including: (1) global data distribution, processing, access and analysis, (2) the coordinated use of massive but still limited computing, storage and network resources, (3) coordinated operation and collaboration within global scientific enterprises each encompassing hundreds to thousands of scientists, and (4) enabling the science programs to make efficient use of the available network and site infrastructures, while simultaneously accommodating and in concert with the network operations required to supporting the worldwide academic and research community across national, regional and transoceanic boundaries.

In order to meet these high-level goals, the teams are working towards deploying dynamic and adaptive programmable software-driven system which coordinates worldwide networks as a first class resource along with computing and storage, across multiple domains. We are following a systems design approach to create a global dynamic fabric that flexibly allocates, balances and conserves the available network resources, while negotiating and working with the site-resident systems that aim to accelerate workflow.

Reinforcement and other model-based machine learning techniques are also planned to be used to optimize system operations, according to objective functions that take priority, policy, responses to network- and site-state changes, workflow objectives and other constraints into account.

An overarching concept is “Consistent Network Operations,” where stable load balanced high throughput workflows crossing optimally chosen network paths, up to preset high water marks to accommodate other traffic, are provided by autonomous site-resident services dynamically interacting with network-resident services, in response to demands from the science programs’ principal data distribution and management systems.

**Parallel (Track 5) / 183**

## **QiboML: a full-stack quantum machine learning framework**

**Author:** Matteo Robbiati<sup>1</sup>

**Co-authors:** Michele Grossi<sup>2</sup>; Sofia Vallecorsa<sup>2</sup>; Stefano Carrazza<sup>3</sup>

<sup>1</sup> *Università degli Studi e INFN Milano (IT)*

<sup>2</sup> *CERN*

<sup>3</sup> *Università degli Studi di Milano (IT)*

**Corresponding Authors:** michele.grossi@cern.ch, matteo.robbiati@cern.ch, stefano.carrazza@cern.ch, sofia.vallecorsa@cern.ch

We present QiboML, an open-source software library for Quantum Machine Learning (QML) integrated with the Qibo quantum computing middleware framework.

QiboML incorporates the most commonly used classical Machine Learning frameworks such as TensorFlow, PyTorch and Jax with Qibo.

This combination enables the users to construct quantum or hybrid classical-quantum models that can be executed on any type of hardware accelerators: CPU, multi-threading CPU, GPU and multi-GPU for quantum simulation on classical hardware (using state-vector and tensor network approaches) and Quantum Processing Units (QPU) for execution on self-hosted quantum devices.

We first present the QiboML core features, including common quantum observable operators and



their gradient representation, models and layers abstraction for training and data encoding, optimizers and standard loss functions for supervised and unsupervised learning. We then present built-in models for QML applications, with particular focus on High-Energy Physics (HEP) problems, where QML has recently shown to be particularly promising. We finally present an application: the determination of the u-quark Parton Distribution Function using a parametric quantum circuit. We discuss the obtained results and how real-time quantum error mitigation routines can be exploited to enhance the training on a self-hosted single-qubit superconducting quantum device.

**Parallel (Track 7) / 184**

## Canadian Belle-II Tier1 Raw Data Centre

**Authors:** Colson Driemel<sup>1</sup>; Marcus Ebert<sup>1</sup>; Randall Sobie<sup>2</sup>; Tristan Sullivan<sup>1</sup>

<sup>1</sup> *University of Victoria*

<sup>2</sup> *University of Victoria (CA)*

**Corresponding Authors:** colsond@uvic.ca, tssulliv@uvic.ca, rsobie@uvic.ca, mebert@uvic.ca

We present our unique approach to host the Canadian share of the Belle-II raw data and the computing infrastructure needed to process the raw data. We will describe the details of the storage system which is a disk-only storage solution based on xrootd and ZFS, as well as TSM for backup purpose. We will also detail the compute that involves starting specialized Virtual Machine (VMs) to process the raw data. We will discuss the merits of our approach as well as the issues observed during deployment and commissioning.

**Poster session / 185**

## Using S3 storage with xrootd vs Dynafed

**Authors:** Colson Driemel<sup>1</sup>; Marcus Ebert<sup>1</sup>; Randall Sobie<sup>2</sup>; Tristan Sullivan<sup>1</sup>

<sup>1</sup> *University of Victoria*

<sup>2</sup> *University of Victoria (CA)*

**Corresponding Authors:** rsobie@uvic.ca, tssulliv@uvic.ca, colsond@uvic.ca, mebert@uvic.ca

The HEP-RC group at UVic used Dynafed intensively to create federated storage clusters for Belle-II and ATLAS; which was used by worker nodes deployed on clouds around the world. Since the end of the DPM development also means the end of the development for Dynafed, xrootd was tested with S3 as backend to replace Dynafed. We will show similarities as well as major differences between the two systems as well as results of tests we run on both, for data transfers, checksum calculations as well as clustering of different endpoints. This may help other to efficiently make use of S3 storage as a WLCG site SE.

**Parallel (Track 7) / 186**

## The BaBar Long Term Data Preservation and Computing Infrastructure

**Authors:** Marcus Ebert<sup>1</sup>; Michael Roney<sup>None</sup>; Randall Sobie<sup>2</sup>

<sup>1</sup> *University of Victoria*

<sup>2</sup> *University of Victoria (CA)*

**Corresponding Authors:** rsobie@uvic.ca, mebert@uvic.ca, mroney@uvic.ca

BaBar stopped data taking in 2008 but its data is still analyzed by the collaboration. In 2021 a new computing system outside of the SLAC National Accelerator Laboratory was developed and major changes were needed to keep the ability to analyze the data by the collaboration, while the user facing front ends all needed to stay the same. The new computing system was put in production in 2022 and we will describe its unique infrastructure, based on cloud compute in Victoria, Canada, data storage at GridKa, Germany, streaming data access, as well as the possibility to analyze any data from anywhere. We will show advantages of the current system and how to run an old and outdated OS in current infrastructures, complications we faced when developing the system, as well as our experience in running and using it for about 2 years. It may be of interest to other groups and experiments when planning data preservation with the ability to continue to analyze the data, even decades after data taking has stopped.

**Plenary session / 187**

## Data Challenges in JUNO distributed computing infrastructure towards JUNO data-taking

**Author:** Xiaomei Zhang<sup>1</sup>

**Co-authors:** Andrea Rendina ; Giuseppe Andronico <sup>2</sup>; Joao Pedro Athayde Marcondes De Andre <sup>3</sup>; Nikolay Kutovskiy <sup>4</sup>; Rachid LEMRANI ALAOUI ; Xuantong Zhang <sup>1</sup>; Yifan Li <sup>5</sup>

<sup>1</sup> *Chinese Academy of Sciences (CN)*

<sup>2</sup> *Universita e INFN, Catania (IT)*

<sup>3</sup> *Centre National de la Recherche Scientifique (FR)*

<sup>4</sup> *JINR*

<sup>5</sup> *Institute of High Energy Physics*

**Corresponding Authors:** jpandre@iphc.cnrs.fr, giuseppe.andronico@cern.ch, andrea.rendina@cnaf.infn.it, rachid.lemrani@cc.in2p3.cnrs.fr, xiaomei.zhang@cern.ch, liyifan@ihep.ac.cn, xuantong.zhang@cern.ch, nikolay.kutovskiy@jinr.ru

The Jiangmen Underground Neutrino Observatory (JUNO) in southern China has set its primary goals as determining the neutrino mass ordering and precisely measuring oscillation parameters. JUNO plans to start data-taking in late 2024, with an expected event rate of approximately 1 kHz at full operation. This translates to around 60 MB of byte-stream raw data being produced every second, resulting in a 2 PB data per year. To address the challenges posed by this massive amount of data, JUNO is conducting data challenges on its distributed computing infrastructure. The data challenges aim to achieve several objectives, including understanding the offline requirements, accurately estimating the necessary resources, identifying potential bottlenecks within the involved systems, and improving overall performance. The ultimate goal is to demonstrate the effectiveness of the JUNO computing model and ensure the smooth operation of the entire data processing chain, encompassing raw data transfer, simulation, reconstruction, and analysis. Furthermore, the data challenges seek to verify the availability and effectiveness of monitoring systems for each activity.

**Parallel (Track 3) / 188**

## Hard scatter vertex identification in ATLAS using graph neural networks

**Authors:** Guglielmo Frattari<sup>1</sup>; Jackson Carl Burzynski<sup>2</sup>

<sup>1</sup> *Brandeis University (US)*

<sup>2</sup> *Simon Fraser University (CA)*

**Corresponding Authors:** [guglielmo.frattari@cern.ch](mailto:guglielmo.frattari@cern.ch), [jackson.carl.burzynski@cern.ch](mailto:jackson.carl.burzynski@cern.ch)

In the ATLAS experiment, colliding proton-proton bunches produce multiple primary vertices per bunch crossing. Typically, the primary vertex with the highest sum of squared transverse momentum of associated tracks is designated as the hard-scatter (HS) vertex, serving as the reference point for all physics objects in the event. However, this method proves suboptimal for scenarios with low track activity in the primary vertex, such as Higgs boson decay into photon pairs or low transverse momentum leptonic signatures. To address this, a novel approach employing graph neural networks has been developed by ATLAS. This new algorithm aims to enhance the efficiency of identifying the correct HS vertex in events by leveraging information from the complete set of reconstructed physics objects in the event. In this presentation, we will discuss the implementation of this algorithm, highlight its performance across various physics processes, and compare its effectiveness to existing algorithms used in ATLAS for HS vertex determination.

**Parallel (Track 1) / 189**

## **Adoption of ROOT RNTuple for the next main event data storage technology in the ATLAS production framework Athena**

**Authors:** Peter Van Gemmeren<sup>1</sup>; Alaettin Serhan Mete<sup>1</sup>; Marcin Nowak<sup>2</sup>; Tatiana Ovsianikova<sup>3</sup>

<sup>1</sup> *Argonne National Laboratory (US)*

<sup>2</sup> *Brookhaven National Laboratory (US)*

<sup>3</sup> *University of Washington (US)*

**Corresponding Authors:** [tatiana.ovsiannikova@cern.ch](mailto:tatiana.ovsiannikova@cern.ch), [peter.van.gemmeren@cern.ch](mailto:peter.van.gemmeren@cern.ch), [marcin.nowak@cern.ch](mailto:marcin.nowak@cern.ch), [alaettin.serhan.mete@cern.ch](mailto:alaettin.serhan.mete@cern.ch)

Since the start of LHC in 2008, the ATLAS experiment has relied on ROOT to provide storage technology for all its processed event data. Internally, ROOT files are organized around TTree structures that are capable of storing complex C++ objects. The capabilities of TTrees developed over the years and are now offering support for advanced concepts like polymorphism, schema evolution and user defined collections and ATLAS makes use of these features to handle its EDM. But some original TTrees concepts, like the POSIX file model and sequential writing, remain unchanged since the beginning and could be an obstacle to achieving the performance required for High Luminosity LHC.

With the HL-LHC performance goals in mind, the ROOT project developed a new storage format - the RNTuple. RNTuple, with its accompanying user API, is now in the final development stage and is planned to be production-ready at the end of 2024. Soon after that the TTree will become a legacy format.

ATLAS intends to have its main Event processing framework Athena ready to use RNTuple in the production environment as early as possible. The work on adopting RNTuple as another ROOT storage technology in Athena started already in 2021 and is now nearly complete. Although the initial goal was to focus on derived-AOD products (PHYS and PHYSLITE), with a little added effort all ATLAS data products: RDO, HITS, ESD, AOD and DAOD can be now stored in RNTuple format and transparently read back.

In this paper we will describe the current state of RNTuple adoption in the Athena framework and explain the ATLAS EDM requirements that had to be met on the ROOT side to successfully integrate both environments. We will demonstrate the ability to run standard ATLAS production workflows, based on RNTuple as the Event data storage technology, and point out key advantages of the new format.

**Parallel (Track 3) / 190**

## Improving Computational Performance of ATLAS GNN Track Reconstruction Pipeline

**Authors:** Alexis Vallier<sup>1</sup>; Alina Lazar<sup>2</sup>; Christophe COLLARD<sup>3</sup>; Daniel Thomas Murnane<sup>4</sup>; Heberth Torres<sup>1</sup>; Jackson Carl Burzynski<sup>5</sup>; Jan Stark<sup>6</sup>; Jared Burleson<sup>7</sup>; Jay Chan<sup>8</sup>; Mark Stephen Neubauer<sup>9</sup>; Minh-Tuan Pham<sup>10</sup>; Paolo Calafiura<sup>8</sup>; Sylvain Caillou<sup>11</sup>; Xiangyang Ju<sup>8</sup>

<sup>1</sup> *L2I Toulouse, CNRS/IN2P3, UT3*

<sup>2</sup> *Youngstown State University (US)*

<sup>3</sup> *Laboratoire des 2 Infinis - Toulouse, CNRS / Univ. Paul Sabatier*

<sup>4</sup> *Niels Bohr Institute, University of Copenhagen*

<sup>5</sup> *Simon Fraser University (CA)*

<sup>6</sup> *Laboratoire des 2 Infinis - Toulouse, CNRS / Univ. Paul Sabatier (FR)*

<sup>7</sup> *University of Illinois at Urbana-Champaign*

<sup>8</sup> *Lawrence Berkeley National Lab. (US)*

<sup>9</sup> *Univ. Illinois at Urbana-Champaign*

<sup>10</sup> *University of Wisconsin Madison (US)*

<sup>11</sup> *Centre National de la Recherche Scientifique (FR)*

**Corresponding Authors:** h.torres@cern.ch, sylvain.caillou@l2it.in2p3.fr, pcalafiura@lbl.gov, jackson.carl.burzynski@cern.ch, xiangyang.ju@cern.ch, jareddb2@illinois.edu, jay.chan@cern.ch, stark@in2p3.fr, msn@uiuc.edu, tuan.minh.pham@cern.ch, christophe.collard@l2it.in2p3.fr, allazar@cern.ch, alexis.vallier@cern.ch, dtmurnane@lbl.gov

Track reconstruction is an essential element of modern and future collider experiments, including the ATLAS detector. The HL-LHC upgrade of the ATLAS detector brings an unprecedented tracking reconstruction challenge, both in terms of the large number of silicon hit cluster readouts and the throughput required for budget-constrained track reconstruction. Traditional track reconstruction techniques often contain steps that scale combinatorially, which could be ameliorated with deep learning approaches. The GNN4ITk project has been shown to apply geometric deep learning algorithms for tracking to a similar level of physics performance with traditional techniques while scaling sub-quadratically. In this contribution, we compare the computational performance of a variety of pipeline configurations and machine learning inference methods. These include heuristic-and-ML-based graph segmentation techniques, GPU-based module map graph construction, and studies of high throughput graph convolutional kernels. In this contribution, we present benchmarks of latency, throughput, memory usage, and power consumption of each pipeline configuration.

**Parallel (Track 3) / 191**

## GNN based 3D calorimeter cell clustering pipeline for the ATLAS detector

**Authors:** Amir Farbin<sup>1</sup>; Debottam Bakshi Gupta<sup>1</sup>

<sup>1</sup> *University of Texas at Arlington (US)*

**Corresponding Authors:** debottam.bakshi.gupta@cern.ch, amir.farbin@cern.ch

We would like to propose a Graph Neural Network (GNN) based 3D calorimeter cell clustering pipeline for the ATLAS detector. The particles from the proton-proton collisions at the Large Hadron Collider (LHC) entering the ATLAS detector are reconstructed based on calorimeter cell energy clustering and inner detector tracks. Both calorimeter clustering and tracking can be viewed as a networking problem where particular cell energy (nodes) and its relation (edges) with the neighboring cells can be established as a graph. GNN has been proven to be a very powerful tool for artificial intelligence (AI) inference as it integrates both input data and the neighboring relationship among them.

Given promising results by Exa. TrkX in the reconstruction of particle tracks using GNN we would like to proceed with a similar approach for calorimeter clustering. The current and stable calorimeter clustering algorithm known as topoclustering has been used in both offline and trigger environments and it has several areas to improve e.g. CPU time, size of the cluster, hard-coded seeding and growth parameters for the cluster, and splitting of final cluster based on local maxima. With this GNN-based clustering approach, we expect to reduce CPU time by running inference in multi-threaded Athena environment (AthenaMT). Since the growth parameter of the clusters will be learned during the training we will have a better control on cluster size and local maxima inside it. The inference capability of this model will be independent of computing platforms, rendering it an ideal choice for integration into the future heterogeneous platform of the ATLAS experiment.

**Parallel (Track 1) / 192**

## **ATLAS WLCG Data Challenge 2024 planning and implementation**

**Authors:** Alessandra Forti<sup>1</sup>; Dimitrios Christidis<sup>2</sup>; Mario Lassnig<sup>2</sup>; Petr Vokac<sup>3</sup>; Shawn Mc Kee<sup>4</sup>; Vincent Garonne<sup>5</sup>

<sup>1</sup> *University of Manchester (GB)*

<sup>2</sup> *CERN*

<sup>3</sup> *Czech Technical University in Prague (CZ)*

<sup>4</sup> *University of Michigan (US)*

<sup>5</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** mario.lassnig@cern.ch, vincent.garonne@cern.ch, petr.vokac@cern.ch, shawn.mckee@cern.ch, alessandra.forti@cern.ch, dimitrios.christidis@cern.ch

ATLAS is participating in the WLCG Data Challenges, a bi-yearly program established in 2021 to prepare for the data rates of the High Luminosity HL-LHC. In each challenge, transfer rates are increased to ensure preparedness for the full rates by 2029. The goal of the 2024 Data Challenge (DC24) was to reach 25% of the HL-LHC expected transfer rates, with each experiment deciding how to execute the challenge based on agreed general guidelines and common dates. The ATLAS challenge was designed to test the ATLAS distributed infrastructure across 66 sites and was carried out over 12 days, with increasing rates and more complex transfer topologies, putting significant strain on the system. It was also the first time the new OAuth 2.0 authorization system was tested at such a large scale. This paper will discuss the planning of the challenge, the tools used to execute it, the agreed-upon transfer rates for the connections, and finally, the achieved results and any unachieved goals, along with an analysis of the bottlenecks. We will then describe how the challenge itself was executed, the results obtained, and the lessons learned. Finally, we will look ahead to the next challenge, currently scheduled for 2026, with 50% of HL-LHC rates.

**Poster session / 193**

## **ATLAS software tools to handle ROOT RNTuple**

**Authors:** Alaettin Serhan Mete<sup>1</sup>; Grigori Rybkin<sup>2</sup>; Marcin Nowak<sup>3</sup>; Peter Van Gemmeren<sup>1</sup>

<sup>1</sup> *Argonne National Laboratory (US)*

<sup>2</sup> *Université Paris-Saclay (FR)*

<sup>3</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** grigori.rybkine@cern.ch, peter.van.gemmeren@cern.ch, marcin.nowak@cern.ch, alaettin.serhan.mete@cern.ch

The software of the ATLAS experiment at the CERN LHC accelerator contains a number of tools to analyze (validate, summarize, peek into etc.) all its official data formats recorded in ROOT files. These tools - mainly written in the Python programming language - handle the ROOT TTree which is currently the main storage object format of ROOT files. However, the ROOT project has developed an alternative to TTree, called RNTuple. The new storage format offers significant improvements and ATLAS plans to adopt it in LHC Run 4. Work is ongoing to enhance the tools in order to handle the RNTuple storage format in addition to TTree in a transparent for the user way. The work is aided by modern and detailed APIs provided by RNTuple. We will present the progress made and lessons learnt.

**Poster session / 194**

## Impact of RNTuple on storage resources for ATLAS production

**Authors:** Alaettin Serhan Mete<sup>1</sup>; Marcin Nowak<sup>2</sup>; Peter Van Gemmeren<sup>1</sup>; Tatiana Ovsianikova<sup>3</sup>

<sup>1</sup> Argonne National Laboratory (US)

<sup>2</sup> Brookhaven National Laboratory (US)

<sup>3</sup> University of Washington (US)

**Corresponding Authors:** marcin.nowak@cern.ch, peter.van.gemmeren@cern.ch, tatiana.ovsiannikova@cern.ch, alaettin.serhan.mete@cern.ch

Over the past years, the ROOT team has been developing a new I/O format called RNTuple to store data from experiments at CERN's Large Hadron Collider. RNTuple is designed to improve ROOT's existing TTree I/O subsystem by improving I/O speed and introducing a more efficient binary data format. It can be stored in both ROOT files and object stores, and it's optimized for modern storage hardware like NVMe SSDs.

The ATLAS experiment plans to use RNTuple as its primary storage container in the upcoming HL-LHC.

There's been significant progress in integrating RNTuple into the ATLAS event processing framework, and now all production ATLAS data output formats support it. Performance studies with open-source data have shown substantial improvements in space resource usage. The reported study examines the I/O throughput and disk-space savings achieved with RNTuple for various ATLAS data output formats, including RDO, ESD, AOD, and various DAOD. These measurements will have an important impact on the computing resource needs of the ATLAS experiment for HL-LHC operation.

**Parallel (Track 3) / 195**

## High Performance Graph Segmentation for ATLAS GNN Track Reconstruction

**Authors:** Alexis Vallier<sup>1</sup>; Alina Lazar<sup>2</sup>; Daniel Thomas Murnane<sup>3</sup>; Daniel Whiteson<sup>4</sup>; Levi Condren<sup>4</sup>; Ryan Liu<sup>5</sup>; Xiangyang Ju<sup>5</sup>

<sup>1</sup> L2I Toulouse, CNRS/IN2P3, UT3

<sup>2</sup> Youngstown State University (US)

<sup>3</sup> Niels Bohr Institute, University of Copenhagen

<sup>4</sup> University of California Irvine (US)

<sup>5</sup> Lawrence Berkeley National Lab. (US)

**Corresponding Authors:** dtmurnane@lbl.gov, allazar@cern.ch, alexis.vallier@cern.ch, xiangyang.ju@cern.ch, lcondren@uci.edu, ryanliu@lbl.gov, daniel.whiteson@cern.ch

Graph neural networks and deep geometric learning have been successfully proven in the task of track reconstruction in recent years. The GNN4ITk project employs these techniques in the context of the ATLAS upgrade ITk detector to produce similar physics performance as traditional techniques, while scaling sub-quadratically. However, one current bottleneck in the throughput and physics performance of graph-based tracking is the final processing of classified graph edges into track candidates. This stage typically requires a trade-off between computational and physics performance. In this contribution, we present a variety of algorithms to ameliorate this trade-off, from heuristic approaches that use GPU-accelerated graph operations, to learned approaches such as hierarchical graph neural networks. Based on dedicated timing studies, we show that these graph segmentation algorithms are also well-suited to online track reconstruction.

**Parallel (Track 9) / 196**

## Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS

**Authors:** Alexander Held<sup>1</sup>; David Jordan<sup>2</sup>; Doug Benjamin<sup>3</sup>; Farnaz Golnaraghi<sup>2</sup>; Fengping Hu<sup>2</sup>; Gordon Watts<sup>4</sup>; Ilija Vukotic<sup>2</sup>; Judith Lorraine Stephen<sup>2</sup>; Lincoln Bryant<sup>2</sup>; Matthew Feickert<sup>1</sup>; Ofer Rind<sup>5</sup>; Robert William Gardner Jr<sup>2</sup>; Wei Yang<sup>6</sup>

<sup>1</sup> *University of Wisconsin Madison (US)*

<sup>2</sup> *University of Chicago (US)*

<sup>3</sup> *Brookhaven National Laboratory (US)*

<sup>4</sup> *University of Washington (US)*

<sup>5</sup> *Brookhaven National Laboratory*

<sup>6</sup> *SLAC National Accelerator Laboratory (US)*

**Corresponding Authors:** rind@bnl.gov, yangw@slac.stanford.edu, fengping@uchicago.edu, robert.w.gardner@cern.ch, judith.lorraine.stephen@cern.ch, djordan66@uchicago.edu, lincoln.bryant@cern.ch, douglas.benjamin@cern.ch, matthew.feickert@cern.ch, ilija.vukotic@cern.ch, gwatts@uw.edu, farnaz@uchicago.edu, alexander.held@cern.ch

The ATLAS experiment is currently developing columnar analysis frameworks which leverage the Python data science ecosystem. We describe the construction and operation of the infrastructure necessary to support demonstrations of these frameworks, with a focus on those from IRIS-HEP. One such demonstrator aims to process the compact ATLAS data format PHYSLITE at rates exceeding 200 Gbps. Various access configurations and setups on different sites are explored, including direct access to a dCache storage system via Xrootd, the use of ServiceX, and the use of multiple XCache servers equipped with NVMe storage devices. Integral to this study was the analysis of network traffic and bottlenecks, worker node scheduling and disk configurations, and the performance of an S3 object store. The system's overall performance was measured as the number of processing cores scaled to over 2,000 and the volume of data accessed in an interactive session approached 200 TB. The presentation will delve into the operational details and findings related to the physical infrastructure that underpins these demonstrators.

**Parallel (Track 9) / 197**

## Building Scalable Analysis Infrastructure for ATLAS

**Authors:** David Jordan<sup>1</sup>; Eric Christian Lancon<sup>2</sup>; Farnaz Golnaraghi<sup>1</sup>; Fengping Hu<sup>1</sup>; Ilija Vukotic<sup>1</sup>; Judith Lorraine Stephen<sup>1</sup>; Lincoln Bryant<sup>1</sup>; Robert William Gardner Jr<sup>1</sup>; Ryan Paul Taylor<sup>3</sup>

<sup>1</sup> *University of Chicago (US)*

<sup>2</sup> *Brookhaven National Laboratory (US)*

<sup>3</sup> *University of Victoria (CA)*

**Corresponding Authors:** judith.lorraine.stephen@cern.ch, lincoln.bryant@cern.ch, elancon@bnl.gov, ilija.vukotic@cern.ch, rptaylor@uvic.ca, djordan66@uchicago.edu, farnaz@uchicago.edu, robert.w.gardner@cern.ch, fengping@uchicago.edu

We explore the adoption of cloud-native tools and principles to forge flexible and scalable infrastructures, aimed at supporting analysis frameworks being developed for the ATLAS experiment in the High Luminosity Large Hadron Collider (HL-LHC) era. The project culminated in the creation of a federated platform, integrating Kubernetes clusters from various providers such as Tier-2 centers, Tier-3 centers, and from the IRIS-HEP Scalable Systems Laboratory, a National Science Foundation project. A unified interface was provided to streamline the management and scaling of containerized applications. Enhanced system scalability was achieved through integration with analysis facilities, enabling spillover of Jupyter/Binder notebooks and Dask workers to Tier-2 resources. We investigated flexible deployment options for a “stretched” (over the wide area network) cluster pattern, including a centralized “lights out management” model, remote administration of Kubernetes services, and a fully autonomous site-managed cluster approach, to accommodate varied operational and security requirements. The platform demonstrated its efficacy in multi-cluster demonstrators for low-latency analyses and advanced workflows with tools such as Coffea, ServiceX, Uproot and Dask, and RDataFrame, illustrating its ability to support various processing frameworks. The project also resulted in a robust user training infrastructure for ATLAS software and computing on-boarding events.

Parallel (Track 7) / 198

## Measurement and Analysis of ATLAS applications (CPU and Storage) energy consumption and data center optimization

**Authors:** Imran Latif<sup>1</sup>; Kevin Casella<sup>1</sup>; Qiulan Huang<sup>2</sup>; Robert Hancock<sup>1</sup>; Vincent Garonne<sup>2</sup>

<sup>1</sup> Brookhaven National Laboratory

<sup>2</sup> Brookhaven National Laboratory (US)

**Corresponding Authors:** qiulan.huang@cern.ch, hancock@bnl.gov, vincent.garonne@cern.ch, kac@bnl.gov, ilatif@bnl.gov

In this contribution, we will present the results of energy studies conducted on HEP payloads using the HEPsScore benchmark and utilizing the computing and storage resources available at BNL’s newly constructed state-of-the-art modern data center. Our objective is to gather various measures and metrics, including power consumption, CPU usage, disk I/O, and cooling temperature, among others, in order to identify additional parameters for fine-tuning. The goal is to reduce energy consumption while upholding the performance standards required by ATLAS.

This analysis will encompass diverse types of CPUs and disks. BNL’s data center offers an optimal setting for conducting such investigations, with power consumption, cooling, and other parameters continually monitored and controlled within each rack through a fully integrated Data Center Infrastructure Management system (DCIM).

The potential impact of this project is substantial. It has the potential to contribute to a reduction in energy consumption, thereby lowering costs and the associated carbon footprint. Moreover, it could furnish ATLAS with valuable insights into potential “green” energy metrics to consider in the development of ATLAS software or operational frameworks.

Parallel (Track 5) / 199

## Generative machine learning for fast silicon detector simulation

**Author:** Tadej Novak<sup>1</sup>



<sup>1</sup> *Jozef Stefan Institute (SI)*

**Corresponding Author:** tadej.novak@cern.ch

Simulation of physics processes and detector response is a vital part of high energy physics research but also representing a large fraction of computing cost. Generative machine learning is successfully complementing full (standard, Geant4-based) simulation as part of fast simulation setups improving the performance compared to classical approaches.

A lot of attention has been given to calorimeters being the slowest part of the full simulation, but their speed becomes comparable with silicon semiconductor detectors when fast simulation is used. This makes silicon detectors the next candidate to make faster, especially with the growing number of channels in future detectors.

This work studies the use of transformer architectures for fast silicon tracking detector simulation. The OpenDataDetector is used as a benchmark detector. Physics performance is estimated comparing reconstructed tracks using the ACTS tracking framework between full simulation and machine learning one.

**Parallel (Track 7) / 200**

## **An implementation of cloud-based grid CE and SE for ATLAS and Belle II**

**Author:** Jonathan Mark Woithe<sup>1</sup>

**Co-authors:** David Dossett ; Marcus Ebert <sup>2</sup>; Martin Sevier <sup>3</sup>; Paul Jackson <sup>4</sup>

<sup>1</sup> *University of Adelaide (AU)*

<sup>2</sup> *University of Victoria*

<sup>3</sup> *University of Melbourne (AU)*

<sup>4</sup> *University of Adelaide*

**Corresponding Authors:** mebert@uvic.ca, jwoithe@atrad.com.au, martin.sevier@cern.ch, paul.douglas.jackson@cern.ch, david.dossett@unimelb.edu.au

The economies of scale realised by institutional and commercial cloud providers make such resources increasingly attractive for grid computing. We describe an implementation of this approach which has been deployed for Australia's ATLAS and Belle II grid sites.

The sites are built entirely with Virtual Machines (VM) orchestrated by an OpenStack [1] instance. The Storage Element (SE) utilises an xrootd-s3 gateway [2][3] with back-end storage provided through an S3-compatible object store from a commercial provider. The provisioning arrangements required the deployment of some site-specific helper modules to ensure all SE interfacing requirements could be met. OpenStack hosts the xrootd redirector and proxy servers in separate VMs.

The Compute Element (CE) comprises virtual machines (VM) within the Openstack instance. Jobs are submitted and managed by HTCondor [4]. A CloudScheduler [5][6] instance is used to coordinate the number of active OpenStack VMs and ensure that VMs run only when there are jobs to run.

Automated configuration of the individual VMs associated with the grid sites is managed using Ansible [7]. This approach was chosen due to its low overheads and the simplicity of deployment.

Performance metrics of the resulting grid sites will be presented to illustrate the viability of this cost-effective approach to resource provisioning for grid computing.

## References

- [1] OpenStack: <https://www.openstack.org/>
- [2] Xrootd: <https://xrootd.slac.stanford.edu/>
- [3] Andrew Hanushevsky and Wei Yang: “Xrootd S3 Gateway for WLCG Storage”, 26th International Conference on Computing in High Energy & Nuclear Physics (CHEP 2023), <https://doi.org/10.1051/epjconf/2024295010>
- [4] HTCondor: <https://htcondor.org/htcondor/overview/>
- [5] CloudScheduler: <https://github.com/hep-gc/cloudscheduler>
- [6] Randall Sobie, F. Berghaus, K. Casteels, C. Driemel, M. Ebert, F. F. Galindo, C. Leavett-Brown, D. MacDonell, M. Paterson, R. Seuster, S. Tolkamp, J. Weldon: “cloudScheduler a VM provisioning system for a distributed compute cloud”, 24th International Conference on Computing in High-Energy and Nuclear Physics (CHEP 2019), <https://doi.org/10.1051/epjconf/202024507031>
- [7] Ansible: <https://www.ansible.com/>

## Poster session / 201

### Leveraging Language Models for Enhanced Code Review in Particle Physics Software Development

**Authors:** Alexey Rybalchenko<sup>1</sup>; Mohammad Al-Turany<sup>1</sup>

<sup>1</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

**Corresponding Authors:** [mohammad.al-turany@cern.ch](mailto:mohammad.al-turany@cern.ch), [alexey.rybalchenko@cern.ch](mailto:alexey.rybalchenko@cern.ch)

Collaborative software development for particle physics experiments demands rigorous code review processes to ensure maintainability, reliability, and efficiency. This work explores the integration of Large Language Models (LLMs) into the code review process, with a focus on utilizing both commercial and open models. We present a comprehensive code review workflow that incorporates LLMs, integrating various enhancements such as multi-agent capabilities and reflection. Furthermore, tools are employed to facilitate the verification of suggested code changes before presentation in the review. By harnessing the capabilities of LLMs, the review process can uncover faults and identify improvements that traditional automated analysis tools may overlook. This integration shows promise for improving code quality, reducing errors, and fostering collaboration among developers in the field of particle physics software development.

## Poster session / 202

### Application of quantum graph neural network on detector data

**Author:** Jihwan Oh<sup>None</sup>

**Corresponding Author:** [jihwan.oh@cern.ch](mailto:jihwan.oh@cern.ch)

We explore applications of quantum graph neural network (QGNN) on physics and non-physics data set. Based on a single quantum circuit architecture, we perform node, edge, and graph-level prediction tasks. Our main example is particle trajectory reconstruction starting from a set of detector data. Along with this, we expand our analysis on artificial helical trajectory data set. Finally, we will check how our quantum algorithm applies for non-physics data set as well by looking at Fingerprint data set in MUTAG, and AIDS data set, which collects molecular compounds graphs, focusing on graph-level task.

## Parallel (Track 3) / 203

## Fast Jet Reconstruction in Julia

**Authors:** Atell-Yehor Krasnopolski<sup>None</sup>; Benedikt Hegner<sup>1</sup>; Graeme A Stewart<sup>1</sup>; Philippe Gras<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> Université Paris-Saclay (FR)

**Corresponding Authors:** delta\_atell@protonmail.com, philippe.gras@cern.ch, graeme.andrew.stewart@cern.ch, benedikt.hegner@cern.ch

Jet reconstruction remains a critical task in the analysis of data from HEP colliders. We describe in this paper a new, highly performant, Julia package for jet reconstruction, JetReconstruction.jl, which integrates into the growing ecosystem of Julia packages for HEP. With this package users can run sequential reconstruction algorithms for jets, In particular, for LHC events, the Anti- $k_T$ , Cambridge/Aachen and Inclusive  $k_T$  algorithms can be used. For FCCee studies the use of alternative algorithms such as the generalised ee- $k_T$  and Durham are also supported.

The full reconstruction history is made available, allowing inclusive and exclusive jets to be retrieved. The package also provides the means to visualise the reconstruction.

The implementation of the package in Julia is discussed, with an emphasis on the features of the language that allow for an easy to work with, ergonomic, code implementation, that achieves high-performance. Julia's ecosystem offers the possibility to vectorise code, using single-instruction-multiple-data processing, in way that is transparent for the developer and more flexible than optimization done via C and C++ compilers. Thanks to this feature, the performance of JetReconstruction.jl is better than the current Fastjet C++ implementation in jet clustering for p-p events produced at the LHC.

Finally, an example of an FCCee analysis using JetReconstruction.jl is shown.

Parallel (Track 1) / 204

## Advancing Large-Scale Scientific Collaborations with Rucio: A Data Management Story

**Author:** Hugo Gonzalez Labrador<sup>1</sup>

**Co-author:** Martin Barisits<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** hugo.gonzalez.labrador@cern.ch, martin.barisits@cern.ch

Managing the data deluge generated by large-scale scientific collaborations is a challenge. The Rucio Data Management platform is an open-source framework engineered to orchestrate the storage, distribution, and management of massive data volumes across a globally distributed computing infrastructure. Rucio meets the requirements of high-energy physics, astrophysics, genomics, and beyond, pioneering new ways to facilitate research at the exabyte-scale.

This presentation introduces Rucio, highlighting its key features and strategic roadmap that underscore its flexibility towards diverse scientific domains, deep diving into concrete operational experience from various EU projects (ESCAPE, DaFab, InterTwin).

A special emphasis will be placed on the contributions of the CERN IT department, whose active engagement with the Rucio project has increased recently and catalysed significant contributions to the core software. This collaboration has not only enhanced Rucio's capabilities but also solidified its role in LHC experiments such as ATLAS and CMS, and provided a path forward for SMEs (Small and Medium experiments) to benefit from a converged data management platform.

**Parallel (Track 3) / 205**

## R&D towards heterogeneous frameworks for Future Experiments

**Authors:** Benedikt Hegner<sup>1</sup>; Mateusz Jakub Fila<sup>1</sup><sup>1</sup> CERN**Corresponding Authors:** benedikt.hegner@cern.ch, mateusz.jakub.fila@cern.ch

With the increasing amount of optimized and specialized hardware such as GPUs, ML cores, etc. HEP applications face the opportunity and the challenge of being enabled to take advantage of these resources, which are becoming more widely available on scientific computing sites. The Heterogeneous Frameworks project aims at evaluating new methods and tools for the support of both heterogeneous computational nodes and multi-node workloads. Based on the experience from the parallel frameworks of the LHC experiments and their ad-hoc support for heterogeneous resources, this project investigates newer libraries and languages that have been developed after the move to parallel frameworks about a decade ago.

This paper will summarize the scope of the problem being tackled, the state of the art of heterogeneous libraries, and the benchmark infrastructure used for the R&D activities. We will as well present some of the tooling developed to extract the benchmark scenarios from existing LHC experiment workflows. First results of using both newer C++ and Julia libraries for parallel execution will be shown.

**Poster session / 206**

## Particle Identification Algorithms Based on Machine Learning for STCF

**Author:** Yuncong Zhai<sup>1</sup>**Co-authors:** Teng LI<sup>2</sup>; Xingtao Huang<sup>1</sup> Shandong University<sup>2</sup> Shandong University, CN**Corresponding Authors:** tengli@sdu.edu.cn, huangxt@sdu.edu.cn, 202016991@mail.sdu.edu.cn

The Super Tau-Charm Facility (STCF) is the new generation  $e^+e^-$  collider aimed at studying tau-charm physics. The particle identification (PID), as one of the most fundamental tools for various physics research in STCF experiment, is crucial for achieving various physics goals of STCF. In the recent decades, machine learning (ML) has emerged as a powerful alternative for particle identification in HEP experiments. ML algorithms, such as neural networks and boosted decision trees, have shown superior performance in handling complex and multi-dimensional data, making them well-suited for integrating particle identification information from multiple sub-detector systems. In this work, we present a powerful PID software based on ML techniques, including a global PID algorithm for charged particles combining information from all sub-detectors, as well as a deep CNN discriminating neutral particles based on calorimeter responses. The preliminary results show the PID models has achieved excellent PID performance, greatly boosting the physics potential of STCF.

**Poster session / 207**

## TechWeekStorage24: Open Source Data Technologies for the Exabyte Era

**Author:** Jakub Moscicki<sup>1</sup>

**Co-authors:** Andreas Joachim Peters <sup>1</sup>; Luca Mascetti <sup>1</sup>; Michael Davis <sup>1</sup>; Oliver Keeble <sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** oliver.keeble@cern.ch, michael.davis@cern.ch, andreas.joachim.peters@cern.ch, luca.mascetti@cern.ch, jakub.moscicki@cern.ch

TechWeekStorage24 was introduced by CERN IT Storage and Data Management group as a new “Center of Excellence” community networking format: a co-located series of events on Open Source Data Technologies, bringing together a wide range of communities, far beyond High Energy Physics and highlighting the wider technology impact of IT solutions born in HEP.

Combining the annual CS3 conference, CERN Storage Day, EOS and CTA Workshops created week-long opportunity for connection, collaboration and discussion on storage services, open source software-defined storage and data management technology, data policies & trends, innovative applications, collaboration platforms, digital sovereignty, FAIR and Open Science, Security and Privacy of Data and more.

This new event format is also environmentally more sustainable: participants from locations such as Brazil, US, China, Japan, Korea had an opportunity to attend multiple related events within a single trip.

<https://techweekstorage.web.cern.ch>

**Poster session / 208**

## dCache-CTA Integration: One Year in Production at DESY

**Authors:** Jacek Chodak<sup>None</sup>; Mwai Karimi<sup>None</sup>; Peter Suchowski<sup>None</sup>; Ralf Lueken<sup>None</sup>; Tigran Mkrtchyan<sup>None</sup>

**Corresponding Authors:** jacek.chodak@desy.de, ralf.lueken@desy.de, tigran.mkrtchyan@desy.de, peter.suchowski@desy.de, mwai.karimi@desy.de

The huge volume of data generated by scientific facilities such as EuXFEL or LHC places immense strain on the data management infrastructure within laboratories. This includes poorly shareable resources of archival storage, typically, tape libraries. Maximising the efficiency of these tape resources necessitates a deep integration between hardware and software components.

CERN’s Tape Archive (CTA) is an open-source storage management system developed by CERN to handle LHC data on tape. Although the primary target of CTA is CERN Tier-0, the Data Management Group considers CTA as the compelling alternative to commercial Hierarchical Storage Management (HSM) systems.

dCache, with its adaptable tape interface allows connectivity to any tape system. Collaborating closely with the CERN Tape Archive team, we have been working on the seamless integration of CTA into the dCache ecosystem.

This work shows the design, current progress, and initial deployment experiences of the dCache-CTA integration at DESY.

**Poster session / 209**

## Evaluating FPGA Acceleration with Intel® oneAPI Toolkit for High-Speed Data Processing

**Author:** Alberto Perro<sup>1</sup>

**Co-authors:** Flavio Pisani<sup>2</sup>; Paolo Durante<sup>2</sup>

<sup>1</sup> *Universite d'Aix-Marseille III (FR)*

<sup>2</sup> *CERN*

**Corresponding Authors:** paolo.durante@cern.ch, alberto.perro@cern.ch

The LHCb Experiment employs GPU cards in its first level trigger system to enhance computing efficiency, achieving a data rate of 40Tb/s from the detector. GPUs were selected for their computational power, parallel processing capabilities, and adaptability.

However, trigger tasks necessitate extensive combinatorial and bitwise operations, ideally suited for FPGA implementation. Yet, FPGA adoption for compute acceleration is hindered by steep learning curves and very different programming paradigms with respect to GPUs and CPUs. In the last few years, interest in high level synthesis has grown because of the possibility of developing FPGA gateware in higher-level languages.

This study assesses the Intel® oneAPI FPGA Toolkit, which aims to simplify the development of FPGA-accelerated workloads by offering a GPU-like programming framework. We detail the integration of a portion of the current pixel clustering algorithm into oneAPI, address common implementation challenges, and compare it against CPU, GPU, and RTL implementations.

Our findings showcase promising outcomes for this emerging technology, potentially facilitating the repurposing of FPGAs in the data acquisition system as compute accelerators during idle data-taking periods.

**Parallel (Track 5) / 210**

## Generalization and self supervision for physics simulation

**Authors:** Anna Zaborowska<sup>1</sup>; Kristina Jaruskova<sup>1</sup>; Piyush Raikwar<sup>1</sup>; Renato Paulo Da Costa Cardoso<sup>1</sup>; Sofia Vallecorsa<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** sofia.vallecorsa@cern.ch, kristina.jaruskova@cern.ch, piyush.raikwar@cern.ch, renato.cardoso@cern.ch, anna.zaborowska@cern.ch

Foundation models are models trained on large and diverse datasets, and capable of adapting to a range of different tasks through self-supervised and transfer learning techniques. These models, characterized by their large number of parameters, excel in generating robust data representations, they make use of transformers, used as a generalized architecture for various data modalities, including text (e.g., BERT, GPT3), time series (e.g., PatchTST), and images (e.g., ViT), thanks to their ability to learn superior representations with ample data, absence of inductive bias, and better modeling of long-range dependencies. Our proposed approach builds upon this premise, introducing a physics simulation model inspired by foundation models, initially trained on a generative task but capable of seamless adaptation to various downstream tasks. Our model not only achieves competitive results in the generative task but also demonstrates relevance in regression tasks, all while operating efficiently without the extensive parameterization previously assumed necessary.

**Poster session / 211**

## Prometheus-Powered Insight: Monitoring Koji's performance

**Author:** Marta Vila Fernandes<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** ana.marta.lisboa.vila.fernandes@cern.ch

Efficient, ideally fully automated, software package building is essential in the computing supply chain of the CERN experiments. With Koji, a very popular package software building system used in the upstream Enterprise Linux communities, CERN IT provides a service to build software and images for the Linux OSes we support. Due to the criticality of the service and the limitations in Koji's built-in monitoring, the CERN Linux team implemented new functionality to allow integration with Prometheus, an open-source monitoring system and time-series database. This contribution will give an overview of Koji and its integration with Prometheus and Grafana, explain the challenges we tackled during the development of the integration, and how we're benefiting from these new metrics to improve the quality of the service.

Poster session / 212

## Managing accelerator control systems with Kubernetes at CERN

**Authors:** Diana Gaponcic<sup>None</sup>; Ricardo Rocha<sup>1</sup>; Spyridon Trigazis<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** diana.gaponcic@cern.ch, ricardo.rocha@cern.ch, spyridon.trigazis@cern.ch

CERN IT has offered a Kubernetes service since 2016, expanding to incorporate multiple other technologies from the cloud native ecosystem over time. Currently the service runs over 500 clusters and thousands of nodes serving use cases from different sectors in the organization.

In 2021 the ATS sector showed interest in looking at a similar setup for their container orchestration effort. A collaboration was started with an initial proof of concept running the CERN IT service inside the control room datacenter, including use cases from multiple teams in the sector. Following a successful initiative that ran over a year, a second phase was launched to bring the service to production.

In this paper we describe the existing CERN IT service and the major changes and improvements that were required to serve accelerator control use cases. We highlight the changes due to running in an isolated, air-gapped network environment, as well as the additional integrations regarding identity, storage and datacenter infrastructure. Finally we detail results from an extensive effort for failure scenario evaluation to comply with the expected service levels, as well as plans for extending the existing infrastructure to new use cases.

Parallel (Track6) / 213

## Extending the Gaudi Software Framework outside of C++

**Author:** Marco Clemencic<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** marco.clemencic@cern.ch

The LHCb Software Framework Gaudi has been developed in C++ since 1998. Over the years it evolved following the changes in the C++ established best practices and the evolution of the C++ standard, even reaching the point of enabling the development of multi-threaded applications. In the past few years there has been several announcements and debates over the so called C++ successor languages and safe alternatives to C++, with Rust leading the way as an example of safe and performing language that can replace C and C++ in a number of cases. This paper explores some ways Rust can be used to extend the Software Framework Gaudi, focusing on how one can leverage on the Rust-C++ interoperability efforts driven by the community. We show how to invoke Rust code from C++ and vice versa, and how Gaudi components could be written completely in Rust. We can use the experience gained in the exercise to evaluate possible integration with other languages or technologies, like WASM.

Parallel (Track 1) / 214

## Archive Metadata for efficient data colocation on tape

**Author:** Julien Leduc<sup>1</sup>

**Co-authors:** Jaroslav Guenther<sup>1</sup>; Joao Afonso<sup>1</sup>; Konstantina Skovola; Michael Davis<sup>1</sup>; Niels Alexander Bugel; Pablo Oliver Cortes<sup>1</sup>; Richard Bachmann<sup>1</sup>; Sergio Alfageme Perez; Vladimir Bahyl<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** julien.leduc@cern.ch, sergio.alfageme@cern.ch, niels.alexander.bugel@cern.ch, vladimir.bahyl@cern.ch, jaroslav.guenther@cern.ch, joao.afonso@cern.ch, konstantina.skovola@cern.ch, michael.davis@cern.ch, richard.bachmann@cern.ch, pablo.oliver.cortes@cern.ch

Due to the increasing volume of physics data being produced, the LHC experiments are making more active use of archival storage. Constraints on available disk storage have motivated the evolution towards the “data carousel” and similar models. Datasets on tape are recalled multiple times for reprocessing and analysis, and this trend is expected to accelerate during the Hi-Lumi era (LHC Run-4 and beyond).

Currently, storage endpoints are optimised for efficient archival, but it is becoming increasingly important to optimise for efficient retrieval. This problem has two dimensions. To reduce unnecessary tape mounts, the spread of each dataset - *the number of tapes containing files which will be recalled at the same time* - should be minimised. To reduce seek times, files from the same dataset should be physically colocated on the tape. The Archive Metadata specification is an agreed format for experiments to provide scheduling and colocation hints to storage endpoints to achieve these goals.

This contribution describes the motivation, the review process with the various stakeholders and the constraints that led to the Archive Metadata proposal. We present the implementation and deployment in the CERN Tape Archive and our preliminary experiences of consuming Archive Metadata at WLCG Tier-0.

Parallel (Track 2) / 215

## An online data processing system for the CMS Level-1 Trigger data scouting demonstrator

**Author:** Matteo Migliorini<sup>1</sup>

<sup>1</sup> *Universita e INFN, Padova (IT)*

**Corresponding Author:** m.migliorini@cern.ch



The CMS Level-1 Trigger Data Scouting (L1DS) introduces a novel approach within the CMS Level-1 Trigger (L1T), enabling the acquisition and processing of L1T primitives at the 40 MHz LHC bunch-crossing (BX) rate. The target for this system is the CMS Phase-2 Upgrade for the High Luminosity phase of LHC, harnessing the improved Phase-2 L1T design, where tracker and high-granularity calorimeter data will be available for the first time. Currently, a L1DS demonstrator is operating during LHC Run3, collecting data from the CMS L1T system. This contribution focuses on the online processing system of the Run3 scouting demonstrator, introduced for the first time in the beginning of 2024. Its function is to aggregate data fragments from all scouting sources (event building), perform online analysis/selection, and generate datasets for offline analysis. Contrary to the standard CMS data-taking, data fragments and processing are based on the LHC orbit, rather than individual events per BX. This allows the system to work at a constant rate of about 11kHz, while opening the possibility of exploring multi-BX correlations. An overview of the L1DS system will be provided, including its architecture, performance, and first results.

Poster session / 216

## PATOF: From the Past To the Future: Legacy Data in Small and Medium-Scale “PUNCH” Experiments - a Blueprint for PUNCH and Other Disciplines

**Author:** Ding-Ze Hu<sup>None</sup>

**Co-authors:** Martin Köhler<sup>1</sup>; UNKNOWN UNKNOWN

<sup>1</sup> DESY

**Corresponding Authors:** thomas.schoerner@desy.de, ding-ze.hu@desy.de, martin.koehler@desy.de

The PATOF project builds on work at MAMI particle physics experiment A4. A4 produced a stream of valuable data for many years which already released scientific output of high quality and still provides a solid basis for future publications. The A4 data set consists of 100 TB and 300 million files of different types (Vague context because of hierarchical folder structure and file format with minimal metadata provided)

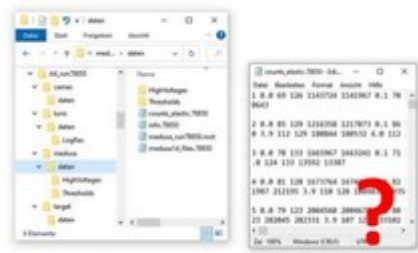


Figure 1: Figure 1:vague context because of hierarchical folder structure and file format with minimal metadata provided

).

In PATOF we would like to build a “FAIR Metadata Factory

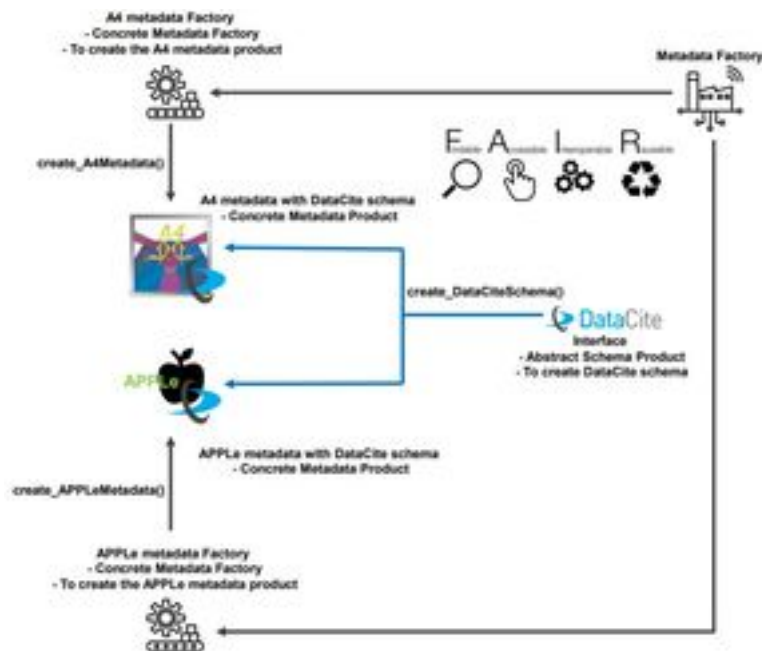


Figure 2: Through the Metadata Factory Pattern, we define a framework that allows us to create related or dependent products that follow a general pattern.

”, i.e. a process to create a naturally evolved metadata schema that can be used across research fields. The first focus will be on creating machine-readable XML files containing metadata from the logbook and other sources and to further enrich them

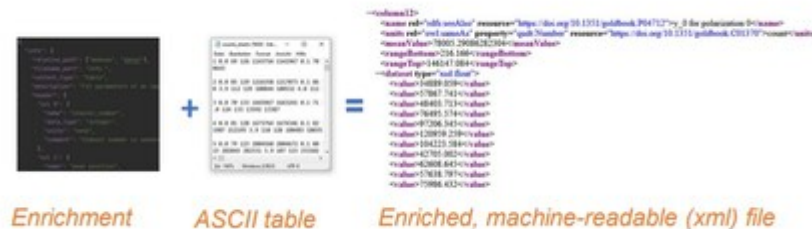


Figure 3: Enriching data from the logbook by providing metadata

In PATOF, we intend to conclude the work on A4 data, to extract the lessons learned there in the form of a cookbook that can capture the methodology for making individual experiment-specific metadata schemas FAIR, and to apply it to four other experiments: The ALPS II axion and dark matter search experiment at DESY. The PRIMA experiment at MAMI in Mainz for measuring the pion transition form factor. The upcoming nuclear physics experiment P2 at MESA in Mainz. Finally, the LUXE experiment at DESY planned to start in 2026. The focus of PATOF is on making these data fully publicly available.

The objectives of the project are i) a FAIR Metadata Factory (i.e. a cookbook of (meta)data management recommendations), and ii) the FAIRification of data from concrete experiments. Both aspects are inherently open in nature so that everybody can profit from PATOF results. The cookbook is expected to be further enhanced with contributions from other experiments even after PATOF (“living cookbook”).

## Poster session / 217

## Facilitating Scientific Reproducibility and Interoperability through CWL Integration in the Dirac Grid Middleware

**Authors:** Alexandre Franck Boyer<sup>1</sup>; Luisa ARRABITO<sup>2</sup>; Nathan PIGOUX<sup>None</sup>

<sup>1</sup> CERN

<sup>2</sup> LUPM IN2P3/CNRS

**Corresponding Authors:** arrabito@in2p3.fr, natthan.pigoux@lupm.in2p3.fr, alexandre.franck.boyer@cern.ch

Dirac, a versatile grid middleware framework, is pivotal in managing computational tasks and workflows across a spectrum of scientific research domains including high energy physics and astrophysics. Historically, Dirac has employed specialized descriptive languages that, while effective, have introduced significant complexities and barriers to workflow interoperability and reproducibility. These challenges have become particularly pressing in light of the reproducibility crisis - an ongoing and pervasive issue that surfaced prominently in the early 2010s, marked by difficulties in replicating scientific results across different studies.

In response to these challenges, the integration of the Common Workflow Language (CWL) into Dirac represents a transformative development. CWL is a specification dedicated to the unambiguous definition and execution of computational workflows, facilitating their shareability and reusability across diverse computing environments. Its adoption within Dirac aims to standardize the description of computational tasks, thereby enhancing both reproducibility and interoperability.

By streamlining the interface for defining computational tasks within Dirac, we enable researchers to effortlessly transition workflows from local to grid-scale environments and foster compatibility with a broader ecosystem of scientific tools. This integration promises not only to mitigate the challenges posed by the reproducibility crisis but also to significantly lower the threshold for engaging with complex computational infrastructures, thus accelerating scientific discovery and innovation across multiple disciplines.

## Parallel (Track 2) / 218

## Architecting software applications in containerized environment for CMS data acquisition

**Authors:** Dainius Simelevicius<sup>1</sup>; Luciano Orsini<sup>2</sup>

**Co-authors:** Albert Corominas I Mariscot<sup>2</sup>; Andrea Bocci<sup>2</sup>; Andrea Petrucci<sup>3</sup>; Antonin Dvorak<sup>2</sup>; Antra Gaile<sup>2</sup>; Attila Racz<sup>2</sup>; Babatunde John Odetayo<sup>4</sup>; Christian Deldicque<sup>2</sup>; Christoph Paus<sup>5</sup>; Christoph Schwick<sup>2</sup>; Christos Emmanouil<sup>2</sup>; Cristina Vazquez Velez<sup>2</sup>; Dinyar Sebastian Rabady<sup>2</sup>; Dominique Gigi<sup>2</sup>; Emilio Meschi<sup>2</sup>; Eric Cano<sup>2</sup>; Frank Glege<sup>2</sup>; Frans Meijers<sup>2</sup>; Georgiana-Lavinia Darlea<sup>5</sup>; Guillermo Gomez-Ceballos<sup>5</sup>; Guillermo Izquierdo Moreno<sup>2</sup>; Hannes Sakulin<sup>2</sup>; Jaafar Alawieh<sup>2</sup>; James Gordon Branson<sup>3</sup>; Jan Andrzej Bugajski<sup>2</sup>; Jeroen Hegeman<sup>2</sup>; Kareen Arutjunjan<sup>2</sup>; Kenneth Long<sup>5</sup>; Marc Dobson<sup>2</sup>; Marco Pieri<sup>3</sup>; Miguel Bacharov Durasov<sup>2</sup>; Patrycja Gorniak<sup>2</sup>; Petr Zejdl<sup>2</sup>; Philipp Maximilian Brummer<sup>2</sup>; Polyneikis Tzanis<sup>2</sup>; Rafal Krawczyk<sup>6</sup>; Sergio Cittolin<sup>3</sup>; Srecko Morovic<sup>3</sup>; Tejeswini Jayakumar<sup>2</sup>; Theodoros Rizopoulos<sup>2</sup>; Thomas Owen James<sup>2</sup>; Ulf Behrens<sup>6</sup>; Wassef Karimeh<sup>2</sup>; Wei Li<sup>6</sup>

<sup>1</sup> CERN, Vilnius University

<sup>2</sup> CERN

<sup>3</sup> UCSD

<sup>4</sup> CERN, University of Benin

<sup>5</sup> MIT

<sup>6</sup> Rice University

The data acquisition (DAQ) system stands as an essential component within the CMS experiment at CERN. It relies on a large network system of computers with demanding requirements on control, monitoring, configuration and high throughput communication. Furthermore, the DAQ system must accommodate various application scenarios, such as interfacing with external systems, accessing custom electronics devices for data readout, and event building. We present a versatile and highly modular programmable C++ framework designed for crafting applications tailored to various needs, facilitating development through the composition and integration of modules to achieve the desired DAQ capabilities. This framework takes advantage of reusable components and readily available off-the-shelf technologies. Applications are structured to seamlessly integrate into a containerized ecosystem, where the hierarchy of components and their aggregation is specified to form the final deployable unit to be used across multiple computers or nodes within an orchestrating environment. The utilization of the framework, along with the containerization of applications, enables coping with the complexity of implementing the CMS DAQ system by providing standardized structures and components to achieve a uniform and consistent architecture.

Parallel (Track 7) / 219

## A Lightweight Analysis & Grid Facility for the DARWIN Experiment

**Author:** Sebastian Brommer<sup>1</sup>

**Co-authors:** Benedikt Maier<sup>2</sup>; Manuel Giffels<sup>1</sup>; Markus Klute<sup>3</sup>; Robin Hofsaess<sup>1</sup>

<sup>1</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>2</sup> *Imperial College (GB)*

<sup>3</sup> *Karlsruhe Inst. of Technology (GER)*

**Corresponding Authors:** sebastian.brommer@cern.ch, robin.hofsaess@cern.ch, manuel.giffels@cern.ch, benedikt.maier@cern.ch, markus.klute@cern.ch

A robust computing infrastructure is essential for the success of scientific collaborations. However, smaller or newly founded collaborations often lack the resources to establish and maintain such an infrastructure, resulting in a fragmented analysis environment with varying solutions for different members. This fragmentation can lead to inefficiencies, hinder reproducibility, and create challenges for the collaboration.

We present an analysis facility for the DARWIN (DARK matter WImp search with liquid xenon) observatory, a new experiment that is currently in its R&D phase. The facility is designed to be lightweight with minimal administrative overhead while providing a common entry point for all DARWIN collaboration members. The setup serves as a blueprint for other collaborations, that want to provide a common analysis facility for their members. Grid computing and storage resources are integrated into the facility, allowing for distributed computing and a common entry point for storage. The authentication and authorization infrastructure for all services is token-based, using an Indigo IAM instance.

This talk will discuss the architecture of the facility, its provided services, first experiences of the DARWIN collaboration, and how it can serve as a sustainable blueprint for other collaborations.

Parallel (Track6) / 220

## Thread-safe N-tuple Writing in Gaudi with TTree and Migration to RNTuple

**Authors:** Marco Clemencic<sup>1</sup>; Silia Taider<sup>2</sup>

<sup>1</sup> CERN<sup>2</sup> CPE (FR)**Corresponding Authors:** silia.taider@cern.ch, marco.clemencic@cern.ch

The software framework of the Large Hadron Collider Beauty (LHCb) experiment, Gaudi, heavily relies on the ROOT framework and its I/O subsystems for data persistence mechanisms. Gaudi internally leverages the ROOT TTree data format, as it is currently used in production by LHC experiments. However, with the introduction and scaling of multi-threaded capabilities within Gaudi, the limitations of TTree as a data storage backend have become increasingly apparent, marking it as a non-negligible bottleneck in data processing workflows.

The following work introduces a comprehensive two-part enhancement to Gaudi to address this challenge. An initial focus is given to optimizing the current n-tuple writing infrastructure to be thread-safe within the constraints of the existing TTree backend, thus maintaining compatibility for users and downstream applications. This phase is then followed by the migration of the n-tuple storage backend from TTree to RNTuple, ROOT's next-generation I/O subsystem for physics data storage. This migration aims at leveraging the thread-safe, asynchronous capabilities of the new data format, thus making Gaudi fit to handle the requirements of HL-LHC computing and beyond.

**Keywords:** LHCb; Gaudi; ROOT; TTree; RNTuple; thread-safety**Parallel (Track 3) / 221**

## Next generation geometry model for Tracking in ACTS

**Authors:** Andreas Salzburger<sup>1</sup>; Paul Gessinger<sup>1</sup><sup>1</sup> CERN**Corresponding Authors:** paul.gessinger@cern.ch, andreas.salzburger@cern.ch

ACTS is an experiment independent toolkit for track reconstruction, which is designed from the ground up for thread-safety and high performance. It is built to accommodate different experiment deployment scenarios, and also serves as community platform for research and development of new approaches and algorithms.

A fundamental component of ACTS is the geometry library. It models a simplified representation of a detector, compared to simulation geometries. It drives the numerical track extrapolation, provides crucial inputs to track finding and fitting algorithms, and is connected to many other geometry libraries in the ecosystem, shipping with multiple plugins.

ACTS' geometry library is historically optimized for symmetric, collider-like detectors and most suitable for arrangements of silicon sensors. An effort has been underway for some amount of time to rewrite large parts of the geometry code.

The goal is to be more flexible to accommodate other detector approaches and simplify the building process, while providing easy conversion to a GPU-optimized geometry for use with the detrax library. Another goal is to allow for a more systematic way to write geometry plugins.

Finally, the navigation logic is delegated to detector regions, so that it can be easily extended for unconventional environments.

This contribution reports on the result of this rewrite, discusses lessons learned from the project and how they were incorporated into a robust geometry modeling solution in ACTS that will be key going forward.

**Parallel (Track 7) / 222**

## Improving overall GPU sharing and usage efficiency with Kubernetes

**Authors:** Dejan Golubovic<sup>None</sup>; Diana Gaponcic<sup>1</sup>; Diogo Filipe Tomas Guerra<sup>2</sup>; Ricardo Rocha<sup>2</sup>

<sup>1</sup> *IT-PW-PI*

<sup>2</sup> *CERN*

**Corresponding Authors:** diogo.filipe.tomas.guerra@cern.ch, digaponc@cern.ch, dejan.golubovic@cern.ch, ricardo.rocha@cern.ch

GPUs and accelerators are changing traditional High Energy Physics (HEP) deployments while also being the key to enable efficient machine learning. The challenge remains to improve overall efficiency and sharing opportunities of what are currently expensive and scarce resources.

In this paper we describe the common patterns of GPU usage in HEP, including spiky requirements with low overall usage for interactive access, as well as more predictable but potentially bursty workloads including distributed machine learning. We then explore the multiple mechanisms to share and partition GPUs, covering time slicing, virtualization, physical partitioning (MIG) and MPS for Nvidia devices.

We conclude with the results of an extensive set of benchmarks for multiple representative HEP use cases, including traditional GPU usage as well as machine learning. We highlight the limitations of each option and the use cases where they fit best. Finally, we cover the deployment aspects and the different options available targeting a centralized GPU pool that can significantly push the overall GPU usage efficiency.

Parallel (Track 4) / 223

## Preliminary findings and recommendations from the Token Trust and Traceability Working Group.

**Authors:** David Kelsey<sup>1</sup>; Linda Ann Cornwall<sup>None</sup>; Maarten Litmaath<sup>2</sup>; Matt Doidge<sup>3</sup>; Tom Dack<sup>None</sup>

**Co-author:** Token Trust and Traceability WG <sup>4</sup>

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

<sup>2</sup> *CERN*

<sup>3</sup> *Lancaster University (GB)*

<sup>4</sup> *WLCG, EGI*

**Corresponding Authors:** linda.cornwall@stfc.ac.uk, token-trust-and-traceability-wg@cern.ch, david.kelsey@stfc.ac.uk, tom.dack@cern.ch, maarten.litmaath@cern.ch, m.doidge@cern.ch

Created in 2023, the Token Trust and Traceability Working Group (TTT) was formed in order to answer questions of policy and best practice with the ongoing move from X.509 and VOMS proxy certificates to token-based solutions as the primary authorisation and authentication method in grid environments. With a remit to act in an investigatory and advisory capacity alongside other working groups in the token space, the TTT is composed of a broad variety of stakeholders to provide a breadth of experience and viewpoints. The requirements of grid sites, users, identity providers and virtual organisations to be able to trace workflows remain largely the same in a token paradigm as when using X.509 certificates, while tokens provide a new set of challenges, requiring a rethink and restructure of the policies and processes that were defined with just X.509 and VOMS in mind.

After providing an overview of the current status of the token trust landscape we will detail the initial findings, future plans and recommendations to be made by the TTT. This will include best practice for sites and identity providers, suggestions for token development, and methodologies for tracing token usage by system administrators within common grid middleware stacks.

**Poster session / 225**

## Portable HCAL reconstruction in the CMS detector using the Alpaka library

**Author:** Ka Hei Martin Kwok<sup>1</sup><sup>1</sup> *Fermi National Accelerator Lab. (US)***Corresponding Author:** ka.hei.martin.kwok@cern.ch

CMS has deployed a number of different GPU algorithms at the High-Level Trigger (HLT) in Run 3. As the code base for GPU algorithms continues to grow, the burden for developing and maintaining separate implementations for GPU and CPU becomes increasingly challenging. To mitigate this, CMS has adopted the Alpaka (Abstraction Library for Parallel Kernel Acceleration) library as the performance portability solution to provide a single-code base for parallel execution on both GPUs and CPUs in CMS software (CMSSW).

A direct CUDA version of HCAL energy reconstruction, called Minimization At Hcal, Iteratively (MAHI), has been deployed at the HLT in the 2022-2023 data taking period. This contribution will describe how the CUDA version is converted into a portable implementation using the Alpaka library. We will discuss the porting experience from CUDA to Alpaka, the validation process and the performance of the Alpaka version in CPU and GPU.

**Parallel (Track 5) / 226**

## AdePT - Enabling GPU electromagnetic transport with Geant4

**Authors:** Andrei Gheata<sup>1</sup>; Benjamin Morgan<sup>2</sup>; John Apostolakis<sup>1</sup>; Jonas Hahnfeld<sup>3</sup>; Juan Gonzalez Caminero<sup>None</sup>; Severin Diederichs<sup>1</sup>; Stephan Hageboeck<sup>1</sup>; Witold Pokorski<sup>1</sup><sup>1</sup> *CERN*<sup>2</sup> *University of Warwick*<sup>3</sup> *CERN & Goethe University Frankfurt***Corresponding Authors:** jonas.hahnfeld@cern.ch, witold.pokorski@cern.ch, andrei.gheata@cern.ch, severin.diederichs@cern.ch, benjamin.morgan@cern.ch, john.apostolakis@cern.ch, stephan.hageboeck@cern.ch, j.gonzalez@cern.ch

An important alternative for boosting the throughput of simulation applications is to take advantage of accelerator hardware, by making general particle transport simulation for high-energy physics (HEP) single-instruction-multiple-thread (SIMT) friendly. This challenge is not yet resolved due to difficulties in mapping the complexity of Geant4 components and workflow to the massive parallelism features exposed by graphics processing units (GPU). The AdePT project is one of the R&D initiatives tackling this limitation and exploring GPUs as potential accelerators for offloading part of the CPU simulation workload. Our main target is the implementation of a complete electromagnetic shower transport engine working on the GPU. A first development phase, allowed us to verify our GPU prototype against the Geant4 simulation for both simplified and complex setups, and to test different Geant4 integration strategies. We have simplified the integration procedure of AdePT as an external library in both standalone applications and experimental frameworks through standard Geant4 mechanisms. The project's current main focus is to provide solutions for the main performance bottlenecks identified so far: inefficient geometry modeling for the GPUs, and a suboptimal CPU-GPU scheduling strategy. We will present the most recent results and conclusions of our work, focusing on the hybrid Geant4-AdePT use case.

**Parallel (Track 7) / 227**

## Towards an Introspective Dynamic Model of Globally Distributed Computing Infrastructures

**Authors:** Adolfo Hoisie<sup>1</sup>; Alexei Klimentov<sup>2</sup>; Frederic Suter<sup>3</sup>; Joseph Boudreau<sup>4</sup>; Norbert Podhorszki<sup>3</sup>; Ozgur Kilic<sup>1</sup>; Paul Nilsson<sup>2</sup>; Scott Klasky<sup>None</sup>; Shengyu Feng<sup>5</sup>; Tadashi Maeno<sup>2</sup>; Tasnuva Chowdhury<sup>6</sup>; Tatiana Korchuganova<sup>4</sup>; Verena Ingrid Martinez Outschoorn<sup>7</sup>; Yihui Ren<sup>2</sup>; Yiming Yang<sup>5</sup>

<sup>1</sup> Brookhaven National Laboratory

<sup>2</sup> Brookhaven National Laboratory (US)

<sup>3</sup> Oak Ridge National Laboratory

<sup>4</sup> University of Pittsburgh (US)

<sup>5</sup> Carnegie Mellon University

<sup>6</sup> University of the Witwatersrand (ZA)

<sup>7</sup> University of Massachusetts (US)

**Corresponding Authors:** alexei.klimentov@cern.ch, ahoisie@bnl.gov, pnilsson@cern.ch, suterf@ornl.gov, tatiana.korchuganova@cern.ch, pnorbert@ornl.gov, ch.tasnuva@cern.ch, joseph.boudreau@cern.ch, yren@bnl.gov, klasky@ornl.gov, shengyuf@andrew.cmu.edu, okilic@bnl.gov, yiming@cs.cmu.edu, verena.martinez@cern.ch, tmaeno@bnl.gov

Large-scale scientific collaborations like ATLAS, Belle II, CMS, DUNE, and others involve hundreds of research institutes and thousands of researchers spread across the globe. These experiments generate petabytes of data, with volumes soon expected to reach exabytes. Consequently, there is a growing need for computation, including structured data processing from raw data to consumer-ready derived data, extensive Monte Carlo simulation campaigns, and a wide range of end-user analysis. To manage these computational and storage demands, centralized workflow and data management systems are implemented. However, decisions regarding data placement and payload allocation are often made disjointly and via heuristic means. A significant obstacle in adopting more effective heuristic or AI-driven solutions is the absence of a quick and reliable introspective dynamic model to evaluate and refine alternative approaches.

In this study, we aim to develop such an interactive system using real-world data. By examining job execution records from the PanDA workflow management system, we have pinpointed key performance indicators such as queuing time, error rate, and the extent of remote data access. The dataset includes six months of activity. Additionally, we are creating a generative AI model to simulate time series of payloads, which incorporate visible features like category, event count, and submitting group, as well as hidden features like the total computational load—derived from existing PanDA records and computing site capabilities. These hidden features, which are not visible to job allocators, whether heuristic or AI-driven, influence factors such as queuing times and data movement.

Parallel (Track 7) / 228

## Exploiting GPU Resources at VEGA for CMS Software Validation

**Authors:** Adriano Di Florio<sup>1</sup>; Andrea Bocci<sup>2</sup>; Antonio Perez-Calero Yzquierdo<sup>3</sup>; Christoph Wissing<sup>4</sup>; Daniele Spiga<sup>5</sup>; Jose Hernandez<sup>6</sup>

<sup>1</sup> CC-IN2P3

<sup>2</sup> CERN

<sup>3</sup> Centro de Investigaciones Energéticas Medioambientales y Tecnológicas

<sup>4</sup> Deutsches Elektronen-Synchrotron (DE)

<sup>5</sup> Università e INFN, Perugia (IT)

<sup>6</sup> CIEMAT



**Corresponding Authors:** jose.hernandez.calama@cern.ch, daniele.spiga@cern.ch, christoph.wissing@desy.de, antonio.perez.calero.yzquierdo@cern.ch, adriano.di.florio@cern.ch

In recent years, the CMS experiment has expanded the usage of HPC systems for data processing and simulation activities. These resources significantly extend the conventional pledged Grid compute capacity. Within the EuroHPC program, CMS applied for a “Benchmark Access” grant at VEGA in Slovenia, an HPC centre that is being used very successfully by the ATLAS experiment. For CMS, VEGA was integrated transparently as a sub-site extension to the Italian Tier-1 site at CNAF. In that first approach, only CPU resources were used, while all storage access was handled via CNAF through the network. Extending Grid sites with HPC resources was an established concept for CMS, however, in this project, HPC resources located in a different country from the Grid site were first integrated. CMS used the allocation primarily to validate a recent CMSSW release regarding its readiness for GPU usage. Former developments in the CMS workload management system that allow the targeting of GPU resources in the distributed infrastructure turned out to be instrumental and jobs could be submitted like any other release validation workflow. The presentation will detail aspects of the actual integration, some required tuning to achieve reasonable GPU utilisation, and an assessment of operational parameters like error rates compared to traditional Grid sites.

Parallel (Track 7) / 229

## First Deployment of XCache for Workflow and Efficiency Optimizations on Opportunistic HPC Resources in Germany

**Authors:** Achim Streit<sup>1</sup>; Andreas Petzold<sup>2</sup>; Artur Il Darovic Gottmann<sup>2</sup>; Gunter Quast<sup>2</sup>; Manuel Giffels<sup>2</sup>; Matthias Schnepf<sup>None</sup>; Robin Hofsaess<sup>2</sup>

<sup>1</sup> Karlsruhe Institute of Technology (KIT)

<sup>2</sup> KIT - Karlsruhe Institute of Technology (DE)

**Corresponding Authors:** manuel.giffels@cern.ch, artur.akhmetshin@cern.ch, robin.hofsaess@cern.ch, andreas.petzold@cern.ch, matthias.schnepf@kit.edu, gunter.quast@cern.ch, achim.streit@kit.edu

The efficient utilization of multi-purpose HPC resources for High Energy Physics applications is increasingly important, in particularly with regard to the upcoming changes in the German HEP computing infrastructure.

In preparation for the future, we are developing and testing an XRootD-based caching and buffering approach for workflow and efficiency optimizations to exploit the full potential of such resources despite the challenges and potential limitations associated with them.

With this contribution, we want to present a first prototype of our approach, deployed for optimizing the utilization of HoreKa, our local HPC cluster at KIT, that is opportunistically integrated into GridKa, the German Tier-1 center.

This includes first experiences and additional benefits for the operation of such sites that come with the additional monitoring capabilities of our setup.

Parallel (Track6) / 230

## The roles of the SCAB Nominations and Activities systems in the ATLAS-CERN speakers selection

**Authors:** Ana Clara Loureiro Cruz<sup>1</sup>; Carolina Niklaus Moreira Da Rocha Rodrigues<sup>1</sup>; Gabriela Lemos Lucidi Pinhao<sup>2</sup>; Leonardo Mira Marins<sup>1</sup>; Pedro Henrique Goes Afonso<sup>1</sup>; Rafaella Lenzi Romano<sup>1</sup>

<sup>1</sup> Federal University of Rio de Janeiro (BR)

<sup>2</sup> LIP - Laboratorio de Instrumentação e Física Experimental de Partículas (PT)

**Corresponding Authors:** gabriela.lemos.lucidi.pinhao@cern.ch, lmarins@cern.ch, ana.clara.loureiro.cruz@cern.ch, rafaella.lenzi.romano@cern.ch, carolina.niklaus.moreira.da.rocha.rodriques@cern.ch, pedro.afonso@cern.ch

Considering CERN's prosperous environment, developing groundbreaking research in physics and pushing technology's barriers, CERN members participate in many talks and conferences every year. However, given that the ATLAS experiment has around 6000 members and more than one could be qualified to present the same talk, the experiment developed metrics to prioritize them.

