


ACAT 2017

Monday 21 August 2017 - Friday 25 August 2017

University of Washington, Seattle



Book of Abstracts

Contents

The upgrade of the LHCb trigger for Run III	1
Data Transfer Node Resource Manager	1
The ALICE O2 common driver for the C-RORC and CRU readout cards	2
Testing the limits of an LVS - GridFTP Cluster as a replacement for Bestman	2
Robust circle reconstruction by a modified Riemann fit	3
Data Knowledge Base for HENP Scientific Collaborations	3
The new ATLAS Fast Calorimeter Simulation	4
Generative Adversarial Networks for Simulation	5
The R* operation and combinatorial challenges at five loops	5
Leveraging the checkpoint-restart technique for optimizing CPU efficiency of ATLAS production applications on opportunistic platforms	5
An ATLAS distributed computing architecture for HL-LHC	6
The ATLAS Data Management System Rucio: Supporting LHC Run-2 and beyond	7
Global heterogeneous resource harvesting: the next-generation PanDA pilot for ATLAS	7
The ATLAS Event Index: The Architecture of the Core Engine	8
Modernising ATLAS Software Build Infrastructure	8
A scalable new mechanism to store and serve the ATLAS detector description through a REST web API	9
A modern approach to HEP visualization - ATLASrift	9
Virtualization of the ATLAS software environment on a shared HPC system	10
GeantV alpha-release preview	11
Shared I/O components for the ATLAS multi-processing framework	11
Background Suppression with the Belle II Neural Network Trigger	12
Event Display in JUNO Experiment	12

Muon g-2 reconstruction and analysis framework	13
Muon reconstruction of JUNO experiment with convolutional neural networks	13
A Conditions Data Management System for HEP Experiments	14
An Online system based (almost exclusively) on off-the-shelf hardware for a trigger-free read-out of the LHCb experiment	14
Using containers with ATLAS offline software	15
ATLAS BigPanDA Monitoring	15
Update of BESIII Event Display System	16
A study on the applicability of Recommender Systems for the Production and Distributed Analysis system PanDA of the ATLAS Experiment	16
Toward real-time data query systems in HEP	17
Designing and prototyping the control system for the Cherenkov Telescope Array	17
Cross-architecture Kalman filter benchmarks on modern hardware platforms	18
Tomographic Reconstruction of LArTPC Events using Wire-cell	19
Deep-learning in jet reconstruction at CMS	19
CMS Analysis and Data Reduction with Apache Spark	20
Track seeding in the Outer Tracker of CMS for HL-LHC	20
Use of GPUs at the CMS High-Level Trigger during Phase-1	21
Continuous software quality analysis for the ATLAS experiment	21
Machine Learning Algorithms for b-jet tagging at the ATLAS experiment	22
Tools for Trigger Rate Monitoring at CMS	22
Parallel Random Number Generation for SIMD/SIMT	23
GQLink: an implementation of Quantized State System (QSS) methods in Geant4	23
High-statistics and GPU-accelerated data analysis in IceCube	24
Deep Neural Networks for Physics Analysis on low-level whole-detector data at the LHC	24
Striped Data Server for Scalable Parallel Data Analysis	25
Machine Learning Photons Separation in the LHCb Calorimeter	25
Parallelized JUNO simulation software based on SNIPEr	26
Four-point function in general kinematics through geometrical splitting and reduction	26
Solving the Signal-Versus-Background Classification Problem Using Neural Network	27

Machine Learning based global particle identification algorithms at LHCb experiment . . .	27
NNLO ntuples	28
Dynamic sharing of tape drives accessing scientific data	28
Convolutional Neural Network for Track Seed Filtering at the CMS HLT	28
The LHCb Software and Computing Upgrade towards LHC Run 3	29
Building a scalable python distribution for HEP data analysis	29
Last developments of the INFN CNAF Long Term Data Preservation (LTDP) project: the CDF data recover and safekeeping	30
Predictive analytics tools to adjust and monitor performance metrics for the ATLAS Pro- duction System	31
Performance of the ATLAS Tau Trigger in Run 2	31
BDTs in the Level 1 Muon Endcap Trigger at CMS	32
The ATLAS Trigger Simulation with Legacy Software	32
Convolution Neural networks in nucleon Decay Searches in Liquid Argon time projection Chambers	33
Experience with SPLUNK for archiving and visualization of operational data in ATLAS TDAQ system	33
Multithreading in the ATLAS High-Level Trigger	34
Test Management Framework for the Data Acquisition of the ATLAS Experiment	35
The VISPA Internet-Platform in Deep Learning Applications	35
Exploiting Apache Spark platform for CMS computing analytics	36
Event reconstruction and simulation in PandaRoot for the PANDA experiment	36
Mastering Opportunistic Computing Resources for HEP	37
The ATLAS Electron and Photon Trigger	38
The Simulation Library of the Belle II Software	38
Computer simulation on homogeneity testing for weighted data sets used in HEP	39
Design and Execution of make-like Distributed Analyses	39
Modeling of modifications induced by jets in the relativistic bulk nuclear matter	40
Supercomputers, Clouds and Grids powered by BigPanDA for Brain studies	41
Virtualization of the PNNL High Energy Physics Computing Infrastructure	41
Real-time alignment and reconstruction: performance and recent developments at the LHCb experiment	42

Application of deep learning to the analysis for $B \rightarrow K^* \text{ Gamma}$ in Belle II	42
Opportunistic data locality for HEP analysis workflows	43
The management of heterogeneous resources in Belle II	43
Applications of A-Common-Tracking-Software (ACTS) in the Future Circular Collider Design Study	44
Direct numerical computation and its application to the higher-order radiative corrections	44
Loopedia, a Database for Loop Integrals	45
CODE-RADE a user centric software delivery system for science	45
A development of an accelerator board dedicated for multi-precision arithmetic operations and its application to Feynman loop integrals II	45
Performance studies of GooFit on GPUs versus RooFit on CPUs while estimating the global statistical significance of a new physical signal	46
Design and implementation of data cache and access system across remote sites	47
High-speed evaluation of loop integrals using lattice rules	47
Weakly Supervised Classifiers in High Energy Physics	48
Parallel computing of SNiPER based on Intel TBB	48
Improving Science Yield for NASA Swift with Automated Planning Technologies.	49
a loosely coupled scalable cloud infrastructure	49
Speeding up experiments software with VecCore, a portable SIMD library	50
Optimizing ROOT I/O for Analysis	50
Machine Learning for Antihydrogen Detection at ALPHA	51
Calculation of One Loop Integrals with Package-X and CollierLink	52
Novel functional and distributed approaches to data analysis available in ROOT	53
Increasing Parallelism in the ROOT I/O subsystem	54
Speeding up prediction performance of the boosting decision trees-based learning models.	54
Learning to Remove Pileup at the LHC with Jet Images	55
Analysis Preservation and Systematic Reinterpretation within the ATLAS Experiment	55
Belle II Conditions Database Interface	55
FELIX: the new detector readout system for the ATLAS experiment	56
Multi-core supports in JUNO distributed computing	56

Identification of Hadronically Decaying W Bosons and Top Quarks using Multivariate Techniques at ATLAS	57
Hydra: A framework for data analysis in massively parallel platforms	57
Enabling High Availability Service with oVirt Virtualization and CephFS	58
Invited parallel: A Machine Learning tool for fast simulation in GeantV	59
Upgrade of the YARR DAQ system for the ATLAS Phase-II Pixel detector readout chip	59
N-jettiness subtraction	60
GooFit 2.0	60
Deep Learning Method for Inferring Cause of Data Anomalies	61
Numerical techniques for 2- and 3-loop integrals	61
Making Containers Lazy With Docker and CernVM-FS	61
A quantitative review of data formats for HEP analyses	62
DAQExpert - An expert system to increase CMS data-taking efficiency	62
Modeling detector digitization and read-out with adversarial networks	63
Parallelized Kalman-Filter-Based Reconstruction of Particle Tracks on Many-Core Architectures	64
Massive Parallel Lattice QCD Computing on FPGA Accelerator with Data-Flow Programming	64
Present and future of PanDA WMS integration with Titan supercomputer at OLCF	65
Gradient reversal for MC/real data calibration	66
Event Reconstruction on NOvA and DUNE with Deep Convolutional Networks	66
Real-time heterogeneous stream processing with NaNet in the NA62 experiment.	67
Automated proton track identification in MicroBooNE using gradient boosted decision trees	67
ROOT at the center of the future analysis ecosystem	68
Alignment and Calibration Framework for the Belle II detector	68
Parallelisation and Vectorisation of ROOT Fitting classes	69
How easily can neural networks learn relativity?	69
A Toolkit to Study Sensitivity of the Geant4 Predictions to the Variations of the Physics Model Parameters	70
NNPDF: Neural Networks for precision PDF determinations	70
High-precision calculation of the 4-loop contribution to the electron g-2 in QED	71

A federated Xrootd cache	71
Exploring end-to-end deep learning solutions for event classification at CMS	71
Networking requirements for Tier 2 sites	72
A container model for resource provision at a WLCG Tier-2	72
The HEP.TrkX Project: deep neural networks for HL-LHC online and offline tracking . . .	73
Federating distributed storage for clouds in ATLAS	74
Numerical Methods and the 4-point 2-loop Higgs amplitudes	74
A Java library to perform S-expansions of Lie algebras	74
Algorithm to find S-related Lie algebras	75
New developments in FORM	75
Deep Learning and Particle Physics	75
Deep Learning usage by large experiments	76
ML in theoretical physics: from PDFs to MC tools	76
Statistics, Data Mining, and Machine Learning in Astronomy	76
Meta: Toward generative C++	76
New results on the $g-2$ calculation	77
FPGA Datacenters – The New Supercomputer	77
Generative models for fast simulation	77
Physics and Computing Challenges in Lattice QCD	77
HEP Software Development in the next decade; the views of the HSF community	78
Development and optimization of the MCFM generator	78
HPC for HEP	78
Deep Learning on the Intel Nervana platform	78
D-wave Quantum Computing	79
Games and loop integrals	79
Survey of how DL is used in a large experiment	79
Machine Learning in HEP	79
Round Table on Machine Learning	79
Deep Learning in Particle physics	80
Data, Software, and Knowledge Preservation of Scientific Results	80

Track reconstruction algorithms in high pile up environments	80
Why Cultural Diversity Matters: Becoming Culturally competent	80
Round Table on Diversity	81
Summary of Track 1	81
Summary of Track 2	81
Summary of Track 3	82
Summary of the Conference	82
Machine Learning at NVIDIA	82
Go-HEP: writing concurrent software with ease and Go	82
On the notion "significance of the difference"	83
Complete one-loop electroweak radiative corrections to polarized Bhabha scattering in SANC	83
THESEUS event generator for collisions at energies available at NICA and FAIR	84
Transfer Learning in Astronomy: A New Machine Learning Paradigm	84
Round table discussion: Analytical vs. numerical methods for NNLO+ computations for LHC	85
Round-Table: Using heterogeneous resources for HEP computing	85
Background Suppression with the Belle II Neural Network Trigger	85
The ATLAS Electron and Photon Trigger	85
Boosted Decision Trees in the Level 1 Muon Endcap Trigger at CMS	86
Tomographic Reconstruction of LArTPC Events using Wire-cell	86
Muon reconstruction of JUNO experiment with convolutional neural networks	86
Implementing the Belle II Conditions Database using Industry-Standard Tools	86
Belle II Conditions Database Overview	86
Machine-Learning techniques for electro-magnetic showers identification in OPERA datasets	87
Final remarks	88
Mastering Opportunistic Computing Resources for HEP	88
Software Citations	88
Welcome from the Department of Physics Chair	89
Welcome To ACAT 20171	89

ACAT Informal Bar Crawl	89
Parllized JUNO simulation software based on SNIPER	89
Paralleization and vectorization of ROOT fitting classes	89
High Statistics and GPU Accerated Data Analysis in ICECUBE	90
Speeding up Software with VECCORE, a portable SIMD library	90
Banquet	90
Awards!	90

Track 1: Computing Technology for Physics Research / 7**The upgrade of the LHCb trigger for Run III**

Authors: Rosen Matev¹; Conor Fitzpatrick²; Sascha Stahl¹; Mika Anton Vesterinen³

¹ CERN

² Ecole Polytechnique Federale de Lausanne (CH)

³ Ruprecht-Karls-Universitaet Heidelberg (DE)

Corresponding Authors: conor.fitzpatrick@cern.ch, sascha.stahl@cern.ch, rosen.matev@cern.ch, mika.vesterinen@cern.ch

The LHCb experiment plans a major upgrade of the detector and DAQ systems in the LHC long shutdown II (2018–2019). For this upgrade, a purely software based trigger system is being developed, which will have to process the full 30 MHz of bunch-crossing rate delivered by the LHC. A fivefold increase of the instantaneous luminosity in LHCb further contributes to the challenge of reconstructing and selecting events in real time. Optimal usage of the trigger output bandwidth will be enabled by the Turbo paradigm, in which only high level reconstructed objects and a subset of the raw event data are persisted. In this talk we discuss the plans and progress towards achieving reconstruction and selection with a 30 MHz throughput by means of efficient utilisation of modern CPU microarchitectures, and strategies for formal testing of the physics content in the trigger, which is a crucial component of a system in which real-time analysis is to be performed.

Track 1: Computing Technology for Physics Research / 8**Data Transfer Node Resource Manager**

Authors: Justas Balcas¹; Dorian Kcira¹; Harvey Newman¹; Azher Mughal¹; Maria Spiropulu¹

¹ California Institute of Technology (US)

Corresponding Authors: azher.mughal@cern.ch, maria.spiropulu@cern.ch, dorian.kcira@cern.ch, harvey.newman@cern.ch, justas.balcas@cern.ch

The Caltech team in collaboration with network, computer science and HEP partners at the DOE laboratories and universities, has developed high-throughput data transfer methods and cost-effective data systems that have defined the state of the art for the last 15 years.

The achievable stable throughput over continental and transoceanic distances using TCP-based open source applications, notably Caltech's Fast Data Transfer application (FDT), has risen by two orders of magnitude over the last decade. This has happened in concert with optimally engineered and configured server systems using the latest generation of motherboards, storage and network interfaces.

These developments constitute the basis for the Data Transfer Nodes (DTNs) used in the SENSE project. The SDN for end-to-end Networked Science at the Exascale (SENSE) project is developing Software Defined Network (SDN) based technologies to enable on-demand end-to-end network services which can be tailored to individual domain science application workflow requirements.

The overarching goal of SENSE is to enable National Labs and universities to request and provision end-to-end intelligent network services for their application workflows leveraging SDN capabilities. The specific areas of Caltech's responsibility and technical innovation within the full scope of the SENSE project are:

End-systems involving DTNs with high throughput capability and instrumentation software aimed at comprehensive end-system monitoring, auto-configuration and tuning Site orchestration and integration of end-to-end flows across the Science DMZ interface.

Real time system monitoring and optimization.

The DTN and its design play a crucial role in the overall SENSE project architecture both as the source and destination of data and as an integral part of the end-to-end network flows. A description of the DTNs and their resource manager (DTN-RM) will be given.