Currently, ATLAS is organized in a tree structure with 260 groups and subgroups, called activities. Each of these activities has responsible members such as the conveners or sub-conveners, project leaders, and activity coordinators. Because of the tree structure mentioned, the member's nomination will work its way up the branches, providing the upper levels with input from the lower ones. Previously, this process was not automated and happened through the exchange of CSVs, not providing these conveners and coordinators with the big picture of the nominations' priorities and reasons.

To improve this process, two systems were developed by the ATLAS Glance team: Activities and SCAB Nominations. The Activities interface provides a user-friendly view to manage the activities tree structure, the coordinators of each activity, and their allowed actions in the nomination process. The SCAB Nominations interface automates the nomination process of the ATLAS Speakers Committee Advisory Board, allowing all the coordinators to give their nominees priorities, and justify them in comments. These two systems contribute to a more holistic process for selecting collaboration members to present at a specific conference. This presentation delves into their specifications.

**Parallel (Track 2) / 231**

## **SYCL-based online data processing framework concept for PANDA**

**Author:** Bartosz Sobol<sup>1</sup>

**Co-author:** Grzegorz Korcyl<sup>1</sup>

<sup>1</sup> *Jagiellonian University*

**Corresponding Authors:** bartosz.sobol@doctoral.uj.edu.pl, grzegorz.korcyl@uj.edu.pl

The PANDA experiment has been designed to incorporate software triggers and online data processing. Although PANDA may not surpass the largest experiments in terms of raw data rates, designing and developing the processing pipeline and software platform for this purpose is still a challenge. Given the uncertain timeline for PANDA and the constantly evolving landscape of computing hardware, our attention is directed towards ensuring the future-proofness of the solutions we develop.

The PandaR2 is a concept for a framework handling online data processing in heterogeneous and distributed HPC environments. It utilizes the SYCL programming model as the primary technology for parallelization and offloading. Being a new and standalone entity, PandaR2 also interfaces with the PANDA's original ROOT-based simulation and analysis framework - PandaRoot, connecting the best of both worlds.

This contribution aims to present an overview of the PandaR2 SYCL-centric architecture. We will share experiences with SYCL during the codebase design process, particularly highlighting its portability across various hardware platforms and compilers. Additionally, we will showcase the performance results of the initial algorithms implemented in PandaR2, focusing on the performance portability of SYCL code and comparison with native programming models for accelerators, such as CUDA or HIP.

Poster session / 232

## Enhancing product management in the ATLAS Management Glance team: challenges and solutions

**Authors:** Ana Clara Loureiro Cruz<sup>1</sup>; Carolina Niklaus Moreira Da Rocha Rodrigues<sup>1</sup>; Gabriela Lemos Lucidi Pinhao<sup>2</sup>; Leonardo Mira Marins<sup>1</sup>; Pedro Henrique Goes Afonso<sup>1</sup>; Rafaella Lenzi Romano<sup>1</sup>

<sup>1</sup> Federal University of Rio de Janeiro (BR)

<sup>2</sup> LIP - Laboratorio de Instrumentação e Física Experimental de Partículas (PT)

**Corresponding Authors:** lmarins@cern.ch, pedro.afonso@cern.ch, carolina.niklaus.moreira.da.rocha.rodrigues@cern.ch, gabriela.lemos.lucidi.pinhao@cern.ch, rafaella.lenzi.romano@cern.ch, ana.clara.loureiro.cruz@cern.ch

CERN has a very dynamic environment and faces challenges such as information centralization, communication between the experiments' working groups, and the continuity of workflows. The solution found for those challenges is automation and, therefore, the Glance project, an essential management software tool for all four large LHC experiments. Its main purpose is to develop and maintain web-based automated solutions that are easy to learn and use and allow collaboration members to perform their tasks quickly.

The ATLAS Management Glance team is a subset of the Glance team focused on attending to the software requests of the ATLAS Spokesperson and deputies. The team maintains 11 systems that allow the management of ATLAS members, appointments, analyses, speaker nomination, and selection, among other tasks. Historically, each Glance developer would be an expert in the requirements of one or more systems, but their product management was inefficient, lacking the mapping of the product vision, goals, business rules, personas, and metrics. Also, the team's roadmap lacked predictability since it had no planned timeline.

In September 2023, the ATLAS Management Glance team adopted the Product Owner role concentrated in one single person recommended (or possibly "required") by the Scrum Guide. This presentation dives into the challenges faced by the Glance Team Product Owner in establishing a strategy for effective product management and roadmap planning and key takeaways from that process.

Parallel (Track 7) / 233

## The Glance project common infrastructure dependencies upgrade from the ATLAS Glance perspective

**Authors:** Ana Clara Loureiro Cruz<sup>1</sup>; Carolina Niklaus Moreira Da Rocha Rodrigues<sup>1</sup>; Gabriela Lemos Lucidi Pinhao<sup>2</sup>; Leonardo Mira Marins<sup>1</sup>; Pedro Henrique Goes Afonso<sup>1</sup>; Rafaella Lenzi Romano<sup>1</sup>

<sup>1</sup> Federal University of Rio de Janeiro (BR)

<sup>2</sup> LIP - Laboratorio de Instrumentação e Física Experimental de Partículas (PT)

**Corresponding Authors:** carolina.niklaus.moreira.da.rocha.rodrigues@cern.ch, pedro.afonso@cern.ch, gabriela.lemos.lucidi.pinhao@cern.ch, ana.clara.loureiro.cruz@cern.ch, rafaella.lenzi.romano@cern.ch, lmarins@cern.ch

The Glance project provides software solutions for managing high-energy physics collaborations' data and workflow. It was started in 2003 and operates in the ALICE, AMBER, ATLAS, CMS, and LHCb CERN experiments on top of CERN common infrastructure. The project develops Web applications using PHP and Vue.js, running on CENTOS virtual machines hosted on the CERN OpenStack private cloud. These virtual machines are built via Puppet for installing and configuring core software while tailoring them to meet each experiment's requirements in a collaborative approach under the Glance Project. This approach minimizes redundant work across experiments while allowing co-operation when responding to operations incidents. In the scenario of the CENTOS 7 end-of-life, the Glance project has chosen to migrate to RHEL9 while undergoing a major upgrade of PHP from (7.3 or 7.4 to 8.2) across the experiments. This presentation will expose the technical and organizational

challenges the Glance project faces on common dependencies upgrades from the perspective of the ATLAS Glance team.

Poster session / 234

## ATLAS Qualification interface refactoring strategy

**Authors:** Ana Clara Loureiro Cruz<sup>1</sup>; Carolina Niklaus Moreira Da Rocha Rodrigues<sup>1</sup>; Gabriela Lemos Lucidi Pinhao<sup>2</sup>; Leonardo Mira Marins<sup>1</sup>; Pedro Henrique Goes Afonso<sup>1</sup>; Rafaella Lenzi Romano<sup>1</sup>

<sup>1</sup> *Federal University of Rio de Janeiro (BR)*

<sup>2</sup> *LIP - Laboratorio de Instrumentação e Física Experimental de Partículas (PT)*

**Corresponding Authors:** gabriela.lemos.lucidi.pinhao@cern.ch, ana.clara.loureiro.cruz@cern.ch, rafaella.lenzi.romano@cern.ch, pedro.afonso@cern.ch, carolina.niklaus.moreira.da.rocha.rodrigues@cern.ch, lmarins@cern.ch

The ATLAS experiment involves over 6000 active members, including students, physicists, engineers, and researchers, and more than 2500 members are authors. This dynamic CERN environment brings up some challenges, such as managing the qualification status of each author. The Qualification system, developed by the Glance team, aims to automate the processes required for monitoring the progress of ATLAS members as they work to achieve author status. Recently, ATLAS modified the policy governing authorship qualification, and updates were necessary to put the changes into effect.

The system's code was originally developed on top of an outdated framework. In order to ease the transition to the new ATLAS authorship qualification policy, the code was updated to a Hexagonal architecture based on Domain Driven Design philosophy. The access to the database has shifted from ORM - Object Relational Mapper - to SQL repositories to align with the team's development stack. The system's quality is ensured with automatic tests as part of an effective refactoring process transparent for the end user. This refactoring strategy enhances our system to meet both previously unaddressed and new requirements, to improve code maintainability, and to increase flexibility to accommodate possible future changes in the qualification policy.

Parallel (Track 8) / 235

## HEP-Help: a first-stop helpline for particle physics software

**Author:** Jim Pivarski<sup>1</sup>

<sup>1</sup> *Princeton University*

**Corresponding Author:** pivarski@princeton.edu

If a physicist needs to ask for help on some software, where should they go? For a specific software package, there may be a preferred website, such as the ROOT Forum or a GitHub/GitLab Issues page, but how would they find this out? What about problems that cross package boundaries? What if they haven't found a tool that would solve their problem yet?

HEP-Help (hep-help.org) is intended as a first-stop helpline for questions about particle physics software. It is not intended to replace established venues, but redirect users to the best place to ask their questions, and possibly help them frame their questions in better ways, such as distinguishing usage questions from bug reports and constructing minimal reproducers.

This project has two parts: one technological and one social. The technical aspect involves collating existing documentation, tutorials, and forum archives to produce a dataset to train an LLM as a first responder. The social aspect involves building a community of part-time responders, people who

take shifts (help-a-thons!) to correct or follow up on the LLM's initial suggestions. This community includes tutorial trainers, developers of particle physics software, and experienced users, all of whom are already invested in helping new users and stand to benefit from a more organized support system.

Poster session / 236

## The evolution of the CMS@Home project

**Authors:** Alan Malta Rodrigues<sup>1</sup>; Christoph Wissing<sup>2</sup>; Daniele Spiga<sup>3</sup>; Federica Fanzago<sup>4</sup>; Ivan Reid<sup>5</sup>; Laurence Field<sup>6</sup>; Marco Mascheroni<sup>7</sup>

<sup>1</sup> *University of Notre Dame (US)*

<sup>2</sup> *Deutsches Elektronen-Synchrotron (DE)*

<sup>3</sup> *INFN Perugia*

<sup>4</sup> *INFN Padova*

<sup>5</sup> *Brunel University London (GB)*

<sup>6</sup> *CERN*

<sup>7</sup> *Univ. of California San Diego (US)*

**Corresponding Authors:** marco.mascheroni@cern.ch, alan.malta@cern.ch, federica.fanzago@pd.infn.it, ivan.reid@brunel.ac.uk, laurence.field@cern.ch, christoph.wissing@desy.de, daniele.spiga@pg.infn.it

Over time, the idea of exploiting voluntary computing resources as additional capacity for experiments at the LHC has given rise to individual initiatives such as the CMS@Home project. With a starting point of R&D prototypes and projects such as “jobs in the Vacuum” and SETI@Home, the experiments have tried integrating these resources into their data production frameworks transparently to the computing infrastructure. Many of these efforts were subsequently rolled into the umbrella LHC@Home project. The use of virtual machines instantiated on volunteer resources, with images created and managed by the experiment according to its needs, provided the opportunity to implement this integration, and virtualization enabled CMS code from a Linux environment to also run on Windows and Macintosh systems, realizing a distributed and heterogeneous computing environment. A prototype of CMS@Home integrated with the CMS workload management CRAB3 was proposed in 2015, demonstrating the possibility of using BOINC as “manager” of volunteer resources and adapting the “vacuum” concept with the HTCondor Glidein system to get CMS pilots and jobs to execute on volunteers’ computers. Since then, the integration of volunteer machines with the CMS workload management WMAgent, the official service dedicated to data production, has been seriously considered. The characteristics of volunteer resources regarding bandwidth capacity, connection behavior, and CPU and RAM capacities make them suitable for low-priority workflows with low I/O demands. The poster describes how the configuration of volunteer resources has evolved to keep pace with the development of the CMS computing infrastructure, including using tokens for resource authentication, exploiting regular expressions to accept workflows, manual glideins to initiate pilots, and other implementation details to achieve successful workflows. Currently volunteers are able to execute task chains also of multicore jobs and, despite their limitations, are contributing to CMS computing capacity with around 600 cores daily.

Parallel (Track 3) / 237

## Application of TRACCC seeding to the CEPC vertex detector

**Author:** Yizhou Zhang<sup>1</sup>

**Co-authors:** Jiaheng Zou<sup>2</sup>; Tao Lin; Teng LI<sup>3</sup>; Weidong Li<sup>2</sup>; Xiaocong Ai<sup>4</sup>; Xingtao Huang

<sup>1</sup> *Institute of High Energy Physics*

<sup>2</sup> *IHEP, Beijing*

<sup>3</sup> *Shandong University, CN*

<sup>4</sup> *Zhengzhou University*

**Corresponding Authors:** huangxt@sdu.edu.cn, liwd@ihep.ac.cn, zoujh@ihep.ac.cn, zhangyz@ihep.ac.cn, lintao@ihep.ac.cn, tengli@sdu.edu.cn, xiaocong.ai@cern.ch

The Circular Electron Positron Collider (CEPC) is a future experiment mainly designed to precisely measure the Higgs boson's properties and search for new physics beyond the Standard Model. In the design of the CEPC detector, the VerTeX detector (VTX) is the innermost tracker playing a dominant role in determining the vertexes of a collision event. The VTX detector is also responsible for providing seeds for the track following algorithms to find tracks in the outer trackers. TRACCC is one of the R&D lines aiming for developing the demonstrator for a full tracking chain for accelerators within the ACTS project.

This contribution will present the implementation of the seeding algorithm for the VTX detector based on TRACCC in the CEPC software (CEPCSW) environment. The integration of TRACCC into the CEPCSW, which is using Gaudi as the underlying framework and employing DD4hep as detector description tool and EDM4hep as the event data model, will be introduced. The CEPC VTX detector has three layers and both sides of each layer are mounted with silicon pixel sensors. To accommodate this specific detector structure, the default seeding algorithm in TRACCC, which creates three-space-points seeds, has been extended for six space points case. This contribution will also describe the solution of using one common memory for both EDM4hep and VecMem to avoid the overhead from data copy. For all the above work, both physics performance and computing performance are measured and will be presented.

**Poster session / 238**

## Tracking for long-lived particles at STCF

**Authors:** Xiaocong Ai<sup>1</sup>; Hao Li<sup>1</sup>; Hang Zhou<sup>None</sup>

<sup>1</sup> *Zhengzhou University*

**Corresponding Authors:** xiaocong.ai@cern.ch, lihao097430@gs.zzu.edu.cn, zh25@mail.ustc.edu.cn

With an electron-positron collider operating at center-of-mass-energy 2~7 GeV and a peak luminosity above  $0.5 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$ , the STCF physics program will provide an unique platform for in-depth studies of hadron structure and non-perturbative strong interaction, as well as probing physics beyond the Standard Model at the  $\tau$ -Charm sector succeeding the present Being Electron-Positron Collider II (BEPICII). To fulfill the physics targets and to further maximize the physics potential at the STCF, not only the particles that decay immediately upon production but also the long-lived particles, e.g. the lambda baryon, which may decay within or outside the inner tracker hence leaving very limited number of hits at the inner tracker, should be reconstructed with good efficiency.

A Common Tracking Software (ACTS) provides a set of performant track reconstruction tools which are agnostic to the details of the detection technologies and magnetic field configuration. Due to its excellent performance, ACTS has been used as a tracking toolkit by various experiments such as ATLAS, sPHENIX, FASER etc. Preliminary results of using ACTS seeding and Combinatorial Kalman Filter algorithms for STCF have been obtained. However, it's found that the tracking performance of ACTS seeding for long-lived particles at STCF is far from satisfactory, due to the fact that the STCF inner tracker has only three layers. Therefore, improving the tracking performance of ACTS for long-lived particles at STCF by combining the global track finding algorithm Hough Transform and the local track following algorithm CKF has been investigated.

In this talk, we will present the tracking performance of ACTS for STCF, which has a tracking system with a three-layer inner tracker and a drift chamber. Improvement of the tracking performance for long-lived particles at STCF using a combined global Hough Transform and the Combinatorial Kalman Filter will be highlighted.

Parallel (Track 5) / 239

## Opticks : GPU ray trace accelerated optical photon simulation

**Author:** simon blyth<sup>1</sup>

<sup>1</sup> *IHEP, CAS*

**Corresponding Author:** simon.c.blyth@gmail.com

Opticks is an open source project that accelerates optical photon simulation by integrating NVIDIA GPU ray tracing, accessed via the NVIDIA OptiX API, with Geant4 toolkit based simulations.

Optical photon simulation times of 14 seconds per 100 million photons have been measured within a fully analytic JUNO GPU geometry auto-translated from the Geant4 geometry when using a single NVIDIA GPU from the first RTX generation.

Optical physics processes of scattering, absorption, scintillator reemission and boundary processes are implemented in CUDA based on Geant4. Wavelength-dependent material and surface

properties as well as inverse cumulative distribution functions for reemission are interleaved into GPU textures providing fast interpolated property lookup or wavelength generation. In this work we describe the application of Opticks to JUNO simulation including new Opticks features that improve performance for complex CSG shapes and torus solids.

Poster session / 240

## ATLAS usage of the Czech national HPC center: HyperQueue, cvmfsexec, and other news

**Authors:** Jiri Chudoba<sup>1</sup>; Michal Svatos<sup>1</sup>; Petr Vokac<sup>2</sup>

<sup>1</sup> *Czech Academy of Sciences (CZ)*

<sup>2</sup> *Czech Technical University in Prague (CZ)*

**Corresponding Authors:** jiri.chudoba@cern.ch, michal.svatos@cern.ch, petr.vokac@cern.ch

The distributed computing of the ATLAS experiment at the Large Hadron Collider (LHC) utilizes computing resources provided by the Czech national High Performance Computing (HPC) center, IT4Innovations. This is done through ARC-CEs deployed at the Czech Tier2 site, pragueulcg2. Over the years, this system has undergone continuous evolution, marked by recent enhancements aimed at improving resource utilization efficiency.

One key enhancement involves the implementation of the HyperQueue meta-scheduler. It enables a division of whole-node jobs into several smaller, albeit longer, jobs, thereby enhancing CPU efficiency. Additionally, the integration of cvmfsexec enables access to the distributed CVMFS filesystem on compute nodes without requiring any special configurations, thereby substantially simplifying software distribution and broadening the range of tasks eligible for execution on the HPC. Another notable change was the migration of the batch system from PBSpro to Slurm.

Poster session / 241

## The new configuration system for ATLAS event processing applications

**Author:** Walter Lamp<sup>1</sup>

<sup>1</sup> *University of Arizona (US)*

**Corresponding Author:** walter.lamp@cern.ch

The offline software framework of the ATLAS experiment (Athena) consists of a large number of components of various types like Algorithm, Tool or Service. To assemble these components into an executable application for event processing, a dedicated configuration step is necessary. The set of components used for a particular job depends on the work-flow (simulation, reconstruction, high-level trigger, overlay, calibration, ...) and the input data (real or simulated data, beam-energy, ...). The configuration step is a program by itself, written in python. Its result (the configuration of the event-processing job) depends on optionally pre-set flags as well as meta-data about the data to be processed.

Between LHC Run 2 and Run 3, the configuration system was entirely redesigned to address problems encountered as the previously used system became more and more complex. Today, all production workflows (except event generation) have been migrated to the new scheme.

This contribution explains the design of the new configuration system (known as ComponentAccumulator) and reports about the migration experience.

**Parallel (Track6) / 242**

## **pip install ROOT: experiences making a complex multi-language package accessible for Python users**

**Authors:** Jonas Rembser<sup>1</sup>; Vincenzo Eduardo Padulano<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** jonas.rembser@cern.ch, vincenzo.eduardo.padulano@cern.ch

ROOT is a software toolkit at the core of LHC experiments and HENP collaborations worldwide, widely used by the community and in continuous development with it. The package is available through many channels that cater different types of users with different needs. This ranges from software releases on the LCG stacks provided via CVMFS for all HENP users to benefit, to pre-built binaries available on the three major platforms (Linux, MacOS, Windows), to more specialised packaging systems such as Homebrew, Snap, Anaconda. The last example is one of the main systems to distribute software to a Python user base, particularly beneficial for complex environments with real-world scientific applications in mind such as those found in HENP. Nonetheless, the standard Python implementation defaults to using pip as a package installer. This technology, together with the Python Package Index (PyPI), distributes many Python packages and has the advantage of providing a lightweight path to downstream development of a package with some upstream Python dependencies. This contribution highlights the steps required towards making pip install ROOT possible, demonstrating its availability as an early-stage release, and discussing some of the unique challenges of delivering a highly-performant multi-language software via the standard Python packaging system.

**Parallel (Track 5) / 243**

## **Zero-overhead training of machine learning models with ROOT data**

**Authors:** Vincenzo Eduardo Padulano<sup>1</sup>; Kristupas Pranckietis<sup>2</sup>; Lorenzo Moneta<sup>1</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *Vilnius University (LT)*



**Corresponding Authors:** lorenzo.moneta@cern.ch, vincenzo.eduardo.padulano@cern.ch, kristupas.pranckietis@cern.ch

The ROOT software framework is widely used in HENP for storage, processing, analysis and visualization of large datasets. With the large increase in usage of ML for experiment workflows, especially lately in the last steps of the analysis pipeline, the matter of exposing ROOT data ergonomically to ML models becomes ever more pressing. This contribution presents the advancements in an experimental component of ROOT that exposes datasets in batches ready for the training phase. This feature avoids the need for intermediate data conversion and can further streamline existing workflows, facilitating direct access of external ML tools to the ROOT input data in particular for the case when it does not fit in memory. The goal is to keep the footprint of using this feature minimal, in fact it represents just an extra line of code in user application. The contribution demonstrates such usage in various examples using different ML training models, also evaluating the performance with key metrics.

Parallel (Track 9) / 244

## Computing Activities at the Spanish Tier-1 and Tier-2s for the ATLAS experiment in the LHC Run3 period and towards the High-Luminosity Phase (HL-LHC)

**Authors:** Andreu Pacheco Pages<sup>None</sup>; Carlos Acosta Silva<sup>1</sup>; Elena Maria Planas Teruel<sup>2</sup>; Emma Torro Pastor<sup>3</sup>; Esther Accion Garcia<sup>2</sup>; Francisco Javier Sanchez Martinez<sup>3</sup>; Helena Burriel Navarro<sup>3</sup>; Jose Del Peso<sup>4</sup>; Jose Enrique Garcia Navarro<sup>3</sup>; Jose Flix Molina<sup>5</sup>; Jose Salt<sup>3</sup>; Josu Cantero<sup>6</sup>; Miguel Villaplana<sup>7</sup>; Pablo Collado Soto<sup>4</sup>; Roberto-Andreeas Uzum<sup>3</sup>; Santiago Gonzalez De La Hoz<sup>3</sup>; Vanessa Acin Portella<sup>8</sup>

<sup>1</sup> PIC

<sup>2</sup> The Barcelona Institute of Science and Technology (BIST) (ES)

<sup>3</sup> Univ. of Valencia and CSIC (ES)

<sup>4</sup> Universidad Autonoma de Madrid (ES)

<sup>5</sup> CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)

<sup>6</sup> IFIC/UV-CSIC (ES)

<sup>7</sup> IFIC - Univ. of Valencia and CSIC (ES)

<sup>8</sup> Institut de Física d'Altes Energies

**Corresponding Authors:** elena.maria.planas.teruel@cern.ch, jose.salt@cern.ch, carlos.acosta.silva@cern.ch, santiago.gonzalezdelahoz@cern.ch, pablo.collado.soto@cern.ch, emma.torro.pastor@cern.ch, javier.sanchez@cern.ch, npo@cern.ch, jose.delpeso@uam.es, robuzum@ific.uv.es, miguel.villaplana.perez@cern.ch, jcantero@cern.ch, jose.flix.molina@cern.ch, jose.enrique.garcia@cern.ch, esthera@pic.es, vacin@pic.es, helena.burriel@uv.es

This work is going to show the Spanish Tier-1 and Tier-2s contribution to the computing of the ATLAS experiment at the LHC during the Run3 period. The Tier-1 and Tier-2 GRID infrastructures, encompassing data storage, processing, and involvement in software development and computing tasks for the experiment, will undergo updates to enhance efficiency and visibility within the experiment.

The fundamental objective of this work is, on one hand, to provide data processing services in a stable manner for 24 hours a day every day of the year with a reliability greater than 95% in the conditions that the experiment needs and, on the other hand, to undertake the resolution of the problems posed by Run3 in which we are fully involved. The potential time interval of validity of this contribution covers a large part of Run3 and the beginning of the Long Shutdown 3 (LS3), namely the period 2022-2026. Central to our efforts is to engage actively with the various challenges inherent in research and development, in preparation for the upcoming, more intricate phase represented by the High-Luminosity LHC (HL-LHC).

We generate billions of simulated events annually for different physics processes. We capitalize on National High Performance Computers like the MareNostrum, part of the Supercomputing Spanish Network. We employ Data Lakes, a versatile paradigm for storing vast amounts of data crucial for the experiment's physics analyses. We belong to the core of GRID centers sufficiently reliable to house critical data and provide a first level of support to local ATLAS physicists.

A new activity in this work is the development and implementation of what we call the "Facility for

Interactive Distributed Analysis”. This initiative aims to facilitate data analysis work for physicists at Spanish centers (IFIC, UAM, and IFAE) by orchestrating the distributed nature of initial analysis phases with subsequent interactive phases involving reduced data files. The ultimate goal is to produce publishable physics results or contributions tailored for workshops and conferences.

The ATLAS Tier-1 and Tier-2 sites in Spain have contributed and will continue to contribute significantly to research and development in computing. These efforts include the evaluation of various models aimed at enhancing computing performance and data storage capacity to meet the demands of the LHC High Luminosity era.

**Parallel (Track 4) / 245**

## Improving the ATLAS Production System for Run 3: Achievements, Lessons, and the Future

**Author:** Misha Borodin<sup>1</sup>

<sup>1</sup> *University of Iowa (US)*

**Corresponding Author:** mikhail.borodin@cern.ch

The ATLAS production system is a key computer information system used for processing and analyzing data from the experiment. It orchestrates hundreds of different processing types across a distributed computing network to meet the experiment’s goals. This includes Monte Carlo simulations, data processing, analysis, and studies of detectors, all dealing with petabytes of data on hundreds of thousands of CPUs daily. The system connects experts in various fields with the distributed computing resources they need. Using modern software and development approaches, especially for web interfaces, is vital for managing the scale of the ATLAS experiment. Drawing on years of experience, the developers have also optimized workflows, saving significant amounts of time and disk space—critical resources for the future. In this paper, we discuss the system’s architecture, development strategies, metrics used, and the achievements of the last five years, along with the lessons learned and future plans to support the demands of the High Luminosity Large Hadron Collider.

**Parallel (Track 4) / 246**

## Enhancements and resource optimisations for ATLAS use of HammerCloud

**Authors:** Alexander Lory<sup>1</sup>; Benjamin Rottler<sup>2</sup>; Guenter Duceck<sup>1</sup>; Michael Boehler<sup>2</sup>

<sup>1</sup> *Ludwig Maximilians Universitat (DE)*

<sup>2</sup> *Albert Ludwigs Universitaet Freiburg (DE)*

**Corresponding Authors:** a.lory@cern.ch, benjamin.rotter@cern.ch, michael.boehler@cern.ch, guenter.duceck@cern.ch

HammerCloud (HC) is a framework for testing and benchmarking resources of the world wide LHC computing grid (WLCG). It tests the computing resources and the various components of distributed systems with workloads that can range from very simple functional tests to full-chain experiment workflows. This contribution concentrates on the ATLAS implementation, which makes extensive use of HC for monitoring global resources, and additionally, has implemented a mechanism to automatically exclude resources if certain critical tests fail. The auto-exclusion mechanism makes it possible to save resources by avoiding sending computationally intensive jobs to non-functioning clusters.

However, in some cases central errors of the distributed computing system lead to massive exclusions of otherwise well-functioning resources. A new feature improves the recovery after such mass-exclusion events. For the auto-exclusion mechanism to be effective and save resources, test jobs need to be sent at a sufficient frequency. This in turn also uses resources. In this contribution,

we give an estimate of the total balance of resources of the auto-exclusion system and explore possible optimisations.

Individual services and scripts have been reorganised as part of a general overhaul including containerisation and the web interface has been given a facelift after more than 10 years of operation. This contribution summarises the work needed to get HC ready for the next decade.

Poster session / 247

## Looking back on 23 years of the ATLAS Metadata Interface (AMI): Reflections on database architecture and ongoing developments in preparation for Run 4

**Authors:** Fabian Lambert<sup>1</sup>; Jerome Henri Fulachier<sup>2</sup>; Jerome Odier<sup>3</sup>; Pierre Antoine Delsart<sup>3</sup>

<sup>1</sup> LPSC Grenoble IN2P3/CNRS (FR)

<sup>2</sup> Centre National de la Recherche Scientifique (FR)

<sup>3</sup> LPSC/CNRS (Grenoble, FR)

**Corresponding Authors:** jerome.fulachier@lpsc.in2p3.fr, jerome.odier@cern.ch, fabian.lambert@lpsc.in2p3.fr, delbart@in2p3.fr

The ATLAS Metadata Interface (AMI) ecosystem has been developed within the context of ATLAS, one of the largest scientific collaborations. AMI is a mature, generic, metadata-oriented ecosystem that has been maintained for over 23 years. This paper briefly describes the main applications of the ecosystem within the experiment, including metadata aggregation for millions of datasets and billions of files, searching for datasets by metadata criteria, and metadata definition for data processing jobs (AMI-tags). The current architecture of the underlying databases will be outlined, in addition to the ongoing developments for preparations for Run 4. Optimizations based on advanced partitioning will also be described, enabling all datasets to be migrated into a new single catalog.

Parallel (Track 8) / 248

## Delivering large scale public engagement open weeks at STFC

**Authors:** Greg Corbett<sup>None</sup>; Ian Collier<sup>1</sup>; Phill Day<sup>2</sup>; Sophy Palmer<sup>None</sup>

<sup>1</sup> Science and Technology Facilities Council STFC (GB)

<sup>2</sup> UKRI/STFC

**Corresponding Authors:** sophy.palmer@stfc.ac.uk, ian.peter.collier@cern.ch, phill.day@stfc.ac.uk, greg.corbett@stfc.ac.uk

The Science and Technology Facilities Council (STFC), part of UK Research and Innovation (UKRI), has a rich tradition of fostering public engagement and outreach, as part of its strategic aim to showcase and celebrate STFC science, technology, and staff, both within its National Laboratories and throughout the broader community.

As part of its wider programme, STFC organised two large scale public engagement open weeks in 2023 and 2024. These events, held at the Sci-Tech Daresbury campus in the North of England and the Harwell Campus in the South of England, home to STFC's largest National Laboratories, collectively welcomed over 17,500 participants.

These open weeks provided an unparalleled opportunity for the public to intimately engage with groundbreaking science and technology. Attendees were immersed in hands-on activities, demonstrations, and enlightening talks spanning various disciplines and catering to all age groups. They

also had the unique opportunity to explore the state-of-the-art facilities on site, and talk with the people who work here.

STFC's Scientific Computing Department (SCD) took the lead in orchestrating and delivering computing outreach initiatives during both open weeks. This paper details STFC's approach to organizing these open weeks, how the open weeks were structured, SCD's planning process for contributing to the events, and delves into the specifics of SCD's impactful contributions. By sharing these insights, this paper aims to offer valuable lessons for the effective execution of large-scale public engagement initiatives within the scientific community.

**Parallel (Track 1) / 249**

## Data Placement Optimization for ATLAS in a Multi-Tiered Storage System within a Data Center

**Authors:** Imran Latif<sup>1</sup>; James Leonardi<sup>1</sup>; Qiulan Huang<sup>2</sup>; Shinjae Yoo<sup>None</sup>; Vincent Garonne<sup>2</sup>

<sup>1</sup> *Brookhaven National Laboratory*

<sup>2</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** [ilatif@bnl.gov](mailto:ilatif@bnl.gov), [qiulan.huang@cern.ch](mailto:qiulan.huang@cern.ch), [sjyoo@bnl.gov](mailto:sjyoo@bnl.gov), [jleonardi@bnl.gov](mailto:jleonardi@bnl.gov), [vincent.garonne@cern.ch](mailto:vincent.garonne@cern.ch)

Scientific experiments and computations, especially in High Energy Physics, are generating and accumulating data at an unprecedented rate. Effectively managing this vast volume of data while ensuring efficient data analysis poses a significant challenge for data centers, which must integrate various storage technologies. This paper proposes addressing this challenge by designing a multi-tiered storage model that employs diverse storage technologies tailored to different data needs, thereby addressing data classification, placement, and migration.

While users and administrators manually optimize storage by migrating data based on simple rules derived from human knowledge, decisions, and basic usage statistics, evaluating the placement of data in different storage classes with I/O-intensive workloads remains a complex task. To overcome this challenge and address existing limitations, we have developed a precise data popularity prediction model utilizing state-of-the-art AI/ML techniques. This model is crafted from the analysis of ATLAS data and access patterns. It enables us to migrate infrequently accessed data to more economical storage media, such as tape drives, while storing frequently accessed data on faster yet costlier storage media like HDD or SSD. This strategic approach ensures data is placed optimally into the appropriate storage classes, thereby maximizing storage capacity while minimizing data access latency for end-users. Additionally, we provide insights and explore potential implementations of an autonomous multi-tiered storage system on the storage infrastructure at BNL, leveraging dCache technology. Furthermore, we will discuss the outcomes and compare different implementation strategies.

**Parallel (Track 7) / 250**

## Using the ATLAS experiment software on heterogeneous resources

**Author:** Johannes Elmsheuser<sup>1</sup>

<sup>1</sup> *Brookhaven National Laboratory (US)*

**Corresponding Author:** [johannes.elmsheuser@cern.ch](mailto:johannes.elmsheuser@cern.ch)

With the large dataset expected from 2029 onwards by the HL-LHC at CERN, the ATLAS experiment is reaching the limits of the current data processing model in terms of traditional CPU resources based on x86\_64 architectures and an extensive program for software upgrades towards the HL-LHC has been set up. The ARM CPU architecture is becoming a competitive and energy efficient

alternative. Accelerators like GPUs are available in any recent HPC. In the past years ATLAS has successfully ported its full data processing and simulation software framework Athena to ARM and has invested significant effort in porting parts of the reconstruction and simulation algorithms to GPUs.

We report on the successful usage of the ATLAS experiment offline and online software framework Athena on ARM and GPUs through the PanDA workflow management system at various WLCG sites. Furthermore we report on performance optimizations of the builds for ARM CPUs and the GPU integration efforts. We will discuss performance comparisons of different ARM and x86 architectures on WLCG resources and Cloud compute providers like GCP and AWS using ATLAS productions workflows as used in the HepScore23 benchmark suite.

Parallel (Track 4) / 251

## Optimisation of ATLAS computing resource usage through a modern HEP Benchmark Suite via HammerCloud and PanDA

**Authors:** Alessandro Di Girolamo<sup>1</sup>; Alexander Lory<sup>2</sup>; Domenico Giordano<sup>1</sup>; Gonzalo Menendez Borge<sup>1</sup>; Ilija Vukotic<sup>3</sup>; Ivan Glushkov<sup>4</sup>; Natalia Diana Szczepanek<sup>1</sup>

<sup>1</sup> CERN

<sup>2</sup> Ludwig Maximilians Universitat (DE)

<sup>3</sup> University of Chicago (US)

<sup>4</sup> Brookhaven National Laboratory (US)

**Corresponding Authors:** domenico.giordano@cern.ch, ilija.vukotic@cern.ch, a.lory@cern.ch, natalia.diana.szczepanek@cern.ch, gonzalo.menendez.borge@cern.ch, alessandro.di.girolamo@cern.ch, ivan.glushkov@cern.ch

In April 2023 HEPscore23, the new benchmark based on HEP specific applications, was adopted by WLCG, replacing HEP-SPEC06. As part of the transition to the new benchmark, the CPU core power published by the sites needed to be compared with the effective power observed while running ATLAS workloads. One aim was to verify the conversion rate between the scores of the old and the new benchmark. The other objective was to understand how the HEPscore performs when run on multi-core job slots, so exactly like the computing sites are being used in the production environment. Our study leverages the HammerCloud infrastructure and the PanDA Workload Management System to collect a large benchmark statistic across 136 computing sites using an enhanced HEP Benchmark Suite. It allows us to collect not only performance metrics, but, thanks to plugins, it also collects information such as machine load, memory usage and other user-defined metrics during the execution and stores it in an OpenSearch database. These extensive tests allow for an in-depth analysis of the actual, versus declared computing capabilities of these sites. The results provide valuable insights into the real-world performance of computing resources pledged to ATLAS, identifying areas for improvement while spotlighting sites that underperform or exceed expectations. Moreover, this helps to ensure efficient operational practices across sites. The collected metrics allowed us to detect and fix configuration issues and therefore improve the experienced performance.

Poster session / 252

## CREST Demonstrator for HL-LHC in ATLAS

**Authors:** Andrea Formica<sup>1</sup>; Elizabeth Gallas<sup>2</sup>; Evgeny Alexandrov<sup>3</sup>; Frank Winklmeier<sup>4</sup>; Guillaume Unal<sup>5</sup>; Marcelo Vogel<sup>6</sup>; Mikhail Mineev<sup>3</sup>; Nurcan Ozturk<sup>6</sup>; Vakho Tsulaia<sup>7</sup>

<sup>1</sup> Université Paris-Saclay (FR)

<sup>2</sup> University of Oxford (GB)

<sup>3</sup> Joint Institute for Nuclear Research (RU)

<sup>4</sup> University of Oregon (US)

<sup>5</sup> CERN

<sup>6</sup> University of Texas at Arlington (US)

<sup>7</sup> Lawrence Berkeley National Lab. (US)

**Corresponding Authors:** vakhtang.tsulaia@cern.ch, nurcan.ozturk@cern.ch, marcelo.vogel@cern.ch, elizabeth.gallas@physics.ox.ac.uk, evgeny.alexandrov@cern.ch, mikhail.mineev@cern.ch, andrea.formica@cern.ch, guillaume.unal@cern.ch, frank.winklmeier@cern.ch

The ATLAS experiment is evolving the Conditions database infrastructure in preparation for LHC Run 4. The new multi-tier architecture called CREST (Conditions REST) is based on a web server that handles all database interactions via a REST API. To test the new infrastructure and its performance, we implemented a demonstrator aimed for the usage in a typical High-Level Trigger (HLT) workflow in the online environment.

This is a very complex workflow and it has some specific characteristics that are quite different from the later stages of the data processing: it requires an efficient caching system and a quick update of new conditions data like luminosity and beam spot simultaneously during raw data processing. In this presentation we describe the setup of the demonstrator and the results of the tests performed.

Poster session / 253

## Performance studies for ATLAS workloads in many-core Grid and HPC environments

**Authors:** Andrej Filipcic<sup>1</sup>; Gianfranco Sciacca<sup>2</sup>; Michele Weber<sup>2</sup>

<sup>1</sup> Jozef Stefan Institute (SI)

<sup>2</sup> Universitaet Bern (CH)

**Corresponding Authors:** gsciacca@mail.cern.ch, andrej.filipcic@ijs.si, michele.weber@cern.ch

Developments in microprocessor technology have confirmed the trend towards higher core counts and decreased amount of memory per core, resulting in major improvements in power efficiency for a given level of performance. Core counts have increased significantly over the past five years for the x86\_64 architecture, which is dominating in the LHC computing environment, and the higher core density is not only a feature of large HPC systems, but is also readily available on commodity hardware preferentially used at Grid sites. The baseline multi-core workloads are however still largely based on 8-cores. The job are sized accordingly in terms of number of events processed. The new multi-threaded AthenaMT framework has been introduced for ATLAS data processing and simulation for Run-3 in order to address the performance limitations of the classic single-threaded Athena when run in parallel in multi-core jobs. In this work, the performance of some ATLAS workloads is investigated when scaling up core counts up to whole node where possible and at different job sizes with the aim of providing input to software developers.

Parallel (Track 4) / 254

## Integrating the Perlmutter HPC system in the ALICE Grid

**Author:** Sergiu Weisz<sup>1</sup>

**Co-authors:** Costin Grigoras<sup>2</sup>; Irakli Chakaberia<sup>3</sup>; Latchezar Betev<sup>2</sup>

<sup>1</sup> Lawrence Berkeley National Lab

<sup>2</sup> CERN

<sup>3</sup> Lawrence Berkeley National Lab. (US)

**Corresponding Authors:** sergiu.weisz@lbl.gov, latchezar.betev@cern.ch, costin.grigoras@cern.ch, irakli.chakaberia@cern.ch

The Perlmutter HPC system is the 9th generation supercomputer deployed at the National Energy Research Scientific Computing Center (NERSC). It provides both CPU and GPU resources, offering 393216 AMD EPYC Milan cores with 4 GB of memory per core, for CPU-oriented jobs and 7168 NVIDIA A100 GPUs. The machine allows connections from the worker nodes to the outside and already mounts CVMFS for users who need to access software from it. These two options make Perlmutter an ideal candidate for integrating into Grid infrastructures.

Due to the specific highly parallel and massive CPU and memory requirements of the native payloads running on supercomputers, there is always an idle part of the computing capacity. Conversely Grid payloads require few CPU cores for a single task and can take advantage of the idle resources. This 'backfill' is advantageous both for the supercomputer operators, increasing the overall use efficiency of the machine and for the Grid users, allowing them to opportunistically use a substantial amount of CPUs. ALICE takes advantage of these conditions, the architecture of the Perlmutter supercomputer, and facilities offered by NERSC by deploying a standard Grid interface to Perlmutter through the NERSC SuperFacility API scheduling tool to submit and monitor normal Grid payloads. Perlmutter has been integrated into the ALICE Grid, running Monte Carlo simulation, with measurements and tests having been made to also integrate analysis jobs connecting to an EOS instance hosted at LBNL shared with the main Tier 2 site. The resulting HPC-based Grid site has proven to be a reliable resource contributor to the ALICE Grid, providing 8000 cores on average, with its only constraints being the short lifetime of jobs and the current time allocation from NERSC.

This paper describes the path taken to integrate Perlmutter in the ALICE Grid and the usual modifications needed to integrate HPC facilities into the standard Grid infrastructure.

Poster session / 255

## Modernizing ATLAS PanDA for a sustainable multi-experiment future

**Authors:** Alexei Klimentov<sup>1</sup>; Edward Karavakis<sup>1</sup>; Fa-Hui Lin<sup>2</sup>; Fernando Harald Barreiro Megino<sup>3</sup>; Kaushik De<sup>2</sup>; Paul Nilsson<sup>1</sup>; Tadashi Maeno<sup>1</sup>; Tatiana Korchuganova<sup>4</sup>; Torre Wenaus<sup>1</sup>; Wen Guan<sup>1</sup>; Xin Zhao<sup>1</sup>; Zhaoyu Yang<sup>1</sup>

<sup>1</sup> Brookhaven National Laboratory (US)

<sup>2</sup> University of Texas at Arlington (US)

<sup>3</sup> University of Texas at Arlington

<sup>4</sup> University of Pittsburgh (US)

**Corresponding Authors:** xzhao@bnl.gov, kaushik@uta.edu, edward.karavakis@cern.ch, tmaeno@bnl.gov, fahui.lin@cern.ch, wenaus@gmail.com, alexei.klimentov@cern.ch, wen.guan@cern.ch, zhaoyu.yang@cern.ch, tatiana.korchuganova@cern.ch, fernando.harald.barreiro.megino@cern.ch, pnilsson@cern.ch

In early 2024, ATLAS undertook an architectural review to evaluate the functionalities of its current components within the workflow and workload management ecosystem. Pivotal to the review was the assessment of the Production and Distributed Analysis (PanDA) system, which plays a vital role in the overall infrastructure.

The review findings indicated that while the current system shows no apparent signs of scalability limitations or critical defects, several issues still require attention. These include areas for improvement, such as cleaning the historical accumulation of code over nearly two decades of continuous operation in ATLAS, further organizing development activities, maximizing the utilization of continuous integration and testing frameworks, bolstering efforts toward cross-experimental outreach, spreading greater awareness of workflows at the core level, expanding support for complex workflows, implementing a more advanced algorithm for workload distribution, optimizing tape and network resource usage, refining interface design, enhancing transparency to showcase system dynamism, ensuring allocation of key developers to R&D projects with clear long-term visions for integration and operation, and accommodating the growing diversity of resources.

In this presentation, we will first highlight the issues identified in the review, exploring their historical and cultural roots. We will then outline the recommendations derived from the review, and present the solutions developed to address these challenges and pave the way to sustainably support multiple experiments.

**Parallel (Track 1) / 256**

## Advancements in the in-file metadata system for the ATLAS experiment

**Authors:** Alaettin Serhan Mete<sup>1</sup>; Attila Krasznahorkay<sup>2</sup>; Maciej Pawel Szymanski<sup>1</sup>; Marcin Nowak<sup>3</sup>; Peter Van Gemmeren<sup>1</sup>

<sup>1</sup> *Argonne National Laboratory (US)*

<sup>2</sup> *CERN*

<sup>3</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** alaettin.serhan.mete@cern.ch, attila.krasznahorkay@cern.ch, peter.van.gemmeren@cern.ch, maciej.szymanski@cern.ch, marcin.nowak@cern.ch

The High-Luminosity upgrade of the Large Hadron Collider (HL-LHC) will increase luminosity and the number of events by an order of magnitude, demanding more concurrent processing. Event processing is trivially parallel, but metadata handling is more complex and breaks that parallelism. However, correct and reliable in-file metadata is crucial for all workflows of the experiment, enabling tasks such as job configuration, decoding trigger information, and keeping track of event selection. Therefore, ATLAS is enhancing its current in-file metadata system to support metadata creation and propagation in more robust ways. This talk presents developments in the areas of the evolution of storage technology for metadata and the redesign of metadata tools.

Firstly, we delve into the investigation of storage technologies tailored for in-file metadata payload, exploring advancements in the ROOT framework, used for storing data collected by the ATLAS experiment. Not only will this work allow ATLAS to utilize modern storage containers (such as RNTuple) for event and metadata, but the goal is to enhance performance and enable seamless handling of metadata. We also discuss whether a challenging process of merging metadata objects could be performed generically, e.g. using RNTuple features.

Furthermore, we introduce a novel approach to metadata tools by developing dual-mode functionality. Such tools, offering both creation and propagation capabilities, improve maintainability and facilitate handling in workflows making use of shared I/O functionality. This work also enables investigation of metadata propagation outside of event processing.

**Parallel (Track 1) / 257**

## Advancing ATLAS DCS Data Analysis with a Modern Data Platform

**Authors:** Andrea Formica<sup>1</sup>; Luca Canali<sup>2</sup>; Michelle Ann Solis<sup>3</sup>

<sup>1</sup> *Université Paris-Saclay (FR)*

<sup>2</sup> *CERN*

<sup>3</sup> *University of Arizona (US)*

**Corresponding Authors:** luca.canali@cern.ch, andrea.formica@cern.ch, michelle.ann.solis@cern.ch

This paper presents a novel approach to enhance the analysis of ATLAS Detector Control System (DCS) data at CERN. Traditional storage in Oracle databases, optimized for WinCC archiver operations, is challenged by the need for extensive analysis across long timeframes and multiple devices, alongside correlating conditions data. We introduce techniques to improve troubleshooting and analysis of ATLAS New Small Wheel (NSW) DAQ links, including data migration to Apache Parquet for efficient storage, and leveraging Big Data technologies like Apache Spark and Apache Hadoop for analysis. Employing Jupyter notebooks on the SWAN service, combined with Spark, Pandas, and the extensive Python ecosystem in general, facilitated a highly efficient analysis workflow. This approach was well-received by NSW experts, allowing them to rapidly gain proficiency and execute advanced analyses within a notably brief period.



Poster session / 258

## Microservices framework and configuration database for ATLAS ITk

**Authors:** Gerhard Immanuel Brandt<sup>1</sup>; Jonas Schmeing<sup>1</sup>; Maren Stratmann<sup>1</sup>; Marvin Emin Geyik<sup>1</sup>; Wolfgang Wagner<sup>1</sup>

<sup>1</sup> *Bergische Universitaet Wuppertal (DE)*

**Corresponding Authors:** jonas.schmeing@cern.ch, maren.stratmann@cern.ch, wolfgang.wagner@cern.ch, marvin.emin.geyik@cern.ch, gerhard.immanuel.brandt@cern.ch

To operate ATLAS ITk system tests and later the final detector, a graphical operation and configuration system is needed. For this a flexible and scalable framework based on distributed microservices has been introduced. Different microservices are responsible for configuration or operation of all parts of the readout chain.

The configuration database microservice provides the configuration files needed to configure the hardware components of the readout chain and perform scans using the DAQ software. It saves the connectivity information and configuration files for the operation of the system in so called runkeys. These runkeys are stored in a flexible, tree-based data structure. This flexible structure allows the storage of specialized runkeys made up of different objects for each of the ITk subdetectors within the same database.

It is investigated whether a single-instance database is sufficient to efficiently serve these files to the subdetectors or if a distributed system of local ConfigDB caches is needed. These caches would each provide only a subset of the runkeys depending on the elements of the readout chain that the specific cache serves.

Poster session / 259

## Updating the software description of the ATLAS Detector

**Authors:** Evgueni Tcherniaev<sup>1</sup>; Johannes Junggeburth<sup>2</sup>; Joseph Boudreau<sup>1</sup>; Marilena Bandieramonte<sup>1</sup>; Riccardo Maria Bianchi<sup>1</sup>; Sarka Todorova<sup>3</sup>; Vakho Tsulaia<sup>4</sup>

<sup>1</sup> *University of Pittsburgh (US)*

<sup>2</sup> *University of Massachusetts (US)*

<sup>3</sup> *Charles University (CZ)*

<sup>4</sup> *Lawrence Berkeley National Lab. (US)*

**Corresponding Authors:** evgueni.tcherniaev@cern.ch, marilena.bandieramonte@cern.ch, joseph.boudreau@cern.ch, vakhtang.tsulaia@cern.ch, riccardo.maria.bianchi@cern.ch, johannes.josef.junggeburth@cern.ch, sarka.todorova@cern.ch

The software description of the ATLAS detector is based on the GeoModel toolkit, developed in-house for the ATLAS experiment but released and maintained as a separate package with few dependencies. A compact SQLite-based exchange format permits the sharing of geometrical information between applications including visualization, clash detection, material inventory, database browsing, and lightweight full simulation. ATLAS simulation, reconstruction, and other elements of standard ATLAS offline workflows are now being adapted to ingest the geometry files which are prepared using platform independent modular geometry plugin code. This represents a major transformation of the ATLAS detector description software, impacting even the development procedures, for which new roles have been invented. During these integration activities, both the GeoModel geometry kernel and the GeoModel toolkit have seen improvements, including volume calculation, material blending, helper classes for simpler memory management, and a richer collection of supported geometrical objects. This talk reports on these activities.

Poster session / 260

## A simplified configuration for common algorithms for ATLAS analysis

**Authors:** Heather Russell<sup>1</sup>; Joseph Earl Lambert<sup>2</sup>; Matthew Kenneth Maroun<sup>3</sup>; Nils Erik Krumnack<sup>4</sup>; Verena Ingrid Martinez Outschoorn<sup>3</sup>; William Axel Leight<sup>5</sup>

<sup>1</sup> *University of Victoria*

<sup>2</sup> *University of Victoria (CA)*

<sup>3</sup> *University of Massachusetts (US)*

<sup>4</sup> *Iowa State University (US)*

<sup>5</sup> *University of Massachusetts Amherst*

**Corresponding Authors:** heather.russell@cern.ch, mmaroun@umass.edu, william.axel.leight@cern.ch, verena.martinez@cern.ch, nils.erik.krumnack@cern.ch, joseph.earl.lambert@cern.ch

In the ATLAS analysis model, users must interact with specialized algorithms to perform a variety of tasks on their physics objects including calibration, identification, and obtaining systematic uncertainties for simulated events. These algorithms have a wide variety of configurations, and often must be applied in specific orders. A user-friendly configuration mechanism has been developed with the goal of improving the user experience from the perspective of both ease-of-use and stability. Users can now configure necessary algorithms via a YAML file, enabled by a physics-oriented python configuration. The configuration mechanism and training will be discussed.

Parallel (Track 4) / 261

## Total Cost of Ownership and Evaluation of Google Cloud Resources for the ATLAS Experiment at the LHC

**Authors:** David South<sup>1</sup>; Gonzalo Merino Arevalo<sup>2</sup>

<sup>1</sup> *Deutsches Elektronen-Synchrotron (DE)*

<sup>2</sup> *The Barcelona Institute of Science and Technology (BIST) (ES)*

**Corresponding Authors:** gonzalo.merino@cern.ch, david.south@cern.ch

The ATLAS Google Project was established as part of an ongoing evaluation of the use of commercial clouds by the ATLAS Collaboration, in anticipation of the potential future adoption of such resources by WLCG grid sites to fulfil or complement their computing pledges. Seamless integration of Google cloud resources into the worldwide ATLAS distributed computing infrastructure was achieved at large scale and for an extended period of time, and hence cloud resources are shown to be an effective mechanism to provide additional, flexible computing capacity to ATLAS. For the first time a Total Cost of Ownership analysis has been performed, to identify the dominant cost drivers and explore effective mechanisms for cost control. Network usage significantly impacts the costs of certain ATLAS workflows, underscoring the importance of implementing such mechanisms. Resource bursting has been successfully demonstrated, whilst exposing the true cost of this type of activity. A follow-up to the project is underway to investigate methods for improving the integration of cloud resources in data-intensive distributed computing environments and reducing costs related to network connectivity, which represents the primary expense when extensively utilising cloud resources.

Poster session / 262

## Enhancing Network Analytics through Machine Learning

**Author:** Petya Vasileva<sup>1</sup>

**Co-authors:** Ilija Vukotic<sup>2</sup>; Marian Babik<sup>3</sup>; Shawn McKee

<sup>1</sup> *University of Michigan (US)*

<sup>2</sup> *University of Chicago (US)*

<sup>3</sup> *CERN*

**Corresponding Authors:** shawn.mckee@cern.ch, petya@cern.ch, marian.babik@cern.ch, ilija.vukotic@cern.ch

The research and education community relies on a robust network in order to access the vast amounts of data generated by their scientific experiments. The underlying infrastructure connects a few hundreds of sites across the world, which require reliable and efficient transfers of increasingly large datasets. These activities demand proactive methods in network management, where potentially severe issues are predicted and circumvented before they can impact the data exchanges. Our ongoing research is focused on leveraging both machine learning (ML) and deep learning (DL) methodologies to find patterns that cause network anomalies, predict key performance metrics, and explore the interconnectivity of paths across the networks.

We explore a diverse set of ML/DL models including a range of strategies suited for time series analysis, anomaly detection, and predictive modeling, where we are continually adjusting and refining our techniques. The goal is to detect subtle indicators of network instability or degradation that could disrupt the scientific workflows. Furthermore, we seek to localize problematic clusters, specific routers, or router-to-router links. This capability could serve not only to inform site administrators of present network health, but to guide upgrades and resource allocation for future network planning.

In this presentation we will share our experiments intended to delve into suitable ML/DL techniques, including ensemble learning and unsupervised models that may capture the complexities inherent in network data. In addition, we will discuss some of the many challenges we encounter, the selected model architectures and achieved results.

**Poster session / 263**

## Infrastructure for deployment and evaluation of LHCb trigger configurations

**Authors:** Luke Grazette<sup>1</sup>; Micol Olocco<sup>2</sup>; Rosen Matev<sup>3</sup>

<sup>1</sup> *University of Warwick (GB)*

<sup>2</sup> *Technische Universitaet Dortmund (DE)*

<sup>3</sup> *CERN*

**Corresponding Authors:** rosen.matev@cern.ch, luke.grazette@cern.ch, micol.olocco@cern.ch

CHEP Track: 6 - Collaborative software and maintainability

The LHCb high-level trigger applications consists of components that run reconstruction algorithms and perform physics object selections, scaling from hundreds to tens of thousands depending on the selection stage. The configuration of the components, the data flow and the control flow are implemented in Python. The resulting application configuration is condensed in the basic form of a list of components with their properties and values.

It is often required to change configuration without deploying new binaries. Moreover, it is essential to be able to reproduce a given production configuration and to be able query it after it has been used. For these reasons, the basic form of the trigger configuration is captured and stored in a Git database.

This contribution is describing the infrastructure around generating and validating the configurations. The process is based on GitLab pipelines that are triggered on user defined specifications

and run several steps ranging from basic checks to performance validation using dedicated runners. Upon merging, the configuration database is deployed on CVMFS. The process as implemented ensures consistency and reproducibility across all selection stages.

This project also aims to take advantage of the query-able nature of the configurations by creating an API that allows probing a single configuration in detail. This is further used to create human-readable summaries and to track changes across configurations to help analysts understand the selections used to collect their datasets.

Poster session / 264

## Investigating Data Access Models for ATLAS: A Case Study with FABRIC Across Borders and ServiceX

**Authors:** David Jordan<sup>1</sup>; Fengping Hu<sup>1</sup>; Ilija Vukotic<sup>1</sup>; Judith Lorraine Stephen<sup>1</sup>; Lincoln Bryant<sup>1</sup>; Robert William Gardner Jr<sup>1</sup>; Shawn Mc Kee<sup>2</sup>

<sup>1</sup> *University of Chicago (US)*

<sup>2</sup> *University of Michigan (US)*

**Corresponding Authors:** [ilija.vukotic@cern.ch](mailto:ilija.vukotic@cern.ch), [lincoln.bryant@cern.ch](mailto:lincoln.bryant@cern.ch), [judith.lorraine.stephen@cern.ch](mailto:judith.lorraine.stephen@cern.ch), [djordan66@uchicago.edu](mailto:djordan66@uchicago.edu), [robert.w.gardner@cern.ch](mailto:robert.w.gardner@cern.ch), [shawn.mckee@cern.ch](mailto:shawn.mckee@cern.ch), [fengping@uchicago.edu](mailto:fengping@uchicago.edu)

This study explores possible enhancements in analysis speed, WAN bandwidth efficiency, and data storage management through an innovative data access strategy. The proposed model introduces specialized “delivery” services for data preprocessing, which include filtering and reformatting tasks executed on dedicated hardware located alongside the data repositories at the CERN Tier-0 or at Tier-1 or Tier-2 facilities. Positioned near the source storage, these services are crucial for limiting redundant data transfers and focus on sending only vital data to distant analysis sites, aiming to optimize network and storage use at those sites. Within the scope of the NSF-funded FABRIC Across Borders (FAB) initiative, we assess this model using an “in-network, edge” computing cluster at CERN, outfitted with substantial processing capabilities (CPU, GPU, and advanced network interfaces). This edge computing cluster features dedicated network peering arrangements that link CERN Tier-0, the FABRIC experimental network, and an analysis center at the University of Chicago, creating a solid foundation for our research.

Central to our infrastructure is ServiceX, an R&D software project under the Data Organization, Management, and Access (DOMA) group of the Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP). ServiceX is a scalable filtering and reformatting service, designed to operate within a Kubernetes environment and deliver output to an S3 object store at an analysis facility. Our study assesses the impact of server-side delivery services in augmenting the existing HEP computing model, particularly evaluating their possible integration within the broader WAN infrastructure. This model could empower Tier-1 and Tier-2 centers to become efficient data distribution nodes, enabling a more cost-effective way to disseminate data to analysis sites and object stores, thereby improving data access and efficiency. This research is experimental and serves as a demonstrator of the capabilities and improvements that such integrated computing models could offer in the HL-LHC era.

Parallel (Track 3) / 265

## Evolution of the ATLAS event data model for the HL-LHC

**Authors:** Attila Krasznahorkay<sup>1</sup>; Carlo Varni<sup>2</sup>; Charles Leggett<sup>3</sup>; Frank Winklmeier<sup>4</sup>; Goetz Gaycken<sup>4</sup>; Nuno Fernandes<sup>None</sup>; Scott Snyder<sup>5</sup>; Tomasz Bold<sup>6</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *University of California Berkeley (US)*

<sup>3</sup> *Lawrence Berkeley National Lab (US)*

<sup>4</sup> *University of Oregon (US)*

<sup>5</sup> *Brookhaven National Laboratory (US)*

<sup>6</sup> *AGH University of Krakow (PL)*

**Corresponding Authors:** attila.krasznahorkay@cern.ch, frank.winklmeier@cern.ch, goetz.gaycken@cern.ch, tomasz.bold@cern.ch, nsfer21@gmail.com, scott.snyder@cern.ch, charles.g.leggett@gmail.com, carlo.varni@cern.ch

Run 4 of the LHC will yield an unprecedented volume of data. In order to process this data, the ATLAS collaboration is evolving its offline software to be able to use heterogenous resources such as GPUs and FPGAs. To reduce conversion overheads, the event data model (EDM) should be compatible with the requirements of these resources. While the ATLAS EDM has long allowed representing data as a structure of arrays, further evolution of the EDM can enable more efficient sharing of data between CPU and GPU resources. Some of this work will be summarized here, including extensions to allow controlling how memory for event data is allocated and implementation of jagged vectors.

Parallel (Track 5) / 266

## NeuroMCT: Fast Monte Carlo Tuning with Generative Machine Learning in the JUNO Experiment

**Authors:** Arsenii Gavrikov<sup>None</sup>; Andrea Serafini<sup>1</sup>

<sup>1</sup> *University of Padova & INFN*

**Corresponding Authors:** andrea. Serafini@infn.it, arsenii.gavrikov@pd.infn.it

The Jiangmen Underground Neutrino Observatory (JUNO) is a neutrino experiment under construction in the Guangdong province of China. The experiment has a wide physics program with the most ambitious goal being the determination of the neutrino mass ordering and the high-precision measurement of neutrino oscillation properties using anti-neutrinos produced in the 50 km distant commercial nuclear reactors of Taishan and Yangjiang.

To reach its aims, the detector features an acrylic sphere of 35.4 meters in diameter filled with 20 kt of liquid scintillator and equipped with 17612 20-inch photomultiplier tubes (PMTs) and 25 600 3-inch PMTs to provide an energy resolution better than 3% at 1 MeV. In addition to the cutting-edge features and performance of the detector, a critical aspect for achieving the physics goals is a deep understanding of such a complicated detector. In this respect, an accurate Monte Carlo (MC) simulation of the detector and the interactions happening inside of it is crucial. The simulation depends on many effective parameters, which must be tuned to accurately describe the data acquired.

In this contribution, we propose a novel machine-learning approach to MC tuning that combines Generative Learning and data acquired during calibration campaigns. We study Generative Adversarial Networks (GAN) as a way to speed up event simulation and as an efficient model to interpolate within the parameter space. We consider three main parameters related to the energy response of the JUNO detector and optimize their value in the MC by comparing calibration data to the GAN simulations. Parameter estimation is performed via Bayesian optimization based on a Nested Sampling algorithm to cope with the wide and complex parameter space.

The presented approach is easily scalable to include more parameters and is general enough to be employed in most modern physics experiments.

Parallel (Track 7) / 267

## Evolving MonALISA: Modernizing the Communication Layer of a Subagent-Based Distributed Monitoring System

**Authors:** Costin Grigoras<sup>1</sup>; Cristian Margineanu<sup>2</sup>; Latchezar Betev<sup>1</sup>; Maria Mateea Popescu<sup>3</sup>; Mihai Carabas<sup>3</sup>

<sup>1</sup> CERN

<sup>2</sup> Universitatea Nationala de Stiinta si Tehnologie Politehnica Bucuresti (RO)

<sup>3</sup> University Politehnica of Bucharest (RO)

**Corresponding Authors:** costin.grigoras@cern.ch, maria.mateea.popescu@cern.ch, mihai.carabas@cern.ch, cristian.margineanu@cern.ch, latchezar.betev@cern.ch

MonALISA serves as the main distributed computing monitoring framework for the ALICE experiment at CERN. Initially developed in collaboration with Caltech, this framework has historically enabled monitoring of services and nodes across computing clusters distributed worldwide. Within this infrastructure, AliMonitor, the ALICE experiment instance of MonALISA, continuously monitors all grid-related jobs, services, nodes, and clusters. AliMonitor operates more than 150 online MonALISA nodes, collectively gathering over 250 thousand values per second and exposing more than 17 million parameters in real time, enabling the prompt and historical monitoring of 200000 cores across the entire Grid and storage servers for 450PB of data.

MonALISA was designed as a cluster of autonomous agent subsystems that registered with a central authority. These agents monitor local infrastructure, advertise their activity to other subagents, and provide monitoring data to clients. This paper outlines a proposal for a complete revamp of the communication layer among these subagents, shifting to a system that utilizes asynchronous commands via WebSocket-based channels. The implementation introduces significant enhancements including authentication and authorization protocols between agents, modernized access methods to the cluster such as web clients, modularization of cluster roles for independent activation, and the establishment of a robust, futureproof communication channel.

Poster session / 268

## Decode the Workload: Training Deep Learning Models for Efficient Compute Cluster Representation

**Author:** Ahmed Mohammed<sup>1</sup>

**Co-authors:** Bryan Hess<sup>1</sup>; Diana McSpadden<sup>1</sup>; Laura Hild<sup>1</sup>; Malachi Schram<sup>1</sup>; Mark Jones<sup>1</sup>; Wesley Moore<sup>1</sup>; Zhenyu Dai<sup>1</sup>

<sup>1</sup> Thomas Jefferson National Accelerator Facility

**Corresponding Authors:** dai@jlab.org, dianam@jlab.org, ahmedm@jlab.org, lsh@jlab.org, bhess@jlab.org, maj@jlab.org, schram@jlab.org, wmoore@jlab.org

Monitoring the status of a high throughput computing (HTC) cluster running computationally-intensive nuclear physics (NP) jobs is a crucial yet challenging system administration task due to the complexity of such systems. To this end, we train (variational) autoencoders (VAE / AE) using the Linux kernel CPU metrics of the cluster. Additionally, we explore assisting these models with graph neural networks (GNNs) to share information across the CPUs within a compute node. The models are compared in terms of their: 1) Ability to produce a compressed latent representation that captures the salient features of the input and, 2) Ability to make distinction between different kinds of NP jobs run at Jefferson Lab. The goal is to have a robust encoder whose compressed embeddings are used by several downstream tasks such as anomaly detection or control in a fully automated setup. We extend this study further by deploying such models in a human-in-the-loop production-based setting for the anomaly detection task and discuss the associated implementation aspects such as continual learning and the criterion to generate alarms. This study represents a first step in the endeavor towards building self-supervised large-scale foundation models for computing centers.

**Parallel (Track6) / 269**

## CERN OSPO's Catalogue Initiative: Amplifying Open Science for the CERN Community

**Authors:** Giacomo Tenaglia<sup>1</sup>; Victoria Stephany Huisman Sigcha<sup>2</sup><sup>1</sup> CERN<sup>2</sup> Saxion University of Applied Scienc (NL)**Corresponding Authors:** giacomo.tenaglia@cern.ch, v.h.s@cern.ch

At the core of CERN's mission lies a profound dedication to open science; a principle that has fueled decades of ground-breaking collaborations and discoveries. This presentation introduces an ambitious initiative: a comprehensive catalogue of CERN's open-source projects, purveyed by CERN's own OSPO. The mission? To spotlight every flag-bearing and nascent project under the CERN umbrella, making them accessible and known to the world.

This catalogue is a testament to CERN's commitment to open science and a tool to highlight all the pros of open source, foster collaboration, and stimulate innovation across the global scientific community. By curating this catalogue, the OSPO aims to not only showcase the breadth and depth of CERN's contributions to open-source software, but also to pave the way for engagement with researchers, external developers, and different institutions.

Discover how we're making open-source projects at CERN visible and why this matters for the future of scientific research. From technical challenges and solutions, to the strategic importance of open source in pushing HE(N)P discoveries forward, the journey so far has been filled with insights and stories that echo the essence that pushes for innovation at CERN. This is not just about showcasing projects; it's about building bridges in the open-source community and contributing to a legacy of open science.

**Parallel (Track 3) / 270**

## Porting and optimizing the performance of LArTPC detector simulations on Intel GPU with various programming models and compilers

**Authors:** Brett Viren<sup>1</sup>; Charles Leggett<sup>2</sup>; FNU Mohammad Atif<sup>1</sup>; Haiwang Yu<sup>None</sup>; Meifeng Lin<sup>None</sup>; Tianle Wang<sup>3</sup>; Zhihua Dong<sup>None</sup><sup>1</sup> Brookhaven National Laboratory<sup>2</sup> Lawrence Berkeley National Lab (US)<sup>3</sup> Brookhaven National Lab**Corresponding Authors:** fmohammad@bnl.gov, hyu@bnl.gov, bv@bnl.gov, twang3@bnl.gov, mlin@bnl.gov, zdong@bnl.gov, charles.g.leggett@gmail.com

There is a significant expansion in the variety of hardware architectures these years, including different GPUs and other specialized computing accelerators. For better performance portability, various programming models are developed across those computing systems, including, for example, Kokkos, SYCL, OpenMP, and std::par. Their supports on CPU and NVIDIA GPU are usually widely tested, however, their performance study on Intel GPU, especially the Intel GPU on leadership HPC clusters, is limited.

As part of the High Energy Physics Center for Computational Excellence (HEP-CCE) project, we investigate how different programming models perform on Intel GPU to see if Intel GPU may be suitable for running experimental HEP workflows with some representative use cases. One of such use cases is the Liquid Argon Time Projection Chamber (LArTPC) simulation which is essential for LArTPC detector design, validation and data analysis. We are going to present the following topics:

- 1). How different programming models are supported on Intel GPU with different compilers and their performance comparison;
- 2). Lesson learned from compiling and building projects on Intel GPU for performance optimization;
- 3). Advantages and disadvantages of different programming models used for LArTPC simulation and other HEP programs.

Parallel (Track 2) / 271

## ML-Assisted Charged Particle Tracking at GlueX

**Authors:** Ahmed Mohammed<sup>1</sup>; Kishan Rajput<sup>1</sup>

**Co-authors:** Malachi Schram<sup>1</sup>; Sergey Furlotov<sup>1</sup>; Simon Taylor<sup>1</sup>

<sup>1</sup> *Thomas Jefferson National Accelerator Facility*

**Corresponding Authors:** furlotov@jlab.org, staylor@jlab.org, schram@jlab.org, ahmedm@jlab.org, kishan@jlab.org

Tracking charged particles resulting from collisions in the presence of strong magnetic field is an important and challenging problem. Reconstructing the tracks from the hits created by those generated particles on the detector layers via ionization energy deposits is traditionally achieved through Kalman filters that scale worse than linearly as the number of hits grow. To improve efficiency there is a need for developing new tracking methods. Machine Learning (ML) has been leveraged in several science applications for both speedups and improved results. To this line, a class of ML algorithms called Graph Neural Networks (GNNs) are explored for charged particle tracking. Each event in the particle tracking data naturally imposes itself as a graph structure with the event hits represented as graph nodes while track segments are represented as a subset of the graph edges that need to be correctly classified by the ML algorithm. We compare three different approaches for tracking at GlueX experiment at Jefferson Lab, namely traditional track finding, GPU-based GNN, and FPGA-based GNN. The comparison is held in terms of inference time and performance results. Beside presenting data processing, graph construction, and the used GNN model, we provide insight into resolving the missing hits issue for GNN training and evaluation. We show that the GNN model can achieve significant speedup by processing multiple events in batches which exploits the high parallel computation capability of GPUs. We present results on real GlueX data in addition to the collective results of the simulation data.

Poster session / 272

## How to make kubernetes easy to use for thousands users

**Author:** Alberto Pimpo<sup>None</sup>

**Corresponding Author:** alberto.pimpo@cern.ch

CERN has a huge demand for computing services. To accommodate this requests, a highly-scalable and highly-dense infrastructure is necessary.

To accomplish this, CERN adopted Kubernetes, an open-source container orchestration platform that automates the deployment, scaling, and management of containerized applications.

This session will discuss the strategies and tooling used to simplify the use of Kubernetes, in particular:

- one-click deployment of any application from a git repository
- zero-config creation of CD pipelines
- specialized managed clusters for common use cases
- dashboards to manage deployments across different clusters



Parallel (Track 3) / 273

## EggNet: An Evolving Graph-based Graph Attention Network for End-to-end Particle Track Reconstruction

**Authors:** Brandon Wang<sup>1</sup>; Jay Chan<sup>2</sup>; Loic Delabrouille<sup>3</sup>; Paolo Calafiura<sup>2</sup>

<sup>1</sup> *University of California, Berkeley*

<sup>2</sup> *Lawrence Berkeley National Lab. (US)*

<sup>3</sup> *École normale supérieure de Rennes*

**Corresponding Authors:** jay.chan@cern.ch, pcalafiura@lbl.gov

Track reconstruction is a crucial task in particle experiments and is traditionally very computationally expensive due to its combinatorial nature. Recently, graph neural networks (GNNs) have emerged as a promising approach that can improve scalability. Most of these GNN-based methods, including the edge classification (EC) and the object condensation (OC) approach, require an input graph that needs to be constructed beforehand. In this work, we consider a one-shot OC approach that reconstructs particle tracks directly from a set of hits (point cloud) by recursively applying graph attention networks with an evolving graph structure. This approach iteratively updates the graphs and can better facilitate the message passing across each graph. Preliminary studies on the TrackML dataset show physics and computing performance comparable to current production algorithms for track reconstruction. We also explore different techniques to reduce constraints on computation memory and computing time.

Poster session / 274

## Xiwu: A basic flexible and learnable LLM for High Energy Physics

**Authors:** Ke Li<sup>None</sup>; Ke Li<sup>1</sup>; Siyang Chen<sup>2</sup>; Yiyu Zhang<sup>3</sup>; Zhengde Zhang<sup>4</sup>

<sup>1</sup> *University of Washington (US)*

<sup>2</sup> *IHEP, China*

<sup>3</sup> *Institute of High Energy Physics*

<sup>4</sup> *中国科学院高能物理研究所*

**Corresponding Authors:** zhangyiyu@ihep.ac.cn, like@ihep.ac.cn, ke.li@cern.ch, zdzhang@ihep.ac.cn

Large Language Models (LLMs) are undergoing a period of rapid updates and changes, with state-of-art model frequently being replaced. When applying LLMs to a specific scientific field it is challenging to acquire unique domain knowledge while keeping the model itself advanced. To address this challenge, a sophisticated large language model system named Xiwu has been developed, allowing switching the most advanced foundation models flexibly and quickly. In this talk, we will discuss one of the best practices of applying LLMs in HEP including some seed fission tools which can collect and clean the HEP dataset quickly, a just-in-time learning system based on vector store technology, and an on-the-fly fine-tuning system. The results show that Xiwu can smoothly switch different models such as LLaMa, Vicuna, chatGLM and Grok-1, and the trained Xiwu model is significantly outperformed the benchmark model on the HEP knowledge in question-and-answering and code generation.

Poster session / 275

## Boost physics study at HEP experiments with Dr. Sai

**Authors:** Ke Li<sup>None</sup>; Ke Li<sup>1</sup>; Yipu Liao<sup>2</sup>; Yiyu Zhang<sup>3</sup>; Zhengde Zhang<sup>4</sup>

<sup>1</sup> *University of Washington (US)*

<sup>2</sup> *IHEP, China*

<sup>3</sup> *Institute of High Energy Physics*

<sup>4</sup> *中国科学院高能物理研究所*

**Corresponding Authors:** like@ihep.ac.cn, ke.li@cern.ch, zdzhang@ihep.ac.cn, zhangyiyu@ihep.ac.cn

The data processing and analyzing is one of the main challenges at HEP experiments, normally one physics result can take more than 3 years to be conducted. To accelerate the physics analysis and drive new physics discovery, the rapidly developing Large Language Model (LLM) is the most promising approach, it have demonstrated astonishing capabilities in recognition and generation of text while most parts of physics analysis can be benefitted. In this talk we will discuss the construction of a dedicated intelligent agent, an AI assistant at BESIII based on LLM, the potential usage to boost hadron spectroscopy study, and the future plan towards a AI scientist.

**Poster session / 276**

## Unsupervised Learning Techniques for Identification of Anomalous LZ Data

**Author:** John Winnicki<sup>1</sup>

<sup>1</sup> *Stanford University*

**Corresponding Author:** winnicki@stanford.edu

LUX-ZEPLIN (LZ) is a dark matter direct detection experiment. Employing a dual-phase xenon time projection chamber, the LZ experiment set a world leading limit for spin-independent scattering at 36 GeV/c<sup>2</sup> in 2022, rejecting cross sections above  $9.2 \times 10^{-48}$  cm<sup>2</sup> at the 90% confidence level. Unsupervised machine learning methods are indispensable tools in working with big data, and have been applied at various stages of LZ analysis for data exploration and anomaly detection. In this work, we discuss an unsupervised dimensionality reduction approach applied to a combination of both PMT waveforms and reconstructed features aiming to identify anomalous events. We examine the tradeoffs in this method, and compare our results to known anomalies in the data, as well as conventional data quality cuts.