Poster Session / 10

The ALICE O2 common driver for the C-RORC and CRU readout cards

Author: Pascal Boeschoten¹

Co-author: Filippo Costa²

¹ *Ministere des affaires etrangeres et europeennes (FR)*

² *CERN*

Corresponding Authors: filippo.costa@cern.ch, pascal.boeschoten@cern.ch

ALICE (A Large Ion Collider Experiment) is the heavy-ion detector designed to study the strongly interacting state of matter realized in relativistic heavy-ion collisions at the CERN Large Hadron Collider (LHC). A major upgrade of the experiment is planned during the 2019-2020 long shutdown. In order to cope with a data rate 100 times higher than during LHC Run 2 and with the continuous readout of the Time Projection Chamber (TPC), it is necessary to upgrade the Online and Offline Computing to a new common system called O2. The O2 readout chain will use commodity x86 Linux servers equipped with custom PCIe FPGA-based readout cards. This paper discusses the driver architecture for the cards that will be used in O2: the PCIe v2 x8, Xilinx Virtex 6 based C-RORC (Common Readout Receiver Card) and the PCIe v3 x16, Intel Arria 10 based CRU (Common Readout Unit). Access to the PCIe cards is provided via three layers of software. Firstly, the low-level PCIe (PCI Express) layer responsible for the userspace interface for low-level operations such as memory mapping the PCIe BAR (Base Address Registers) and creating scatter-gather lists, which is provided by the PDA (Portable Driver Architecture) library developed by the Frankfurt Institute for Advanced Studies (FIAS). Above that sits our userspace driver which implements synchronization, controls the readout card –e.g. resetting and configuring the card, providing it with bus addresses to transfer data to and checking for data arrival –and presents a uniform, high-level C++ interface that abstracts over the differences between the C-RORC and CRU. This interface –of which direct usage is principally intended for high-performance readout processes –allows users to configure and use the various aspects of the readout cards, such as configuration, DMA transfers and commands to the front-end. The top layer consists of a Python wrapper and command-line utilities that are provided to facilitate scripting and executing tasks from a shell, such as card resetting; performing benchmarks; reading or writing registers; and running test suites. Additionally, the paper presents the results of benchmarks in various test environments. Finally, we present our plans for future development, testing and integration.

Track 1: Computing Technology for Physics Research / 11

Testing the limits of an LVS - GridFTP Cluster as a replacement for Bestman

Authors: Brian Paul Bockelman¹; Clifton Pottberg²; Edgar Fajardo Hernandez³; Terrence Martin⁴; Garhan Attebury⁵; Frank Wuerthwein⁴

¹ *University of Nebraska-Lincoln (US)*

² *University of California San Diego*

³ *Univ. of California San Diego (US)*

⁴ *UCSD*

⁵ *University of Nebraska-Lincoln*

Corresponding Authors: brian.bockelman@cern.ch, cpottber@ucsd.edu, atebury@cse.unl.edu, fkw@fnal.gov, tmartin@physics.ucsd.edu, edgar.mauricio.fajardo.hernandez@cern.ch

The Worldwide LHC Computing Grid (WLCG) is the largest grid computing infrastructure in the world pooling the resources of 170 computing centers (sites). One of the advantages of grid computing is that multiple copies of data can be stored at different sites allowing user access that is independent of that site's geographic location, unique operating systems, and software. Each site is able to communicate using software stacks collectively referred to as "middleware". One of the middleware pieces is the storage element (SE) which manages data access between sites.

The middleware distributed by the Open Science Grid (OSG) previously used a storage resource manager (SRM) allowing for sites to expose their SEs for access by off-site compute elements (CEs) via the Grid File Transfer Protocol (GridFTP). OSG is eliminating the use of an SRM entirely and transitioning towards a solution based solely on GridFTP and Linux Virtual Server (LVS). LVS is a core component of the Linux kernel, so this change increases both maintainability and interoperability. In this document, we outline our methodologies and results from the large scale testing of an LVS+GridFTP cluster for data reads. Additionally, we discuss potential optimizations to the cluster to maximize total throughput.

Poster Session / 12

Robust circle reconstruction by a modified Riemann fit

Authors: Rudolf Fruhwirth¹; Are Strandlie²

¹ *Austrian Academy of Sciences (AT)*

² *NTNU Gjøvik, Norway*

Corresponding Authors: rudolf.fruhwirth@oeaw.ac.at, are.strandlie@ntnu.no

Circle finding and fitting is a frequent problem in the data analysis of high-energy physics experiments. In a tracker immersed in a homogeneous magnetic field, tracks with sufficiently high momentum are close to perfect circles if projected to the bending plane. In a ring-imaging Cherenkov detector, a circle of photons around the crossing point of a charged particles has to be found and its radius has to be estimated. In both cases, non-negligible background may be present that tends to complicate the circle finding and to bias the circle fit. In this contribution we present a robust circle finder/fitter based on a modified Riemann fit that significantly reduces the effect of background hits. As in the standard Riemann fit, the measured points are projected to the Riemann sphere or paraboloid, and a plane is fitted to the projected points. The fit is made robust by replacing the usual least squares estimate of the plane by a least median of squares (LMS) estimate. Because of the high breakdown point of the LMS estimator, the fit is insensitive to background points which can then be eliminated from the sample. This constitutes the finding stage. The plane is then refitted with the remaining points by an M-estimator in order to suppress eventual remaining outliers and to obtain the final circle fit. The method is demonstrated on artificial data with points on a circle plus up to 100% background points, with points on two overlapping circles with additional background, and with points obtained by the simulation of a generic inner tracking system with mirror hits and additional background. The results show high circle finding efficiency and small contamination of the final fitted circles.

Poster Session / 13

Data Knowledge Base for HENP Scientific Collaborations

Authors: Maria Grigoryeva¹; Marina Golosova¹; Torre Wenaus²; Alexei Klimentov²

¹ *Institute for Theoretical and Experimental Physics (RU)*

² *Brookhaven National Laboratory (US)*

Corresponding Authors: siarhei.padolski@cern.ch, alexei.klimentov@cern.ch, marina.golosova@cern.ch, maria.grigorieva@cern.ch, wenaus@gmail.com

Scientific collaborations operating on modern facilities generate vast volumes of data and auxiliary metadata, and the information is constantly growing. High energy physics data is a long term investment and contains the potential for physics results beyond the lifetime of a collaboration or/and experiment. Many existing HENP experiments are concluding their physics programs, and looking for ways to preserve their data heritage.

The Run1,2 estimated LHC experiments RAW data volume is 15+ PB/year and ~130 PB/year is expected by the end of Run 3 in 2022 and 200 PB/year for High-Luminosity LHC Run. Even today the managed data volume of the ATLAS experiment is close to 300PB.

Data Preservation HEP working group announced Data Preservation Model in May 2012. This model includes the preservation of real and simulated data, the analysis level, reconstruction and simulation software, and the preservation of documentation (such as internal notes, wikis, etc). However, existing long-term preservation resources are loosely connected and don't provide tools for the automatic reproduction of connections/links between various data and auxiliary metadata from storage subsystems. Data Knowledge Base R&D Project, started in 2016, is aimed at developing of the software environment and providing a coherent view/representation of the basic information preservation objects/components. The present architecture of DKB is based on the ontological model of HENP studies. The central storage –OpenLink Virtuoso –consolidates the basic metadata about scientific papers and internal documents, experimental environment and data samples, used in physical analysis. Specific services –metadata export/extraction/import tools, aggregation/integration modules – are organized as workflows run by Apache Kafka, providing nonstop data processing. One of the most challenging tasks is to establish and keep connectivity between data samples and scientific publications, internal notes and conference talks. Scientific publications and notes have information to establish the above connectivity. (Meta)Information could be extracted from papers and would be used to connect data samples and analysis, and then import these links into Virtuoso in accordance with the ontological model. As a result, all data samples, used in the data analysis described in the document of interest, can be obtained with a simple SPARQL request to Virtuoso. In the nearest future DKB architecture is planned to be enhanced with the SPARQL-endpoint services for InSpire HEP, CERN Document Server and Production Systems, thus providing virtual integration of these storage resources.

Track 1: Computing Technology for Physics Research / 15

The new ATLAS Fast Calorimeter Simulation

Author: Jana Schaarschmidt¹

¹ *University of Washington (US)*

Corresponding Author: jana.schaarschmidt@cern.ch

Producing the very large samples of simulated events required by many physics and performance studies with the ATLAS detector using the full GEANT4 detector simulation is highly CPU intensive. Fast simulation tools are a useful way of reducing CPU requirements when detailed detector simulations are not needed. During the LHC Run-1, a fast calorimeter simulation (FastCaloSim) was successfully used in ATLAS. FastCaloSim provides a simulation of the particle energy response at the calorimeter read-out cell level, taking into account the detailed particle shower shapes and the correlations between the energy depositions in the various calorimeter layers. It is interfaced to the standard ATLAS digitization and reconstruction software, and it can be tuned to data more easily than GEANT4. It is 500 times faster than full simulation in the calorimeter system.

Now an improved version of FastCaloSim is in development, incorporating the experience with the version used during Run-1. The new FastCaloSim makes use of machine learning techniques, such as principal component analysis and neural networks, to optimise the amount of information stored in the ATLAS simulation

infrastructure. This allows for further performance improvement by reducing the I/O time and the memory usage during the simulation job. A first prototype is available and is being tested and validated now. ATLAS plans to use this new FastCaloSim parameterization to simulate several billion

events in the upcoming LHC runs. It will be combined with other fast tools used in the ATLAS production chain. In this Fast Chain the simulation, digitisation and reconstruction of the events are handled by fast tools. In this talk, we will describe this new FastCaloSim parametrisation and the current status of the ATLAS Fast Chain.

Track 2: Data Analysis - Algorithms and Tools / 16

Generative Adversarial Networks for Simulation

Authors: Luke Percival De Oliveira^{None}; Michela Paganini¹; Ben Nachman²

¹ *Yale University (US)*

² *Lawrence Berkeley National Lab. (US)*

Corresponding Authors: luke.percival.de.oliveira@cern.ch, benjamin.philip.nachman@cern.ch, michela.paganini@cern.ch

We introduce the first use of deep neural network-based generative modeling for high energy physics (HEP). Our novel Generative Adversarial Network (GAN) architecture is able cope with the key challenges in HEP images, including sparsity and a large dynamic range. For example, our Location-Aware Generative Adversarial Network learns to produce realistic radiation patterns inside high energy jets - collimated sprays of particles resulting from quarks and gluons produced at high energy. The pixel intensities of the GAN-generated images faithfully span many orders of magnitude and reproduce the distributions of important low-dimensional physical properties (e.g. jet mass, n -subjettiness, etc.). We provide many visualizations of what the GAN has learned, to build additional confidence in the algorithm. Our results demonstrate that high-fidelity, fast simulation through GANs is a promising application of deep neural networks for solving one of the most important challenges facing HEP today.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 17

The R^* operation and combinatorial challenges at five loops

Authors: Ben Ruijl¹; Franz Herzog²; Takahiro Ueda³; Jozef Anton Vermaseren^{None}; Andreas Vogt⁴

¹ *Nikhef*

² *CERN*

³ *KEK*

⁴ *University of Liverpool*

Corresponding Authors: andreas.vogt@liv.ac.uk, benruyl@gmail.com, t68@nikhef.nl, tueda@nikhef.nl, franz.herzog@cern.ch

We show how an extended version of the R^* operation, a method to remove UV and soft IR divergences, can be used to calculate the poles of Feynman diagrams with arbitrary tensor structure from diagrams with fewer loops. We discuss solutions to combinatorial problems we encountered during the computation of the five loop QCD beta function, such as postponing Feynman rule substitutions, integral isomorphisms, and efficient tensor reductions.

Track 1: Computing Technology for Physics Research / 19

Leveraging the checkpoint-restart technique for optimizing CPU efficiency of ATLAS production applications on opportunistic platforms

Authors: Johannes Elmsheuser¹; Lukas Alexander Heinrich²; Wim Lavrijsen³; Vakho Tsulaia³

¹ *Brookhaven National Laboratory (US)*

² *New York University (US)*

³ *Lawrence Berkeley National Lab. (US)*

Corresponding Authors: lukas.heinrich@cern.ch, vakhtang.tsulaia@cern.ch, johannes.elmsheuser@cern.ch, wim.lavrijsen@cern.ch

Data processing applications of the ATLAS experiment, such as event simulation and reconstruction, spend considerable amount of time in the initialization phase. This phase includes loading a large number of shared libraries, reading detector geometry and condition data from external databases, building a transient representation of the detector geometry and initializing various algorithms and services. In some cases the initialization step can take as long as 10-15 minutes. Such slow initialization, being inherently serial, has a significant negative impact on overall CPU efficiency of the production job, especially when the job is executed on opportunistic, often short-lived, resources such as commercial clouds or volunteer computing. In order to improve this situation, we can take advantage of the fact that ATLAS runs large numbers of production jobs with similar configuration parameters (e.g. jobs within the same production task). This allows us to checkpoint one job at the end of its configuration step and then use the generated checkpoint image for rapid startup of thousands of production jobs. By applying this technique we can bring the initialization time of a job from tens of minutes down to just a few seconds. In addition to that we can leverage container technology for restarting checkpointed applications on the variety of computing platforms, in particular of platforms different from the one on which the checkpoint image was created.

We will describe the mechanism of creating checkpoint images of Geant4 simulation jobs with AthenaMP (the multi-process version of the ATLAS data simulation, reconstruction and analysis framework Athena) and the usage of these images for running ATLAS Simulation production jobs on volunteer computing resources (ATLAS@Home)

Poster Session / 20

An ATLAS distributed computing architecture for HL-LHC

Authors: Simone Campana¹; Torre Wenaus²

¹ *CERN*

² *Brookhaven National Laboratory (US)*

Corresponding Authors: wenaus@gmail.com, simone.campana@cern.ch

The ATLAS collaboration started a process to understand the computing needs for the High Luminosity LHC era. Based on our best understanding of the computing model input parameters for the HL-LHC data taking conditions, results indicate the need for a larger amount of computational and storage resources with respect of the projection of constant yearly budget for computing in 2026. Filling the gap between the projection and the needs will be one of the challenges in preparation for LHC Run-4. While the gains from improvements in offline software will play a crucial role in this process, a different model for data processing, management, access and bookkeeping should also be envisaged to optimise resource usage. In this contribution we will describe a straw man of this model, founded on basic principles such as single event level granularity for data processing and virtual data. We will explain how the current architecture will evolve adiabatically into the future distributed computing system, through the prototyping of building blocks that would be integrated in the production infrastructure as early as possible, so that specific use cases can be covered much earlier with respect of the HL-LHC time scale. We will also discuss how such system would adapt to and drive the evolution of the WLCG infrastructure in terms of facilities and services.