**Parallel (Track 5) / 277**

## Storage of nuclear waste suitable for non-invasive monitoring using muon scattering tomography

**Author:** Purba Bhattacharya<sup>1</sup>

**Co-authors:** Indira Mukherjee<sup>1</sup>; Nayana Majumdar<sup>2</sup>; Piyush Pallav<sup>1</sup>; Promita Roy<sup>3</sup>; Sridhar Tripathy<sup>4</sup>; Subhendu Das<sup>5</sup>; Supratik Mukhopadhyay<sup>5</sup>

<sup>1</sup> *Adamas University, Kolkata, India*

<sup>2</sup> *Saha Institute of Nuclear Physics, Kolkata, India and A CI of Homi Bhabha National Institute, Kolkata, India*

<sup>3</sup> *Virginia Polytechnic Institute & State University, United States*

<sup>4</sup> *University of California, Davis, United States*

<sup>5</sup> *Applied Nuclear Physics Division, Saha Institute of Nuclear Physics, A CI of Homi Bhabha National Institute, Kolkata, India*

**Corresponding Authors:** supratik.mukhopadhyay@saha.ac.in, nayana.majumdar@saha.ac.in, promita@vt.edu, piyush1.pallav@adamasuniversity.ac.in, subhendudas456038@gmail.com, indira.mukherjee@stu.adamasuniversity.ac.in, sritripathy@ucdavis.edu, purba.bhattacharya1985@gmail.com

Over the past few decades, there has been a noticeable surge in muon tomography research, also referred to as muography. This method, falling under the umbrella of Non-Destructive Evaluation (NDE), constructs a three-dimensional image of a target object by harnessing the interaction between cosmic ray muons and matter, akin to how radiography utilizes X-rays. Essentially, muography entails scanning a target object by analyzing its interaction with muons, with the interaction mode contingent upon the energy of the incident muon and the characteristics of the medium involved. As cosmic muons interact electromagnetically with atoms within the target medium, their trajectories are likely to deviate prior to reaching the position sensitive detectors placed at suitable locations around the object under study. These deviations serve as a rich source of data that can be used to generate images and infer the material composition of the target.

In this study, a numerical simulation has been conducted using the GEANT4 framework to assess the efficacy of various position sensitive charged particle detectors in muography. The feasibility of detectors with a broad range of position resolutions has been tested, particularly in the context of developing an imaging algorithm to monitor drums containing nuclear waste. The Cosmic Ray Shower Library (CRY) has been employed to simulate muon showers on the detector-target system. The reconstruction of muon tracks, crucial for analyzing muon scattering, has been achieved by collecting hits from all detector layers. Incoming muon tracks have been reconstructed using hits from the upper set of detectors, while outgoing muon tracks have been reconstructed using hits from the lower set. In this presentation, the discussion will center on track reconstruction algorithms, emphasizing the use of efficient single scattering point algorithms like Point of Closest Approach (PoCA) for simplified implementation and fast computation. To enhance material discrimination confidence, a Support Vector Machine (SVM) based algorithm has been applied, utilizing features such as scattering vertices density ( $\langle N \rangle$ ) and average deviation angle ( $\langle \theta \rangle$ ) as inputs. SVM hyperplanes have been generated to segregate various material classes, and corresponding confusion matrices have been obtained. Additionally, for analyzing the shape of materials within nuclear waste drums, an algorithm based on Pattern Recognition Method (PRM) has been employed. This presentation will delve into studies of track reconstruction algorithms applied to GEANT4 data for particle detectors with varying position resolutions, followed by shape and image analysis based on the PRM with the motivation of optimizing storage of nuclear waste that can be efficiently monitored by techniques such as muography.

Parallel (Track 5) / 278

## End-to-end event simulation with Flow Matching and generator Oversampling

**Authors:** Filippo Cattafesta<sup>1</sup>; Francesco Vaselli<sup>1</sup>; Patrick Asenov<sup>2</sup>; Andrea Rizzi<sup>2</sup>

<sup>1</sup> *Scuola Normale Superiore & INFN Pisa (IT)*

<sup>2</sup> *Universita & INFN Pisa (IT)*

**Corresponding Authors:** filippo.cattafesta@cern.ch, andrea.rizzi@cern.ch, patrick.asenov.asenov@cern.ch, francesco.vaselli@cern.ch

The event simulation is a key element for data analysis at present and future particle accelerators. We show that novel machine learning algorithms, specifically Normalizing Flows and Flow Matching, can be effectively used to perform accurate simulations with several orders of magnitude of speed-up compared to traditional approaches when only analysis level information is needed. In such a case it is indeed feasible to skip the whole simulation chain and directly simulate analysis observables from generator information (end-to-end simulation). We simulate jets features to compare discrete and continuous Normalizing Flows models. The models are validated across a variety of metrics to select the best ones. We discuss the scaling of performance with the increase in training data, as well as the generalization power of these models on physical processes different from the training one. We investigate sampling multiple times from the same inputs, a procedure we call oversampling, and we show that it can effectively reduce the statistical uncertainties of a sample. This class of ML

algorithms is found to be highly expressive and useful for the task of simulation. Their speed and accuracy, coupled with the stability of the training procedure, make them a compelling tool for the needs of current and future experiments.

1 arXiv:2402.13684

Parallel (Track 3) / 279

## End-to-end ML-based reconstruction

**Authors:** Brieuc Francois<sup>1</sup>; Dolores Garcia<sup>1</sup>; Michele Selvaggi<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** maria.dolores.garcia.marti@cern.ch, brieuc.francois@cern.ch, michele.selvaggi@cern.ch

We present an ML-based end-to-end algorithm for adaptive reconstruction in different FCC detectors. The algorithm takes detector hits from different subdetectors as input and reconstructs higher-level objects. For this, it exploits a geometric graph neural network, trained with object condensation, a graph segmentation technique. We apply this approach to study the performance of pattern recognition in the IDEA detector using hits from the pixel vertex detector and the drift chamber. We also build particle candidates from detector hits and tracks in the CLD detector. Our algorithm outperforms current baselines in efficiency and energy reconstruction and allows pattern recognition in the IDEA detector. This approach is easily adaptable to new geometries and therefore opens the door to reconstruction performance-aware detector optimization.

Poster session / 280

## ATLAS Tile Calorimeter Software Tools for Data Quality Monitoring

**Author:** Danijela Bogavac<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** danijela.bogavac@cern.ch

The ATLAS Tile Calorimeter (TileCal) is the central hadronic calorimeter of the ATLAS detector at the Large Hadron Collider at CERN. It plays an important role in the reconstruction of jets, hadronically decaying tau leptons and missing transverse energy, and also provides information to the dedicated calorimeter trigger. The TileCal readout is segmented into nearly 10000 channels that are calibrated using the dedicated calibration systems such as laser, charge injection, integrator and Cesium source.

Data quality assurance is paramount, with collision and calibration data subject to rigorous scrutiny. Automated checks are performed on predefined histograms, and the results are summarized on dedicated web pages. Operators use a suite of tools to further inspect the data and identify any issues or irregularities. The TileCal conditions data, including calibration constants and channel statuses, are therefore regularly updated in databases. These databases are used for data reprocessing and are also crucial for maintenance work during the technical stops.

In this talk, we will discuss the software tools used for data quality monitoring, emphasizing recent advancements and our pursuit of consolidating multiple tools into a more streamlined web application. Our overarching goal is to optimize the efficiency of the shifters responsible for monitoring data quality while simultaneously simplifying the entire process.

**Parallel (Track 3) / 281**

## Optimizing cut-based algorithms to specific physics acceptance regions

**Author:** Benjamin Huber<sup>1</sup>**Co-authors:** Dinyar Rabady<sup>2</sup>; Elias Leutgeb<sup>1</sup>; Gabriele Bortolato<sup>3</sup>; Hannes Sakulin<sup>2</sup>; Jaana Heikkilae<sup>2</sup>; Maria Cepeda<sup>4</sup><sup>1</sup> *Technische Universitaet Wien (AT)*<sup>2</sup> *CERN*<sup>3</sup> *Universita e INFN, Padova (IT)*<sup>4</sup> *CIEMAT***Corresponding Authors:** gabriele.bortolato@cern.ch, benjamin.huber@cern.ch, jaana.kristiina.heikkilae@cern.ch, hannes.sakulin@cern.ch, elias.leutgeb@cern.ch, dinyar.rabady@cern.ch, maria.cepeda@cern.ch

Developments of the new Level-1 Trigger at CMS for the High-Luminosity Operation of the LHC are in full swing. The Global Trigger, the final stage of this new Level-1 Trigger pipeline, is foreseen to evaluate a menu of over 1000 cut-based algorithms, each of which targeting a specific physics signature or acceptance region. Automating the task of tailoring individual algorithms to specific physics regions would be a significant time saver while ensuring flexibility to adapt swiftly to evolving run conditions. This task essentially resembles a multi-objective optimization problem, where the goal is to strike a balance between the trigger rate and the trigger efficiency of the desired physics region.

We present the idea of leveraging achievement scalarization, a technique to turn the two objective functions into a scalar function with a minimum closest to a reference point chosen by a decision maker. An iterative gradient descent approach can then be employed to minimize this function, each iteration slightly modifying the cut parameters in the direction of descent. The decision maker in this context can be a single person designing parts of the menu or a collective group like CERN's data performance group agreeing on specific goals for upcoming data-taking sessions.

Preliminary results of using this procedure in targeting B meson decays have demonstrated promising outcomes. Ongoing efforts involve exploring alternative minimization techniques like evolutionary algorithms and extending the method to other physics signatures.

**Parallel (Track 1) / 282**

## Data Movement Model for the Vera C. Rubin Observatory

**Authors:** Andrew Hanushevsky<sup>1</sup>; Fabio Hernandez<sup>2</sup>; George Beckett<sup>None</sup>; Kian-Tat Lim<sup>3</sup>; Peter Love<sup>4</sup>; Stephen R. Pietrowicz<sup>5</sup>; Tim Jenness<sup>6</sup>; Timothy John Noble<sup>7</sup>; Wei Yang<sup>8</sup><sup>1</sup> *Stanford University/SLAC*<sup>2</sup> *IN2P3 / CNRS computing centre*<sup>3</sup> *SLAC National Accelerator Laboratory, USA*<sup>4</sup> *Lancaster University (GB)*<sup>5</sup> *National Center for Supercomputing Applications, USA*<sup>6</sup> *Vera C. Rubin Observatory, USA*<sup>7</sup> *Science and Technology Facilities Council STFC (GB)*<sup>8</sup> *SLAC National Accelerator Laboratory (US)*

**Corresponding Authors:** [tjenness@lsst.org](mailto:tjenness@lsst.org), [ktl@slac.stanford.edu](mailto:ktl@slac.stanford.edu), [fabio.hernandez@cern.ch](mailto:fabio.hernandez@cern.ch), [george.beckett@ed.ac.uk](mailto:george.beckett@ed.ac.uk), [timothy.noble@stfc.ac.uk](mailto:timothy.noble@stfc.ac.uk), [srp@illinois.edu](mailto:srp@illinois.edu), [abh@slac.stanford.edu](mailto:abh@slac.stanford.edu), [peter.love@cern.ch](mailto:peter.love@cern.ch), [yangw@slac.stanford.edu](mailto:yangw@slac.stanford.edu)

The set of sky images recorded nightly by the camera mounted on the telescope of the Vera C. Rubin Observatory will be processed in facilities located on three continents. Data acquisition will happen in Cerro Pachón in the Andes mountains in Chile where the observatory is located. A first copy of the raw image data set is stored at the summit site of the observatory and immediately transferred through dedicated network links to the archive site and US Data Facility hosted at SLAC National Laboratory in California, USA. After an embargo period of a few days, the full image set is copied to the UK and French Data Facilities where a third copy is located.

During its 10 years in operation starting late 2025, annual processing campaigns across all images taken to date will be jointly performed by the three facilities, involving sophisticated algorithms to extract the physical properties of the celestial objects and producing science-ready images and catalogs. Data products resulting from the processing campaigns at each facility will be sent to SLAC and combined to create a consistent Data Release which is served to the scientific community for its science studies via Data Access Centers in the US and Chile and Independent Data Access Centers elsewhere.

In this contribution we present an overall view of how we leverage the tools selected for managing the movement of data among the Rubin processing and serving facilities, including Rucio and FTS3. We will also present the tools we developed to integrate Rucio's data model and Rubin's Data Butler, the software abstraction layer that mediates all access to storage by the pipeline tasks which implement the science algorithms.

**Poster session / 283**

## Surrogate Modeling for Scalable Evaluation of Distributed Computing Systems for HEP applications

**Authors:** Larissa Schmid<sup>1</sup>; Maximilian Maria Horzela<sup>2</sup>; Valerii Zhyla<sup>1</sup>; Manuel Giffels<sup>2</sup>; Gunter Quast<sup>2</sup>; Anne Koziol<sup>1</sup>

<sup>1</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>2</sup> *KIT - Karlsruhe Institute of Technology (DE)*

**Corresponding Authors:** [larissa.schmid@kit.edu](mailto:larissa.schmid@kit.edu), [manuel.giffels@cern.ch](mailto:manuel.giffels@cern.ch), [gunter.quast@cern.ch](mailto:gunter.quast@cern.ch), [valerii.zhyla@student.kit.edu](mailto:valerii.zhyla@student.kit.edu), [maximilian.horzela@cern.ch](mailto:maximilian.horzela@cern.ch), [koziol@kit.edu](mailto:koziol@kit.edu)

The sheer volume of data generated by LHC experiments presents a computational challenge, necessitating robust infrastructure for storage, processing, and analysis. The Worldwide LHC Computing Grid (WLCG) addresses this challenge by integrating global computing resources into a cohesive entity. To cope with changes in the infrastructure and increased demands, the compute model needs to be adapted. Simulations of different compute models present a feasible approach for evaluating different design candidates. However, running these simulations incurs a trade-off between accuracy and scalability. For example, while the simulator DCSim can provide accurate results, it falls short on scalability when increasing the size of the simulated platform. Generative Machine Learning as a surrogate is successfully used to overcome these limitations in other domains that exhibit similar trade-offs between scalability and accuracy, such as the simulation of detectors.

In our work, we evaluate the usage of three different machine learning models as surrogate models for the simulation of distributed computing systems and assess their ability to generalize to unseen jobs and platforms. We show that those models can predict the simulated platforms' main observables derived from the execution traces of compute jobs with approximate accuracy. Potential for further improving the predictions lies in using other machine learning models and different encodings of the platform-specific information to achieve better generalizability for unseen platforms.

**Parallel (Track 3) / 284**

## Towards end-to-end particle flow in high granularity detectors

**Author:** Philipp Zehetner<sup>1</sup>**Co-authors:** Jan Kieseler<sup>2</sup>; Shah Rukh Qasim<sup>3</sup><sup>1</sup> *Ludwig Maximilians Universitat (DE)*<sup>2</sup> *KIT - Karlsruhe Institute of Technology (DE)*<sup>3</sup> *University of Zurich (CH)***Corresponding Authors:** philipp.zehetner@cern.ch, shah.rukh.qasim@cern.ch, jan.kieseler@cern.ch

We present an end-to-end reconstruction algorithm for highly granular calorimeters that includes track information to aid the reconstruction of charged particles. The algorithm starts from calorimeter hits and reconstructed tracks, and outputs a coordinate transformation in which all shower objects are well separated from each other, and in which clustering becomes trivial. Shower properties such as particle ID and energy are predicted from representative points within showers. This is achieved using an extended version of the object condensation loss, a graph segmentation technique that allows the clustering of a variable number of showers in every event while simultaneously performing regression and classification tasks. The backbone is an architecture based on a newly-developed translation-equivariant version of GravNet layers. These dynamically build learnable graphs from input data to exchange information along their edges. The model is trained on data from a simulated detector that matches the complexity of the CMS high-granularity calorimeter (HGCAL).

**Parallel (Track 5) / 285**

## EDM4hep - The common event data model for the Key4hep project

**Authors:** Andre Sailer<sup>1</sup>; Benedikt Hegner<sup>1</sup>; Brieuc Francois<sup>1</sup>; Frank-Dieter Gaede<sup>2</sup>; Juan Miguel Carceller<sup>1</sup>; Juraj Smiesko<sup>1</sup>; Mateusz Jakub Fila<sup>1</sup>; Thomas Madlener<sup>3</sup><sup>1</sup> *CERN*<sup>2</sup> *Deutsches Elektronen-Synchrotron (DE)*<sup>3</sup> *Deutsches Elektronen-Synchrotron (DESY)***Corresponding Authors:** frank-dieter.gaede@cern.ch, mateusz.jakub.fila@cern.ch, andre.philippe.sailer@cern.ch, benedikt.hegner@cern.ch, juraj.smiesko@cern.ch, j.m.carcell@cern.ch, thomas.madlener@cern.ch, brieuc.francois@cern.ch

The common and shared event data model EDM4hep is a core part of the Key4hep project. It is the component that is used to not only exchange data between the different software pieces, but it also serves as a common language for all the components that belong to Key4hep. Since it is such a central piece, EDM4hep has to offer an efficient implementation. On the other hand, EDM4hep has to be flexible enough in order to allow for new developments in detector technology and reconstruction. In order to meet these challenges EDM4hep is using the podio EDM toolkit to generate its implementation from a high level description.

In this talk we give an overview of EDM4hep emphasizing the most recent developments that were tackled on the way to a first stable release. We use this opportunity to also highlight the latest developments in the podio toolkit that were required by the latest EDM4hep features. These include the introduction of type erased interface types, and a new generic RDataSource to support the full data model API in RDataFrame.

Poster session / 286

## OMS Data Aggregation and Management in the CMS Experiment

**Author:** Andrea Petrucci<sup>1</sup>

**Co-authors:** Albert Corominas I Mariscot<sup>2</sup>; Andrea Bocci<sup>3</sup>; Antonin Dvorak<sup>3</sup>; Antra Gaile<sup>4</sup>; Attila Racz<sup>3</sup>; Babatunde John Odetayo<sup>5</sup>; Christian Deldicque<sup>3</sup>; Christoph Paus<sup>6</sup>; Christoph Schwick<sup>3</sup>; Christos Emmanouil; Cristina Vazquez Velez<sup>3</sup>; Dainius Simelevicius<sup>7</sup>; Dinyar Rabady<sup>3</sup>; Dominique Gigi<sup>3</sup>; Emilio Meschi<sup>3</sup>; Eric Cano<sup>3</sup>; Frank Glege<sup>3</sup>; Frans Meijers<sup>3</sup>; Georgiana Lavinia Darlea<sup>6</sup>; Guillermo Gomez Ceballos Retuerto<sup>6</sup>; Guillermo Izquierdo Moreno<sup>3</sup>; Hannes Sakulin<sup>3</sup>; Jaafar Alawieh<sup>8</sup>; James Gordon Branson<sup>1</sup>; Jan Andrzej Bugajski<sup>9</sup>; Jeroen Hegeman<sup>3</sup>; Kareen Arutjunjan<sup>3</sup>; Kenneth Long<sup>6</sup>; Luciano Orsini<sup>3</sup>; Marc Dobson<sup>3</sup>; Marco Pieri<sup>1</sup>; Miguel Bacharov Durasov<sup>3</sup>; Patrycja Ewa Gorniak<sup>3</sup>; Petr Zejdl<sup>3</sup>; Philipp Brummer<sup>3</sup>; Polyneikis Tzanis<sup>3</sup>; Rafal Dominik Krawczyk<sup>10</sup>; Sergio Cittolin<sup>1</sup>; Srecko Morovic<sup>1</sup>; Tejeswini Jayakumar<sup>3</sup>; Theodoros Rizopoulos<sup>3</sup>; Thomas Owen James<sup>3</sup>; Ulf Behrens<sup>10</sup>; Wassef Karimeh<sup>3</sup>; Wei Li<sup>10</sup>

<sup>1</sup> Univ. of California San Diego (US)

<sup>2</sup> University of Girona UdG (ES)

<sup>3</sup> CERN

<sup>4</sup> Riga Technical University (LV)

<sup>5</sup> University of Benin (NG)

<sup>6</sup> Massachusetts Inst. of Technology (US)

<sup>7</sup> Vilnius University (LT)

<sup>8</sup> American University of Beirut (LB)

<sup>9</sup> AGH University of Krakow (PL)

<sup>10</sup> Rice University (US)

**Corresponding Authors:** miguel.durasov.bacharov@cern.ch, luciano.orsini@cern.ch, cristina.vazquez@cern.ch, antonin.dvorak@cern.ch, dominique.gigi@cern.ch, tejeswini.jayakumar@cern.ch, andrea.petrucci@cern.ch, philipp.maximilian.brummer@cern.ch, jan.andrzej.bugajski@cern.ch, branson@ucsd.edu, polyneikis.tzanis@cern.ch, emilio.meschi@cern.ch, marc.dobson@cern.ch, theodoros.rizopoulos@cern.ch, albert.corominas.i.mariscot@cern.ch, rafal.dominik.krawczyk@cern.ch, attila.racz@cern.ch, marco.pieri@cern.ch, antra.gaile@cern.ch, ulf.behrens@cern.ch, frans.meijers@cern.ch, sergio.cittolin@cern.ch, pagornia@cern.ch, frank.glege@cern.ch, christian.deldicque@cern.ch, petr.zejdl@cern.ch, christoph.schwick@cern.ch, jeroen.hegeman@cern.ch, babatunde.john.odetayo@cern.ch, jaafar.alawieh@cern.ch, wei.li@cern.ch, kenneth.long@cern.ch, guillermo.gomez.cebaldos@cern.ch, kareen.arutjunjan@cern.ch, wassef.karimeh@cern.ch, guillermo.izquierdo.moreno@cern.ch, christos.emmanouil@cern.ch, dinyar.rabady@cern.ch, hannes.sakulin@cern.ch, lavinia.darlea@cern.ch, thomas.owen.james@cern.ch, eric.cano@cern.ch, dainius.simelevicius@cern.ch, andrea.bocci@cern.ch, paus@mit.edu, srecko.morovic@cern.ch

The OMS data warehouse (DWH) constitutes the foundation of the Online Monitoring System (OMS) architecture within the CMS experiment at CERN, responsible for the storage and manipulation of non-event data within ORACLE databases. Leveraging on PL/SQL code, the DWH orchestrates the aggregation and modification of data from several sources, inheriting and revamping code from the previous project known as Web Based Monitoring to meet evolving requirements. The main goals of the DWH restructuring were: the modernization of inherited PL/SQL code, necessitating the creation of new aggregation tables and the implementation of enhancements such as standardized naming conventions; improved development workflows; and continuous integration strategies. DWH is composed of multiple Oracle schemas and integrates external PL/SQL libraries, in particular the CERN Beams Common4Oracle library, which consolidates common functionalities from various CERN Beams department databases into a unified codebase for widespread application. This article delves into the architecture and development strategies employed within the OMS data warehouse, underscoring its role in facilitating efficient data aggregation and management within the OMS project in the CMS experiment at CERN.

Poster session / 287

## Lessons Learned from Writing My Dissertation With Quarto

**Author:** Eli Mizrachi<sup>1</sup>



<sup>1</sup> *SLAC National Accelerator Laboratory*

**Corresponding Author:** emiz@berkeley.edu

Taking on a new project as a graduate student is a daunting task, especially if the only documentation is someone else's dissertation or talk, and their code is nowhere to be found. In deciding how I would write my own dissertation, I wanted to branch out from a typical Overleaf/LaTeX workflow, and take advantage of cutting-edge software that would integrate more with the code I wrote for my own analyses. Ideally, I would be able to write in markdown, and directly reference Jupyter notebooks and/or Python scripts instead of having to export plots repeatedly.

In my search for a documentation system that would resolve these problems, I stumbled upon Quarto. Quarto 1.0 was announced in 2022 as an evolution of R Markdown. It is built on top of Pandoc, with first-class support for R, Python, Julia, and Observable. It is capable of tightly weaving together code with long-form documentation, and renders PDFs, as well as HTML content from the same source files. In this talk, I will discuss putting Quarto to the test of producing my dissertation and defense slides. I will introduce Quarto, discuss the best practices I learned from using it, and summarize my outlook for introducing it to other members of the LZ dark matter search experiment.

**Parallel (Track 1) / 288**

## Validation of Shoveler XRootD monitoring

**Authors:** Alessandra Forti<sup>1</sup>; Borja Garrido Bear<sup>2</sup>; Derek Weitzel<sup>3</sup>; Julia Andreeva<sup>2</sup>; Katy Ellis<sup>4</sup>; Latchezar Betev<sup>2</sup>

<sup>1</sup> *University of Manchester (GB)*

<sup>2</sup> *CERN*

<sup>3</sup> *University of Nebraska Lincoln (US)*

<sup>4</sup> *Science and Technology Facilities Council STFC (GB)*

**Corresponding Authors:** borja.garrido.bear@cern.ch, katy.ellis@stfc.ac.uk, alessandra.forti@cern.ch, derek.weitzel@cern.ch, latchezar.betev@cern.ch, julia.andreeva@cern.ch

The Large Hadron Collider (LHC) experiments rely heavily on the XRootD software suite for data transfer and streaming across the Worldwide LHC Computing Grid (WLCG) both within sites (LAN) and across sites (WAN). While XRootD offers extensive monitoring data, there's no single, unified monitoring tool for all experiments. This becomes increasingly critical as network usage grows, and with the High-Luminosity LHC (HL-LHC) demanding even higher bandwidths.

The "Shoveler" system addresses this challenge by providing a platform to collect and visualize XRootD traffic data from all four LHC experiments, separated by type, direction and locality of the traffic. This contribution explores the Shoveler plus Collector architecture, its current deployment status at WLCG sites, and validates its collected information by comparing it with data from individual experiment monitoring frameworks.

**Parallel (Track 7) / 289**

## Heterogeneous computing at INFN-T1

**Authors:** Alessandro Pascolini<sup>1</sup>; Andrea Chierici<sup>1</sup>; DIEGO MICHELOTTO<sup>2</sup>; Daniele Lattanzio<sup>None</sup>; Francesco Noferini<sup>1</sup>; Giusy Sergi<sup>3</sup>

<sup>1</sup> *Universita e INFN, Bologna (IT)*

<sup>2</sup> *INFN - National Institute for Nuclear Physics*

<sup>3</sup> *INFN - CNAF*

**Corresponding Authors:** andrea.chierici@cern.ch, francesco.noferini@cern.ch, giusy.sergi@inf.n.cnaf.it, daniele.lattanzio@cnaf.infn.it, alessandro.pascolini@cern.ch, diego.michelotto@cnaf.infn.it

At INFN-T1 we recently acquired some ARM nodes: initially they were given to LHC experiments to test workflow and submission pipelines. After some time, they were given as standard CPU resources, since the stability both of the nodes and of the code was production quality ready.

In this presentation we will describe all the activities that were necessary to enable users to run on ARM and will give some figures on performance, compared to x86 counterpart. In the end we will try to describe our point of view for the possible massive adoption of this architecture in Tier1 data centers.

Poster session / 290

## The Big Data Processing Infrastructure for monitoring and analysing the ATLAS experiment processing activities at INFN-CNAF Tier-1

**Authors:** Aksieniia Shtimmerman<sup>None</sup>; Antonio Falabella<sup>1</sup>; DIEGO MICHELOTTO<sup>2</sup>; Enrico Fattibene<sup>3</sup>; Giacomo Levri<sup>4</sup>; Giusy Sergi<sup>3</sup>

<sup>1</sup> INFN CNAF

<sup>2</sup> INFN - National Institute for Nuclear Physics

<sup>3</sup> INFN - CNAF

<sup>4</sup> Università e INFN, Bologna (IT)

**Corresponding Authors:** giusy.sergi@inf.n.cnaf.it, aksieniia.shtimmerman@cnaf.infn.it, antonio.falabella@cnaf.infn.it, enrico.fattibene@cnaf.infn.it, diego.michelotto@cnaf.infn.it, giacomo.levrini@cern.ch

The modern data centers provide the efficient Information Technologies (IT) infrastructure needed to deliver resources, services, monitoring systems and collected data in a timely fashion. At the same time, data centres have been continuously evolving, foreseeing large increase of resources and adapting to cover multifaced niches.

The CNAF group at INFN (National Institute for Nuclear Physics) has implemented a Big Data Platform (BDP) infrastructure, designed for the collection and the indexing of log reports from CNAF facilities. The infrastructure is an ongoing project at CNAF and it is at service of the Italian groups working in high energy physics experiments. Within this framework, the first data pipeline was established for the ATLAS experiment, using input from the ATLAS Distributed Computing system PanDa.

This pipeline focuses on the ATLAS computational job data processed by the Italian INFN Tier-1 computing farm. The system has been operational and effective for several years, marking our initiative as the first to integrate job information directly with the infrastructure. Following the finalization of data transmission, our objective is to conduct an analysis and surveillance of the PanDA Jobs data. This will involve examining the performance metrics of the machines and identifying the log errors that lead to job failures.

Parallel (Track 2) / 291

## Real-time compression of high-speed Streaming Readout signals with Artificial Intelligence algorithms.

**Authors:** Fabio Rossi<sup>None</sup>; Marco Battaglieri<sup>1</sup>; Edoardo Ragusa<sup>2</sup>; Paolo Gastaldo<sup>2</sup>; Gagik Gavalian<sup>3</sup>

<sup>1</sup> INFN

<sup>2</sup> UNIGE

<sup>3</sup> Jefferson Lab

**Corresponding Authors:** battaglieri@ge.infn.it, fabio.rossi@ge.infn.it

The new generation of high-energy physics experiments plans to acquire data in streaming mode. With this approach, it is possible to access the information of the whole detector (organized in time slices) for optimal and lossless triggering of data acquisition. Each front-end channel sends data to the processing node via TCP/IP when an event is detected. The data rate in large detectors is often very high and the network is likely to be a bottleneck for the system. On the other hand, the network devices do not need to know the signal shape but only the hit timing (time-stamp) and the address of the front end that generated it. This is a key point to implement a compression algorithm: it is possible to compress signal samples after the front end and decompress them when data are needed for high-level analysis.

To achieve a high compression ratio and fast inference time on hardware an AI-supported algorithm, an autoencoder, is chosen. Autoencoder is an unsupervised machine learning algorithm composed of two parts: an encoder and a decoder that reduces the size of the input and reconstructs the original input from the encoded representation, respectively.

This contribution describes the compression algorithm and the Streaming Readout DAQ (SRO) system prototype developed to test it. The SRO prototype is designed with three separate nodes connected to the same network. We use a PC as a proxy of the final high-level analysis node and two Raspberry Pi single-board computers as signal generators and data processing units (compressors). The architecture was made such that each node could/will be easily replaced with faster hardware (eg. an FPGA).

Consideration concerning the compression algorithm complexity, loss during compression, and execution time are taken into account to achieve the best tradeoff. Results of autoencoder training and timing of some implemented configurations are reported.

Poster session / 292

## Evaluating a File-Based Event Builder to enhance the Data Acquisition in the CMS Experiment

**Authors:** Andrea Petrucci<sup>1</sup>; Srecko Morovic<sup>1</sup>

**Co-authors:** Albert Corominas I Mariscot<sup>2</sup>; Andrea Bocci<sup>3</sup>; Antonin Dvorak<sup>3</sup>; Antra Gaile<sup>4</sup>; Attila Racz<sup>3</sup>; Babatunde John Odetayo<sup>5</sup>; Christian Deldicque<sup>3</sup>; Christoph Paus<sup>6</sup>; Christoph Schwick<sup>3</sup>; Christos Emmanouil; Cristina Vazquez Velez<sup>3</sup>; Dainius Simelevicius<sup>7</sup>; Dinyar Rabady<sup>3</sup>; Dominique Gigi<sup>3</sup>; Emilio Meschi<sup>3</sup>; Eric Cano<sup>3</sup>; Frank Glege<sup>3</sup>; Frans Meijers<sup>3</sup>; Georgiana Lavinia Darlea<sup>6</sup>; Guillermo Gomez Ceballos Retuerto<sup>6</sup>; Guillermo Izquierdo Moreno<sup>3</sup>; Hannes Sakulin<sup>3</sup>; Jaafar Alawieh<sup>8</sup>; James Gordon Branson<sup>1</sup>; Jan Andrzej Bugajski<sup>9</sup>; Jeroen Hegeman<sup>3</sup>; Kareen Arutjunjan<sup>3</sup>; Kenneth Long<sup>6</sup>; Luciano Orsini<sup>3</sup>; Marc Dobson<sup>3</sup>; Marco Pieri<sup>1</sup>; Miguel Bacharov Durasov<sup>3</sup>; Patrycja Ewa Gorniak<sup>3</sup>; Petr Zejdl<sup>3</sup>; Philipp Brummer<sup>3</sup>; Polyneikis Tzanis<sup>3</sup>; Rafal Dominik Krawczyk<sup>10</sup>; Sergio Cittolin<sup>1</sup>; Tejeswini Jayakumar<sup>3</sup>; Theodoros Rizopoulos<sup>3</sup>; Thomas Owen James<sup>3</sup>; Ulf Behrens<sup>10</sup>; Wassef Karimeh<sup>3</sup>; Wei Li<sup>10</sup>

<sup>1</sup> Univ. of California San Diego (US)

<sup>2</sup> University of Girona UdG (ES)

<sup>3</sup> CERN

<sup>4</sup> Riga Technical University (LV)

<sup>5</sup> University of Benin (NG)

<sup>6</sup> Massachusetts Inst. of Technology (US)

<sup>7</sup> Vilnius University (LT)

<sup>8</sup> *American University of Beirut (LB)*

<sup>9</sup> *AGH University of Krakow (PL)*

<sup>10</sup> *Rice University (US)*

**Corresponding Authors:** wassef.karimeh@cern.ch, guillermo.gomez.ceballos@cern.ch, antonin.dvorak@cern.ch, dinyar.rabady@cern.ch, philipp.maximilian.brunner@cern.ch, dainius.simelevicius@cern.ch, marco.pieri@cern.ch, christian.deldicque@cern.ch, kenneth.long@cern.ch, marc.dobson@cern.ch, jeroen.hegeman@cern.ch, rafal.dominik.krawczyk@cern.ch, albert.corominas.i.mariscot@cern.ch, thomas.owen.james@cern.ch, christoph.schwick@cern.ch, hannes.sakulin@cern.ch, frans.meijers@cern.ch, sergio.cittolin@cern.ch, emilio.meschi@cern.ch, frank.glege@cern.ch, cristina.vazquez@cern.ch, antra.gaile@cern.ch, miguel.durasov.bacharov@cern.ch, theodoros.rizopoulos@cern.ch, pagornia@cern.ch, jan.andrzej.bugajski@cern.ch, babatunde.john.odetayo@cern.ch, lavinia.darlea@cern.ch, guillermo.izquierdo.moreno@cern.ch, polyneikis.tzani@cern.ch, kareen.arutjunjan@cern.ch, christos.emmanouil@cern.ch, dominique.gigi@cern.ch, andrea.bocci@cern.ch, luciano.orsini@cern.ch, ulf.behrens@cern.ch, srecko.morovic@cern.ch, jaafar.alawieh@cern.ch, tejeswini.jayakumar@cern.ch, wei.li@cern.ch, andrea.petrucchi@cern.ch, petr.zejdl@cern.ch, attila.racz@cern.ch, paus@mit.edu, branson@ucsd.edu, eric.cano@cern.ch

The event builder in the Data Acquisition System (DAQ) of the CMS experiment at the CERN Large Hadron Collider (LHC) is responsible for assembling events at a rate of 100 kHz during the current LHC run 3, and 750 kHz for the upcoming High Luminosity LHC, scheduled to start in 2029. Both the current and future DAQ architectures leverage on state-of-the-art network technologies, employing Ethernet switches capable of supporting RoCE protocols. The DAQ Front-end hardware is custom-designed, utilizing a reduced TCP/IP protocol implemented in FPGA for reliable data transport between custom electronics and commercial computing hardware.

An alternative architecture for the event builder, known as the File-based Event Builder (FEVB), is under evaluation. The FEVB comprises two separate systems: the Super-Fragment Builder (SFB) and the Builder File-based Filter Farm (BF<sup>3</sup>).

A super-fragment consists of the event data read by one or more Front-End Drivers and corresponding to the same L1 accept, and the SFB constructs multiple super-fragments corresponding to the number of Read-Unit (RU) machines in the DAQ system, storing them in local RAM disks. Subsequently, the BF<sup>3</sup> accesses super-fragments from all RU machines via the Network File System (NFS) over Ethernet and builds complete events within the High Level Trigger process.

This paper describes the first prototype of the FEVB and presents preliminary performance results obtained within the DAQ system for LHC Run 3.

Parallel (Track 5) / 293

## ML-based classification of photons and neutral mesons for direct photon measurement in ALICE

**Author:** Abhishek Nath<sup>1</sup>

<sup>1</sup> *Heidelberg University (DE)*

**Corresponding Author:** aabhishek.naath@gmail.com

Direct photons are unique probes to study and characterize the quark-gluon plasma (QGP) as they leave the collision medium mostly unscathed. Measurements at top Large Hadron Collider (LHC) energies at low pT reveal a very small thermal photon signal accompanied by considerable systematic uncertainties. Reduction of such uncertainties, which arise from the  $\pi^0$  and  $\eta$  measurements, as well as the photon identification, is crucial for the comparison of the results with the theoretical calculations that are available.

To address these challenges, a novel approach employing machine learning (ML) techniques has been implemented for the classification of photons and neutral mesons. An open-source set of frameworks comprising hipe4ml, scikit-learn, and ONNX packages is chosen for training, validation, and testing the model on a part of Run2 Pb–Pb data at  $\sqrt{s_{NN}} = 5.02$  TeV collision energy.

In this talk, the performance of the novel approach in comparison to the standard cut-based analysis is presented. Initial findings employing gradient-boosted decision trees demonstrate a substantial enhancement in photon purity while preserving efficiency levels comparable to those of the standard cut-based method. Strategies for addressing highly imbalanced data sets, including techniques

like feature reduction during training and the implementation of scaled penalty factors to enhance discrimination between signal and background are also addressed. Finally, the feasibility of incorporating such ML methods into the main workflow of direct photon analysis is also presented.

Parallel (Track 1) / 294

## A Tape RSE for Extremely Large Data Collection Backups

**Author:** Andrew Bohdan Hanushevsky<sup>1</sup>

**Co-authors:** Guangwei Che <sup>2</sup>; Lance Nakata <sup>2</sup>; Wei Yang <sup>1</sup>

<sup>1</sup> SLAC National Accelerator Laboratory (US)

<sup>2</sup> SLAC National Accelerator Laboratory

**Corresponding Authors:** lnakata@slac.stanford.edu, yangw@slac.stanford.edu, andrew.bohdan.hanushevsky@cern.ch, gche01@slac.stanford.edu

The Vera Rubin Observatory is a very ambitious project. Using the world's largest ground-based telescope, it will take two panoramic sweeps of the visible sky every three nights using a 3.2 Giga-pixel camera. The observation products will generate 15 PB of new data each year for 10 years. Accounting for reprocessing and related data products the total amount of critical data will reach several hundred PB. Because the camera consists of 201 CCD panels, the majority of the data products will consist of relatively small files in the low megabyte range, impacting data transfer performance. Yet, all of this data needs to be backed up in offline storage and still be easily retrievable not only for groups of files but also for individual files. This paper describes how SLAC is building a Rucio-centric specialized Tape Remote Storage Element (TRSE) that automatically creates a copy of a Rucio dataset as a single indexed file avoiding transferring many small files. This not only allows high-speed transfer of the data to tape for backup and dataset restoral, but also simple retrieval of individual dataset members in order to restore lost files. We describe the design and implementation of the TRSE and how it relates to current data management practices. We also present performance characteristics that make backups of extremely large scale data collections practical.

Parallel (Track 5) / 295

## Reconstruction of Full Decays using Transformers and Hyperbolic Embedding at Belle II

**Author:** Boyang Yu<sup>None</sup>

**Co-authors:** Nikolai Hartmann <sup>1</sup>; Thomas Kuhr <sup>1</sup>

<sup>1</sup> Ludwig Maximilians Universitat (DE)

**Corresponding Authors:** thomas.kuhr@cern.ch, nikolai.marcel.hartmann@cern.ch, boyang.yu@physik.uni-muenchen.de

In analyses conducted at Belle II, it is often beneficial to reconstruct the entire decay chain of both B mesons produced in an electron-positron collision event using the information gathered from detectors. The currently used reconstruction algorithm, starting from the final state particles, consists of multiple stages that necessitate manual configurations and suffers from low efficiency and a high number of wrongly reconstructed candidates.

Within this project, we are developing a software with the goal of automatically reconstructing B decays at Belle II with both high efficiency and accuracy. The trained models should be capable of accommodating rare decays with very small branching ratios, or even those that are unseen during the training phase.

To ensure optimal performance, the project is divided into the steps embedding of particles, particle reconstruction, and link prediction. Drawing inspiration from recent advancements in computer science, transformers and hyperbolic embedding are employed as fundamental components, with metric learning serving as the primary training technique.

Poster session / 296

## Integrated Research Infrastructure with JIRIAF

**Authors:** Christopher Larrieu<sup>None</sup>; David Lawrence<sup>None</sup>; Graham Heyes<sup>1</sup>; Jeng-Yuan Tsai<sup>2</sup>; Patrick Meagher<sup>2</sup>; Vardan Gyurjyan<sup>None</sup>

<sup>1</sup> *Jefferson Lab*

<sup>2</sup> *JLAB*

**Corresponding Authors:** pmeagher@jlab.org, larrieu@jlab.org, heyese@jlab.org, tsai@jlab.org, davidl@jlab.org, gurjyan@jlab.org

In this study, we introduce the JIRIAF (JLAB Integrated Research Infrastructure Across Facilities) system, an innovative prototype of an operational, flexible, and widely distributed computing cluster, leveraging readily available resources from Department of Energy (DOE) computing facilities. JIRIAF employs a customized Kubernetes orchestration system designed to integrate geographically dispersed resources into a unified, elastic distributed cluster. This system operates without the need for additional infrastructure investments by resource providers. Notably, JIRIAF has demonstrated a capability to process data streams at rates up to 100 Gbps, facilitating real-time data-stream processing across vast distances.

Furthermore, we developed a digital representation of workflows using a Bayesian probability graph model. This model utilizes a standard joint probability distribution to represent various probabilities associated with the digital state, including relevant quantities and potential rewards, all derived from observed actions and data. The determination of these quantities and rewards employs queueing theory, focusing on two critical metrics: the rate of workflow input and the processing rate. Our results confirm the efficacy of the JIRIAF digital twin in managing and orchestrating highly distributed workflows, showcasing its potential to significantly enhance computational resource utilization and process efficiency in complex environments.

Parallel (Track 1) / 297

## So FAIR, so good: the INFN strategy for Data Stewardship

**Authors:** Concezio Bozzi<sup>1</sup>; Daniele Bonacorsi<sup>2</sup>; Francesca Marchegiani<sup>None</sup>; Irene Piergentili<sup>3</sup>; Lorenzo Rinaldi<sup>4</sup>; Luca dell'Agnello<sup>3</sup>; Luciano Gaido<sup>None</sup>; Stefano Bianco<sup>5</sup>; Stefano Dal Pra<sup>3</sup>

<sup>1</sup> *INFN Ferrara*

<sup>2</sup> *University of Bologna / INFN*

<sup>3</sup> *INFN*

<sup>4</sup> *Universita e INFN, Bologna (IT)*

<sup>5</sup> *INFN e Laboratori Nazionali di Frascati (IT)*

**Corresponding Authors:** luciano.gaido@to.infn.it, daniele.bonacorsi@bo.infn.it, concezio.bozzi@cern.ch, irene.piergentili@lnf.infn.it, stefano.bianco@cern.ch, lorenzo.rinaldi@cern.ch, francesca.marchegiani@lngs.infn.it, luca.dellagnello@cern.ch, stefano.dalpra@cnaf.infn.it

Large international collaborations in the field of Nuclear and Subnuclear Physics have been leading the implementation of FAIR principles for managing research data. These principles are essential when dealing with large volumes of data over extended periods and involving scientists from multiple countries. Recently, smaller communities and individual experiments have also started adopting

these principles.

Many universities and research institutions are creating teams of Data Stewards in order to promote the implementation of FAIR principles (such as writing good Data Management Plans), particularly for smaller research units.

This contribution outlines the strategy adopted by the Italian National Institute for Nuclear Physics (INFN) to define the Data Steward profile and make its role suitable for the organization. This initiative is also linked to the Skills4EOSC project, which aims at establishing a network of competence centers for training European researchers and developing new professional roles for managing scientific data.

**Parallel (Track 2) / 298**

## **Design and developing status of the data acquisition and the slow control systems in the Mu2e experiment at Fermilab**

**Author:** Antonio Gioiosa<sup>1</sup>

<sup>1</sup> *University of Molise & INFN Roma Tor Vergata*

**Corresponding Author:** antonio.gioiosa@roma2.infn.it

The Mu2e experiment at Fermilab aims to observe coherent neutrinoless conversion of a muon to an electron in the field of an aluminum nucleus, with a sensitivity improvement of 10,000 times over current limits.

The Mu2e Trigger and Data Acquisition System (TDAQ) uses `otsdaq` framework as the online Data Acquisition System (DAQ) solution.

Developed at Fermilab, `otsdaq` integrates several framework components, an `artdaq` based DAQ, an `art` based event processing, and an EPICS-based detector control system (DCS), and provides a uniform multi-user interface to its components through a web browser.

The Mu2e tracker and calorimeter data streams are processed by a one-level software trigger implemented within the `art` framework.

Events accepted by the trigger have their data combined, post-trigger, with the separately read-out data from the Mu2e Cosmic Ray Veto system.

The Mu2e DCS is built on EPICS (Experimental Physics and Industrial Control System), an open-source platform for monitoring, controlling, alarming, and archiving.

A prototype of the TDAQ and DCS systems has been built and tested over the last three years at Fermilab's Feynman Computing Center. Now, the production system installation is underway.

This report covers our project's development progress, especially the web-based user interface, and slow control implementation.

**Parallel (Track 3) / 299**

## **Adoption of the alpaka performance portability library in the CMS software**

**Authors:** Andrea Bocci<sup>1</sup>; Matti Kortelainen<sup>2</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** matti.kortelainen@cern.ch, andrea.bocci@cern.ch

To achieve better computational efficiency and exploit a wider range of computing resources, the CMS software framework (CMSSW) has been extended to offload part of the physics reconstruction to NVIDIA GPUs. To support additional back-ends, as well to avoid the need to write, validate and

maintain a separate implementation of the reconstruction algorithms for each back-end, CMS has adopted the Alpaka performance portability library.

Alpaka (Abstraction Library for Parallel Kernel Acceleration) is a header-only C++ library that provides performance portability across different back-ends, abstracting the underlying levels of parallelism. It supports serial and parallel execution on CPUs, and extremely parallel execution on NVIDIA, AMD and Intel GPUs.

This contribution will show how Alpaka is used in the CMS software to develop and maintain a single code base; to use different toolchains to build the code for each supported back-end, and link them into a single application; to seamlessly select the best back-end at runtime, and implement portable reconstruction algorithms that run efficiently on CPUs and GPUs from different vendors. It will describe the validation and deployment of the Alpaka-based implementation in the CMS High Level Trigger, and highlight how it achieves near-native performance.

Parallel (Track 3) / 300

## Adapting to Change: A look at the evolution of ALICE's Quality Control framework

**Authors:** Barthelemy Von Haller<sup>1</sup>; Piotr Konopka<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** barthelemy.von.haller@cern.ch, piotr.jan.konopka@cern.ch

Since the mid-2010s, the ALICE experiment at CERN has seen significant changes in its software, especially with the introduction of the Online-Offline (O<sup>2</sup>) computing system during Long Shutdown 2. This evolution required continuous adaptation of the Quality Control (QC) framework responsible for online Data Quality Monitoring (DQM) and offline Quality Assurance (QA).

After a general overview of the system, this talk delves into the evolving user requirements that shaped the QC framework from its initial prototyping phase to its current state. We will explore the changing landscape of performance needs and feature demands, highlighting which initial requirements persisted, which emerged later, and which features ultimately proved unnecessary.

Additionally, we will trace the framework's development in relation to other software components within the ALICE ecosystem, offering valuable insights and lessons learned throughout the process. Finally, we will also discuss the challenges encountered in balancing development team resources with the evolving project scope.

Parallel (Track 2) / 301

## A graph neural network based cosmic muon trigger for the Mu3e experiment

**Authors:** Andre Schoening<sup>1</sup>; David Karres<sup>1</sup>; Sebastian Dittmeier<sup>2</sup>; Tamasi Kar<sup>1</sup>

<sup>1</sup> Heidelberg University (DE)

<sup>2</sup> Ruprecht-Karls-Universitaet Heidelberg (DE)

**Corresponding Authors:** tamasi.kar@cern.ch, karres@physi.uni-heidelberg.de, schoening@mail.desy.de, sebastian.dittmeier@cern.ch

The Mu3e experiment at the Paul-Scherrer-Institute will be searching for the charged lepton flavor violating decay  $\mu^+ \rightarrow e^+ e^- e^+$ . To reach its ultimate sensitivity to branching ratios in the order



of  $10^{-16}$ , an excellent momentum resolution for the reconstructed electrons is required, which in turn necessitates precise detector alignment. To compensate for weak modes in the main alignment strategy based on electrons and positrons from muon decays, the exploitation of cosmic ray muons is proposed.

The trajectories of cosmic ray muons are so different from the decays of stopped muons in the experiment that they cannot be reconstructed using the same method in the online filter farm. For this reason and in view of their comparatively rare occurrence, a special cosmic muon trigger is being developed. A study on the application of graph neural networks to classify events and to identify cosmic muon tracks will be presented.

Parallel (Track 3) / 302

## Reconstruction in Key4hep using Gaudi

**Authors:** Alvaro Tolosa-Delgado<sup>1</sup>; Andre Sailer<sup>1</sup>; Benedikt Hegner<sup>1</sup>; Briec Francois<sup>1</sup>; Frank-Dieter Gaede<sup>2</sup>; Gerardo Ganis<sup>1</sup>; Graeme A Stewart<sup>1</sup>; Jiaheng Zou<sup>3</sup>; Juan Miguel Carceller<sup>1</sup>; Juraj Smiesko<sup>1</sup>; Leonhard Reichenbach<sup>4</sup>; Mateusz Jakub Fila<sup>1</sup>; Sang Hyun Ko<sup>5</sup>; Swathi Sasikumar<sup>1</sup>; Tao Lin<sup>None</sup>; Teng LI<sup>6</sup>; Thomas Madlener<sup>7</sup>; Weidong Li<sup>3</sup>; Wenxing Fang<sup>None</sup>; Xiaomei Zhang<sup>8</sup>; Xingtao Huang<sup>None</sup>

<sup>1</sup> CERN

<sup>2</sup> Deutsches Elektronen-Synchrotron (DE)

<sup>3</sup> IHEP, Beijing

<sup>4</sup> University of Bonn (DE)

<sup>5</sup> Seoul National University (KR)

<sup>6</sup> Shandong University, CN

<sup>7</sup> Deutsches Elektronen-Synchrotron (DESY)

<sup>8</sup> Chinese Academy of Sciences (CN)

**Corresponding Authors:** thomas.madlener@cern.ch, liwd@ihep.ac.cn, swathi.sasikumar@cern.ch, frank-dieter.gaede@cern.ch, mateusz.jakub.fila@cern.ch, lintao@ihep.ac.cn, j.m.carceller@cern.ch, xiaomei.zhang@cern.ch, fangwx@ihep.ac.cn, sanghyun.ko@cern.ch, leonhard.reichenbach@cern.ch, zoujh@ihep.ac.cn, briec.francois@cern.ch, graeme.andrew.stewart@cern.ch, andre.philippe.sailer@cern.ch, huangxt@sdu.edu.cn, gerardo.ganis@cern.ch, benedikt.hegner@cern.ch, alvaro.tolosa.delgado@cern.ch, juraj.smiesko@cern.ch, tengli@sdu.edu.cn

Key4hep, a software framework and stack for future accelerators, integrates all the steps in the typical offline pipeline: generation, simulation, reconstruction and analysis. The different components of Key4hep use a common event data model, called EDM4hep. For reconstruction, Key4hep leverages Gaudi, a proven framework already in use by several experiments at the LHC, to orchestrate configuration and execution of reconstruction algorithms.

In this contribution, a brief overview of Gaudi is given. The specific developments built to make Gaudi work seamlessly with EDM4hep (and therefore in Key4hep) are explained, as well as other improvements requested by the Key4hep community. The list of developments includes a new IO service to run algorithms that read or write EDM4hep files in multithreading in a thread-safe way and a possibility to easily switch the EDM4hep I/O to the new ROOT RNTuple format for reading or writing. We show that both native (algorithms that use EDM4hep as input and output) and non-native algorithms from the ILC community can run together in Key4hep, picking up on knowledge and software developed over many years. A few examples of algorithms that have been created or ported to Key4hep recently are given, featuring the usage of Key4hep-specific features.

Parallel (Track6) / 303

## Building the Key4hep Software Stack with Spack

**Authors:** Andre Sailer<sup>1</sup>; Juan Miguel Carceller<sup>1</sup>; Thomas Madlener<sup>2</sup>; Wouter Deconinck<sup>None</sup>

<sup>1</sup> CERN

<sup>2</sup> Deutsches Elektronen-Synchrotron (DESY)

**Corresponding Authors:** thomas.madlener@cern.ch, wouter.deconinck@umanitoba.ca, j.m.carcell@cern.ch, andre.philippe.sailer@cern.ch

The Key4hep software stack enables studies for future collider projects. It provides a full software suite for doing event generation, detector simulation as well as reconstruction and analysis. In the Key4hep stack, over 500 packages are built using the spack package manager and deployed via the cvmfs software distribution system. In this contribution, we explain the current setup for building nightly builds and stable releases that are made every few months or as needed. These builds are made available to users, who have access to a full and consistent software stack via a simple setup script. Different operating systems and compilers are supported and some utilities are provided to make development on top of the Key4hep builds easier. Both the benefits of the community-driven approach followed in spack and the issues found along the way are discussed.

Poster session / 304

## Particle Identification at LHCb combining Neural Networks and evolutionary computation

**Author:** Sergi Bernet Andres<sup>1</sup>

**Co-authors:** Alvaro Garcia Piquer<sup>2</sup>; Miriam Calvo Gomez<sup>1</sup>; Xavier Vilasis Cardona<sup>1</sup>

<sup>1</sup> La Salle, Ramon Llull University (ES)

<sup>2</sup> La Salle

**Corresponding Authors:** sergi.bernet.andres@cern.ch, alvaro.garcia@salle.url.edu, xavier.vilasis.cardona@cern.ch, miriam.calvo@cern.ch

Particle identification (PID) in the LHCb experiment is performed by combining information from several subdetectors of the experiment, including RICH, calorimeter and muon systems. Information is acquired and processed in real time at the trigger level, where it is combined and used for particle identification. LHCb employs 2 methods for particle identification: a global likelihood approach which is the traditional method and a neural network based approach called probNN. For LHC Run 3, subdetectors has been upgraded to operate under higher luminosity conditions compared to Run 2. These differences require an update of many algorithms, including the probNN method. To address these updated and take advantage of it, a depth multivariate analysis combined with evolutionary algorithms is developed to optimize neural networks and variable selection. As result, architectures have been reduced by up to 80% of their original size, which directly improves model inference by a factor 4 compared to its predecessor used during Run 2 with an equivalent PID performance.

Poster session / 305

## Experience with ARM WNs at the WLCG Tier1 GridKa

**Author:** Armin Alfredo Krull<sup>None</sup>

**Co-authors:** Andreas Petzold<sup>1</sup>; Matthias Jochen Schnepf; Max Fischer<sup>2</sup>

<sup>1</sup> KIT - Karlsruhe Institute of Technology (DE)

<sup>2</sup> Karlsruhe Institute of Technology

**Corresponding Authors:** max.fischer@kit.edu, matthias.jochen.schnepf@cern.ch, ulmze@student.kit.edu, andreas.petzold@cern.ch

Computing Centers always look for new server systems that can reduce operational costs, especially power consumption, and provide higher performance.

ARM-CPU promise higher energy efficiency than x86-CPU.

Therefore, the WLCG Tier1 center GridKa will partially use worker nodes with ARM-CPU and has already carried out various power consumption and performance tests based on the HEPscore23 benchmark.

Various system settings, such as maximum CPU frequency, were studied to determine the best performance and highest energy efficiency of the ARM-CPU systems.

GridKa will provide the HEP community with several ARM-CPU worker nodes in their batch farm. We present the results of these benchmarks on systems with ARM-CPU compared to benchmarks of current x86-CPU worker nodes at GridKa and the status of provisioning ARM-CPU worker nodes to the community.

Parallel (Track 7) / 306

## Prévessin Data Centre Powers Up

**Authors:** Joel Murray Davies<sup>1</sup>; Max Dupuis<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** joel.murray.davies@cern.ch, max.dupuis@cern.ch

CERN's state-of-the-art Prévessin Data Centre (PDC) is now operational, complementing CERN's Meyrin Data Centre Tier-0 facility to provide additional and sustainable computing power to meet the needs of High-Luminosity LHC in 2029 (expected to be ten times greater than today). In 2019, it was decided to tender the design and construction of a new, modern, energy-efficient (PUE of  $\leq 1.15$ ) Data Centre with a total of 12 MW IT capacity spread across six IT rooms. As it stands, two out of six IT rooms are production ready with a combined 4MW of IT capacity, with the remaining to be commissioned in two phases over the next ten years.

To begin, we will guide you through the commissioning of the Data Centre, with explanations of the various steps taken to equip the IT rooms. We will outline the acceptance process and the comprehensive trial operation tests which ensured a smooth transition into O&M (Operations and Maintenance) mode. O&M will be handled by Service Provider, Equans, in a collaborative partnership with the CERN IT department. This approach is the first of its kind at CERN and in this talk we'll delve into how the contract was established.

We will finish by providing an overview of our progress in the first operational year with a look forward to scalable growth through the phased deployment of the remaining four IT rooms that will meet the anticipated need for physics computing into Run4.

Poster session / 307

## The End of an Era: Termination of the Meyrin Data Centre Console Service

**Author:** Joel Murray Davies<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** joel.murray.davies@cern.ch

For nearly five decades, Data Centre Operators have provided critical support to the CERN Meyrin Data Centre, from its infancy, until spring 2024. However, advancements in Data Centre technology and resilience built into IT services have rendered the Console Service obsolete.

In the early days of the Meyrin Data Centre, day to day operations relied heavily on the expertise and manual operations of Console Operators. A crucial function of the Console operator was the manual management of the massive tape libraries used for data storage. More recently, operators have been ensuring the uptime of Data Centre equipment, responding to local power incidents, and coordinating response to incidents in the data centre.

In this talk we'll cover the principles that underpin the decision for terminating the service, a view into the gradual wind down of the service and the transfer of critical functions to operations teams at CERN.

**Poster session / 308**

## Automated Template Testing and Management with GitLab Continuous Integration for AMS Offline Computing

**Authors:** Baosong Shan<sup>1</sup>; Deyuan Ma<sup>2</sup>

**Co-authors:** Vitali Choutko<sup>3</sup>; Alexander Egorov<sup>3</sup>; Alexandre Eline<sup>3</sup>

<sup>1</sup> *Beihang University (CN)*

<sup>2</sup> *Shandong Institute of Advanced Technology (SDIAT) (CN)*

<sup>3</sup> *Massachusetts Inst. of Technology (US)*

**Corresponding Authors:** alexandre.eline@cern.ch, baosong.shan@cern.ch, alexander.egorov@cern.ch, deyuan.ma@cern.ch, vitali.choutko@cern.ch

The Alpha Magnetic Spectrometer (AMS) is a particle physics experiment installed and operating aboard the International Space Station (ISS) from May 2011 and expected to last through 2030 and beyond. Data reconstruction and Monte-Carlo simulation are two major production activities in AMS offline computing, and templates are defined as a collection of data cards to describe different reconstruction and simulation tasks and to provide necessary input parameters. This paper presents how we use the Continuous Integration mechanism in GitLab to better manage the production datasets and templates, including syntax checking, functionality testing, performance testing, and the integration with existing production statistics monitoring system. The system also uses the pipeline schedules to periodically check the completing status of the production tasks and send warning messages to the administrators if the production progress is not as expected.

**Parallel (Track 1) / 309**

## Recent Experience with the CMS Data Management System

**Authors:** Andres Manrique Ardila<sup>1</sup>; Andrew Wightman<sup>2</sup>; Christos Emmanouil<sup>3</sup>; Dmytro Kovalskyi<sup>4</sup>; Eric Vaandering<sup>5</sup>; Hasan Ozturk<sup>3</sup>; Panos Paparrigopoulos<sup>3</sup>; Rahul Chauhan<sup>3</sup>

<sup>1</sup> *University of Wisconsin Madison (US)*

<sup>2</sup> *University of Nebraska Lincoln (US)*

<sup>3</sup> *CERN*

<sup>4</sup> *Massachusetts Inst. of Technology (US)*

<sup>5</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** dmytro.kovalskyi@cern.ch, andrew.steven.wightman@cern.ch, h.ozturk@cern.ch, rahul.chauhan@cern.ch, panos.paparrigopoulos@cern.ch, ewv@fnal.gov, andres.manrique.ardila@cern.ch

The CMS experiment manages a large-scale data infrastructure, currently handling over 200 PB of disk and 500 PB of tape storage and transferring more than 1 PB of data per day on average between various WLCG sites. Utilizing Rucio for high-level data management, FTS for data transfers,

and a variety of storage and network technologies at the sites, CMS confronts inevitable challenges due to the system's growing scale and evolving nature. Key challenges include managing transfer and storage failures, optimizing data distribution across different storages based on production and analysis needs, implementing necessary technology upgrades and migrations, and efficiently handling user requests. The data management team has established comprehensive monitoring to supervise this system and has successfully addressed many of these challenges. The team's efforts aim to ensure data availability and protection, minimize failures and manual interventions, maximize transfer throughput and resource utilization, and provide reliable user support. This paper details the operational experience of CMS with its data management system in recent years, focusing on the encountered challenges, the effective strategies employed to overcome them and the ongoing challenges as we prepare for future demands.

**Parallel (Track 1) / 310**

## The Fermilab experience transitioning to CERN Tape Archive (CTA)

**Authors:** Dan Szkola<sup>1</sup>; David Alexander Mason<sup>1</sup>; Dennis Box<sup>None</sup>; Dmitry Litvintsev<sup>1</sup>; Eric Vaandering<sup>1</sup>; Jeff Derbyshire<sup>None</sup>; Leo Muñoz<sup>2</sup>; Lucas Trestka<sup>2</sup>; Rafael Arturo Rocha Vidaurri<sup>None</sup>; Ren Bauer<sup>1</sup>; Robert Illingworth<sup>1</sup>; Scarlet Norberg<sup>None</sup>; Tammy Walton<sup>None</sup>; Terrance Jones<sup>3</sup>; Tim Messer<sup>1</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

<sup>2</sup> *Fermilab*

<sup>3</sup> *FL/CD/SCD/SDS/SSA*

**Corresponding Authors:** [dbox@fnal.gov](mailto:dbox@fnal.gov), [victormf@fnal.gov](mailto:victormf@fnal.gov), [norberg@ou.edu](mailto:norberg@ou.edu), [dmason@fnal.gov](mailto:dmason@fnal.gov), [ewv@fnal.gov](mailto:ewv@fnal.gov), [litvinse@fnal.gov](mailto:litvinse@fnal.gov), [twalton@fnal.gov](mailto:twalton@fnal.gov), [jonest@fnal.gov](mailto:jonest@fnal.gov), [jeffderb@fnal.gov](mailto:jeffderb@fnal.gov), [dszkola@fnal.gov](mailto:dszkola@fnal.gov), [rrochavi@fnal.gov](mailto:rrochavi@fnal.gov), [renbauer@fnal.gov](mailto:renbauer@fnal.gov), [ltrestka@fnal.gov](mailto:ltrestka@fnal.gov), [timothy.alan.messer@cern.ch](mailto:timothy.alan.messer@cern.ch), [illingwo@fnal.gov](mailto:illingwo@fnal.gov)

During the next year Fermilab will fully transition to CTA (CERN Tape Archive) for tape storage management from its in-house legacy software, Enstore.

We will describe the testing and evaluation process including a test replicating 10% of the peak bandwidth of the current system which required less than 10% of the current resources with CTA. This test was used to verify performance and make a choice between EOS and dCache as the front end file system.

Fermilab has special considerations, which a replacement system needs to address, including maintaining good performance with smaller files (under 10 MB), reading existing archives of small files, and reading legacy formatted tapes. We will detail the developments needed for reading the legacy tapes and small file archives. We will also present a plan for the transition of the remaining Enstore installations to CTA.

**Parallel (Track 2) / 311**

## ALICE Event Display - from the legacy ROOT-based visualization to the web-based application

**Author:** Julian Myrcha<sup>1</sup>

<sup>1</sup> *Warsaw University of Technology (PL)*

**Corresponding Author:** [julian.wojciech.myrcha@cern.ch](mailto:julian.wojciech.myrcha@cern.ch)

The architecture of the existing ALICE Run 3 on-line real time visualization solution was designed for easy modification of the visualization method used. In addition to the existing visualization based

on the desktop application, a version using browser-based visualization has been prepared. In this case, the visualization is computed and displayed on the user's computer. There is no need to install any software on the user's computer. The overall visualization architecture used allows for a smooth switch to a new version of visualization: for a transition period both solutions (traditional desktop and web) can be used simultaneously.

ALICE visualization requires loading information about the displayed tracks (which may be several dozen thousand). This type of visualizations differs from visualizations typically used in computer graphics, where high efficiency of motion representation is achieved by modifying transformations describing the motion of already loaded models. In event visualization, the description of the tracks (models) changes with each view. Achieving high display performance requires the use of a number of optimization solutions.

The data downloaded by the web application is already pre-processed and prepared to be loaded to the graphics card, thanks to which the calculations in the browser are significantly simplified and the performance of the browser visualization is comparable to the visualization in the desktop application.

When creating a new visualization, a component approach to building a web application was used: individual components are responsible for various functions (e.g. data retrieval, different visualizations, interaction with the user). This construction of blocks allows for easy rearrangement by replacing or adding new components. The solution testing process is therefore significantly simplified because each component can be tested independently.

**Parallel (Track 5) / 312**

## **ROOT RNTuple implementation in Julia programming language**

**Authors:** Jerry Ling<sup>1</sup>; Tamas Gal<sup>2</sup>

<sup>1</sup> *Harvard University (US)*

<sup>2</sup> *ECAP, FAU Erlangen-Nürnberg*

**Corresponding Authors:** [tamas.gal@fau.de](mailto:tamas.gal@fau.de), [jerry.ling@cern.ch](mailto:jerry.ling@cern.ch)

At the LHC experiments, RNTuple is emerging as the primary data storage solution, and will be ready for production next year. In this context, we introduce the latest development in UnROOT.jl, a high-performance and thread-safe Julia ROOT I/O package that facilitates both the reading and writing of RNTuple data.

We briefly share insights gained from implementing RNTuple Reader twice: first in Python, and then in Julia. We discuss the composability of the RNTuple type system and demonstrate how Julia's multiple dispatch feature has been effectively employed to realize this concisely.

Regarding the implementation of RNTuple Writer, we outline the current capabilities and illustrate how they support end-user analyses. Furthermore, we present a roadmap for future development aimed at achieving seamless data I/O interoperability across various programming languages and libraries, including C++, Python, and Julia.

Lastly, we showcase the capabilities and performance of our Julia implementation with real examples. We highlight how our solution facilitates interactive analysis for end-users utilizing RNTuple.

**Parallel (Track 5) / 313**

## **GIL-free scaling of Uproot with Python 3.13 subinterpreters**

**Author:** Jim Pivarski<sup>1</sup>

<sup>1</sup> *Princeton University*

**Corresponding Author:** pivarski@princeton.edu

Uproot is a Python library for ROOT I/O that uses NumPy and Awkward Array to represent and perform computations on bulk data. However, Uproot uses pure Python to navigate through ROOT's data structures to *find* the bulk data, which can be a performance issue in metadata-intensive I/O: (a) many small files, (b) many small TBaskets, and/or (c) low compression overhead. Worse, these performance issues can't be alleviated by multithreading because Python imposes a thread-lock between each instruction on its virtual machine, infamously known as the Global Interpreter Lock (GIL).

Python 3.13, released this month, introduces a fundamental new feature: a single Python process can run multiple interpreters, each in its own thread, each with its own (thread-local!) GIL. Subinterpreters are an intermediate choice between share-everything threads and share-nothing processes. Subinterpreters can only share Python objects through FIFO Queues (or, equivalently, Channels), and not by reference. However, they can freely operate on shared array data. Similar solutions can be cobbled together with multiple Python processes, using multiprocessing.Queue and multiprocessing.SharedMemory, but these rely on POSIX pipes and shared memory, depend on ulimit settings, and are much slower than subinterpreter communication.

In this talk, I'll show how Uproot takes advantage of subinterpreters to improve scaling for metadata-intensive I/O.

**Parallel (Track 3) / 314**

## Offline data processing software for the High Energy cosmic-Radiation Detection facility

**Author:** qianqian Shi<sup>None</sup>

**Co-authors:** Xingtao Huang ; Teng LI <sup>1</sup>

<sup>1</sup> *Shandong University, CN*

**Corresponding Authors:** shiqq@mail.sdu.edu.cn, huangxt@sdu.edu.cn, tengli@sdu.edu.cn

The High Energy cosmic-Radiation Detection facility (HERD) is a scientific instrument planned for deployment on the Chinese Space Station, aimed at indirectly detecting dark matter and conducting gamma-ray astronomical research. HERD Offline Software (HERDOS) is developed for the HERD offline data processing, including Monte Carlo simulation, calibration, reconstruction and physics analysis tasks. HERDOS is developed based on SNI<sub>PER</sub>, a lightweight framework designed for HEP experiments, as well as a few state-of-the-art software in the HEP community, such as Detector Description Toolkit (DD4hep), the plain-old-data I/O (podio) and Intel Thread Building Blocks (TBB), etc.

This contribution will provide an overview of the design and implementation details of HERDOS, and in particular, the following details will be addressed:

1. The design of the Event Data Model (EDM) based on Podio, and the implementation of data management system (DMS) through the integration of Podio and SNI<sub>PER</sub>.
2. The parallelized DMS based on SNI<sub>PER</sub> and TBB, specifically the development of GlobalStore based on the Podio to enable concurrent data access and data I/O.
3. The parallelized detector simulation based on MT-SNI<sub>PER</sub>, including both event-level and track-level parallelism.
4. The geometry management system based on DD4hep that provides consistent detector description, an easy-to-use interface to retrieve detector description information.

At present, HERDOS is operating effectively to support the design of the detector, as well as the exploration of its physics potential.

Parallel (Track 3) / 315

## Implementation of an Open Drift Chamber in ACTS

**Authors:** Xiaocong Ai<sup>1</sup>; Xingtao Huang<sup>None</sup>; Hao Li<sup>None</sup>; Teng Li<sup>2</sup>; Weidong Li<sup>3</sup>; Tao Lin<sup>None</sup>; Xiaoshuai Qin<sup>4</sup>

<sup>1</sup> Zhengzhou University

<sup>2</sup> Shandong University, CN

<sup>3</sup> IHEP, Beijing

<sup>4</sup> Shandong University, Qingdao

**Corresponding Authors:** qinxs@sdu.edu.cn, huangxt@sdu.edu.cn, lihao097430@gs.zzu.edu.cn, xiaocong.ai@cern.ch, tengli@sdu.edu.cn, liwd@ihep.ac.cn, lintao@ihep.ac.cn

A Common Tracking Software (ACTS) is an open-source experiment-independent and framework-independent tracking software for both current and future particle and nuclear physics experiments. It provides a set of high-level detector-agnostic track reconstruction tools, which are initially developed and validated with a few example detectors, e.g. the Open Data Detector (ODD), which is an open-access HL-LHC style detector for algorithmic development and benchmarking, and has been well integrated and supported within the ACTS toolkit. Current implementation of the ODD includes a full silicon-based tracking system and calorimetry.

So far, ACTS has already been deployed for data production at experiments such as ATLAS, sPHENIX, FASER etc., where the applications are focusing on silicon-based trackers. Recently, ACTS has been successfully extended to gaseous trackers and the new development is being validated with uRWell-based tracker and drift chambers for future HEP experiments such as Super Tau Charm Facility (STCF) and Circular Electron Positron Collider (CEPC).

In this contribution, we will introduce the progress we have made in implementing drift chamber as a sub-detector of the ODD. With the newly added drift chamber, the ODD can be configured either as a full silicon tracker or a combined tracker consisting of a pixel tracker and a drift chamber. The tracking performance of ODD with the two detector configurations are studied using ACTS. Our experience with ODD and ACTS hopefully will be beneficial to the potential ACTS clients who are interested to evaluate the performance of ACTS algorithms.

Parallel (Track 2) / 316

## Implementation and development of a DAQ system DELILA at ELI-NP

**Author:** Sohichiroh Aogaki<sup>1</sup>

**Co-authors:** Dimiter Loukanov Balabanski<sup>1</sup>; Mihai Cuciuc<sup>1</sup>; Radu Corbu<sup>2</sup>; Stefan Niculae<sup>3</sup>; Vlad Andrei Toma<sup>4</sup>

<sup>1</sup> Extreme Light Infrastructure-Nuclear Physics (ELI-NP)/Horia Hulubei National Institute for Physics and Nuclear Engineering (IFIN-HH), Str. Reactorului 30, Bucharest-Măgurele 077125, Romania

<sup>2</sup> Extreme Light Infrastructure-Nuclear Physics (ELI-NP)/Horia Hulubei National Institute for Physics and Nuclear Engineering (IFIN-HH), Str. Reactorului 30, Bucharest-Măgurele 077125, Romania. Universitatea Națională de Știință și Tehnologie Politehnica București (UNSTPB), Splaiul Independentei no. 313, sector 6, Bucharest, ROMANIA

<sup>3</sup> Extreme Light Infrastructure-Nuclear Physics (ELI-NP)/Horia Hulubei National Institute for Physics and Nuclear Engineering (IFIN-HH), Str. Reactorului 30, Bucharest-Măgurele 077125, Romania. Faculty of Mathematics and Computer Science, University of Bucharest, 90, Panduri Street, Sector 5, Bucharest, Romania

<sup>4</sup> Extreme Light Infrastructure-Nuclear Physics (ELI-NP)/Horia Hulubei National Institute for Physics and Nuclear Engineering (IFIN-HH), Str. Reactorului 30, Bucharest-Măgurele 077125, Romania. Faculty of Physics, University of Bucharest, Atomistilor 405, 077125 Bucharest-Măgurele, Romania



**Corresponding Authors:** sohichiroh.aogaki@eli-np.ro, vlad.toma@eli-np.ro, stefan.niculae@eli-np.ro, dimiter.balabanski@eli-np.ro, mihai.cuciuc@eli-np.ro, radu.corbu@eli-np.ro

Digital ELI-NP List-mode Acquisition (DELILA) is a data acquisition (DAQ) system for the Variable Energy GAMMA (VEGA) beamline system at Extreme Light Infrastructure –Nuclear Physics (ELI-NP), Magurele, Romania 1. ELI-NP has been implementing the VEGA beamline and entirely operate the beamline in 2026. Several different detectors/experiments (e.g. High Purity Ge (HPGe) detectors, Si detectors and scintillator detectors) will be placed at the VEGA beam line and read out by CAEN digitizers, Mesytec ADC and TDC, and some other electronics 2. DELILA has been developed using mainly DAQ-Middleware and CAEN digitizer libraries to fit the experiments and the read-out electronics 3. The main requirements are network transparency and synchronized time stamps. DAQ-Middleware allows us to fetch data from different electronics and computers to a data merger via Ethernet.

DELILA uses two databases to record experimental information: MongoDB for the run information and InfluxDB for event rates. DELILA uses ROOT libraries for online monitoring and recording experiment data.

The DAQ system has been used for several experiments at IFIN-HH 9MV and 3MV tandem beamlines in Romania [4]. The presenter will present the implementation and results of DELILA.

1 S. Gales, K.A. Tanaka et al., Rep. Prog. Phys. 81 094301 (2018)

2 N.V. Zamfir et al., Romanian Reports in Physics 68 Supplement, S3–S945 (2016)

3 Y. Yasu et al., J. Phys.: Conf. Ser. 219 022025 (2010)

[4] S.Aogaki et al. Nuclear Inst. and Methods in Physics Research, A 1056 (2023) 168628

**Parallel (Track 1) / 317**

## Carbon costs of storage: a UK perspective.

**Authors:** Alison Packer<sup>1</sup>; Alison Packer<sup>None</sup>

**Co-author:** Samuel Cadellin Skipsey

<sup>1</sup> STFC - Science & Technology Facilities Council (GB)

**Corresponding Authors:** alison.packer@cern.ch, samuel.cadellin.skipsey@cern.ch, alison.packer@stfc.ac.uk

In order to achieve the higher performance year on year required by the 2030s for future LHC upgrades at a sustainable carbon cost

to the environment, it is essential to start with accurate measurements of the state of play. Whilst there have been a number of studies

of the carbon cost of compute for WLCG workloads published, rather less has been said on the topic of storage, both nearline and archival.

We present a study of the embedded and ongoing carbon costs of storage in multiple configurations, from Tape farms through to SSDs, within the UK Tier-1 and Tier-2s and discuss how this directs future policy.

**Poster session / 318**

## Implementing a Domain-Specific Language (DSL) for Physics Data Analysis in RootInteractive

**Authors:** Marian I Ivanov<sup>1</sup>; Marian Ivanov<sup>2</sup>

<sup>1</sup> GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)

<sup>2</sup> Comenius University (SK)

**Corresponding Authors:** marian.i@cern.ch, marian.ivanov@cern.ch

We propose a Domain-Specific Language (DSL) embedded within RootInteractive specifically designed for physics data analysis, particularly within the ALICE experiment at CERN. This DSL addresses the need for concise and expressive code, allowing physicists to efficiently manipulate and visualize data reflecting the unique computational patterns encountered in physics research.

Our DSL offers functionalities tailored for physics data analysis, including querying reconstruction and calibration data, performing data joins, and computing rolling statistics. A functional prototype demonstrates significant efficiency gains in performing joins and calculating rolling statistics.

Future development plans include integration with RDataFrame workflows, enabling the handling of diverse data containers and facilitating complex queries on multidimensional data and time series analysis. Additionally, we plan to explore transitioning to awkward arrays for Python analysis and investigate hybrid approaches for flexible data manipulation.

Furthermore, our DSL simplifies the handling of relationships between elements like tracks, V0s, cascades, and collisions. This enhances the versatility and depth of data analysis within the ALICE experiment framework.

In conclusion, our DSL represents a significant advancement in physics data analysis. It empowers physicists with a powerful and intuitive tool for efficient and accurate analysis of complex datasets within the ALICE experiment.