Poster Session / 21

The ATLAS Data Management System Rucio: Supporting LHC Run-2 and beyond

Authors: Martin Barisits¹; Thomas Beermann¹; Vincent Garonne²; Tomas Javurek³; Mario Lassnig¹; Cedric Serfon²; Simone Campana¹

¹ CERN

² University of Oslo (NO)

³ Albert-Ludwigs-Universitaet Freiburg (DE)

Corresponding Authors: simone.campana@cern.ch, thomas.beermann@cern.ch, martin.barisits@cern.ch, mario.lassnig@cern.ch, vincent.garonne@cern.ch, tomas.javurek@cern.ch, cedric.serfon@cern.ch

With this contribution we present the recent developments made to Rucio, the data management system of the High-Energy Physics Experiment ATLAS. Already managing 260 Petabytes of both official and user data, Rucio has seen incremental improvements throughout LHC Run-2, and is currently laying the groundwork for HEP computing in the HL-LHC era. The focus of this contribution are (a) the automations that have been put in place such as data rebalancing or dynamic replication of user data, as well as their supporting infrastructures such as real-time networking metrics or transfer time predictions; (b) the flexible approach towards inclusion of heterogeneous storage systems, including object stores, while unifying the potential access paths using generally available tools and protocols; (c) the improvements made to the real time monitoring of the system to alleviate the work of our human shifters; and (d) the adoption of Rucio for two other experiments, AMS and Xenon1t. We conclude by presenting operational numbers and figures to quantify these improvements, and extrapolate the necessary changes and developments for future LHC runs.

Poster Session / 22

Global heterogeneous resource harvesting: the next-generation PanDA pilot for ATLAS

Authors: Daniil Drizhuk¹; Paul Nilsson²; Mario Lassnig³; Wen Guan⁴; Danila Oleynik⁵; Alexey Anisenkov⁶; Pavlo Svirin⁷

¹ National Research Centre Kurchatov Institute (RU)

² Brookhaven National Laboratory (US)

³ CERN

⁴ University of Wisconsin (US)

⁵ Joint Institute for Nuclear Research (RU)

⁶ Budker Institute of Nuclear Physics (RU)

⁷ National Academy of Sciences of Ukraine (UA)

Corresponding Authors: ruslan.mashinistov@cern.ch, daniil.drizhuk@cern.ch, danila.oleynik@cern.ch, paul.nilsson@cern.ch, pavlo.svirin@cern.ch, wen.guan@cern.ch, mario.lassnig@cern.ch, alexey.anisenkov@cern.ch

The Production and Distributed Analysis system (PanDA), used for workload management in the ATLAS Experiment for over a decade, has in recent years expanded its reach to diverse new resource types such as HPCs, and innovative new workflows such as the event service. PanDA meets the heterogeneous resources it harvests in the PanDA pilot, which has embarked on a next-generation reengineering to efficiently integrate and exploit the new platforms and workflows. The new modular architecture is the product of a year of design and prototyping in conjunction with the design of a completely new component, Harvester, that will mediate a richer flow of control and information between pilot and PanDA. Harvester will enable more intelligent and dynamic matching between processing tasks and resources, with an initial focus on HPCs, simplifying the operator and user view of a PanDA site but internally leveraging deep information gathering on the resource to accrue

detailed knowledge of a site's capabilities and dynamic state to inform the matchmaking. This talk will give an overview of the new pilot architecture, how it will be used in and beyond ATLAS, its relation to Harvester, and the work ahead.

Poster Session / 23

The ATLAS Event Index: The Architecture of the Core Engine

Authors: Julius Hrivnac¹; Dario Barberis²; Andrea Favareto³; Fedor Prokoshin⁴; Grigori Rybkin¹; Rainer Toebicke⁵; Ruijun Yuan¹

¹ *Universite de Paris-Sud 11 (FR)*

² *Università e INFN Genova (IT)*

³ *Università degli Studi e INFN Genova*

⁴ *Federico Santa Maria Technical University (CL)*

⁵ *CERN*

Corresponding Authors: dario.barberis@cern.ch, andrea.favareto@cern.ch, julius.hrivnac@cern.ch, rainer.toebicke@cern.ch, fedor.prokoshin@cern.ch, grigori.rybkine@cern.ch, ruijun.yuan@cern.ch

The global view of the ATLAS Event Index system has been presented in the last ACAT. This talk will concentrate on the architecture of the system core component. This component handles the final stage of the event metadata import, it organizes its storage and provides a fast and feature-rich access to all information. A user is able to interrogate metadata in various ways, including by executing user-provided code on the data to make selections and to interpret the results. A wide spectrum of clients is available, from a set of linux-like commands to an interactive graphical Web Service. The stored event metadata contain the basic description of the related events, the references to the experiment event storage, the full trigger record and can be extended with other event characteristics. Derived collections of events can be created. Such collections can be annotated and tagged with further information. This talk will describe all system sub-components and their development evolution, which lead into the choices in the current architecture. The system performance, the development, the runtime environment and the interoperation with other ATLAS software components will be also described. The problems and mistakes made during the development will be explained. And the lessons for the future evolution of the Event Index software and the general data analysis framework will be summarised.

Track 1: Computing Technology for Physics Research / 24

Modernising ATLAS Software Build Infrastructure

Authors: Graeme Stewart¹; Attila Krasznahorkay²; Christian Gumpert²; Edward Moyses³; Elmar Ritsch²; Emil Obreshkov⁴; Frank Winklmeier⁵; Goetz Gaycken⁶; Karolos Potamianos⁷; Walter Lampl⁸; Alexander Undrus⁹

¹ *University of Glasgow (GB)*

² *CERN*

³ *University of Massachusetts (US)*

⁴ *University of Texas at Arlington (US)*

⁵ *University of Oregon (US)*

⁶ *University of Bonn (DE)*

⁷ *Lawrence Berkeley National Lab. (US)*

⁸ *University of Arizona (US)*

⁹ *Brookhaven National Laboratory (US)*

Corresponding Authors: undrus@bnl.gov, graeme.andrew.stewart@cern.ch, edward.moyse@cern.ch, emil.obreshkov@cern.ch, goetz.gaycken@cern.ch, karolos.potamianos@cern.ch, elmar.ritsch@cern.ch, frank.winklmeier@cern.ch, christian.gumpert@cern.ch, walter.lampl@cern.ch, attila.krasznahorkay@cern.ch

In the last year ATLAS has radically updated its software development infrastructure hugely reducing the complexity of building releases and greatly improving build speed, flexibility and code testing. The first step in this transition was the adoption of CMake as the software build system over the older CMT. This required the development of an automated translation from the old system to the new, followed by extensive testing and improvements. This resulted in a far more standard build process that was married to the method of building ATLAS software as a series of 12 separate projects from SVN.

We then proceeded with a migration of its code base from SVN to git. As the SVN repository had been structured to manage each package more or less independently there was no simple mapping that could be used to manage the migration into git. Instead a specialist set of scripts that captured the software changes across official software releases was developed. With some clean up of the repository and the policy of only migrating packages in production releases, we managed to reduce the repository size from 62GB to 220MB.

After moving to git we took the opportunity to introduce continuous integration so that now each code change from developers is built and tested before being approved.

With both CMake and git in place we also dramatically simplified the build management of ATLAS software. Many heavyweight homegrown tools were dropped and the build procedure was reduced to a single bootstrap of some external packages, followed by a full build of the rest of the stack. This has reduced the time for a build by a factor of 2. It is now easy to build ATLAS software, freeing developers to test compile intrusive changes or new platform ports with ease. We have also developed a system to build lightweight ATLAS releases, for simulation, analysis or physics derivations which can be built from the same branch.

Poster Session / 25

A scalable new mechanism to store and serve the ATLAS detector description through a REST web API

Authors: Riccardo Maria Bianchi¹; Ilija Vukotic²

¹ *University of Pittsburgh (US)*

² *University of Chicago (US)*

Corresponding Authors: riccardo.maria.bianchi@cern.ch, ilija.vukotic@cern.ch

Until now, geometry information for the detector description of HEP experiments was only stored in online relational databases integrated in the experiments' frameworks or described in files with text-based markup languages. In all cases, to build and store the detector description, a full software stack was needed.

In this paper we present a new and scalable mechanism to store the geometry data and to serve the detector description data through a REST web-based API. This new approach decouples the geometry information from the experiment's framework. Moreover, it provides new functionalities to users, who can now search for specific volumes and get partial detector description, or filter geometry data based on custom criteria.

We present two approaches to build a REST API to serve geometry data, based on two different technologies used in other fields and communities: the search engine Elasticsearch and the graph database Neo4j. We describe their characteristics and we compare them using real-world usage tests to test their speed and scalability, targeted to a HEP usage.

A modern approach to HEP visualization - ATLASrft

Authors: Ilija Vukotic¹; Riccardo Maria Bianchi²

¹ *University of Chicago (US)*

² *University of Pittsburgh (US)*

Corresponding Authors: ilija.vukotic@cern.ch, riccardo.maria.bianchi@cern.ch

At the times when HEP computing needs were mainly fulfilled by mainframes, graphics solutions for event and detector visualizations were necessarily hardware as well as experiment specific and impossible to use anywhere outside of HEP community. A big move to commodity computing did not precipitate a corresponding move of graphics solutions to industry standard hardware and software. In this paper, we list functionalities expected from contemporary tools and describe their implementation by a specific application: ATLASrft.

We start with a basic premise that HEP visualization tools should be open in practice and not only in intentions. This means that a user should not be limited to specific and little used platforms, HEP-only software packages, or experiment-specific libraries. Equally important is that no special knowledge or special access rights are needed. Using industry standard frameworks brings not only sustainability, but also good support, a lot of community contributed tools, and a possibility of community input in a form of feedback or direct help. Moreover, an ideal visualization tool should be accessible for non-expert people (for example for outreach and education purposes) as well as offering all the functionalities needed by the experts.

Next we share our experience developing the ATLASrft application. The application is based on the Unreal Engine, currently a gold standard in the field of interactive visualization. This gives us wide platform coverage (Linux, Windows, iOS, Android, Web browser, all VR platforms) and a seamless integration with diverse online platforms for both application delivery and multi-user support (ie. Oculus, Steam, Amazon). It makes integration of outreach oriented multimedia content - 4Pi coverage photos, guided tours, photos and videos - straightforward. We describe the usage of a web-based service to import the detector description geometry, in order to ensure high performance even for the most demanding platforms. We present the ATLASrft user interface, which proves that the expert and the outreach functionalities of the application are not mutually exclusive.

Poster Session / 27

Virtualization of the ATLAS software environment on a shared HPC system

Author: Ulrike Schnoor¹

Co-authors: Markus Schumacher²; Felix Bührer¹; Anton Josef Gamel¹; Konrad Meier³

¹ *Albert-Ludwigs-Universitaet Freiburg (DE)*

² *Physikalisches Institut-Albert-Ludwigs-Universitaet Freiburg-Unk*

³ *University of Freiburg*

Corresponding Authors: ulrike.schnoor@cern.ch, anton.josef.gamel@cern.ch, felix.buehrer@cern.ch, markus.schumacher@cern.ch, konrad.meier@rz.uni-freiburg.de

High-Performance Computing (HPC) and other research cluster computing resources provided by universities can be useful supplements to the collaboration's own WLCG computing resources for data analysis and production of simulated event samples. The shared HPC cluster "NEMO" at the University of Freiburg has been made available to local ATLAS users through the provisioning of virtual machines incorporating the ATLAS software environment analogously to a WLCG center. The talk describes the concept and implementation of virtualizing the ATLAS software environment to run both data analysis and production on the HPC host system which is connected to the existing Tier-3 infrastructure. Main challenges include the integration into the NEMO and Tier-3 schedulers in a dynamic, on-demand way, the scalability of the OpenStack infrastructure, as well as the automatic generation of a fully functional virtual machine image providing access to the local user environment, the dCache storage element and the parallel file system. The performance in the virtualized environment is evaluated for typical High-Energy Physics applications.

Track 1: Computing Technology for Physics Research / 28**GeantV alpha-release preview**

Authors: Guilherme Amadio¹; Ananya²; John Apostolakis¹; Marilena Bandieramonte¹; Abhijit Bhattacharyya³; Rene Brun¹; Philippe Canal⁴; Federico Carminati¹; Gabriele Cosmo¹; Laurent Duhem^{None}; Victor Daniel Elvira⁴; Andrei Gheata¹; Mihaela Gheata⁵; Ilias Goulas¹; Soon Yung Jun⁴; Jose Guilherme Lima⁶; Farah Hariri¹; Mihaly Novak¹; Witold Pokorski¹; Alberto Ribon¹; Raman Sehgal^{None}; Oksana Shadura⁷; Sofia Vallecorsa⁸; Sandro Christian Wenzel¹

¹ CERN

² IIT- Indian Institute of Technology (IN)

³ Bhabha Atomic Research Centre (IN)

⁴ Fermi National Accelerator Lab. (US)

⁵ Institute of Space Science (RO)

⁶ FermiLab (US)

⁷ National Academy of Sciences of Ukraine (UA)

⁸ Gangneung-Wonju National University (KR)

Corresponding Authors: amadio@cern.ch, andrei.gheata@cern.ch, rene.brun@cern.ch, lima@fnal.gov, daniel@fnal.gov, john.apostolakis@cern.ch, sandro.wenzel@cern.ch, farah.hariri@cern.ch, sofia.vallecorsa@cern.ch, mihaly.novak@cern.ch, federico.carminati@cern.ch, alberto.ribon@cern.ch, gabriele.cosmo@cern.ch, syjun@fnal.gov, pcanal@fnal.gov, mihaela.gheata@cern.ch, marilena.bandieramonte@cern.ch, ananya.ananya@cern.ch, oksana.shadura@cern.ch, sc.ramanshgal@gmail.com, witold.pokorski@cern.ch, laurent.duhem@intel.com, abhijit.bhattacharyya@cern.ch, ilias.goulas@cern.ch

GeantV went through a thorough community discussion in the fall 2016 reviewing the project's status and strategy for sharing the R&D benefits with the LHC experiments and with the HEP simulation community in general. Following up to this discussion GeantV has engaged onto an ambitious 2-year road-path aiming to deliver a beta version that has most of the performance features of the final product, partially integrated with some of the experiment's frameworks.

The initial GeantV prototype has been re-cast into a vector-aware concurrent framework able to deliver high-density floating point computation for most of the performance-critical components such as propagation in field and physics models. Electromagnetic physics models were adapted for the specific GeantV requirements, aiming for the full demonstration of shower physics performance in the alpha release this fall. We have revisited and formalized GeantV user interfaces and helper protocols, allowing to connect to user code and provide recipes to access efficiently MC truth and generate user data in a concurrent environment.

The presentation will give a preview of the features available in the alpha release, including a new R&D for ML-driven fast simulation engine, up to date performance figures compared to Geant4 and the status of the co-processor integration.

Poster Session / 29**Shared I/O components for the ATLAS multi-processing framework**

Authors: Peter Van Gemmeren¹; David Malon¹; Marcin Nowak²; Vakho Tsulaia³

¹ Argonne National Laboratory (US)

² Brookhaven National Laboratory (US)

³ Lawrence Berkeley National Lab. (US)

Corresponding Authors: malon@anl.gov, peter.van.gemmeren@cern.ch, vakhtang.tsulaia@cern.ch, marcin.nowak@cern.ch

ATLAS uses its multi-processing framework AthenaMP for an increasing number of workflows, including simulation, reconstruction and event data filtering (derivation). After serial initialization, AthenaMP forks worker processes that then process events in parallel, with each worker reading data individually and producing its own output. This mode, however, has inefficiencies: 1) The worker no longer reads events sequentially, which negatively affects data caching strategies at the storage backend. 2) For its non-RAW data ATLAS uses ROOT and compresses across 10-100 events. Workers will only need a subsample of these events, but have to read and decompress the complete buffers. 3) Output files from the individual workers need to be merged in a separate, serial process. 4) Propagating metadata describing the complete event sample through several workers is nontrivial.