**Parallel (Track6) / 319**

## **Evolution of the ATLAS TRT Gas Gain Stabilization System Software**

**Author:** Bartosz Mindur<sup>1</sup>

<sup>1</sup> *AGH University of Krakow (PL)*

**Corresponding Author:** mindur@agh.edu.pl

I will be presenting the history of the design, implementation, testing, and release of the production version of a C++-based software for the Gas Gain Stabilization System (GGSS) used in the TRT detector at the ATLAS experiment. This system operates 24/7 in the CERN Point1 environment under the control of the Detector Control System (DCS) and plays a crucial role in delivering reliable data during the LHC's stable beams.

The uniqueness of this software lies in its initial release around 2004, followed by subsequent refactoring, improvements, and implementation for the Run1 period of the LHC in 2008. Another significant change occurred during Long Shutdown 1 when the operating system transitioned from Windows to Linux for Run2 in 2015. More recently, there have been frequent updates and upgrades to the operating system and external libraries.

My aim is to present the evolution of the software, highlighting changes introduced from an external perspective due to shifts in the environment or requirements. Additionally, I'll discuss the evolution of the C++ standard, compiler changes, security considerations, and modifications to the build and test environment. During the conference, I will focus on the most compelling and significant milestones, as well as key aspects relevant to the lifecycle of this software.

**Parallel (Track 5) / 320**

## **Advancements in Computing and Simulation Techniques for the HIBEAM-NNBAR Experiment**

**Author:** Bernhard Meirose<sup>1</sup>

<sup>1</sup> *Chalmers University of Technology + Lund University*

**Corresponding Author:** bernhard.meirose@hep.lu.se

The HIBEAM-NNBAR experiment at the European Spallation Source is a multidisciplinary two-stage program of experiments that includes high-sensitivity searches for neutron oscillations, searches for sterile neutrons, searches for axions, as well as the search for exotic decays of the neutron. The computing framework of the collaboration includes diverse software, from particle generators to Monte Carlo transport codes, which are uniquely interfaced together. Significant advances have been made in computing and simulation for HIBEAM-NNBAR, particularly with machine learning applications and with the introduction of fast parametric simulations in Geant4. A summary of the simulation steps of the experiment, including beamline, cosmic veto system, as well as detector simulations and estimation of the background processes, will be presented.

**Poster session / 321**

## AI-based approach for provider selection in the INDIGO PaaS Orchestration System of INFN Cloud

**Author:** Luca Giommi<sup>1</sup>

**Co-authors:** Alessandro Costantini <sup>2</sup>; Francesco Debiase <sup>3</sup>; Giacinto Donvito <sup>4</sup>; Gioacchino Vino <sup>5</sup>; Giovanni Savarese <sup>6</sup>; Marica Antonacci <sup>3</sup>

<sup>1</sup> *INFN CNAF*

<sup>2</sup> *INFN-CNAF*

<sup>3</sup> *INFN*

<sup>4</sup> *INFN-Bari*

<sup>5</sup> *INFN Bari (IT)*

<sup>6</sup> *INFN-BA*

**Corresponding Authors:** giacinto.donvito@ba.infn.it, marica.antonacci@ba.infn.it, giovanni.savarese@ba.infn.it, francesco.debiase@ba.infn.it, alessandro.costantini@cnafe.infn.it, luca.giommi@cern.ch, gioacchino.vino@cern.ch

The Italian National Institute for Nuclear Physics (INFN) has recently launched the INFN Cloud initiative, aimed at providing a federated Cloud infrastructure and a dynamic portfolio of services to scientific communities supported by the Institute. The federative middleware of INFN Cloud is based on the INDIGO PaaS orchestration system, consisting of interconnected open-source microservices. Among these, the INDIGO PaaS Orchestrator receives high-level deployment requests in the form of TOSCA templates and coordinates the process of creating deployments on the IaaS platforms made available by the federated providers.

Through internal projects like INFN DataCloud and European initiatives such as interTwin and AI4EOSC, INFN is working to improve the orchestration system by integrating artificial intelligence to optimize deployment scheduling. This contribution outlines the preparatory work to identify the key features and their sources (e.g., databases, logs, monitoring tools), followed by the data preprocessing necessary for in-depth analysis of different AI techniques. The first implemented approach involves the design of two models: one for the deployment success/failure classification and another for the deployment time regression. The combination of the output of the two models trained on recent and sliding time windows aims to define the ordered list of providers that the orchestrator can use for deployment submission.

An alternative solution for an AI-based Orchestrator could involve a Reinforcement Learning approach, in which an agent is trained as if it has to win a game and learns which provider is best suited to user demand.

Parallel (Track 2) / 322

## Tracking and vertexing downstream the LHCb magnet at the first stage of the trigger

**Authors:** Arantza De Oyanguren Campos<sup>1</sup>; Brij Kishor Jashal<sup>2</sup>; Jiahui Zhuo<sup>1</sup>; Valerii Kholoimov<sup>3</sup>; Volodymyr Svintozelskyi<sup>1</sup>

<sup>1</sup> *Univ. of Valencia and CSIC (ES)*

<sup>2</sup> *RAL, TIFR and IFIC*

<sup>3</sup> *Instituto de Física Corpuscular (Univ. of Valencia)*

**Corresponding Authors:** brij.kishor.jashal@cern.ch, volodymyr.svintozelskyi@cern.ch, arantza.de.oyanguren.campos@cern.ch, jiahui.zhuo@cern.ch, valerii.kholoimov@cern.ch

A new algorithm, called “Downstream”, has been developed and implemented at LHCb, which is able to reconstruct and select very displaced vertices in real time at the first level of the trigger (HLT1). It makes use of the Upstream Tracker (UT) and the Scintillator Fiber detector (SciFi) of LHCb and it is executed on GPUs inside the Allen framework. In addition to an optimized strategy, it utilizes a Neural Network (NN) implementation to increase the track efficiency and reduce the ghost rates, with very high throughput and limited time budget. Besides serving to reconstruct Ks and Lambda vertices to calibrate and align the detectors, the Downstream algorithm and the associated two-track vertexing will largely increase the LHCb physics potential for detecting long-lived particles during the Run3.

Poster session / 323

## A Streamlined Neural Model for Real-Time Analysis at the First Level of the LHCb Trigger

**Authors:** Arantza De Oyanguren Campos<sup>1</sup>; Brij Kishor Jashal<sup>2</sup>; Jiahui Zhuo<sup>1</sup>; Valerii Kholoimov<sup>3</sup>; Volodymyr Svintozelskyi<sup>1</sup>

<sup>1</sup> *Univ. of Valencia and CSIC (ES)*

<sup>2</sup> *RAL, TIFR and IFIC*

<sup>3</sup> *Instituto de Física Corpuscular (Univ. of Valencia)*

**Corresponding Authors:** jiahui.zhuo@cern.ch, arantza.de.oyanguren.campos@cern.ch, brij.kishor.jashal@cern.ch, valerii.kholoimov@cern.ch, volodymyr.svintozelskyi@cern.ch

One of the most significant challenges in tracking reconstruction is the reduction of “ghost tracks,” which are composed of false hit combinations in the detectors. When tracking reconstruction is performed in real-time at 30 MHz, it introduces the difficulty of meeting high efficiency and throughput requirements. A single-layer feed-forward neural network (NN) has been developed and trained to address this challenge. The simplicity of the NN allows for parallel evaluation of many track candidates to filter ghost tracks using CUDA within the Allen framework. This capability enables us to run this type of NN at the first level of the trigger (HLT1) in the LHCb experiment. This neural network approach is already utilized in several HLT1 algorithms and is becoming an essential tool for Run 3. Details of the implementation and performance of this strategy will be presented in this talk.

Poster session / 324

## Six months of Single Sign On on Kubernetes

**Author:** Antonio Nappi<sup>1</sup>

**Co-authors:** Paul Van Uytvinck<sup>1</sup>; Sebastian Lopienski<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** paul.van.uytvinck@cern.ch, antonio.nappi@cern.ch, sebastian.lopienski@cern.ch

The CERN Single Sign On (SSO) hosting infrastructure underwent a major reconstruction in 2023 in an effort to increase service reliability and operational efficiency.

This session will outline how the Cloud Native Computing Foundation (CNCf) tools facilitate that, with particular attention to the key decisions, difficulties, and architectural concerns for this critical IT service

Parallel (Track 2) / 325

## BuSca: a Buffer Scanner at 30 MHz for New Long-Lived Particle Searches at LHCb

**Authors:** Arantza De Oyanguren Campos<sup>1</sup>; Brij Kishor Jashal<sup>2</sup>; Jiahui Zhuo<sup>1</sup>; Valerii Kholoimov<sup>3</sup>; Volodymyr Svintozelskyi<sup>1</sup>

<sup>1</sup> Univ. of Valencia and CSIC (ES)

<sup>2</sup> RAL, TIFR and IFIC

<sup>3</sup> Instituto de Física Corpuscular (Univ. of Valencia)

**Corresponding Authors:** jiahui.zhuo@cern.ch, arantza.de.oyanguren.campos@cern.ch, brij.kishor.jashal@cern.ch, volodymyr.svintozelskyi@cern.ch, valerii.kholoimov@cern.ch

In this presentation, we introduce BuSca, a prototype algorithm designed for real-time particle searches, leveraging the enhanced parallelization capabilities of the new LHCb trigger scheme implemented on GPUs. BuSca is focused on downstream reconstructed tracks, detected exclusively by the UT and SciFi detectors. By projecting physics candidates onto 2D histograms of flight distance and mass hypotheses at a remarkable 30 MHz rate, BuSca identifies hot spots indicative of potential candidates of new particles, thereby providing strategic guidance for the development of new trigger lines. Additionally, BuSca offers an Armenteros-Podolanski representation, providing insights into the mass hypotheses of the decay products associated with the new particle. The performance of BuSca, including the outcomes of its initial prototype on simulated data, will be elucidated in this talk.

Parallel (Track 7) / 326

## Heterogeneous Computing and Power Efficiency in HEP

**Author:** Emanuele Simili<sup>1</sup>

**Co-authors:** Albert Gyorgy Borbely<sup>1</sup>; David Britton<sup>1</sup>; Gordon Stewart; Samuel Cadellin Skipsey

<sup>1</sup> University of Glasgow (GB)

**Corresponding Authors:** gordon.stewart@glasgow.ac.uk, samuel.cadellin.skipsey@cern.ch, albert.borbely@cern.ch, david.britton@cern.ch, emanuele.simili@glasgow.ac.uk

The Glasgow ScotGrid facility is now a truly heterogeneous site, with over 4k ARM cores representing 20% of our compute nodes, which has enabled large-scale testing by the experiments and more detailed investigations of performance in a production environment. We present here a number of

updates and new results related to our efforts to optimise power efficiency for High Energy Physics (HEP) research.

We will show updated benchmark results, including a new figure-of-merit designed to characterise the power usage during the execution of the HEPscore benchmark. Previously, community measurements have used either the average or maximum power, neither of which is a good estimator. We expand our HEP-Score/Watt comparison to include additional machines such as Ampere Altra Q80 and M80, NVidia Grace, and the most recent AMD EPYC chips. We also introduce a Frequency Scan methodology to better characterize performance/watt trade-offs, potentially informing strategies like frequency scaling during peak hours to optimize power efficiency.

In addition, we present a comparison of single-socket versus dual-socket performance, revealing consistent findings that dual-socket configurations exhibit performance degradation compared to two single-socket machines, though of varying magnitudes. Leveraging HEPscore jobs and the ‘taskset’ command to target specific core configurations, we explore performance variations across core groups within the same socket or across dual sockets. Preliminary results show that same-CPU cores have better performance, confirming the importance of workload optimization strategies, such as fine-tuning the job scheduler to prioritize same-socket core utilization.

Our findings contribute to advancing heterogeneous computing strategies and power efficiency optimizations in HEP, paving the way toward more sustainable hardware solutions.

**Parallel (Track 7) / 327**

## Taking on RISC for Energy-Efficient Computing in HEP

**Authors:** Emanuele Simili<sup>1</sup>; Tommaso Boccali<sup>2</sup>

**Co-authors:** Albert Gyorgy Borbely<sup>1</sup>; David Britton<sup>1</sup>; Gordon Stewart; Samuel Cadellin Skipsey

<sup>1</sup> *University of Glasgow (GB)*

<sup>2</sup> *INFN Sezione di Pisa*

**Corresponding Authors:** emanuele.simili@glasgow.ac.uk, tommaso.boccali@cern.ch, david.britton@cern.ch, albert.borbely@cern.ch, samuel.cadellin.skipsey@cern.ch, gordon.stewart@glasgow.ac.uk

In pursuit of energy-efficient solutions for computing in High Energy Physics (HEP) we have extended our investigations of non-x86 architectures beyond the ARM platforms that we have previously studied. In this work, we have taken a first look at the RISC-V architecture for HEP workloads, leveraging advancements in both hardware and software maturity.

We introduce the Pioneer Milk-V, a 64-core RISC-V machine running Fedora Linux, as our new testbed, available at ScotGrid Glasgow (UK) and INFN Bologna (Italy). Despite this early stage of RISC-V adoption in HEP, significant progress has been made in software compatibility. Standard frameworks such as ROOT, Geant4, CVMFS, and XRootD can be successfully built and run on the RISC-V platform, showcasing the evolving ecosystem. Additionally, efforts are underway to port CMSSW, promising further integration of HEP experiment software.

In this first study, we assess performance and power efficiency, and we leverage various benchmarking tools to compare the RISC-V system with existing ARM and x86 architectures. Although it is not yet possible to run the HEPscore suite, we have conducted ROOT tests and benchmarks, along with DB12 and HS06 benchmarks, demonstrating promising performance-per-watt on the RISC-V platform.

These early results suggest that RISC-V architecture holds potential for advancing energy-efficient computing in HEP, offering decent performance and significantly better power efficiency, while contributing to an increasingly heterogeneous computing landscape.

**Poster session / 328**

## Parsing TTreeFormula in Python

**Author:** Aryan Roy<sup>None</sup>**Co-author:** Jim Pivarski<sup>1</sup><sup>1</sup> *Princeton University***Corresponding Authors:** aryanroy5678@gmail.com, pivarski@princeton.edu

Uproot can read ROOT files directly in pure Python but cannot (yet) compute expressions in ROOT's TTreeFormula expression language. Despite its popularity, this language has only one implementation and no formal specification. In a package called "formulate," we defined the language's syntax in standard BNF and parse it with Lark, a fast and modern parsing toolkit in Python. With formulate, users can now convert ROOT TTreeFormula expressions into NumExpr and Awkward Array manipulations.

In this contribution, we describe BNF notation and the Look Ahead Left to Right (LALR) parsing algorithm, which scales linearly with expression length. We also present the challenges with interpreting TTreeFormula expressions as a functional language; some function-like forms can't be expressed as true functions. We also describe the design of the abstract syntax tree that facilitates conversion between the three languages. The formulate package has zero package dependencies, so we are adding it as one of Uproot's dependencies so that Uproot will be able to use TTreeFormula expressions, whether they are hand-written or embedded in a ROOT file as TTree aliases.

**Poster session / 329**

## CppInterOp: Advancing Interactive C++ for High Energy Physics

**Authors:** Aaron Jomy<sup>1</sup>; Baidyanath Kundu<sup>1</sup>; Alexander Penev<sup>2</sup>; Wim Lavrijsen<sup>3</sup>; Vassil Vasilev<sup>1</sup><sup>1</sup> *Princeton University (US)*<sup>2</sup> *Paisii Hilendarski University of Plovdiv*<sup>3</sup> *Lawrence Berkeley National Laboratory (LBL)***Corresponding Authors:** vassil.vassilev@cern.ch, baidyanath.kundu@cern.ch, apenev@uni-plovdiv.bg, aaron.jomy@cern.ch, wlavrijsen@lbl.gov

The Cling C++ interpreter has transformed language bindings by enabling incremental compilation at runtime. This allows Python to interact with C++ on demand and lazily construct bindings between the two. The emergence of Clang-REPL as a potential alternative to Cling within the LLVM compiler framework highlights the need for a unified framework for interactive C++ technologies.

We present CppInterOp, a C++ Interoperability library, which leverages Cling and LLVM's Clang-REPL, to provide a minimalist and backward-compatible API facilitating seamless language interoperability. This provides downstream interactive C++ tools with the compiler as a service by embedding Clang and LLVM as libraries in their codebases. By enabling dynamic Python interactions with static C++ codebases, CppInterOp enhances computational efficiency and rapid development in high-energy physics. The library offers primitives enabling cppy (PyROOT), an automatic, runtime, Python-C++ bindings generator. We also demonstrate CppInterOp's utility in diverse computing environments through its adoption as the runtime engine for xeus-cpp, a Jupyter kernel designed for C++.

CppInterOp is a general-purpose library inspired by the developments in the ROOT framework which pushed the frontiers of interactive C++. It aims to extend this approach and serve as an integral component of ROOT, enhancing both speed and resilience. This talk introduces CppInterOp to

the HEP community and showcases how it optimizes cross-language execution and computational tasks in high-energy physics, making it a valuable tool for researchers and developers.

**Parallel (Track 9) / 330**

## **On-Grid GPU development via interactive HTCondor jobs and Analysis Facility style workflows**

**Author:** Albert Gyorgy Borbely<sup>1</sup>

**Co-authors:** David Britton<sup>1</sup>; Emanuele Simili<sup>1</sup>; Gordon Stewart; Samuel Cadellin Skipsey

<sup>1</sup> *University of Glasgow (GB)*

**Corresponding Authors:** samuel.cadellin.skipsey@cern.ch, emanuele.simili@glasgow.ac.uk, albert.borbely@cern.ch, david.britton@cern.ch, gordon.stewart@glasgow.ac.uk

Over the last few years, an increasing number of sites have started to offer access to GPU accelerator cards but in many places they remain underutilised. The experiment collaborations are gradually increasing the fraction of their code that can exploit GPUs, driven in many cases by developments of specific reconstruction algorithms to exploit the HLT farms when data is not being taken. However, there is no wide-spread usage of GPUs on the Grid and no mechanism, yet, to pledge GPU resources. Whilst the experiments gradually make progress porting their production code, and external projects such as Celeritas and AdePT tackle key common tasks such as the acceleration of E/M calorimeter simulation as a plug-in for GEANT, there is no easy way for smaller groups or individual developers to develop GPU usage in a way that is easily transferred to the Grid environment. Currently, a user typically develops code on a local GPU in an interactive manner but there is significant overhead in subsequently containerising this work and moving it to the Grid environment. Indeed, many user jobs are not big enough to benefit from this last step and many sites must then maintain GPUs that are not integrated with the Grid infrastructure.

We have developed a proof-of-principle solution to enable interactive user access to Grid GPUs, enabling the initial development to take place on-Grid. This will ensure the development and production environments are identical and enable sites to move more GPUs to the Grid. An interactive development environment has been implemented with interactive HTCondor jobs and Apptainer containers. GPUs are split into MIG instances to allow for simultaneous multi-user utilisation. Users can install packages on the fly, giving them control over package versions as well as use what's available on CVMFS. Once development is done the sandbox container can be made immutable and submitted to either the local batch style GPU queue or sent to the rest of the GPUs available on the Grid. The nature of interactive development means many hurdles had to be overcome such as: User authentication, security considerations, data replication to other sites, as well as management tools to allowing users to keep track of their environments and jobs. We will report for the first time on the current status of this project.

**Poster session / 331**

## **User sharing of computational workflows in the REANA reproducible analysis platform**

**Authors:** Marco Donadoni<sup>1</sup>; Daan Eduard Rosendal<sup>2</sup>; Giuseppe Steduto<sup>3</sup>; Tibor Simko<sup>1</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *Christelijke Hogeschool Windesheim (NL)*

<sup>3</sup> *Politecnico di Milano (IT)*



**Corresponding Authors:** daan.rosendal@cern.ch, tibor.simko@cern.ch, giuseppe.steduto@cern.ch, marco.donadoni@cern.ch

We present the new user-sharing feature of the REANA reproducible analysis platform. The researchers are allowed to share their selected workflow runs, job logs, and output files with colleagues. The analyst retains the full read-write access to the workflow and may opt for granting individual read-only access to colleagues for a possibly-limited period of time. The workflow sharing feature was developed to answer the needs of physics teams using REANA computational workflow platform and is available for all supported CWL, Serial, Snakemake, and Yadage workflow systems. The feature is available in the REANA command-line client and on the REANA web interface. The contribution describes the main use cases, presents the architecture and the implementation details, as well as comments on the challenges of supporting a variety of external Identity and Access Management systems holding user information for customising REANA deployments.

**Parallel (Track 9) / 332**

## **Benchmarking massively-parallel Analysis Grand Challenge workflows using Snakemake and REANA**

**Authors:** Marco Donadoni<sup>1</sup>; Matthew Feickert<sup>2</sup>; Alexander Held<sup>2</sup>; Andrii Povsten<sup>3</sup>; Oksana Shadura<sup>4</sup>; Tibor Simko<sup>1</sup>

<sup>1</sup> CERN

<sup>2</sup> University of Wisconsin Madison (US)

<sup>3</sup> Princeton University (US)

<sup>4</sup> University of Nebraska Lincoln (US)

**Corresponding Authors:** matthew.feickert@cern.ch, andrii.povsten@cern.ch, oksana.shadura@cern.ch, tibor.simko@cern.ch, marco.donadoni@cern.ch, alexander.held@cern.ch

We have created a Snakemake computational analysis workflow corresponding to the IRIS-HEP Analysis Grand Challenge (AGC) example studying  $t\bar{t}$  production channels in the CMS open data. We describe the extensions to the AGC pipeline that allowed porting of the notebook-based analysis to Snakemake. We discuss the applicability of the Snakemake multi-cascading paradigm for running massively-parallel RECAST-compatible physics analysis workflows where the analysis process may run over numerous independent data samples with large number of independent individual data files in a fully concurrent manner. The created Snakemake workflow example was run on the REANA reproducible analysis platform. We describe the improvements brought to the REANA job scheduling, tracking and termination processes for massively-parallel Snakemake workflows. We present results of several numerical experiments running the same workflow on the Kubernetes cluster with increasing number of identical nodes. We infer on the feasibility of REANA to schedule numerous concurrent jobs from the same Snakemake workflow rule, study the importance of cluster node size from the point of view of the job memory requirements, as well as estimate the overhead of dispatching workload to many cluster nodes. The results demonstrate the applicability of Snakemake for even massively-parallel RECAST-compatible physics analysis workflows.

**Parallel (Track 1) / 333**

## **Reading Tea Leaves - Understanding internal events and addressing performance issues within a CephFS/XRootD Storage Element.**

**Authors:** Gerard Hand<sup>1</sup>; Matt Doidge<sup>1</sup>; Peter Love<sup>1</sup>; Steven Simpson<sup>2</sup>

<sup>1</sup> Lancaster University (GB)

<sup>2</sup> Lancaster University

**Corresponding Authors:** m.doidge@cern.ch, peter.love@cern.ch, s.simpson@lancaster.ac.uk, gerard.hand@cern.ch

Erasure-coded storage systems based on Ceph have become a mainstay within UK Grid sites as a means of providing bulk data storage whilst maintaining a good balance between data safety and space efficiency. A favoured deployment, as used at the Lancaster Tier-2 WLCG site, is to use CephFS mounted on frontend XRootD gateways as a means of presenting this storage to grid users.

These storage systems are complex and self-correcting, but despite access to a myriad of metrics the inner workings of the storage tend to be opaque to the storage admin. One of the common problems seen within Ceph based systems are “Slow Ops” - instances of operations that take longer than expected, that are also often blocking in nature, impacting the overall performance and reliability of the system. These could be caused by, for example, intensive client side usage, internal CEPH data movement or hardware and/or network issues. Identifying the causes of a slow operation can provide a means to prevent or reduce the impact of future occurrences, leading to an increase in performance and reliability.

We detail the Lancaster Grid Site’s attempts to understand the causes of and mitigate against these “Slow Ops” and other performance bottlenecks within our storage system, with a focus on deletions as a case study on operations with a potential high-impact for the Ceph backend. We endeavour to bring together a holistic monitoring model, utilising Ceph metrics, detailed XRootD monitoring streams and client-side logging, in order to understand how data-management events impact the health of the storage.

**Poster session / 334**

## **EvtGen –on its first steps towards thread safety**

**Authors:** Fernando Abudinén<sup>1</sup>; John James Back<sup>1</sup>; Michal Kreps<sup>1</sup>; Thomas Latham<sup>1</sup>

<sup>1</sup> *University of Warwick (GB)*

**Corresponding Authors:** j.j.back@warwick.ac.uk, fernando.abudinen@cern.ch, michal.kreps@cern.ch, t.latham@warwick.ac.uk

The EvtGen generator, an essential tool for the simulation of heavy-flavour hadron decays, has recently gone through a modernisation campaign aiming to implement thread safety. A first iteration of this concluded with an adaptation of the core software, where we identified possibilities for future developments to further exploit the capabilities of multi-threaded processing. However, the current main limitations stem from external dependencies that are not yet thread safe, such as the simulation of final state radiation (FSR). Along with thread safety, we have recently implemented alternatives for FSR simulation which open new possibilities for systematic studies.

**Parallel (Track 4) / 335**

## **CMS Token Transition**

**Authors:** Alan Malta Rodrigues<sup>1</sup>; Brian Paul Bockelman<sup>2</sup>; Chan-Anun Runghitakchai<sup>2</sup>; Dave Dykstra<sup>3</sup>; Diego Ciangottini<sup>4</sup>; Edita Kizinevic<sup>5</sup>; Eric Vaandering<sup>3</sup>; Marco Mascheroni<sup>6</sup>; Panos Paparrigopoulos<sup>5</sup>; Rahul Chauhan<sup>5</sup>; Sarun Nuntaviriyaikul<sup>7</sup>; Stephan Lammel<sup>3</sup>; Vaiva Zokaite<sup>8</sup>

<sup>1</sup> *University of Notre Dame (US)*

<sup>2</sup> *University of Wisconsin Madison (US)*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

<sup>4</sup> *INFN, Perugia (IT)*

<sup>5</sup> *CERN*

<sup>6</sup> *Univ. of California San Diego (US)*

<sup>7</sup> *Chulalongkorn University (TH)*

<sup>8</sup> *Vilnius University (LT)*

**Corresponding Authors:** edita.kizinevic@cern.ch, panos.paparrigopoulos@cern.ch, sarun.nuntaviriyakul@cern.ch, brian.bockelman@cern.ch, lammel@fnal.gov, dwd@fnal.gov, rahul.chauhan@cern.ch, vaiva.zokaite@cern.ch, ewv@fnal.gov, marco.mascheroni@cern.ch, diego.ciangottini@cern.ch, alan.malta@cern.ch, rungphitakch@wisc.edu

Within the LHC community, a momentous transition has been occurring in authorization. For nearly 20 years, services within the Worldwide LHC Computing Grid (WLCG) have authorized based on mapping an identity, derived from an X.509 credential, or a group/role derived from a VOMS extension issued by the experiment. A fundamental shift is occurring to capabilities: the credential, a bearer token, asserts the authorizations of the bearer, not the identity.

By the HL-LHC era, the CMS experiment plans for the transition to tokens, based on the WLCG Common JSON Web Token profile, to be complete. Services in the technology architecture include the INDIGO Identity and Access Management server to issue tokens; a HashiCorp Vault server to store and refresh access tokens for users and jobs; a managed token bastion server to push credentials to the HTCondor CredMon service; and HTCondor to maintain valid tokens in long-running batch jobs. We will describe the transition plans of the experiment, current status, configuration of the central authorization server, lessons learned in commissioning token-based access with sites, and operational experience using tokens for both job submissions and file transfers.

**Poster session / 336**

## Heterogeneous reconstruction of hadronic Particle Flow clusters with the Alpaka Portability Library

**Author:** Jonathan Samudio<sup>1</sup>

<sup>1</sup> *Baylor University (US)*

**Corresponding Author:** jonathan.jacob.samudio@cern.ch

In response to increasing data challenges, CMS has adopted the use of GPU offloading at the High-Level Trigger (HLT). However, GPU acceleration is often hardware specific, and increases the maintenance burden on software development. The Alpaka (Abstraction Library for Parallel Kernel Acceleration) portability library offers a solution to this issue, and has been implemented into the CMS software (CMSSW) for use online at HLT.

A portion of the final-state particle candidate reconstruction algorithm, Particle Flow, has been ported to Alpaka and deployed at HLT for 2024 data taking. The formation of hadronic Particle Flow clusters represented a target for increased performance through parallel operation. We will discuss the port of hadronic Particle Flow clustering to Alpaka, and the validation of physics and performance at HLT.

**Parallel (Track 9) / 338**

## Efficiency, Reproducibility, and Portability in HEP Machine Learning Training - ML Training Facility at Vanderbilt University

**Author:** Jethro Taylor Gaglione<sup>1</sup>

<sup>1</sup> *Vanderbilt University (US)*

The success and adoption of machine learning (ML) approaches to solving HEP problems has been widespread and fast. As useful a tool as ML has been to the field, the growing number of applications,

larger datasets, and increasing complexity of models creates a demand for both more capable hardware infrastructure and cleaner methods of reproducibility and deployment. We have developed a prototype ML Training facility (MLTF) with the goal of meeting these demands. The proof-of-concept MLTF is based at ACCRE, Vanderbilt's computing cluster, with sufficient GPU storage and networking to efficiently test very large models. The software component of MLTF is developed with an eye on reproducibility and portability. We adapt MLflow as an end-to-end ML solution for its capabilities as a user-friendly job submission interface; as a tracking server for model and run details, arbitrary metrics logging, and system diagnostics logging; and as an inference server.

**Parallel (Track 3) / 339**

## **A Multi-objective Graph Neural Network for enhanced HGICAL hadronic reconstruction**

**Author:** Jekaterina Jaroslavceva<sup>1</sup>

<sup>1</sup> *Czech Technical University in Prague (CZ)*

In response to the increased data complexity anticipated with the upcoming upgrade of the Large Hadron Collider (LHC), the Compact Muon Solenoid (CMS) at LHC is developing an advanced endcap High-Granularity Calorimeter (HGICAL) capable of enduring the more demanding conditions of the High-Luminosity LHC with about 200 overlapping proton-proton collisions in a single bunch-crossing, resulting in several hundred thousands of hits in each endcap. During the particle shower reconstruction phase in HGICAL, 3D graph structures called tracksters are generated. These tracksters connect energy deposits across each layer of the detector, representing clusters of energy believed to originate from the same physics object. However, the inhomogeneous geometry of the detector, coupled with the lumpy nature of hadronic showers, particle overlaps and preceding algorithm cuts being tuned for high purity, often leads to full particle showers being fragmented into multiple tracksters. This effect compromises the quality of the reconstruction and requires addressing through an additional trackster linking step. This study delves into a machine learning approach leveraging Graph Neural Network (GNN) models with attention mechanisms to enhance the multi-purpose tasks of calorimetric event reconstruction, including trackster linking, energy regression, and particle identification. In this work, we show the result of applying the proposed model when evaluating hadronic showers data in the dense environment of HGICAL, with the network fully integrated into the CMS Software.

**Poster session / 341**

## **Object storage model for CMS data**

**Author:** CMS Collaboration<sup>None</sup>

**Corresponding Author:** borislav.pavlov@cern.ch

In CMS, data access and management is organized around the data-tier model: a static definition of what subset of event information is available in a particular dataset, realized as a collection of files. In previous works, we have proposed a novel data management model that obviates the need for data tiers by exploding files into individual event data product objects. We present here a study of the fraction of event data products per data-tier actively read by CMS users as collected by CRAB3, to estimate the storage savings CMS could realize by adopting such a model. We also present an update on ongoing work to interface RNTuple with a Ceph object store through the S3 protocol.

**Parallel (Track 2) / 344**

## Anomaly detection for data quality monitoring of the Muon system at CMS

**Author:** Marco Buonsante<sup>1</sup>

<sup>1</sup> *Universita e INFN, Bari (IT)*

Ensuring the quality of data in large HEP experiments such as CMS at the LHC is crucial for producing reliable physics outcomes. The CMS protocols for Data Quality Monitoring (DQM) are based on the analysis of a standardized set of histograms offering a condensed snapshot of the detector's condition. Besides the required personpower, the method has a limited time granularity, potentially hiding temporary anomalies. Unsupervised machine learning models such as auto encoders and convolutional neural networks have been recently deployed for anomaly detection with per-lumisection granularity. Nevertheless, given the diversity of detector technologies, geometries and physics signals characterizing each subdetector, different tools are developed in parallel and maintained by the sub detector experts. In this contribution, we discuss the development of an automated DQM for the online monitoring of the CMS Muon system, offering a flexible tool for the different muon subsystems based on deep learning models trained on occupancy maps. The potential flexibility and extensibility to different detectors, as well as the effort towards the integration of per-lumisection monitoring in the DQM workflow will be discussed.

**Parallel (Track 3) / 345**

## Enhancing GRB Detection: Machine Learning Optimization of Triggerless Data Analysis Algorithms for LHAASO-WCDA

**Authors:** Abdulhafiz Ahmed Mustofa<sup>1</sup>; Yao Dong Cheng<sup>2</sup>; Zhiguo Yao<sup>3</sup>

<sup>1</sup> *IHEP, UCAS*

<sup>2</sup> *IHEP*

<sup>3</sup> *IHEP, Beijing*

**Corresponding Authors:** yaozg@ihep.ac.cn, mustofa@ihep.ac.cn, chyd@ihep.ac.cn

Detecting Gamma-Ray Burst (GRB) signals from triggerless data poses significant challenges due to high noise levels, a problem similarly encountered in the Large High Altitude Air Shower Observatory's Water Cherenkov Detector Array (LHAASO-WCDA) triggerless data analysis. This research aims to enhance the GRB triggerless data algorithm which leverages the distinct spatial properties of gamma-ray showers. By incorporating advanced machine learning techniques such as Bayesian optimization, we refine the algorithm to more effectively detect GRB signals within noisy background signals. Preliminary findings indicate a marked improvement in the detection of GRB events, suggesting that machine learning methods can substantially enhance existing astrophysical data analysis techniques. These methods could lead to more accurate and reliable identification of GRB signals, thereby contributing to our understanding of these cosmic phenomena.

**Parallel (Track 7) / 346**

## A RoCE-based network framework for science workloads in HEPS data center

**Authors:** Fazhi QI<sup>1</sup>; Tao CUI<sup>1</sup>; 曾珊 zengshan<sup>1</sup>

<sup>1</sup> *IHEP*

**Corresponding Author:** zengshan@ihep.ac.cn

According to the estimated data rates, it is predicted that 800 TB raw experimental data will be produced per day from 14 beamlines at the first stage of the High-Energy Photon Source (HEPS) in China, and the data volume will be even greater with the completion of over 90 beamlines at the second stage in the future. Therefore, designing a high-performance, scalable network architecture plays a crucial role in the efficient output of scientific tasks. We designed a RoCE-based network framework for science workloads in HEPS data center, which provides high-performance network connectivity between HPES Data Acquisition system(DAQ) and HEPS data center, as well as the compute and storage system within the HEPS data center. The test results show that the performance of the RoCE-based network framework of the HEPS data center can be comparable to that of the IB-based network framework, and is better than the TCP/IP-based network framework.

**Poster session / 347**

## Keep-up Production in JUNO's Offline Data Processing

**Author:** 尹维卿 yinwq<sup>None</sup>

**Corresponding Author:** yinwq@ihep.ac.cn

On behalf of JUNO collaboration.

The Jiangmen Underground Neutrino Observatory (JUNO), located in Southern China, is a neutrino experiment aiming to determine the neutrino mass ordering (NMO) and precisely measure neutrino oscillation parameters. JUNO is expected to operate over 20-30 years, generating approximately 2PB of raw data annually. Offline Data Processing Workflow involves data transfer, reconstruction, grid computing, and long-term data preservation. Keep Up Production (KUP) pipeline addresses the need for a pipeline-driven approach to handle the intricate and interdependent steps of raw data processing without human intervention. The message-driven architecture decouples subsystems of pipeline and allows subsystems to process data in an asynchronous manner, which means that each processing step can be performed independently. KUP will automatically create and submit the reconstruction job. It uses YML files as job templates, which provides flexibility, allowing for easy modification and adjustment of the template without the need to modify the system's code. KUP also provides the real-time tracking and monitoring feature of job running status. It enhances job stability and reliability, minimizing the impact of job failures.

**Parallel (Track 1) / 348**

## Enhancing XRootD Load Balancing for High-Throughput transfers

**Authors:** Thomas Byrne<sup>None</sup>, Thomas Jyothish<sup>1</sup>

**Co-author:** James William Walder<sup>2</sup>

<sup>1</sup> STFC

<sup>2</sup> Science and Technology Facilities Council STFC (GB)

**Corresponding Authors:** james.william.walder@cern.ch, tom.byrne@stfc.ac.uk, jyothish.thomas@stfc.ac.uk

To address the need for high transfer throughput for projects such as the LHC experiments, including the upcoming HL-LHC, it is important to make optimal and sustainable use of our available capacity. Load balancing algorithms play a crucial role in distributing incoming network traffic across multiple servers, ensuring optimal resource utilization, preventing server overload, and enhancing performance and reliability. At the Rutherford Appleton Laboratory (RAL), the UK's Tier-1 centre for the Worldwide LHC Computing Grid (WLCG), we started with a DNS round robin then moved to XRootD's cluster management service component, which has an active load balancing algorithm to distribute traffic across 26 servers, but encountered its limitations when the system as a

whole is under heavy load. We describe our tuning of the configuration of the existing algorithm before proposing a new tuneable, dynamic load-balancer based on a weighted random selection algorithm.

Parallel (Track 3) / 349

## Primary vertex timing reconstruction with the LHCb Ring Imaging Cherenkov detectors

**Author:** Lorenzo Malentacca<sup>1</sup>

<sup>1</sup> *Cern, Milano-Bicocca*

**Corresponding Author:** lorenzo.malentacca@cern.ch

During LHC High-Luminosity phase, the LHCb RICH detector will face challenges due to increased particle multiplicity and high occupancy. Introducing sub-100ps time information becomes crucial for maintaining excellent particle identification (PID) performance. The LHCb RICH collaboration plans to anticipate the introduction of timing through an enhancement program during the third LHC Long Shutdown. In the RICH detector, Cherenkov photons from a track arrive nearly simultaneously at the detector plane, allowing precise hit time prediction. The RICH reconstruction algorithm computes track and photon time-of-flight and estimates where photons are expected on the photodetector plane. Determining the primary vertex time (PV  $T_0$ ) is crucial in predicting the time of arrival of photons on the photodetector plane. Adding time information allows applying a software time gate around the predicted time per track to enhance signal-to-background ratio and PID performance. This contribution describes how to estimate the PV  $T_0$  using RICH information only, a novel approach for LHCb. The proposed algorithm computes a reconstructed PV time for every photon from hit time and tracking information. The PV  $T_0$  is extracted by averaging this reconstructed time for all photons belonging to the PV. The challenge lies in correctly associate photons to their PV, which is a two-step process: PV-track and track-photon associations, both presenting inefficiencies. Results compare the estimated PV time resolution with Monte Carlo simulations. This contribution aims to describe the integration of fast-timing in the RICH detector, illustrating the impact of the PV time estimation method on PID performance.

Poster session / 350

## Research on Wide Area Network Performance Anomaly Detection Technology Based on Machine Learning

**Authors:** 曾珊 zengshan<sup>None</sup>, 李骋 chengli<sup>None</sup>

**Corresponding Authors:** chengli@ihep.ac.cn, zengshan@ihep.ac.cn

As a WLCG prototype T1 site, IHEP's network performance directly impacts the site's reliability. The current primary method for measuring network performance is implemented through Perfsonar, which actively measures performance metrics such as bandwidth, connection status, one-way and two-way latency, packet loss rate, and jitter between IHEP and other sites. However, there is a lack of relatively efficient network performance issue detection capabilities, posing significant challenges for network operations personnel when addressing network performance problems. This paper proposes a machine learning-based network anomaly detection algorithm, utilizing performance metric data obtained from both Perfsonar and third-party network monitoring tools. By integrating network protocol analysis and network traffic analysis techniques, the algorithm achieves network communication anomaly detection and alerting. Ultimately, this enhances the ability to detect network performance issues, helping network operations personnel to provide a more efficient network environment more effectively and quickly.

**Parallel (Track 5) / 351**

## Taking derivatives of Geant4 - closer than you might think?

**Authors:** Max Aehle<sup>None</sup>; Max Aehle<sup>1</sup>**Co-authors:** David Lange ; David Lange<sup>2</sup>; Lukas Alexander Heinrich<sup>3</sup>; Michael Kagan<sup>4</sup>; Mihaly Novak<sup>5</sup>; Nicolas Gauger ; Vassil Vasilev<sup>2</sup>; Vassil Vasilev<sup>2</sup><sup>1</sup> *Fachhochschule Worms (DE)*<sup>2</sup> *Princeton University (US)*<sup>3</sup> *Technische Universitat Munchen (DE)*<sup>4</sup> *SLAC National Accelerator Laboratory (US)*<sup>5</sup> *CERN***Corresponding Authors:** max.aehle@scicomp.uni-kl.de, max.aehle@cern.ch, michael.aaron.kagan@cern.ch, david.lange@princeton.edu, mihaly.novak@cern.ch, vasil.georgiev.vasilev@cern.ch, nicolas.gauger@scicomp.uni-kl.de, lukas.heinrich@cern.ch, vassil.vasilev@cern.ch, david.lange@cern.ch

Built on algorithmic differentiation (AD) techniques, differentiable programming allows to evaluate derivatives of computer programs. Such derivatives are useful across domains for gradient-based design optimization and parameter fitting, among other applications. In high-energy physics, AD is frequently used in machine learning model training and in statistical inference tasks such as maximum likelihood estimation. Recently, AD has begun to be explored for the end-to-end optimization of particle detectors, with potential applications ranging from HEP to medical physics to astrophysics. To that end, the ability to estimate derivatives of the Geant4 simulator for the passage of particles through matter would be a huge step forward.

The complexity of Geant4, its programmatic control flow, and its underlying stochastic sampling processes, introduce challenges that cannot all be addressed by current AD tools. As such, the application of current AD tools to Geant4-like simulations can provide invaluable insights into the accuracy and errors of the AD gradient estimates and into how to address remaining challenges.

In this spirit, we have applied the operator-overloading AD tool CoDiPack to the compact G4HepEm/HepEmShow package for the simulation of electromagnetic showers in a simple sampling calorimeter. Our AD-enabled simulator allows to estimate derivatives of energy depositions with respect to properties of the geometry and the incoming particles. The derivative estimator comes with a small bias, which however proved unproblematic in a simple optimization study. In this talk, we will report on our methodology and encouraging results, and demonstrate how a next-generation AD tool, Derivgrind, can be used to bring these results to the scale of Geant4.

**Parallel (Track 3) / 352**

## TrackHHL: A Quantum Computing Algorithm for Track Reconstruction at the LHCb

**Author:** Xenofon Chiotopoulos<sup>1</sup>**Co-authors:** Davide Nicotra<sup>1</sup>; Jacco A. de Vries<sup>1</sup>; Kurt Driessens<sup>1</sup>; Marcel Merk<sup>1</sup>; Mark H.M. Winands<sup>1</sup>; Miriam Lucio Martinez<sup>1</sup><sup>1</sup> *Maastricht University***Corresponding Authors:** xenofon.chiotopoulos@maastrichtuniversity.nl, m.luciomartinez@maastrichtuniversity.nl, m.merk@maastrichtuniversity.nl, kurt.driessens@maastrichtuniversity.nl, m.winands@maastrichtuniversity.nl, d.nicotra@maastrichtuniversity.nl, jacco.devries@maastrichtuniversity.nl

With the future high-luminosity LHC era fast approaching high-energy physics faces large computational challenges for event reconstruction. Employing the LHCb vertex locator as our case study



we are investigating a new approach for charged particle track reconstruction. This new algorithm hinges on minimizing an Ising-like Hamiltonian using matrix inversion. Performing this matrix inversion classically achieves reconstruction efficiency akin to the current state-of-the-art algorithms but is hindered by worse time complexity. Exploiting the Harrow-Hassadim-Lloyd (HHL) quantum algorithm for linear systems holds the promise of an exponential speedup in the number of input hits over its classical counterpart. Contingent upon the following conditions: efficient quantum phase estimation (QPE) and an intuitive way to read out the algorithm's output. This contribution builds on previous work (DOI 10.1088/1748-0221/18/11/P11028) and strives to fulfil these conditions and streamlines the proposed algorithm's circuit depth, by a factor up to  $10^4$ . We propose a modified version of the HHL algorithm by restricting QPE precision to two bits. Enabling us to introduce a novel post-processing algorithm, which estimates event Primary Vertices (PVs), then efficiently computes all event tracks through an Adaptive Hough Transform. This alteration significantly reduces circuit depth and addresses HHL's readout issue, bringing the reconstruction of small events closer to current hardware implementation. The findings presented here aim to further illuminate the potential of harnessing quantum computing for the future of particle track reconstruction in high-energy physics.

Parallel (Track 2) / 353

## TrackNET: Deep Learning-Based Track Recognition in Pixel and Strip-Based Particle Detectors

**Author:** Pavel Goncharov<sup>None</sup>

**Co-authors:** Daniil Rusov<sup>1</sup>; Anastasiya Nikolskaya<sup>1</sup>; Gennady Ososkov<sup>1</sup>; Alexey Zhemchugov<sup>1</sup>

<sup>1</sup> *Joint Institute for Nuclear Research*

**Corresponding Authors:** gososkov@gmail.com, pgoncharov13@gmail.com

The reconstruction of charged particle trajectories in tracking detectors is crucial for analyzing experimental data in high-energy and nuclear physics. Processing of the vast amount of data generated by modern experiments requires computationally efficient solutions to save time and resources. In response, we introduce TrackNET, a recurrent neural network specifically designed for track recognition in pixel and strip-based particle detectors. TrackNET acts as a scalable alternative to the Kalman filter, exemplifying local tracking methods by independently processing each track-candidate. We rigorously tested TrackNET using the TrackML dataset and simulated data from the straw tracker of the SPD experiment at JINR, Dubna. Our results demonstrate significant improvements in processing speed and accuracy. The paper concludes with a comprehensive analysis of TrackNET's performance and a discussion on its limitations and potential enhancements.

Parallel (Track 3) / 354

## Guided Quantum Compression for Higgs Identification

**Authors:** Vasilis Belis<sup>1</sup>; Denis-Patrick Odagiu<sup>1</sup>; Michele Grossi<sup>2</sup>; Florentin Reiter<sup>None</sup>; Guenther Dissertori<sup>1</sup>; Sofia Vallecorsa<sup>2</sup>

<sup>1</sup> *ETH Zurich (CH)*

<sup>2</sup> *CERN*

**Corresponding Authors:** sofia.vallecorsa@cern.ch, freiter@ethz.ch, guenther.dissertori@cern.ch, podagiu@student.ethz.ch, vasileios.belis@cern.ch, michele.grossi@cern.ch

Quantum machine learning provides a fundamentally novel and promising approach to analyzing data. However, many data sets are too complex for currently available quantum computers. Consequently, quantum machine learning applications conventionally resort to dimensionality reduction

algorithms, e.g., auto-encoders, before passing data through the quantum models.

We show that using a classical auto-encoder as an independent preprocessing step can significantly decrease the classification performance of a quantum machine learning algorithm. To ameliorate this issue, we design an architecture that unifies the preprocessing and quantum classification algorithms into a single trainable model: the guided quantum compression model. The utility of this model is demonstrated by using it to identify the Higgs boson in proton-proton collisions at the LHC, where the conventional approach proves ineffective. Conversely, the guided quantum compression model excels at solving this classification problem, achieving a good accuracy. Additionally, the model developed herein shows better performance compared to the classical benchmark when using only low-level kinematic features.

**Parallel (Track 1) / 355**

## **Evolving StoRM WebDAV: delegation of file transfers to NGINX and support for SciTags**

**Authors:** Enrico Vianello<sup>1</sup>; Francesco Giacomini<sup>2</sup>; Luca Bassi<sup>None</sup>

**Co-authors:** Federica Agostini<sup>3</sup>; Jacopo Gasparetto<sup>3</sup>; Roberta Miccoli<sup>2</sup>; Stefano Enrico Zotti

<sup>1</sup> *INFN-CNAF*

<sup>2</sup> *INFN CNAF*

<sup>3</sup> *CNAF*

**Corresponding Authors:** roberta.miccoli@cnaf.infn.it, federica.agostini@cnaf.infn.it, francesco.giacomini@cern.ch, enrico.vianello@cern.ch, jacopo.gasparetto@cnaf.infn.it, stefano.zotti@cnaf.infn.it, luca.bassi@cnaf.infn.it

After the deprecation of the open-source Globus Toolkit used for GridFTP transfers, the WLCG community has shifted its focus to the HTTP protocol. The WebDAV protocol extends HTTP to create, move, copy and delete resources on web servers. StoRM WebDAV provides data storage access and management through the WebDAV protocol over a POSIX file system. Mainly designed to be used by the WLCG community, StoRM WebDAV supports authentication through X.509 certificates, VOMS proxies and JWT tokens. Moreover, Third-Party Copies (an extension of the WebDAV COPY verb to support copies between data centers) are supported.

With the aim of improving data transfer performance, this contribution describes the changes made to StoRM WebDAV in order to delegate file transfers to the external reverse proxy NGINX, decoupling them from the internal Java implementation. To even more simplify the StoRM WebDAV codebase, also the validation of VOMS proxies and JWT tokens is delegated to NGINX, augmented with specific modules developed by us. Even with this solution, authorization is still enforced by StoRM WebDAV.

Following the effort of the WLCG community to have better metrics about data flows, this contribution also describes the work done in order to support SciTags, an initiative promoting identification of the science domains and their high-level activities at the network level.

**Parallel (Track 5) / 356**

## **Benchmark Studies of ML Inference with TMVA SOFIE**

**Authors:** Ioanna Maria Panagou<sup>None</sup>; Lorenzo Moneta<sup>1</sup>; SANJIBAN SENGUPTA<sup>None</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** lorenzo.moneta@cern.ch, ioanna.maria.panagou@cern.ch, sanjiban.sg@gmail.com

Within the ROOT/TMVA project, we have developed a tool called SOFIE, that takes externally trained deep learning models in ONNX format or Keras and PyTorch native formats and generates C++ code that can be easily included and invoked for fast inference of the model. The code has a minimal dependency and can be easily integrated into the data processing and analysis workflows of the HEP experiments.

This study presents a comprehensive benchmark analysis of SOFIE and prominent machine learning frameworks for model evaluation such as PyTorch, TensorFlow XLA and ONNXRuntime. Our research focuses on evaluating the performance of these tools in the context of HEP, with an emphasis on their application with typical models used, such as Graph Neural Networks for jet tagging and Variational auto-encoder and GAN for fast simulation. We assess the tools based on several key parameters, including computational speed, memory usage, scalability, and ease of integration with existing HEP software ecosystems. Through this comparative study, we aim to provide insights that can guide the HEP community in selecting the most suitable framework for their specific needs.

Parallel (Track 7) / 357

## Front-End RDMA Over Converged Ethernet, lightweight RoCE endpoint

**Author:** Gabriele Bortolato<sup>1</sup>

**Co-authors:** Andrea Triossi<sup>1</sup>; Antonio Bergnoli<sup>1</sup>; Damiano Bortolato<sup>2</sup>; Daniele Mengoni<sup>1</sup>; Fabio Montecassiano<sup>1</sup>; Jacopo Pazzini<sup>3</sup>; Marco Zanetti<sup>1</sup>; Matteo Migliorini<sup>1</sup>; Sandro Ventura<sup>1</sup>

<sup>1</sup> *Universita e INFN, Padova (IT)*

<sup>2</sup> *INFN - National Institute for Nuclear Physics*

<sup>3</sup> *Università e INFN, Padova (IT)*

**Corresponding Authors:** sandrov@pd.infn.it, gabriele.bortolato@cern.ch, daniele.mengoni@cern.ch, andrea.trioffi@cern.ch, antonio.bergnoli@cern.ch, fabio.montecassiano@cern.ch, damiano.bortolato@lnl.infn.it, m.migliorini@cern.ch, marco.zanetti@cern.ch, jacopo.pazzini@cern.ch

In a DAQ system a large fraction of CPU resources is engaged in networking rather than in data processing. The common network stacks that take care of network traffic usually manipulate data through several copies performing expensive operations. Thus, when the CPU is asked to handle networking, the main drawbacks are throughput reduction and latency increase due to the overhead added to the data transmission process. Networking with zero-copy can be achieved by adding a Remote Direct Memory Access (RDMA) layer to the network stack and making dedicated hardware take care of the burden of the stack handling. Considering the ever-growing demand of larger bandwidth for big data systems, many works point in the direction of implementing network stacks on custom hardware. FPGAs are the natural target for reducing time to market and keeping a low entry-barrier. In this work implementation of RDMA directly on the front-end electronics is explored, in this way it is possible to free part of the computing farm's CPU resources. RDMA over Converged Ethernet (RoCE) is the industry-standard Ethernet-based RDMA solution with a multi-vendor ecosystem, making it the natural choice. This work focuses on the hardware implementation of a stripped-down version of RoCEv2 implementing only the transmitter part of the protocol, enabling its deployment in small FPGA such as the rad-hard parts used in the detector front-end. Preliminary results of resource usage, latency and throughput will be shown.

Parallel (Track 2) / 358

## Use of topological correlations in ML-based conditions for the CMS Level-1 Global Trigger upgrade for the HL-LHC

**Author:** Gabriele Bortolato<sup>1</sup>

**Co-authors:** Benjamin Huber<sup>2</sup>; Dinyar Rabady<sup>3</sup>; Elias Leutgeb<sup>2</sup>; Hannes Sakulin<sup>3</sup>; Jaana Heikkilae<sup>3</sup>

<sup>1</sup> *Universita e INFN, Padova (IT)*

<sup>2</sup> *Technische Universitaet Wien (AT)*

<sup>3</sup> *CERN*

**Corresponding Authors:** dinyar.rabady@cern.ch, jaana.kristiina.heikkilae@cern.ch, elias.leutgeb@cern.ch, hannes.sakulin@cern.ch, benjamin.huber@cern.ch, gabriele.bortolato@cern.ch

The High-Luminosity LHC upgrade will have a new trigger system that utilizes detailed information from the calorimeter, muon and track finder subsystems at the bunch crossing rate, which enables the final stage of the Level-1 Trigger, the Global Trigger (GT), to use high-precision trigger objects. In addition to cut-based algorithms, novel machine-learning-based algorithms will be employed in the trigger system to achieve higher selection efficiency and detect unexpected signals. The focus of this study is the comparison of different machine learning architecture models, including Boosted Decision Trees, Deep Neural Networks and Auto-Encoders. The trigger system will be implemented in FPGAs, benefiting from the performance of the employed AMD Ultrascale+ parts and an increased latency budget available in the new trigger system the utilization of topological correlations as inputs to these novel algorithms will be explored. Notable topological correlations employed are two-objects  $\Delta R$  and invariant masses (e.g. di-jets, di-muons, di-electrons). The effective FPGA hardware implementation and its optimization will play a key role in this study.

**Parallel (Track 3) / 359**

## The CMS Phase-2 High Granularity Calorimeter Endcap Event Reconstruction with the TICL Framework

**Authors:** Alessandro Tarabini<sup>1</sup>; Alexander Schmidt<sup>2</sup>; Aurora Perego<sup>3</sup>; Felice Pantaleo<sup>4</sup>; Florian Beaudette<sup>5</sup>; Jekaterina Jaroslavceva<sup>6</sup>; Kenichi Hatakeyama<sup>7</sup>; Marco Rovere<sup>4</sup>; Mohamed Darwish<sup>7</sup>; Shamik Ghosh<sup>5</sup>; Theo Cuisset<sup>8</sup>; Wahid Redjeb<sup>9</sup>

<sup>1</sup> *ETH Zurich (CH)*

<sup>2</sup> *RWTH Aachen (DE)*

<sup>3</sup> *Universita & INFN, Milano-Bicocca (IT)*

<sup>4</sup> *CERN*

<sup>5</sup> *Centre National de la Recherche Scientifique (FR)*

<sup>6</sup> *Czech Technical University in Prague (CZ)*

<sup>7</sup> *Baylor University (US)*

<sup>8</sup> *LLR / École Polytechnique (FR)*

<sup>9</sup> *Rheinisch Westfaelische Tech. Hoch. (DE)*

**Corresponding Authors:** kenichi.hatakeyama@cern.ch, alexander.schmidt@cern.ch, marco.rovere@cern.ch, mohamed.anwar@cern.ch, wahid.redjeb@cern.ch, aurora.perego@cern.ch, shamik.ghosh@cern.ch, theo.cuisset@polytechnique.edu, felice.pantaleo@cern.ch, florian.beaudette@cern.ch, alessandro.tarabini@cern.ch, jejarosl@cern.ch

The imminent high-luminosity era of the LHC will pose unprecedented challenges to the CMS detector. To meet these challenges, the CMS detector will undergo several upgrades, including replacing the current endcap calorimeters with a novel High-Granularity Calorimeter (HGCAL). A dedicated reconstruction framework, The Iterative Clustering (TICL), is being developed within the CMS Software (CMSSW). This new framework is designed to fully exploit the high spatial resolution and precise timing information provided by HGCAL, as well as the information from other subdetectors (e.g., Tracker and Mip-Timing-Detector). Its reconstruction capabilities aim to provide the final global event interpretation while mitigating the effects of the dense pile-up environment. The TICL framework, crafted with heterogeneous computing in mind, is a unique solution to the computing challenges of the HL-LHC phase. Data structures and algorithms have been developed for massively parallel architectures using the Alpaka performance portability library. The framework reconstructs particle candidates starting from the hundreds of thousands of energy deposits left in the calorimeter. Dedicated clustering algorithms have been developed to retain the physics information while reducing the problem complexity by order of magnitudes. Pattern recognition algorithms aim to reconstruct particle showers in the 3-dimensional space, striving for high efficiency and cluster purity,

keeping the pile-up contamination as low as possible. The high-purity requirements, together with detector inhomogeneity, lead to fragmented 3D clusters. An additional linking step is available to recover the fragmentation. In this step, several algorithms are adopted to target different types of particle shower reconstruction. A SuperClustering linking plugin has been developed for electron and photon reconstruction, while a geometrical linking is used to target the hadron reconstruction. The final charged candidates are built by linking Tracks with the HGAL 3D clusters, exploiting timing information from both HGAL and MTD. This presentation will introduce the TICL framework. Its physics and computational performance will be highlighted, showcasing the approach adopted to face the challenges of HL-LHC.

**Parallel (Track 4) / 360**

## Fermilab's Transition to Token Authentication

**Authors:** Dave Dykstra<sup>1</sup>; Marc Mengel<sup>1</sup>; Mine Altunay<sup>1</sup>; Shreyas Bhat<sup>1</sup>; Stephen White<sup>1</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** [mengel@fnal.gov](mailto:mengel@fnal.gov), [swhite@fnal.gov](mailto:swhite@fnal.gov), [sbhat@fnal.gov](mailto:sbhat@fnal.gov), [dwd@fnal.gov](mailto:dwd@fnal.gov), [maltunay@fnal.gov](mailto:maltunay@fnal.gov)

Fermilab is the first High Energy Physics institution to transition from X.509 user certificates to authentication tokens in production systems. All of the experiments that Fermilab hosts are now using JSON Web Token (JWT) access tokens in their grid jobs. Many software components have been either updated or created for this transition, and most of the software is available to others as open source. The tokens are defined using the WLCG Common JWT Profile. Token attributes for all the tokens are stored in the Fermilab FERRY system which generates the configuration for the CILogon token issuer. High security-value refresh tokens are stored in Hashicorp Vault configured by `htvault-config`, and JWT access tokens are requested by the `htgettoken` client through its integration with HTCondor. The Fermilab job submission system `jobsub` was redesigned to be a lightweight wrapper around HTCondor. For automated job submissions a managed tokens service was created to reduce duplication of effort and knowledge of how to securely keep tokens active. The existing Fermilab file transfer tool `ifdh` was updated to work seamlessly with tokens, as well as the Fermilab POMS (Production Operations Management System) which is used to manage automatic job submission and the RCDS (Rapid Code Distribution System) which is used to distribute analysis code via the CernVM FileSystem. The `dCache` storage system was reconfigured to accept tokens for authentication in place of X.509 proxy certificates. As some services and sites have not yet implemented token support, proxy certificates are still sent with jobs for backwards compatibility but some experiments are beginning to transition to stop using them. There have been some glitches and learning curve issues but in general the system has been performing well and is being improved as operational problems are addressed.

**Poster session / 361**

## Automation and Job Management for LZ Simulations at NERSC

**Author:** Jacopo Siniscalco<sup>None</sup>

**Corresponding Author:** [jacopo.siniscalco.22@ucl.ac.uk](mailto:jacopo.siniscalco.22@ucl.ac.uk)

The LUX-ZEPLIN (LZ) experiment is a world-leading direct dark matter detection experiment, implementing a dual-phase Xe Time Projection Chamber (TPC) design. The success of the experiment necessitates an in-depth characterization of the pertinent backgrounds, which in turn implies a heavy simulations burden. In this talk, I will present the infrastructure that was developed to allocate and manage the simulations workload on Perlmutter, NERSC's most recent HPC facility. The pipeline includes a system to automatically generate production configurations based on requests from the simulations team, along with utilities to monitor job progress and success. A RabbitMQ queue is used to coordinate job dispatchement amongst a selection of workers running on specially

allocated compute nodes, allowing for fine-grained control over the use of computational resources available.

**Parallel (Track 3) / 362**

## **tracc: GPU track reconstruction library for HEP experiments**

**Authors:** Andreas Salzburger<sup>1</sup>; Attila Krasznahorkay<sup>1</sup>; Beomki Yeo<sup>2</sup>; Heather Gray<sup>3</sup>; Joana Niermann<sup>4</sup>; Stephen Nicholas Swatman<sup>5</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *University of California Berkeley (US)*

<sup>3</sup> *UC Berkeley/LBNL*

<sup>4</sup> *Georg August Universitaet Goettingen (DE)*

<sup>5</sup> *University of Amsterdam (NL)*

**Corresponding Authors:** andreas.salzburger@cern.ch, beom.ki.yeo@cern.ch, attila.krasznahorkay@cern.ch, joana.niermann@cern.ch, stephen.nicholas.swatman@cern.ch, heather.gray@cern.ch

GPUs are expected to be a key solution to the data challenges posed by track reconstruction in future high energy physics experiments. *tracc*, an R&D project within the ACTS track reconstruction toolkit, aims to demonstrate tracking algorithms in GPU programming models including CUDA and SYCL without loss of physical accuracy such as tracking efficiency and fitted parameter resolution. We discuss the current status and demonstrate the performance of the full track reconstruction chain with GPUs for the first time. The physical and computational performance are studied using events simulated in the Open Data Detector, which is an open-source tracking geometry. The benchmark result shows that a GPU is faster than a CPU for pp collision events with pileups higher than 140, corresponding to the data size of the HL-LHC. We also explore its potential as an experimental-independent toolkit for other high energy physics experiments such as ATLAS and CEPC.

**Poster session / 363**

## **Packaging HEP heterogeneous mini-apps for portable benchmarking and facility evaluation on modern HPCs**

**Authors:** Charles Leggett<sup>1</sup>; FNU Mohammad Atif<sup>2</sup>; Ka Hei Martin Kwok<sup>3</sup>; Pengfei Ding<sup>None</sup>

<sup>1</sup> *Lawrence Berkeley National Lab (US)*

<sup>2</sup> *Brookhaven National Laboratory*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** ka.hei.martin.kwok@cern.ch, fmohammad@bnl.gov, charles.g.leggett@gmail.com, pding@lbl.gov

High energy physics experiments are making increasing use of GPUs and GPU dominated High Performance Computer facilities. Both the software and hardware of these systems are rapidly evolving, creating challenges for experiments to make informed decisions as to where they wish to devote resources. In its first phase, the High Energy Physics Center for Computational Excellence (HEP-CCE) produced portable versions of a number of heterogeneous HEP mini-apps, such as *p2r*, *FastCaloSim*, *Patatrack* and the *WireCell Toolkit*, that exercise a broad range of GPU characteristics, enabling cross platform and facility benchmarking and evaluation. However, these mini-apps still require a significant amount of manual intervention to deploy on a new facility.

We present our work in developing turn-key deployments of these mini-apps, where by means of containerization and automated configuration and build techniques such as *spack*, we are able to

quickly test new hardware, software, environments and entire facilities with minimal user intervention, and then track performance metrics over time.

Poster session / 364

## HEP-CCE Phase 2 Overview and Outlook

**Authors:** Charles Leggett<sup>1</sup>; Meifeng Lin<sup>None</sup>; Paolo Calafiura<sup>2</sup>; Peter Van Gemmeren<sup>3</sup>; Salman Habib<sup>4</sup>; Seth Johnson<sup>5</sup>; Thomas ds.LDAP.sn Not Supplied<sup>6</sup>; Walter Hopkins<sup>3</sup>

<sup>1</sup> *Lawrence Berkeley National Lab (US)*

<sup>2</sup> *Lawrence Berkeley National Lab. (US)*

<sup>3</sup> *Argonne National Laboratory (US)*

<sup>4</sup> *Argonne National Laboratory*

<sup>5</sup> *Oak Ridge National Laboratory (US)*

<sup>6</sup> *Oak Ridge National Laboratory*

**Corresponding Authors:** pcalafiura@lbl.gov, johnsonsr@ornl.gov, charles.g.leggett@gmail.com, peter.van.gemmeren@cern.ch, mlin@bnl.gov, habib@anl.gov, evanstm@ornl.gov, walter.hopkins@cern.ch

In its highly successful first three years, the High Energy Physics Center for Computational Excellence explored the intersectionality of High Energy Physics computing and next generation, heterogeneous High Performance Computing facilities, focussing on issues of portability, data models and I/O, accelerating Event Generators, and exploring complex workflows. In its second phase, we seek to apply lessons learned from Phase 1, and will focus on these subjects:

- **Storage Optimization:** the project will add focus on aiding the adoption of new and more efficient storage technologies such as ROOT7 RNTuple, enabling the deployment of intelligent/lossy compression and reducing data duplication by deploying Object Stores
- **Workflow Portability:** HEP workflows on HPCs are usually very site specific, limiting portability to other facilities. We will investigate portability layers to enable cross facility workflow execution. The portable GPU testbeds developed in CCE Phase 1 will also be containerized for rapid deployment in order to test and benchmark new facilities and hardware.
- **Scaling ML:** we will focus on scaling up the training, inference, and optimization of HEP ML models on HPC resources. These models will be identified in collaboration with experiments and individual physicists. Specifically, we will study and optimize distributed inference using NVIDIA Triton and the HEP-specific SONIC framework based on it. We will compare the performance of various distributed training mechanisms (data parallel initially) on resource-intensive architectures like Graph networks.
- **Simulation:** Detector simulation is focused on developing GPU-accelerated, platform portable capabilities for experiments. The two objectives are to integrate new optical photon transport into experiment frameworks for luminescent detectors, and to develop a new geometry representation, ORANGE, optimized for GPU particle navigation and ray tracing. The optical photon throughput improvements will be demonstrated in the LUX-ZEPLIN experiment, targeting a factor of 20–25× for the 2nd generation and 50× for the 3rd generation, using the new platform portable geometry.

<br>

In this presentation we will report on the current project status and outlook for HEP/CCE in its second phase.

Parallel (Track 2) / 365

## Reconstruction Framework Advancements for Streaming Readout for the ePIC Experiment at the EIC

**Author:** Nathan Brei<sup>1</sup>

**Co-author:** David Lawrence<sup>1</sup>

<sup>1</sup> *Thomas Jefferson National Accelerator Facility*

**Corresponding Authors:** davidl@jlab.org, nbrei@jlab.org

The ePIC collaboration adopted the JANA2 framework to manage its reconstruction algorithms. This framework has since evolved substantially in response to ePIC's needs. There have been three main design drivers: integrating cleanly with the PODIO-based data models and other layers of the key4hep stack, enabling external configuration of existing components, and supporting timeframe splitting for streaming readout. The result is a unified component model featuring a new declarative interface for specifying inputs, outputs, parameters, services, and resources. This interface enables the user to instantiate, configure, and wire components via an external file. One critical new addition to the component model is a hierarchical decomposition of data boundaries into levels such as Run, Timeframe, PhysicsEvent, and Subevent. Two new component abstractions, Folder and Unfolder, are introduced in order to traverse this hierarchy, e.g. by splitting or merging. The pre-existing components can now operate at different event levels, and JANA2 will automatically construct the corresponding parallel processing topology. This means that a user may write an algorithm once, and configure it at runtime to operate on timeframes or on physics events. Overall, these changes mean that the user requires less knowledge about the framework internals, obtains greater flexibility with configuration, and gains the ability to reuse the existing abstractions in new streaming contexts.

Parallel (Track 3) / 366

## Particle Identification at STCF DTOF Detector Based on Classical/Quantum Convolutional Neural Network

**Author:** Zhipeng Yao<sup>None</sup>

**Co-authors:** Teng LI<sup>1</sup>; Xingtao Huang

<sup>1</sup> *Shandong University, CN*

**Corresponding Authors:** tengli@sdu.edu.cn, yaozp@mail.sdu.edu.cn, huangxt@sdu.edu.cn

The Super Tau Charm Facility (STCF) is a future electron-positron collider proposed with a center-of-mass energy ranging from 2 to 7 GeV and a peak luminosity of  $0.5 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$ . In STCF, the identification of high-momentum hadrons is critical for various physics studies, therefore two Cherenkov detectors (RICH and DTOF) are designed to boost the PID performance. In this work, targeting the pion/kaon identification at STCF, we developed a PID algorithm based on the convolutional neural network (CNN) for the DTOF detector, which combines the hit channel and arrival time of Cherenkov photons at multi-anode microchannel plate photomultipliers. The current performance meets the physics requirements of STCF, with a pion identification efficiency exceeding 97% along with a kaon misidentification rate of less than 2% at  $p = 2\text{GeV}/c$ . In addition, based on classical CNN, we conducted a proof-of-concept study on quantum convolutional neural networks (QCNN) to explore potential quantum advantages and feasibility. Preliminary results indicate that QCNN has a promising potential to outperform classical CNN on a same dataset.

Parallel (Track 3) / 367

## Anomaly Detection using Autoencoders on Fundamental LZ Signals



**Author:** Tyler Anderson<sup>1</sup>

<sup>1</sup> SLAC

**Corresponding Author:** tylerja@stanford.edu

Searching for anomalous data is especially important in rare event searches like that of the LUX-ZEPLIN (LZ) experiment's hunt for dark matter. While LZ's data processing provides analyzer-friendly features for all data, searching for anomalous data after minimal reconstruction allows one to find anomalies which may not have been captured by reconstructed features and allows us to avoid any reconstruction errors. Autoencoders can be used to probe for anomalous light-detecting PMT waveforms resulting from ionization signals (S2) and have found unresolved S2s resulting from multiple scatter interactions. In addition to comparing results to waveform-shape template-fitting methods, these techniques can be extended by applying them to PMT waveforms from prompt scintillation light (S1) and S2 heatmaps which capture positional information. Results from such methods are discussed and compared to known anomalies.

**Parallel (Track 1) / 368**

## Enhancing CMS XCache efficiency: A comparative study of Machine Learning techniques and LRU mechanisms

**Author:** Jose Flix Molina<sup>1</sup>

**Co-authors:** Anna Sikora<sup>2</sup>; Antonio Delgado Peris<sup>1</sup>; Antonio Perez-Calero Yzquierdo<sup>3</sup>; Francisco Javier Rodriguez-Calonge ; Jose Hernandez<sup>4</sup>; Paula Serrano Sierra<sup>2</sup>

<sup>1</sup> CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)

<sup>2</sup> UAB

<sup>3</sup> Centro de Investigaciones Energéticas Medioambientales y Tecnológicas

<sup>4</sup> CIEMAT

**Corresponding Authors:** antonio.delgado.peris@cern.ch, calonge@ciemat.es, paula.serranos@autonoma.cat, anna.sikora@uab.cat, antonio.perez.calero.yzquierdo@cern.ch, jose.flix.molina@cern.ch, jose.hernandez.calama@cern.ch

The Large Hadron Collider (LHC) at CERN in Geneva is preparing for a major upgrade that will improve both its accelerator and particle detectors. This strategic move comes in anticipation of a tenfold increase in proton-proton collisions, expected to kick off by 2029 in the upcoming high-luminosity phase. The backbone of this evolution is the World-Wide LHC Computing Grid, crucial for handling the flood of data from these collisions. Therefore, expanding and adapting it is vital to meet the demands of the new phase, all while working within a tight budget. Many research and development projects are in progress to keep future resources manageable and cost-effective in managing the growing data. One area of focus is Content Delivery Network (CDN) techniques, which promise data access and resource use optimization, improving task performance by caching input data close to users. A comprehensive study has been conducted to assess how beneficial it would be to implement data caching for the Compact Muon Solenoid (CMS) experiment. This study, with a focus on Spanish computing facilities, shows that user analysis tasks are the ones that can benefit the most from CDN techniques. As a result, a data cache has been introduced in the region to understand these benefits better. In this contribution, we analyze remote data access from users in Spanish CMS sites to figure out the best size and network connectivity requirements for a data cache serving the whole Spanish region. Exploration of machine learning techniques, along with comparisons to traditional LRU mechanisms, allow for the identification and preservation of frequently accessed datasets within the cache. This approach aims to optimize storage usage efficiently, while prioritizing accessibility to the most popular data.

**Parallel (Track 5) / 370**

## R&D Adoption and Progress in Full Simulation of the CMS experiment

**Authors:** Phat Srimanobhas<sup>1</sup>; Vladimir Ivantchenko<sup>2</sup>

**Co-authors:** Kevin Pedro<sup>3</sup>; Malik Shahzad Muzaffar<sup>2</sup>; Natascha Krammer<sup>4</sup>; Sunanda Banerjee<sup>5</sup>

<sup>1</sup> *Chulalongkorn University (TH)*

<sup>2</sup> *CERN*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

<sup>4</sup> *Austrian Academy of Sciences (AT)*

<sup>5</sup> *University of Wisconsin Madison (US)*

**Corresponding Authors:** vladimir.ivantchenko@cern.ch, pedrok@cern.ch, natascha.hoermann@cern.ch, sunanda.banerjee@cern.ch, norraphat.srimanobhas@cern.ch, shahzad.malik.muzaffar@cern.ch

For the start of Run-3 CMS Full Simulation was based on Geant4 10.7.2. In this work we report on evolution of usage of Geant4 within CMSSW and adaptation of the newest Geant4 11.2.1, which is expected to be used for CMS simulation production in 2025. Physics validation results and results on CPU performance are reported.

For the Phase-2 simulation several R&D are carried out. A significant update for CMS geometry description is performed using DD4hep and VecGeom tools, modifications of the CMS geometry concern a new tracker, a new timing detector, an extended muon system, and a new endcap high granular calorimeter. Different aspects of geometry description and physics simulation for the new detectors will be discussed. Progress on R&D efforts for the Phase-2 simulation will be presented, which includes reports on experience of application of G4HepEm, Seleritas, and AdePT external libraries.

**Poster session / 371**

## Time Based Clustering for Muon Chamber(MuCh) in the CBM experiment at FAIR, Germany

**Authors:** Radha Krishna Garg<sup>1</sup>; Vikas Singhal<sup>2</sup>

**Co-authors:** Chandrasekhar Ghosh<sup>3</sup>; Jogender Saini<sup>2</sup>; Subhasis Chattopadhyay<sup>2</sup>

<sup>1</sup> *Vellore Institute of Technology*

<sup>2</sup> *Department of Atomic Energy (IN)*

<sup>3</sup> *Variable Energy Cyclotron Centre*

**Corresponding Authors:** subhasis.chattopadhyay@cern.ch, jogender.saini@cern.ch, krishna.21bce7371@vitapstudent.ac.in, vikas.singhal@cern.ch, c.ghosh@vecc.gov.in

The soon-to-be-realised Compressed Baryonic Matter (CBM) experiment at the Facility for Antiproton and Ion Research (FAIR) is engineered to handle data acquisition during nuclear collisions at an exceptional high interaction rate, reaching up to 10 MHz. This state-of-the-art setup incorporates a self-triggered readout system, eliminating the need for conventional hardware triggers. It creates a new challenge for time based clustering for the Muon Chamber (MuCh) subsystem of CBM. Our paper delves into the first pivotal process of the reconstruction, namely the formation of clusters and hits from digitized data “digis” gathered from CBM data acquisition system. These clusters emerge from a meticulous fusion of temporal and spatial characteristics. The work is dedicated to an exclusive exploration of the methodology and the pragmatic execution of cluster formation for MuCh within the CBM experiment, highlighting its profound importance in and pertinence to this specialized experimental framework. This paper highlights the advancements for the Muon Chamber (MuCh) clustering algorithm, which is utilised for organising detector data within the CBM setup and introducing dynamic adaptations to temporal patterns. The algorithm achieves precise and efficient clustering. Furthermore, enhancements in spatial separation and hit formation contribute to improved accuracy and fidelity in identifying and analysing hits within the detector system to

handle a 10MHz interaction rate. As a result, the algorithmic performance will be derived from simulated Monte Carlo (MC) points across multiple interaction rates varying from  $10^4$  to  $10^7$  with the time based CBM framework.

Poster session / 372

## Di-Muon cocktail reconstruction using Machine Learning technique in CBM experiment at FAIR SIS100 energies

**Author:** Pawan Kumar Sharma<sup>1</sup>

**Co-authors:** Raktim Mukerjee<sup>1</sup>; Abhishek Sharma<sup>2</sup>; Partha Pratim Bhaduri<sup>1</sup>; Apar Agarwal<sup>1</sup>; Anand K Dubey<sup>1</sup>; Subhashish Chattopadhyay<sup>1</sup>

<sup>1</sup> VECC Kolkata

<sup>2</sup> AMU Aligarh

**Corresponding Authors:** parthabhp@gmail.com, raktimmukherjee711@gmail.com, anandkdb@gmail.com, sub.chattopadhyay@gmail.com, apar.2010big@gmail.com, pawan75sikhwal1998@gmail.com

The CBM experiment at FAIR-SIS100 will investigate strongly interacting matter at high baryon density and moderate temperature. One of proposed key observable is the measurement of the low mass vector mesons(LMVMs), which can be detected via their di-lepton decay channel. As the decayed leptons leave the dense and hot fireball without further interactions, they can provide unscathed information about the fireball, produced in energetic nuclear collisions.