To address these shortcomings, ATLAS has developed shared reader and writer components presented in this paper. With the shared reader, a single process reads the data and provides objects to the workers on demand via shared memory. The shared writer uses the same mechanism to collect output objects from the workers and write them to disk. Disk I/O and compression / decompression of data are therefore localized only in these components, event access (by the shared reader) remains sequential and a single output file is produced without merging. Still for object data, which can only be passed between processes as serialized buffers, the efficiency gains depend upon the storage backend functionality.

Track 2: Data Analysis - Algorithms and Tools / 30

Background Suppression with the Belle II Neural Network Trigger

Authors: Sebastian Skambraks¹; Sara Neuhaus^{None}; Christian Kiesling²

¹ *Technische Universität München*

² *Werner-Heisenberg-Institut*

Corresponding Authors: sebastian.skambraks@tum.de, sneuhaus@mpp.mpg.de, cmk@mppmu.mpg.de

Neural networks are going to be used in the pipelined first level trigger of the upgraded flavor physics experiment Belle II at the high luminosity B factory SuperKEKB in Tsukuba, Japan. A luminosity of $\mathcal{L} = 8 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$ is anticipated, 40 times larger than the world record reached with the predecessor KEKB. Background tracks, with vertices displaced along the beamline (z -axis), are expected to be severely increased due to the high luminosity. Using input from the central drift chamber, the main tracking device of Belle II, the online neural network trigger provides 3D track reconstruction within the fixed latency of the first level trigger. In particular, the robust estimation of the z -vertices allows a significantly improved suppression of the machine background. Based on a Monte Carlo background simulation, the high event rate faced by the first level trigger is analyzed and the benefits of the neural network trigger are evaluated.

Track 1: Computing Technology for Physics Research / 31

Event Display in JUNO Experiment

Authors: Jiang Zhu¹; Zhengyun You²; Yumei Zhang²

¹ *Sun Yat-Sen University*

² *Sun Yat-sen University*

Corresponding Authors: zhuj38@mail2.sysu.edu.cn, youzhy5@mail.sysu.edu.cn, zhangym26@mail.sysu.edu.cn

The current event display module is based on the ROOT EVE package in Jiangmen Underground Neutrino Observatory (JUNO). we use Unity, a multiplatform game engine, to improve its performance and make it available in different platform. Compared with ROOT, Unity can give a more

vivid demonstration of high energy physics experiments and it can be transplanted into another platform easily. We build a tool for event display in JUNO with Unity. It provides us an intuitive way to observe the detector model, the particle trajectory and the hit time distribution. We also build a demo of Circular Electron-Positron Collider(CEPC), for detector optimization and visualization purpose.

Poster Session / 32

Muon $g-2$ reconstruction and analysis framework

Author: Kim Siang Khaw¹

¹ University of Washington

Corresponding Author: khaw84@uw.edu

The Muon $g-2$ experiment at Fermilab will begin beam and detector commissioning in summer 2017 to measure the muon anomalous magnetic moment to an unprecedented level of 140 ppb. To deal with incoming data projected to be around tens of petabytes, a robust data reconstruction and analysis framework, built on Fermilab's *art* event-processing framework, is developed. In this workshop, we report the current status of the framework, together with its novelty features such as multi-threaded reconstruction chain for fast-turnaround operation (nearline) and online data quality monitor (DQM) based on *art*, MIDAS, ZeroMQ, and Node.js. We will also discuss the performance of the framework during the commissioning run.

Track 2: Data Analysis - Algorithms and Tools / 33

Muon reconstruction of JUNO experiment with convolutional neural networks

Authors: Lu Wang¹; Kun Zhang²; Miao HE²; Tao Lin²; Weidong Li²

¹ Computing Center, Institute of High Energy Physics, CAS

² Institute of High Energy Physics, CAS

Corresponding Authors: lintao@ihep.ac.cn, wanglu@ihep.ac.cn, zhangkun@ihep.ac.cn, hem@ihep.ac.cn, liwd@ihep.ac.cn

The Jiangmen Underground Neutrino Observatory (JUNO) is a multiple purpose neutrino experiment to determine neutrino mass hierarchy and precisely measure oscillation parameters. The experimental site is under a 286m mountain, and the detector will be at -480m depth. Twenty thousand ton liquid scintillator (LS) is contained in a spherical container of radius of 17.7 m as the central detector (CD). The light emitted by the LS is watched by about 17,000 20-inch PMTs.

For such a large LS detector, the rate of cosmic muons reaching the inner detector is about 3 Hz. The muon induced background is one of the main backgrounds for the experiment. The most effective approach to reject this background is to define a sufficient detector volume along the muon trajectory and then to veto the events with their vertexes lying in the defined region within a time window. Precise reconstruction of the muon track can reduce unnecessary vetoes and therefore improve the efficiency of the neutrino detection. The traditional reconstruction methods based on the theoretical optical model can only leverage the first-hit-time (FHT) signals of PMTs, and meanwhile it is very difficult to model impacts such as reflection and refraction of optical photons, latency of the light scintillation, and time resolution of the PMT. Additional corrections to the FHT bias are sometimes necessary for these methods.

In this paper, we propose a novel approach of muon reconstruction with convolutional neural networks (CNNs). The main idea is to treat the CD as a 2D image and PMTs as pixels of the image, then

use methods of object detection in computer vision to predict parameters of the muon trajectory. This method can leverage both the charge quantity and time signals of PMTs and bypass the thorny task of optical modeling of the CD. Preliminary results show that by using a 5-layers CNN model (3 convolutional layers and 2 fully connected layers) trained with 50k MC events, we achieved a slightly better performance compared to the traditional method. With 10K testing MC events, the mean error of injection angle is ~ 0.5 degree and the mean error of injecting point is ~ 8 cm. We will further present the improvements by increasing the complexity of CNN models, enlarging the training dataset and optimizing the PMT arrangement in the 2D image.

Poster Session / 34

A Conditions Data Management System for HEP Experiments

Authors: Paul James Laycock¹; Dave Dykstra²; Andrea Formica³; Giacomo Govi²; Andreas Pfeiffer¹; Shaun Roe¹; Roland Sipos⁴

¹ CERN

² Fermi National Accelerator Lab. (US)

³ CEA/IRFU, Centre d'étude de Saclay Gif-sur-Yvette (FR)

⁴ Eotvos Lorand University (HU)

Corresponding Authors: roland.sipos@cern.ch, andrea.formica@cern.ch, andreas.pfeiffer@cern.ch, dwd@fnal.gov, giacomo.govi@cern.ch, shaun.roe@cern.ch, paul.james.laycock@cern.ch

Conditions data infrastructure for both ATLAS and CMS have to deal with the management of several Terabytes of data. Distributed computing access to this data requires particular care and attention to manage request-rates of up to several tens of kHz. Thanks to the large overlap in use cases and requirements, ATLAS and CMS have worked towards a common solution for conditions data management with the aim of using this design for data-taking in Run 3. In the meantime other experiments, including NA62, have expressed an interest in this cross-experiment initiative. For experiments with a smaller payload volume and complexity, there is particular interest in simplifying the payload storage.

The conditions data management model is implemented in a small set of relational database tables. A prototype access toolkit consisting of an intermediate web server has been implemented, using standard technologies available in the Java community. Access is provided through a set of REST services for which the API has been described in a generic way using standard Open API specifications, implemented in Swagger. Such a solution allows the automatic generation of client code and server stubs and further allows changes in the backend technology transparently. An important advantage of using a REST API for conditions access is the possibility of caching identical URLs, addressing one of the biggest challenges that large distributed computing solutions impose on conditions data access, avoiding direct DB access by means of standard web proxy solutions.

Track 1: Computing Technology for Physics Research / 35

An Online system based (almost exclusively) on off-the-shelf hardware for a trigger-free read-out of the LHCb experiment

Authors: Niko Neufeld¹; Tommaso Colombo¹

¹ CERN

Corresponding Author: niko.neufeld@cern.ch

Abstract:

LHCb has decided to optimise its physics reach by removing the first level hardware trigger for LHC

Run3 and beyond. In addition to requiring fully redesigned front-end electronics this design creates interesting challenges for the data-acquisition and the rest of the Online computing system. Such a system can only be realized within realistic cost using as much off-the-shelf hardware as possible. Relevant technologies evolve very quickly and thus the Online system design is architecture-centered and tries to avoid to depend too much on specific features.

In this paper I will describe the design, the motivations for various choices and the current favored options for the implementation, and the status of the R&D. I will cover the back-end readout, which contains the only non-COTS building block, the event-building, the high-level trigger infrastructure and storage and I will also discuss plans for data-flow and the control-system, which will be put in place to configure, control and monitor the entire hard- and software infrastructure.

Poster Session / 36

Using containers with ATLAS offline software

Authors: Johannes Elmsheuser¹; Marcelo Vogel²; Lukas Alexander Heinrich³; Graeme Stewart⁴

¹ *Brookhaven National Laboratory (US)*

² *Bergische Universitaet Wuppertal (DE)*

³ *New York University (US)*

⁴ *University of Glasgow (GB)*

Corresponding Authors: johannes.elmsheuser@cern.ch, lukas.heinrich@cern.ch, marcelo.vogel@cern.ch, graeme.andrew.stewart@cern.ch

This paper describes the deployment of ATLAS offline software in containers for software development and the use in production jobs on the grid - such as event generation, simulation, reconstruction and physics derivations - and in physics analysis. For this we are using Docker and Singularity which are both lightweight virtualization technologies to encapsulates a piece of software inside a complete file system.

The deployment of offline releases via containers removes the interdependence between the runtime environment needed for job execution and the configuration of a computing site's worker nodes. Once the two are decoupled from each other, sites can upgrade their nodes whenever and however they see fit. Docker or Singularity will provide a uniform runtime environment for the grid. The ATLAS software is distributed to the containers either via the CernVM File System (CVMFS) or with a full standalone installation.

For software development, splitting the build and runtime environment from the development environment allows users to take advantage of many modern code development tools that may not be available in production runtime setups like SLC6. It also frees developers from a dependence on resources like lxplus at CERN and allows any reasonable laptop to be used for ATLAS code development.

We document here a comprehensive comparison of the performance of the different deployment options in different host operating systems, e.g. Ubuntu, OS X and CC7, using minimal Cern Scientific Linux 6 base installations.

Poster Session / 37

ATLAS BigPanDA Monitoring

Authors: Torre Wenaus¹; Siarhei Padolski²; Tatiana Korchuganova³; Alexei Klimentov¹

Co-author: Aleksandr Alekseev³

¹ *Brookhaven National Laboratory (US)*

² BNL

³ National Research Tomsk Polytechnic University (RU)

Corresponding Authors: siarhei.padolski@cern.ch, wenaus@gmail.com, aleksandr.alekseev@cern.ch, alexei.klimentov@cern.ch, tatiana.korchuganova@cern.ch

BigPanDA monitoring is a web based application which provides various processing and representation of the Production and Distributed Analysis (PanDA) system objects states. Analyzing hundreds of millions of computation entities such as an event or a job BigPanDA monitoring builds different scale and levels of abstraction reports in real time mode. Provided information allows users to drill down into the reason of a concrete event failure or observe system bigger picture such as tracking the computation nucleus and satellites performance or the progress of whole production campaign. PanDA system was originally developed for the ATLAS experiment and today effectively managing more than 2 million jobs per day distributed over 170 computing centers worldwide. BigPanDA is its core component commissioned in the middle of 2014 and now is the primary source of information for ATLAS users about state of their computations and the source of decision support information for shifters, operators and managers. In this work we describe evolution of the architecture, current status and plans for development of the BigPanDA monitoring.

Poster Session / 38

Update of BESIII Event Display System

Authors: Shuhui Huang¹; Zhengyun You¹

¹ Sun Yat-sen University

Corresponding Authors: youzhy5@mail.sysu.edu.cn, huangshh28@mail2.sysu.edu.cn

BESIII, the detector of BEPCII accelerator, has accomplished a big upgrade on the endcaps of TOF detector, to make more precise measurement. As a result, BesVis system for event display on BESIII experiment needs to be updated. We used ROOT Geometry package to build up geometrical structure and display system. BesVis system plays a significantly important role in DAQ system, reconstruction algorithm as well as data analysis. By using new BesVis system, we can display fired points inside the new ETOF detector and reconstructed tracks more precisely.

Poster Session / 40

A study on the applicability of Recommender Systems for the Production and Distributed Analysis system PanDA of the ATLAS Experiment

Authors: Mikhail Titov¹; Gergely Zaruba²; Kaushik De²; Alexei Klimentov³; Shantenu Jha⁴

¹ National Research Centre Kurchatov Institute (RU)

² University of Texas at Arlington (US)

³ Brookhaven National Laboratory (US)

⁴ Rutgers University (US)

Corresponding Authors: shantenu.jha@rutgers.edu, kaushik@uta.edu, alexei.klimentov@cern.ch, zaruba@uta.edu, mikhail.titov@cern.ch

Scientific computing has advanced in a way of how to deal with massive amounts of data, since the production capacities have increased significantly for the last decades. Most large science experiments require vast computing and data storage resources in order to provide results or predictions

based on the data obtained. For scientific distributed computing systems with hundreds of petabytes of data and thousands of users it is important to keep track not just of how data is distributed in the system, but also of individual user's interests in the distributed data (reveal implicit interconnection between user and data objects). This however requires the collection and use of specific statistics such as correlations between data distribution, the mechanics of data distribution, and mainly user preferences.

This work focuses on user activities (specifically, data usages) and interests in such a distributed computing system, namely PanDA (Production ANd Distributed Analysis system). PanDA is a high-performance workload management system originally designed to meet production and analyses requirements for a data-driven workload at the Large Hadron Collider Computing Grid for the ATLAS Experiment hosted at CERN (the European Organization for Nuclear Research). In this work we are going to investigate whether data collection that was gathered in the past in PanDA shows any trends indicating that users could have mutual interests that would be kept for the next data usages (i.e., data usage patterns), with using data mining techniques such as association analysis, sequential pattern mining, and basics of the recommender system approach. We will show that such common interests between users indeed exist and thus could be used to provide recommendations (in terms of the collaborative filtering) to help users with their data selection process.

Track 1: Computing Technology for Physics Research / 41

Toward real-time data query systems in HEP

Authors: Jim Pivarski¹; Peter Elmer²

¹ *Princeton University*

² *Princeton University (US)*

Corresponding Authors: peter.elmer@cern.ch, pivarski@fnal.gov

Exploratory data analysis must have a fast response time, and some query systems used in industry (such as Impala, Kudu, Dremel, Drill, and Ibis) respond to queries about large (petabyte) datasets on a human timescale (seconds). Introducing similar systems to HEP would greatly simplify physicists' workflows. However, HEP data are most naturally expressed as objects, not tables. In particular, analysis-ready data consists of arbitrary-length lists of particles, possibly containing nested lists of detector measurements. Manipulations of these structures, such as applying quality cuts to particles, not just events, selecting pairs for invariant masses, or matching generator-level data to reconstructed data, are difficult or impossible in SQL-like query languages.

To enable fast querying in HEP, we are developing Femtocode, a language designed for real-time plotting of HEP-scale datasets. We use the same techniques as modern big data query systems, such as performing operations on memory-cached, homogeneous columns of data, rather than each event individually, but adapt them to the scope of manipulations required by HEP. In this talk, I will describe key aspects of the language, how object-oriented expressions are translated into vectorized statements, and how computations are distributed in a fault-tolerant way. I will also show preliminary performance results, which suggest that a thousand-core cluster would be capable of real-time analysis of large LHC datasets. The new capabilities offered by this system may also find application outside of HEP.