We report, simulation results for the reconstruction of di-muon continuum spectra for AuAu 8AGeV central collisions using machine learning(ML) techniques for selection of muon track candidates. The results from various ML models have been compared with the traditional selection cuts for  $\omega(\omega)$ ,  $\eta(\eta)$ ,  $\phi(\phi)$ ,  $\rho(\rho)$  mesons and full di-muon cocktail spectra.

We have attempted to reconstruct LMVM ( $\omega, \eta, \phi, \rho$ ) in the event by event mode using standard reconstruction software. Background of central Au-Au collisions at 8 AGeV was generated using UrQMD event generator, whereas for LMVMs signals PLUTO event generator was used. Single LMVM decaying into  $\mu^+ + \mu^-$  was embedded into each background event. The particles are then transported through the experimental setup including upgraded Muon Chamber(MuCh) setup, using the GEANT3 transport engine. Various ML algorithms like Gradient boosted decision trees (BDTG), KNN, MLP, HMatrix etc. from the TMVA class have been employed for the present study. Based on the the simulation results, improvement in di-muon performance is reported. For comparable S/B ratio, the pair reconstruction efficiency and significance is observed to be increased substantially for  $\omega, \eta, \phi$  mesons using ML techniques.

Parallel (Track 1) / 373

## DUNE Rucio development and monitoring

**Authors:** Brandon White<sup>1</sup>; James Perry<sup>None</sup>; Wenlong Yuan<sup>2</sup>

<sup>1</sup> FNAL

<sup>2</sup> The University of Edinburgh (GB)

**Corresponding Authors:** wenlong.yuan@cern.ch, bjwhite@fnal.gov, j.perry@epcc.ed.ac.uk

The Deep Underground Neutrino Experiment (DUNE) is scheduled to start running in 2029, expected to record 30 PB/year of raw data. To handle this large-scale data, DUNE has adopted and deployed

Rucio, the next-generation Data Replica service originally designed by the ATLAS collaboration, as an essential component of its Distributed Data Management system.

DUNE's use of Rucio has demanded the addition of various features to the Rucio code base, both specific functionality for DUNE alone, and more general functionality that is crucial for DUNE whilst being potentially useful for other experiments. As part of our development work, we have introduced a "policy package" system allowing experiment-specific code to be maintained separately from the core Rucio code, as well as creating a DUNE policy package containing algorithms such as logical to physical filename translation, and special permission checks. We have also developed other features such as improved object store support, and customisable replica sorting. A DUNE-specific test suite that will run on GitHub Actions is currently under development.

Recently, DUNE has deployed new internal monitoring to Rucio, enabling us to extract more useful information from core Rucio servers, and daemons such as transmogrifier, reaper, etc. Additionally, DUNE has implemented monitoring for Rucio transfer and deletion activities which are sent to a Message Queue via Rucio Hermes daemon. Information such as data location, accounting, and storage summary is extracted from the Rucio internal database and dumped into Elasticsearch for visualisation. The visualisation platforms utilised are based at Fermilab and Edinburgh. This monitoring is crucial for the ongoing DUNE data transfers and management development.

**Parallel (Track 3) / 374**

## **GPU-Accelerated Point Clouds Library for GNN-based Reconstruction and Beyond**

**Author:** Shah Rukh Qasim<sup>1</sup>

**Co-authors:** Jan Kieseler<sup>2</sup>; Nicola Serra<sup>1</sup>; Oleksandr Viazlo

<sup>1</sup> *University of Zurich (CH)*

<sup>2</sup> *KIT - Karlsruhe Institute of Technology (DE)*

**Corresponding Authors:** oleksandr.viazlo@cern.ch, shah.rukh.qasim@cern.ch, nicola.serra@cern.ch, jan.kieseler@cern.ch

High quality particle reconstruction is crucial to data acquisition at large CERN experiments. While the classical algorithms have been successful so far, in recent years, the use of pattern recognition has become more and more necessary due to increasing complexity of the modern detectors. Graph Neural Network based approaches have been recently proposed to tackle challenges such as non-uniformity and high level of sparsity. They have been shown to work well not only for calorimetric reconstruction<sup>1</sup> but also for tracking<sup>2</sup>. These GNN based approaches require fast GPU execution of certain operations (such as kNN computation). We present a fast GPU-enabled software library which significantly outperforms existing approaches including those present in large machine learning frameworks (such as TensorFlow and PyTorch). In light of these optimizations, we discuss the performance and computational requirements of GNN based reconstruction algorithms. Finally, we also discuss how the application of these techniques can be extended beyond particle physics.

<sup>1</sup> Qasim, Shah Rukh, et al. "End-to-end multi-particle reconstruction in high occupancy imaging calorimeters with graph neural networks." *The European Physical Journal C* 82.8 (2022): 1-15.

<sup>2</sup> <https://indico.jlab.org/event/459/contributions/11761/>

**Parallel (Track 2) / 375**

## **Performance of the parallelized General Triplet Track Fit implemented on the GPU**

**Authors:** Abhirikshma Nandi<sup>1</sup>; André Schöning<sup>1</sup>; Christof Sauer<sup>1</sup>; Sebastian Dittmeier<sup>2</sup>

<sup>1</sup> Heidelberg University (DE)

<sup>2</sup> Ruprecht-Karls-Universitaet Heidelberg (DE)

**Corresponding Authors:** schoning@mail.desy.de, christof.sauer@cern.ch, sebastian.dittmeier@cern.ch, abhirikshma.nandi@cern.ch

The General Triplet Track Fit (GTTF) is a generalization of the Multiple Scattering Triplet Fit [NIMA 844 (2017) 135] to additionally take hit uncertainties into account. This makes it suitable for use in collider experiments, where the position uncertainties of hits dominate for high momentum tracks. Since the GTTF is based on triplets of hits that can be processed independently, the fit is particularly suitable for acceleration with parallel hardware such as GPUs, and can therefore be used for fast track fitting in online reconstruction. The performance of the track fit and its acceleration is studied using the OpenDataDetector in tracc, a demonstrator tracking chain designed for hardware accelerators under the umbrella of the ACTS track reconstruction framework, and the results will be presented.

Parallel (Track 7) / 376

## Efficient and fast container execution using image snapshotters

**Authors:** Amal Santosh Thundiyil<sup>1</sup>; Clemens Lange<sup>2</sup>; Derek Feichtinger<sup>2</sup>; Jakob Blomer<sup>3</sup>; Max Fatouros<sup>2</sup>; Valentin Volk<sup>3</sup>

<sup>1</sup> Sardar Patel Institute Technology of Technology (SPIT) (IN)

<sup>2</sup> Paul Scherrer Institute (CH)

<sup>3</sup> CERN

**Corresponding Authors:** max.fatouros@psi.ch, derek.feichtinger@cern.ch, clemens.lange@cern.ch, valentin.volk@cern.ch, amal.santosh.thundiyil@cern.ch, jakob.blomer@cern.ch

A large fraction of computing workloads in high-energy and nuclear physics is executed using software containers. For physics analysis use, such container images often have sizes of several gigabytes. Executing a large number of such jobs in parallel on different compute nodes efficiently, demands the availability and use of caching mechanisms and image loading techniques to prevent network saturation and significantly reduce startup time. Using the industry-standard containerd container runtime for pulling and running containers, enables the use of various so-called snapshotter plugins that “lazily” load container images. We present a quantitative comparison of the performance of the CVMFS, SOCI, and Stargz snapshotter plugins. Furthermore, we also evaluate the user-friendliness of such approaches and discuss how such seamlessly containerised workloads contribute to the reusability and reproducibility of physics analyses.

Poster session / 377

## FPGA implementation of the General Triplet Track Fit

**Authors:** Abhirikshma Nandi<sup>1</sup>; André Schöning<sup>1</sup>; Christof Sauer<sup>1</sup>; Kadir Murat Tastepe<sup>1</sup>; Sebastian Dittmeier<sup>2</sup>

<sup>1</sup> Heidelberg University (DE)

<sup>2</sup> Ruprecht-Karls-Universitaet Heidelberg (DE)

**Corresponding Authors:** schoning@mail.desy.de, christof.sauer@cern.ch, sebastian.dittmeier@cern.ch, abhirikshma.nandi@cern.ch, kadir.tastepe@cern.ch

Online reconstruction of charged particle tracks is one of the most computationally intensive tasks within current and future filter farms of large HEP experiments, requiring clever algorithms and appropriate hardware choices for its acceleration. The General Triplet Track Fit is a novel track-fitting algorithm that offers great potential for speed-up by processing triplets of hits independently. FPGAs, with their inherent parallelism, power efficiency and reconfigurability, are becoming increasingly attractive as co-processors for large data centres, such as the filter farms, to meet the challenges of increasing throughput and computational complexity.

We present an FPGA implementation of the General Triplet Track Fit suitable for future use in heterogeneous online farms. The algorithm is implemented on AMD FPGAs using high-level synthesis. We discuss algorithmic optimisation strategies to exploit the full potential of the device.

**Parallel (Track6) / 378**

## **XKIT for GridPP: (XRootD Kubentes Integration Testing for GridPP)**

**Authors:** Robert Andrew Currie<sup>1</sup>; Wenlong Yuan<sup>1</sup>

<sup>1</sup> *The University of Edinburgh (GB)*

**Corresponding Authors:** rcurrie@cern.ch, wenlong.yuan@cern.ch

XRootD is a robust, scalable service that supports globally distributed data management for diverse scientific communities. Within GridPP in the UK, XRootD is used by the Astronomy, High-Energy Physics (HEP) and other communities to access >100PB of storage. The optimal configuration for XRootD varies significantly across different sites due to unique technological frameworks and site-specific factors.

XRootD's adaptability has made it a cornerstone of the national data-management strategy for GridPP. Given its high-profile role, new releases, and features of XRootD undergo rigorous testing and verification before national deployment. Historically, this process involved manual integration testing and dedicated test deployments, which required substantial input from both local site administrators and remote support teams. This approach has placed considerable demands on support staff, requiring extensive technical expertise and significant time for verification.

To support the storage community within GridPP, we have developed a system that automates the deployment of a virtual grid using Kubernetes for XRootD testing, "XKIT". Using a container-based approach this system enables high-level integration tests to be performed automatically and reproducibly. This not only simplifies the support process but also significantly reduces the time staff need to dedicate to repetitive testing for new deployments.

We have identified >20 unique XRootD configurations necessary for XKIT. By deploying each of these setups on our platform, we aim to provide the GridPP community with a consistent suite of functional tests tailored to various site topologies.

This presentation will explore the development of the XKIT platform, discuss the challenges we encountered, and highlight the advantages this system offers to GridPP and the wider community.

**Parallel (Track 1) / 379**

## **Achieving 100Gb/s data rates with XRootD - Preparing for HL-LHC and SKA**

**Authors:** James William Walder<sup>1</sup>; Jyothish Thomas<sup>2</sup>; Thomas Byrne<sup>None</sup>

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

<sup>2</sup> STFC**Corresponding Authors:** jyothish.thomas@cern.ch, tom.byrne@stfc.ac.uk, james.william.walder@cern.ch

To address the needs of forthcoming projects such as the Square Kilometre Array (SKA) and the HL-LHC, there is a critical demand for data transfer nodes (DTNs) to realise O(100)Gb/s of data movement. This high-throughput can be attained through combinations of increased concurrency of transfers and improvements in the speed of individual transfers. At the Rutherford Appleton Laboratory (RAL), the UK's Tier-1 centre for the Worldwide LHC Computing Grid, and initial site for the UK SKA Regional Centre (SRC), we have provisioned 100GbE XRootD servers in preparation for SKA development and operations. This presentation details the efforts undertaken to reach 100Gb/s data ingress and egress rates using the WebDAV protocol through XRootD endpoints, including the use of a novel XRootD plug-in designed to assess XRootD performance independently of physical storage backend. Results are also presented for transfer tests against a CephFS storage backend under different configuration settings (e.g. via tunings to file layouts). We discuss the challenges encountered, bottlenecks identified, and insights gained, along with a description of the most effective solutions developed to date and areas of future activities.

Parallel (Track 5) / 380

## Fair Universe HiggsML Uncertainty Challenge

**Authors:** Aishik Ghosh<sup>1</sup>; Ben Nachman<sup>2</sup>; Christopher Harris<sup>3</sup>; David Rousseau<sup>4</sup>; Ihsan Ullah<sup>5</sup>; Isabelle Guyon<sup>None</sup>; Paolo Calafiura<sup>2</sup>; Peter Nugent<sup>6</sup>; Ragansu Chakkappai<sup>4</sup>; Sascha Diefenbacher<sup>2</sup>; Shih-Chieh Hsu<sup>7</sup>; Steven Farrell<sup>8</sup>; Wahid Bhimji<sup>None</sup>; Yuan-Tang Chou<sup>9</sup>; Yulei Zhang<sup>10</sup>

<sup>1</sup> *University of California Irvine (US)*<sup>2</sup> *Lawrence Berkeley National Lab. (US)*<sup>3</sup> *Unknown*<sup>4</sup> *IJCLab-Orsay*<sup>5</sup> *Chalearn*<sup>6</sup> *Berkeley*<sup>7</sup> *University of Washington Seattle (US)*<sup>8</sup> *Lawrence Berkeley National Laboratory*<sup>9</sup> *University of Washington (US)*<sup>10</sup> *Shanghai Jiao Tong University (CN) & APC-Paris (FR)*

**Corresponding Authors:** ragansu.chakkappai@cern.ch, pcalafiura@lbl.gov, wbhimji@lbl.gov, schsu@uw.edu, yuan-tang.chou@cern.ch, sfarrell@lbl.gov, benjamin.philip.nachman@cern.ch, christopher.harris@cern.ch, yulei.zhang@cern.ch, guyon@clopinnet.com, rousseau@lal.in2p3.fr, sdiefenbacher@lbl.gov, aishik.ghosh@cern.ch

The Fair Universe project is organising the HiggsML Uncertainty Challenge, which will/has run from June to October 2024.

This HEP and Machine Learning competition is the first to strongly emphasise uncertainties: mastering uncertainties in the input training dataset and outputting credible confidence intervals.

The context is the measurement of the Higgs to tau+ tau- cross section like in HiggsML challenge on Kaggle in 2014, from a dataset of the 4-momentum signal state. Participants should design an advanced analysis technique that can not only measure the signal strength but also provide a confidence interval, from which correct coverage will be evaluated automatically from pseudo-experiments.

The confidence interval should include statistical and systematic uncertainties (concerning detector calibration, background levels, etc...). It is expected that advanced analysis techniques that can control the impact of systematics will perform best, thereby pushing the field of uncertainty-aware AI techniques for HEP and beyond.

The challenge is hosted on Codabench (an evolution of the popular Codalab platform); the significant resources needed (to run the thousands of pseudo-experiments needed) are possible thanks to using NERSC infrastructure as a backend.

The competition will have ended just before CHEP 2024 so that a first glimpse of the competition results could be made public for the first time.

**Parallel (Track 4) / 381**

## **The path to exabyte astronomy: SRCNet v0.1 for the Square Kilometre Array**

**Authors:** Ian Collier<sup>1</sup>; James William Walder<sup>1</sup>

**Co-author:** SRCNet Members <sup>2</sup>

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

<sup>2</sup> *SRCNet*

**Corresponding Authors:** [ian.peter.collier@cern.ch](mailto:ian.peter.collier@cern.ch), [james.william.walder@cern.ch](mailto:james.william.walder@cern.ch)

The Square Kilometre Array (SKA) is set to be the largest and most sensitive radio telescope in the world. As construction advances, the managing and processing of data on an exabyte scale becomes a paramount challenge to enable the SKA science community to process and analyse their data. To address this, the SKA Regional Centre Network (SRCNet) has been established to provide the necessary computational and storage resources, and to facilitate access to SKA data for the global scientific community, providing a science archive for public data engagement.

SRCNet v0.1 marks the prototype deployment phase of this architecture, scheduled for early 2025, and is a critical step in the roadmap towards achieving full operational capability of SRCNet. This report outlines the specific requirements for SRCNet v0.1, including major architectural components and a review of the Distributed Data Management (DDM) software selection process, which includes tooling originally developed for the HEP Community, namely: Rucio, FTS, and perfSONAR and Storage Endpoint technologies. The primary goals of SRCNet v0.1 are presented, including topics of authentication and authorisation and interactive data access, and the planned data challenge campaigns.

**Parallel (Track 2) / 382**

## **High-throughput data distribution for CBM online computing**

**Author:** Jan de Cuveland<sup>1</sup>

**Co-authors:** Dirk Hutter <sup>1</sup>; Volker Lindenstruth <sup>1</sup>

<sup>1</sup> *Goethe University Frankfurt (DE)*

**Corresponding Authors:** [voli@compeng.de](mailto:voli@compeng.de), [jan.de.cuveland@cern.ch](mailto:jan.de.cuveland@cern.ch), [hutter@compeng.uni-frankfurt.de](mailto:hutter@compeng.uni-frankfurt.de)

The CBM experiment, currently being constructed at GSI/FAIR, aims to investigate QCD at high baryon densities. The CBM First-level Event Selector (FLES) serves as the central event selection system of the experiment. It functions as a high-performance computer cluster tasked with the online analysis of physics data, including full event reconstruction, at an incoming data rate which exceeds 1 TByte/s.

The CBM detector systems operate in a free-running and self-triggered manner, delivering time-stamped data streams. Without inherent event separation, timeslice building replaces global event



building. The FLES HPC system integrates data from around 5000 input links into self-contained, overlapping processing intervals and distributes these to the compute nodes.

Using a combination of RDMA and zero-copy techniques, timeslices can be built efficiently over a high-throughput InfiniBand network and distributed to available online computing resources for a full online event reconstruction and analysis in a heterogeneous HPC cluster system. A new IPC online interface to timeslice data utilizes a Posix shared memory governed by a reference-counting item distributor. This design combines maximum performance and flexibility with minimum memory consumption. These new developments have already been successfully field-tested in production at the CBM predecessor experiment mCBM at the GSI/FAIR SIS18.

This work is supported by BMBF (05P21RFFC1).

#### Poster session / 383

## Scientific NREC Cluster - a fast and flexible solution for virtual computing clusters on OpenStack

**Author:** Emmanuel Moutoussamy<sup>1</sup>

**Co-authors:** Alexander Oltu<sup>1</sup>; Matthias Richter<sup>2</sup>; Tomasz Stokowy<sup>1</sup>

<sup>1</sup> *University of Bergen, Norway*

<sup>2</sup> *University of Bergen (NO)*

**Corresponding Authors:** matthias.richter@cern.ch, tomasz.stokowy@uib.no, alexander.oltu@uib.no, emmanuel.moutoussamy@uib.no

Research groups at scientific institutions have an increasing demand for computing and storage resources. The national High-Performance Computing (HPC) systems usually have a high threshold to come in and cloud solutions could be challenging and demand a high learning curve.

Here we introduce the Scientific NREC Cluster (SNC), which leverages the Norwegian Research and Education Cloud (NREC). NREC operates on an Infrastructure-as-a-Service (IaaS) model, offering users full control over host administration, installation, and upgrade options for provided virtual instances. The SNC project aims to bridge the gap between the foundational NREC infrastructure and user requirements, providing easy access to flexible cluster resources and storage while maintaining elevated levels of security and flexibility. Nevertheless, the solution is built from the data-centric point of view, where it gives easy access to the campus storage.

SNC offers a SLURM queueing system with the following functionalities: access to the central, shared and secure storage solution; centralized user authentication from the campus Active Directory; provisioning based on NREC's Infrastructure-as-a-Service policy; monitoring solution with the Prometheus/Grafana ecosystem, including both metrics and log messages of the actual user jobs; access to the scientific software stack EESSI - European Environment for Scientific Software Installations.

As an initial release, SNC has been launched as a compact solution tailored for research groups at the University of Bergen. Examples of user stories, usability, and scaling studies will be highlighted in the presentation.

#### Parallel (Track 2) / 384

## The NextGen Triggers project: overview, plans and first actions

**Authors:** Alberto Di Meglio<sup>1</sup>; Alberto Pace<sup>1</sup>; Axel Naumann<sup>1</sup>; Cristina Botta<sup>1</sup>; Giovanna Lehmann Miotto<sup>1</sup>; Marco Rovere<sup>1</sup>; Markus Elsing<sup>1</sup>; Michelangelo Mangano<sup>1</sup>; Stefano Veneziano<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> INFN e Università Roma Sapienza

**Corresponding Authors:** alberto.di.meglio@cern.ch, michelangelo.mangano@cern.ch, giovanna.lehmann@cern.ch, stefano.veneziano@cern.ch, markus.elsing@cern.ch, marco.rovere@cern.ch, alberto.pace@cern.ch, cristina.botta@cern.ch, axel.naumann@cern.ch

The Next Generation Triggers project (NextGen in short) is a five-year collaboration across ATLAS and CMS (with contributions from LHCb and ALICE) and the Experimental Physics, Theoretical Physics, and Information Technology Departments of CERN to research and develop new ideas and technologies for the experiment trigger systems for HL-LHC and beyond. After more than a year of preparation in 2022-2023, the project started in January 2024 and involves the effort of more than 100 researchers and engineers working on four interacting areas: (1) online data processing, modern computing architectures, novel algorithmic concepts, machine learning and the direct interplay of experimental approaches and theory simulation; (2) enhancing the ATLAS trigger and data acquisition to focus on improved and accelerated filtering and exotic signature detection; (3) rethinking the CMS real-time data processing to design a novel AI-powered real-time processing workflow to analyze every single collision produced in the LHC; and (4) designing novel education and training programmes to support the experiment research plans. Investigations of the use of ML from front-end systems inference to development of workflows for large-scale training on local, cloud and HPC systems are among the objectives of the project. Explorations of novel quantum-inspired methods for event generators, optimization in data structures, compression, processes and pipelines are all areas in scope of the research programme. This presentation describes the overall concepts and objectives of the project and the preliminary results and lessons learned after NextGen's first 10 months.

Parallel (Track 7) / 385

## Allocating Carbon Costs to Computing Payloads across Heterogeneous Infrastructures.

**Author:** Richard Alexander Owen<sup>None</sup>

**Co-authors:** Daniel Peter Traynor ; Daohai Li ; Dimitrios Toroz ; Jonathan Hays <sup>1</sup>

<sup>1</sup> University of London (GB)

**Corresponding Authors:** daohai.li@qmul.ac.uk, d.traynor@qmul.ac.uk, d.toroz@qmul.ac.uk, r.a.owen@qmul.ac.uk, j.hays@qmul.ac.uk

As UKRI moves towards a NetZero Digital Research Infrastructure 1 an understanding of how carbon costs of computing infrastructures can be allocated to individual scientific payloads will be required. The IRIS community 2 forms a multi-site heterogenous infrastructure so is a good testing ground to develop carbon allocation models with wide applicability.

The IRISCAST Project [3,4] developed methods to measure carbon costs for a facility. Building on that work the IRIS Carbon Mapping Project [5] has developed models to allocate carbon costs to individual payloads. These models were developed with a learning by doing approach and have been applied to both batch and cloud resources. We present our key findings, lessons learned, and recommendations.

1 <http://doi.org/10.5281/zenodo.8199984>

2 <https://iris.ac.uk>

3 <http://doi.org/10.5281/zenodo.7692451>

[4] <http://doi.org/10.1051/epjconf/202429507029>

[5] <http://doi.org/10.5281/zenodo.10966001>

Parallel (Track 5) / 386

## GlitchFlow, a Digital Twin for transient noise in Gravitational Wave Interferometers

**Author:** Federica Legger<sup>1</sup>

**Co-authors:** Francesco Sarandrea<sup>2</sup>; Lorenzo Asprea<sup>3</sup>; Sara Vallero

<sup>1</sup> *Università e INFN Torino (IT)*

<sup>2</sup> *Università di Torino*

<sup>3</sup> *INFN*

**Corresponding Authors:** svallero@to.infn.it, francesco.sarandrea@unito.it, federica.legger@cern.ch, lorenzo.asprea@gmail.com

Gravitational Waves (GW) were first predicted by Einstein in 1918, as a consequence of his theory of General Relativity published in 1915. The first direct GW detection was announced in 2016 by the LIGO and Virgo collaborations. Both experiments consist of a modified Michelson-Morley interferometer that can measure deformations of the interferometer arms of about 1/1,000 the width of a proton. The sensitivity of GW interferometers is limited by noise. Non-Gaussian transient noise artifacts, also known as glitches, are particularly challenging due to their similarity with astrophysical signals in the time and frequency domains. Noise reduction and subtraction is one of the most important and challenging activities in GW research.

InterTwin is a EU-funded project with the aim of building Digital Twins (DT) for various scientific use cases based on a co-designed blueprint architecture. Within InterTwin, we are developing GlitchFlow, a pipeline for modeling and generating glitches for GW interferometers using deep generative algorithms. In this contribution, we present results of the glitch generation using several Neural Network (NN) architectures, and describe the implementation of the pipeline as execution of DAGs (Directed Acyclic Graph) with an Apache Airflow instance deployed on Kubernetes. We show how the most computing intensive tasks such as model training can be off-loaded to Vega, the EuroHPC, using InterLink, a module developed within the InterTwin framework.

**Poster session / 387**

## Design and construction of High Energy Photon Source (HEPS) scientific data storage system

**Author:** Yaosong Cheng<sup>1</sup>

<sup>1</sup> *Institute of High Energy Physics Chinese Academy of Sciences, IHEP*

**Corresponding Author:** chengys@ihep.ac.cn

The High Energy Photon Source (HEPS) in China will become one of the world's fourth-generation synchrotron light sources with the lowest emittance and highest brightness. The 14 beamlines for the phase I of HEPS will produce about 300PB/year raw data, posing significant challenges in data storage, data access, and data exchange. In order to balance the cost-effectiveness of storage devices and realize the high reliability of data storage, a three-tier storage is designed for storing experimental data, including beamline storage, central storage, and tape. Raw data and processed data are stored on the beamline storage for a maximum of 7 days, on the central storage for a maximum of 90 days, and only the raw data are archived to tape for long-term storage with two copies. Of course, this data storage policy could be adjusted according to the actual data volume and funding situation of HEPS. The beamline storage utilizes a distributed all-flash SSD array to achieve high data input/output speeds. The central storage utilizes a distributed high-density HDD array to achieve medium to high-speed data IO. The tape storage complies with the LTO9 standard.

In addition, we have conducted some personalized optimizations based on the requirements of the HEPS project, such as adapting the Lustre file system to the Roce network protocol, and mapping permissions for users from AD domain control and LDAP domain control.

**Parallel (Track 7) / 388**

## Software defined network control for LHC Experiments

**Authors:** Aashay Arora<sup>1</sup>; Diego Davila Foyo<sup>1</sup>; Justas Balcas<sup>2</sup>; Thomas Lehman<sup>3</sup>; Xi Yang<sup>4</sup><sup>1</sup> *Univ. of California San Diego (US)*<sup>2</sup> *California Institute of Technology (US)*<sup>3</sup> *ESnet*<sup>4</sup> *LBNL***Corresponding Authors:** tlehman@es.net, xiyang@es.net, justas.balcas@cern.ch, aashay.arora@cern.ch

The Large Hadron Collider (LHC) experiments rely on a diverse network of National Research and Education Networks (NRENs) to distribute their data efficiently. These networks are treated as “best-effort” resources by the experiment data management systems. Following the High Luminosity upgrade, the Compact Muon Solenoid (CMS) experiment is projected to generate approximately 0.5 exabytes of data annually. The seamless operation of NRENs is crucial for the success of CMS and other LHC experiments. However, challenges arise during data movement as NRENs lack awareness of data transfer priorities, importance, or quality of service requirements and NRENs operators can not ensure predictable data flow rates across multi-domain networks.

Our work focuses on SENSE, The Software-defined network for End-to-end Networked Science at Exascale, and Rucio, data management software used by multiple experiments, to allocate and prioritize specific data transfers across the wide area network. In this paper, we will showcase our advancements since the last publication, sharing insights gained and detailing the enhancements made to the software stack. These improvements enable science experiments to treat networks as first-class citizens and effectively utilize, prioritize, and manage wide area networks to sites.

**Parallel (Track 1) / 389**

## Distributed Data Management with Rucio for the Einstein Telescope

**Authors:** Andres Jorge Tanasijczuk<sup>1</sup>; Federica Legger<sup>2</sup>; Gabriele Gaetano Fronze<sup>3</sup>; Lia Lavezzi<sup>None</sup>; et al.<sup>None</sup><sup>1</sup> *Universite Catholique de Louvain (UCL) (BE)*<sup>2</sup> *Universita e INFN Torino (IT)*<sup>3</sup> *INFN Torino (IT) and LIGO-Virgo-Kagra Collaboration (US/IT/JP)***Corresponding Authors:** federica.legger@cern.ch, andres.jorge.tanasijczuk@cern.ch, gabriele.gaetano.fronze@cern.ch, lia.lavezzi@to.infn.it

Modern physics experiments are often led by large collaborations including scientists and institutions from different parts of the world. To cope with the ever increasing computing and storage demands, computing resources are nowadays offered as part of a distributed infrastructure. Einstein Telescope (ET) is a future third-generation interferometer for gravitational wave (GW) detection, and is currently in the process of defining a computing model to sustain ET physics goals. A critical challenge for present and future experiments is an efficient and reliable data distribution and access system. Rucio is a framework for data management, access and distribution. It was originally developed by the ATLAS experiment and has been adopted by several collaborations within the high energy physics domain (CMS, Belle II, DUNE) and outside (ESCAPE, SKA, CTA). In the GW community Rucio is used by the second-generation interferometers LIGO and Virgo, and is currently being evaluated for ET. ET will observe a volume of the Universe about one thousand times larger than LIGO and Virgo, and this will reflect on a larger data acquisition rate. In this contribution, we briefly describe Rucio usage in current GW experiments, and outline the on-going R&D activities for integration of Rucio within the ET computing infrastructure, which include the setup of an ET Data Lake based on Rucio for future Mock Data Challenges. We discuss the customization of Rucio features for the GW community: in particular we describe the implementation of RucioFS, a POSIX-like

filesystem view to provide the user with a more familiar structure of the Rucio data catalogue, and the integration of the ET Data Lake with mock Data Lakes belonging to other experiments within the astrophysics and GW communities. This is a critical feature for astronomers and GW data analysts since they often require access to open data from other experiments for sky localisation and multi-messenger analysis.

Parallel (Track 5) / 390

## GPU-friendly surface model for Monte-Carlo detector simulations

**Authors:** Severin Diederichs<sup>1</sup>; Andrei Gheata<sup>1</sup>; Juan Gonzalez Caminero<sup>1</sup>; Guilherme Lima<sup>2</sup>

<sup>1</sup> CERN

<sup>2</sup> FermiLab (US)

**Corresponding Authors:** guilherme.lima@cern.ch, andrei.gheata@cern.ch, j.gonzalez@cern.ch, severin.diederichs@cern.ch

The demands for Monte-Carlo simulation are drastically increasing with the high-luminosity upgrade of the Large Hadron Collider, and expected to exceed the currently available compute resources. At the same time, modern high-performance computing has adopted powerful hardware accelerators, particularly GPUs. AdePT is one of the projects aiming to address the demanding computational needs by leveraging these heterogeneous compute architectures. While AdePT has successfully ported realistic detector simulations to GPUs using the VecGeom library, the complexity of geometry modeling emerged as a bottleneck. Thread divergence and high register usage were impeding the GPU performance. Therefore, a new, GPU-friendly surface-based model has been introduced in the VecGeom library that decomposes the divergent code of the 3D primitive solids into simpler and more balanced surface algorithms. In this work, we present the latest performance results, in particular on complex setups like the CMS Phase-2 geometry. Additionally, we explore techniques such as mixed precision and bounding volume hierarchies to further accelerate simulations.

Parallel (Track 7) / 391

## Prospects (and Hurdles) of the Computing Facility for FAIR Research

**Author:** Johan Messchendorp<sup>None</sup>

**Corresponding Author:** j.messchendorp@gsi.de

A major milestone has been attained as the civil construction of the Facility for Antiproton and Ion Research (FAIR) nears completion, paving the way for early and first science campaigns in the years ahead. Its computing infrastructure will play a pivotal role in facilitating groundbreaking studies across diverse research communities in atomic, nuclear, hadron, heavy-ion, and accelerator physics. Anchored by its TIER-0 green-IT HPC center, it foresees supporting a large dynamic range in data rates and volumes, including large-scale monolithic experiments with free-streaming data processing schemes, modular setups with moderate data rates, and offline computations dedicated to precision studies in experiment, theory, and accelerator science. The distributed computing model will be based on a federated concept that upholds the F.A.I.R. principles while being seamlessly integrated within EOSC. This presentation delves into the conceptual design of the computing infrastructure, shedding light on both the challenges it faces and the opportunities it presents.

Parallel (Track 1) / 392

## Data Movement Manager (DMM) for the SENSE-Rucio Interoperation Prototype

**Authors:** Aashay Arora<sup>1</sup>; Diego Davila Foyo<sup>1</sup>; Justas Balcas<sup>2</sup>; Thomas Lehman<sup>3</sup>; Xi Yang<sup>4</sup>

<sup>1</sup> *Univ. of California San Diego (US)*

<sup>2</sup> *California Institute of Technology (US)*

<sup>3</sup> *ESnet*

<sup>4</sup> *LBNL*

**Corresponding Authors:** tlehman@es.net, aashay.arora@cern.ch, justas.balcas@cern.ch, diego.davila@cern.ch, xiyang@es.net

The data movement manager (DMM) is a prototype interface between the CERN developed data management software Rucio and the software defined networking (SDN) service SENSE by ESNet. It allows for SDN enabled high energy physics data flows using the existing worldwide LHC computing grid infrastructure. In addition to the key feature of DMM, namely transfer-priority based bandwidth allocation for optimal network usage; it also allows for the identification of the exact cause of underperforming flows using end-to-end monitoring of the data flows by having access to host (network interface) level throughput metrics and transfer-tool (FTS) data transfer job level metrics. This paper describes the design and implementation of DMM.

Parallel (Track 5) / 393

## Declarative paradigms for data analysis description and implementation

**Authors:** Alberto Annovi<sup>1</sup>; Andrea Rizzi<sup>2</sup>; Paolo Mastrandrea<sup>2</sup>; Tommaso Boccali<sup>1</sup>

<sup>1</sup> *INFN Sezione di Pisa*

<sup>2</sup> *Universita & INFN Pisa (IT)*

**Corresponding Authors:** andrea.rizzi@cern.ch, alberto.annovi@pi.infn.it, paolo.mastrandrea@cern.ch, tommaso.boccali@cern.ch

The software toolbox used for “big data” analysis in the last few years is rapidly changing. The adoption of software design approaches able to exploit the new hardware architectures and improve code expressiveness plays a pivotal role in boosting data processing speed, resources optimisation, analysis portability and analysis preservation.

The scientific collaborations in the field of High Energy Physics (e.g. the LHC experiments, the next-generation neutrino experiments, and many more) are devoting increasing resources to the development and implementation of bleeding-edge software technologies in order to cope effectively with always growing data samples, pushing the reach of the single experiment and of the whole HEP community.

The introduction of declarative paradigms in the analysis description and implementation is growing interest and support in the main collaborations. This approach can simplify and speed-up the analysis description phase, support the portability of the analyses among different datasets/experiments and strengthen the preservation and reproducibility of the results.

Furthermore this approach, providing a deep decoupling between the analysis algorithm and back-end implementation, is a key element for present and future processing speed, potentially even with back-ends not existing today.

In the landscape of the approaches currently under study, an activity is ongoing in the ICSC (Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing, Italy) which focuses on the development of a framework characterised by a declarative paradigm for the analysis description and able to operate on datasets from different experiments.

The existing NAIL (Natural Analysis Implementation Language 1) Python package, developed in the context of the CMS data analysis for the event processing, is used as a building base for the

development of a demonstrator able to provide a general and effective interface characterised by a declarative paradigm and targeted to the description and implementation of a full analysis chain for HEP data, with support for different data formats.

Status and development plan of the demonstrator will be discussed.

[1https://indico.cern.ch/event/769263/contributions/3413006/attachments/1840145/3016759/NAIL\\_Project\\_Natural\\_Analysis\\_In](https://indico.cern.ch/event/769263/contributions/3413006/attachments/1840145/3016759/NAIL_Project_Natural_Analysis_In)

**Parallel (Track 4) / 394**

## Optimization of distributed compute resources utilization in the CMS Global Pool

**Author:** Antonio Perez-Calero Yzquierdo<sup>1</sup>

**Co-authors:** Florian Von Cube<sup>2</sup>; Hyunwoo Kim<sup>3</sup>; Marco Mascheroni<sup>4</sup>; Vaiva Zokaite<sup>5</sup>

<sup>1</sup> *Centro de Investigaciones Energéticas Medioambientales y Tecnológicas*

<sup>2</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

<sup>4</sup> *Univ. of California San Diego (US)*

<sup>5</sup> *Vilnius University (LT)*

**Corresponding Authors:** marco.mascheroni@cern.ch, vaiva.zokaite@cern.ch, hyunwoo@fnal.gov, antonio.perez.calero.yzquierdo@cern.ch, ralf.florian.von.cube@cern.ch

Efficient utilization of vast amounts of distributed compute resources is a key element in the success of the scientific programs of the LHC experiments. The CMS Submission Infrastructure is the main computing resource provisioning system for CMS workflows, including data processing, simulation and analysis. Resources geographically distributed across numerous institutions, including Grid, HPC and cloud providers, are joined into a set of federated resource pools, supervised by HTCondor and GlideinWMS services. The CMS Submission Infrastructure team is responsible for acquiring and managing this aggregated computing power, with a total capacity of about 500k CPU cores, and assigning it to CMS workloads according to their requirements and the priorities defined by the collaboration.

The scheduling strategies implemented for this purpose need to be flexible enough to support a number of concurrent workload types, taking into account the availability of resources from diverse providers, as well as the evolving resource requirements of the processing campaigns that the system needs to manage concurrently and consecutively. This complex system needs to be optimized in order to maximize the resource utilization efficiency, thus harnessing the full potential of our distributed compute resources.

This contribution will describe the systematic investigation by the CMS Submission Infrastructure team aimed at identifying, classifying, and minimizing inefficiencies in the use of the CMS distributed resources resulting from our workload management and scheduling algorithms. Additionally, our presentation will include certain strategies devised and implemented to compensate for other sources of inefficiency, thereby optimizing resource utilization and enhancing the overall CMS computational throughput.

**Parallel (Track 9) / 395**

## Operational experience from the Spanish CMS Analysis Facility at CIEMAT

**Author:** Antonio Perez-Calero Yzquierdo<sup>1</sup>

**Co-authors:** Antonio Delgado Peris <sup>2</sup>; Cecilia Maria Morcillo Perez <sup>2</sup>; Francisco Javier Rodriguez-Calonge ; Jaime Leon Holgado <sup>3</sup>; Jose Flix Molina <sup>2</sup>; Jose Hernandez <sup>4</sup>; Miguel Cardenas Montes <sup>5</sup>

<sup>1</sup> *Centro de Investigaciones Energéticas Medioambientales y Tecnológicas*

<sup>2</sup> *CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)*

<sup>3</sup> *Imperial College (GB)*

<sup>4</sup> *CIEMAT*

<sup>5</sup> *Cent.de Investigac.Energeticas Medioambientales y Technol. (CIEMAT)*

**Corresponding Authors:** jose.flix.molina@cern.ch, antonio.delgado.peris@cern.ch, antonio.perez.calero.yzquierdo@cern.ch, cecilia.maria.morcillo.perez@cern.ch, jaime.leon.holgado@cern.ch, miguel.cardenas.montes@cern.ch, jose.hernandez.calama@cern.ch, calonge@ciemat.es

The anticipated surge in data volumes generated by the LHC in the coming years, especially during the High-Luminosity LHC phase, will reshape how physicists conduct their analysis. This necessitates a shift in programming paradigms and techniques for the final stages of analysis. As a result, there's a growing recognition within the community of the need for new computing infrastructures tailored to these evolving demands. To meet this need, the recently established Analysis Facility at the CIEMAT institute is already providing crucial support to the local analysis community. This contribution will describe the diverse resources and functionalities provided by the new facility, its expansion to complementary resources also available at CIEMAT, as well as the important feedback gained from the operational experience by the users.

**Poster session / 396**

## **Benchmarking XRootD-HTTPS on 400Gbps Links with Variable Latencies**

**Authors:** Aashay Arora<sup>1</sup>; Diego Davila Foyo<sup>1</sup>; Justas Balcas<sup>2</sup>; Thomas Lehman<sup>3</sup>; Xi Yang<sup>4</sup>

<sup>1</sup> *Univ. of California San Diego (US)*

<sup>2</sup> *California Institute of Technology (US)*

<sup>3</sup> *ESnet*

<sup>4</sup> *LBNL*

**Corresponding Authors:** xiyang@es.net, justas.balcas@cern.ch, aashay.arora@cern.ch, diego.davila@cern.ch, tlehman@es.net

In anticipation of the High Luminosity-LHC era, there's a critical need to oversee software readiness for upcoming growth in network traffic for production and user data analysis access. This paper looks into software and hardware required improvements in US-CMS Tier-2 sites to be able sustain and meet the projected 400 Gbps bandwidth demands, while tackling the challenge posed by varying latencies between sites. Specifically, our study focuses on identifying the performance of XRootD HTTP third-party copies across multiple 400 Gbps links and exploring different host and transfer configurations.

Our approach involves systematic testing with variations in the number of origins per cluster and CPU, Memory allocations for each origin. By replicating real network conditions and creating network "loops" that traverse multiple switches across the wide area network, we are able to replicate authentic network conditions.

**Parallel (Track 5) / 397**

## **Simulation Comparison for the mSTS Geometry based on Primitive ROOT/TGeo Solids and Tessellated Solids**



**Author:** Mehulkumar Shiroya<sup>1</sup>

**Co-authors:** Alberica Toia<sup>2</sup>; Christoph Blume<sup>2</sup>; Eoin Clerkin<sup>3</sup>; Florian Uhlig<sup>4</sup>; Omveer Singh<sup>2</sup>; Volker Friese<sup>4</sup>

<sup>1</sup> *GSI Helmholtzzentrum für Schwerionenforschung GmbH, Goethe Universität, Frankfurt am Main, Germany, Helmholtz Forschungsakademie Hessen für FAIR, Frankfurt am Main, Germany*

<sup>2</sup> *Goethe Universität, Frankfurt am Main, Germany*

<sup>3</sup> *Facility for Antiproton and Ion Research in Europe, Darmstadt*

<sup>4</sup> *GSI Helmholtzzentrum für Schwerionenforschung GmbH*

**Corresponding Authors:** v.friese@gsi.de, f.uhlig@gsi.de, e.clerkin@gsi.de, a.toia@gsi.de, mehul0113@gmail.com, blume@ikf.uni-frankfurt.de, omvir.ch@gmail.com

The Compressed Baryonic Matter (CBM) is an under-construction heavy-ion physics experiment for exploring the QCD phase diagram at high  $\mu_B$  which will use the new SIS-100 accelerator at the Facility for Anti-Proton and Ion Research (FAIR) in Darmstadt, Germany. The Silicon Tracking System (STS) is to be the main detector for tracking and momentum determination. A scaled-down prototype of various detector systems including mini STS (mSTS) is undergoing meticulous testing in the mini CBM (mCBM) experiment at the existing SIS-18 accelerator at GSI, Helmholtzzentrum für Schwerionenforschung in Darmstadt. This initiative seeks to comprehensively assess both hardware and software components, ensuring their efficacy in online capturing, processing and analyzing the intricate topological data generated by real events detected by the detector sub-systems.

In recent years, much effort has been put into a better and more accurate description of the detector geometries to better model the background. The direct conversion of Computer-Aided Design (CAD) based geometry model to Geometry Description Markup Language (GDML), XML-based format using different software toolkits has attracted considerable attention. The solids extracted from CAD models and represented in GDML format typically consist of triangular or quadrilateral facets. TGDMLParser functionality in the ROOT and G4GDMLParser in the GEANT facilitate the reading of different volumes from the GDML file and the creation of volume assemblies. However, this approach leads to an increase in simulation computation run-time.

We will present a comparative analysis of simulation studies with two distinct representations of the mSTS geometry: one employing simplified primitive ROOT/TGeo solids and the other utilizing Tessellated solid-based geometry, including secondary particle production, the significance of passive volumes, computation time; as well as a comparison of simulation data with real data measured with Ni-Ni collisions at 1.93 AGeV.

**Parallel (Track 4) / 398**

## **The Cherenkov Telescope Array Observatory Production System Status and Development**

**Author:** Nathan PIGOUX<sup>None</sup>

**Co-author:** Luisa ARRABITO<sup>1</sup>

<sup>1</sup> *LUPM IN2P3/CNRS*

**Corresponding Authors:** natthan.pigoux@lupm.in2p3.fr, arrabito@in2p3.fr

The Cherenkov Telescope Array Observatory (CTAO) is the next-generation instrument in the very-high energy gamma ray astronomy domain. It will consist of tens of Cherenkov telescopes deployed in 2 CTAO array sites at La Palma (Spain) and Paranal (ESO, Chile) respectively. Currently under construction, CTAO will start operations in the coming years for a duration of about 30 years. During operations CTAO is expected to produce about 2 PB of raw data per year plus 5-20 PB of Monte Carlo data as well as very high processing needs of the order of hundreds of millions of CPU HS06 hours per year. These computing resources will be distributed across the 4 official CTAO Data Centers. To handle these simulations and data processing, we have developed a production system prototype

based on the DIRAC interware. We will present the current status of this prototype, the underlying infrastructure and the used technologies, as well as recent developments regarding workflows interface and failure management, along with future perspectives.

Parallel (Track 7) / 399

## Commissioning and exploitation of the MareNostrum5 cluster at the Barcelona Supercomputing Center for CMS computing

**Author:** Antonio Perez-Calero Yzquierdo<sup>1</sup>

**Co-authors:** Antonio Delgado Peris<sup>2</sup>; Jose Flix Molina<sup>2</sup>; Jose Hernandez<sup>3</sup>

<sup>1</sup> *Centro de Investigaciones Energéticas Medioambientales y Tecnológicas*

<sup>2</sup> *CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)*

<sup>3</sup> *CIEMAT*

**Corresponding Authors:** antonio.delgado.peris@cern.ch, jose.hernandez.calama@cern.ch, antonio.perez.calero.yzquierdo@cern.ch, jose.flix.molina@cern.ch

The MareNostrum 5 (MN5) is the new 750k-core general-purpose cluster recently deployed at the Barcelona Supercomputing Center (BSC). MN5 presents new opportunities for the execution of CMS data processing and simulation tasks but suffers from the same stringent network connectivity limitations as its predecessor, MN4. The innovative solutions implemented to navigate these constraints and effectively leverage the resources within the CMS distributed computing environment need to be revisited. First, the new worker nodes have increased their processor core count, and are thus capable of handling larger multicore CPU-bound CMS simulation tasks. Furthermore, the provisioning of larger disk storage capacity for MN5 broadens the spectrum of CMS workload types that can be accommodated at BSC. This storage space could, for example, be used to temporarily host large datasets required as input for CMS tasks, such as the pile-up samples, usually accessed by proton collision simulation jobs at runtime from remote grid sites' storages. These tasks were previously unsuitable for execution, given the connectivity limitations from BSC to remote storages. Enhanced network bandwidth between MN5 and the Port d'Informació Científica (PIC) can also facilitate the expansion of BSC capabilities by provisioning input for CMS data processing tasks at BSC, thus expanding the role of this resource in the CMS computing landscape. This contribution will provide an overview of the commissioning efforts and the results of the subsequent exploitation of MN5 for CMS, showcasing the new transformative capacities introduced by the MN5 cluster.

Parallel (Track 7) / 400

## Cache Rules Everything Around Me: Building ePIC Containers With Spack

**Author:** Wouter Deconinck<sup>None</sup>

**Corresponding Author:** wouter.deconinck@umanitoba.ca

The ePIC collaboration is working towards the realization of the first detector at the upcoming Electron-Ion Collider. As part of our computing strategy, we have settled on containers for the distribution of our modular software stacks using spack as the package manager. Based on abstract definitions of multiple mutually consistent software environments, we build dedicated containers on each commit of every pull request for the software projects under our purview. This is only possible through judicious caching from container layers, over downloaded artifacts and binary builds, down to individual compiled files. These containers are subsequently used for our benchmark and validation workflows. Our container build infrastructure runs with redundancy between GitHub and self-hosted GitLab resources, and can take advantage of cloud-based resources in periods of peak demand. In this talk, I will discuss our experiences with newer features of spack, including storing

build products as OCI layers and inheritance of previously concretized environments for software stack layering.

**Parallel (Track 4) / 401**

## Towards the Einstein Telescope Computing Model

**Author:** Stefano Bagnasco<sup>1</sup>

<sup>1</sup> *Istituto Nazionale di Fisica Nucleare, Torino*

**Corresponding Author:** bagnasco@to.infn.it

The Einstein Telescope is the proposed European next-generation ground-based gravitational-wave observatory, that is planned to have a vastly increased sensitivity with respect to current observatories, particularly in the lower frequencies. This will result in the detection of far more transient events, which will stay in-band for much longer, such that there will nearly always be at least one transient signal within the detector sensitivity. Besides the technological challenges that an underground, cryogenic instrument poses, many current data analyses cannot be trivially scaled without the required computing power growing beyond reasonably available resources. Furthermore, the detection and characterisation of events needs to be carried out with the minimum possible latency to guarantee the timely distribution of public alerts for multimessenger science follow-up. The Einstein Telescope Collaboration is carrying out a series of Mock Data Challenges with the aim, besides developing scientific algorithms and techniques, to test and evaluate the technological components that will form the ET distributed computing infrastructure. The strategy is to iteratively evolve the tools available to manage data and workloads and get early feedback from the scientific user community, possibly on different, competing implementations offering the same functionalities. We will discuss the preliminary outcomes of the first Mock Data Challenge and plans for subsequent ones, and on the status of the preparation of the Einstein Telescope Computing Model in general.

**Poster session / 402**

## Just-in-time workflow management for DUNE

**Author:** Andrew McNab<sup>1</sup>

<sup>1</sup> *University of Manchester*

**Corresponding Author:** andrew.mcnab@cern.ch

We describe the justIN workflow management system developed by DUNE to address its unique requirements and constraints. The DUNE experiment will start running in 2029, recording 30 PB/year of raw data from the detectors, with typical readouts at the scale of gigabytes, but with regular supernova candidate readouts of several hundred terabytes. DUNE benefits from the rich heritage of neutrino experiments at Fermilab, including the use of the SAM system to manage data, metadata, and processing campaigns. Due to the increase in scale required for DUNE, SAM's metadata database has been replaced by a new system, MetaCat, and its data management role is now taken up by Rucio. A new workflow system, justIN, has been developed since 2021 to replace the remaining functionality of SAM, and to tie together MetaCat, Rucio, and the GlideInWMS job management system to allow data processing campaigns involving hundreds of thousands of files and jobs using CPU and storage on four continents. Crucial to the design of justIN is an evolution of SAM's just-in-time philosophy, in which running jobs ask the central service for the optimal file to process next given their location. This model and the SAM interface are directly supported by the liquid argon data processing application used by DUNE, and justIN's design allows us to continue to support both. justIN goes a step further by assigning the workflows themselves to running GlideInWMS pilot jobs based on the locations of unprocessed files in active workflows at that time. This architecture allows the system to rapidly respond to problems such as downtimes, and to the sudden appearance of higher priority tasks such as processing supernova candidates.

**Plenary session / 403****The 200Gbps Challenge: Imagining HL-LHC analysis facilities**

**Authors:** Alexander Held<sup>1</sup>; Andrew Wightman<sup>2</sup>; Benjamin Tovar Lopez<sup>3</sup>; Brian Paul Bockelman<sup>1</sup>; Carl Lundstedt<sup>2</sup>; Derek Weitzel<sup>2</sup>; Fengping Hu<sup>4</sup>; Garhan Attebury<sup>5</sup>; Gordon Watts<sup>6</sup>; Ilija Vukotic<sup>4</sup>; Jim Pivarski<sup>7</sup>; John Thiltges<sup>2</sup>; Kenneth Bloom<sup>2</sup>; Kyle Stuart Cranmer<sup>1</sup>; Kyungeon Choi<sup>8</sup>; Lincoln Bryant<sup>4</sup>; Lindsey Gray<sup>9</sup>; Matthew Feickert<sup>1</sup>; Nick Smith<sup>9</sup>; Oksana Shadura<sup>2</sup>; Peter Onyisi<sup>8</sup>; Robert William Gardner Jr<sup>4</sup>; Sam Albin<sup>10</sup>

<sup>1</sup> *University of Wisconsin Madison (US)*

<sup>2</sup> *University of Nebraska Lincoln (US)*

<sup>3</sup> *University of Notre Dame*

<sup>4</sup> *University of Chicago (US)*

<sup>5</sup> *University of Nebraska-Lincoln*

<sup>6</sup> *University of Washington (US)*

<sup>7</sup> *Princeton University*

<sup>8</sup> *University of Texas at Austin (US)*

<sup>9</sup> *Fermi National Accelerator Lab. (US)*

<sup>10</sup> *UNL*

**Corresponding Authors:** john.thiltges@cern.ch, attebury@cse.unl.edu, nick.smith@cern.ch, kyungeon.choi@cern.ch, derek.weitzel@cern.ch, lindsey.gray@cern.ch, lincoln.bryant@cern.ch, fengping@uchicago.edu, matthew.feickert@cern.ch, sam.albin@unl.edu, btovar@nd.edu, clundstedt@unl.edu, ilija.vukotic@cern.ch, gwatts@uw.edu, kyle.cranmer@cern.ch, robert.w.gardner@cern.ch, pivarski@princeton.edu, alexander.held@cern.ch, ponyisi@utexas.edu, brian.bockelman@cern.ch, kenbloom@unl.edu, oksana.shadura@cern.ch, andrew.steven.wightman@cern.ch

The IRIS-HEP software institute, as a contributor to the broader HEP Python ecosystem, is developing scalable analysis infrastructure and software tools to address the upcoming HL-LHC computing challenges with new approaches and paradigms, driven by our vision of what HL-LHC analysis will require. The institute uses a “Grand Challenge” format, constructing a series of increasingly large, complex, and realistic exercises to show the vision of HL-LHC analysis. Recently, the focus has been demonstrating the IRIS-HEP analysis infrastructure at scale and evaluating technology readiness for production.

As a part of the Analysis Grand Challenge activities, the institute executed a “200 Gbps Challenge”, aiming to show sustained data rates into the event processing of multiple analysis pipelines. The challenge integrated teams internal and external to the institute, including operations and facilities, analysis software tools, innovative data delivery and management services, and scalable analysis infrastructure. The challenge showcases the prototypes—including software, services, and facilities—built to process around 200 TB of data in both the CMS NanoAOD and ATLAS PHYSLITE data formats with test pipelines.

The teams were able to sustain the 200 Gbps target across multiple pipelines. The pipelines focusing on event rate were able to process at over 30MHz. These target rates are demanding; the activity revealed considerations for future testing at this scale and changes necessary for physicists to work at this scale in the future. The 200 Gbps challenge has established a baseline on today’s facilities, setting the stage for the next exercise at twice the scale.

**Poster session / 404****Orchestrated columnar-based analysis with columnflow**

**Authors:** Bogdan Wiederspan<sup>1</sup>; Daniel Savoie<sup>1</sup>; Johannes Haller<sup>1</sup>; Marcel Rieger<sup>1</sup>; Mathis Frahm<sup>1</sup>; Matthias Schroeder<sup>1</sup>; Nathan Provoust<sup>1</sup>; Peter Schleper<sup>1</sup>; Philip Keicher<sup>1</sup>

<sup>1</sup> *Hamburg University (DE)*

**Corresponding Authors:** peter.schleper@physik.uni-hamburg.de, philip.daniel.keicher@cern.ch, mathis.frahm@cern.ch, bogdan.wiederspan@uni-hamburg.de, daniel.savoiu@cern.ch, nathan.prouvost@uni-hamburg.de, johannes.haller@uni-hamburg.de, marcel.rieger@cern.ch, matthias.schroeder@uni-hamburg.de

To study and search for increasingly rare physics processes at the LHC, a staggering amount of data needs to be analyzed with progressively complex methods. Analyses involving tens of billions of recorded and simulated events, multiple machine learning algorithms for different purposes, and an amount of 100 or more systematic variations are no longer uncommon. These conditions impose a complex data flow on an analysis workflow and render its steering and bookkeeping a serious challenge. For this purpose, a toolkit for columnar HEP analysis, called columnflow, has been developed. It is written in Python, experiment agnostic in its core, and supports any flat file format, such as ROOT-based trees or Parquet files. Leveraging on the vast Python ecosystem, vectorization and convenient physics objects representation can be achieved through NumPy, awkward arrays and other libraries. Based upon the Luigi Analysis Workflow (law) package, columnflow provides full analysis automation over arbitrary, distributed computing resources. Despite the end-to-end nature, this approach allows for persistent, intermediate outputs for purposes of debugging, caching, and exchange with collaborators. Job submission to various batch systems, such as HTCondor, Slurm, or CMS-CRAB, is natively supported. Remote files can be seamlessly accessed via various protocols using either the Grid File Access Library (GFAL2) or the fsspec file system interface. In addition, a sandboxing mechanism can encapsulate the execution of parts of a workflow into dedicated environments, supporting subshells, virtual environments, and containers. This contribution introduces the key components of columnflow and highlights the benefits of a fully automated workflow for complex and large-scale HEP analyses.

Parallel (Track 5) / 405

## Novel Fitting Approach Based on a Neural Network for JUNO

**Author:** Yury Malyshkin<sup>1</sup>

<sup>1</sup> GSI / Forschungszentrum Jülich

**Corresponding Author:** y.malyshkin@gsi.de

JUNO (Jiangmen Underground Neutrino Observatory) is a neutrino experiment being built in South China. Its primary goals are to resolve the order of the neutrino mass eigenstates and to precisely measure the oscillation parameters  $\sin^2 \theta_{12}$ ,  $\Delta m_{21}^2$ , and  $\Delta m_{31(32)}^2$  by observing the oscillation pattern of electron antineutrinos produced in eight reactor cores of two commercial nuclear power plants at a distance of 52.5 km. A crucial stage in the data analysis is to fit the observed spectrum to the expected one under different oscillation scenarios taking into account realistic detector response, backgrounds, and all relevant uncertainties. This task becomes computationally challenging when a full Monte Carlo simulation of the detector is directly used to predict the detector response instead of otherwise used empirical models. It is proposed using a neural network to precisely predict the detector spectrum as a function of oscillation parameters and a set of detector response parameters. This approach drastically reduces computation time and makes it possible to fit a spectrum within one second. The contribution presents the details, performance, and limitations of the method.

Parallel (Track 2) / 406

## A high-throughput input interface for the CBM FLES

**Author:** Dirk Hutter<sup>1</sup>

**Co-authors:** Jan de Cuveland<sup>1</sup>; Volker Lindenstruth<sup>1</sup>

<sup>1</sup> *Goethe University Frankfurt (DE)*

**Corresponding Authors:** voli@compeng.de, jan.de.cuveland@cern.ch, hutter@compeng.uni-frankfurt.de

The CBM First-level Event Selector (FLES) serves as the central data processing and event selection system for the upcoming CBM experiment at FAIR. Designed as a scalable high-performance computing cluster, it facilitates online analysis of unfiltered physics data at rates surpassing 1 TByte/s.

As the input to the FLES, the CBM detector subsystems deliver free-streaming, self-triggered data to the common readout interface (CRI), which is a custom FPGA PCIe board installed in the FLES entry nodes. A subsystem-specific part of the FPGA design time-partitions the input streams into context-free packages. The FLES interface module (FLIM), a component of the FPGA design, acts as the interface between the subsystem-specific readout logic and the generic FLES data distribution. It transfers the packed detector data to the host's memory using a low-latency, high-throughput PCIe DMA engine. This custom design enables a shared-memory-based, true zero-copy data flow.

A fully implemented FLIM for the CRI board is currently in use within CBM test setups and the FAIR Phase-0 experiment mCBM. We present an overview of the FLES input interface architecture and provide performance evaluations under synthetic as well as real-world conditions.

This work is supported by BMBF (05P21RFFC1).

**Parallel (Track 3) / 407**

## **Towards an open-source hybrid quantum operating system**

**Author:** Edoardo Pedicillo<sup>None</sup>

**Co-authors:** Andrea Pasquale <sup>1</sup>; Matteo Robbiati <sup>2</sup>; Stefano Carrazza <sup>3</sup>

<sup>1</sup> *University of Milan*

<sup>2</sup> *Università degli Studi e INFN Milano (IT)*

<sup>3</sup> *CERN*

**Corresponding Authors:** edoardo.pedicillo@unimi.it, stefano.carrazza@cern.ch, andrea.pasquale@unimi.it, matteo.robbiati@cern.ch

Over the last 20 years, thanks to the development of quantum technologies, it has been possible to deploy quantum algorithms and applications that before were only accessible through simulation on real quantum hardware.

The current devices available are often referred to as noisy intermediate-scale quantum (NISQ) computers, and they require calibration protocols in order to obtain consistent results.

In this context, we present Qibo, an open-source framework for quantum computing. Qibo was initially born as a tool for simulating quantum circuits.

Through its modular layout for backend abstraction, it is possible to change effortlessly between different backends, including a simulator based on just-in-time compilation, Qibojit. Thanks to the Rust-based library Qibo-core, Qibo offers multi-language support for its backends.

The Qibo ecosystem also offers packages for specific applications, like Qiboml, a suite to deploy quantum and hybrid classical-quantum machine learning models that can be executed on any type of hardware accelerators.

In order to enable the execution and calibration of self-hosted quantum hardware, we have developed two open-source libraries integrated with the Qibo framework: Qibolab and Qibocal.

Qibolab provides the software layer required to automatically execute circuit-based algorithms on custom self-hosted quantum hardware platforms.

Qibocal is based on a modular QPU (Quantum Processor Unit) platform-agnostic approach and introduces tools that support the calibration and characterization of QPUs. Qibocal provides a code library to rapidly develop protocols for different hardware abstraction layers.

Poster session / 408

## Data-flow parallelism for high-energy and nuclear physics frameworks

**Authors:** Kyle Knoepfel<sup>1</sup>; Marc Paterno<sup>None</sup>; Saba Sehrish<sup>2</sup>

<sup>1</sup> *Fermi National Accelerator Laboratory*

<sup>2</sup> *Fermilab*

**Corresponding Authors:** [knoepfel@fnal.gov](mailto:knoepfel@fnal.gov), [paterno@fnal.gov](mailto:paterno@fnal.gov), [ssehrish@fnal.gov](mailto:ssehrish@fnal.gov)

The processing tasks of an event-processing workflow in high-energy and nuclear physics (HENP) can typically be represented as a directed acyclic graph formed according to the data flow—i.e. the data dependencies among algorithms executed as part of the workflow. With this representation, an HENP framework can optimally execute a workflow, exploiting the parallelism inherent among independent tasks. Despite such a natural description of a workflow, most HENP frameworks do not make use of technologies that provide concurrent execution of graph-based tasking structures.

In this talk, we describe Fermilab efforts to adopt a graph-based technology (specifically Intel's oneTBB flow graph) for meeting the framework needs of its experiments, notably DUNE. Building on the Meld project as presented at CHEP2023, we demonstrate that all common processing idioms supported by current frameworks can naturally be supported by oneTBB's data-flow technology, optimally leveraging the concurrent capabilities of the machine. In addition, we discuss collaborative efforts between Fermilab and the Intel oneTBB development team, who is considering improvements to the flow-graph technology to better support HENP use cases.

Parallel (Track6) / 409

## Multi-package development at Fermilab with Spack

**Author:** Kyle Knoepfel<sup>1</sup>

<sup>1</sup> *Fermi National Accelerator Laboratory*

**Corresponding Author:** [knoepfel@fnal.gov](mailto:knoepfel@fnal.gov)

The Spack package manager has been widely adopted in the supercomputing community as a means of providing consistently built on-demand software for the platform of interest. Members of the high-energy and nuclear physics (HENP) community, in turn, have recognized Spack's strengths, used it for their own projects, and even become active Spack developers to better support HENP needs. Code development in a Spack context, however, can be challenging as the provision of external software via Spack must integrate with the developed packages' build systems. Spack's own development features can be used for this task, but they tend to be inefficient and cumbersome.

We present a solution pursued at Fermilab called MPD (multi-package development). MPD aims to facilitate the development of multiple Spack-based packages in concert without the overhead of Spack's own development facilities. In addition, MPD allows physicists to create multiple development projects with an interface that insulates users from the many commands required to use Spack well.

**Parallel (Track 7) / 410**

## Checkpoint-Restart for HPC

**Author:** Madan Timalina<sup>1</sup>**Co-authors:** Johannes Blaschke<sup>1</sup>; Lisa Gerhardt<sup>1</sup>; Nicholas Tyler<sup>1</sup>; Urjoshi Sinha<sup>1</sup>; William Arndt,<sup>1</sup><sup>1</sup> NERSC/LBNL**Corresponding Author:** mtimalina@lbl.gov

This presentation delves into the implementation and optimization of checkpoint-restart mechanisms in High-Performance Computing (HPC) environments, with a particular focus on Distributed MultiThreaded CheckPointing (DMTCP). We explore the use of DMTCP both within and outside of containerized environments, emphasizing its application on NERSC Perlmutter, a cutting-edge supercomputing system. The discussion highlights the benefits of checkpoint-restart (C/R) techniques in managing complex, long-duration computations, showcasing the efficiency and reliability of these methods. Based on Geant4, a crucial tool for High Energy and Nuclear Physics, these techniques have been thoroughly tested and have passed the assessments. We further examine the integration of HPC containers, such as Shifter and Podman-HPC, which enhance computational task management and ensure consistent performance across various environments. Through real-world application examples, we illustrate the advantages of DMTCP in multi-threaded and distributed computing scenarios. Additionally we present the methods and results, demonstrating the impact of C/R on resource utilization, the future directions of this research, and its potential across various scientific domains.

**Parallel (Track 2) / 411**

## Machine Learning for Optimized Polarization at Jefferson Lab

**Authors:** Armen Kasparian<sup>None</sup>; Cristiano Fanelli<sup>1</sup>; David Lawrence<sup>None</sup>; James Maxwell<sup>2</sup>; Jiawei Guo<sup>3</sup>; Malachi Schram<sup>None</sup>; Naomi Jarvis<sup>3</sup>; Patrick Moran<sup>None</sup>; Thomas Britton<sup>None</sup>; Torri Jeske<sup>None</sup><sup>1</sup> *William & Mary*<sup>2</sup> *Jefferson Lab*<sup>3</sup> *Carnegie Mellon University***Corresponding Authors:** gjwei@jlab.org, schram@jlab.org, armenk@jlab.org, moranp@jlab.org, cfanelli@wm.edu, nsj@cmu.edu, tbritton@jlab.org, jmaxwell@jlab.org, roark@jlab.org, davidl@jlab.org

Polarized cryo-targets and polarized photon beams are widely used in experiments at Jefferson Lab. Traditional methods for maintaining the optimal polarization involve manual adjustments throughout data taking— an approach that is prone to inconsistency and human error. Implementing machine learning-based control systems can improve the stability of the polarization without relying on human intervention. The cryo-target polarization is influenced by temperature, microwave energy, the distribution of paramagnetic radicals, as well as operational conditions including the radiation dose. Diamond radiators are used to generate linearly polarized photons from a primary electron beam. The energy spectrum of these photons can drift over time due to changes in the primary electron beam conditions and diamond degradation. As a first step towards automating the continuous optimization and control processes, uncertainty aware surrogate models have been developed to predict the polarization based on historical data. This talk will provide an overview of the use cases and models developed, highlighting the collaboration between data scientists and physicists at Jefferson Lab.

**Parallel (Track 1) / 412**



## Efficient metadata management with the AMI ecosystem

**Authors:** Fabian Lambert<sup>1</sup>; Jerome Henri Fulachier<sup>2</sup>; Jerome Odier<sup>3</sup>; Pierre Antoine Delsart<sup>3</sup>

<sup>1</sup> LPSC Grenoble IN2P3/CNRS (FR)

<sup>2</sup> Centre National de la Recherche Scientifique (FR)

<sup>3</sup> LPSC/CNRS (Grenoble, FR)

**Corresponding Authors:** jerome.fulachier@lpsc.in2p3.fr, fabian.lambert@lpsc.in2p3.fr, delsart@in2p3.fr, jerome.odier@cern.ch

The ATLAS Metadata Interface (AMI) is a comprehensive ecosystem designed for metadata aggregation, transformation, and cataloging. With over 20 years of feedback in the LHC context, it is particularly well-suited for scientific experiments that generate large volumes of data.

This presentation explains, in a general manner, why managing metadata is essential regardless of the experiment's scale. It then presents the different AMI ecosystem's components and their main functionalities, particularly the Web interfaces for searching data based on metadata criteria. Finally, it discusses the deployment of a functional demo, its subsequent scaling up, and how to integrate it into a data production system.

**Parallel (Track 7) / 413**

## Near-Data Computing Model for Accelerating LHC Data Filtering

**Authors:** Diego Davila Foyo<sup>1</sup>; Frank Wuerthwein<sup>1</sup>; Jonathan Guiang<sup>1</sup>; Narangerelt Batsoyol<sup>2</sup>; Philip Chang<sup>3</sup>; Steven Swanson<sup>2</sup>

<sup>1</sup> Univ. of California San Diego (US)

<sup>2</sup> University of California San Diego

<sup>3</sup> University of Florida (US)

**Corresponding Authors:** frank.wuerthwein@cern.ch, nbatsoyo@ucsd.edu, sjswanson@ucsd.edu, philip.chang@cern.ch, jonathan.guiang@cern.ch, diego.davila@cern.ch

The data reduction stage is a major bottleneck in processing data from the Large Hadron Collider (LHC) at CERN, which generates hundreds of petabytes annually for fundamental particle physics research. Here, scientists must refine petabytes into only gigabytes of relevant information for analysis. This data filtering process is limited by slow network speeds when fetching data from globally dispersed storage facilities, which leads to thousands of wasted CPU hours waiting for data to arrive.

We demonstrate a near-data computing model that optimizes data access and enhances performance by filtering LHC data close to its storage before transmission over the slow network. This model is designed to be implemented with minimal change in the existing data layout and seamless integration with the underlying storage infrastructure, ensuring compatibility and ease of adoption for current systems.

We achieve this by deploying Data Processing Units (DPUs) within the storage cluster. Our model leverages DPU's high-bandwidth connections to perform fast data retrieval and filtering near storage, significantly improving overall data processing speeds and freeing up compute node CPUs for more important tasks. Additionally, it streamlines the workflow by removing coding complexities and making programming accessible for end users. We demonstrate that our model significantly outperforms current methods using real physics data and a realistic data reduction workflow.

**Parallel (Track 3) / 414**

## Automatic and effective discovery of quantum kernels

**Authors:** Alessandra Di Piero<sup>None</sup>; Daniele Lizzio Bosco<sup>None</sup>; Francesco Martini<sup>None</sup>; Giuseppe Serra<sup>None</sup>; Massimiliano Incudini<sup>1</sup>; Michele Grossi<sup>2</sup>

<sup>1</sup> *Universita e INFN, Padova (IT)*

<sup>2</sup> *CERN*

**Corresponding Authors:** [mincudin@cern.ch](mailto:mincudin@cern.ch), [michele.grossi@cern.ch](mailto:michele.grossi@cern.ch)

Quantum computing can empower machine learning models by enabling kernel machines to leverage quantum kernels for representing similarity measures between data. Quantum kernels are able to capture relationships in the data that are not efficiently computable on classical devices. However, there is no straightforward method to engineer the optimal quantum kernel for each specific use case. While recent literature has focused on exploiting the potential offered by the presence of symmetries in the data to guide the construction of quantum kernels, we adopt here a different approach, which employs optimization techniques, similar to those used in neural architecture search and AutoML, to automatically find an optimal kernel in a heuristic manner. The algorithm we present constructs a quantum circuit implementing the similarity measure as a combinatorial object, which is evaluated based on a cost function and is then iteratively modified using a meta-heuristic optimization technique. The cost function can encode many criteria ensuring favorable statistical properties of the candidate solution, such as the rank of the Dynamical Lie Algebra. Importantly, our approach is independent of the optimization technique employed. The results obtained by testing our approach on a high-energy physics problem demonstrate that, in the best-case scenario, we can either match or improve testing accuracy with respect to the manual design approach, showing the potential of our technique to deliver superior results with reduced effort.

Parallel (Track 2) / 415

## Overview of the data acquisition system architecture for the DarkSide-20k experiment

**Author:** Marek Walczak<sup>1</sup>

<sup>1</sup> *Gran Sasso Science Institute*

**Corresponding Author:** [marek.walczak@gssi.it](mailto:marek.walczak@gssi.it)

The DarkSide-20k detector is now under construction in the Gran Sasso National Laboratory (LNGS) in Italy, the biggest underground physics facility. It is designed to directly detect dark matter by observing weakly interacting massive particles (WIMPs) scattering off the nuclei in 20 tonnes of underground-sourced liquid argon in the dual-phase time projection chamber (TPC). Additionally two layers of veto detectors allow operating with virtually zero instrumental background in the region of interest, leaving only irreducible neutrino interaction. When operating the DarkSide-20k experiment is expected to lead the field for high mass WIMPs searches in the next decade and due to the low background will have a high discovery potential. Thanks to its size and sensitivity the detector will allow a broad physics program including supernova neutrino detection.

The light generated during the interactions in the liquid argon is detected by custom silicon photomultipliers (SiPMs) assemblies of size 20 cm by 20 cm. The units installed in the veto detectors are equipped with application specific integrated circuits (ASICs) coupled to SiPMs allowing linear signal response up to 100 photons and signal to noise ratio of 6 for a single photon, while those for the TPC employ a discrete element front-end with similar performances.

The data acquisition system (DAQ) for the DarkSide-20k experiment is designed to acquire signals from the 2720 channels of these photosensors in a triggerless mode. The data rate from the TPC alone is expected to be at the level of 2.5 GB/s and will be acquired by 36 newly available commercial VX2745 CAEN 16 bit, 125 MS/s, high channel density (64 ch.) waveform digitizers. The Veto detector is readout by an additional 12 modules. The data is first transferred to 24 Frontend Processor machines for filtering and reduction. Finally the data stream is received by another set of Time Slice

Processor computers where the whole detector data is assembled in fixed length time series, analysed and stored for offline use. These operations will be supervised by a Maximum Integration Data Acquisition System (MIDAS) developed in the Paul Scherrer Institute in Switzerland and TRIUMF laboratory in Canada.

Parallel (Track 2) / 416

## Hydra: An AI-Based Framework for Interpretable and Portable Data Quality Monitoring

**Authors:** Nataliia Matsiuk<sup>1</sup>; Raiqa Rasool<sup>1</sup>; Thomas Britton<sup>None</sup>; Torri Jeske<sup>None</sup>

<sup>1</sup> *JLab*

**Corresponding Authors:** tbritton@jlab.org, roark@jlab.org, rasool@jlab.org

Hydra is an advanced framework designed for training and managing AI models for near real time data quality monitoring at Jefferson Lab. Deployed in all four experimental halls, Hydra has analyzed over 2 million images and has extended its capabilities to offline monitoring and validation. Hydra utilizes computer vision to continually analyze sets of images of monitoring plots generated 24/7 during experiments. Generally, these sets of images are produced at a rate and quantity that is exceedingly difficult for shift crews to effectively monitor. Significant effort has been devoted to enhancing Hydra's user interface, to ensure that it provides clear, actionable insights for shift workers and other users. Gradient Weighted Class Activation Maps (GradCAM) provide added interpretability, allowing users to visualize important regions of the image for classification. Hydra has been containerized to enable the creation of portable demos and seamless integration with container-based technologies such as Kubernetes and Docker. With the user interface enhancements and containerization, Hydra can be rapidly deployed for new use cases and experiments. This talk will describe the Hydra framework, its user interface and experience, and the challenges inherent in its design and deployment.

Poster session / 417

## Event Workflow Management System - A SaaS Solution for Massively Divisible and Distributed Workflows

**Author:** Ric Evans<sup>1</sup>

**Co-authors:** Benedikt Riedel<sup>2</sup>; Brian Aydemir<sup>3</sup>; Brian Paul Bockelman<sup>4</sup>; David Schultz<sup>2</sup>; MIRON LIVNY

<sup>1</sup> *Wisconsin IceCube Particle Astrophysics Center*

<sup>2</sup> *University of Wisconsin-Madison*

<sup>3</sup> *Morgridge Institute for Research*

<sup>4</sup> *University of Wisconsin Madison (US)*

**Corresponding Authors:** david.schultz@icecube.wisc.edu, briedel@icecube.wisc.edu, brian.bockelman@cern.ch, baydemir@morgridge.org, eevans@icecube.wisc.edu, miron@cs.wisc.edu

How does one take a workload, consisting of millions or billions of tasks, and group it into tens of thousands of jobs? Partitioning the workload into a workflow of long-running jobs minimizes the use of scheduler resources; however, smaller, more fine-grained jobs allow more efficient use of computing resources. When the runtime of a task averages a minute or less, severe scaling challenges due to scheduling overhead can surface. Employing jobs that run for several hours, each with a large input file comprising a bundle of tasks, is effective in ideal situations. However, given the heterogeneity of available distributed resources and limited control of task-job matching, runtimes can vary widely.

The Event Workflow Management System (EWMS) augments HTCondor to solve this issue. EWMS implements a pilot-based paradigm where each worker, running inside an HTCondor execution point, connects to a message broker and executes many individual fine-grained tasks. This adaptive design increases task throughput while incorporating additional fail-safe features. In addition, EWMS manages workflow scheduling, enables real-time worker scaling, and exports a public-facing interface for user accessibility. Here, we outline the EWMS technique, detail science driver workflows from the IceCube experiment, and provide system usage metrics.