This project is being developed in association with FNAL-LDRD-2016-032.

Poster Session / 42

Designing and prototyping the control system for the Cherenkov Telescope Array

Authors: Igor Oya¹; Matthias Fuessling¹; Pablo Oliveira Antonino²; Mauricio Araya³; Jerzy Borkowski⁴; Andrea Bulgarelli⁵; Jaume Castroviejo⁶; Josep Colome⁶; Vito Conforti⁵; Alvaro Garcia⁶; Lars Hagge⁷; Etienne Lyard⁸; Andreas Morgenstern²; Michael Mayer⁹; David Melkumyan¹; Thomas Murach¹; Leonardo Pizarro³; Iftach Sadeh¹; Torsten Schmidt¹; Ullrich Schwanke⁹; Gerrit Spengler⁹; Joseph Schwarz¹⁰; Gino Tosti¹¹; Roland Walter⁸; Peter Wegner¹²

¹ *Deutsches Elektronen-Synchrotron (DESY)*

² *Fraunhofer IESE*

³ *CCTVal / Universidad Técnica Federico Santa María*

⁴ *N. Copernicus Astronomical Center*

⁵ *INAF/IASF Bologna*

⁶ *Institut de Ciències de l'Espai (IEEC-CSIC)*

⁷ *DESY*

⁸ *ISDC Data Centre for Astrophysics, Observatory of Geneva, University of Geneva*

⁹ *Institut für Physik, Humboldt-Universität zu Berlin*

¹⁰ *INAF*

¹¹ *Università di Perugia, Dip. Fisica*

¹² *Deutsches Elektronen-Synchrotron (DESY)*

Corresponding Authors: roland.walter@unige.ch, igor.oya.vallejo@desy.de, imsadeh@gmail.com, bulgarelli@iasfbo.inaf.it, lars.hagge@desy.de, peter.wegner@cern.ch, schwanke@physik.hu-berlin.de, gino2007@gmail.com, etienne.lyard@unige.ch, jubork@ncac.torun.pl

The Cherenkov Telescope Array (CTA) is the next-generation atmospheric Cherenkov gamma-ray observatory. CTA will consist of two installations, one in the southern (Cerro Armazones Chile) and the other in the northern hemisphere (La Palma, Spain). The two sites will contain dozens of telescopes of different sizes, constituting one of the largest astronomical installation under development. The CTA observatory will implement simultaneous automatic operation of multiple sub-arrays. It will be capable of quick re-scheduling of observations (within a few seconds), in order to allow observations of elusive transient events. The array control and data acquisition (ACTL) team within the CTA project is designing and prototyping the software to execute the observations and to handle the acquisition of scientific data at GB/s rates. The operation, control, and monitoring of the distributed multi-telescope CTA arrays is inherently complex. As such, they pose new challenges in scientific instrumentation control systems and in particular in the context of gamma-ray astronomy. In this contribution we present an update on the ACTL system design as it is being modeled via a tailored approach using elements from both the Unified Modelling Language (UML) and the Systems Modeling Language (SysML) formalisms. In addition, we present the status of the associated prototyping activities.

Track 1: Computing Technology for Physics Research / 45

Cross-architecture Kalman filter benchmarks on modern hardware platforms

Authors: Daniel Hugo Campora Perez¹; Omar Awile²; Othman Bouizi³

Co-author: Niko Neufeld²

¹ *Universidad de Sevilla (ES)*

² *CERN*

³ *Intel*

Corresponding Authors: othman.bouizi@intel.com, daniel.hugo.campora.perez@cern.ch, omar.awile@cern.ch

The 2020 upgrade of the LHCb detector will vastly increase the rate of collisions the Online system needs to process in software in order to filter events in real time. 30 million collisions per second will pass through a selection chain where each step is executed conditional to its prior acceptance.

The Kalman filter is a process of the event reconstruction that, due to its time characteristics and early execution in the selection chain, consumes 40% of the whole reconstruction time in the current trigger software. This makes it a time-critical component as the LHCb trigger evolves into a full software trigger in the Upgrade.

The algorithm Cross Kalman allows execution and performance tests across a variety of architectures, including multi and many core platforms and has been successfully integrated and validated in the LHCb codebase. Since its inception, new hardware architectures have become available exposing features that require fine-grained tuning in order to fully utilize their resources.

In this paper we present performance benchmarks and explore the Intel Skylake and latest generation Intel Xeon Phi architectures in depth. We determine the performance gain over previous architectures and show that the efficiency of our implementation is close to the maximum attainable given the mathematical formulation of our problem.

Track 2: Data Analysis - Algorithms and Tools / 46

Tomographic Reconstruction of LArTPC Events using Wire-cell

Author: Chao Zhang¹

¹ *Brookhaven National Laboratory*

Corresponding Author: chao@bnl.gov

The Liquid Argon Time Projection Chamber (LArTPC) is an exciting detector technology that is undergoing rapid development. Due to its high density, low diffusion, and excellent time and spatial resolutions, the LArTPC is particularly attractive for applications in neutrino physics and nucleon decay, and is chosen as the detector technology for the future Deep Underground Neutrino Experiment (DUNE). However, the event reconstruction in LArTPC is challenging due to the intrinsic degeneracy introduced by the wire-readout scheme, typically enforced to satisfy the power dissipation constraint presented by implementing electronics inside liquid argon. In this talk, we present a new 3D reconstruction method “Well-cell” based on the physics principle of charge conservation among wire planes. It then transfers the problem into a set of linear equations closely resembling the concept of tomography. We will further discuss the technique of compressed sensing with L1 regularization to efficiently reconstructing the sparse images. Finally, we will show a web-based event display developed with the WebGL and modern Javascript technologies for interactive 3D visualization.

Track 2: Data Analysis - Algorithms and Tools / 49

Deep-learning in jet reconstruction at CMS

Author: Markus Stoye¹

¹ *CERN*

Corresponding Author: markus.stoye@cern.ch

Deep learning for jet-tagging and jet calibration have recently been increasingly explored. For jet-flavor tagging CMS's most performant tagger for 2016 data (DeepCSV) was based on a deep neural network. The input was a set of standard tagging variables of pre-selected objects. For 2017 improved algorithms are implemented that start from particle candidates without much preselection, i.e. much more raw and unfiltered data. Significantly better tagging is achieved especially in the boosted regime of high transverse momentum. The presentation will include flavor tagging and

further newest public results on deep-learning for jet tagging/calibration. The presenter will discuss the neural network structures used that capture the structure of CMS jet-data and lead to the performance boost.

Track 2: Data Analysis - Algorithms and Tools / 50

CMS Analysis and Data Reduction with Apache Spark

Authors: Collaboration CMS^{None}; Oliver Gutsche¹

¹ *Fermi National Accelerator Lab. (US)*

Corresponding Authors: oliver.gutsche@cern.ch, silvano.tosi@cern.ch

Experimental Particle Physics has been at the forefront of analyzing the world's largest datasets for decades. The HEP community was among the first to develop suitable software and computing tools for this task. In recent times, new toolkits and systems for distributed data processing, collectively called "Big Data" technologies have emerged from industry and open source projects to support the analysis of Petabyte and Exabyte datasets in industry. While the principles of data analysis in HEP have not changed (filtering and transforming experiment-specific data formats), these new technologies use different approaches and tools, promising a fresh look at analysis of very large datasets that could potentially reduce the time-to-physics with increased interactivity. Moreover these new tools are typically actively developed by large communities, often profiting of industry resources, and under open source licensing. These factors result in a boost for adoption and maturity of the tools and for the communities supporting them, at the same time helping in reducing the cost of ownership for the end-users. In this talk, we are presenting studies of using Apache Spark for end user data analysis. We are studying the HEP analysis workflow separated into two thrusts: the reduction of centrally produced experiment datasets and the end-analysis up to the publication plot. Studying the first thrust, CMS is working together with CERN openlab and Intel on the CMS Big Data Reduction Facility. The goal is to reduce 1 PB of official CMS data to 1 TB of ntuple output for analysis. We are presenting the progress of this 2-year project with first results of scaling up Spark-based HEP analysis. Studying the second thrust, we are presenting studies on using Apache Spark for a CMS Dark Matter physics search, investigating Spark's feasibility, usability and performance compared to the traditional ROOT-based analysis.

Track 2: Data Analysis - Algorithms and Tools / 51

Track seeding in the Outer Tracker of CMS for HL-LHC

Author: Erica Brondolin¹

Co-authors: Rudolf Fruhwirth¹; Nikolaus Frohner²

¹ *Austrian Academy of Sciences (AT)*

² *HEPHY*

Corresponding Authors: erica.brondolin@cern.ch, e0526087@student.tuwien.ac.at, rudolf.fruhwirth@oeaw.ac.at

The CMS experiment is in the process of designing a complete new tracker for the high-luminosity phase of LHC. The latest results of the future tracking performance of CMS will be shown as well as the latest developments exploiting the new outer tracker possibilities. In fact, in order to allow for a track trigger, the modules of the new outer tracker will produce stubs or vector hits containing both position and direction information. In this contribution we present algorithms for finding track seeds in the outer tracker without using any information from the pixel tracker. This is particularly important for finding tracks from displaced vertices, but also helps to mitigate the effects of missing pixel hits. We compare the performance of a simple combinatorial search with various clustering

methods employing multi-layer perceptrons and recurrent neural networks, both in terms of efficiency and computational cost. We also present results from neural networks trained to reduce the combinatorics by pre-filtering of vector hits.

Track 1: Computing Technology for Physics Research / 52

Use of GPUs at the CMS High-Level Trigger during Phase-1

Author: Felice Pantaleo¹

¹ CERN

Corresponding Author: felice.pantaleo@cern.ch

Starting from 2017, during CMS Phase-I, the increased accelerator luminosity with the consequently increased number of simultaneous proton-proton collisions (pile-up) will pose significant new challenges for the CMS experiment.

The primary goal of the HLT is to apply a specific set of physics selection algorithms and to accept the events with the most interesting physics content. To cope with the incoming event rate, the online reconstruction of a single event for the HLT has to be done within 220ms on average.

The increasing complexity of events will make track reconstruction especially challenging. For this reason, reconstruction of Pixel Tracks is not executed for every event or is executed in ROIs.

The quest of retaining those events which are potentially interesting for searches of new physics phenomena, led to the evaluation of GPUs for the enhancement of the existing computing infrastructure used at High-Level Trigger (HLT).

We will show the results of the effort in reducing the effect of pile-up in CMS Tracking by redesigning the seeding with novel algorithms which are intrinsically parallel and executing these new algorithms on massively parallel architectures.

We will also show how Pixel Tracks can be evaluated globally for every event on GPUs.

Track 1: Computing Technology for Physics Research / 57

Continuous software quality analysis for the ATLAS experiment

Author: Andrew John Washbrook¹

¹ University of Edinburgh (GB)

Corresponding Author: andrew.washbrook@cern.ch

The regular application of software quality tools in large collaborative projects is required to reduce code defects to an acceptable level. If left unchecked the accumulation of defects invariably results in performance degradation at scale and problems with the long-term maintainability of the code. Although software quality tools are effective for identification there remains a non-trivial sociological challenge to resolve defects in a timely manner. This is an ongoing concern for the ATLAS software which has evolved over many years to meet the demands of Monte Carlo simulation, detector reconstruction and data analysis. At present over 3.8 million lines of C++ code (and close to 6 million total lines of code) are maintained by a community of hundreds of developers worldwide. It is therefore preferable to address code defects **before** they are introduced into a widely used software release.

Recent wholesale changes to the ATLAS software infrastructure have provided an ideal opportunity to apply software quality evaluation as an integral part of the new code review workflow. The

results from static code analysis tools - such as *cppcheck* and *Coverity* - can now be inspected by a rota of review shifters as part of a continuous integration (CI) process. Social coding platforms (e.g. *Gitlab*) allow participants in a code review to consider identified defects and to decide upon any action before code changes are accepted into a production release. A complete audit trail on software quality considerations is thus provided for free as part of the review.

The methods employed to incorporate software quality tools into the ATLAS software CI process will be presented. The implementation of a container-based software quality evaluation platform designed to emulate the live infrastructure will be described with a consideration of how continuous software quality analysis can be optimised for large code bases. We will then conclude with a preview of how analytics on test coverage and code activity - useful in steering the prioritisation of defect resolution - can be incorporated into this new workflow.

Track 2: Data Analysis - Algorithms and Tools / 58

Machine Learning Algorithms for b-jet tagging at the ATLAS experiment

Authors: Salvador Marti I Garcia¹; Richard Hawkings²; Andrea Coccaro³; Michela Paganini⁴

¹ *IFIC-Valencia (UV/EG-CSIC)*

² *CERN*

³ *University of Geneva*

⁴ *Yale University (US)*

Corresponding Authors: michela.paganini@cern.ch, martis@ific.uv.es, andrea.coccaro@cern.ch, richard.hawkings@cern.ch

The separation of b-quark initiated jets from those coming from lighter quark flavours (b-tagging) is a fundamental tool for the ATLAS physics program at the CERN Large Hadron Collider. The most powerful b-tagging algorithms combine information from low-level taggers exploiting reconstructed track and vertex information using a multivariate classifier. The potential of modern Machine Learning techniques such as Recurrent Neural Networks and Deep Learning is explored using simulated events, and compared to that achievable from more traditional classifiers such as boosted decision trees.

Poster Session / 61

Tools for Trigger Rate Monitoring at CMS

Authors: Geoffrey Nathan Smith¹; Andrew Wightman¹; Charles Nicholas Mueller¹

¹ *University of Notre Dame (US)*

Corresponding Authors: andrew.steven.wightman@cern.ch, gsmith15@nd.edu, charles.mueller@cern.ch

In 2017, we expect the LHC to deliver an instantaneous luminosity of roughly $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ to the CMS experiment, with about 60 simultaneous proton-proton collisions (pileup) per event. In these challenging conditions, it is important to be able to intelligently monitor the rate at which data is being collected (the trigger rate). It is not enough to simply look at the trigger rate; we need to know if what we are seeing is what we expect. We will present a set of software tools that have been developed to accomplish this. The tools include a real-time component - a script (run in the CMS control room) that monitors the rates of individual triggers during data-taking, and activates an alarm if rates deviate significantly from expectation. Fits are made to previously collected data and extrapolated to higher pileup. The behavior of triggers as a function of pileup is then monitored as data is collected - plots are automatically produced on an hourly basis and uploaded to a web area

for inspection. We will discuss how this same set of tools is also used offline in data certification, as well as in more complex offline analysis of trigger behavior.

Poster Session / 62

Parallel Random Number Generation for SIMD/SIMT

Author: Soon Yung Jun¹

Co-authors: John Apostolakis²; Philippe Canal¹

¹ *Fermi National Accelerator Lab. (US)*

² *CERN*

Corresponding Authors: syjun@fnal.gov, john.apostolakis@cern.ch, pcanal@fnal.gov

Sequences of pseudorandom numbers of high statistical quality and their efficient generation are critical for the use of Monte Carlo simulation in many areas of computational science. As high performance parallel computing systems equipped with wider vector pipelines or many-cores technologies become widely available, a variety of parallel pseudo-random number generators (PRNGs) are being developed for specific hardware architectures such as SIMD or GPU. However, portable libraries of random number services which can be used across different architectures and in hybrid computing models are not commonly available. We report on an initial implementation of library portable across serial CPUs, vector CPUs, accelerators and GPUs which rely on a common source code implementation for robustness and which use efficient implementations for most operations on each category of hardware. Results of preliminary performance evaluation and statistical tests are presented as well.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 64

GQLink: an implementation of Quantized State System (QSS) methods in Geant4

Author: Lucio Santi¹

Co-authors: Soon Yung Jun²; Krzysztof Genser³; Federico Bergero⁴; Victor Daniel Elvira²; Rodrigo Daniel Castro¹

¹ *Universidad de Buenos Aires (AR), ICC-CONICET*

² *Fermi National Accelerator Lab. (US)*

³ *Fermi National Accelerator Laboratory (US)*

⁴ *Universidad Nacional de Rosario (AR), CIFASIS-CONICET*

Corresponding Authors: syjun@fnal.gov, genser@fnal.gov, rcastro@dc.uba.ar, bergero@cifasis-conicet.gov.ar, lsanti@dc.uba.ar, daniel@fnal.gov

Simulation in high energy physics (HEP) requires the numerical solution of ordinary differential equations (ODE) to determine the trajectories of charged particles in a magnetic field when particles move throughout detector volumes. Each crossing of a volume interrupts the underlying numerical method that solves the equations of motion, triggering iterative algorithms to estimate the intersection point within a given accuracy. The computational cost of this procedure can grow significantly depending on the application at hand. Quantized State System (QSS) is a novel family of asynchronous discrete-event driven numerical methods exhibiting attractive features for this type

of problems. QSS offers native dense output (sequences of polynomial segments updated only by accuracy-driven events) and lightweight detection and handling of volume crossings. In previous works we verified the potential of QSS to offer speedups in HEP simulations, in particular in scenarios with heavy volume crossing activity. Yet, our studies were limited to compare two standalone simulation toolkits: Geant4 (with its default Runge-Kutta method) and QSS Solver (with optimized implementations of QSS methods). A salient limitation of this approach is that physics processes were turned off for comparability purposes, restricting the comparisons to simple setups conceived as baselines. In this work we present a proof-of-concept integration of QSS within Geant4, unleashing the capability to evaluate performance in realistic HEP applications. We developed a Geant4 to QSS Link (GQLink), which is an interface for co-simulation that orchestrates robustly and transparently the interaction between QSS Solver and aspects such as geometry definition and physics processes that are kept under the control of Geant4. Results of GQLink for a simple case study to prove the correctness of the method and for a realistic HEP application using the CMS detector will be discussed along with their computing performance.

Poster Session / 66

High-statistics and GPU-accelerated data analysis in IceCube

Author: Philipp Eller¹

¹ Penn State University

Corresponding Author: pde3@psu.edu

The IceCube neutrino observatory is a cubic-kilometer scale ice Cherenkov detector located at the South Pole. The low energy analyses, that are for example used to measure neutrino oscillations, exploit shape differences in very high-statistics datasets. We present newly-developed tools to estimate reliable event rate distributions from limited statistics simulation and very fast algorithms to produce these. We also ported several features to run on GPUs (CUDA) to considerably speed up data analyses and render it possible to run a more sophisticated treatment of statistical and systematic uncertainties. Advancements are also being made in the reconstruction of low energy events, which are intrinsically difficult to deal with due to the weak signals they produce in the detector.

Track 2: Data Analysis - Algorithms and Tools / 67

Deep Neural Networks for Physics Analysis on low-level whole-detector data at the LHC

Authors: Wahid Bhimji¹; Steven Andrew Farrell¹; Thorsten Kurth²; Evan Racah³; Prabhat⁴; Michela Paganini⁵

¹ Lawrence Berkeley National Lab. (US)

² Unknown

³ LBL

⁴ Lawrence Berkeley National Laboratory

⁵ Yale University (US)

Corresponding Authors: michela.paganini@cern.ch, wahid.bhimji@cern.ch, prabhat@lbl.gov, eracah@lbl.gov, thorsten.kurth@cern.ch, steven.farrell@cern.ch

There has been considerable recent activity applying deep convolutional neural nets (CNNs) to data from particle physics experiments. Current approaches on ATLAS/CMS have largely focussed on

a subset of the calorimeter, and for identifying objects or particular particle types. We explore approaches that use the entire calorimeter, combined with track information, for directly conducting physics analyses: i.e. classifying events as known-physics background or new-physics signals.

We use an existing RPV-Supersymmetry analysis as a case study and evaluate different approaches to make whole-detector deep-learning tractable. We explore CNNs and alternative architectures on multi-channel, high-resolution sparse images: applied on GPU and multi-node CPU architectures (including Knights Landing (KNL) Xeon Phi nodes) on the Cori supercomputer at NERSC.

We compare statistical performance of our approaches with both selections on high-level physics variables from the current physics analyses, and shallow classifiers trained on those variables. We also compare time-to-solution performance of CPU (scaling to multiple KNL nodes) and GPU implementations.

Poster Session / 68

Striped Data Server for Scalable Parallel Data Analysis

Authors: Igor Mandrichenko¹; James Pivarski²; Jin Chang¹; Igor Vasilyevich Mandrichenko^{None}

¹ *FNAL*

² *Princeton University*

Corresponding Authors: ivm@cern.ch, pivarski@fnal.gov, ivm@fnal.gov, jchang@fnal.gov

Columnar data representation is known to be an efficient way to store and access data, specifically in cases when the analysis is often done based only on a small fragment of the available data structure. Data representations like Apache Parquet, on the other hand, split data horizontally to allow for easy parallelization of data analysis. Based on the general idea of columnar data storage, working on the LDRD Project FNAL-LDRD-2016-032, we have developed Striped data representation, which, we believe, is better suited to the needs of High Energy Physics data analysis.

Traditional columnar approach allows for efficient analysis of complex data structures. While keeping all the benefits of columnar data representation, striped mechanism goes further by enabling efficient parallelization of computations and flexible distribution of data analysis.

We present simple and efficient striped data representation model based on Numpy arrays and unified API, which have been implemented for a range of different types of physical storage from local file system to distributed no-SQL database. We further demonstrate a Python-based analysis application platform, which leverages the striped data representation.

We have also implemented Striped Data Server (SDS) as a web service, which hides storage implementation details from the end user and exposes data to WAN users via the web service. Such web service can be deployed as a part of the enterprise computing facility or as a cloud service.

We plan to explore SDS as an enterprise scale data analysis platform for High Energy Physics community and hope to expand it to the other areas that require similar high performance analysis with massive datasets. We have been testing this architecture with 2TB CMS dark matter search dataset and plan to expand it to full CMS public dataset, which is close to 10PB in size.

Track 2: Data Analysis - Algorithms and Tools / 69

Machine Learning Photons Separation in the LHCb Calorimeter

Authors: Fedor Ratnikov¹; Victoria Checalina²; Alexander Panin²

¹ *Yandex School of Data Analysis (RU)*

² *Yandex School of Data Analysis*

Corresponding Authors: sayankotor@yandex-team.ru, fedor.ratnikov@gmail.com, jheuristic@yandex-team.ru

Reconstruction and identification in calorimeters of modern High Energy Physics experiments is a complicated task. Solutions are usually driven by a priori knowledge about expected properties of reconstructed objects. Such an approach is also used to distinguish single photons in the electromagnetic calorimeter of the LHCb detector on LHC from overlapping photons produced from high momentum π^0 decays. We studied an alternative solution based on applying machine learning techniques to primary calorimeter information, that are energies collected in individual cells around the energy cluster.

Constructing such a discriminator from “first principles” allowed improve separation performance from 80% to 93%, that means reducing primary photons fake rate by factor of two.

In presentation we discuss different approaches to the problem, architecture of the classifier, its optimization, and compare performance of the ML approach with classical one.

Poster Session / 70

Parallelized JUNO simulation software based on SNI_{PER}

Authors: Tao Lin¹; Jiaheng Zou¹

Co-authors: Weidong Li²; Zyan Deng³; Guofu Cao ; Xingtao Huang⁴; Zhengyun You⁵

¹ *IHEP*

² *IHEP, Beijing*

³ *Institute of High Energy Physics, Beijing*

⁴ *Shandong University*

⁵ *Sun Yat-sen University*

Corresponding Authors: youzhy5@mail.sysu.edu.cn, dengzy@ihep.ac.cn, huangxt@sdu.edu.cn, liwd@ihep.ac.cn, zoujh@ihep.ac.cn, lintao@ihep.ac.cn, caogf@ihep.ac.cn

The Jiangmen Underground Neutrino Observatory (JUNO) is a neutrino experiment to determine neutrino mass hierarchy. It has a central detector used for neutrino detection, which consists of a spherical acrylic vessel containing 20 kt liquid scintillator (LS) and about 18,000 20-inch photomultiplier tubes (PMT) to collect light from LS.

As one of the important parts in JUNO offline software, the single-threaded simulation framework is developed based on SNI_{PER}. It is in charge of physics generator, detector simulation, event mixing and digitization. However Geant4 based detector simulation of such a large detector is time-consuming and challenging. It is necessary to take full advantages of parallel computing to speedup simulation. Starting from version 10.0, Geant4 supports event-level parallelism. Even though based on pthread, it could be extended with other libraries such as Intel TBB and MPI. It is possible to parallelize JUNO simulation framework via integrating Geant4 and SNI_{PER}.

In this proceeding, we present our progress in developing parallelized simulation software. The SNI_{PER} framework can run in sequential mode, Intel TBB mode or other modes. In SNI_{PER}, the SNI_{PER} task component is in charge of event loop, which is like a simplified application manager. Two types of tasks are introduced in the simulation framework, one is global task and another is worker task. The global task will run only once to initialize detector geometry and physics processes before any other tasks spawned. Later it is accessed by other tasks passively. The worker tasks will be spawned after global task is done. In each worker task, a Geant4 run manager is invoked to do real simulation. Therefore the simulation framework and underlying TBB are decoupled. Finally, the software performance of parallelized JUNO simulation software is also presented.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 71**Four-point function in general kinematics through geometrical splitting and reduction**

Author: Andrei Davydychev¹

¹ *Moscow State University and Schlumberger*

Corresponding Author: adavyd@gmail.com

It is shown how the geometrical splitting of N -point Feynman diagrams can be used to simplify the parametric integrals and reduce the number of variables in the occurring functions. As an example, a calculation of the dimensionally-regulated one-loop four-point function in general kinematics is presented.

Poster Session / 72**Solving the Signal-Versus-Background Classification Problem Using Neural Network**

Author: Rui Li¹

¹ *Sun Yat-sen University*

Corresponding Author: lirui47@mail2.sysu.edu.cn

In order to find the rare particles generated from the collisions at high-energy particle colliders, we need to solve the signal-versus-background classification problems. It turns out neural network can be used here to improve the performance without any manually constructed inputs.

This is the content of my oral report:

1. A brief introduction to neural network
2. Detail information about the neural network I used to solve this problem
3. Performance on the benchmark datasets

Poster Session / 74**Machine Learning based global particle identification algorithms at LHCb experiment**

Authors: Denis Derkach¹; Mikhail Hushchyn¹; Tatiana Likhomanenko²; Aleksei Rogozhnikov¹; Fedor Ratnikov¹

¹ *Yandex School of Data Analysis (RU)*

² *National Research Centre Kurchatov Institute (RU), Yandex School of Data Analysis (RU)*

Corresponding Authors: fedor31415@gmail.com, fedor.ratnikov@gmail.com, mikhail.hushchyn@cern.ch, alex.rogozhnikov@cern.ch, tatiana.likhomanenko@cern.ch, denis.derkach@cern.ch

One of the most important aspects of data processing at LHC experiments is the particle identification (PID) algorithm. In LHCb, several different sub-detector systems provide PID information:

the Ring Imaging Cherenkov (RICH) detector, the hadronic and electromagnetic calorimeters, and the muon chambers. To improve charged particle identification, several neural networks including a deep architecture and gradient boosting have been applied to data. These new approaches provide higher identification efficiencies than existing implementations for all charged particle types. It is also necessary to achieve a flat dependency between efficiencies and spectator variables such as particle momentum, in order to reduce systematic uncertainties during later stages of data analysis. For this purpose, “flat” algorithms that guarantee the flatness property for efficiencies have also been developed. This talk presents this new approach based on machine learning and its performance.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 75

NNLO ntuples

Author: Daniel Maitre¹

¹ *IPPP*

Corresponding Author: daniel.maitre@durham.ac.uk

The storage of computationally intensive matrix elements for NLO processes have proven a good solution for high-multiplicity processes. In this presentation I will present the challenges of extending this method to calculation at NLO and offer some ways of alleviating them.

Poster Session / 77

Dynamic sharing of tape drives accessing scientific data

Authors: Enrico Fattibene¹; Alessandro Cavalli^{None}; Andrea Prosperini²; Daniele Cesini³; Vladimir Sapunenko⁴; Lucia Morganti²

¹ *INFN - National Institute for Nuclear Physics*

² *INFN - CNAF*

³ *Universita e INFN, Bologna (IT)*

⁴ *INFN-CNAF (IT)*

Corresponding Authors: daniele.cesini@cnafe.infn.it, lucia.morganti@cnafe.infn.it, enrico.fattibene@cnafe.infn.it, vladimir.sapunenko@cern.ch, andrea.prosperini@cnafe.infn.it, alessandro.cavalli@cnafe.infn.it

The data management infrastructure operated at CNAF, the central computing and storage facility of INFN (Italian Institute for Nuclear Physics), is based on both disk and tape storage resources. About 40 Petabytes of scientific data produced by LHC (Large Hadron Collider at CERN) and other experiments in which INFN is involved are stored on tape. This is the higher latency storage tier within HSM (Hierarchical Storage Management) environment. Writing and reading requests on tape media are satisfied through a set of Oracle-StorageTek T10000D tape drives, shared among different scientific communities. In the next years, the usage of tape drives will become more intense due to the growing amount of scientific data to manage and the trend to increase the reading traffic rate from tape, announced by the main user communities. In order to reduce hardware purchases, a key point is to minimize the inactivity periods of tape drives. In this paper we present a study of drive resources access patterns in case of concurrent requests and a software solution designed to optimize the efficiency of the shared usage of tape drives in our environment.

Track 2: Data Analysis - Algorithms and Tools / 78

Convolutional Neural Network for Track Seed Filtering at the CMS HLT

Author: Adriano Di Florio¹

Co-authors: Maurizio Pierini²; Felice Pantaleo²

¹ *Universita e INFN, Bari (IT)*

² *CERN*

Corresponding Authors: adriano.diflorio@ba.infn.it, maurizio.pierini@cern.ch, felice.pantaleo@cern.ch

Starting with Run II, future development projects for the Large Hadron Collider will constantly bring nominal luminosity increase, with the ultimate goal of reaching a peak luminosity of $5 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for ATLAS and CMS experiments planned for the High Luminosity LHC (HL-LHC) upgrade. 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 the CMS detector and, specifically, for track reconstruction in the Silicon Pixel Tracker.

One of the first steps of the track finding workflow is the creation of track seeds, i.e. compatible pairs of hits from different detector layers, that are subsequently fed to higher level pattern recognition steps. However the set of compatible hit pairs is highly affected by combinatorial background resulting in the next steps of the tracking algorithm to process a significant fraction of fake doublets.