**Parallel (Track 4) / 418**

## Development and Testing of DUNE's Production System

**Author:** Jacob Michael Calcutt<sup>1</sup>

<sup>1</sup> *Oregon State University (US)*

**Corresponding Author:** calcuttj@oregonstate.edu

The DUNE experiment will start running in 2029 and record 30 PB/year of raw waveforms from Liquid Argon TPCs and photon detectors. The size of individual readouts can range from 100 MB to a typical 8 GB full readout of the detector to extended readouts of up to several 100 TB from supernova candidates. These data then need to be cataloged, stored and then distributed for processing worldwide. This large massive amount of data and heterogeneous computing environment necessitates powerful and robust distributed computing infrastructure. In the process of building up that infrastructure, DUNE's production system has recently undergone an overhaul, in which it has integrated 1) a new workflow management system (justIN) 2) a new data catalog (MetaCat) and 3) a state-of-the-art data management system (Rucio). Simulations of DUNE's Far Detector and its prototypes ProtoDUNE-HD (PDHD) and ProtoDUNE-VD, as well as data from PDHD serve as the first tests of this infrastructure.

**Poster session / 419**

## Application of linear and non-linear constraints in a brute-force based alignment approach for CBM

**Author:** Nora Bluhme<sup>1</sup>

**Co-authors:** Sergey Gorbunov<sup>1</sup>; Volker Lindenstruth<sup>1</sup>

<sup>1</sup> *Goethe University Frankfurt (DE)*

**Corresponding Authors:** bluhme@fias.uni-frankfurt.de, voli@compeng.de, sergey.gorbunov@cern.ch

The future Compressed Baryonic Matter experiment (CBM), which is currently being planned and will be realised at the Facility for Antiproton and Ion Research (FAIR), is dedicated to the investigation of heavy-ion collisions at high interaction rates. For this purpose, a track-based software alignment is necessary to determine the precise detector component positions with sufficient accuracy. This information is crucial as it enables adequate utilisation of the high intrinsic accuracy of the sensors.

The alignment parameters to be determined are typically translations and rotations of individual sensors in relation to their intended nominal positions. They are usually determined by minimising a  $\chi^2$  function of a set of high quality reconstructed tracks.

To complement the available alignment tools, an additional approach is being developed that is based on brute-force  $\chi^2$  minimisation. On the one hand, this approach should allow different parameters

to be treated individually and, on the other hand, it opens up the possibility of integrating different types of constraints into the minimisation, such as inequality and non-linear constraints.

This contribution presents the latest developments in the application of brute-force alignment within the CBM project. The question of how the results of optical detector measurements, which usually precede software alignment, can be taken into account in this procedure is also addressed.

This work is supported by BMBF (05P21RFFC1).

Parallel (Track 8) / 420

## LHCb Open Data Ntupling Service: On-demand production and publishing of custom LHCb Open Data

**Authors:** Adam Morris<sup>1</sup>; Christine Aidala<sup>2</sup>; Daan Rosendal<sup>None</sup>; Dillon Fitzgerald<sup>2</sup>; Eduardo Rodrigues<sup>3</sup>; Franz Ludwig Kramer<sup>4</sup>; Kai Sebastian Habermann<sup>4</sup>; Marco Donadoni<sup>1</sup>; Piet Nogga<sup>4</sup>; Sebastian Neubert<sup>4</sup>; Tibor Simko<sup>1</sup>

<sup>1</sup> CERN

<sup>2</sup> University of Michigan (US)

<sup>3</sup> University of Liverpool (GB)

<sup>4</sup> University of Bonn (DE)

**Corresponding Authors:** kai.sebastian.habermann@cern.ch, s6frkrae@uni-bonn.de, christine.angela.aidala@cern.ch, adam.morris@cern.ch, daanrosendal@outlook.com, eduardo.rodrigues@cern.ch, dillfitz@umich.edu, sebastian.neubert@cern.ch, tibor.simko@cern.ch, marco.donadoni@cern.ch, pnogga@uni-bonn.de

The Large Hadron Collider Beauty (LHCb) experiment offers an excellent environment to study a broad variety of modern physics topics. Its data from the major physics campaigns (Run 1 and 2) at the Large Hadron Collider (LHC) has accumulated over 600 scientific publications. In accordance with the CERN Open Data Policy, LHCb announced the release of the full Run 1 dataset gathered from proton-proton collisions, amounting to approximately 800 terabytes. The Run 1 data was released on the CERN Open Data portal in 2023. However, due to the large amount of data collected during Run 2, it is no longer feasible to make the reconstructed data accessible to the public in the same way.

We have, therefore, developed a new and innovative approach to publishing Open Data by means of a dedicated LHCb Ntupling Service which allows third-party users to query the data collected by LHCb and request custom samples in the same columnar data format used by LHCb physicists. These samples are called Ntuples and can be individually customized in the web interface using LHCb standard tools for saving measured or derived quantities of interest. The configuration output is kept in a pure data structure format (YAML) and is interpreted by internal parsers generating the necessary Python scripts for the LHCb Ntuple production job. In this way, the LHCb Ntupling Service serves as a gateway for third-party users for preparing custom Ntuple jobs eliminating the need for real-time interaction with the LHCb database and solving potential access control and computer security issues related to opening LHCb internal tools to the public.

The LHCb Ntupling Service was developed as a collaborative effort by LHCb and the CERN Open Data team from the CERN Department of Information Technology. The service consists of the web interface frontend allowing users to create Ntuple production requests, the backend application processing the user requests and storing them in the GitLab repositories, offering vetting capabilities to the LHCb Open Data team, and automatically dispatching user requests to the LHCb Ntuple production systems after the approval. The produced Ntuples are then collected and exposed back to the users by the frontend web interface.

This talk is a joint presentation by LHCb and CERN IT and will elaborate on the LHCb Ntupling Service system infrastructure as well as its typical use case scenarios allowing to query and study the LHCb open data.

**Parallel (Track 3) / 421**

## **Sprucing and Analysis Productions: Offline data processing in LHCb without the pain**

**Authors:** Chris Burr<sup>1</sup>; Nicole Skidmore<sup>2</sup><sup>1</sup> *CERN*<sup>2</sup> *University of Warwick***Corresponding Authors:** christopher.burr@cern.ch, nicola.skidmore@cern.ch

After two successful physics runs the LHCb experiment underwent a comprehensive upgrade to enable LHCb to run at five times the instantaneous luminosity for Run 3 of the LHC. With this upgrade, LHCb is now the largest producer of data at the LHC. A new offline dataflow was developed to facilitate fast time-to-insight whilst respecting constraints from disk and CPU resources. The Sprucing is an offline data processing step that further refines the selections and persistency of physics channels coming out of the trigger system. In addition, the Sprucing splits the data into multiple streams, which are written in a format that facilitates more efficient compression. Next, Analysis Productions provide LHCb analysts with a declarative approach to tupling this data, efficiently exploiting WLCG resources in a centralised way.

The Sprucing and Analysis Productions offline chain provides analysts with their customised tuples within days of the data being taken by the LHCb experiment.

This talk will present the development of this offline data processing chain with a focus on performance results gathered during operations in 2024.

**Poster session / 422**

## **IceSONIC - Network AI Inference on Coprocessors for IceCube Offline Processing**

**Authors:** Alec Sheperd<sup>1</sup>; Benedikt Riedel<sup>None</sup>; David Schultz<sup>1</sup><sup>1</sup> *University of Wisconsin-Madison***Corresponding Authors:** briedel@wisc.edu, david.schultz@icecube.wisc.edu, alec.sheperd@icecube.wisc.edu

An Artificial Intelligence (AI) model will spend “90% of its lifetime in inference.” To fully utilize coprocessors, such as FPGAs or GPUs, for AI inference requires O(10) CPU cores to feed to work to the coprocessors. Traditional data analysis pipelines will not be able to effectively and efficiently use the coprocessors to their full potential. To allow for distributed access to coprocessors for AI inference workloads, the LHC’s Compact Muon Solenoid (CMS) experiment has developed the concept of Services for Optimized Network Inference on Coprocessors (SONIC) using NVIDIA’s Triton Inference Servers. We have extended this concept for the IceCube Neutrino Observatory by deploying NVIDIA’s Triton Inference Servers in local and external Kubernetes clusters, integrating an NVIDIA Triton Client with IceCube’s data analysis framework, and deploying an OAuth2-based HTTP authentication service in front of the Triton Inference Servers. We will describe the setup and our experience adding this to IceCube’s offline processing system.

**Poster session / 423**

## **Online Data Processing and monitoring for DarkSide-20k**

**Author:** Maria Adriana Sabia<sup>None</sup>**Corresponding Author:** maria.adriana.sabia@roma1.infn.it

Liquid Argon (LAr) Time Projection Chambers (TPC)s are promising detectors for dark matter search, due to their response uniformity, scalability to large target masses, and suitability for extremely low background operations. The Darkside-20k (DS-20k) experiment is a new dark matter detector under construction at INFN LNGS that aims to push the sensitivity for Weakly Interacting Massive Particles (WIMP) detection down to the neutrino floor.

DS-20k will employ a triggerless Data Acquisition system (DAQ) that continuously captures signals from SiPMs-based photosensors. The expected interaction rate is about 100 Hz, with a dark count rate of approximately 20 Hz per channel (2720 in total).

The DS-20k Data Acquisition System must be able to identify signals as small as 1 photoelectron with event rates as large as 200 Hz. Signals are digitized by the VME VX2745 CAEN module 16 bit, 125 MS/s ADC). The Waveform Digitizer (WFD) data are read out by a set of Frontend Processors (FEP), which filter signal waveforms and reduce them by identifying hits before passing the data to the Time Slice Processor (TSP).

Live data monitoring is crucial during the detector run phase for ensuring data integrity, identifying anomalies, and optimizing detector performance. This talk will cover the analysis performed using a prototype of the DS-20k monitoring system called “vertical slice” and developed at TRIUMF laboratory. The system consists of 10 machines, and will allow us to stress-test a fraction of the DAQ architecture, to test the digitizers readout, and the same diagnostics tools that will be used for DS-20k.

Additionally, the talk will cover a proposal for an online prompt analysis based on 1s time slices to address peculiar events like S2-only events, multiple scatterings and Supernova neutrino bursts.

#### Poster session / 424

### The Real-Time Data Workflow of LZ Dark Matter Experiment at NERSC

**Author:** Maris Arthurs<sup>None</sup>

**Corresponding Author:** amaris@slac.stanford.edu

LUX-ZEPLIN (LZ) is a dark matter direct detection experiment using a dual-phase xenon time projection chamber with a 7-ton active volume. In 2022, LZ collaboration published a world leading limit on WIMP dark matter interactions with nucleons. The success of the LZ experiment hinges both on the resilient design of its hardware and software infrastructures. This talk will give an overview of the offline software infrastructure of the LZ experiment, which includes the automated movement of the data and real time processing at NERSC, using its foremost HPC machine, Perlmutter. Additionally, I will talk about the monitoring tools and web services that enable the management, and operation of LZ's data workflow and cataloging.

#### Parallel (Track 5) / 425

### Quantum-Assisted Generative AI for Simulation of the Calorimeter Response

**Authors:** Abhishek Abhishek<sup>None</sup>; Abhishek Abhishek<sup>None</sup>; Alison Lister<sup>1</sup>; Colin Warren Gay<sup>1</sup>; Deniz Sogutlu<sup>2</sup>; Eric Paquet<sup>3</sup>; Hao Jia<sup>1</sup>; J. Quetzalcoatl Toledo-Marin<sup>4</sup>; Maximilian J Swiatlowski<sup>5</sup>; Roger Melko<sup>None</sup>; Sebastian Gonzalez<sup>4</sup>; Sehmimul Hoque<sup>6</sup>; Soren Andersen<sup>7</sup>; Wojtek Fedorko<sup>4</sup>

<sup>1</sup> University of British Columbia (CA)

<sup>2</sup> TRIUMF, UBC

<sup>3</sup> National Research Council Canada, University of Ottawa

<sup>4</sup> TRIUMF

<sup>5</sup> TRIUMF (CA)

<sup>6</sup> University of Waterloo, Perimeter Institute

<sup>7</sup> TRIUMF, Lund University

**Corresponding Authors:** rgmelko.uwaterloo@gmail.com, wfedorko@triumf.ca, akajal@triumf.ca, j.toledo.mx@gmail.com, eric.paquet@nrc-cnrc.gc.ca, cgay@physics.ubc.ca, soeren.andersen@hotmail.fr, alison.lister@cern.ch, maximilian.j.swiatlowski@cern.ch, abhishek@myumanitoba.ca, s4hoque@uwaterloo.ca, haojia@phas.ubc.ca, bastigonza2000@gmail.com, dsogutlu@triumf.ca

As CERN approaches the launch of the High Luminosity-LHC Large Hadron Collider (HL-LHC) by the decade's end, the computational demands of traditional simulations have become untenably high. Projections show millions of CPU-years required to create simulated datasets - with a substantial fraction of CPU time devoted to calorimetric simulations. This presents unique opportunities for breakthroughs in computational physics. We show how Quantum-assisted Generative AI can be used for the purpose of creating synthetic, realistically scaled calorimetry dataset. The model is constructed by combining D-Wave's Quantum Annealer processor with a Deep Learning architecture, increasing the timing performance with respect to first principles simulations and Deep Learning models alone, while maintaining current state-of-the-art data quality.

Poster session / 426

## Comprehensive Description and Orchestration of Complex Data Processing Pipelines

**Author:** Vardan Gyurjyan<sup>None</sup>

**Corresponding Author:** gurjyan@jlab.org

Managing and orchestrating complex data processing pipelines require advanced systems capable of handling diverse and collaborative components, such as data acquisition, streaming, aggregation, event identification, distribution, detector calibration, processing, analytics, and archiving. This paper introduces a data processing workflow description and orchestration system designed to facilitate the coordination and operation of these components using both centralized orchestration and decentralized choreography approaches. Our system employs a decentralized actor-based model used in the data acquisition system and data stream processing framework at Jefferson Lab (JLAB) to create component-specific configurations for the effective choreography of component actors. Simultaneously, the centralized orchestration provides global control and management of the entire data processing pipeline from acquisition to final processing. Our system's core is an ontology language developed explicitly for serializing data processing pipeline descriptions. A user-friendly graphical interface also enables seamless data pipeline composition and real-time monitoring. This integrated approach ensures efficient deployment, management, and orchestration of data processing workflows, ensuring robustness and flexibility in handling complex scientific data workflows.

Poster session / 427

## Development of a WebRTC-Based Remote Desktop for Neutron Scattering Data Analysis on Openstack

**Author:** 李亚康 liyk<sup>None</sup>

**Corresponding Author:** liyk@ihep.ac.cn

In neutron scattering experiments, the complexity of data analysis and the demand for computational resources have significantly increased. To address these challenges, we have developed a remote desktop system for neutron scattering data analysis based on the Openstack platform. This system leverages WebRTC technology to build a push-pull streaming service system, which includes the development of modules for screen capture, stream transmission, and GPU-based screen rendering.



By integrating and optimizing these technical modules, we have realized an interactive web-based remote desktop tailored for neutron scattering users.

Firstly, we set up a virtualized environment on Openstack to provide flexible resource management and high availability. Utilizing WebRTC technology, we achieved efficient screen capture and transmission, ensuring real-time and smooth remote desktop operations. Furthermore, the GPU-based screen rendering module significantly enhances the performance of image processing and display, meeting the high computational demands of neutron scattering data analysis.

The final implementation of this system not only provides an efficient, interactive remote data analysis platform but also significantly improves the efficiency of neutron scattering data analysis. Users can perform real-time data processing and analysis through a web interface without the need for local high-performance computing devices. This innovative solution offers new possibilities for neutron scattering data analysis and lays a solid foundation for future scientific research.

Poster session / 428

## Likelihood analysis methods for the PID system of STCF experiment

Author: Binbin Qi<sup>1</sup>

<sup>1</sup> University of Science and Technology of China (USTC)

Corresponding Author: qbb180@ustc.edu.cn

The Super Tau-Charm Facility (STCF) is a proposed electron-positron collider in China, designed to achieve a peak luminosity exceeding  $0.5 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$  and a center-of-mass energy ranging from 2 to 7 GeV. To meet the particle identification (PID) requirements essential for the physics goals of the STCF experiment, a dedicated PID system is proposed to identify  $\pi/K$  at momenta up to 2 GeV/c. This system comprises a Ring Imaging Cherenkov (RICH) detector for the barrel region and a time-of-flight detector using internally reflected Cherenkov light (DToF) for the endcap region. In this report, we introduce likelihood analysis methods to evaluate the PID performance of both the RICH and DToF detectors within the STCF offline software framework. These methods utilize photon 2D hit maps, with 2D spatial positions for the RICH and time-position patterns for the DToF, to distinguish particles. Furthermore, we present two distinct analytical approaches tailored for each detector, facilitating the rapid extraction of photon hit patterns for comprehensive likelihood analysis.

Parallel (Track 4) / 429

## Preparation for Multi-Site Processing at the Vera C. Rubin Observatory

Authors: Brian Yanny<sup>1</sup>; Edward Karavakis<sup>2</sup>; Fabio Hernandez<sup>3</sup>; Jennifer Adelman-Mccarthy<sup>4</sup>; Michelle Gower<sup>None</sup>; Peter Love<sup>5</sup>; Stephen R. Pietrowicz<sup>6</sup>; Tim Jenness<sup>7</sup>; Timothy John Noble<sup>8</sup>; Wei Yang<sup>9</sup>; Wen Guan<sup>2</sup>; Zhaoyu Yang<sup>2</sup>; kian tat lim<sup>10</sup>; richard dubois<sup>None</sup>

<sup>1</sup> Fermilab

<sup>2</sup> Brookhaven National Laboratory (US)

<sup>3</sup> IN2P3 / CNRS computing centre

<sup>4</sup> Fermi National Accelerator Lab. (Fermilab)-Unknown-Unknown

<sup>5</sup> Lancaster University (GB)

<sup>6</sup> National Center for Supercomputing Applications, USA

<sup>7</sup> Vera C. Rubin Observatory, USA

<sup>8</sup> Science and Technology Facilities Council STFC (GB)

<sup>9</sup> *SLAC National Accelerator Laboratory (US)*

<sup>10</sup> *SLAC National Accelerator Lab/Vera C. Rubin Observatory*

**Corresponding Authors:** richard@slac.stanford.edu, timothy.noble@stfc.ac.uk, peter.love@cern.ch, edward.karavakis@cern.ch, zhaoyu.yang@cern.ch, mgower@illinois.edu, ktl@slac.stanford.edu, yangw@slac.stanford.edu, jennifer.adelman-mccarthy@cern.ch, wen.guan@cern.ch, srp@illinois.edu, yanny@fnal.gov, fabio.hernandez@cern.ch, tjenness@lsst.org

After several years of focused work, preparation for Data Release Production (DRP) of the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST) at multiple data facilities is taking its shape. Rubin Observatory DRP features both complex, long workflows with many short jobs, and fewer long jobs with sometimes unpredictably large memory usage. Both of them create scaling issues that need to be addressed in order to meet the annual processing timeline.

This paper summarizes the infrastructure and services deployed at Rubin data facilities to support multi-site data processing. Rubin selected PanDA (Production and Distributed Analysis) to orchestrate its complex workflow and to manage its distributed workload. We address the interface between workflow/workload management system and Rubin's campaign management system, as well as the associated analytics platform, and the interface to the observatory's data management system.

Rubin has already exercised this infrastructure to process data from other observatories as well as simulated data. The experience of those processing campaigns is summarized in this paper. Finally, this paper outlines future plans, including providing the campaign management team a higher level view on ongoing campaigns and analyzing finished campaigns as well as using PanDA to support end users' need for batch processing from within a "hybrid" cloud approach to data hosting.

**Poster session / 430**

## ProtoDUNE Run Conditions Database

**Author:** Nilay Bostan<sup>1</sup>

**Co-author:** Ana Paula Vizcaya Hernandez<sup>2</sup>

<sup>1</sup> *University of Iowa (US) & TENMAK (TR)*

<sup>2</sup> *Colorado State University (US)*

**Corresponding Authors:** ana.vizcaya\_hernandez@colostate.edu, nilay.bostan@cern.ch

The DUNE experiment will produce vast amounts of metadata, which describe the data coming from the read-out of the primary DUNE detectors. Various databases will collect the metadata from different sources. The conditions data, which is the subset of all the metadata that is accessed during the offline reconstruction and analysis, will be stored in a dedicated database. ProtoDUNE at CERN is the largest DUNE far detector prototype, and as such serves to prove database solutions and schemas for DUNE.

The ProtoDUNE Run Conditions Database is a PostgreSQL relational database that stores the conditions metadata coming from sources such as: DAQ, Slow Control, and Beam databases. This contribution will summarize the Run Conditions Database infrastructure which consists of a python rest API users' interface, a C++ interface, an Art interface (which is the framework used for the offline LArTPC data processing), and a plug in to the new data catalog (MetaCat). We will present how the conditions data, coming from the slow controls database, is retrieved, studied, and stored in a convenient format.

**Parallel (Track6) / 431**

## LHCb Stripping Project: Continuing to Fully and Efficiently Utilize Legacy Data

**Authors:** Aravindhan Venkateswaran<sup>1</sup>; Eduardo Rodrigues<sup>2</sup>; Federico Leo Redi<sup>3</sup>; Mark Smith<sup>4</sup>; Nathan Grieser<sup>5</sup>; Nicole Skidmore<sup>6</sup>; Niladri Sahoo<sup>7</sup>; Shuqi Shenq<sup>8</sup>

<sup>1</sup> EPFL - Ecole Polytechnique Federale Lausanne (CH)

<sup>2</sup> University of Liverpool (GB)

<sup>3</sup> Università degli Studi e INFN Milano (IT)

<sup>4</sup> Imperial College (GB)

<sup>5</sup> University of Cincinnati (US)

<sup>6</sup> University of Warwick

<sup>7</sup> University of Birmingham (GB)

<sup>8</sup> IHEP (CN)

**Corresponding Authors:** niladri.sahoo@cern.ch, federico.redi@cern.ch, eduardo.rodrigues@cern.ch, nicola.skidmore@cern.ch, mark.elliott.smith@cern.ch, nathan.allen.grieser@cern.ch, aravindhan.venkateswaran@cern.ch

The LHCb collaboration continues to primarily utilize the Run 1 and Run 2 legacy datasets well into Run 3. As the operational focus shifts from the legacy data to the live Run 3 samples, it is vital that a sustainable and efficient system is in place to allow analysts to continue to profit from the legacy datasets. The LHCb Stripping project is the user-facing offline data-processing stage that allows analysts to select their physics candidates of interest simply using a Python-configurable architecture. After physics selections have been made and validated, the full legacy datasets are then reprocessed in small time windows known as Stripping campaigns.

Stripping campaigns at LHCb are characterized by a short development window with a large portion of collaborators, often junior researchers, directly developing a wide variety of physics selections; the most recent campaign dealt with over 900 physics selections. Modern organizational tools, such as GitLab Milestones, are used to track all of the developments and ensure the tight schedule is adhered to by all developers across the physics working groups. Additionally, continuous integration is implemented within GitLab to run functional tests of the physics selections, monitoring rates and timing of the different algorithms to ensure operational conformity. Outside of these large campaigns the project is also subject to nightly builds, ensuring the maintainability of the software when parallel developments are happening elsewhere.

Parallel (Track 3) / 432

## Advancements in Offline Analysis Software for Run III at LHCb

**Author:** Abhijit Mathad<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Author:** abhijit.mathad@cern.ch

As the Large Hadron Collider progresses through Run 3, the LHCb experiment has made significant strides in upgrading its offline analysis framework and associated tools to efficiently handle the increasing volumes of data generated. Numerous specialised algorithms have been developed for offline analysis, with a central innovation being FunTuple—a newly developed component designed to effectively compute and store offline data. Built upon the robust Gaudi functional framework, FunTuple merges a user-friendly Python interface with a flexible templated design. This modern architecture supports a wide range of data types, including both reconstructed and simulated events, facilitating processing of event-level and decay-level information. Crucially, FunTuple is primed for future enhancements to integrate new event models, optimising vectorised data processing across heterogeneous resources.

A pivotal feature of FunTuple is its capability to align trigger-computed observables with those analysed offline, crucial for maintaining data integrity across LHCb analyses. This alignment is achieved through Throughput Oriented (ThOr) functors, specifically crafted to meet the high throughput demands of the trigger system. Moreover, FunTuple offers comprehensive customisation options, enabling users to define and store tailored observables within ROOT files in anticipation of future increases in data volumes. FunTuple has undergone rigorous testing, including numerous unit tests and pytest evaluations. In 2024, it is undergoing a comprehensive stress test by hundreds of analysts to validate its reliability in managing and validating the quality of data recorded by LHCb.

This presentation will delve into the design, user interface, and integration of FunTuple alongside other analysis components, showcasing their efficiency and reliability through detailed performance metrics in managing large-scale data.

Parallel (Track 8) / 433

## Leveraging workflow engines and computing frameworks for physics analysis scalability and reproducibility

**Author:** Mindaugas Sarpis<sup>1</sup>

<sup>1</sup> *University of Manchester (GB)*

**Corresponding Author:** mindaugas.sarpis@cern.ch

With the onset of ever more data collected by the experiments at the LHC and the increasing complexity of the analysis workflows themselves, there is a need to ensure the scalability of a physics data analysis. Logical parts of an analysis should be well separated - the analysis should be modularized. Where possible, these different parts should be maintained and reused for other analyses or reinterpretation of the same analysis.

Also, having an analysis prepared in such a way helps to ensure its reproducibility and preservation in the context of good data and analysis code management practices following the FAIR principles. In this talk, a few different topics on analysis modularization are discussed. An analysis on searches for pentaquarks within the LHCb experiment at CERN is used as an example.

Parallel (Track 7) / 434

## Provisioning of Grid computing resources in the Norwegian Research and Education Cloud

**Author:** Matthias Richter<sup>1</sup>

**Co-authors:** Emmanuel Moutoussamy ; Maiken Pedersen <sup>2</sup>; Raymond Kristiansen <sup>3</sup>; Terje Furenes

<sup>1</sup> *University of Bergen (NO)*

<sup>2</sup> *University of Oslo (NO)*

<sup>3</sup> *University of Bergen, Norway*

**Corresponding Authors:** emmanuel.moutoussamy@uib.no, matthias.richter@cern.ch, maiken.pedersen@fys.uio.no, terje.furenes@uib.no, raymond.kristiansen@uib.no

Norwegian contributions to the WLCG consist of computing and storage resources in Bergen and Oslo for the ALICE and ATLAS experiments. The increasing scale and complexity of Grid site infrastructure and operation require integration of national WLCG resources into bigger shared installations. Traditional HPC resources often come with restrictions with respect to software, administration, and accessibility. Furthermore, expensive HPC infrastructure like fast interconnects is hardly used by grid workload.

As a cost-efficient solution, the Norwegian Grid resources are operated as two platforms within NREC, the Norwegian Research and Education Cloud, which is a cloud computing service operated by the Universities of Oslo and Bergen. It aims to provide easily accessible computing and storage infrastructure for national academic and scientific applications.

By using cloud technology instead of traditional HPC resources, WLCG installations benefit from a high degree of accessibility, flexibility, and scalability while the service provider ensures reliable and secure operation of infrastructure and network.

Orchestration of the virtual instances is based on the Infrastructure-as-a-service paradigm and implemented as declarative configuration files in Terraform. All custom host configuration, software deployment and cluster configuration are implemented as YAML code and deployed using Ansible.

This concept allows for the delivery of high-quality WLCG services with key features such as: fixed and opportunistic computing resources; ARC and JAliEn grid middleware; Slurm and HTCondor backend; CEPH disk storage integrated into Neic NDGF dCache; integrated tape storage; monitoring and alerting based on Prometheus/Grafana ecosystem; fully controlled setup by site admin; scalable extension; quick failover and recovery.

This presentation describes the capabilities of the Norwegian Research and Education Cloud and the strategy for provisioning of Grid computing and storage using the IaaS approach. Details on cluster management and monitoring as a service, flexible cluster orchestration, scalability and performance studies will be highlighted in the presentation.

Poster session / 435

## Track classification in Active Target Time Projection Chamber using supervised machine learning methods

**Author:** Pralay Kumar Das<sup>None</sup>

**Co-authors:** Nayana Majumdar<sup>1</sup>; Supratik Mukhopadhyay

<sup>1</sup> *Saha Institute of Nuclear Physics*

**Corresponding Authors:** [nayana.majumdar@saha.ac.in](mailto:nayana.majumdar@saha.ac.in), [daspralay11@gmail.com](mailto:daspralay11@gmail.com), [supratik.mukhopadhyay@saha.ac.in](mailto:supratik.mukhopadhyay@saha.ac.in)

In low-energy nuclear physics experiments, an Active Target Time Projection Chamber (AT-TPC) 1 can be advantageous for studying nuclear reaction kinematics. The  $\alpha$ -cluster decay of  $^{12}\text{C}$  is one such reaction requiring careful investigation due to its vital role in producing heavy elements through astrophysical processes 2. The breakup mechanism of the Hoyle state, a highly  $\alpha$ -clustered state at 7.65 MeV in  $^{12}\text{C}$ , has long been an important case of study using different experimental techniques to investigate its decaying branch ratio. The direct decay of the Hoyle state into three  $\alpha$ -particles and the sequential decay via the  $^8\text{Be}$  ground state into three  $\alpha$ -particles can be identified by tracking the  $\alpha$ -particles and measuring their energies, which can be accomplished with AT-TPC. In this work, a numerical model using Hough transformation and neural network models for track separation and classification has been developed for identifying the breakup of Hoyle state  $^{12}\text{C}$  into three  $\alpha$ -particles from the background scattering events in the active volume. The reaction kinematics of the decay of  $^{12}\text{C}$  have been determined using low-energy non-relativistic scattering of  $\alpha$ -particles on  $^{12}\text{C}$ . The event tracks in the active gas medium of the AT-TPC have been generated through primary ionization created by the  $\alpha$ -particles produced in the aforementioned reaction. This has been carried out using the Geant4 3 simulation framework. The three  $\alpha$ -tracks have then been individually separated from the others and the background (uniform random noise generated) using the Hough transformation. Each  $\alpha$ -track has been fitted with a CrystalBall function to extract the features useful for training the Artificial Neural Network (ANN) model. A Convolutional Neural Network (CNN) model has also been developed to identify all possible scattering events of  $^4\text{He}$  and  $^{12}\text{C}$  in the lab frame which are labeled as background events. These events have been further labeled by binary and multi-class classification models, which have been trained using simulated data. For this purpose, a fully connected ANN and CNN with hidden layers have been implemented using

the high-level deep learning library Keras, written in Python [4]. The model has been tested on the events generated using simulation. Thus, it has been possible to precisely classify the Hoyle state  $\alpha$ -particles from background events. This model can also be beneficial as an automated analysis framework for tagging and separating events from experimental data.

1 D. Suzuki et al., NIM A 691(2012) 39–54

2 M. Freer, Reports on Progress in Physics 70, 2149 (2007)

3 GEANT4 collaboration, GEANT4: A Simulation toolkit 2003 NIM A 506 (2003) 250-303

[4] F. Chollet et al., Keras, <https://keras.io>, 2015.

**Parallel (Track 9) / 436**

## Tuning the CMS Coffea-casa facility for 200 Gbps Challenge

**Authors:** Andrew Wightman<sup>1</sup>; Benjamin Tovar Lopez<sup>2</sup>; Brian Paul Bockelman<sup>3</sup>; Carl Lundstedt<sup>1</sup>; Derek Weitzel<sup>1</sup>; Garhan Attebury<sup>1</sup>; John Thiltges<sup>1</sup>; Kenneth Bloom<sup>1</sup>; Oksana Shadura<sup>1</sup>; Sam Albin<sup>4</sup>

<sup>1</sup> *University of Nebraska Lincoln (US)*

<sup>2</sup> *University of Notre Dame*

<sup>3</sup> *University of Wisconsin Madison (US)*

<sup>4</sup> *UNL*

**Corresponding Authors:** btovar@nd.edu, oksana.shadura@cern.ch, clundstedt@unl.edu, kenbloom@unl.edu, andrew.steven.wightman@cern.ch, brian.bockelman@cern.ch, john.thiltges@cern.ch, garhan.attebury@cern.ch, sam.albin@unl.edu, derek.weitzel@cern.ch

As a part of the IRIS-HEP “Analysis Grand Challenge” activities, the Coffea-casa AF team executed a “200 Gbps Challenge”. One of the goals of this challenge was to provide a setup for execution of a test notebook-style analysis on the facility that could process a 200 TB CMS NanoAOD dataset in 20 minutes.

We describe the solutions we deployed at the facility to execute the challenge tasks. The facility was configured to provide 2000+ cores for quick turn-around, low-latency analysis. To reach the highest event processing rates we tested different scaling backends, both scaling over HTCondor and Kubernetes resources and using Dask and Taskvine schedulers. This configuration also allowed us to compare two different services for managing Dask clusters, Dask labextension, and Dask Gateway server, under extreme conditions.

A robust set of XCache servers with a redirector were deployed in Kubernetes to cache the dataset to minimize wide-area network traffic. The XCache servers were backed with solid-state NVME drives deployed within the Kubernetes cluster nodes. All data access was authenticated using scitokens and was transparent to the user.

To ensure we could track and measure data throughput precisely, we used our existing Prometheus monitoring stack to monitor the XCache pod throughput on the Kubernetes network layer. Using the rate query across all of the 8 XCache pods we were able to view a stacked cumulative graph of the total throughput for each XCache. This monitoring setup allowed us to ensure uniform data rates across all nodes while verifying we had reached the 200 Gbps benchmark.

**Poster session / 437**

## Integration of A Supercomputing Center in LHAASO Distributed Computing System

**Author:** Xiaowei Jiang<sup>1</sup>

**Co-authors:** Chaoqi Guo <sup>2</sup>; Jingyan Shi <sup>1</sup>; Yaodong CHENG <sup>3</sup>; Yaodong Cheng <sup>4</sup>; Yujiang BI <sup>5</sup>

<sup>1</sup> *Chinese Academy of Sciences (CN)*

<sup>2</sup> *the Institute of High Energy Physics (IHEP) of Chinese Academy of Sciences*

<sup>3</sup> *IHEP, Beijing*

<sup>4</sup> *IHEP, CAS*

<sup>5</sup> *Institute of High Energy Physics, Chinese Academy of Sciences*

**Corresponding Authors:** jingyan.shi@cern.ch, biyujiang@ihep.ac.cn, guocq@ihep.ac.cn, yaodong.cheng@cern.ch, xiaowei.jiang@cern.ch, ycheng@cern.ch

LHAASO experiment is a new generation multi-component experiment designed to study cosmic rays and gamma-ray astronomy. The data volume from LHAASO are currently reaching to ~40PB and ~11PB of new data will be generated every year in the future. Such scale of data needs a big scale of computing resources to process. For LHAASO experiment, there are several types of computing sites to join the LHAASO distributed computing system, that a supercomputing center is one of them. To adding a supercomputing center, several components are developed: 1. a site job agent based on HTCondor is deployed in front of the batch system of the site; 2. two proxy services are respectively responsible to job schedule and data transfer from central side to the site across the firewall of supercomputing center; 3. a service synchronizes the necessary data to a shared filesystem of the site, including the software. 4. job running environment is wrapped inside a container. The services have been deployed in a Chengdu Supercomputing Center and the test results indicate an acceptable performance on job schedule (thousands of jobs) and data transfer (tens of TB).

**Parallel (Track 4) / 438**

## Designing Operational Security systems: People, Processes and Technology

**Authors:** David Crooks<sup>1</sup>; Liviu Valsan<sup>2</sup>; Jeny Lucia Teheran Sierra<sup>3</sup>

<sup>1</sup> *UKRI STFC*

<sup>2</sup> *CERN*

<sup>3</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** david.crooks@cern.ch, jteheran@fnal.gov, liviu.valsan@cern.ch

The risk of cyber attack against members of the research and education sector remains persistently high, with several recent high visibility incidents including a well-reported ransomware attack against the British Library. As reported previously, we must work collaboratively to defend our community against such attacks, notably through the active use of threat intelligence shared with trusted partners both within and beyond our sector.

We discuss the development of capabilities to defend sites across the WLCG and other research and education infrastructures, with a particular focus on sites other than Tier1s which may have fewer resources available to implement full-scale security operations processes. These capabilities include a discussion of the pDNSSOC software which enables a lightweight and flexible means to correlate DNS logs with threat intelligence, and an examination of the use of Endpoint Detection and Response (EDR) tools in a high throughput context.

This report will include an important addition to the work of the Security Operations Centre Working Group; while this group had previously focused primarily on the technology stacks appropriate for use in deploying fine-grained security monitoring services, the people and processes involved with such capabilities are equally important.

Defending as a community requires a strategy that brings people, processes and technology together. We suggest approaches to support organisations and their computing facilities to defend against a wide range of threat actors. While a robust technology stack plays a significant role, it must be guided and managed by processes that make their cybersecurity strategy fit their environment.

**Poster session / 439**

## Simulation of a MPGD-based hadronic calorimeter for Muon Collider experiments

**Authors:** Raffaella Radogna<sup>1</sup>; Anna Stammera<sup>1</sup><sup>1</sup> *Universita e INFN, Bari (IT)***Corresponding Authors:** raffaella.radogna@cern.ch, a.stammera@studenti.uniba.it

The proposal to create a multi-Tev Muon Collider presents an unprecedented opportunity for advancing high energy physics research and offers the possibility to accurately measure the Higgs couplings with other Standard Model particles and search for new physics at TeV scale.

This demands for accurate full event reconstruction and particle identification. However, this is complicated by the beam induced background (BIB), originating from the decays of the muons of the beam, that represents one of the major challenges for the experiment design and that poses potential limitations on the detector performances and requirements on radiation hardness. The discrimination of signal showers from the BIB requires high granularity, superb energy resolution and precise timing. The calorimeter should thus provide 5D measurement (3D position, time and energy).

To address these challenges, an innovative hadronic calorimeter has been designed that utilizes Micro Pattern Gas Detectors (MPGDs) as active layers. MPGDs are ideal for high radiation environments and offer high granularity for precise spatial measurements. The response of such MPGD-based HCAL to the incoming particles is studied and presented in this contribution with Monte Carlo simulations performed using GEANT4, comparing the performance of a digital and semi-digital readout and considering the energy resolution as a figure of merit.

This contribution details the design and optimization of the MPGD-based hadronic calorimeter, and shows a comparison of the simulated performance with preliminary experimental data. This project is endorsed by the International Muon Collider Collaboration.

**Poster session / 440**

## Belle II Network Performance Analysis in the Context of the WLCG Data Challenge 2024

**Author:** Silvio Pardi<sup>1</sup>**Co-authors:** Cedric Serfon<sup>2</sup>; Hideki Miyake<sup>3</sup>; Hiroaki Ono ; I Ueda<sup>4</sup>; Michel Hernandez Villanueva<sup>2</sup>; Takanori Hara<sup>5</sup><sup>1</sup> *University Federico II and INFN, Naples (IT)*<sup>2</sup> *Brookhaven National Laboratory (US)*<sup>3</sup> *KEK*<sup>4</sup> *KEK IPNS*<sup>5</sup> *High Energy Accelerator Research Organization (JP)***Corresponding Authors:** cedric.serfon@cern.ch, ueda@post.kek.jp, silvio.pardi@cern.ch, takanori.hara@kek.jp, ono@ngt.ndu.ac.jp, mhernande1@bnl.gov, hideki.miyake@kek.jp

The Belle II experiment relies on a distributed computing infrastructure spanning 19 countries and over 50 sites. It is expected to generate approximately 40TB/day of raw data in 2027, necessitating distribution from the High Energy Accelerator Research Organization (KEK) in Japan to six Data Centers across the USA, Europe, and Canada. Establishing a high-quality network has been a priority since 2012 to address the challenge of transferring data across long distances in high-latency environments. This effort has included joining LHCONE and conducting periodic Data Challenges to assess network performance following significant changes in infrastructure or experiment schedules.



In February 2024 Belle II joined the WLCG Data Challenge, performed together with LHC experiments with the goal to test network performance under stress, particularly due to the anticipated increase in traffic from the experiments with the High Luminosity LHC program at CERN.

In this work, we will present a comprehensive overview of the tests conducted by Belle II. We will start with the test design, define the goals, and outline the preliminary steps taken. We will then describe the working environment, the tests performed, and the tools employed. Furthermore, we will discuss the results achieved in detail. Finally, we will outline the future steps for subsequent tests.

**Parallel (Track 1) / 441**

## The Belle II Raw Data Transfer System

**Authors:** Dhiraj Kalita<sup>None</sup>; Matt Barrett<sup>1</sup>; Takanori Hara<sup>2</sup>; Tristan Bloomfield<sup>3</sup>

<sup>1</sup> KEK

<sup>2</sup> High Energy Accelerator Research Organization (JP)

<sup>3</sup> KEK IPNS

**Corresponding Authors:** takanori.hara@kek.jp, dhirajk@post.kek.jp, tristan.bloomfield@kek.jp, mattbarrett31@gmail.com

The Belle II raw data transfer system is responsible for transferring raw data from the Belle II detector to the local KEK computing centre, and from there to the GRID. The Belle II experiment recently completed its first Long Shutdown period - during this time many upgrades were made to the detector and tools used to handle and analyse the data. The Belle II data acquisition (DAQ) systems received significant improvements, necessitating changes in the processing steps for raw data. Furthermore, experience gained during Run 1 identified areas where the scalability of the system could be improved to better handle the expected increase in data rates in future years.

To address these issues, extensive upgrades were made to the raw data transfer system, including: utilisation of the DIRAC framework for all data transfers; a change in the protocol used to communicate with the DAQ systems; and retirement of the previously used file format conversion component of the system. This talk will describe these changes and improvements in detail, and give an overview of the current state of the Belle II raw data transfer system.

**Poster session / 442**

## Kalman filter for muon reconstruction in the CMS Phase-2 endcap calorimeter

**Authors:** Erica Brondolin<sup>1</sup>; Marco Rovere<sup>1</sup>; Mark Nicholas Matthewman<sup>2</sup>; Wolfgang Waltenberger<sup>3</sup>

<sup>1</sup> CERN

<sup>2</sup> HEPHY

<sup>3</sup> Austrian Academy of Sciences (AT)

**Corresponding Authors:** wolfgang.waltenberger@gmail.com, mark.matthewman@cern.ch, erica.brondolin@cern.ch, marco.rovere@cern.ch

The High Luminosity phase of the LHC (HL-LHC) will offer a greatly increased number of events for more precise standard model measurements and BSM searches. To cope with the harsh environment created by numerous simultaneous proton-proton collisions, the CMS Collaboration has begun construction of a new endcap calorimeter, the High-Granularity Calorimeters (HGCal). As part of this project, a new reconstruction framework, TICL, is being developed, aiming to exploit the possibilities of heterogeneous computing, and employing machine learning elements.

While TICL has shown impressive results for particle shower reconstruction in HGAL, the proposed calorimeters' high granularity can be used to track muons. Precise tracking of externally identified muons through the calorimeter allows them to be used for the crucial task of following the evolving inter-cell relative response, and calibrating it, in order to maintain good energy resolution.

In this contribution, we propose to integrate a Kalman Filter into the TICL framework for dedicated muon reconstruction. Furthermore, we present a comprehensive performance evaluation of the algorithm under various conditions akin to those at the HL-LHC. Finally, we discuss the capabilities and limitations of the Kalman Filter as a tool for inter-cell calibration.

**Parallel (Track 8) / 443**

## Enhancing Data Management in Nuclear Physics and Cross-Domain Metadata Integration

**Authors:** Andrew Kishor Mistry<sup>1</sup>; Ivan Knezevic<sup>2</sup>

<sup>1</sup> *GSI Helmholtzzentrum für Schwerionenforschung GmbH(GSI)*

<sup>2</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

**Corresponding Authors:** ivan.knezevic@cern.ch, a.k.mistry@gsi.de

The metadata schema for experimental nuclear physics project aims to facilitate data management and data publication under the FAIR principles in the experimental Nuclear Physics communities, by developing a cross-domain metadata schema and generator, tailored for diverse datasets, with the possibility of integration with other, similar fields of research (i.e. Astro and Particle physics).

Our project focuses on creating a standardized, adaptable framework that enhances data Findability, Accessibility, Interoperability, and Reusability (FAIR principles). By creating a comprehensive and adaptable metadata schema, the project ensures scalable integration of both machine and human-readable metadata, thereby improving the efficiency of data discovery and utilization.

A pivotal component of the project is its nodal, multi-layered schema structure, allowing metadata enrichment from multiple domains while maintaining essential overlaps for enhanced versatility. This comprehensive approach supports the unification of data standards across various research institutions, promoting interoperability and collaboration on a European scale. Our efforts also extend to the development of a user-friendly frontend generator, designed not only to facilitate metadata input but also to allow users to specify field-specific attributes, customize generic names to suit their needs, and export schemas in various formats such as JSON and XML, adhering to different nomenclatures.

The project involves world-class RIs and ESFRIs, and leverages synergies from existing Open Science initiatives like EOSC, ESCAPE, EURO-LABS, and PUNCH4NFDI. In this contribution, we will present an overview of the project, detailing the development steps, key features of the metadata schema, and the functionality of the frontend generator.

**Poster session / 444**

## Exploiting Kubernetes to Simplify the Deployment and Management of the Multi-purpose CMS Pilot Job Factory

**Authors:** Jeffrey Michael Dost<sup>1</sup>; Marco Mascheroni<sup>1</sup>

**Co-authors:** Brian Paul Bockelman<sup>2</sup>; Colby Walsworth<sup>3</sup>; Edita Kizinevic<sup>4</sup>; John Thiltges<sup>5</sup>; Merina Albert<sup>6</sup>; Vaiva Zokaite<sup>7</sup>

<sup>1</sup> *Univ. of California San Diego (US)*

<sup>2</sup> *University of Wisconsin Madison (US)*

<sup>3</sup> *University of California San Diego*

<sup>4</sup> *CERN*

<sup>5</sup> *University of Nebraska Lincoln (US)*

<sup>6</sup> *Fermi National Accelerator Lab. (US)*

<sup>7</sup> *Vilnius University (LT)*

**Corresponding Authors:** merina@fnal.gov, edita.kizinevic@cern.ch, jdost@ucsd.edu, marco.mascheroni@cern.ch, vaiva.zokaite@cern.ch, colby@sdsc.edu, brian.bockelman@cern.ch, john.thiltges@cern.ch

GlideinWMS, a widely utilized workload management system in high-energy physics (HEP) research, serves as the backbone for efficient job provisioning across distributed computing resources. It is utilized by various experiments and organizations, including CMS, OSG, Dune, and FIFE, to create HTCondor pools as large as 600k cores. In particular, a shared factory service historically deployed at UCSD has been configured to interface with more than 500 routes to compute clusters.

As part of our team's initiative to modernize infrastructure and enhance scalability, we undertook the migration of the glideinWMS factory service into the Kubernetes environment. Leveraging the flexibility and orchestration capabilities of Kubernetes, we successfully deployed the factory service within the OSG Tiger Kubernetes cluster. The major benefits Kubernetes gives us is it streamlines the management and monitoring of the factory infrastructure, and improves fault tolerance through its resilient deployment strategies.

Through this case study, we aim to share insights, challenges, and best practices encountered during the migration process. Our experience underscores the benefits of embracing containerization and Kubernetes orchestration for HEP computing infrastructure, paving the way for scalability and resilience in distributed computing environments.

**Parallel (Track 8) / 445**

## Geometry Simplification Methods for Virtual Reality Applications

**Author:** Alexander Sharmazanashvili<sup>1</sup>

**Co-authors:** Giorgi Mirziashvili<sup>1</sup>; Kote Tsutskiridze<sup>1</sup>; Ela Abramovi<sup>1</sup>; Avtandil Khelashvili<sup>2</sup>

<sup>1</sup> *Georgian Technical University (GTU)*

<sup>2</sup> *Georgian Technical University*

**Corresponding Authors:** avtandil.khelashvili@cern.ch, giorgi.mirziashvili@cern.ch, lasha.sharmazanashvili@cern.ch, kote.tsutskiridze@cern.ch, abramovi.e@gtu.ge

Virtual Reality (VR) applications play an important role in HEP Outreach & Education. They make it possible to organize virtual tours of the experimental infrastructure by virtually interacting with detector facilities, describing their purpose and functionalities. However, nowadays VR applications require expensive hardware, like the Oculus headset or MS Hololense, and powerful computers. As a result, this reduces the reach of VR application implementation and makes their benefits questionable. An important improvement to VR development is thus to facilitate the usage of inexpensive hardware, like Google cardboard and phones with average computational power.

Requirements to use inexpensive hardware and achieve quality and performance close to the advanced hardware bring challenges to the VR application developers. One of these challenges concerns the geometry of the 3D VR scenes. Geometry defines the quality of the 3D scenes and at the same time causes big loads on the GPU. Therefore, development methods of the geometry make it possible to find a good balance between the quality and performance of the VR applications.

The paper describes methods of the simplification of the "as-built" geometry descriptions; ways to

reduce the number of facets to meet the GPU limitations in performance, and ensure the smooth movement in the VR scenes.

**Parallel (Track 3) / 446**

## Quantum Machine Learning for track reconstruction

**Authors:** Laura Cappelli<sup>1</sup>; Matteo Argenton<sup>2</sup>

**Co-authors:** Concezio Bozzi<sup>1</sup>; Enrico Calore; Sebastiano Fabio Schifano<sup>3</sup>; Valentina Amitrano<sup>4</sup>

<sup>1</sup> *INFN Ferrara*

<sup>2</sup> *Universita e INFN, Ferrara (IT)*

<sup>3</sup> *University and INFN of Ferrara*

<sup>4</sup> *INFN*

**Corresponding Authors:** enrico.calore@fe.infn.it, laura.cappelli@fe.infn.it, matteo.argenton@cern.ch, schifano@fe.infn.it, concezio.bozzi@cern.ch, valentina.amitrano@fe.infn.it

Tracking charged particles in high-energy physics experiments is a computationally intensive task. With the advent of the High Luminosity LHC era, which is expected to significantly increase the number of proton-proton interactions per beam collision, the amount of data to be analysed will increase dramatically. As a consequence, local pattern recognition algorithms suffer from scaling problems.

In this work, we investigate the possibility of using machine learning techniques in combination with quantum computing. In particular, we represent particle trajectories as a graph data structures and train a quantum graph neural network to perform global pattern recognition. We show recent results on the application of this method, with scalability tests for increasing pileup values. We discuss the critical points and give an outlook of potential improvements and alternative approaches.

We also provide insights into various aspects of code development in different quantum programming frameworks such as PennyLane and IBM Qiskit.

**Parallel (Track 9) / 447**

## HEPS scientific computing system design for interactive data analysis scenarios

**Authors:** Qingbao Hu<sup>1</sup>; Wei Zheng<sup>1</sup>; Xiaofei Yan<sup>2</sup>; 吉平 xujp<sup>3</sup>

<sup>1</sup> *IHEP*

<sup>2</sup> *Institute of High Energy Physics*

<sup>3</sup> *中国科学院高能物理研究所 (IHEP)*

**Corresponding Authors:** zhengw@ihep.ac.cn, huqb@ihep.ac.cn, xujp@ihep.ac.cn, yanxf@ihep.ac.cn

China's High-Energy Photon Source (HEPS), the first national high-energy synchrotron radiation light source, is under design and construction. HEPS computing center is the principal provider of high-performance computing and data resources and services for science experiments of HEPS. The mission of HEPS scientific computing platform is to accelerate the scientific discovery for the characteristics of light source experiments through high-performance computing and data analysis. In order to meet the diverse analysis needs of data analysis in light source disciplines, we have built a scientific computing platform that can provide desktop analysis, interactive analysis, batch analysis and other types of computing services, and support scientists to access the computing environment

through the web anytime and anywhere, quickly analyze experimental data. In this article, a scientific computing platform for HEPS's diverse analysis requirements is designed. First, the diverse analysis requirements of HEPS is introduced. Second, the challenges faced by the HEPS scientific computing system. Third, the architecture and service process of the scientific computing platform are described from the perspective of user, and some key technical implementations will be introduced in detail. Finally, the application effect of computing platforms will be demonstrated.

Poster session / 448

## RICH ring reconstruction using machine learning for CBM

Author: Martin Beyer<sup>1</sup>

<sup>1</sup> *Justus-Liebig-Universität Gießen*

Corresponding Author: martin.beyer@physik.uni-giessen.de

The Compressed Baryonic Matter experiment (CBM) at FAIR is designed to explore the QCD phase diagram at high baryon densities with interaction rates up to 10 MHz using triggerless free-streaming data acquisition. For the overall PID, the CBM Ring Imaging Cherenkov detector (RICH) contributes by identifying electrons from lowest momenta up to 10 GeV/c, with a pion suppression of  $> 100$ . The RICH reconstruction combines a local Cherenkov ring finding with a ring-track matching of found rings and extrapolated tracks from the Silicon Tracking System (STS).

The existing conventional algorithms were revised and optimized, and alternative machine learning approaches were investigated. Methods based on CNN/GNN architectures were developed for ring finding, noise suppression and ring-track matching while taking into account the latency and data format (space and time, i.e. 3+1) constraints of the triggerless free-streaming readout. The methods were tested and validated on simulations taking into account the time data stream and on data from the prototype mini-RICH (mRICH) in the mini-CBM (mCBM) experiment, which shares the same free-streaming readout concept as the future CBM experiment.

Parallel (Track 4) / 449

## Distributed Computing Infrastructure for HERD experiment

Authors: Xiaowei Jiang<sup>1</sup>; Xuantong Zhang<sup>1</sup>

<sup>1</sup> *Chinese Academy of Sciences (CN)*

Corresponding Authors: xiaowei.jiang@cern.ch, xuantong.zhang@cern.ch

The High Energy cosmic-Radiation Detection (HERD) facility is an under construction space astronomy and particle astrophysics experiment in collaboration between China and Italy, and will run on the China Space Station for more than 10 years since 2027. HERD is designed to search for dark matter with unprecedented sensitivity, investigate the century-old mystery of the origin of cosmic rays, conduct high-sensitivity surveys and monitoring of high-energy gamma rays, and explore new methods of pulsar navigation.

Once operational, HERD experiment is expected to generate and distribute more than 90 PB of data over a span of 10 years. To share the data and computing resource for data processing and analysis, a distributed computing system based on grid computing technology is developed with the supports of computing and storage sites in China and Italy.

HERD distributed computing infrastructure (DCI) integrates a distributed computing system managing data processing jobs based on DIRAC and dHTC, a distributed data management system distributing raw and produced data based on Rucio, and other grid computing middleware including IAM, FTS3, etc. Furthermore, HERD DCI is designed to deeply involved with HERD Offline Software

(HERDOS) and data production workflow, providing grid computing services without normal and production users being aware of it.

Poster session / 450

## The CloudVeneto's Container-as-a-Service ecosystem

**Author:** Lisa Zangrando<sup>None</sup>

**Co-authors:** Antonino Troja ; Daniel Lupu ; Davide Marcato <sup>1</sup>; Federica Fanzago <sup>2</sup>; Marco Verlato <sup>3</sup>; Massimo Sgaravatto <sup>3</sup>; Sergio Traldi

<sup>1</sup> *Istituto Nazionale di Fisica Nucleare*

<sup>2</sup> *INFN Padova*

<sup>3</sup> *Universita e INFN, Padova (IT)*

**Corresponding Authors:** daniel.lupu@cern.ch, massimo.sgaravatto@pd.infn.it, davide.marcato@lnl.infn.it, antonino.troja@pd.infn.it, federica.fanzago@pd.infn.it, lisa.zangrando@pd.infn.it, marco.verlato@pd.infn.it, sergio.traldi@gmail.com

CloudVeneto is a distributed private cloud, which harmonizes the resources of two INFN units and the University of Padua. Tailored to meet the specialized scientific computing needs of user communities within these organizations, it promotes collaboration and enhances innovation. CloudVeneto basically implements an OpenStack based IaaS (Infrastructure-as-a-Service) cloud. However users are also provided with some higher level services, such as the CloudVeneto Container-as-a-Service (CaaS).

Unlike the user-managed Kubernetes-as-a-Service (KaaS) model, CaaS offers a fully managed orchestration platform, eliminating administrative overhead for users. CloudVeneto has developed a CaaS solution designed to provide a secure, multi-tenant Kubernetes platform, where users deploy application containers to our service without the burden of administrative tasks. Our solution features a centrally managed Kubernetes control plane, allowing users to create and customize nodes on demand using CloudVeneto resources within their projects. These nodes, deployed as Virtual Machines, integrate seamlessly into the cluster, giving users the flexibility to maintain node privacy or share them within their OpenStack projects, while delegating deployment and monitoring tasks to CaaS. By implementing CaaS, CloudVeneto meets the diverse requirements of the user community, including those of Quantum, ISOLPHARM and SPES, underlining its adaptability within the CloudVeneto ecosystem. Furthermore, we were able to successfully demonstrate that it is possible to offload part of the workload submitted to a remote external Kubernetes cluster to the CloudVeneto CaaS service using the interLink-sidecar solution implemented in the context of the EU interTwin project, effectively extending the Virtual Kubelet concept.

Parallel (Track 5) / 451

## BAT.jl, the Bayesian Analysis Toolkit in Julia

**Authors:** Allen Christopher Caldwell<sup>1</sup>; Cornelius Grunwald<sup>2</sup>; Kevin Alexander Kroening<sup>2</sup>; Oliver Schulz<sup>1</sup>; Salvatore La Cagnina<sup>3</sup>

<sup>1</sup> *Max-Planck-Institut für Physik*

<sup>2</sup> *Technische Universität Dortmund*

<sup>3</sup> *Technische Universität Dortmund*

**Corresponding Authors:** salvatore.lacagnina@udo.edu, oliver.schulz@cern.ch, allen.christopher.caldwell@cern.ch, kevin.alexander.kroening@cern.ch, cornelius.grunwald@cern.ch

The Bayesian Analysis Toolkit in Julia (BAT.jl) is an open source software package that provides user-friendly tooling to tackle statistical problems encountered in Bayesian (an not just Bayesian) inference.

BAT.jl succeeds the very successful BAT-C++ (over 500 citations) using modern Julia language. We chose Julia because of its high performance, native automatic differentiation, support for parallel CPU/GPU computing and state-of-the-art package management. BAT.jl is developed in cooperation with the Julia HEP community and is also the basis for the prototype Julia implementation of the HEP Statistics Serialization Standard (HS3), opening a path for direct compatibility with RooFit, HistFactory and pyHF models. EFTFitter is based on BAT.jl as well.

BAT.jl is intended both for quick-and-easy inference but also for use cases that cannot easily be expressed in a domain-specific language (like STAN) or that are computationally costly or require interfacing with existing C/FORTRAN models. A recent application has been inference of parton PDFs from HERA data, combining BAT.jl in Julia with the FORTRAN QCDNUM package.

BAT.jl provides a range of posterior sampling algorithms like Metropolis-MCMC, HMC, MGVI and nested sampling under a common API, as well as methods for evidence estimation, with added tooling for quick plotting, reporting and exporting results.

Recently, we have added more concepts from measure theory, building on BATs capability of using pre-computed normalizing flows to transform problems into spaces optimized for each given algorithm. We also have preliminary support for machine-learned normalizing flows as additional posterior-transform and -approximation tools.

**Parallel (Track 4) / 452**

## Evolving INDIGO IAM towards the next challenges

**Authors:** ALESSANDRA CASALE<sup>1</sup>; Davide Marcato<sup>2</sup>; Enrico Vianello<sup>3</sup>; Federica Agostini<sup>4</sup>; Francesco Giacomini<sup>5</sup>; Jacopo Gasparetto<sup>4</sup>; Manoj Garai<sup>None</sup>; Roberta Miccoli<sup>5</sup>; Saiteja Reddy Vennapusa<sup>None</sup>; Stefano Enrico Zotti<sup>None</sup>

<sup>1</sup> LNGS-INFN

<sup>2</sup> Istituto Nazionale di Fisica Nucleare

<sup>3</sup> INFN-CNAF

<sup>4</sup> CNAF

<sup>5</sup> INFN CNAF

**Corresponding Authors:** jacopo.gasparetto@cnaf.infn.it, alessandra.casale@lngs.infn.it, federica.agostini@cnaf.infn.it, stefano.zotti@cnaf.infn.it, davide.marcato@lnl.infn.it, manoj.garai@stfc.ac.uk, francesco.giacomini@cern.ch, roberta.miccoli@cnaf.infn.it, enrico.vianello@cern.ch, saiteja-reddy.vennapusa@stfc.ac.uk

INDIGO IAM (Identity and Access Management) is a comprehensive service that enables organizations to manage and control access to their resources and systems effectively. It implements a standard OAuth2 Authorization Service and OpenID Connect Provider and it has been chosen as the AAI solution by the WLCG community for the transition from VOMS proxy-based authorization to JSON web tokens.

This contribution describes the recent updates introduced by the latest IAM releases and the current roadmap for its evolution. In the near future, a primary focus is on avoiding to store access tokens in the database, to enhance the performance of both token issuance and token deletion. Another important milestone is the integration of a Multi-Factor Authentication mechanism. Additionally, substantial effort will be dedicated to migrating from outdated frameworks, such as MITREid Connect and AngularJS, to more stable and robust solutions based on Spring Security and React respectively. As a consequence, a new dashboard is also being developed, aligned with the latest advances in User Interface design.

This abstract highlights the progress made in the development roadmap described above, not forgetting the general auditing and performance improvements introduced with the latest releases or planned, such as the use of Open Policy Agent to re-implement the internal mechanism of the Scope Policy API.

Parallel (Track 2) / 453

## The new hardware trigger processor at NA62 experiment: Status of the System and First Results

**Authors:** Alessandro Lonardo<sup>1</sup>; Andrea Biagioni<sup>1</sup>; Carlotta Chiarini<sup>2</sup>; Cristian Rossi<sup>2</sup>; Francesca Lo Cicero<sup>1</sup>; Francesco Simula<sup>1</sup>; Luca Pontisso<sup>1</sup>; Matteo Turisini<sup>3</sup>; Mauro Raggi<sup>4</sup>; Michele Martinelli<sup>1</sup>; Ottorino Frezza<sup>1</sup>; Paolo Cretaro<sup>None</sup>; Piero Vicini<sup>5</sup>; Pierpaolo Perticaroli<sup>1</sup>; Roberto Ammendola<sup>6</sup>; Roberto Piandani<sup>7</sup>

<sup>1</sup> INFN, Roma I (IT)

<sup>2</sup> Sapienza Universita e INFN, Roma I (IT)

<sup>3</sup> Cineca

<sup>4</sup> CERN and Sapienza

<sup>5</sup> INFN Rome Section

<sup>6</sup> INFN Roma Tor Vergata (IT)

<sup>7</sup> Univ. Autonoma de San Luis Potosi

**Corresponding Authors:** francesca.locicero@roma1.infn.it, francesco.simula@roma1.infn.it, luca.pontisso@roma1.infn.it, cristian.rossi@roma1.infn.it, paolo.cretaro@roma1.infn.it, mauro.raggi@cern.ch, piero.vicini@roma1.infn.it, roberto.piandani@cern.ch, ottorino.frezza@roma1.infn.it, andrea.biagioni@roma1.infn.it, alessandro.lonardo@roma1.infn.it, pierpaolo.perticaroli@roma1.infn.it, roberto.ammendola@roma2.infn.it, carlotta.chiarini@roma1.infn.it, m.turisini@cineca.it, michele.martinelli@roma1.infn.it

The NA62 experiment is designed to study kaon's rare decays using a decay-in-flight technique. Its Trigger and Data Acquisition (TDAQ) system is multi-level, making it critically dependent on the performance of the inter-level network.

To manage the enormous amount of data produced by the detectors, three levels of triggers are used. The first level L0TP, implemented using an FPGA device, has been in operation since the start of data taking in 2016.

To increase the efficiency of the system and implement additional algorithms, an upgraded system (L0TP+) was developed starting in 2018. This upgrade utilizes a high-end FPGA available on the market, offering more computing power, larger local memory, and higher transmission bandwidth.

We have planned tests for a new trigger algorithm that implements quadrant-based logic for the veto systems. This new approach is expected to improve the main trigger efficiency by several percent.

Extensive tests were conducted using a parasitic setup that included a set of Network TAPs and a commodity server, allowing for proficient comparison of trigger decisions on an event-by-event basis. The experience gained from this parasitic mode operation can be leveraged for the next data-taking period as a development setup to implement additional features, thereby accelerating the TDAQ upgrade.

After the testing period, the new system has been adopted as the online processor since 2023. Preliminary results on the efficiency of the new system will be reported. Integration with the new AI-based FPGA-RICH system, which performs online partial particle identification, will also be discussed.

Parallel (Track 1) / 454

## Distributed management and processing of ALICE monitoring data with Onedata

**Author:** Michał Orzechowski<sup>1</sup>

**Co-authors:** Bartosz Baliś<sup>1</sup>; Łukasz Dutka<sup>2</sup>; Jacek Kitowski<sup>1</sup>



<sup>1</sup> *AGH University of Krakow, Faculty of Computer Science, Poland*

<sup>2</sup> *Academic Computer Centre Cyfronet AGH, Krakow, Poland*

**Corresponding Authors:** lukasz.dutka@cyfronet.pl, kito@agh.edu.pl, morzech@agh.edu.pl, balis@agh.edu.pl

Onedata 1 platform is a high-performance data management system with a distributed, global infrastructure that enables users to access heterogeneous storage resources worldwide. It supports various use cases ranging from personal data management to data-intensive scientific computations. Onedata has a fully distributed architecture that facilitates the creation of a hybrid cloud infrastructure with private and commercial cloud resources. Users can collaborate, share, and publish data, as well as perform high-performance computations on distributed data using different interfaces.

Within the ALICE project, we are designing an architecture that live-streams monitoring data from MonALISA data and stores it in Onedata, utilising S3 storage. This data is accessible through the POSIX filesystem HPC, cloud infrastructures, and external MLOps systems (via S3 or REST API). When a computational task requires the data, it is seamlessly transferred and cached at the task's location. Onedata's distributed and multi-protocol nature facilitates the creation of a hybrid data processing infrastructure, where Onedata functions as the data plane. The platform also includes robust security features to safeguard data and metadata from unauthorised changes, ensuring the integrity of the datasets during the final preparation stages. Additionally, Onedata enables long-term archiving of datasets, preserving crucial information for future reference. Data can be structured hierarchically within Onedata, and datasets are annotated with metadata, simplifying the organisation and retrieval of specific information.

Currently, Onedata is used in European projects: EUreka3D 3, EuroScienceGateway [4], DOME [5], InterTwin [6] where it provides a data transparency layer for managing large, distributed datasets on dynamic hybrid cloud containerised environments.

**Acknowledgements:** This work is co-financed in part supported by the Ministry of Science and Higher Education (Agreement Nr 2023/WK/07) and by the program of the Ministry of Science and Higher Education entitled "PMW".

**References:**

1 Onedata project website. <https://onedata.org>.

2 ALICE - A Large Ion Collider Experiment. <https://alice-collaboration.web.cern.ch>.

3 EUreka3D: European Union's REKconstructed in 3D. <https://eureka3d.eu>.

[4] EuroScienceGateway project: open infrastructure for data-driven research. <https://galaxyproject.org/projects/esg/>.

[5] DOME: A Distributed Open Marketplace for Europe Cloud and Edge Services. <https://dome-marketplace.eu>.

[6] InterTwin: Interdisciplinary Digital Twin Engine for Science. <https://intertwin.eu>.

**Parallel (Track 7) / 455**

## **Revolutionising Radio Astronomy: The UK's Role in SKA's SRC-Net Deployment**

**Authors:** Ian Collier<sup>1</sup>; James William Walder<sup>1</sup>; John Garbutt<sup>2</sup>

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

<sup>2</sup> *StackHPC*

**Corresponding Authors:** john.garbutt@stackhpc.com, james.william.walder@cern.ch, ian.peter.collier@cern.ch

The Square Kilometre Array (SKA) is set to revolutionise radio astronomy and will utilise a distributed network of compute and storage resources, known as SRCNet, to store, process and analyse the data at the exoscale. The United Kingdom plays a pivotal role in this initiative, contributing a significant portion of the SRCNet infrastructure. SRCNet v0.1, scheduled for early 2025, will prototype the movement and management of data, demonstrating both ingestion and dissemination processes (using Rucio and related tools from HEP). It will also demonstrate access to data through science platforms and interactive analysis tools, building on synergies with discussions in the HEP

community with Analysis Facilities.

Azimuth, a self-service portal optimised for high-performance computing in scientific applications, will be integral to this effort. It simplifies the complex management of cloud resources, making it an ideal tool for the heterogeneous compute environments provided by UK HPC sites. Federating access to UK and SKA members across these resources, and ensuring efficient use of network and storage placement will be critical for science exploitation by enhancing the experience and productivity of researchers. This work outlines the UK's architectural vision and roadmap, detailing implementation strategies for SRCNet v0.1 and beyond.

**Poster session / 456**

## Designing the Identity Infrastructure for the Square Kilometre Array Regional Centre Network

**Authors:** Ian Collier<sup>1</sup>; Jens Jensen<sup>None</sup>; Tom Dack<sup>None</sup>

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

**Corresponding Authors:** tom.dack@cern.ch, ian.peter.collier@cern.ch, jens.jensen@stfc.ac.uk

The Square Kilometre Array Observatory (SKAO) is a next-generation radio astronomy-driven Big Data facility that will revolutionise our understanding of the Universe and the laws of fundamental physics. SKAO headquarters is located at Jodrell Bank in the UK, with the main telescope sites located in South Africa and Australia, supported by a global network of SKA Regional Centres (SRCs) located in its member states. Collectively these SRCs will form a federated science ecosystem, the SRCNet, which will support global access to SKA data –with data rates projected to be around 700PB per year.

To deliver on these goals, SRCNet requires a scalable and robust federated authorisation, authentication, and identity management infrastructure. Design and prototyping has been underway for two years, led by UKSRC members at UKRI-STFC, together with international collaborators from other SRCs and the wider community. This work is informed by WLCG's transition to token based authorization and the Authentication and Authorisation for Research and Collaboration (AARC) Projects.

The AARC Blueprint Architecture (BPA) has been central to the prototyping process, however it has been found that the architecture needed for SRCNet does not map directly to the existing BPA. The SRCNet team is in turn supporting the evolution AARC BPA as part of the recently commenced AARC-TREE project.

This presentation will provide an overview of the prototyping work done so far, as well as share principles and assertions for the design of the final infrastructure and how these both relate and steer changes to the AARC BPA.

**Parallel (Track 7) / 457**

## A Successful Data Centre Refurbishment Project

**Author:** Daniel Peter Traynor<sup>None</sup>

**Co-authors:** Daohai Li ; Dimitrios Toroz ; Jonathan Hays <sup>1</sup>; Richard Alexander Owen

<sup>1</sup> *University of London (GB)*

**Corresponding Authors:** j.hays@qmul.ac.uk, d.toroz@qmul.ac.uk, d.traynor@qmul.ac.uk, r.a.owen@qmul.ac.uk, daohai.li@qmul.ac.uk

Queen Mary University of London (QMUL) has recently finished refurbishing its data centre that house our computing cluster supporting the WLCG project. After 20 years of operation the original data centre had significant cooling issues and increases in energy prices have all driven the need for refurbishment amid growing awareness of climate change. In addition there is a need to increase the capacity (from 150KW) to cope with the expected increased needs to the high luminosity LHC and new astronomy projects such as the LSST and SKA observatories.

A summary of the project is presented covering the project time line and solutions implemented (in row cooling, hot aisle containment, heat pumps and dry air coolers). Experiences and lessons learnt in the design, building and use of the data centre (covering choices in power supply, rack density, storage space, floor type, lighting, monitoring, etc...) are discussed. Effects of budget constraints and project rescoping due to inflation are also discussed.

First data from the energy use and heat recovery are presented and estimates of the energy and carbon saving over time are given.

**Parallel (Track 4) / 458**

## **CVMFS: Pushing performance on highly parallel, many-core clients**

**Authors:** Jakob Blomer<sup>1</sup>; Laura Promberger<sup>1</sup>; Matt Harvey<sup>2</sup>; Reza Naghibi<sup>2</sup>; Valentin Volkl<sup>1</sup>

<sup>1</sup> *CERN*

<sup>2</sup> *Jump Trading*

**Corresponding Authors:** laura.promberger@cern.ch, valentin.volkl@cern.ch, jakob.blomer@cern.ch, rnaghibi@jumptrading.com, mharvey@jumptrading.com

The CernVM File System (CVMFS) is an efficient distributed, read-only file system that streams software and data on demand. Its main focus is to distribute experiment software and conditions data to the world-wide LHC computing infrastructure. In WLCG, more than 5 billion files are distributed via CVMFS and its read-only file system client is installed on more than 100,000 worker nodes. Recent hardware trends have increased the usage of CVMFS in highly parallel environments. Nodes with more than 64 physical cores running concurrent workloads are common. These highly parallel, many-core workloads have exposed specific bugs and limitations of CVMFS in such environments. This contribution reports on the developments that address these issues, and presents new performance benchmarks on machines with 256 (virtual) cores.

**Parallel (Track 2) / 459**

## **Machine learning based event reconstruction for the MUonE experiment**

**Author:** Milosz Zdybal<sup>1</sup>

**Co-authors:** Anna Driutti<sup>2</sup>; Damian Mizera<sup>3</sup>; Emma Hess<sup>2</sup>; Marcin Kucharczyk<sup>1</sup>; Marcin Wolter<sup>1</sup>; Mateusz Jacek Goncerz<sup>1</sup>; Patrick Asenov<sup>2</sup>

<sup>1</sup> *Polish Academy of Sciences (PL)*

<sup>2</sup> *Universita & INFN Pisa (IT)*

<sup>3</sup> *Cracow University of Technology (PL)*

**Corresponding Authors:** mateusz.goncerz@cern.ch, damian.mizera@outlook.com, ehess@cern.ch, marcin.wolter@cern.ch, marcin.kucharczyk@cern.ch, milosz.zdybal@cern.ch, anna.driutti@unipi.it, patrick.asenov@cern.ch

The evergrowing amounts of data produced by the high energy physics experiments create a need for fast and efficient track reconstruction algorithms. When storing all incoming information is not feasible, online algorithms need to provide reconstruction quality similar to their offline counterparts. To achieve it, novel techniques need to be introduced, utilizing acceleration offered by the highly parallel hardware platforms, like GPUs. Artificial Neural Networks are a natural candidate here, thanks to their good pattern recognition abilities, non-iterative execution, and easy implementation on hardware accelerators.

The MUonE experimenting, searching for the signs of New Physics in the sector of anomalous magnetic moment of a muon, is investigating the use of the machine learning techniques in data processing. Works related to the ML-based track reconstruction will be presented. The first attempt used deep multilayer perceptron network to predict parameters of the tracks in the detector. Neural network was used as the base of the algorithm that proved to be as accurate as the classical approach but substituting the tedious step of iterative CPU-based pattern recognition. Further works included implementation of the Graph Neural Network for classification of track segment candidates.

Poster session / 460

## Implementing XRootD/SciToken-Based Access to Lustre Storage at GSI: A First Step Toward Data Federation for FAIR

**Authors:** Anar Manafov<sup>1</sup>; Mohammad Al-Turany<sup>2</sup>; Rouven Spreckels<sup>None</sup>; Soren Lars Gerald Fleischer<sup>1</sup>; Thorsten Kollegger<sup>1</sup>

<sup>1</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

<sup>2</sup> *CERN*

**Corresponding Authors:** t.kollegger@gsi.de, s.fleischer@gsi.de, r.spreckels@gsi.de, a.manafov@gsi.de, mohammad.al-turany@cern.ch

The implementation of a federated access system for GSI's local Lustre storage using XRootD and HTTP(s) protocols will be presented. It aims at ensuring a secure and efficient data access for the diverse scientific communities at GSI. This prototype system is a key step towards integrating GSI/FAIR into a federated data analysis model. We use Keycloak for authentication, which issues SciTokens through OpenID Connect, while LDAP manages local users. After successful login, a JSON Web Token (JWT) is created with appropriate read and write permissions. This token is passed to XRootD's multiuser plugin, which performs the requested operations as the specified user. We also developed an easy-to-use web interface to improve the user experience. This federated access model enhances the security, scalability, and usability of GSI's storage systems, making it a strong solution for modern data management needs.

Poster session / 461

## LbMCSubmit: Streamlined production and submission of LHCb MC requests

**Author:** Emir Muhammad<sup>1</sup>

**Co-authors:** Adam Morris<sup>2</sup>; Chris Burr<sup>2</sup>; Michal Kreps<sup>1</sup>; Qiuchan Lu<sup>3</sup>

<sup>1</sup> *University of Warwick (GB)*

<sup>2</sup> *CERN*

<sup>3</sup> *Centre National de la Recherche Scientifique (FR)*

**Corresponding Authors:** adam.morris@cern.ch, michal.kreps@cern.ch, christopher.burr@cern.ch, emir.muhammad@cern.ch, qiuchan.lu@cern.ch

The LHCb experiment requires a wide variety of Monte Carlo simulated samples to support its physics programme. LHCb's centralised production system operates on the DIRAC backend of the WLCG; users interact with it via the DIRAC web application to request and produce samples.

To simplify this procedure, LbMCSsubmit was introduced, automating the generation of request configurations from a basic parameterisation of the desired samples, eliminating the need for defining hundreds of static models as was done in the past. However, this submission process is still clunky, time-consuming, and lacks a conducive platform for discussion.

We present a more streamlined approach to testing and submitting requests via a GitLab repository. This involves creating a simulation request by opening a GitLab merge request, which is then processed via LbMCSsubmit to generate the necessary configuration, and then automatically tested using Continuous Integration on a few events. Upon approval, productions are automatically submitted and run on the WLCG via DIRAC, with samples accessible in the bookkeeping.

This approach is faster and simpler for users, ensuring efficient resource utilisation through testing. Using GitLab has the added benefits of collating discussion and reviewing requests, and leveraging GitLab's powerful approval feature with the custom API to enforce approvals by certain experts in requests where necessary.

Introduced in spring 2023 alongside LbMCSsubmit and continually updated since, this system has become the standard procedure for simulation requests within LHCb, processing nearly five hundred requests to date.

**Poster session / 462**

## Versioning of the LHCb DD4Hep-based Detector Description

**Author:** Menglin Xu<sup>1</sup>

**Co-authors:** Ben Couturier<sup>2</sup>; Emir Muhammad<sup>1</sup>; Thomas Latham<sup>1</sup>

<sup>1</sup> *University of Warwick (GB)*

<sup>2</sup> *CERN*

**Corresponding Authors:** ben.couturier@cern.ch, emir.muhammad@cern.ch, t.latham@warwick.ac.uk, menglin.xu@cern.ch

The LHCb Detector project is home to the detector description of the LHCb experiment. It is used in all data processing applications, from simulation to reconstruction. It is based on the DD4hep package relying on the combination of XML files and C++ code. The need to support different versions of the detector layout in different data taking periods, on top of the DD4hep detector description format, necessitated a suite of recent custom developments.

The detector descriptions are identified by the C++ code and folders containing their description files (in the DD4hep CompactXML format). To support concurrent detectors layout, by convention, the components' description has to be self-contained and individually versioned within their respective component folders. There is no dependency between the components. The description of the overall LHCb detector combines specific versions of the components for the period being described; the versions for a specific data taking period or potential layouts to be simulated are in corresponding folders. When a sub-detector is identical in two LHCb layouts it can be loaded by both, avoiding code duplication and facilitating future upgrade simulation studies. A convention for naming the versions consistent between the components and the whole detector has been set up. Dependency and conventions checks as well as geometry monitoring are enforced by use of GitLab's CI testing. They include verifications that changes introduced do not affect existing versions.

**Parallel (Track 7) / 463**

## Towards an IPv6-only WLCG: more successes in reducing IPv4

**Author:** David Kelsey<sup>1</sup>

**Co-authors:** Andrea Sciabà<sup>2</sup>; Bruno Heinrich Hoefft<sup>3</sup>; Christopher Walker; Costin Grigoras<sup>2</sup>; Dimitrios Christidis<sup>2</sup>; Duncan Rand<sup>4</sup>; Edoardo Martelli<sup>2</sup>; Francesco Prelz<sup>5</sup>; Hironori Ito<sup>6</sup>; Jiri Chudoba<sup>7</sup>; Jose Flix Molina<sup>8</sup>; Kars Ohrenberg; Maria Del Carmen Misa Moreira<sup>2</sup>; Marian Babik<sup>2</sup>; Martin Bly<sup>9</sup>; Nick Buraglio<sup>10</sup>; Philip Demar; Raja Nandakumar<sup>1</sup>; Shawn Mc Kee<sup>11</sup>; Tim Chown; Tim Skirvin

<sup>1</sup> *Science and Technology Facilities Council STFC (GB)*

<sup>2</sup> *CERN*

<sup>3</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>4</sup> *Imperial College (GB)*

<sup>5</sup> *Università degli Studi e INFN Milano (IT)*

<sup>6</sup> *Brookhaven National Laboratory (US)*

<sup>7</sup> *Acad. of Sciences of the Czech Rep. (CZ)*

<sup>8</sup> *CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)*

<sup>9</sup> *STFC-RAL*

<sup>10</sup> *ESnet*

<sup>11</sup> *University of Michigan (US)*

**Corresponding Authors:** dimitrios.christidis@cern.ch, edoardo.martelli@cern.ch, francesco.prelz@mi.infn.it, andrea.sciaba@cern.ch, hito@bnl.gov, bruno.hoefft@kit.edu, tskirvin@fnal.gov, shawn.mckee@cern.ch, buraglio@es.net, chudobaj@cern.ch, tim.chown@jisc.ac.uk, costin.grigoras@cern.ch, duncan.rand@cern.ch, marian.babik@cern.ch, martin.bly@stfc.ac.uk, demar@fnal.gov, jose.flix.molina@cern.ch, christopher.walker@jisc.ac.uk, david.kelsey@stfc.ac.uk, carmen.misa.moreira@cern.ch, kars.ohrenberg@desy.de, raja.nandakumar@cern.ch

The Worldwide Large Hadron Collider Computing Grid (WLCG) community's deployment of dual-stack IPv6/IPv4 on its worldwide storage infrastructure is very successful and has been presented by us at earlier CHEP conferences. Dual-stack is not, however, a viable long-term solution; the HEPiX IPv6 Working Group has focused on studying where and why IPv4 is still being used, and how to flip such traffic to IPv6. The agreed end goal is to turn IPv4 off and run IPv6-only over the wide area network, to simplify both operations and security management.

This paper reports our work since the CHEP2023 conference. Firstly, we present our campaign to encourage the deployment of IPv6 on CPU services and Worker Nodes, with a deadline of end of June 2024. Then, the WLCG Data Challenge (DC24) performed during two weeks of February 2024 was an excellent opportunity to observe the percentage of data transfers which were carried by IPv6. We present our observation of the predominance of IPv6 in data transfers during DC24 and the opportunity we had to understand yet more reasons for the use of IPv4 and to get them removed.

The paper ends with the working group's proposed plans and timescale for moving WLCG to "IPv6-only". We have continued to test IPv6-only clusters as a way of confirming the readiness of the LHC experiments for an IPv6-only environment. Another aspect of the plan is the possible use of IPv6-only clients configured with a customer-side translator, or CLAT, together with a deployment of 464XLAT using what is often known as "IPv6-mostly" as in IETF RFCs 6877/8925. This will enable IPv6-only sites to continue to connect to non-WLCG IPv4-only services.

Parallel (Track 2) / 464

## A Comprehensive Bandwidth Testing Framework for the LHCb Upgrade Trigger System

**Authors:** Ella Godiva Noomen<sup>1</sup>; Luke Gazette<sup>2</sup>; Mika Anton Vesterinen<sup>2</sup>; Nicole Skidmore<sup>3</sup>; Ross John Hunter<sup>2</sup>; Sascha Stahl<sup>4</sup>; Shunan Zhang<sup>5</sup>

<sup>1</sup> *Nikhef National institute for subatomic physics (NL)*

<sup>2</sup> *University of Warwick (GB)*

<sup>3</sup> *University of Warwick*

<sup>4</sup> *CERN*

<sup>5</sup> *University of Oxford (GB)*

**Corresponding Authors:** luke.grazette@cern.ch, shunan.zhang@cern.ch, ross.john.hunter@cern.ch, nicola.skidmore@cern.ch, e.noomen@alumni.maastrichtuniversity.nl, sascha.stahl@cern.ch, mika.vesterinen@cern.ch

The LHCb experiment at CERN has undergone a comprehensive upgrade. In particular, its trigger system has been completely redesigned into a hybrid-architecture, software-only system that delivers ten times more interesting signals per unit time than its predecessor. This increased efficiency - as well as the growing diversity of signals physicists want to analyse - makes conforming to crucial operational targets on bandwidth and storage capacity ever more challenging. To address this, a comprehensive, automated testing framework has been developed that emulates the entire LHCb trigger and offline-processing software stack on simulated and real collision data. Scheduled both nightly and on-demand by software testers during development, these tests measure the online- and offline-processing's key operational performance metrics (such as rate and bandwidth), for each of the system's 3500 distinct physics selection algorithms, and their cumulative totals. The results are automatically delivered via concise summaries - to GitLab merge requests and instant messaging channels - that further link to an extensive dashboard of per-algorithm information. The dashboard and pages therein (categorised by physics working group) facilitate exploratory data analysis and test-driven trigger development by 100s of physicists, whilst the concise summaries enable efficient, data-driven decision-making by management and software maintainers. Altogether, this novel and performant bandwidth testing framework has been helping LHCb build an operationally-viable trigger and data-processing system whilst maintaining the efficiency to satisfy its physics goals.

Poster session / 465

## Prospects for novel track reconstruction algorithms based on Graph Neural Network models using telescope detector testbed

**Author:** Wojciech Gomulka<sup>1</sup>

**Co-authors:** Piotr Kowalski<sup>2</sup>; Tomasz Bold<sup>1</sup>; Tomasz Szumlak<sup>1</sup>

<sup>1</sup> *AGH University of Krakow (PL)*

<sup>2</sup> *AGH University*

**Corresponding Authors:** tomasz.bold@cern.ch, wgomulka@agh.edu.pl, pkowal@agh.edu.pl, szumlak@agh.edu.pl

The upcoming upgrades of LHC experiments and next-generation FCC (Future Circular Collider) machines will again change the definition of big data for the HEP environment. The ability to effectively analyse and interpret complex, interconnected data structures will be vital. This presentation will delve into the innovative realm of Graph Neural Networks (GNNs). This powerful tool extends traditional deep learning techniques to handle graph-structured data and may provide new and fast algorithms for track reconstruction in both the 3D and 4D domains.

Projecting challenging task of track reconstruction, especially challenging in harsh hadronic environment, into non-Euclidean domain of GNNs may leverage the intrinsic structure of graph data to extract addition crucial features and patterns that are either difficult or impossible for traditional statistical or intelligent reconstruction algorithms.

We present our initial studies using various GNN models implemented within the ACTS (A Common Tracking Software Project) framework. In our studies, we created a telescope detector that resembles an LHCb silicon vertex locator and used toy-generated data with truth information. Using such simulated setup, we were able to successfully train several GDN models to perform track reconstruction tasks. Based on these initial results, we performed preliminary studies to obtain efficiencies and resolutions for selected kinematical variables.

Our preliminary studies are very promising and show significant potential for using GDNs models as track reconstruction engines for future LHC upgrades and beyond.

Parallel (Track 4) / 466

## SPECTRUM: towards a Computing Strategy for Data-intensive Science Infrastructures in Europe

**Authors:** Chiara Ferrari<sup>1</sup>; Corentin Lefevre<sup>2</sup>; Eric Wulff<sup>3</sup>; Hans-Christian Hoppe<sup>4</sup>; Jean-Pierre Vilotte<sup>5</sup>; Luis Cifuentes<sup>4</sup>; Maria Girone<sup>3</sup>; Patricia Ruiz<sup>6</sup>; Sergio Andreozzi<sup>6</sup>; Tommaso Boccali<sup>7</sup>; Xavier Salazar<sup>6</sup>

<sup>1</sup> *Observatoire de la Cote D'Azur (OCA)*

<sup>2</sup> *Neovia Innovation*

<sup>3</sup> *CERN*

<sup>4</sup> *Forschungszentrum Jülich GmbH*

<sup>5</sup> *CNRS-INSU*

<sup>6</sup> *EGI Foundation*

<sup>7</sup> *INFN Sezione di Pisa*

**Corresponding Authors:** eric.wulff@cern.ch, patricia.ruiz@egi.eu, vilotte@ipgp.fr, sergio.andreozzi@egi.eu, chiara.ferrari@oca.eu, l.cifuentes@fz-juelich.de, xavier.salazar@egi.eu, tommaso.boccali@cern.ch, corentin.lefevre@neovia-innovation.eu, h.hoppe@fz-juelich.de, maria.girone@cern.ch

The amount of data gathered, shared and processed in frontier research is set to increase steeply in the coming decade, leading to unprecedented data processing, simulation and analysis needs.

In particular, the research communities in High Energy Physics and Radio Astronomy are preparing to launch new instruments that require data and compute infrastructures several orders of magnitude larger than what is currently available and entering in the Exascale era.

To meet these requirements, new data-intensive architectures, heterogeneous resource federation models, and IT frameworks will be needed, including large-scale compute and storage capacity to be procured and made accessible at the pan-European level.

Additionally, the emergence of high-end Exascale HPC and Quantum computing systems provides new opportunities for accelerating discoveries and complementing the capabilities of existing research HTC and Cloud facilities.

Addressing key questions around scalability, performance, energy efficiency, portability, interoperability and cybersecurity is crucial to ensuring the successful integration of these heterogeneous systems.

In this context, the SPECTRUM project (<https://spectrumproject.eu/>) aims to deliver a Strategic Research, Innovation and Deployment Agenda (SRIDA) and a Technical Blueprint for a European compute and data continuum.

With a consortium composed of leading European science organisations in High Energy Physics and Radio Astronomy, and leading e-Infrastructure providers covering HTC, HPC, Cloud and Quantum technologies, the project will work with a Community of Practice composed of external experts.

The ultimate goal is to pave the way towards data-intensive scientific collaborations with access to a federated European Exabyte-scale research data federation and compute continuum.

The contribution is going to show how the project is going to operate, the results already obtained and the roadmap to the end of the project.

Parallel (Track6) / 467

## The ePIC Simulation Campaign Workflow on the Open Science Grid



**Author:** Sakib Rahman<sup>1</sup>

<sup>1</sup> *Brookhaven National Laboratory*

**Corresponding Author:** srahman1@bnl.gov

The ePIC collaboration is realizing the first experiment of the future Electron-Ion Collider (EIC) at the Brookhaven National Laboratory that will allow for a precision study of the nucleon and the nucleus at the scale of sea quarks and gluons through the study of electron-proton/ion collisions. This talk will discuss the current workflow in place for running centralized simulation campaigns for ePIC on the Open Science Grid infrastructure. This involves monthly releases of ePIC software and container deployments to CVMFS, generation of input datasets in HepMC format according to collaboration-defined policy, using Snakemake in CI/CD for validation and benchmarking, and submitting jobs to the Open Science Grid condor scheduler for opportunistic running on available resources. File transfers utilize XrootD, and RUCIO is used for data management. The workflow is being continuously refined to improve daily throughput (currently ~50-100k core hours per day) and minimize job failures. Since May 2023, monthly simulation campaigns employing the workflow have cumulatively used over ~10 million core hours on the Open Science Grid and produced over ~280 TB of simulation data. The campaigns incorporate simulations for the broad science program of the EIC and are actively used for the detector and physics studies in preparation of the Technical Design Review.

**Parallel (Track 7) / 468**

## Latest developments of the PUNCH4NFDI compute and storage infrastructures

**Authors:** Arman Khalatyan<sup>1</sup>; Benoit Roland<sup>2</sup>; Christoph Wissing<sup>3</sup>; Harry Enke<sup>1</sup>; Manuel Giffels<sup>4</sup>; Matthias Hoeff<sup>5</sup>; Michael Huebner<sup>6</sup>; Oliver Freyermuth<sup>6</sup>

<sup>1</sup> *AIP - Leibniz-Institut für Astrophysik Potsdam (DE)*

<sup>2</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>3</sup> *Deutsches Elektronen-Synchrotron (DE)*

<sup>4</sup> *KIT - Karlsruhe Institute of Technology (DE)*

<sup>5</sup> *TLS - Thüringer Landessternwarte (DE)*

<sup>6</sup> *University of Bonn (DE)*

**Corresponding Authors:** benoit.roland@kit.edu, michael.huebner@uni-bonn.de, christoph.wissing@desy.de, henke@aip.de, akhalatyan@aip.de, oliver.freyermuth@cern.ch, hoeft@tls-tautenburg.de, manuel.giffels@cern.ch

The PUNCH4NFDI consortium, funded by the German Research Foundation for an initial period of five years, gathers various physics communities - particle, astro-, astroparticle, hadron and nuclear physics - from different institutions embedded in the National Research Data Infrastructure initiative. The overall goal of PUNCH4NFDI is the establishment and support of FAIR data management solutions for all users of the participating communities.

The federated compute and storage infrastructures made available to the PUNCH4NFDI consortium, Compute4PUNCH and Storage4PUNCH, will be presented. These infrastructures, comprising a variety of heterogeneous compute and storage systems provided by the participating institutions, are managed by an HTCondor overlay batch system and COBalD/TARDIS metaschedulers. The TARDIS manager dynamically integrates the various compute resources into one overlay batch system based on HTCondor, while the COBalD workload balancer optimizes the distribution of the tasks to be performed. The standardized access to the federated compute and storage resources is managed by a token-based authentication and authorization infrastructure. The refreshment of short-lived access tokens is automated in a transparent monitoring and renewal mechanism making use of the HTCondor credential manager in combination with the MyToken service. Login nodes are defining single entry points to the federation, while a virtualized and scalable software environments provisioning is ensured by the use of containers and the CERN Virtual Machine File System.

The latest developments of Compute4PUNCH and Storage4PUNCH will be presented, including the newly developed automated token management using HTCondor and Mytoken. In addition, the integration of Compute4PUNCH as a compute backend into the REANA reproducible analysis platform developed at CERN, with an instance hosted and managed by the PUNCH4NFDI consortium, will be shown.

Parallel (Track 2) / 469

## Performance of the LHCb heterogeneous software trigger

**Author:** Alessandro Scarabotto<sup>1</sup>

<sup>1</sup> *Technische Universitaet Dortmund (DE)*

**Corresponding Author:** alessandro.scarabotto@cern.ch

Since 2022, the LHCb detector is taking data with a full software trigger at the LHC proton-proton collision rate, implemented in GPUs in the first stage and CPUs in the second stage. This setup allows to perform the alignment & calibration online and to perform physics analyses directly on the output of the online reconstruction, following the real-time analysis paradigm. This talk will give a detailed overview of the LHCb trigger implementation and its underlying computing infrastructure, discuss challenges of using a heterogeneous architecture and report its performance in nominal data taking conditions during 2024 after two commissioning years.

Parallel (Track 3) / 470

## Track reconstruction in the STCF detector

**Authors:** Hang Zhou<sup>None</sup>; Hao Li<sup>1</sup>; Xiaocong Ai<sup>1</sup>; Xiaoqian Jia<sup>None</sup>; Xiaoshuai Qin<sup>2</sup>

<sup>1</sup> *Zhengzhou University*

<sup>2</sup> *Shandong University, Qingdao*

**Corresponding Authors:** jiaxq@mail.sdu.edu.cn, qinxs@sdu.edu.cn, lihao097430@gs.zzu.edu.cn, zh25@mail.ustc.edu.cn, xiaocong.ai@cern.ch

The Super Tau-Charm Facility (STCF) proposed in China is an electron-positron collider designed to operate in a center-of-mass energy range from 2 to 7 GeV with peak luminosity above  $0.5 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$ . The STCF will provide a unique platform for studies of hadron physics, strong interactions and searches for new physics beyond the Standard Model in the tau-charm region. To fulfill the physics goals of the STCF experiment, good performance of reconstruction of charged tracks with momentum down to 50 MeV is required. The tracking system of the STCF detector consists of an inner tracker (ITK) with several independent layers and a large cylindrical drift chamber (main drift chamber, MDC). A track reconstruction software framework has been developed for the baseline STCF detector design. A global track finding algorithm based on Hough Transform, which handles the hits from the ITK and MDC together, hence less sensitive to missing local hits, has been implemented and optimized in this framework. The tracking toolkit of A Common Tracking Software (ACTS) is also explored for tracking in the STCF detector, where the Combinatorial Kalman Filter (CKF) implemented in ACTS is used to find the tracks based on the track seeds composed of ITK hits. In addition, innovative machine learning techniques have been explored to use in track reconstruction in the STCF tracking detectors. Different tracking methods and techniques have been compared to achieve good overall track reconstruction performance for the STCF experiment. In this contribution, we present the tracking system of the STCF detector and development and implementation of afore-mentioned various tracking algorithms in the STCF offline software and the simulated tracking performance.

**Poster session / 471**

## **KKMCee: Multiphoton MC for lepton and quark pair production at lepton colliders**

**Author:** Andrzej Konrad Siodmok<sup>1</sup><sup>1</sup> *Jagiellonian University (PL)***Corresponding Author:** a.siodmok@cern.ch

We present an overview of the Monte Carlo event generator for lepton and quark pair production for the high-energy electron-positron annihilation process. We note that it is still the most sophisticated event generator for such processes. Its entire source code is rewritten in the modern C++ language. We checked that it reproduces all features of the older code in Fortran 77. We discuss a number of improvements both in the MC algorithm and in its various interfaces, such as those to parton showers and detector simulation.

**Parallel (Track 5) / 472**

## **Recent developments in the Gaussino simulation software**

**Author:** Adam Morris<sup>1</sup>**Co-authors:** Gloria Corti<sup>1</sup>; Marco Clemencic<sup>1</sup>; Michał Mazurek<sup>2</sup>; Ruben Pozzi<sup>1</sup>; Witold Pokorski<sup>1</sup><sup>1</sup> *CERN*<sup>2</sup> *National Centre for Nuclear Research (PL)***Corresponding Authors:** marco.clemencic@cern.ch, ruben.pozzi@cern.ch, witold.pokorski@cern.ch, gloria.corti@cern.ch, adam.morris@cern.ch, michal.mazurek@cern.ch

Gaussino is an experiment-independent simulation package built upon the Gaudi software framework. It provides generic core components and interfaces for a complete HEP simulation application: event generation, detector simulation, geometry, monitoring and output of the simulated data. This makes it suitable for use as a standalone application for early prototyping, testbeam setups etc as well as a toolkit with which to build a dedicated simulation application for a full-scale experiment.

In this talk we give an overview of recent developments in Gaussino.

The generator phase is adapted to separate the handling of the hard scattering from the parton shower and hadronisation. This allows easier integration of Matrix Elements generators, like MadGraph or Powheg with the general purpose generators like Pythia. A new interface for specifying generic generator-level cuts at configuration time is also introduced.

A new dedicated component for machine-learning based fast simulations that facilitates the integration of fast simulation hooks in Geant4 with machine learning frameworks, all based on Gaudi's scheduling and data processing tools.

Support for selectively offloading electromagnetic particle transport to GPUs is investigated. The toolbox for customisation, verification, visualisation and debugging of the geometry is extended.

The user configuration is made more robust and flexible.

We will provide examples in the use of the new features in the context of the LHCb Simulation.

**Plenary session / 473**

## Julia in HEP

**Authors:** Alexander Moreno Briceño<sup>1</sup>; Allen Christopher Caldwell<sup>2</sup>; Benedikt Hegner<sup>3</sup>; Graeme A Stewart<sup>3</sup>; Jerry Ling<sup>4</sup>; Mikhail Mikhasenko<sup>5</sup>; Oliver Schulz<sup>2</sup>; Pere Mato Vila<sup>3</sup>; Philippe Gras<sup>6</sup>; Samuel Cadellin Skipsey<sup>None</sup>; Tamas Gal<sup>7</sup>; Uwe Hernandez Acosta<sup>8</sup>

<sup>1</sup> *Universidad Antonio Nariño*

<sup>2</sup> *Max-Planck-Institut für Physik*

<sup>3</sup> *CERN*

<sup>4</sup> *Harvard University (US)*

<sup>5</sup> *Ruhr Univeristy Bochum*

<sup>6</sup> *Université Paris-Saclay (FR)*

<sup>7</sup> *ECAP, FAU Erlangen-Nürnberg*

<sup>8</sup> *Helmholtz-Zentrum Dresden-Rossendorf*

**Corresponding Authors:** u.hernandez@hzdr.de, philippe.gras@cern.ch, allen.christopher.caldwell@cern.ch, jerry.ling@cern.ch, samuel.cadellin.skipsey@cern.ch, graeme.andrew.stewart@cern.ch, tamas.gal@fau.de, benedikt.hegner@cern.ch, pere.mato@cern.ch, oliver.schulz@cern.ch, mikhail.mikhasenko@cern.ch, alexander.moreno@uan.edu.co

Julia is a mature general-purpose programming language, with a large ecosystem of libraries and more than 10000 third-party packages, which specifically targets scientific computing. As a language, Julia is as dynamic, interactive, and accessible as Python with NumPy, but achieves run-time performance on par with C/C++. In this paper, we describe the state of adoption of Julia in HEP, where momentum has been gathering over a number of years.

HEP-oriented Julia packages can, via UnROOT.jl, already read HEP's major file formats, including TTree and RNTuple formats. Interfaces to some of HEP's major software packages, such as through Geant4.jl, are available too. Jet reconstruction algorithms in Julia show excellent performance. A number of full HEP analyses have been performed in Julia.

We show how, as the support for HEP has matured, developments have benefited from Julia's core design choices, which makes reuse from and integration with other packages easy. In particular, libraries developed outside HEP for plotting, statistics, fitting, and scientific machine learning are extremely useful.

We believe that the powerful combination of flexibility and speed, the wide selection of scientific programming tools, and support for all modern programming paradigms and tools, make Julia the ideal choice for a future language in HEP.

Parallel (Track 5) / 474

## Generative AI for fast simulations in LHCb

**Author:** Michał Mazurek<sup>1</sup>

**Co-authors:** Gloria Corti<sup>2</sup>; Mateusz Kmiec<sup>1</sup>

<sup>1</sup> *National Centre for Nuclear Research (PL)*

<sup>2</sup> *CERN*

**Corresponding Authors:** gloria.corti@cern.ch, michal.mazurek@cern.ch, mateusz.rafal.kmiec@cern.ch

In high energy physics, fast simulation techniques based on machine learning could play a crucial role in generating sufficiently large simulated samples. Transitioning from a prototype to a fully deployed model usable in a full scale production is a very challenging task.

In this talk, we introduce the most recent advances in the implementation of fast simulation for calorimeter showers in the LHCb simulation framework based on Generative AI. We use a novel

component in Gaussino to streamline the incorporation of generic machine learning models. It leverages on the use of fast simulation hooks from Geant4 and machine learning backends such as PyTorch and ONNXRuntime.

Using this infrastructure the first implementation of selected ML models is trained and validated on the LHCb calorimeters. We will show a Variational Autoencoder (VAE) equipped with a custom sampling mechanism, as well as a transformer-based diffusion model (DiT). Both are compatible with the setup used in the CaloChallenge initiative, a collaborative effort aimed at training generic models for calorimeter shower simulation. We will share insights gained from the validation of these models on dedicated physics samples, including how to cope with handling and versioning multiple ML models in production in a distributed environment.

**Poster session / 475**

## A point library for the fast simulation of the LHCb Calorimeter

**Authors:** Matteo Rama<sup>1</sup>; Sergey Kholodenko<sup>2</sup>

**Co-authors:** Michal Kreps<sup>3</sup>; Michał Mazurek<sup>4</sup>

<sup>1</sup> INFN Pisa (IT)

<sup>2</sup> Università & INFN Pisa (IT)

<sup>3</sup> University of Warwick (GB)

<sup>4</sup> National Centre for Nuclear Research (PL)

**Corresponding Authors:** matteo.rama@cern.ch, michal.kreps@cern.ch, michal.mazurek@cern.ch, sergey.kholodenko@cern.ch

The simulation of physics events in the LHCb experiment uses the majority of the distributed computing resources available to the experiment. Notably, around 50% of the overall CPU time in the Geant4-based detailed simulation of physics events is spent in the calorimeter system. This talk presents a solution implemented in the LHCb simulation software framework to accelerate the calorimeter simulation.

During the Geant4 transport, the simulation of particles entering the calorimeter is stopped and the corresponding showers are generated using libraries of pre-simulated energy deposits. The energy deposits subsequently undergo a series of transformations to improve the accuracy of the simulation without increasing the library size. This technique reduces the computation time of the calorimeter to a negligible level.

The use of machine learning techniques in conjunction with the libraries to further enhance the simulation accuracy is also discussed, and a comparison between the outputs of the fast and detailed simulations is shown.

**Parallel (Track 3) / 476**

## A Geometry Agnostic Heterogenous Framework for Track Reconstruction for HEP Experiments

**Author:** Adriano Di Florio<sup>1</sup>

<sup>1</sup> CC-IN2P3

**Corresponding Author:** adriano.di.florio@cern.ch

The future development projects for the Large Hadron Collider towards HL-LHC will constantly bring nominal luminosity increase, with the ultimate goal of reaching, e.g., a peak luminosity of  $5 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$  for ATLAS and CMS experiments. This rise in luminosity will directly result in an

increased number of simultaneous proton collisions (pileup), up to 200, that will pose new challenges for track reconstruction in such a dense environment.

In response to these challenges, many experiments have started rewriting an increasing fraction of their track reconstruction software to run on heterogeneous architectures. While very successful in some cases, most of the time these efforts have stayed confined to single experiment projects.

In this work we will show the potentiality of a unique standalone software, running on multiple backends (CPUs, NVIDIA GPUs and AMD GPUs) aiming at the reconstruction of the tracker detector of multiple HEP experiments with a cylindrical geometry. We will discuss both the physics and computational performance for different detectors.

This represents the first step towards a unique standalone tool capable of carrying out the reconstruction of a model detector for HL-LHC leveraging on heterogeneous resources. A detector defined solely by its constituent elements: a silicon tracker, at least one calorimeter and a muon detector.

Poster session / 477

## MEASUREMENT OF PROPAGATION TIME WITH THE USE OF FPGA DEVICES IN ALICE EXPERIMENT AT CERN

**Authors:** Jacek Kitowski<sup>None</sup>; Jacek Tomasz Otwinowski<sup>1</sup>; Oleksandr Savchenko<sup>2</sup>; Pawel Grzegorz Russek<sup>2</sup>; Sebastian Koryciak<sup>3</sup>

<sup>1</sup> *Institute of Nuclear Physics Polish Academy of Sciences (PL)*

<sup>2</sup> *AGH University of Krakow (PL)*

<sup>3</sup> *AGH University of Science and Technology*

**Corresponding Authors:** russek@agh.edu.pl, koryciak@agh.edu.pl, jacek.otwinowski@ifj.edu.pl, kito@agh.edu.pl, osavchenko@student.agh.edu.pl

The poster presents the first experiments with the time-to-digital converter (TDC) for the Fast Interaction Trigger detector in ALICE experiment at CERN. It is implemented in Field-Programmable Gate Array (FPGA) technology and uses Serializer and Deserializers (ISERDES) with multiple-phase clocks.

The input pulse is a standard differential input signal. The signal is sampled with eight evenly spaced phase-shifted clock pulses generated by Mixed-Mode Clock Manager (MMCM). Before reaching ISERDES units the input signal is first buffered and then divided into two complementary outputs. The two ISERDES units are set up in oversample mode, enabling them to capture 2-phase DDR data. One ISERDES is synchronized with clocks at 0° and 90° angles, whereas the other ISERDES synchronizes with clocks at 45° and 135° angles. Additional four clocks are generated by locally inverting logic within each ISERDES unit. Since the FPGA contains a large number of ISERDES blocks, it will allow us to create multi-channel systems in a single FPGA chip.

Also, we consider another TDC implementation, where the phase shift between ISERDES is performed using the IDELAY blocks. This solution makes it possible to use only two clock signals 0 and 90°. IDELAYs have adjustable delays from 39 ps to 78 ps depending on the clocking frequency and has built-in temperature-dependent compensation.

Poster session / 478

## Up-scaling for measuring the spatial distribution of radiation dose for applications in the preparation of individual patient treatment plans

**Authors:** Bartłomiej Rachwał<sup>1</sup>; Jakub Hajduga<sup>2</sup>; Kamila Kalecińska<sup>None</sup>; Maciej Kalka<sup>2</sup>

**Co-authors:** Bartosz Mindur<sup>3</sup>; Bartłomiej Łach<sup>3</sup>; Dagmara Kulig<sup>2</sup>; Jakub Moron<sup>1</sup>; Maciej Kopeć; Piotr Kazimierz Wiacek<sup>1</sup>; Stefan Zenon Koperny<sup>4</sup>; Tomasz Andrzej Fiutowski<sup>1</sup>; Tomasz Szumlak<sup>1</sup>; damian Kabat<sup>5</sup>

<sup>1</sup> AGH University of Krakow (PL)

<sup>2</sup> AGH University of Krakow

<sup>3</sup> AGH University of Science and Technology

<sup>4</sup> AGH University of Science and Technology (PL)

<sup>5</sup> NIO-PIB, Krakow

**Corresponding Authors:** stefan.koperny@agh.edu.pl, kalka@agh.edu.pl, damian.kabat@krakow.nio.gov.pl, bartosz.mindur@cern.ch, kamkal@agh.edu.pl, tomasz.fiutowski@cern.ch, piotr.kazimierz.wiacek@cern.ch, bartlomiej.lach@fis.agh.edu.pl, bartlomiej.rachwal@cern.ch, jakub.moron@cern.ch, maciej.kopec@agh.edu.pl, jhajduga@agh.edu.pl, szumlak@agh.edu.pl, dkulig@agh.edu.pl

The super-resolution (SR) techniques are often used in the up-scaling process to add-in details that are not present in the original low-resolution image. In radiation therapy the SR can be applied to enhance the quality of medical images used in treatment planning. The Dose3D detector measuring spatial dose distribution<sup>1</sup>, the dedicated set of ML algorithms for SR has been proposed to perform final dose distribution up-scaling. Despite the significant advancements in image processing, the task of three-dimensional (3D) image upscaling remains a formidable challenge and has not gained widespread popularity due to the inherent complexities associated with preserving spatial consistency and accurately interpolating volumetric pixel intensities.

In our project, as the SR technique, the SRCNN 3 architecture has been adjusted. The training and validation data being produced with Geant4 MC simulation with in-house developed application and with two different scoring resolutions. Extra features related to the beam shape have been defined. The input data resolution is the one coming from the measurement (1cc) and the target data resolution is defined at the level of the CT image. Our research's latest breakthroughs and advancements will feature at the conference.

References:

1 <https://dose3d.fis.agh.edu.pl>,

2 M. Kopeć, et al. A reconfigurable detector for measuring the spatial distribution of radiation dose for applications in the preparation of individual patient treatment plans. Nuclear Inst. and Methods in Physics Research, A 1048 (2023) 167937

3 Dong, C., Loy, C.C., He, K., Tang, X. (2014). Learning a Deep Convolutional Network for Image Super-Resolution. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds) Computer Vision –ECCV 2014. ECCV 2014. Lecture Notes in Computer Science, vol 8692. Springer, Cham. [https://doi.org/10.1007/978-3-319-10593-2\\_13](https://doi.org/10.1007/978-3-319-10593-2_13)

**Poster session / 479**

## Experiences from the CBM collaboration: CAD to ROOT conversion for detector geometries

**Authors:** Simon Nicklas Neuhaus<sup>None</sup>; Omveer Singh<sup>None</sup>; MEHULKUMAR SHIROYA<sup>1</sup>; Patrick Dahm<sup>2</sup>; Eoin Clerkin<sup>3</sup>

<sup>1</sup> GSI Helmholtzzentrum für Schwerionenforschung GmbH(GSI)

<sup>2</sup> GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)

<sup>3</sup> FAIR - Facility for Antiproton and Ion Research in Europe, Darmstadt

**Corresponding Authors:** e.clerkin@gsi.de, sneuhaus@uni-wuppertal.de, m.shiroya@gsi.de, omvir.ch@gmail.com, p.dahm@gsi.de

Fully automated conversion from CAD geometries directly into their ROOT geometry equivalents has been a hot topic of conversation at CHEP conferences. Today multiple approaches for CAD to ROOT conversion exist. Many appear not to work well. In this paper, we report on three separate

and distinct successful efforts from within the CBM collaboration, namely from our Silicon Tracking System team, from our Transition Radiation Detector team, from our Ring-imaging Cherenkov Detector team and from our Beam-Monitor Assembly team on their experiences. These studies benefit from being semi-independent and were reported upon during our biannual CBM collaboration meetings during 2023 and 2024.

In these eventually successful investigative studies, we discuss conversion from CAD in portable STEP format to meshed solids in STL format. Methods which form these into tessellated shapes are discussed. We reported on tessellation based upon TGeoArbN classes with due regard to computational costs occurred in conversion and in simulation. We conduct computation efficiency tests for transport and reconstruction simulations using different subsystems after converting them to VG shapes through a VecGeom converter command in ROOT. We bench tested different navigators and discuss geometry complexity with computational cost.

The main purpose of this contribution is to assess these investigative studies, in order to fix a plan on whether and how CAD to ROOT methods could and should be used by the CBM experiment. As the experiment is in its final design stage, we need to discuss this topics frankly.

**Parallel (Track 2) / 480**

## **Data-driven efficiencies of the LHCb High Level Trigger in Run 3**

**Authors:** Alessandro Scarabotto<sup>1</sup>; Jamie Gooding<sup>1</sup>; Johannes Albrecht<sup>1</sup>

<sup>1</sup> *Technische Universitaet Dortmund (DE)*

**Corresponding Authors:** alessandro.scarabotto@cern.ch, johannes.albrecht@cern.ch, jamie.gooding@cern.ch

Ahead of Run 3 of the LHC, the trigger of the LHCb experiment was redesigned. The L0 hardware stage present in Runs 1 and 2 was removed, with detector readout at 30 MHz passing directly into the first stage of the software-based High Level Trigger (HLT), run on GPUs. Additionally, the second stage of the upgraded HLT makes extensive use of the Turbo event model, wherein only those candidates required for a trigger decision are saved. As the LHCb detector records only events selected by the trigger system, an absolute trigger efficiency cannot be evaluated. The TISTOS method provides a solution to this by evaluating the signal trigger efficiency on a trigger-selected sub-sample independent of signal. Events can be classified as having triggered on signal (TOS), triggered independent of signal (TIS), or both (TISTOS). Efficiencies are then calculated by a tag-and-probe approach, in which TIS and TISTOS events are used as tag and probe, respectively. This approach was applied successfully in Runs 1 and 2; however, in saving only candidates required for trigger decision, all such candidates are TOS by default. The TISTOS method has thus been specified in terms of the stage of selection below each stage of interest to define meaningful efficiencies. This contribution presents the development and performance of the TISTOS method for the upgraded trigger and event model, and an overview of the HLT trigger efficiencies evaluated in 2024 LHCb proton-proton collision data.

**Parallel (Track 4) / 481**

## **HEPCloud Facility Operations at Fermilab—The First Six Years**

**Author:** Steven Timm<sup>1</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Author:** timmsteve@yahoo.com



The HEPCloud Facility at Fermilab has now been in operation for six years. This facility is used to give a unified provisioning gateway to high performance computing centers, including NERSC, ORNL, and ALCF, other large supercomputers run by the NSF, and commercial clouds. HEPCloud delivers hundreds of millions of core-hours yearly for CMS. HEPCloud also serves other Fermilab experiments including DUNE, Mu2E, Muon g-2, and NOvA. In this paper we present the practical considerations of operating a distributed facility such as HEPCloud. We also mention some of the interesting research and development that HEPCloud has been used for including GPU-based machine learning inference servers, and tests of Quantum Computing.

Parallel (Track 5) / 482

## Development of the platform for simulation of spatial distribution for therapeutic dose in Dose-3D phantom

**Authors:** Bartłomiej Rachwał<sup>1</sup>; Jakub Hajduga<sup>2</sup>

**Co-authors:** Bartosz Mindur<sup>1</sup>; Bartłomiej Łach<sup>3</sup>; Dagmara Kulig<sup>2</sup>; Jakub Moron<sup>1</sup>; Kamila Kalecińska; Maciej Kopeć; Piotr Kazimierz Wiacek<sup>1</sup>; Stefan Zenon Koperny<sup>4</sup>; Tomasz Andrzej Fiutowski<sup>1</sup>; Tomasz Szumlak<sup>1</sup>; Damian Kabat<sup>5</sup>

<sup>1</sup> AGH University of Krakow (PL)

<sup>2</sup> AGH University of Krakow

<sup>3</sup> AGH University of Science and Technology

<sup>4</sup> AGH University of Science and Technology (PL)

<sup>5</sup> NIO-PIB, Krakow

**Corresponding Authors:** tomasz.fiutowski@cern.ch, kamkal@agh.edu.pl, bartlomiej.rachwal@cern.ch, dkulig@agh.edu.pl, stefan.koperny@agh.edu.pl, maciej.kopec@agh.edu.pl, piotr.kazimierz.wiacek@cern.ch, damian.kabat@krakow.nio.gov.pl, szumlak@agh.edu.pl, jhajduga@agh.edu.pl, mindur@agh.edu.pl, bartlomiej.lach@fis.agh.edu.pl, jakub.moron@cern.ch

Reconfigurable detector for the measurement of spatial radiation dose distribution for applications in the preparation of individual patient treatment plans 1 was a research and development project aimed at improving radiation dose distribution measurement techniques for therapeutic applications. The main idea behind the initiative was to change the current radiation dose distribution measurement methods used in preparing individual plans for radiotherapy. To this end, a prototype of a fully three-dimensional filled with 3D-printed plastic scintillators 2 was designed and built to measure the spatial distribution of radiation dose in real time for treatment planning.

In parallel with the construction of the phantom, software tools for dose simulation and data analysis are being developed. The Monte Carlo simulations developed for the Dose-3D project were a crucial part of its success. These simulations were used to create data reflecting the apparatus used in the test beam on our simulation platform called G4RT. This simulation data was instrumental in optimizing the Dose-3D detector cell and ensuring the proper calibration of the prototype phantom. It was also used in the design and planning of the test beam, taking into account the effects of beam reflection from the therapeutic table and beam absorption at depth in phantom.

From the software architecture point of view our work presents a novel approach leveraging the Geant4 toolkit. Unlike existing software solutions that are predominantly designed for standard geometries or applications, our proposed toolkit offers unparalleled flexibility, enabling researchers and practitioners to rapidly implement and model custom geometries configurations and experimental scenarios with precision. Additionally, it details the modifications made to the instance of the G4IAEAphspReader library that we use.

1 M. Kopeć, et al. A reconfigurable detector for measuring the spatial distribution of radiation dose for applications in the preparation of individual patient treatment plans. Nuclear Inst. and Methods in Physics Research, A 1048 (2023) 167937

2 D. Kulig, et al. Comparison of cell casted and 3D-printed plastic scintillators for dosimetry applications Radiation Protection Dosimetry, Volume 199, Issue 15-16, October 2023, Pages 1824–1828,

Parallel (Track 9) / 483

## Evolution and Broadening of the National Analysis Facility at DESY

**Authors:** Christian Voss<sup>None</sup>; Christoph Beyer<sup>None</sup>; Elena Gapon<sup>None</sup>; Johannes Reppin<sup>None</sup>; Krunoslav Sever<sup>None</sup>; Martin Flemming<sup>None</sup>; Rachel Wolf<sup>None</sup>; Stefan Dietrich<sup>None</sup>; Thomas Hartmann<sup>1</sup>; Yves Kemp<sup>None</sup>

<sup>1</sup> *Deutsches Elektronen-Synchrotron (DE)*

**Corresponding Authors:** martin.flemming@desy.de, krunoslav.sever@desy.de, johannes.reppin@desy.de, stefan.dietrich@desy.de, christian.voss@desy.de, yves.kemp@desy.de, thomas.hartmann@cern.ch, christoph.beyer@desy.de

The National Analysis Facility at DESY has been in production for nearly 15 years. Over various stages of development, experiences gained in continuous operations have continuously been feed and integrated back into the evolving NAF. As a “living” infrastructure, one fundamental constituent of the NAF is the close contact between NAF users, NAF admins and storage admins & developers. Since the NAF is used by a wide field of physics groups and users with different levels of expertise, the social component has shown to be more crucial for the operations and success of the NAF as large scale tool for science compared to plain technology.

While the NAF’s focus has been on HEP communities and their workflows, the NAF itself is part of the more comprehensive Interdisciplinary Data and Analysis Facility (IDAF) at DESY. ((As an example beyond HEP for emerging user communities from light particle or photon physics,)) the ALPS II experiment has chosen the NAF as their reconstruction and analysis infrastructure. The Any Light Particle Search II (ALPS II) experiment itself is a light-shining-through-a-wall (LSW) experiment based at DESY in Hamburg, Germany, that will search for axions and axion-like particles down to the coupling of the axion to two photons for masses below 0.1meV. With an eager but limited team, ALPS is focused on the physics exploitation with little overhead available for generic computing tasks.

While the computational needs of the ALPS experiment are on a smaller scale than, e.g., the LHC experiments, unspecific data processing tasks like managing computing, storage and data management for online data taking, long term storage, and analyses still need to be handled. Utilizing the experiences from experiment data handling and processing, ALPS II and DESY IT collaborate to offload generic tasks from the scientists and establish a blueprint for other experiments as well. Eyeing for their scientists needs, the ALPS team and DESY IT have been collaborating in establishing modern tools and approaches like concentrating on CI/CD pipelines and flexible storage access methods.

On the administrative side, with a broadening user community an increased focus on security has become paramount. Since an Analysis Facility presents a large attack surface due to the large number of users with varying experiences and backgrounds, hardening an AF has to be one of the core design and operations aims. In addition to the user support view, we will present as well hardening measures taken over the past two years and our future plans for keeping the NAF secure.

Parallel (Track 2) / 484

## HIGH PERFORMANCE ALGORITHMS FOR LOW POWER SUSTAINABLE HARDWARE IN HEP AT VALENCIA

**Authors:** Alvaro Fernandez Casani<sup>1</sup>; Arantza De Oyanguren Campos<sup>1</sup>; Brij Kishor Jashal<sup>2</sup>; Luca Fiorini<sup>1</sup>

<sup>1</sup> *Univ. of Valencia and CSIC (ES)*

<sup>2</sup> *RAL, TIFR and IFIC*

**Corresponding Authors:** brij.kishor.jashal@cern.ch, alvaro.fernandez.casani@cern.ch, luca.fiorini@cern.ch, arantza.de.oyanguren.c

In this talk we present the HIGH-LOW project, which addresses the need to achieve sustainable computational systems and to develop new Artificial Intelligence (AI) applications that cannot be implemented with the current hardware solutions due to the requirements of high-speed response and power constraints. In particular we are focused on the several computing solutions at the Large Hadron Collider (LHC), where most of the computational systems are based on CPUs, but those solutions are not scalable for the future upgrade at high luminosity HL-LHC. Experiments at Future

Colliders (FC) will have to face similar computational challenges and will benefit from our outcomes. The HIGH-LOW project develops the next generation of AI algorithms for high-energy physics, specifically tailored to run in real time and in an energy efficient way. This can also help to provide computational solutions for industrial and real life applications, like autonomous cars, autonomous drones, robotics, wearable medical devices, industrial production, visual inspection of production lines and surveillance.

Our project tackles both the ecological transition to low-power hardware and the shift from serial to highly parallel computing using GPUs, FPGAs, and IPUs. This dual approach enhances the efficiency and sustainability of computational resources in high-energy physics and beyond.

**Poster session / 485**

## Operations Developments at the National Analysis Facility at DESY

**Authors:** Christoph Beyer<sup>None</sup>; Thomas Hartmann<sup>1</sup>; Yves Kemp<sup>None</sup>

**Co-authors:** Christian Voss ; Krunoslav Sever ; Martin Flemming

<sup>1</sup> *Deutsches Elektronen-Synchrotron (DE)*

**Corresponding Authors:** martin.flemming@desy.de, thomas.hartmann@cern.ch, krunoslav.sever@desy.de, christian.voss@desy.de, christoph.beyer@desy.de, yves.kemp@desy.de

The National Analysis Facility (NAF) at DESY is a multi-purpose compute cluster available to a broad community of high-energy particle physics, astro particle physics as well as other communities. Being continuously in production for about 15 years now, the NAF evolved through a number of hardware and software revisions. A constant factor however has been the human factor, as the broad set of user experiences and user interactions with compute and storage infrastructures requires a ongoing support endeavor.

While utilizing compute resources, i.e., CPU cycles, has shown to be the more easy part in running an Analysis Facility, setting up and operating storage as well as authentication/authorization infrastructures and models have always caused the most labour intense work loads. As such easing the operational load is paramount to run an AF efficiently. While solutions for monitoring of the different dimensions of an AF exists, i.e., the compute part or the storage systems, and integrated monitoring has been missing and bridging the information silos has been cumbersome. Thus the need for an integrated monitoring like per user and job file and network socket handles on the compute part being mapped to I/O information on the storage instances.

Furthermore, as an AF is an exposed system with regard to security constraints, care has to be taken to set up a sufficient safe service model in the various steps of designs and operations.

**Poster session / 486**

## Modular Experiment Control System packages for the CBM experiment

**Author:** Pierre-Alain Loizeau<sup>1</sup>

<sup>1</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

**Corresponding Author:** p.-a.loizeau@gsi.de

The Compressed Baryonic Matter (CBM) experiment at FAIR will explore the QCD phase diagram at high net-baryon densities through heavy-ion collisions, using the beams provided by the SIS100 synchrotron in the energy range of 4.5-11 AGeV/c (fully stripped gold ions). This physics program strongly relies on rare probes with complex signatures, for which high interaction rates and a strong selection are needed to achieve the necessary statistics.

These requirements led to the technical decision for a self-triggered and free-streaming data acquisition, followed by an online full reconstruction and selection chain. Such a system can operate reliably and efficiently only with a performant Experiment Control System (ECS) to ensure the synchronization and data quality of all sub-systems.

After looking at existing solutions, the development of a Python based solution focused only on the Experiment Control and on the upper layer of Detector Controls (state and configuration propagation) was instead chosen for CBM. To allow maximal quality checks of the core functions, it will be divided in three levels, from an experiment independent modular core to various user interfaces.

This contribution presents the design choices for this ECS, the technical core package, the CBM ECS core package and the demonstrator GUI packages based on it, as well as the checks and tests done with them.

Poster session / 487

## label-based virtual directories in the dCache storage system

**Authors:** Christopher Green<sup>1</sup>; Dmitry Litvintsev<sup>1</sup>; Krishnaveni Chitrapu<sup>2</sup>; Lea Morschel<sup>None</sup>; Marina Sahakyan<sup>None</sup>; Svenja Meyer<sup>None</sup>; Tigran Mkrtchyan<sup>3</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

<sup>2</sup> *National Supercomputer Centre in Sweden*

<sup>3</sup> *DESY*

**Corresponding Authors:** litvinse@fnal.gov, lea.morschel@desy.de, greenc@fnal.gov, marina.sahakyan@desy.de, tigran.mkrtchyan@desy.de, svenja.meyer@desy.de, krishnaveni@nsc.liu.se

Traditional filesystems organize data in directories. The directories are typically a collection of files whose grouping is based on one criteria, i.e., the starting date of the experiment, experiment name, beamline ID, measurement device, or instrument. However, each file in a directory can belong to different logical groups, such as a special event type, experiment condition, or a part of a selected dataset.

The dCache is a storage system developed to store large amounts of scientific data, used by many HEP and Photon Science experiments.

With recent developments in dCache, we have introduced a concept of file tagging, which dynamically groups files with the same label into virtual directories. The file labels can be added, removed, renamed, and deleted through the admin interface or via Rest API. The files in the virtual directories are exposed through all protocols supported by dCache.

This contribution will describe the details of the implementation for the file tagging in dCache and present our future development plans on automatic metadata extractions, a feature that will significantly simplify data management. Additionally, we are exploring the use of virtual directories as a way to translate scientific data catalogs into filesystem views for direct data analysis

Parallel (Track 5) / 488

## Ultra-Fast Geometry-Independent Highly-Granular Calorimeter Simulation with Diffusion Point Clouds

**Authors:** Anatolii Korol<sup>1</sup>; Erik Buhmann<sup>2</sup>; Frank-Dieter Gaede<sup>3</sup>; Gregor Kasieczka<sup>2</sup>; Katja Kruger<sup>3</sup>; Peter McKeown<sup>4</sup>; William Korcar<sup>2</sup>

<sup>1</sup> *Deutsches Elektronen-Synchrotron (DESY)*

<sup>2</sup> *Hamburg University (DE)*

<sup>3</sup> *Deutsches Elektronen-Synchrotron (DE)*

<sup>4</sup> CERN

**Corresponding Authors:** gregor.kasieczka@cern.ch, erik.buhmann@desy.de, akorol98.17@gmail.com, peter.mckeown@cern.ch, frank-dieter.gaede@cern.ch, katja.kruger@cern.ch, william.korcari@cern.ch

Fast simulation of the energy depositions in high-granular detectors is needed for future collider experiments with ever increasing luminosities. Generative machine learning (ML) models have been shown to speed up and augment the traditional simulation chain. Many previous efforts were limited to models relying on fixed regular grid-like geometries leading to artifacts when applied to highly granular calorimeters with realistic cell layouts. We present CaloClouds III, a novel point cloud diffusion model that allows for high-speed generation of realistic electromagnetic showers due to the distillation into a consistency model. The model is conditioned on incident energy and impact angles and implemented into a realistic DD4hep based simulation model of the ILD detector concept for a future Higgs factory. This is done with the DDFastShowerML library which has been developed to allow for easy integration of generative fast simulation models into any DD4hep based detector model. With this it is possible to benchmark the performance of a generative ML model using fully reconstructed physics events by comparing them against the same events simulated with Geant4, thereby ultimately judging the fitness of the model for application in an experiment's Monte Carlo.

Poster session / 489

## A Managed Tokens Service for Securely Keeping and Distributing Grid Tokens

**Authors:** Shreyas Bhat<sup>1</sup>; Dave Dykstra<sup>2</sup>

<sup>1</sup> Fermi National Accelerator Laboratory (US)

<sup>2</sup> Fermi National Accelerator Lab. (US)

**Corresponding Authors:** dwd@fnal.gov, sbhat@fnal.gov

Fermilab is transitioning authentication and authorization for grid operations to using bearer tokens based on the WLCG Common JWT (JSON Web Token) Profile. One of the functionalities that Fermilab experimenters rely on is the ability to automate batch job submission, which in turn depends on the ability to securely refresh and distribute the necessary credentials to experiment job submit points. Thus, with the transition to using tokens for grid operations, we needed to create a service that would obtain, refresh, and distribute tokens for experimenters' use. This service would avoid the need for experimenters to be experts in obtaining their own tokens and would better protect the most sensitive long-lived credentials. Further, the service needed to be widely scalable, as we are currently keeping credentials active for approximately 15 experiments, each with 1-3 different credentials, and distributing those credentials to 2-20 submit points per experiment, with those numbers steadily increasing. To address these issues, we created and deployed a *Managed Tokens* service. The service is written in Go, taking advantage of that language's native concurrency primitives to easily be able to scale operations as we onboard experiments. The service uses as its first credentials a set of kerberos keytabs, stored on the same secure machine that the *Managed Tokens* service runs on. These kerberos credentials allow the service to use *htgettoken* via `<i>condor_vault_storer</i>` to store vault tokens in the *HTCondor* credential managers (credDs) that run on the batch system scheduler machines (*HTCondor* schedds); as well as downloading a local, shorter-lived copy of the vault token. The kerberos credentials are then also used to distribute copies of the locally-stored vault tokens to experiment submit points. When experimenters schedule jobs to be submitted, these distributed vault tokens are used to access a *Hashicorp Vault* instance (run separately from the *Managed Tokens* service), and previously-stored refresh tokens there are used to obtain the bearer token that is submitted with the job. We will discuss here the design of the *Managed Tokens* service, including elaborating on certain choices we made with regards to concurrent operations, configuration, monitoring, and deployment.

Parallel (Track 5) / 490

## zfit: general likelihood model fitting in Python

**Authors:** Albert Puig Navarro<sup>1</sup>; Iason Krommydas<sup>2</sup>; Jonas Eschle<sup>3</sup>; Matthieu Marinangeli<sup>4</sup>; Rafael Silva Coutinho<sup>3</sup>

<sup>1</sup> *Universität Zürich (CH)*

<sup>2</sup> *Rice University (US)*

<sup>3</sup> *Syracuse University (US)*

<sup>4</sup> *EPFL - Ecole Polytechnique Federale Lausanne (CH)*

**Corresponding Authors:** rafael.silva.coutinho@cern.ch, albert.puig@cern.ch, ik23@rice.edu, jonas.eschle@cern.ch, matthieu.marinangeli@cern.ch

Model fitting using likelihoods is a crucial part of many analyses in HEP.

zfit started over five years ago with the goal of providing this capability within the Python analysis ecosystem by offering a variety of advanced features and high performance tailored to the needs of HEP.

After numerous iterations with users and a continuous development, zfit reached a maturity stage with a stable core and feature set.

In this talk, we will highlight the latest developments. We will discuss its comprehensive feature set, which includes binned and unbinned fits, advanced model building and the ability to create custom models and a variety of available minimizers. Additionally, the talk will cover current and future backend strategies, leveraging TensorFlow and JAX, to deliver state-of-the-art performance on both CPUs and GPUs through extensive optimizations. Furthermore, we will explore the seamless integration of zfit into the broader Python HEP ecosystem, primarily with Scikit-HEP libraries, and its capability to serialize likelihoods in a human-readable format.

Poster session / 491

## Moving a data center keeping availability at the top

**Authors:** Alessandro Pascolini<sup>1</sup>; Andrea Chierici<sup>1</sup>; DIEGO MICHELOTTO<sup>2</sup>; Daniele Cesini<sup>1</sup>; Daniele Lattanzio<sup>None</sup>; Giusy Sergi<sup>3</sup>; Luigi Scarponi<sup>4</sup>; Stefano Zani<sup>None</sup>; Vladimir Sapunenko<sup>5</sup>

<sup>1</sup> *Universita e INFN, Bologna (IT)*

<sup>2</sup> *INFN - National Institute for Nuclear Physics*

<sup>3</sup> *INFN - CNAF*

<sup>4</sup> *INFN*

<sup>5</sup> *INFN-CNAF*

**Corresponding Authors:** v.sapunenko@gmail.com, andrea.chierici@cern.ch, giusy.sergi@infn.cnaf.it, diego.michelotto@cnaf.infn.it, luigi.scarponi@cnaf.infn.it, daniele.cesini@cnaf.infn.it, alessandro.pascolini@cern.ch, stefano.zani@cnaf.infn.it, daniele.lattanzio@cnaf.infn.it

We are moving INFN-T1 data center to a new location. In this presentation we will describe all the steps taken to complete the task without decreasing the general availability of the site and of all the services provided.

We will also briefly describe the new features of our new data center compared to the current one.

Poster session / 492

## Adaptive Hough Transform for Charged Particles Tracking at the LHC

**Authors:** Piotr Libucha<sup>1</sup>; Stefan Michal Horodenski<sup>2</sup>

**Co-author:** Tomasz Bold<sup>2</sup>

<sup>1</sup> AGH University of Krakow

<sup>2</sup> AGH University of Krakow (PL)

**Corresponding Authors:** stefan.michal.horodenski@cern.ch, tomasz.bold@cern.ch

The Adaptive Hough Transform (AHT) is a variant of the Hough transform for particle tracking. Compared to other solutions using Hough Transforms, the benefit of the described algorithm is a shifted balance between memory usage and computation, which could make it more suitable for computational devices with less memory that can be accessed very fast. In addition, the AHT algorithm's flexibility is explored to suppress the number of false positives while maintaining high efficiency.

The algorithm's efficiency has been tested on single muon and pion events as well as high pile-up simulated data consistent with the High Luminosity LHC experiment using the ODD detector available in the ACTS toolkit. The AHT for single muons events yielded an efficiency of over 99% with an average of 9.9 reconstructed particles for a single truth particle. Filtering methods reduced the number of reconstructed particles to 1.8 while maintaining very high tracking efficiency. For pile-up cases, efficiency is above 98%. Additional peak filtering cuts the number of reconstructed tracks over 12 times.

**Parallel (Track6) / 493**

## AccGPT: A CERN Knowledge Retrieval Chatbot

**Authors:** Florian Rehm<sup>1</sup>; Juan Manuel Guijarro<sup>1</sup>; Sofia Vallecorsa<sup>1</sup>; Verena Kain<sup>1</sup>

<sup>1</sup> CERN

**Corresponding Authors:** verena.kain@cern.ch, juan.manuel.guijarro@cern.ch, florian.matthias.rehm@cern.ch, sofia.vallecorsa@cern.ch

In the vast landscape of CERN's internal documentation, finding and accessing relevant detailed information remains a complex and time-consuming task. To address this challenge, the AccGPT project proposes the development of an intelligent chatbot leveraging Natural Language Processing (NLP) technologies. The primary objective is to harness open-source Large Language Models (LLMs) to create a purpose-built chatbot for text knowledge retrieval, with the potential to serve as an assistant for code development and other features in the future.

This initiative was driven by the growing demand at CERN for access to LLMs, not only for building AI Chatbots but also for various other use cases, including Transcription and Translation as a Service (TTaaS), CDS and Zenodo Information Categorization, HR selection processes, and many others. Providing easy and efficient access to LLMs is crucial for the adoption of Generative AI across numerous processes at CERN.

A promising first prototype has already been developed in the realm of knowledge retrieval. It demonstrates a sufficient understanding of user inquiries and provides comprehensive responses utilizing a Retrieval Augmented Generation (RAG) pipeline. However, there is room for improvement to further increase the precision of the responses, which can be achieved by enhancing the retrieval pipeline, considering more powerful and larger LLMs, or fine-tuning the LLMs with more relevant scientific data.

The user interface design and overall user experience of the current prototype chatbot are being iteratively improved, and preparations are underway to make AccGPT available to the community for testing. Automated data scraping and preprocessing pipelines are also being developed to update the chatbot's knowledge base fully autonomously.

Parallel (Track6) / 494

## Collaborative software and maintainability for ePIC experiment at EIC

**Authors:** Dmitry Kalinkin<sup>None</sup>; Markus Diefenthaler<sup>None</sup>; Torre Wenaus<sup>1</sup>; Wouter Deconinck<sup>None</sup>

<sup>1</sup> *Brookhaven National Laboratory (US)*

**Corresponding Authors:** wouter.deconinck@umanitoba.ca, veprbl+cern@gmail.com, wenaus@gmail.com, mdiefent@jlab.org

The ePIC collaboration is working towards realizing the primary detector for the upcoming Electron-Ion Collider (EIC). As ePIC approaches critical decision milestones and moves towards future operation, software plays a critical role in systematically evaluating detector performance and laying the groundwork for achieving the scientific goals of the EIC project. The scope and schedule of the project require a balanced approach between near-term priorities, such as preparing the Technical Design Report, and long-term objectives for the future construction, commissioning, and operational phases. ePIC leverages an agile development process with high-level milestones to ensure continuous real-world testing of the software through monthly production campaigns and CI-driven benchmarks. The ePIC software stack embraces cutting-edge, sustainable community software tools and avoids the “not invented here” syndrome by building on top of well-supported and actively developed frameworks like the key4HEP stack (DD4hep, PODIO, EDM4hep) and ACTS. This collaborative development approach fosters an elevated standard of quality based on lessons learned by the nuclear physics and high energy physics communities. This talk will explore our setup for a collaborative development process and how it integrates with our vision for Software & Computing in the future ePIC experiment.

Parallel (Track 7) / 495

## Monitoring large-scale dCache installations with storage events using Kafka streams

**Authors:** Christian Voss<sup>None</sup>; Marina Sahakyan<sup>None</sup>; Tigran Mkrtchyan<sup>1</sup>

**Co-authors:** Alexander Trautsch<sup>1</sup>; Christian Sperl<sup>1</sup>; Felix Christians<sup>1</sup>; Juergen Hannappel<sup>2</sup>; Sandro Grizzo<sup>1</sup>; Thomas Hartmann<sup>3</sup>

<sup>1</sup> *DESY*

<sup>2</sup> *DESY It*

<sup>3</sup> *Deutsches Elektronen-Synchrotron (DE)*

**Corresponding Authors:** juergen.hannappel@desy.de, sandro.grizzo@desy.de, christian.voss@desy.de, alexander.trautsch@desy.de, felix.christians@desy.de, christian.sperl@desy.de, tigran.mkrtchyan@desy.de, thomas.hartmann@cern.ch, marina.sahakyan@desy.de

DESY operates multiple dCache storage instances for multiple communities. As each community has different workflows and workloads, their dCache installations range from very large instances with more than 100 PB of data, to instances with up to billions of files or instances with significant LAN and WAN I/O.

To successfully operate all instances and quickly identify issues and performance bottlenecks, DESY IT relies for monitoring heavily on dCache own storage events. Each atomic operation in the distributed storage instances triggers a storage event with details to the corresponding transfer or service status change.

These events are collected and parsed through an Apache Kafka event streaming bus. From the Kafka event stream, the events are aggregated in an Elastic Search+Lucene based database and search engine for on the fly operational diagnostics and analytics. Beyond day to day operations, an on demand Apache Spark cluster on top of the National Analysis Facility at DESY is used for in detail analyses of operational data to extract information over a wide time span and number of storage events. In a similar fashion, all dCache logging messages are also processed through Kafka stream



allowing to employ a passive monitoring waiting for specific signature to raise an alarm. In the future ML and AI algorithms for predictive maintenance are in the development pipeline. Furthermore, additional matrices are collecting from the dCache pools themselves and also pushed to Kafka to generate an almost complete picture of the dCache instances.

In this talk, we present our aggregation and analyses pipelines and workflows and how they are enabling DESY IT to scale out dCache storages for heterogeneous user groups and use cases.

Parallel (Track 2) / 496

## FPGA-RICH: A low-latency, high-throughput online particle identification system for the NA62 experiment

**Author:** Pierpaolo Perticaroli<sup>1</sup>

**Co-authors:** Alessandro Lonardo<sup>1</sup>; Andrea Biagioni<sup>2</sup>; Carlotta Chiarini<sup>3</sup>; Cristian Rossi<sup>3</sup>; Francesca Lo Cicero<sup>1</sup>; Francesco Simula<sup>4</sup>; Luca Pontisso<sup>1</sup>; Matteo Turisini<sup>3</sup>; Mauro Raggi<sup>5</sup>; Michele Martinelli<sup>2</sup>; Ottorino Frezza<sup>1</sup>; Paolo Cretaro<sup>6</sup>; Piero Vicini<sup>7</sup>; Roberto Ammendola<sup>8</sup>; Roberto Piandani<sup>9</sup>; Andrea Ciardiello

<sup>1</sup> INFN, Roma I (IT)

<sup>2</sup> INFN

<sup>3</sup> Sapienza Università e INFN, Roma I (IT)

<sup>4</sup> Istituto Nazionale di Fisica Nucleare - sezione di Roma 1

<sup>5</sup> CERN and Sapienza

<sup>6</sup> INFN - National Institute for Nuclear Physics

<sup>7</sup> INFN Rome Section

<sup>8</sup> INFN e Università Roma Tor Vergata (IT)

<sup>9</sup> Univ. Autonoma de San Luis Potosi

**Corresponding Authors:** roberto.piandani@cern.ch, andrea.biagioni@roma1.infn.it, michele.martinelli@roma1.infn.it, mauro.raggi@cern.ch, andrea.ciardiello@uniroma1.it, paolo.cretaro@roma1.infn.it, pierpaolo.perticaroli@roma1.infn.it, roberto.ammendola@roma2.infn.it, francesca.locicero@roma1.infn.it, cristian.rossi@roma1.infn.it, carlotta.chiarini@roma1.infn.it, alessandro.lonardo@roma1.infn.it, ottorino.frezza@roma1.infn.it, francesco.simula@roma1.infn.it, matteo.turisini@cern.ch, luca.pontisso@roma1.infn.it, piero.vicini@roma1.infn.it

This work presents FPGA-RICH, an FPGA-based online partial particle identification system for the NA62 experiment utilizing AI techniques. Integrated between the readout of the Ring Imaging Cherenkov detector (RICH) and the low-level trigger processor (L0TP+), FPGA-RICH implements a fast pipeline to process in real-time the RICH raw hit data stream, producing trigger-primitives containing elaborate physics information, such as the number of charged particles in a physics event, that the L0TP+ can use to improve trigger decision efficiency.

The system is deployed on a single FPGA device and uses both classical online processing methods and a compact Neural Network algorithm to achieve efficient event classification while managing NA62's challenging throughput requirements ( $\approx 10$  MHz). The streaming pipeline guarantees low latency ( $\sim 1 \mu\text{s}$ ), comparable to the other NA62 sub-detectors that send trigger-primitives to the L0TP+, allowing seamless integration in the existing TDAQ setup as a new detector running in parallel with the RICH.

The system development leverages High Level Synthesis (HLS) programming language and the open-source *hls4ml* software-hardware codesign workflow for fast, flexible and relatively simple reprogramming, debugging and feature enhancements.

Currently integrated in parasitic mode in the experiment TDAQ, we will present and discuss our experience with the system's design and deployment along with the results obtained during the 2024 data taking.

The work highlights the strength and maturity of modern computing solutions, programming paradigms and machine learning algorithms even within the challenging context of modern HEP experiments' online data acquisition and analysis.

Poster session / 497

## Collaborative Tools for the ePIC Experiment

Author: Maxim Potekhin<sup>1</sup><sup>1</sup> Brookhaven National Laboratory (US)

Corresponding Author: potekhin@bnl.gov

The ePIC Collaboration is actively working on the Technical Design Report (TDR) for its future detector at the Electron Ion Collider to be built at Brookhaven National Laboratory within the next decade. The development of the TDR by an international Collaboration with over 850 members requires a plethora of physics and detector studies that need to be coordinated. An effective set of Collaborative Tools and an open, collaborative environment are instrumental for the success of this effort. This includes the Collaboration Web presence, modern digital repositories (Zenodo), collaborative document development (Google Docs, Overleaf), file sharing (XRootD, Google Drive), communication (Zoom and Mattermost). ePIC is leveraging GitHub for its shared development environment and code version control and validation. Current activities of the ePIC Collaboration in this area are informed in part by the previous successful Data and Analysis Preservation effort of the PHENIX Collaboration at RHIC. This included a complete redesign of the Collaboration's public website with the goal of simplified long-term maintenance, preservation of its software environment using containerization techniques, and migration of its various research materials to Zenodo. In this presentation we describe the technology choices and progress in the area of the ePIC Collaborative Tools, covering the emerging web presence to replace the existing Wiki, communications and collaborative software development, digital repositories and routes to possible future data migration to new document development workflow systems.

Parallel (Track 3) / 498

## From Hope to Heuristic: Realistic Runtime Estimates for Quantum Optimisation in NHEP

Authors: Maja Franz<sup>1</sup>; Manuel Schönberger<sup>1</sup>; Melvin Strobl<sup>2</sup>; Eileen Kuehn<sup>2</sup>; Achim Streit<sup>3</sup>; Pia Zurita<sup>4</sup>; Markus Diefenthaler<sup>None</sup>; Wolfgang Mauerer<sup>1</sup><sup>1</sup> Technical University of Applied Sciences, Regensburg<sup>2</sup> Karlsruhe Institute of Technology<sup>3</sup> Karlsruhe Institute of Technology (KIT)<sup>4</sup> Universidad Complutense de Madrid

Corresponding Authors: melvin.strobl@kit.edu, marzurit@ucm.es, wolfgang.mauerer@othr.de, manuel.schoenberger@othr.de, achim.streit@kit.edu, mdiefent@jlab.org, eileen.kuehn@kit.edu, maja.franz@othr.de

Noisy intermediate-scale quantum (NISQ) computers, while limited by imperfections and small scale, hold promise for near-term quantum advantages in nuclear and high-energy physics (NHEP) when coupled with co-designed quantum algorithms and special-purpose quantum processing units.

Developing co-design approaches is essential for near-term usability, but inherent challenges exist due to the fundamental properties of NISQ algorithms.

In this contribution we therefore investigate the core algorithms, which can solve optimisation problems via the abstraction layer of a quadratic Ising model or general unconstrained binary optimisation problems (QUBO), namely quantum annealing (QA) and the quantum approximate optimisation algorithm (QAOA).

Applications in NHEP utilising QUBO formulations range from particle track reconstruction, over job scheduling on computing clusters to experimental control.

While QA and QAOA do not inherently imply quantum advantage, QA runtime for specific problems can be determined based on the physical properties of the underlying Hamiltonian, albeit it is a computationally hard problem itself.

Our primary focus is on two key areas:

Firstly, we estimate runtimes and scalability for common NHEP problems addressed via QUBO formulations by identifying minimum energy solutions of intermediate Hamiltonian operators encountered during the annealing process.

Secondly, we investigate how the classical parameter space in the QAOA, together with approximation techniques such as a Fourier-analysis based heuristic, proposed by Zhou et al. (2018), can help to achieve (future) quantum advantage, considering a trade-off between computational complexity and solution quality.

Our computational analysis of seminal optimisation problems suggests that only lower frequency components in the parameter space are of significance for deriving reasonable annealing schedules, indicating that heuristics can offer improvements in resource requirements, while still yielding near-optimal results.

Parallel (Track 9) / 499

## Reshaping Analysis for Fast Turnaround

**Authors:** Austin Townsend<sup>1</sup>; Barry Sly-Delgado<sup>2</sup>; Benjamin Tovar Lopez<sup>2</sup>; Connor Moore<sup>1</sup>; Douglas Thain<sup>2</sup>; Kevin Patrick Lannon<sup>1</sup>

<sup>1</sup> *University of Notre Dame (US)*

<sup>2</sup> *University of Notre Dame*

**Corresponding Authors:** btovar@nd.edu, cmoore24@nd.edu, dthain@cse.nd.edu, austin.townsend@cern.ch, barysly@gmail.com, klannon@nd.edu

In the data analysis pipeline for LHC experiments, a key aspect is the step in which small groups of researchers—typically graduate students and postdocs—reduce the smallest, common-denominator data format down to a small set of specific histograms suitable for statistical interpretation. Here, we will refer to this step as “analysis” with the recognition that in other contexts, “analysis” might include other pieces, such as the actual computation required to extract statistical interoperation from the histograms. Analysis is a very important part of the pipeline as it is the step where individual researchers exercise their creativity in trying new ideas in the pursuit of discovery. Therefore, a critical metric for the analysis step is turnaround time because it determines how rapidly researchers can explore their space of ideas. We demonstrate our experience reshaping late-stage analysis applications on thousands of nodes with the goal of minimizing turnaround time. It is not enough merely to increase scale: it is necessary to make changes throughout the stack, including storage systems, data management, task scheduling, and application design. We demonstrate these changes when applied to CMS analysis applications built using the Coffea framework, leveraging Dask and TaskVine to scale out to distributed resources. We evaluate the performance of the applications on opportunistic campus clusters, showing effective scaling up to 7200 cores, thus producing significant improvement in turnaround time.

Parallel (Track 7) / 500

## Hardware and software design of APEnetX: a custom FPGA-based NIC for scientific computing

**Authors:** Alessandro Lonardo<sup>1</sup>; Andrea Biagioni<sup>2</sup>; Carlotta Chiarini<sup>3</sup>; Cristian Rossi<sup>3</sup>; Elena Pastorelli<sup>2</sup>; Francesca Lo Cicero<sup>1</sup>; Francesco Simula<sup>4</sup>; Luca Pontisso<sup>1</sup>; Michele Martinelli<sup>2</sup>; Ottorino Frezza<sup>1</sup>; Paolo Cretaro<sup>5</sup>; Pier Stanislao Paolucci<sup>2</sup>; Piero Vicini<sup>6</sup>; Roberto Ammendola<sup>7</sup>

<sup>1</sup> *INFN, Roma I (IT)*

<sup>2</sup> *INFN*

<sup>3</sup> *Sapienza Universita e INFN, Roma I (IT)*

<sup>4</sup> *Istituto Nazionale di Fisica Nucleare - sezione di Roma 1*

<sup>5</sup> INFN - National Institute for Nuclear Physics

<sup>6</sup> INFN Rome Section

<sup>7</sup> INFN e Universita Roma Tor Vergata (IT)

**Corresponding Authors:** michele.martinelli@roma1.infn.it, pier.paolucci@roma1.infn.it, ottorino.frezza@roma1.infn.it, cristian.rossi@roma1.infn.it, paolo.cretaro@roma1.infn.it, luca.pontisso@roma1.infn.it, alessandro.lonardo@roma1.infn.it, francesco.simula@roma1.infn.it, andrea.biagioni@roma1.infn.it, piero.vicini@roma1.infn.it, roberto.ammendola@roma2.infn.it, elena.pastorelli@roma1.infn.it, carlotta.chiarini@roma1.infn.it, francesca.locicero@roma1.infn.it

In the High-Performance Computing (HPC) field, fast and reliable interconnects remain pivotal in delivering efficient data access and analytics.

In recent years, several interconnect implementations have been proposed, targeting optimization, reprogrammability and other critical aspects. Custom Network Interface Cards (NIC) have emerged as viable alternatives to commercially available products, which often come with high price tags and limited or no customization options.

In this field, the APEnet project has been and continues to be engaged in developing custom FPGA-based NICs tailored for toroidal interconnection systems dedicated to scientific computing and simulations: leveraging a custom network protocol and being easily portable and reconfigurable, it ensures adaptability across various scientific domains spanning from High Energy Physics to Brain Simulation; it implements a 3D direct torus interconnect, which nested in a multi-tier topology, enables high path diversity, short cabling at low dimension and high efficiency.

In this work, we present the latest advancements for the APEnet NIC, APEnetX, which integrates cutting-edge Xilinx Ultrascale+ technologies with custom hardware and software components to enable Remote Direct Memory Access (RDMA) functionalities targeting both the remote hosts and accelerators such as GPUs. A custom network protocol is used, accompanied by Quality-of-Service (QoS) functionalities, to ensure efficient data transfers between nodes even in the event of critical congestion states. Finally, we developed the necessary libraries to replicate APEnetX in a simulated environment (Omnet++): the emulation of the network at large scale enables us to tailor the architecture for specific scientific applications.

Parallel (Track 9) / 501

## Supporting the development of Machine Learning for fundamental science in a federated Cloud with the AI\_INFNO platform

**Authors:** Lucio Anderlini<sup>1</sup>; Matteo Barbetti<sup>2</sup>; Giulio Bianchini<sup>None</sup>; Diego Ciangottini<sup>3</sup>; Carmelo Pellegrino<sup>None</sup>; Rosa Petrini<sup>4</sup>; Daniele Spiga<sup>5</sup>

<sup>1</sup> *Universita e INFNO, Firenze (IT)*

<sup>2</sup> *INFNO CNAF*

<sup>3</sup> *INFNO, Perugia (IT)*

<sup>4</sup> *INFNO Sezione di Pisa, Universita' e Scuola Normale Superiore, P*

<sup>5</sup> *Universita e INFNO, Perugia (IT)*

**Corresponding Authors:** giulio.bianchini@studenti.unipg.it, carmelo.pellegrino@cnafe.infn.it, rosa.petrini@cern.ch, lucio.anderlini@cern.ch, matteo.barbetti@cern.ch, daniele.spiga@cern.ch, diego.ciangottini@cern.ch

Machine Learning (ML) is driving a revolution in the way scientists design, develop, and deploy data-intensive software. However, the adoption of ML presents new challenges for the computing infrastructure, particularly in terms of provisioning and orchestrating access to hardware accelerators for development, testing, and production.