A possible way of reducing this effect is taking into account the shape of the hit pixel cluster to check the compatibility between two hits. To each doublet is attached a collection of two images built with the ADC levels of the pixels forming the hit cluster. Thus the task of fake rejection can be seen as an image classification problem for which Convolutional Neural Networks (CNNs) have been widely proven to provide reliable results.

Track 1: Computing Technology for Physics Research / 79

The LHCb Software and Computing Upgrade towards LHC Run 3

Authors: Stefan Roiser¹; Concezio Bozzi²

¹ *CERN*

² *CERN and INFN Ferrara*

Corresponding Authors: stefan.roiser@cern.ch, bozzi@fe.infn.it

LHCb is planning major changes for its data processing and analysis workflows for LHC Run 3. Removing the hardware trigger, a software only trigger at 30 MHz will reconstruct events using final alignment and calibration information provided during the triggering phase. These changes pose a major strain on the online software framework which needs to improve significantly. The foreseen changes in the area of the core framework include a re-design of the event scheduling, introduction of concurrent processing, optimisations in processor cache accesses and code vectorisation. Further changes in the areas of event model, conditions data and detector description are foreseen. The changes in the data processing workflow will allow an unprecedented amount of signal events to be selected and therefore increase the load on the experiment's simulation needs. Several areas of improvement for fast simulation are currently being investigated together with improvements needed in the area of distributed computing. Finally the amount of data stored needs to be reflected in the analysis computing model where individual user analysis on distributed computing resources will become inefficient. This contribution will give an overview of the status of those activities and future plans in the different areas from the perspective of the LHCb computing project.

Track 2: Data Analysis - Algorithms and Tools / 80**Building a scalable python distribution for HEP data analysis****Author:** David Lange¹¹ *Princeton University (US)***Corresponding Author:** david.lange@cern.ch

There are numerous approaches to building analysis applications across the high-energy physics community. Among them are Python-based, or at least Python-driven, analysis workflows. We aim to ease the adoption of a Python-based analysis toolkit by making it easier for non-expert users to gain access to Python tools for scientific analysis. Experimental software distributions and individual user analysis have quite different requirements. One tends to worry most about stability, usability and reproducibility, while the other usually strives to be fast and nimble. We discuss how we built and now maintain a python distribution built for analysis while satisfying requirements both a large software distribution (in our case, that of CMSSW) and user, or laptop, level analysis. We pursued the integration of tools use in use by the broader data science community as well as HEP developed (e.g., histogrammar, root_numpy) Python packages. We discuss concepts we investigated for package integration and testing, as well as issues we encountered through this process. Distribution and platform support are important topics. We discuss our approach and progress towards a sustainable infrastructure for supporting this Python stack for the CMS user community and for the broader HEP user community.

Track 1: Computing Technology for Physics Research / 81**Last developments of the INFN CNAF Long Term Data Preservation (LTDP) project: the CDF data recover and safekeeping****Author:** Pier Paolo Ricci¹¹ *INFN CNAF***Corresponding Author:** pierpaolo.ricci@cnafe.infn.it

The INFN CNAF Tier-1 has become the Italian national data center for the INFN computing activities since 2005. As one of the reference sites for data storage and computing provider in the High Energy Physics (HEP) community it offers resources to all the four LHC experiments and many other HEP and non-HEP collaborations. The CDF experiment has used the INFN Tier-1 resources for many years and, after the end of data taking in 2011, it faced the challenge to both preserve the large amount of scientific data produced and give the possibility to access and reuse the whole information in the future using the specific computing model. For this reason starting from the end of 2012 the CDF Italian collaboration, together with the INFN CNAF and Fermilab (FNAL), introduced a Long Term Data Preservation (LTDP) project at our Tier-1 with the purpose of preserve and share all the CDF data and the related analysis framework and knowledge. This is particularly challenging since part of the software releases is no longer supported and the amount of data to be preserved is rather large. The first objective of the collaboration was the copy of all the CDF RUN-2 raw data and user level ntuples (about 4 PB) from FNAL to the INFN CNAF tape library backend using a dedicated network link. This task was successfully accomplished during the last years and, in addition, a system for implementing regular integrity check of data has been developed. This system ensures that all the data are completely accessible and it can automatically retrieve an identical copy of problematic or corrupted file from the original dataset at FNAL. The setup of a dedicated software framework which allows users to access and analyze the data with the complete CDF analysis chain was also carried out with detailed users and system administrators documentation for the long-term future. Furthermore a second and more ambitious objective emerged during 2016 with a feasibility study for reading the first CDF RUN-1 dataset now stored as an unique copy in a huge amount (about 4000) of old Exabyte tape cartridges. With the installation of compatible refurbished tape drive autoloaders an initial test bed was completed and the first phase of the Exabyte tapes reading activity started. In the present article, we will illustrate the state of the art of the LTDP project with a particular attention to the

technical solutions adopted in order to store and maintain the CDF data and the analysis framework and to overcome the issues that have arisen during the recent activities. The CDF model could also prove useful for designing new data preservation projects for other experiments or use cases.

Poster Session / 83

Predictive analytics tools to adjust and monitor performance metrics for the ATLAS Production System

Authors: Alexei Klimentov¹; Dmitri Golubkov²; Fernando Harald Barreiro Megino³; Maksim Gubin⁴; Maria Grigoryeva⁵; Mikhail Titov⁶; Misha Borodin⁷; Siarhei Padolski⁸; Tadashi Maeno¹; Tatiana Korchuganova⁴

¹ Brookhaven National Laboratory (US)

² Institute for High Energy Physics (IHEP)-Unknown-Unknown

³ University of Texas at Arlington

⁴ National Research Tomsk Polytechnic University (RU)

⁵ Institute for Theoretical and Experimental Physics (RU)

⁶ National Research Centre Kurchatov Institute (RU)

⁷ University of Iowa (US)

⁸ BNL

Corresponding Authors: maksim.gubin@cern.ch, mikhail.borodin@cern.ch, tatiana.korchuganova@cern.ch, tmaeno@bnl.gov, dmitri.golubkov@cern.ch, siarhei.padolski@cern.ch, maria.grigorieva@cern.ch, mikhail.titov@cern.ch, alexei.klimentov@cern.ch, fernando.harald.barreiro.megino@cern.ch

Every scientific workflow involves an organizational part which purpose is to plan an analysis process thoroughly according to defined schedule, thus to keep work progress efficient. Having such information as an estimation of the processing time or possibility of system outage (abnormal behaviour) will improve the planning process, provide an assistance to monitor system performance and predict its next state.

The ATLAS Production System is an automated scheduling system that is responsible for central production of Monte-Carlo data, highly specialized production for physics groups, as well as data pre-processing and analysis using such facilities as grid infrastructures, clouds and supercomputers. With its next generation (ProdSys2) the processing rate is around 2M tasks per year that is more than 365M jobs per year. ProdSys2 evolves to accommodate a growing number of users and new requirements from the ATLAS Collaboration, physics groups and individual users. ATLAS Distributed Computing in its current state is the aggregation of large and heterogenous facilities, running on the WLCG, academic and commercial clouds, and supercomputers. This cyber-infrastructure presents computing conditions in which contention for resources among high-priority data analysis happens routinely, that might lead to significant workload and data handling interruptions. The lack of the possibility to monitor and predict the behaviour of the analysis process (its duration) and system's state itself caused to focus on design of the built-in situational awareness analytic tools.

Proposed suite of tools aims to estimate completion time (so called "Time To Complete", TTC) for every (production) task (i.e., prediction of the task duration), completion time for a chain of tasks, and to predict the failure state of the system (e.g., based on "abnormal" task processing times). Its implementation is based on Machine Learning methods and techniques, and besides the historical information about finished tasks it uses ProdSys2 job execution information and resources usage state (real-time parameters and metrics to adjust predicted values according to the state of the computing environment).

The WLCG ML R&D project started in 2016. Within the project the first implementation of the TTC Estimator (for production tasks) was developed, and its visualization was integrated into the ProdSys Monitor.

Poster Session / 84**Performance of the ATLAS Tau Trigger in Run 2****Author:** Mariel Pettee¹¹ *Yale University (US)***Corresponding Author:** mariel.pettee@gmail.com

Tau leptons are used in a range of important ATLAS physics analyses, including the measurement of the SM Higgs boson coupling to fermions, searches for Higgs boson partners, and heavy resonances decaying into pairs of tau leptons. Events for these analyses are provided by a number of single and di-tau triggers, as well as triggers that require a tau lepton in combination with other objects.

The luminosity of proton-proton collisions at the LHC during Run 2 exceeds the design of $10^{34} \text{cm}^{-2} \text{s}^{-1}$. Therefore, sophisticated triggering strategies have been developed to maintain reasonably low trigger thresholds. The main developments to allow a large programme of physics analyses with tau leptons include topological selections at the first trigger level, fast tracking algorithms, and improved identification requirements.

The ATLAS tau trigger strategy and its performance during the 2015 and 2016 data taking will be presented. The investigations for further developments for future data-taking period will also be discussed.

Track 2: Data Analysis - Algorithms and Tools / 86**BDTs in the Level 1 Muon Endcap Trigger at CMS****Authors:** Andrew Mathew Carnes¹; Andrew Brinkerhoff¹; Darin Acosta¹; Ivan Kresimir Furic¹; Bobby Scurlock²; Khristian Kotov¹; Wei Shi³; Alexander Madorsky¹¹ *University of Florida (US)*² *University of Florida*³ *Rice University (US)***Corresponding Authors:** andrew.wilson.brinkerhoff@cern.ch, khristian.kotov@cern.ch, carnesandnoble@gmail.com, ws13@rice.edu, ivan.kresimir.furic@cern.ch, bobby.scurlock@cern.ch, darin.acosta@cern.ch, madorsky@phys.ufl.edu

The first implementation of Machine Learning inside a Level 1 trigger system at the LHC is presented. The Endcap Muon Track Finder at CMS uses Boosted Decision Trees to infer the momentum of muons based on 25 variables. All combinations of variables represented by 2^{30} distinct patterns are evaluated using regression BDTs, whose output is stored in 2 GB look-up tables. These BDTs take advantage of complex correlations between variables, the inhomogeneous magnetic field, and non-linear effects to distinguish high momentum signal muons from the overwhelming low-momentum background. The new algorithm reduced the background rate by a factor of two compared to the previous analytic algorithm, with further improvements foreseen.

Poster Session / 87**The ATLAS Trigger Simulation with Legacy Software**

Authors: Catrin Bernius¹; Adam Edward Barton²

¹ *SLAC National Accelerator Laboratory (US)*

² *Lancaster University (GB)*

Corresponding Authors: catrin.bernius@cern.ch, adam.edward.barton@cern.ch

Physics analyses at the LHC which search for rare physics processes or measure Standard Model parameters with high precision require accurate simulations of the detector response and the event selection processes. The accurate simulation of the trigger response is crucial for determination of overall selection efficiencies and signal sensitivities. For the generation and the reconstruction of simulated event data, generally the most recent software releases are used to ensure the best agreement between simulated data and real data. For the simulation of the trigger selection process, however, the same software release with which real data were taken should be ideally used. This requires potentially running with software dating many years back, the so-called legacy software. Therefore having a strategy for running legacy software in a modern environment becomes essential when data simulated for past years start to present a sizeable fraction of the total. The requirements and possibilities for such a simulation scheme within the ATLAS software framework were examined and a proof-of-concept simulation chain has been successfully implemented. One of the greatest challenges was the choice of a data format which promises long term compatibility with old and new software releases. Over the time periods envisaged, data format incompatibilities are also likely to emerge in databases and other external support services. Software availability may become an issue, when e.g. the support for the underlying operating system might stop. The encountered problems and developed solutions will be presented, and proposals for future development will be discussed. Some ideas reach beyond the retrospective trigger simulation scheme in ATLAS as they also touch more generally aspects of data preservation.

Track 2: Data Analysis - Algorithms and Tools / 88

Convolution Neural networks in nucleon Decay Searches in Liquid Argon time projection Chambers

Author: Kevin Wierman¹

¹ *Pacific Northwest National Laboratory*

Corresponding Author: kevin.wierman@pnnl.gov

Liquid argon time projection chambers (LArTPCs) are an innovative technology used in neutrino physics measurements that can also be utilized in establishing lifetimes on several partial lifetimes for proton and neutron decay. Current analyses suffer from low efficiencies and purities that arise from the misidentification of nucleon decay final states as background processes and vice-versa. One solution is to utilize convolutional neural networks (CNNs) to identify decay topologies in LArTPC data. In this study, CNNs are trained on Monte Carlo simulated data, labeled by truth, and then assessed by out-of-sample simulation. Currently running LArTPCs play an instrumental role in establishing the capabilities of this technology. Simultaneously, the next generation tens-of-kilotons flagship LArTPC experiment – one of whose main charges is to search for nucleon decay – is planning on using this technology in the future. We discuss analysis possibilities and further, a potential application of proton decay-sensitive CNN-enabled data acquisition.

Poster Session / 89**Experience with SPLUNK for archiving and visualization of operational data in ATLAS TDAQ system****Author:** Kevin Thomas Bauer¹¹ *University of California Irvine (US)***Corresponding Author:** andrei.kazarov@cern.ch

The ATLAS Trigger and Data Acquisition (TDAQ) is a large, distributed system composed of several thousands interconnected computers and tens of thousands software processes (applications). Applications produce a large amount of operational messages (at the order of $O(10^4)$ messages per second), which need to be reliably stored and delivered to TDAQ operators in a realtime manner, and also be available for post-mortem analysis by experts.

We have selected SPLUNK, a commercial solution by Splunk Inc, as a all-in-one solution for storing different types of operational data in an indexed database, and a web-based framework for searching and presenting the indexed data and for rapid development of user-oriented dashboards accessible in a web browser.

The paper describes capabilities of Splunk framework, use cases, applications and web dashboards developed for facilitating the browsing and searching of TDAQ operational data by TDAQ operators and experts.

Track 1: Computing Technology for Physics Research / 92**Multithreading in the ATLAS High-Level Trigger****Author:** Adam Edward Barton¹¹ *Lancaster University (GB)***Corresponding Author:** adam.edward.barton@cern.ch

Over the next decade of LHC data-taking the instantaneous luminosity will reach up 7.5 times the design value with over 200 interactions per bunch-crossing and will pose unprecedented challenges for the ATLAS trigger system.

With the evolution of the CPU market to many-core systems, both the ATLAS offline reconstruction and High-Level Trigger (HLT) software will have to transition from a multi-process to a multithreaded processing paradigm in order not to exhaust the available physical memory of a typical compute node. The new multithreaded ATLAS software framework, AthenaMT, has been designed from the ground up to support both the offline and online use-cases with the aim to further harmonize the offline and trigger algorithms. The latter is crucial both in terms of maintenance effort and to guarantee the high trigger efficiency and rejection factors needed for the next two decades of data-taking.

We report on an HLT prototype in which the need for HLT specific components has been reduced to a minimum while retaining the key

aspects of trigger functionality including regional reconstruction and early event rejection. We report on the first experience of migrating trigger algorithms to this new framework and present the next steps towards a full implementation of the ATLAS trigger within AthenaMT.