The INFNO-funded project AI\_INFNO ("Artificial Intelligence at INFNO") aims at fostering the adoption of ML techniques within INFNO use cases by providing support on multiple aspects, including the provision of AI-tailored computing resources. It leverages cloud-native solutions in the context of INFNO Cloud, to share hardware accelerators as effectively as possible, ensuring the diversity of the Institute's research activities is not compromised.

In this contribution, we provide an update on the commissioning of a Kubernetes platform designed to ease the development of GPU-powered data analysis workflows and their scalability on heterogeneous, distributed computing resources, possibly federated as Virtual Kubelets with the interLink provider.

Finally we showcase the deployment of the training and validation infrastructure for the flash-simulation pipeline of the LHCb experiment, known as Lamarr, providing a practical example of how our infrastructure supports complex ML workflows in high-energy physics.

Parallel (Track 2) / 502

## An optimized C++ software for the management of Timepix4 data acquisition and analysis

**Authors:** Alberto Gianoli<sup>1</sup>; Massimiliano Fiorini<sup>1</sup>; Nicolo Vladi Biesuz<sup>1</sup>; Riccardo Bolzonella<sup>2</sup>; Sebastiano Schifano<sup>1</sup>; Viola Cavallini<sup>1</sup>; Xavi Llopart Cudie<sup>3</sup>

<sup>1</sup> *Universita e INFN, Ferrara (IT)*

<sup>2</sup> *University of Ferrara and INFN*

<sup>3</sup> *CERN*

**Corresponding Authors:** riccardo.bolzonella@cern.ch, sebastiano.schifano@cern.ch, xavier.llopart@cern.ch, massimiliano.fiorini@cern.ch, nicolo.vladi.biesuz@cern.ch, viola.cavallini@cern.ch, alberto.gianoli@cern.ch

Timepix4 is an innovative multi-purpose ASIC developed by the Medipix4 Collaboration at CERN for fundamental and applied physics detection systems. It is composed by a  $\sim 7\text{cm}^2$  area matrix with about 230k independent pixels, each one with a charge integration circuit, a discriminator and a time-to-digital converter that allows to measure Time-of-Arrival with 195 ps width bins and Time-over-Threshold with 1.56 ns width bins. Timepix4 can produce up to 160 Gbps of output data, so a strong software counterpart is needed for fast and efficient data processing.

We developed an open-source multi-thread C++ framework to manage the Timepix4 ASIC, regardless of which control board is used for communication with the server. The software can configure Timepix4 through low and high level functions, depending on the final user's expertise and his need for customization. Those methods also allow the user to easily perform complex routines, like pixel matrix equalization and calibration, with user-friendly C++ scripts.

When the acquisition starts, some read-out threads can safely store Timepix4 data on disk. Offline post-acquisition classes can be used to analyze the data, using a custom clustering algorithm that can process more than 1M events/s and, if needed, an ad-hoc convolutional neural network for particle track identification. If the acquisition rate is lower than 1M events/s, clustering can be performed online, exploiting a dedicated thread, connected to read-out ones, that runs the same algorithm. Moreover, an online monitor thread can be connected to clustering object to view up to  $O(100)\text{kEvents/s}$ , showing a hit-map and real-time statistics like cluster dimension and energy.

In this contribution we will present the software architecture, its performances and some results obtained during acquisitions using radioactive sources, X-ray tubes and monochromatic synchrotron X-ray beams.

Parallel (Track 4) / 503

## Supporting medium/small-sized experiments in the transition from X.509 to JWTs

**Authors:** Aksieniiia Shtimmerman<sup>None</sup>; Alessandro Pascolini<sup>1</sup>; Carmelo Pellegrino<sup>None</sup>; Carmen Giugliano<sup>None</sup>

**Co-authors:** Andrea Rendina ; Daniele Lattanzio ; Lucia Morganti ; Matteo Barbetti <sup>2</sup>

<sup>1</sup> *Universita e INFN, Bologna (IT)*

<sup>2</sup> *INFN CNAF*

**Corresponding Authors:** andrea.rendina@cnaif.infn.it, carmen.giugliano@cern.ch, lucia.morganti@cnaif.infn.it, aksieniia.shtimmerman@cnaif.infn.it, daniele.lattanzio@cnaif.infn.it, alessandro.pascolini@cern.ch, carmelo.pellegrino@cnaif.infn.it, matteo.barbetti@cern.ch

X.509 certificates and VOMS proxies are still widely used by various scientific communities for authentication and authorization (authN/Z) in Grid Storage and Computing Elements. Although this has contributed to improve the scientific collaboration worldwide, X.509 authN/Z comes with some interoperability issues with modern Cloud-based tools and services.

The Grid computing communities have decided to migrate to token-based authentication, a new web technology that has proved to be flexible and secure.

The model being recently adopted by the communities is based on industrial standards such as OAuth2 and OpenID-Connect and exploits JSON Web Tokens (JWT): a compact way to securely transmit information as JSON objects.

JWT are usually short-lived and provide fine-grained authorization, based on “scopes”, to perform specific actions.

These scopes are embedded into the token and are specified during the request procedure so they last only until token expiration time. Scopes can be requested based on user groups and permission thus providing the possibility of restricting a group to perform only a subset of actions.

These characteristics make up to a more secure alternative to X.509 proxies.

Being largely used in industries, JWTs are also easily integrated in services not specifically developed for the scientific community, such as calendars, Sync and Share services, collaborative software development platforms, and more.

As such, JWTs suit the many heterogeneous demands of Grid communities and some of them already started the transition in 2022.

In the Italian WLCG Tier-1, located in Bologna and managed by INFN - CNAF, several computing resources are hosted and made available to scientific collaborations in the fields of High-Energy Physics, Astroparticle Physics, Gravitational Waves, Nuclear Physics and many others.

Although LHC experiments at CERN are the main users of CNAF resources, many other communities and experiments are being supported in their computing activities.

While the main LHC experiments have already planned their own transition from X.509 to token-based authN/Z, many medium/small-sized collaborations struggle to put effort into it.

The Tier-1 User Support unit has the duty of guiding users towards efficient and modern computing techniques and workflows involving data and computing resources access.

As such, the User Support group is playing a central role in preparing documentation, tools and services to ease the transition from X.509 to JWTs.

The foreseen support strategy and the related tools will be presented. Future workflow plans in view of the complete transition will also be presented.

**Parallel (Track 5) / 504**

## **Efficient Tracking Algorithm Evaluations through Multi-Level Reduced Simulations**

**Author:** Uraz Odyurt<sup>1</sup>

**Co-authors:** Ana-Lucia Varbanescu<sup>2</sup>; Sascha Caron<sup>1</sup>

<sup>1</sup> *Nikhef National institute for subatomic physics (NL)*

<sup>2</sup> *University of Twente*

**Corresponding Authors:** a.l.varbanescu@utwente.nl, sascha.caron@cern.ch, u.odyurt@nikhef.nl

Subatomic particle track reconstruction (tracking) is a vital task in High-Energy Physics experiments. Tracking, in its current form, is exceptionally computationally challenging. Fielded solutions, relying on traditional algorithms, do not scale linearly and pose a major limitation for the HL-LHC era. Machine Learning (ML) assisted solutions are a promising answer.

Current ML model design practice is predominantly ad hoc. We aim for a methodology for automated search of model designs, consisting of complexity reduced descriptions of the main problem, forming a complexity spectrum. As the main pillar of such a method, we provide the REDUCED VIRTUAL DETECTOR (REDVID) as a complexity-aware detector model and particle collision event simulator. Through a multitude of configurable dimensions, REDVID is capable of simulations throughout the complexity spectrum. REDVID can also act as a simulation-in-the-loop, to both generate synthetic data efficiently and to simplify the challenge of ML model design evaluation. With REDVID, starting from the simplistic end of the complexity spectrum, lesser designs can be eliminated in a systematic fashion, early on. REDVID is not bound by real detector geometries and can be considered for simulations involving arbitrary detector designs.

As a simulation and a generative tool for ML-assisted solution design, REDVID is highly flexible, reusable and open-source. Reference data sets generated with REDVID are publicly available. Data generated using REDVID has enabled rapid development of multiple novel ML model designs, which is currently ongoing.

Poster session / 505

## Identification of muon-electron elastic events using Graph Neural Networks for precision measurements.

**Authors:** Anna Driutti<sup>1</sup>; Damian Mizera<sup>2</sup>; Emma Hess<sup>1</sup>; Marcin Kucharczyk<sup>3</sup>; Marcin Wolter<sup>3</sup>; Mateusz Jacek Goncerz<sup>3</sup>; Milosz Zdybal<sup>3</sup>; Patrick Asenov<sup>1</sup>

<sup>1</sup> *Universita & INFN Pisa (IT)*

<sup>2</sup> *Cracow University of Technology (PL)*

<sup>3</sup> *Polish Academy of Sciences (PL)*

**Corresponding Authors:** damian.mizera@outlook.com, anna.driutti@unipi.it, marcin.kucharczyk@cern.ch, ehess@cern.ch, milosz.zdybal@cern.ch, mateusz.goncerz@cern.ch, patrick.asenov.asenov@cern.ch, marcin.wolter@cern.ch

Precision measurements of fundamental properties of particles serve as stringent tests of the Standard Model and search for new physics. These experiments require robust particle identification and event classification capabilities, often achievable through machine learning techniques. This presentation introduces a Graph Neural Network (GNN) approach tailored for identifying outgoing particles in elastic events where a muon beam interacts with the atomic electrons of thin low- $Z$  targets in a series of tracking stations containing silicon strip modules. The processes include, among others, ionization and pair production (resulting in  $e^+e^-$  pairs) caused by muons. We illustrate the application of the developed technique through a case study utilizing simulated data of a reduced geometrical configuration of the MUonE experiment, which aims to precisely measure the leading hadronic contribution to the muon magnetic moment anomaly at CERN.

Parallel (Track 2) / 506

## Free-streaming online tracking in CBM

**Authors:** Sergei Zharko<sup>1</sup>; Sergey Gorbunov<sup>1</sup>; Valentina Akishina<sup>2</sup>

<sup>1</sup> *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

<sup>2</sup> *Goethe University Frankfurt (DE)*

**Corresponding Authors:** zharkosergey94@gmail.com, sergey.gorbunov@cern.ch, valentina.akishina@cern.ch

The event reconstruction in the CBM experiment is challenging. There will be no simple hardware trigger due to the novel concepts of free-streaming data and self-triggered front-end electronics.

Thus, there is no a priori association of signals to physical events.

CBM will operate at interaction rates of 10 MHz, unprecedented for heavy ion experiments.

At this rate, collisions overlap in time and are to be resolved in software by reconstruction algorithms.

These complications made the speed and quality of the data reconstruction crucial.

The core of the track reconstruction is the Cellular Automaton (CA) based algorithm used for the Silicon Tracking System (STS).

It digests free-streaming data both online and offline, taking large time slices of the hit measurements as input with non-a priori-defined physical collisions.

The data is reconstructed in time portions by applying a nonmerging sliding window algorithm, which achieves almost constant time per event regardless of the time slice size.

The algorithm was successfully applied to run online for the mini-CBM experiment during the March 2024 data-taking campaign.

Parallel (Track 1) / 507

## New GridKa Tape Storage System –from design to production deployment

**Author:** Dorin-Daniel Lobontu<sup>None</sup>

**Co-authors:** Andreas Petzold<sup>1</sup>; Artur Il Darovic Gottmann<sup>1</sup>; Doris Ressmann; Haykuhi Musheghyan<sup>2</sup>; Preslav Konstantinov<sup>1</sup>; Samuel Ambroj Perez; Xavier Mol<sup>1</sup>

<sup>1</sup> KIT - Karlsruhe Institute of Technology (DE)

<sup>2</sup> Georg August Universitaet Goettingen (DE)

**Corresponding Authors:** xavier.mol@cern.ch, dorin-daniel.lobontu@kit.edu, artur.akhmetshin@cern.ch, samuel.perez@kit.edu, andreas.petzold@cern.ch, preslav.konstantinov@cern.ch, doris.ressmann@kit.edu, haykuhi.musheghyan@cern.ch

Storing the ever-increasing amount of data generated by LHC experiments is still inconceivable without making use of the cost effective, though inherently complex, tape technology. GridKa tape storage system used to rely on IBM Spectrum Protect (SP). Due to a variety of limitations and to meet the even higher requirements of HL-LHC project, GridKa decided to switch from SP to High Performance Storage System (HPSS).

Even though HPSS is a highly scalable and performant tape management software, it required special adjustments to fulfill all GridKa requirements. Based on the experience gained with the former tape system, the implementation team developed specific stress scenarios. Running these tests and interpreting their results allowed a successful adaptation of HPSS and made it the core component of GridKa tape storage system.

To increase the performance the architecture of the system was reshaped and the stored data has been collocated in a more appropriate, tape-oriented way to match the requirements of every single experiment and HPSS demands. In total 70 PB of data and 40 million files were migrated from the legacy to the new tape system at GridKa.

This contribution presents the internal architecture of the new tape storage system, the implementation and migration process, the encountered issues, the achieved results and ongoing work on open items.

Parallel (Track 5) / 508



## Accelerating detector simulations with Celeritas: performance improvements and new capabilities

**Author:** Seth Johnson<sup>1</sup>

**Co-authors:** Amanda Lund<sup>2</sup>; Benjamin Morgan<sup>2</sup>; Elliott Biondo<sup>3</sup>; Guilherme Lima<sup>4</sup>; Hayden Richard Hollenbeck<sup>5</sup>; Julien Esseiva<sup>6</sup>; Philippe Canal<sup>7</sup>; Soon Yung Jun<sup>7</sup>; Stefano Tognini<sup>3</sup>; Tom Evans<sup>3</sup>

<sup>1</sup> *Oak Ridge National Laboratory (US)*

<sup>2</sup> *University of Warwick*

<sup>3</sup> *Oak Ridge National Laboratory*

<sup>4</sup> *FermiLab (US)*

<sup>5</sup> *University of Virginia (US)*

<sup>6</sup> *Lawrence Berkeley National Lab. (US)*

<sup>7</sup> *Fermi National Accelerator Lab. (US)*

**Corresponding Authors:** soon.yung.jun@cern.ch, guilherme.lima@cern.ch, kcr8wx@virginia.edu, philippe.canal@cern.ch, benjamin.morgan@cern.ch, alund@anl.gov, julien.esseiva@cern.ch, stognini@gmail.com, tmdelellis@gmail.com, biondoed@ornl.gov, johnsonsr@ornl.gov

Celeritas is a rapidly developing GPU-enabled detector simulation code aimed at accelerating the most computationally intensive problems in high energy physics. This presentation will highlight exciting new performance results for complex subdetectors from the CMS and ATLAS experiments using EM secondaries from hadronic interactions. The performance will be compared on both Nvidia and AMD GPUs as well as multicore CPUs, made possible by a new native Celeritas geometry representation of Geant4 geometry objects. This new surface-based geometry, ORANGE, provides a robust and efficient navigation engine fundamentally different from existing detector simulation models. Finally, we introduce two new physics capabilities to Celeritas, optical photon tracking and extended EM models, that demonstrate the code's extensibility and promise potential applications beyond LHC detectors.

Parallel (Track 8) / 509

## Open Data at ATLAS: Bringing TeV collisions to the World

**Authors:** António Costa<sup>1</sup>; Kate Shaw<sup>2</sup>; Steven Goldfarb<sup>3</sup>

<sup>1</sup> *IST*

<sup>2</sup> *University of Sussex (GB)*

<sup>3</sup> *University of Melbourne (AU)*

**Corresponding Authors:** steven.goldfarb@cern.ch, antonio.jacques.costa@tecnico.ulisboa.pt, kate.shaw@cern.ch

ATLAS Open Data for Education delivers proton-proton collision data from the ATLAS experiment at CERN to the public along with open-access resources for education and outreach. To date ATLAS has released a substantial amount of data from 8 TeV and 13 TeV collisions in an easily-accessible format and supported by dedicated documentation, software, and tutorials to ensure that everyone can access and exploit the data for different educational objectives. Along with datasets, ATLAS also provides data visualisation tools and interactive web based applications for studying the data, along with Jupyter Notebooks and downloadable code enabling users to further analyse data for known and unknown physics cases. The Open Data educational platform which hosts the data and tools is used by tens of thousands of students worldwide, and we present the project development, lessons learnt, impacts, and future goals.

Parallel (Track 8) / 510

## Recent Updates to the Popular ATLAS Virtual Visit Programme

**Authors:** Muhammad Alhroob<sup>1</sup>; Steven Goldfarb<sup>2</sup>

<sup>1</sup> *University of Warwick*

<sup>2</sup> *University of Melbourne (AU)*

**Corresponding Authors:** steven.goldfarb@cern.ch, muhammad.alhroob@cern.ch

Virtual Visits have been an integral component of the ATLAS Education and Outreach programme since their inception in 2010. Over the years, collaboration members have hosted visits for tens of thousands of visitors located all over the globe. In 2024, alone there have already been 59 visits through the month of May. Visitors in classrooms, festivals, events or even at home have a unique opportunity to engage with scientists located either underground in the ATLAS experimental cavern or in front of the control room, to learn about the goals and achievements of the collaboration. As part of the renovation of the ATLAS Visitor Centre at LHC Point 1, a new installation was constructed to facilitate Virtual Visits during the running of LHC. We present the overall programme, the new installation and discuss recent initiatives to expand our reach, including Open Visits on Zoom, Facebook, YouTube and TikTok Live.

**Parallel (Track 7) / 511**

## Unlocking the compute continuum: scaling out from cloud to HPC and HTC resources

**Authors:** Daniele Spiga<sup>1</sup>; Diego Ciangottini<sup>2</sup>

**Co-authors:** Ahmed Shiraz Memon<sup>3</sup>; Andrea Manzi; Andrej Filipcic<sup>4</sup>; Antonino Troja; Federica Fanzago<sup>5</sup>; Giulio Bianchini; Massimo Sgaravatto<sup>6</sup>; Teo Prica<sup>7</sup>; Tommaso Boccali; Tommaso Tedeschi<sup>1</sup>

<sup>1</sup> *Universita e INFN, Perugia (IT)*

<sup>2</sup> *INFN, Perugia (IT)*

<sup>3</sup> *Forschungszentrum Juelich GmbH (DE)*

<sup>4</sup> *Jozef Stefan Institute (SI)*

<sup>5</sup> *INFN Padova*

<sup>6</sup> *Universita e INFN, Padova (IT)*

<sup>7</sup> *IZUM*

**Corresponding Authors:** tommaso.tedeschi@cern.ch, diego.ciangottini@cern.ch, teo.prica@izum.si, andrej.filipcic@ijs.si, daniele.spiga@cern.ch, andrea.manzi@cern.ch, antonino.troja@pd.infn.it, tboccali@cern.ch, giulio.bianchini@studenti.unipg.it, a.memon@fz-juelich.de, federica.fanzago@pd.infn.it, massimo.sgaravatto@pd.infn.it

In a geo-distributed computing infrastructure with heterogeneous resources (HPC and HTC and possibly cloud), a key to unlock an efficient and user-friendly access to the resources is being able to offload each specific task to the best suited location. One of the most critical problems involve the logistics of wide-area with multi stage workflows back and forth multiple resource providers. We envision a model where such a challenge can be addressed enabling a “transparent offloading” of containerized payloads using the Kubernetes API primitives creating a common cloud-native interface to access any number of external hardware machines and type of backends. Thus we created the interLink project, an open source extension to the concept of Virtual-Kubelet with a design that aims for a common abstraction over heterogeneous and distributed backends. interLink is developed by INFN in the context of interTwin, an EU funded project that aims to build a digital-twin platform (Digital Twin Engine) for sciences, and the ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing in Italy. In this talk we first provide a comprehensive overview of the key features and the technical implementation. We showcase our major case studies such as the scale out of an analysis facility, and the distribution of ML training processes. We focus on the impacts of being able to seamlessly exploit world-class EuroHPC supercomputers with such a technology.

Poster session / 512

## Evolution of Regional, Age and Gende Demographics in the ATLAS Collaboration

**Authors:** Flavia De Almeida Dias<sup>1</sup>; Maria Teresa Dova<sup>2</sup>; Simon Connell<sup>3</sup>; Steven Goldfarb<sup>4</sup>

<sup>1</sup> *Nikhef National institute for subatomic physics (NL)*

<sup>2</sup> *National University of La Plata (AR)*

<sup>3</sup> *University of Johannesburg (ZA)*

<sup>4</sup> *University of Melbourne (AU)*

**Corresponding Authors:** simonhconnell@gmail.com, maria.teresa.dova@cern.ch, flavia.dias@cern.ch, steven.goldfarb@cern.ch

The ATLAS Collaboration consists of around 6000 members from over 100 different countries. Regional, age and gender demographics of the collaboration are presented, including the time evolution over the lifetime of the experiment. In particular, the relative fraction of women is discussed, including their share of contributions, recognition and positions of responsibility, including showing how these depend on other demographic measures.

Poster session / 513

## A Cloud oriented platform to exploit ANSYS application

**Authors:** Ahmad Alkhansa<sup>None</sup>; Alessandro Costantini<sup>None</sup>; Antonino La Mattina<sup>1</sup>; Barbara Martelli<sup>None</sup>; Cristina Curreli<sup>1</sup>; Francesco Sinisi<sup>None</sup>; Giusy Sergi<sup>None</sup>; Jacopo Gasparetto<sup>None</sup>; Letizia Magenta<sup>None</sup>

<sup>1</sup> *IRCCS Istituto Ortopedico Rizzoli*

**Corresponding Authors:** francesco.sinisi@cnaf.infn.it, ahmad.alkhansa@cnaf.infn.it, alessandro.costantini@cnaf.infn.it, letizia.magenta@cnaf.infn.it, jacopo.gasparetto@cnaf.infn.it, giusy.sergi@cnaf.infn.it, cristina.curreli@ior.it, barbara.martelli@cnaf.infn.it, antoninoamedeo.lamattina@ior.it

Cloud computing technologies are becoming increasingly important to provide a variety of services able to serve different communities' needs. This is the case of the DARE project (Digital Lifelong Prevention), a four-year initiative, co-financed by the Italian Ministry of University and Research as part of the National Plan of Complementary Investments to the PNRR. The project aims to develop prevention and digital health in Italy through the complete valorization of the health data chain, including data relevant to health. Within DARE, INFN the Italian National Institute for Nuclear Physics (INFN) is leading the technology scouting and integration.

In particular, we present the activities aimed at using ANSYS software, moving from a local batch solution to a cloud-enabled platform. As a result, we deployed a microservices-based environment using the solutions and services made available within INFN Cloud, the Cloud infrastructure of INFN, adopting consolidated technologies like Kubernetes and CEPH and integrating services like Nextflow to improve the interoperability of the presented solution.

In such respect, Kubernetes offers a dynamic and adaptable system for launching and overseeing containerized applications, making it an excellent option for handling intricate workloads within the Cloud environment. CEPH offers a distributed storage environment and the possibility to provide persistent storage for Kubernetes as well as Object storage that can be made available to services and users to store their outcome and analyze their data. Similarly, Nextflow, a tool for managing workflows, was selected for its ability to seamlessly incorporate different software packages and systems for environment management. Additionally, it streamlines the process of creating and implementing computational pipelines that handle large amounts of data.

The objective of the current activity is to identify benefits and explore potential enhancements through the use of a cloud-based approach. Additionally, due to handling patient information and being subject to GDPR regulations, the workflow must be carried out within a secure infrastructure. Having worked on projects involving personal data at CNAF, we have acquired experience in the security sector, as well as in hardening tools like RKE2+CIS, and infrastructures such as EPIC. Ultimately, our goal is to integrate these workflows into a microservices-based environment within a

secure Cloud infrastructure, with the intention of easily reproducing this process for future collaborations in the biomedical field.

**Parallel (Track 3) / 514**

## Efficient ML-Assisted Particle Track Reconstruction Designs

**Authors:** Antonio Ferrer Sánchez<sup>None</sup>; Nadezhda Dobрева<sup>None</sup>; Roberto Ruiz De Austri<sup>None</sup>; Sascha Caron<sup>1</sup>; Uraz Odyurt<sup>1</sup>; Yue Zhao<sup>None</sup>; Zef Wolffs<sup>None</sup>

<sup>1</sup> *Nikhef National institute for subatomic physics (NL)*

**Corresponding Authors:** antonio.ferrer-sanchez@uv.es, sascha.caron@cern.ch, zef.wolffs@cern.ch, yue.zhao@surf.nl, u.odyurt@nikhef.nl, rruiz@ific.uv.es, nadetodobreva@gmail.com

Track reconstruction, a.k.a., tracking, is a crucial part of High Energy Physics experiments. Traditional methods for the task, relying on Kalman Filters, scale poorly with detector occupancy. In the context of the upcoming High Luminosity-LHC, solutions based on Machine Learning (ML) and deep learning are very appealing. We investigate the feasibility of training multiple ML architectures to infer track-defining parameters from detector signals, for the application of offline reconstruction. We study and compare three Transformer model designs, as well as a U-Net model design. Firstly, we consider an autoregressive Transformer with the original encoder-decoder architecture, reconstructing a particle's trajectory given a few initial hits. Secondly, we employ an encoder-only model with the purpose of regressing track parameter values for each hit in an event, followed by a clustering step. Next, an encoder-only model design as a classifier is considered, producing class labels for each hit in an event, given pre-defined bins within the track parameter-space. Lastly, similar to the third Transformer design, a U-Net model for pixel classification into pre-defined classes is evaluated.

The models are benchmarked for physics performance and inference speed on methodically simplified datasets, generated by the recently developed simulation framework, REDuced VIRTUAL Detector (REDVID). Our second batch of simplified datasets are derived from the TrackML dataset. Our preliminary results show promise for the application of such deep learning techniques on more realistic data for tracking, as well as efficient elimination of solutions.

**Parallel (Track 2) / 515**

## An online GPU hit finder for the STS detector in the CBM experiment

**Author:** Felix Weiglhofer<sup>1</sup>

**Co-author:** Volker Lindenstruth<sup>1</sup>

<sup>1</sup> *Goethe University Frankfurt (DE)*

**Corresponding Authors:** felix.weiglhofer@cern.ch, voli@compeng.de

The CBM experiment is expected to run with a data rate exceeding 500 GB/s even after averaging. At this rate storing raw detector data is not feasible and an efficient online reconstruction is instead required. GPUs have become essential for HPC workloads. The higher memory bandwidth and parallelism of GPUs can provide significant speedups over traditional CPU applications. These properties also make them a promising target for the planned online processing in CBM.

We present an online hit finder for the STS detector capable of running on GPUs. It consists of four steps using STS digis (timestamped detector messages) as input. Digis first are sorted by sensor

and then for each sensor, they are sorted by channel and their timestamp. Neighboring digis are combined into clusters. Finally, after time sorting clusters on each sensor are combined into hits. Each of those steps is trivially parallel across STS sensors or even sensor sides. To fully utilize GPU hardware, we modify the algorithms to be parallel on digi or cluster level. This includes a custom implementation of parallel merge sort allowing full parallelism within GPU blocks.

Our implementation achieves a speedup of 24 on mCBM data compared to the same code on a single CPU core. The exact achieved throughput will be shown and discussed during the presentation.

This work is supported by BMBF (05P21RFFC1).

Parallel (Track 3) / 516

## Energy-efficient graph-based algorithm for tracking at the HL-LHC

**Author:** Heberth Torres<sup>1</sup>

<sup>1</sup> *L2I Toulouse, CNRS/IN2P3, UT3*

**Corresponding Author:** h.torres@cern.ch

Graph neural networks represent a potential solution for the computing challenge posed by the reconstruction of tracks at the High Luminosity LHC [1, 2, 3]. The graph concept is convenient to organize the data and to split up the tracking task itself into the subtasks of identifying the correct hypothetical connections (edges) between the hits, subtasks that are easy to parallelize and process efficiently, for example using GPUs.

We will describe an algorithm that benefits from the graph advantages, but instead of using neural networks, it consists of direct geometric comparisons of neighboring edge pairs, testing the hypothesis of both edges corresponding to the same particle, to build up hit triplets and track candidates. The compatibility of edge pairs is tested based on two observables: the  $\eta$  direction, and an estimator of the transverse momentum of the particle hypothetically associated with each edge. Before this step, the tracking algorithm includes graph construction with a modified version of the Module Map method described at 2. In this Module Map version, the hits are organized in a two dimension map of the modules in the longitudinal detector plane ( $r, z$ ), and next the edges are built based on a list of possible connections, in combination with a  $\Delta\phi$  cut. At each module, the hits are ordered based on their  $\phi$  position, which serves to significantly reduce the combinatorics when applying that  $\Delta\phi$  cut.

We will present the track reconstruction performance of this algorithm, estimated using the Open Data Detector [4], as well as its computing efficiency. For this tracking chain executed in a single CPU core, the time required to process an HL-LHC  $t\bar{t}$  event with 200  $pp$  interaction pileup per bunch crossing is of the order of a second, which makes it a rather energy efficient tracking algorithm. This is the processing time from the collection of reconstructed hits to the track candidates, targeting primary particles with transverse momentum above 1 GeV. The algorithm is highly parallelizable; executed on a GPU, the processing time is expected to decrease at least by one order of magnitude.

1 S. Farrell et al., Novel deep learning methods for track reconstruction, 2018, arXiv: 1810.06111 [hep-ex], url: <https://arxiv.org/abs/1810.06111>.

2 C. Biscarat et al., Towards a realistic track reconstruction algorithm based on graph neural networks for the HL-LHC, EPJ Web Conf. 251 (2021) 03047, url: <https://doi.org/10.1051/epjconf/202125103047>.

3 H. Torres et al. on behalf of the ATLAS Collaboration, Physics Performance of the ATLAS GNN4ITk Track Reconstruction Chain, Proceedings of the CTD 2023, ATL-SOFT-PROC-2023-047, url: <http://cds.cern.ch/record/2882507>.  
[4] <https://gitlab.cern.ch/acts/OpenDataDetector>.

**Parallel (Track 5) / 517**

## Event generation with quantum computers through particle-oriented simulation

**Author:** Yutaro Iiyama<sup>1</sup><sup>1</sup> *University of Tokyo (JP)***Corresponding Author:** yutaro.iiyama@cern.ch

Quantum computers may revolutionize event generation for collider physics by allowing calculation of scattering amplitudes from full quantum simulation of field theories. Although rapid progress is being made in understanding how best to encode quantum fields onto the states of quantum registers, most formulations are lattice-based and would require an impractically large number of qubits when applied to scattering events at colliders with a wide momentum dynamic range. In this regard, the single-particle digitization approach of Barata et al. (Phys. Rev. A 103) is highly attractive for its qubit efficiency and strong association with scattering. Since the original work established the digitization scheme on the scalar  $\phi^4$  theory, we explore its extensions to fermion fields and other types of interactions. We then implement small-scale scattering simulations on both real quantum computers and a statevector calculator run on HPCs. A possible roadmap toward realizing the ultimate goal of performing collider event generation from quantum computers will be discussed.

**Poster session / 518**

## Hadronization effects using String and Cluster Models in Herwig 7

**Authors:** Andrzej Konrad Siodmok<sup>1</sup>; Michaela Divisova<sup>None</sup>; Miroslav Myska<sup>2</sup>; Pratixan Sarmah<sup>None</sup><sup>1</sup> *Jagiellonian University (PL)*<sup>2</sup> *Czech Technical University in Prague (CZ)***Corresponding Authors:** miroslav.myska@cern.ch, divismi5@fjfi.cvut.cz, a.siodmok@cern.ch, pratixan.sarmah@doctoral.uj.edu.pl

Monte Carlo Event Generators contain several free parameters that cannot be inferred from first principles and need to be tuned to better model the data. With increasing precision of perturbative calculations to higher orders and hence decreasing theoretical uncertainties, it becomes crucial to study the systematics of non-perturbative phenomenological models. A recent attempt was made at tuning the combination of the angular-ordered parton shower in Herwig7 and Lund string hadronization model to LEP data with a new approach called Autotunes 1. However, the results showed worse performance to important observables like LEP event shapes when compared to the previous tunes of the cluster hadronization model. Since the angular ordered parton shower and the string hadronization model perform well independently with Herwig7 and Pythia8, we would naively expect them to perform better together and thus investigate further by tuning the setup with the Professor approach adopted in 2. I present the results of our tune and compare it with the Herwig7 default, Pythia 8 and Autotunes tunes.

1 High dimensional parameter tuning for event generators, J.Belm, L.Gellersen, Eur.Phys.J.C 80 (2020) 1, 54

2 Systematic event generator tuning for the LHC, A.Buckley et al, Eur.Phys.J.C 65 (2010) 331-357

**Poster session / 519**

## GlideinBenchmark: collecting resource information to optimize provisioning

**Authors:** Marco Mambelli<sup>1</sup>; Shrijan Swaminathan<sup>2</sup>

<sup>1</sup> *Fermilab (US)*

<sup>2</sup> *Purdue University*

**Corresponding Authors:** swamina7@purdue.edu, marcom@fnal.gov

Choosing the right resource can speedup jobs completion, better utilize the available hardware and visibly reduce costs, especially when renting computers on the cloud. This was demonstrated in earlier studies on HEPCloud. But the benchmarking of the resources proved to be a laborious and time-consuming process. This paper presents GlideinBenchmark, a new Web application leveraging the pilot infrastructure of GlideinWMS to benchmark resources, and shows how to use the data collected and published by GlideinBenchmark to automate the optimal selection of resources.

An experiment can select the benchmark or the set of benchmarks that most closely evaluate the performance of its workflows. With GlideinBenchmark and the help of the GlideinWMS Factory it controls the benchmark execution. Finally, a scheduler like HEPCloud's Decision Engine can use the results to optimize resource provisioning.

**Parallel (Track 8) / 520**

## Open Science and Compressed Baryonic Matter experiment

**Author:** Eoin Clerkin<sup>1</sup>

<sup>1</sup> *FAIR - Facility for Antiproton and Ion Research in Europe, Darmstadt*

**Corresponding Author:** e.clerkin@gsi.de

In recent years, there has been significant political and administrative interest in “Open Science”, which on one hand has led to additional obligations but also to significant financial backing. For institutes and scientific collaborations, the funding opportunities may have brought some focus on these topics, but there is also a significant hope, though engagement in open science infrastructure and culture, a possible multiplying effect on scientific output through the sharing of knowledge among and between scientists and citizens.

The Facility for AntiProton and Ion Research in Europe (FAIR) is a particle accelerator just outside Darmstadt in Germany, which is under final construction at a site adjacent to the GSI Helmholtz Centre for Heavy Ion Research. One of its five scientific pillars is the Compressed Baryonic Matter (CBM) experiment, which is now prioritised and expected to receive its first beam in 2028. For CBM, as a leading international scientific collaboration, an active open science policy is an imperative.

In this contribution, we outline our fully-formed policy towards “Open Software” and describe how we overcame difficulties to facilitate a F.A.I.R.-level of openness. We discuss the internally controversial issue of “Open Data” and the availability to technically test data policies at the prototype experiment mini-CBM, before application to our more important physics-rich data coming from our future world-class experiment. Lastly we discuss what it means to be an “Open Collaboration” and how engagement in open science strategy within the collaboration could facilitate a plethora of new citizen science projects and help progress our research and the open science agenda.

**Poster session / 521**

## Parallel Photon Simulation for IceCube In C++

**Author:** Kevin Meagher<sup>None</sup>

**Corresponding Author:** kjmeagher@wisc.edu

The IceCube Neutrino Observatory instruments one cubic kilometer of glacial ice at the geographic South Pole. Cherenkov light emitted by charged particles is detected by 5160 photomultiplier tubes embedded in the ice. Deep antarctic ice is extremely transparent, resulting in absorption lengths exceeding 100m. However, yearly variations in snow deposition rates on the glacier over the last 100 thousand years have created roughly horizontal layers which vary significantly in scattering and absorption coefficients. These variations must be taken into account when simulating IceCube events. In addition, anisotropies in photon propagation have been observed and recently described by deflection by birefringent polycrystals. Modeling of ice properties remains one of the largest sources of systematic uncertainties in IceCube analyses, requiring intensive studies of the ice. Despite the fact that photon tracking is highly parallelizable and is an ideal case for GPUs, the limiting constraint for these studies is time spent simulating photon propagation. In order to efficiently and accurately perform these simulations, custom software has been developed and optimized for our specific use case. IceCube's current production simulation code CLSim is based on OpenCL and is tightly coupled to the IceCube's simulation stack and is in need of modernization. This talk will discuss the current requirements for Photon tracking code in IceCube and the effort to transition the code to new C++ frameworks which uses `std::par`.

**Poster session / 522**

## Geant4 models for nuclear de-excitation

**Author:** Nikita Chalyi<sup>1</sup>

**Co-author:** Vladimir Ivantchenko<sup>2</sup>

<sup>1</sup> *Tomsk State University (RU)*

<sup>2</sup> *CERN*

**Corresponding Authors:** vladimir.ivantchenko@cern.ch, ch.online.edu@gmail.com

Geant4 hadronic physics sub-library includes a wide variety of models for high and low-energy hadronic interactions. We report on recent progress in development of the Geant4 nuclear de-excitation module. This module is used by many Geant4 models for sampling of de-excitation of nuclear recoil produced in nuclear reactions. Hadronic shower shape and energy deposition are sensitive to these processes. We will present comparisons of Geant4 predictions for the thin target experiments and will discuss CPU efficiency of Geant4 de-excitation module.

**Poster session / 523**

## Zero Degree Calorimeter fast simulation with normalizing flows

**Author:** Emilia Majerz<sup>1</sup>

**Co-author:** Witold Dzwiniel<sup>2</sup>

<sup>1</sup> *AGH University of Krakow (PL)*

<sup>2</sup> *AGH University of Science and Technology (PL)*

**Corresponding Authors:** witold.dzwiniel@cern.ch, emilia.maria.majerz@cern.ch

Simulating the Large Hadron Collider detectors, particularly the Zero Degree Calorimeter (ZDC) of the ALICE experiment, is computationally expensive. This process uses the Monte Carlo approach, which demands significant computational resources, and involves many steps. However, recent



advances in generative deep learning architectures present promising methods for speeding up these simulations.

In this work, we apply normalizing flows to the simulation of ZDC neutron detector responses, thus obtaining high-fidelity surrogates of numerical models, and achieving competitive results on the GEANT4 dataset. We also provide and compare post-processing techniques for enhancing the results. Moreover, we check if the reasoning of the networks is physically relevant by employing state-of-the-art explainability techniques. This we see as a vital step in deciding whether our model is ready to replace the current simulation engine.

**Parallel (Track 3) / 524**

## Status of DUNE Offline Computing

**Author:** Michael Hudson Kirby<sup>1</sup>

<sup>1</sup> *Brookhaven National Laboratory (US)*

**Corresponding Author:** kirby@cern.ch

We summarize the status of the Deep Underground Neutrino Experiment (DUNE) software and computing development. The DUNE Collaboration has been successfully operating the DUNE prototype detectors at both Fermilab and CERN, and testing offline computing services, software, and infrastructure using the data collected. We give an overview of results from end-to-end testing of systems needed to acquire, catalog, reconstruct, simulate and analyze the beam data from ProtoDUNE Horizontal Drift (PDHD) and Near Detector 2x2 Demonstrator, and cosmic data from ProtoDUNE Vertical Drift (PDVD). These tests included reconstruction and simulation of data from all prototype detector runs utilizing a variety of distributed computing and HPC resources. The results of these studies help define the development path of DUNE core software and computing to support the physics goals of precision measurements of neutrino oscillation parameters, detection of astrophysical neutrinos, measurement of neutrino interaction properties and searches for physics beyond the Standard Model. The data from the full DUNE far and near detectors, expected in 2029 and 2031 respectively, will present significant challenges in terms of data product memory management, optimized use of parallel processing for reconstruction and simulation, and management of large individual trigger data volumes. DUNE will present plans for future development to accommodate the requirements of the larger DUNE far and near detectors, and the timeline for future data challenges leading to data taking at the end of the decade.

**Poster session / 525**

## Numerical studies of space charge effect on particle tracking in a small TPC

**Author:** Pralay Kumar das<sup>1</sup>

**Co-authors:** Jaydeep Datta<sup>2</sup>; Nayana Majumdar<sup>3</sup>; Subhendu Das<sup>4</sup>; Supratik Mukhopadhyay

<sup>1</sup> *Saha Institute Of Nuclear Physics*

<sup>2</sup> *Stony Brook University, CFNS*

<sup>3</sup> *Saha Institute of Nuclear Physics*

<sup>4</sup> *Saha Institute of Nuclear Physics (SINP), India*

**Corresponding Authors:** jaydeep.datta@stonybrook.edu, subhendudas456038@gmail.com, nayana.majumdar@saha.ac.in, supratik.mukhopadhyay@saha.ac.in, daspralay11@gmail.com

In the realm of low-energy nuclear physics experiments, the Active Target Time Projection Chamber (AT-TPC) can be advantageous for studying nuclear reaction kinematics, such as the alpha cluster

decay of  $^{12}\text{C}$ , by tracking the reaction products produced in the active gas medium of the TPC. The tracking capability of the TPC is strongly influenced by the homogeneity of the electric field applied across its drift medium, which is affected by the space charge produced by low-energy projectiles and reaction products in the active gas medium through the ionization process. In this work, we implemented a mathematical model based on a hydrodynamic approach to simulate the space charge effect caused by the alpha beam on the TPC performance using the platform of a commercial Finite Element Method (FEM) package available in COMSOL Multiphysics. This novel approach is computationally less expensive compared to the particle model. The primary ionization caused by the alpha particles was simulated using Geant4, and the electron transport parameters for the active gas were obtained from MAGBOLTZ. The effect of space charge on the applied electric field and the angular resolution of the TPC with beam currents ranging from a few pA to 20 pA have been reported. The same model was also utilized to simulate the temporal evolution of an alpha track in two different active gas mediums,  $\text{He} + \text{CO}_2$  and  $\text{He} + \text{C}_4\text{H}_{10}$ , in the volumetric ratios of 90:10 and 93:7, respectively. Different readout geometries of the TPC were studied to find the optimum strip width and number of strips at the TPC end cap to properly resolve the alpha particle tracks. A tracking algorithm had been developed to distinguish between the multiple tracks between scattered events and the  $^{12}\text{C}$  breakup. We are designing a 64 channel Micromegas based prototype TPC on the basis of the simulation results.

Poster session / 526

## Towards more efficient job scheduling in ALICE: predicting job execution time using machine learning

**Authors:** Bartosz Balis<sup>1</sup>; Costin Grigoras<sup>2</sup>; Marcin Kurdziel<sup>3</sup>; Michał Faciszewski<sup>3</sup>; Mikołaj Zasada<sup>3</sup>; Sara Świętek<sup>3</sup>; Tomasz Marcin Lelek<sup>1</sup>; on behalf of the ALICE Collaboration<sup>None</sup>

<sup>1</sup> AGH University of Krakow (PL)

<sup>2</sup> CERN

<sup>3</sup> AGH

**Corresponding Authors:** tomekl007@gmail.com, costin.grigoras@cern.ch, bartosz.balis@cern.ch

The ALICE Grid processes up to one million computational jobs daily, leveraging approximately 200,000 CPU cores distributed across about 60 computing centers. Enhancing the prediction accuracy for job execution times could significantly optimize job scheduling, leading to better resource allocation and increased throughput of job execution. We present results of applying machine learning techniques to predicting the execution time of ALICE computational jobs. To this end, we focus on the following main challenges in this prediction task:

- (1) Feature extraction and selection: extracting the relevant features from the collected data and selecting the ones that are most important for model training and inference
- (2) Model selection: identifying an ML model that is accurate and robust for our prediction problem.
- (3) Model decay: making sure that the model accuracy does not deteriorate in time as new data arrives, possibly from an evolving data distribution.
- (4) Near-real-time processing: predictions need to be made in near-real-time.

Our goal is to develop a solution capable of predicting job execution times for batches of hundreds of elements in less than 100 milliseconds, without compromising accuracy or hindering continuous learning. This requires striking a delicate balance between computational complexity and real-time performance. By addressing these challenges within the ALICE CERN experiment framework, we can enhance job scheduling efficiency and optimize resource allocation, ultimately advancing scientific research in particle physics.

**Acknowledgements.** This work is co-financed in part supported by the Ministry of Science and Higher Education (Agreement Nr 2023/WK/07) and by the program of the Ministry of Science and Higher Education entitled PMW

Parallel (Track6) / 527

## A data Quality-Assurance framework for online and offline applications for the CBM experiment

Author: Sergei Zharko<sup>1</sup>

<sup>1</sup> GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany

Corresponding Author: s.zharko@gsi.de

A data quality assurance (QA) framework is being developed for the CBM experiment. It provides flexible tools for monitoring of reference quantity distributions for different detector subsystems and data reconstruction algorithms. This helps to identify software malfunctions and calibration status, to prepare a setup for the data taking and to prepare data for the production. A modular structure of the QA framework allows to keep independent QA units for different steps of the data reconstruction.

Since the offline and the online scenarios of data reconstruction need to meet different requirements, the QA framework is implemented differently for those two regimes. In the offline scenario, the data QA software is based on the FairRoot framework and is used to track the effects on data in the continuous development of the reconstruction algorithms as well as to check the data quality on the production stage. The QA software for the online reconstruction scenario utilizes the standard and boost C++ libraries and provides a real-time monitoring of detector and algorithm performance. This was successfully applied to the data taking at the mini-CBM experiment in May 2024.

Parallel (Track 8) / 528

## A technical overview of industry-science R&D projects for the High Luminosity LHC under CERN openlab

Author: Thomas Owen James<sup>1</sup>

Co-authors: Antonio Nappi<sup>1</sup>; Luca Atzori<sup>1</sup>; Luca Mascetti<sup>1</sup>; Maria Girone<sup>1</sup>

<sup>1</sup> CERN

Corresponding Authors: maria.girone@cern.ch, luca.mascetti@cern.ch, antonio.nappi@cern.ch, thomas.owen.james@cern.ch, luca.atzori@cern.ch

CERN openlab is a unique resource within CERN that works to establish strategic collaborations with industry, fuel technological innovation and expose novel technologies to the scientific community.

ICT innovation is needed to deal with the unprecedented levels of data volume and complexity generated by the High Luminosity LHC. The current CERN openlab Phase VIII is designed to tackle these challenges on a number of fronts, including, but not limited to: heterogeneous computing, platforms, and infrastructures; novel storage, compression, and data management solutions; emerging low-latency interconnect and link protocols; and the exploitation of artificial intelligence algorithms across a multitude of domains, including edge devices for real-time event selection and triggering. The evaluation and adoption of these technologies are being accelerated by ongoing collaborations between industrial leaders in the relevant fields and the scientific community at CERN. The work of ongoing focussed projects in these areas will be summarised, and results demonstrating their impact will be shown. Incubator projects on emerging technologies such as digital twins and generative AI will be presented, as well as the next steps in these R&D efforts.

Poster session / 529

## Real-time Level-1 Trigger Data Scouting at CMS using CXL Memory Lake

**Author:** Giovanna Lazzari Miotto<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Author:** giovanna.lazzari.miotto@cern.ch

Level-1 Data Scouting (L1DS) is a novel data acquisition subsystem at the CMS Level-1 Trigger (L1T) that exposes the L1T event selection data primitives for online processing at the LHC's 40 MHz bunch-crossing rate, enabling unbiased and unconventional analyses. An L1DS demonstrator has been operating since Run 3, relying on a ramdisk for ephemeral storage of incoming and intermediate data, accessible by the system's units through NFS. With the HL-LHC and CMS'Phase 2 upgrade projected to enhance trigger resolutions, a high-performance shared memory system is key to retain real-time processing capabilities in Run 4. For this, we leverage the emerging Compute Express Link (CXL) open standard, which provides uniform, cache-coherent memory access from heterogeneous processing units, targeting a streamlined pipeline with minimized data movement over a memory lake shared among CPUs and GPUs. In this contribution, we present the integration of CXL-compliant shared memory into the L1DS demonstrator at CMS, including an overview of our approach's design, benefits, and limitations. Furthermore, we evaluate CXL-based L1DS performance through analyses in heterogeneous contexts, supporting a discussion of the memory lake model and its use cases for the CMS community.

**Poster session / 530**

## Harnessing the power of neural networks to identify and locate primary vertices in proton-proton collision data

**Authors:** Lauren Alexandra Tompkins<sup>1</sup>; Rocky Bala Garg<sup>1</sup>

<sup>1</sup> *Stanford University (US)*

**Corresponding Authors:** rocky.bala.garg@cern.ch, laurenat@stanford.edu

Extensive research has been conducted on deep neural networks (DNNs) for the identification and localization of primary vertices (PVs) in proton-proton collision data from ATLAS/ACTS. Previous studies focused on locating primary vertices in simulated ATLAS data using a hybrid methodology. This approach began with the derivation of kernel density estimators (KDEs) from the ensemble of charged track parameters, employing an analytical probability density estimation technique. These KDEs were subsequently utilized as input for two neural network (NN) architectures, namely UNet and UNet++, alongside the truth PV positions extracted in the form of target histograms from the simulated data. Through these investigations, a proof-of-concept was demonstrated, achieving performance comparable to the ATLAS Adaptive Multi-Vertex Finder (AMVF) algorithm, while also enhancing the vertex position resolution.

The current studies transition from analytical KDE computation to a fully NN-based implementation, presenting an end-to-end primary vertex finder algorithm driven by neural networks. A comprehensive analysis of this approach, including a comparative assessment of its performance against the AMVF algorithm, will be presented.

**Parallel (Track6) / 531**

## Exploring new directions in enhancing the ACTS parameter optimization suite

**Authors:** Lauren Alexandra Tompkins<sup>1</sup>; Rocky Bala Garg<sup>1</sup>

<sup>1</sup> *Stanford University (US)*

**Corresponding Authors:** rocky.bala.garg@cern.ch, laurenat@stanford.edu

Particle tracking is among the most sophisticated and complex parts of the full event reconstruction chain. Various reconstruction algorithms work in sequence to build trajectories from detector hits. Each of these algorithms requires numerous configuration parameters that need fine-tuning to properly account for the detector/experimental setup, the available CPU budget, and the desired physics performance. To automate and optimize the tuning of these parameters, automatic parameter optimization techniques were implemented in “A Common Tracking Software”(ACTS) framework, the open-source track reconstruction software framework. These techniques allow users to flexibly choose tunable parameters and define a cost/benefit function for optimizing the full reconstruction chain. Since their implementation, these techniques have been greatly beneficial for researchers across various experiments.

The current study discusses recent developments in these optimization techniques. It explores novel optimization methods that provide a more gradual and systematic approach to parameter tuning. Specifically, we have investigated the potential application of surrogate methods for tuning tracking algorithm parameters and discussed the improvements achieved over the original techniques.

**Poster session / 532**

## **Addressing tokens dynamic generation, propagation, storage and renewal to secure the GlideinWMS pilot based jobs and system.**

**Authors:** Bruno Moreira Coimbra<sup>1</sup>; Marco Mambelli<sup>2</sup>

<sup>1</sup> *Fermi National Accelerator Lab. (US)*

<sup>2</sup> *Fermilab (US)*

**Corresponding Authors:** marcom@fnal.gov, bruno.moreira.coimbra@cern.ch

GlideinWMS has been one of the first middleware in the WLCG community to transition from X.509 to support also tokens. The first step was to get from the prototype in 2019 to using tokens in production in 2022. This paper will present the challenges introduced by the wider adoption of tokens and the evolution plans for securing the pilot infrastructure of GlideinWMS and supporting the new requirements.

In the last couple of years, the GlideinWMS team supported the migration to tokens of experiments and resources. Inadequate support in the current infrastructure, more stringent requirements, and the higher spatial and temporal granularity forced GlideinWMS to revisit once more how credentials are generated, used, and propagated.

The new credential modules have been designed to be used in multiple systems (GWMS, HC) and use a model where credentials have type, purpose, and different flows.

Credentials are dynamically generated in order to customize the duration and limit the scope to the targeted resource. This allows to enforce the least privilege principle. Finally, we also considered adding credential storage, renewal, and invalidation mechanisms within the GlideinWMS infrastructure to serve better the experiments' needs.

**Parallel (Track 2) / 533**

## **Real-time pattern recognition with FPGA at LHCb, an O(n) complexity architecture**

**Authors:** Andrea Contu<sup>1</sup>; Ao Xu<sup>2</sup>; Arantza De Oyanguren Campos<sup>3</sup>; Brij Kishor Jashal<sup>4</sup>; Federico Lazzari<sup>5</sup>; Francesco Terzuoli<sup>6</sup>; Giovanni Punzi<sup>2</sup>; Giulia Tuci<sup>7</sup>; Jiahui Zhuo<sup>3</sup>; Jibo He<sup>8</sup>; Lorenzo Pica<sup>9</sup>; Maurizio Martinelli<sup>10</sup>; Michael J. Morello<sup>11</sup>; Qi Shi<sup>8</sup>; Riccardo Fantechi<sup>2</sup>

<sup>1</sup> INFN

<sup>2</sup> Universita & INFN Pisa (IT)

<sup>3</sup> Univ. of Valencia and CSIC (ES)

<sup>4</sup> RAL, TIFR and IFIC

<sup>5</sup> Universita di Pisa & INFN Pisa (IT)

<sup>6</sup> Università di Siena & INFN Pisa (IT)

<sup>7</sup> Heidelberg University (DE)

<sup>8</sup> University of Chinese Academy of Sciences (CN)

<sup>9</sup> SNS & INFN Pisa (IT)

<sup>10</sup> Universita & INFN, Milano-Bicocca (IT)

<sup>11</sup> SNS and INFN-Pisa (IT)

**Corresponding Authors:** riccardo.fantechi@cern.ch, michael.joseph.morello@cern.ch, ao.xu@cern.ch, maurizio.martinelli@cern.ch, jibo.he@cern.ch, giovanni.punzi@cern.ch, lorenzo.pica@cern.ch, andrea.contu@cern.ch, qi.shi@cern.ch, federico.lazzari@cern.ch, francesco.terzuoli@cern.ch, giulia.tuci@cern.ch, jiahui.zhuo@cern.ch, arantza.de.oyanguren.campos@cern.ch, brij.kishor.jashal@cern.ch

**Abstract:** The LHCb collaboration is planning an upgrade (LHCb “Upgrade-II”) to collect data at an increased instantaneous luminosity (a factor of 7.5 larger than the current one). LHCb relies on a complete real-time reconstruction of all collision events at LHC-Point 8, which will have to cope with both the luminosity increase and the introduction of correspondingly more granular and complex detectors.

After an intensive R&D programme, LHCb proposed to build an FPGA-based system to reconstruct tracks in the SciFi detector during Run 4, as an intermediate step towards a system that could be extended to other tracking detectors. Based on an extremely parallel architecture, the so-called ‘artificial retina’, this system has an  $O(n)$  complexity, which is a crucial feature for high luminosity scenarios.

In this talk we describe why this system scales linearly with luminosity and how much it can accelerate the LHCb High Level Trigger in Run 4.

Poster session / 534

## Real-time monitoring of LHCb interaction region with a fast trackless methodology

**Authors:** Daniele Passaro<sup>1</sup>; Elena Graverini<sup>2</sup>; Federico Lazzari<sup>3</sup>; Francesco Terzuoli<sup>4</sup>; Giovanni Bassi<sup>1</sup>; Giovanni Punzi<sup>5</sup>; Giulio Cordova<sup>5</sup>; Michael J. Morello<sup>6</sup>

<sup>1</sup> SNS & INFN Pisa (IT)

<sup>2</sup> EPFL - Ecole Polytechnique Federale Lausanne (CH)

<sup>3</sup> Universita di Pisa & INFN Pisa (IT)

<sup>4</sup> Università di Siena & INFN Pisa (IT)

<sup>5</sup> Universita & INFN Pisa (IT)

<sup>6</sup> SNS and INFN-Pisa (IT)

**Corresponding Authors:** michael.joseph.morello@cern.ch, giovanni.bassi@cern.ch, elena.graverini@cern.ch, giulio.cordova@cern.ch, federico.lazzari@cern.ch, francesco.terzuoli@cern.ch, daniele.passaro@cern.ch, giovanni.punzi@cern.ch

The increasing computing power and bandwidth of FPGAs opens new possibilities in the field of real-time processing of HEP data. LHCb now uses a cluster-finder FPGA architecture to reconstruct hits in the VELO pixel detector on-the-fly during readout. In addition to its usefulness in accelerating HLT1 reconstruction by providing it with pre-reconstructed data, this system enables further opportunities. Thanks to the triggerless readout architecture of LHCb, these reconstructed hit positions

are available for every collision, amounting to a flow of  $10^{11}$  hits per second, that can be exploited further.

We have implemented a set of programmable counters, counting the hit rate at many locations in the detector volume simultaneously on the fly. We use this data to measure, and continuously track, both the motion of the interaction region and the relative position of the detector elements, with precisions of  $O(\mu m)$  and time granularity  $O(ms)$ . To this purpose we use a linearized calculation, based on a principal component analysis (PCA) of these low-level counters, that can be executed online with minimal computational effort. This methodology, being based on just the raw hit positions, differs in a substantial way from methodologies commonly in use, that rely on reconstructing of particle trajectories and their origin vertex, and are therefore computationally much more complex and more prone to biases due to internal detector misalignments. We report results obtained with real data in the 2024 run of LHCb.

Poster session / 537

## Refining FastSim with Machine Learning

Author: CMS Collaboration<sup>None</sup>

Corresponding Author: borislav.pavlov@cern.ch

A growing reliance on the fast Monte Carlo (FastSim) will accompany the high luminosity and detector granularity expected in Phase 2. FastSim is roughly 10 times faster than equivalent GEANT4-based full simulation (FullSim). However, reduced accuracy of the FastSim affects some analysis variables and collections. To improve its accuracy, FastSim is refined using regression-based neural networks trained with ML. The status of FastSim refinement is presented. The results show improved agreement with the FullSim output and an improvement in correlations among output observables and external parameters.

Parallel (Track 7) / 538

## Simulating the Carbon Cost of Grid Sites

Authors: Albert Gyorgy Borbely<sup>1</sup>; David Britton<sup>1</sup>; Dwayne Spiteri<sup>2</sup>; Emanuele Simili<sup>1</sup>; Gordon Stewart<sup>None</sup>; Samuel Cadellin Skipsey<sup>None</sup>; Steve Lloyd<sup>3</sup>

<sup>1</sup> University of Glasgow (GB)

<sup>2</sup> University of Glasgow

<sup>3</sup> University of London (GB)

Corresponding Authors: emanuele.simili@glasgow.ac.uk, s.l.lloyd@qmul.ac.uk, dwayne.patrick.spiteri@cern.ch, albert.borbely@cern.ch, gordon.stewart@glasgow.ac.uk, david.britton@cern.ch, samuel.cadellin.skipsey@cern.ch

We present first results from a new simulation of the WLCG Glasgow Tier-2 site, designed to investigate the potential for reducing our carbon footprint by reducing the CPU clock frequency across the site in response to a higher-than-normal fossil-fuel component in the local power supply. The simulation uses real (but historical) data for the UK power-mix, together with measurements of power consumption made at Glasgow on a variety of machines, and is designed to provide a tool to inform future procurements and the operation of sites. The output of the simulation, combined with considerations of embedded carbon, can also be used to inform and optimise the policy for replacing older hardware with more energy efficient devices. The rate of transition to more energy efficient hardware must be balanced against the embedded carbon in the manufacture of new machines, and frequency modulation must be balanced against both the loss of site throughput and the accounting of embedded carbon. Frequency modulation can also be used to reduce power requirements to address short-term supply issues, irrespective of the carbon content.

**Poster session / 539**

## Quantifying the computational speedup with MG4GPU for CMS workflows

**Authors:** CMS Collaboration<sup>None</sup>; Jin Choi<sup>1</sup><sup>1</sup> *Seoul National University (KR)***Corresponding Author:** borislav.pavlov@cern.ch

We will present the first analysis of the computational speedup achieved through the use of the GPU version of Madgraph, known as MG4GPU. Madgraph is the most widely used event generator in CMS. Our work is the first step toward benchmarking the improvement obtained through the use of its GPU implementation. In this presentation, we will show the timing improvement achieved without affecting physics performance, for a wide range of physics processes that are of general interest in CMS, quantified both by gridpack-generation and event-generation. Preliminary results demonstrate a speedup of a factor of three in matrix element calculation and a factor of 2.5 in full gridpack production for one of the most computationally intensive processes: Drell-Yan with four additional emissions. The workflows have been tested with diverse computational resources, including CUDA-enabled NVIDIA GPUs and modern vectorized CPUs from Intel and AMD, accessible via CERN resources and HPCs.

**Poster session / 540**

## Bamboo: A high-level HEP analysis library for ROOT::RDataFrame

**Authors:** Jindrich Lidrych<sup>1</sup>; Oğuz Güzel<sup>2</sup><sup>1</sup> *Universite Catholique de Louvain (UCL) (BE)*<sup>2</sup> *Université Catholique de Louvain (BE)***Corresponding Authors:** jindrich.lidrych@cern.ch, oguz.guzel@cern.ch

This poster presents an overview and features of a bamboo framework designed for HEP data analysis. The bamboo framework defines a domain-specific language, embedded in python, that allows to concisely express the analysis logic in a functional style. The implementation based on ROOT's RDataFrame and cling C++ JIT compiler approaches the performance of dedicated native code. Bamboo is currently being used for several CMS Run 2 and Run 3 analyses that rely on the NanoAOD data format for which many reusable components are included.

**Parallel (Track 8) / 542**

## Using containers to speed up development, to run integration tests and to teach about distributed systems

**Authors:** Bruno Moreira Coimbra<sup>1</sup>; Marco Mambelli<sup>2</sup><sup>1</sup> *Fermi National Accelerator Lab. (US)*<sup>2</sup> *Fermilab (US)***Corresponding Authors:** bruno.moreira.coimbra@cern.ch, marcom@fnal.gov



GlideinWMS is a workload manager provisioning resources for many experiments including CMS and DUNE. The software is distributed both as native packages and specialized production containers. Following an approach used in other communities like web development

we built our workspaces, system-like containers to ease development and testing.

Developers can change the source tree or check out a different branch and quickly reconfigure the services to see the effect of their changes.

In this paper, we'll talk about what differentiates workspaces from other containers.

We'll describe our base system composed of three containers. A one-node cluster including a compute element and a batch system. A GlideinWMS Factory controlling pilot jobs. And a scheduler and Frontend, to submit jobs and provision resources. Additional containers can be used for optional components. This system can easily run on a laptop and we'll share our evaluation of different container runtimes, with an eye for ease of use and performance.

Finally, we'll talk about our experience as developers and with students.

The GlideinWMS workspaces are easily integrated with IDEs like VS Code, simplifying debugging and allowing development and testing of the system also when offline.

They simplified the training and onboarding of new team members and Summer interns.

And they were useful in workshops where students could have first-hand experience with the mechanisms and components that, in production, run millions of jobs.

**Poster session / 546**

## ProtoDUNE HD Offline Computing Infrastructure

**Author:** Barnali Chowdhury<sup>1</sup>

<sup>1</sup> *Argonne National Laboratory*

**Corresponding Author:** bchowdhury@anl.gov

The Deep Underground Neutrino Experiment (DUNE), hosted by the U.S. Department of Energy's Fermilab, is expected to begin operations in the late 2020s. The validation of one far detector module design for DUNE will come from operational experience gained from deploying offline computing infrastructure for the ProtoDUNE (PD) Horizontal Drift (HD) detector. The computing infrastructure of PD HD is developed to achieve the physics goals by storing, globally distributing, cataloging, reconstructing, simulating, and analyzing data. Offline computing activities start with raw data in Hierarchical Data Format (HDF5) collected by DUNE's data acquisition (DAQ) system. The data is conveyed from the neutrino platform at CERN to the host labs using a data pipeline that includes Rucio for replica management and FTS3 for transport. The database system for PD HD comprises several backend relational databases (run configuration, beam instrumentation, slow controls, and calibration) that are fed into a Master Store database, an unstructured database, containing all of the information collected from the backend databases. A subset of this information, needed by offline users, is then moved into a lighter-weight relational database that users can access via an API or by an art service during the reconstruction and analysis stage. The primary software strategy for event reconstruction has been the adoption of algorithms that are flexible and accessible enough to support creative software solutions and advanced algorithms as HEP computing evolves. The collaboration anticipates making substantial use of HPCs for ProtoDUNE HD algorithms.

**Poster session / 547**

## FORM, a Fine-grained Object Reading/Writing Model for DUNE

**Author:** Barnali Chowdhury<sup>1</sup>

<sup>1</sup> *Argonne National Laboratory*

**Corresponding Author:** bchowdhury@anl.gov

DUNE's current processing framework (art) was branched from the event processing framework of CMS, a collider-physics experiment. Therefore art is built on event-based concepts as its fundamental processing unit. The "event" concept is not always helpful for neutrino experiments, such as DUNE. DUNE uses trigger records that are much larger than collider events (several GB vs. MB). Therefore, to avoid allocating large chunks of memory, DUNE is developing a framework that is able to break apart trigger records into smaller segments for more granular processing, and then stitch those chunks back together into an event.

In order for such an event-processing framework to work it needs to be integrated with input/output infrastructure that allows for fine-grained storage and I/O. FORM (Fine-grained Object Reading/Writing Model), a DUNE project, which intends to explore fine grained I/O and storage framework to store data in finer containers. Data objects need to be partitioned into segments to serve the fine grained processing, and be stored separately at accessible locations. Thus enabling I/O to read/write segmented data objects individually to avoid excessive memory consumption caused by reading large storage objects. The details of data storage and I/O should be encapsulated by the framework and transparent to client code such as algorithms. The persistence framework, FORM, is designed to write and read data in multiple smaller entries/events resulting in improved concurrency.

Parallel (Track 5) / 548

## Navigating Phase Space for Event Generation –interfacing Sherpa with BAT.jl

**Authors:** Cornelius Grunwald<sup>1</sup>; Kevin Alexander Kroeninger<sup>1</sup>; Salvatore La Cagnina<sup>None</sup>; Steffen Schumann<sup>None</sup>; Timo Janssen<sup>None</sup>

<sup>1</sup> *Technische Universitaet Dortmund (DE)*

**Corresponding Authors:** kevin.alexander.kroeninger@cern.ch, cornelius.grunwald@cern.ch, steffen.schumann@phys.uni-goettingen.de, salvatore.lacagnina@udo.edu, timo.janssen@theorie.physik.uni-goettingen.de

The generation of Monte Carlo events is a crucial step for all particle collider experiments. Accurately simulating the hard scattering processes is the foundation for subsequent steps, such as QCD parton showering, hadronization, and detector simulations. A major challenge in event generation is the efficient sampling of the phase spaces of hard scattering processes due to the potentially large number and complexity of Feynman diagrams and their interference and divergence structures.

In this presentation, we address the challenges of efficient Monte Carlo event generation and demonstrate improvements that can be achieved through the application of advanced sampling techniques. We highlight that using the algorithms implemented in BAT.jl for sampling the phase spaces given by Sherpa offers great flexibility in the choice of sampling algorithms and has the potential to significantly enhance the efficiency of event generation.

By interfacing BAT.jl, a package designed for Bayesian analyses that offers a collection of modern sampling algorithms, with the Sherpa event generator, we aim to improve the efficiency of phase space exploration and Monte Carlo event generation. We combine the physics-informed multi-channel sampling approach of Sherpa with advanced sampling techniques such as Markov Chain Monte Carlo (MCMC) and Nested Sampling. Additionally, we investigate the potential of novel machine learning-enhanced sampling methods to optimize phase space mappings and accelerate the event generation process. The current prototype interface between Sherpa and BAT.jl features a modular design that offers full flexibility in selecting target processes and provides detailed control over the sampling algorithms. It also allows for a simple integration of innovative sampling techniques such as normalizing flow-enhanced MCMC.

Parallel (Track6) / 549

## Bridging Reproducibility Gaps in Partial Wave Analysis

**Authors:** Ilya Segal<sup>1</sup>; Mikhail Mikhasenko<sup>2</sup>; Remco De Boer<sup>None</sup>

<sup>1</sup> *Ruhr-Universitaet Bochum (DE)*

<sup>2</sup> *Ruhr Univeristy Bochum*

**Corresponding Authors:** remco.deboer@rub.de, ilya.segal@cern.ch, mikhail.mikhasenko@cern.ch

Nearly none of the models from partial wave analysis can be reproduced based on published papers due to omitted nuances and implementation details. This issue affects progress and reliability in high-energy physics. Our project addresses this by standardizing the serialization of amplitude models into a lightweight, human-readable format, starting with three-body decay analyses. This standardization ensures accurate model verification, addressing common issues found even in published research.

We have developed a centralized repository containing a collection of models from LHCb and COMPASS analyses using `ThreeBodyDecays.jl`, `CompPWA`, and `TFAnalysis`. The serialized models facilitate community reanalysis, enable the use of models in MC generators, and provide data for testing and benchmarking new frameworks. We employ the `Pixi` setup for reproducible package management across platforms and `Quarto` for multilanguage support. Julia and Python notebooks run different frameworks, enhancing analysis and visualization.

As a fresh development, we see potential for broader adoption, including a possible extension to ROOT via the HS3 initiative. Success requires community support and collaboration with framework developers, advancing transparent, reproducible research.

Further details are available on our project webpage.

**Parallel (Track 5) / 550**

## **New RooFit PyROOT interfaces for connections with Machine Learning**

**Authors:** Jonas Rembser<sup>1</sup>; Lorenzo Moneta<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Authors:** jonas.rembser@cern.ch, lorenzo.moneta@cern.ch

With the growing datasets of HE(N)P experiments, statistical analysis becomes more computationally demanding, requiring improvements in existing statistical analysis algorithms and software. One way forward is to use Machine Learning (ML) techniques to approximate the otherwise untractable likelihood ratios. Likelihood fits in HEP are often done with RooFit, a C++ framework for statistical modelling that is part of ROOT. This contribution demonstrates how learned likelihood ratios can be used in RooFit analyses, showcasing new RooFit features that were developed for that purpose. Since ML models are often created with Python libraries, this necessitated new RooFit pythonizations, e.g. for using Python functions as RooFit functions in general. Some of these pythonizations were only possible by a major PyROOT upgrade that was undertaken this year. Therefore, this contribution will also summarize the new PyROOT features, resulting in a presentation that will promote both new RooFit and PyROOT features for the benefit of the users of the most recent ROOT versions.

**Parallel (Track 5) / 551**

## **Monte Carlo efficiency via negative weight reduction in Herwig**

**Author:** James Christopher Whitehead<sup>None</sup>

**Co-author:** Andrzej Konrad Siodmok <sup>1</sup>

<sup>1</sup> *Jagiellonian University (PL)*

**Corresponding Authors:** james.whitehead@ifj.edu.pl, a.siodmok@cern.ch

The generation of large event samples with Monte Carlo Event Generators is expected to be a computational bottleneck for precision phenomenology at the HL-LHC and beyond. This is due in part to the computational cost incurred by negative weights in ‘matched’ calculations combining NLO perturbative QCD with a parton shower: for the same target uncertainty, a larger sample must be generated.

We summarise two approaches taken to tackle this problem in Herwig: the development of the KrkNLO matching method, which uses a redefinition of the PDF factorisation scheme to guarantee positive weights by construction, and the restructuring of the Matchbox module to reduce the fraction of negative weights for MC@NLO matching.

**Poster session / 552**

## Multi-Experiment dCache Storage Service at BNL

**Author:** Carlos Fernando Gamboa<sup>1</sup>

**Co-authors:** Iris Wu <sup>2</sup>; Jason Smith <sup>2</sup>; Justin Spradley ; Ognian Novakov ; Qiulan Huang <sup>1</sup>; Robert Hancock ; Shigeki Misawa <sup>1</sup>; Tim Chou ; Zhenping Liu

<sup>1</sup> *Brookhaven National Laboratory (US)*

<sup>2</sup> *Brookhaven National Laboratory*

**Corresponding Authors:** hancock@bnl.gov, novakov@bnl.gov, zhliu@bnl.gov, cgamboa@bnl.gov, iriswu@rcf.rhic.bnl.gov, tchou@bnl.gov, misawa@bnl.gov, jspradley@bnl.gov, qiulan.huang@cern.ch, smithj4@bnl.gov

The Scientific Data and Computing Center (SDCC) at Brookhaven National Laboratory (BNL) is dedicated to providing access to storage services for a diverse array of Nuclear and High Energy Physics (HEP) experiments, including LHC-ATLAS, Belle2 and DUNE. With an aggregate of 143 PiBs of data, distributed and managed by independent storage instances dedicated to each experiment, approximately 30% of the data is stored on disk, while the rest resides on tape.

The underlying software technology supporting this storage infrastructure is dCache1.

This presentation aims to provide a comprehensive overview of the BNL storage service, delving into operational challenges, recent advancements, and key insights gained from experience.

1 <https://www.dcache.org/>

**Parallel (Track 8) / 553**

## Leveraging Language Models to Navigate Conference Abstracts: An Open-Source Approach

**Author:** Gordon Watts<sup>1</sup>

<sup>1</sup> *University of Washington (US)*

**Corresponding Author:** gwatts@uw.edu

Large Language Models (LLMs) have emerged as a transformative tool in society and are steadily working their way into scientific workflows. Despite their known tendency to hallucinate, rendering them perhaps unsuitable for direct scientific pipelines, LLMs excel in text-related tasks, offering a unique solution to manage the overwhelming volume of information presented at large conferences such as ACAT, ICHEP, and CHEP. This poster presents an innovative open-source application that harnesses the capabilities of an LLM to rank conference abstracts based on a user's specified interests. By providing a list of interests to the LLM, it can sift through a multitude of abstracts, identifying those most relevant to the user, effectively helping to tailor the conference experience. The LLM, in this context, serves an assistant role, aiding conference attendees in navigating the deluge of information typical of large conferences. The poster will detail the workings of this application, provide prompts to optimize its use, and discuss potential future directions for this type of application.

**Parallel (Track 8) / 554**

## **Saving through Sharing - Reflections on Sustainability in HEP**

**Authors:** Kristin Lohwasser<sup>1</sup>; Peter Millington<sup>2</sup>; Rakhi Mahbubani<sup>3</sup>

<sup>1</sup> *University of Sheffield (GB)*

<sup>2</sup> *University of Manchester*

<sup>3</sup> *Rudjer Boskovic Institute (HR)*

**Corresponding Authors:** kristin.lohwasser@cern.ch, rakhi@cern.ch, peter.millington@manchester.ac.uk

The climate crisis is a serious threat to prosperity and fundamental science in the mid-term future and to avert it will require significant effort. Closing the energy gap between projected energy usage and the available carbon-free energy will require expanding CO<sub>2</sub>-free energy production by a factor of 12, or reducing energy consumption by a factor of 2, by 2030.

The HECAP+ report 1 reflects on the current practices in HEP. This talk discusses some of the highlights of the report before reflecting on energy savings through data sharing. Theorists and experimentalists often work with large simulated datasets, generated independently for each experiment/theoretical study. We identify opportunities for energy savings through sharing of simulated data, estimate the GHG mitigation potential of these measures, and present potential advantages and drawbacks. We discuss the infrastructure necessary to facilitate effective data sharing across HECAP+ fields.

1 Environmental sustainability in basic research - A perspective from HECAP+, Sustainable HECAP+ Initiative, <https://sustainable-hecap-plus.github.io/>, accepted by JINST

**Poster session / 555**

## **An automation framework for the calibration of the CMS Precision Proton Spectrometer**

**Author:** Andrea Bellora<sup>1</sup>

**Co-author:** Tomasz Dariusz Ostafin<sup>2</sup>

<sup>1</sup> *Universita e INFN Torino (IT)*

<sup>2</sup> *AGH University of Krakow (PL)*

**Corresponding Authors:** tomasz.dariusz.ostafin@cern.ch, andrea.bellora@cern.ch

The Precision Proton Spectrometer (PPS) is a near-beam spectrometer that utilizes timing and tracking detectors to measure scattered protons surviving collisions at the CMS interaction point (IP). It is installed on both sides of CMS, approximately 200 meters from the IP, within mechanical structures called Roman Pots. These special beam pockets enable the detectors to approach the LHC beam within a few millimeters of its center. Due to the challenging environment, PPS detectors require frequent calibrations and close monitoring. This talk will introduce an automation software framework designed to streamline the calibration process, reducing the time users spend on these tasks, facilitating their implementation, and enhancing the monitoring of their execution and results. Developed alongside other CMS subsystems, the framework supports multi-stage calibrations that leverage CERN's distributed computing resources to run containerized tasks. Industry-grade tools such as Jenkins, InfluxDB, and Grafana are employed for monitoring the calibration execution and storing results, which can further be processed to identify anomalies in the data quality.

**Poster session / 556**

## Testing abstract

**Author:** Bartłomiej Rachwał<sup>1</sup>

<sup>1</sup> *AGH University of Krakow (PL)*

**Corresponding Author:** bartlomiej.rachwal@cern.ch

Hello

**Plenary session / 557**

## The State of Event Generators for the HL-LHC

**Author:** Olivier Mattelaer<sup>1</sup>

<sup>1</sup> *UCLouvain*

**Corresponding Author:** olivier.mattelaer@uclouvain.be

High-Luminosity LHC will provide an unprecedented amount of experimental data. The improvement in experimental precision needs to be matched with an increase of accuracy in the theoretical predictions, stressing our compute capability.

In this talk, I will focus on the current and future precision needed by LHC experiments and how those needs are supplied by Event Generators. I will focus on how computationally intensive these computations are. And finally, describe the various software development efforts underway (or in progress) to tackle such challenges, both in the direction of AI development and better usage of GPUs.

560

## Large Language Models in Physics

**Author:** Sarah Heim<sup>1</sup>

<sup>1</sup> *Deutsches Elektronen-Synchrotron (DE)*

**Corresponding Authors:** gregor.kasieczka@cern.ch, sarah.heim@cern.ch

Recent Large Language Models like ChatGPT show impressive capabilities, e.g. in the automated generation of text and computer code. These new techniques will have long-term consequences, including for scientific research in fundamental physics. In this talk I present the highlights of the first Large Language Model Symposium (LIPS) which took place in Hamburg earlier this year. I will focus on high energy physics and will also give an outlook towards future developments and applications.