Poster Session / 93

Test Management Framework for the Data Acquisition of the ATLAS Experiment

Author: Kevin Thomas Bauer¹

¹ *University of California Irvine (US)*

Corresponding Author: andrei.kazarov@cern.ch

Data Acquisition (DAQ) of the ATLAS experiment is a large distributed and inhomogeneous system: it consists of thousands of interconnected computers and electronics devices that operate coherently to read out and select relevant physics data. Advanced diagnostics capabilities of the TDAQ control system are a crucial feature which contributes significantly to smooth operation and fast recovery in case of the problems and, finally, to the high efficiency of the whole experiment.

The base layer of the verification and diagnostic functionality is a test management framework. We have developed a flexible test management system that allows the experts to define and configure tests for different components, indicate follow-up actions to test failures and describe inter-dependencies between DAQ or detector elements. This development is based on the experience gained with the previous test system that was used during the first three years of the data taking. We discovered that experts in different domains or of different components of the system must have more flexibility to configure the verification and diagnostic capabilities of the controls framework, such that later it is used in an automated manner.

In this paper we describe the design and implementation of the test management system and also some aspects of its exploitation during the ATLAS data taking in the LHC Run 2.

Track 2: Data Analysis - Algorithms and Tools / 95

The VISPA Internet-Platform in Deep Learning Applications

Authors: Benjamin Fischer¹; Christian Glaser²; Christoph Welling^{None}; David Walz^{None}; Erik Geiser³; Gero Lutz Mueller⁴; Marcel Rieger¹; Martin Erdmann¹; Martin Urban⁵; Ralf Florian Von Cube¹; Robert Fischer¹; Thorben Quast¹

¹ *Rheinisch-Westfaelische Tech. Hoch. (DE)*

² *RWTH Aachen*

³ *III. Physics Institute A - RWTH Aachen University*

⁴ *III. Physikalisches Institut (B)-Rheinisch-Westfaelische Tech. H*

⁵ *RWTH Aachen University*

Corresponding Authors: martin.erdmann@cern.ch, robert.fischer@cern.ch, martin.urban@physik.rwth-aachen.de, marcel.rieger@cern.ch, ralf.florian.von.cube@cern.ch, erik.geiser@rwth-aachen.de, benjamin.fischer@cern.ch, thorben.quast@rwth-aachen.de, gero.lutz.mueller@cern.ch, glaser@physik.rwth-aachen.de, walz@physik.rwth-aachen.de

Latest developments in many research fields indicate that deep learning methods have the potential to significantly improve physics analyses.

They not only enhance the performance of existing algorithms but also pave the way for new measurement techniques that are not possible with conventional methods.

As the computation is highly resource-intensive both dedicated hardware and software are required to obtain results in a reasonable time which poses a substantial entry barrier.

We provide direct access to this technology after a revision of the internet platform VISPA to serve the needs of researches as well as students.

VISPA equips its users with working conditions on remote computing resources comparable to a local computer through a standard web browser.

For providing the required hardware resources for deep learning applications we extend the CPU infrastructure with a GPU cluster consisting of 20 GeForce GTX 1080 cards.

Direct access through VISPA, preinstalled analysis software and a workload management system allowed us on one hand to support more than 100 participants in a workshop on deep learning and in corresponding university classes and on the other hand to achieve significant progress in particle and astroparticle research.

We present the setup of the system and report on the performance and achievements in the above mentioned usecases.

Track 1: Computing Technology for Physics Research / 96

Exploiting Apache Spark platform for CMS computing analytics

Authors: Daniele Bonacorsi¹; Valentin Y Kuznetsov²; Tommaso Boccali³; Marco Meoni³; Luca Menichetti⁴

¹ *University of Bologna*

² *Cornell University (US)*

³ *INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P*

⁴ *CERN*

Corresponding Authors: tommaso.boccali@cern.ch, luca.menichetti@cern.ch, daniele.bonacorsi@bo.infn.it, marco.meoni@cern.ch, valentin.kuznetsov@cern.ch

The CERN IT provides a set of Hadoop clusters featuring more than 5 PB of raw storage. Different open-source user-level tools are installed for analytics purposes. For this reason, since early 2015, the CMS experiment has started to store a large set of computing metadata, including e.g. a massive number of dataset access log.. Several streamers have registered some billions traces from heterogeneous providers. These trace logs represent a valuable yet scarcely investigated set of information that needs to be cleansed, categorized and correlated; in the case of the CMS dataset access information, this work may lead to discover useful patterns to enhance the overall efficiency of the distributed infrastructure in terms of CPU utilization and task completion time. This work presents an evaluation of Apache Spark platform for CMS needs. We demonstrate a few use-cases how to efficiently process metadata information stored on CERN HDFS system in a scalable manner by harnessing a variety of languages of choice. Among them, Scala and Python offer the best approach to CMS use cases for executing extremely I/O intensive queries that leverage in-memory and persistence Spark API as well as assess streamlining predictive models that can learn dataset properties using machine learning approaches.

Poster Session / 97

Event reconstruction and simulation in PandaRoot for the PANDA experiment

Authors: Dominik Steinschaden¹; on behalf of the PANDA Collaboration^{None}

¹ *Stefan Meyer Institute*

Corresponding Author: dominik.steinschaden@oeaw.ac.at

The \bar{P} PANDA experiment, currently under construction at the Facility for Antiproton and Ion Research (FAIR) in Darmstadt, Germany, addresses fundamental questions in hadron and nuclear physics via interactions of antiprotons with a proton or nuclei, e.g. light and charm exotics, multi-strange baryons and hadrons in nuclei. It will be installed at the High Energy Storage Ring (HESR), which will provide an antiproton beam with a momentum range of 1.5 - 15 GeV/c and enables an high average interaction rate on the fixed target of 2×10^7 events/s.

The \bar{P} PANDA experiment adopts a triggerless, continuous data acquisition. The data rate without any suppression will be in the order of 200 GB/s. With an online software-based data selection system a data reduction of a factor 100 - 1000 has to be achieved. This demands a highly advanced online analysis due to the high interaction rate which has to deal also with overlapping event data. Scalability and parallelization of the reconstruction algorithms are therefore a particular focus in the development process. A simulation framework called PandaRoot is used to develop and evaluate different reconstruction algorithms for event building, tracking and particle identification as well as further optimization of the detector performance. In a novel approach PandaRoot is able to run time-based simulations which allows to simulate the continuous data stream and the mixing of events in addition to the standard event-based simulation. It utilizes the common software framework for the future FAIR experiments, FairRoot, which is based on ROOT and Virtual MonteCarlo with Geant3 and Geant4.

This contribution will give an overview about PandaRoot, the requirements on the event reconstruction algorithms and present the status of a reconstruction and tracking algorithm currently under development.

Track 1: Computing Technology for Physics Research / 98

Mastering Opportunistic Computing Resources for HEP

Authors: Matthias Jochen Schnepf¹; Max Fischer²; Manuel Giffels¹; Christoph Heidecker¹; Adreas Heiss¹; Eileen Kuehn¹; Andreas Petzold¹; Gunter Quast¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

² *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

Corresponding Authors: andreas.heiss@kit.edu, max.fischer@cern.ch, matthias.jochen.schnepf@cern.ch, gunter.quast@cern.ch, christoph.heidecker@cern.ch, manuel.giffels@cern.ch, eileen.kuehn@kit.edu, andreas.petzold@cern.ch

As results of the excellent LHC performance in 2016, more data than expected has been recorded leading to a higher demand for computing resources. It is already foreseeable that for the current and upcoming run periods a flat computing budget and the expected technology advance will not be sufficient to meet the future requirements. This results in a growing gap between supplied and demanded resources. Physics is likely to be limited by the available computing and storage resources.

One option to reduce the emerging lack of computing resources is the utilization of opportunistic resources such as local university clusters, public and commercial cloud providers, HPC centers and volunteer computing. However, to use opportunistic resources additional challenges have to be tackled.

The traditional HEP Grid computing approach leads to a complex software framework that has special dependencies in operation system and software requirements, which currently prevents HEP from using these additional resources. To overcome these obstacles the concept of pilot jobs in

combination with virtualization and/or container technology is the way to go. In this case the resource providers only needs to operate the “Infrastructure as a Service”, whereas HEP manages its complex software environment and the on-demand resource allocation. This approach allows us to utilize additional resources in a dynamically fashion on different kind of opportunistic resource providers.

Another challenge that has to be addressed is that not all workflows are suitable for opportunistic resources. For the HEP workflows the deciding factor is mainly the external network usage. To identify suitable workflows that can be outsourced to external resource providers, we propose an online clustering of workflows to identify those with low external network usage. This class of workflows can than be transparently outsourced to opportunistic resources dependent on the local site utilization.

Our approach to master opportunistic resources for the HEP community in Karlsruhe is currently evaluated and refined. Since the general approach is not tailored to HEP, it can be easily adapted by other communities as well.

Track 2: Data Analysis - Algorithms and Tools / 101

The ATLAS Electron and Photon Trigger

Author: Samuel David Jones¹

¹ *University of Sussex (GB)*

Corresponding Author: samuel.david.jones@cern.ch

Electron and photon triggers covering transverse energies from 5 GeV to several TeV are essential for signal selection in a wide variety of ATLAS physics analyses to study Standard Model processes and to search for new phenomena. Final states including leptons and photons had, for example, an important role in the discovery and measurement of the Higgs boson. Dedicated triggers are also used to collect data for calibration, efficiency and fake rate measurements. The ATLAS trigger system is divided in a hardware-based Level-1 trigger and a software-based high-level trigger, both of which were upgraded during the LHC shutdown in preparation for Run-2 operation. To cope with the increasing luminosity and more challenging pile-up conditions at a center-of-mass energy of 13 TeV, the trigger selections at each level are optimized to control the rates and keep efficiencies high. To achieve this goal multivariate analysis techniques are used. The ATLAS electron and photon triggers and their performance with Run 2 data will be presented.

Poster Session / 102

The Simulation Library of the Belle II Software

Authors: Doris Yangsoo Kim¹; Tadeas Bilka²; Alexander Bobrov³; Casarosa Casarosa⁴; Chilikin Chilikin⁵; Torben Ferber⁶; Romulus Godang⁷; Igal Jaegle⁸; Jakub Kandra⁹; Peter Kodys¹⁰; Thomas Kuhr¹¹; Peter Kvasnicka²; Hiroyuki Nakayama¹²; Leo Piilonen¹³; Christian Pulvermacher¹²; Martin Ritter^{None}; Luka Santenlj¹²; Benjamin Schwenker¹⁴; Alexei Sibidanov¹⁵; Yuri Soloviev¹⁶; marko staric^{None}; Timofey Uglov⁵

¹ *Soongsil University*

² *Charles University (CZ)*

³ *Budker Institute of Nuclear Physics (BINP)*

⁴ *Johannes Gutenberg University of Mainz*

⁵ *P.N. Lebedev Physics Institute*

⁶ *University Of British Columbia*

⁷ *University of South Alabama*

⁸ *University of Florida*

⁹ *Charles University in Prague*

¹⁰ *Charles University*

¹¹ *Ludwig Maximilians University Muenchen*

¹² *KEK*

¹³ *Virginia Tech*

¹⁴ *Göttingen University*

¹⁵ *University of Victoria*

¹⁶ *DESY*

Corresponding Authors: dorisykim@ssu.ac.kr, peter.kvasnicka@cern.ch, jakub.kandra@karlov.mff.cuni.cz, tadeas.bilka@gmail.com, luka.santelj@kek.jp, thomas.kuhr@lmu.de, gjaegle@gmail.com, marko.staric@ijs.si, chilikin@lebedev.ru, soloviev@mail.desy.de, godang@southalabama.edu, sibirid@uvic.ca, hiroyuki.nakayama@kek.jp, christian.pulvermacher@kek.jp, alexander.bobrov@cern.ch, casarosa@uni-mainz.de, ferber@physics.ubc.ca, piilonen@vt.edu, peter.kodys@mff.cuni.cz, martin.ritter@lmu.de, uglov@lebedev.ru, benjamin.schwenker@phys.uni-goettingen.de

The Belle II experiment / the SuperKEKB collider at KEK is a next generation B factory. Phase I of the experiment has been just finished, during which extensive beam studies were conducted. The collaboration is preparing for the physics run in 2018 with the full detector setup. The simulation library of the Belle II experiment is based on the Geant4 package. In this talk, we will summarize the various aspects of the Belle II simulation, including geometry, magnetic field, beam background handling, and validation.

Poster Session / 104

Computer simulation on homogeneity testing for weighted data sets used in HEP

Authors: Petr Bour¹; Václav Kůs¹

¹ *FNSPE CTU Prague*

Corresponding Authors: vaclav.kus@jfifi.cvut.cz, petr.bour@jfifi.cvut.cz

We introduce several modifications of classical statistical tests applicable to weighted data sets in order to test homogeneity of weighted and unweighted samples, e.g. Monte Carlo simulations compared to the real data measurements. Specifically, we deal with the Kolmogorov-Smirnov, Anderson-Darling and f-divergence homogeneity tests. The asymptotic approximation of p-value and power of our weighted variants of homogeneity tests are investigated by means of simulation experiments. The simulation is performed for various sample sizes and statistical distributions of samples and weights. Finally, our methods of homogeneity testing are applied to Monte Carlo samples and real DATA sets measured at the particle accelerator Tevatron in Fermilab at DZero experiment originating from top-antitop quark pair production in two decay channels (electron, muon) with 2, 3, or 4+ jets detected. Consequently, the final variable selection for these 6 variants of decay channels is carried out and the resulting subsets chosen from 46 dimensional physical parameters are recommended for further top quark cross section analysis. Our variable selections differ from the ROOT TMVA based selection for the top-antitop (electron and muon) decay channels.

Track 1: Computing Technology for Physics Research / 105

Design and Execution of make-like Distributed Analyses

Authors: Benjamin Fischer¹; Marcel Rieger¹; Martin Erdmann¹; Ralf Florian Von Cube¹; Robert Fischer¹

¹ *Rheinisch-Westfaelische Tech. Hoch. (DE)*

Corresponding Authors: ralf.florian.von.cube@cern.ch, benjamin.fischer@cern.ch, martin.erdmann@cern.ch, robert.fischer@cern.ch, marcel.rieger@cern.ch

In particle physics, workflow management systems are primarily used as tailored solutions in dedicated areas such as Monte Carlo production. However, physicists performing data analyses are usually required to steer their individual workflows manually, which is time-consuming and often leads to undocumented relations between particular workloads. We present a generic analysis design pattern that copes with the sophisticated demands of end-to-end HEP analyses. The approach presents a paradigm shift from executing parts of the analysis to defining the analysis. The clear interface and dependencies between individual workloads then enables a make-like execution.

Our tools allow to specify arbitrary workloads and dependencies between them in a lightweight and scalable structure. Further features are multi-user support, automated dependency resolution and error handling, central scheduling, and status visualization. The WLCG infrastructure is supported including CREAM-CE, DCAP, SRM and GSIFTP. Due to the open structure, additional computing resources, such as local computing clusters or Dropbox storage, can be easily added and supported. Computing jobs execute their payload, which may be any executable or script, in a dedicated software environment. Software packages are installed as required, and input data is retrieved on demand.

The management system is explored by a team performing ttbb and ttH cross section measurements.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 107

Modeling of modifications induced by jets in the relativistic bulk nuclear matter

Authors: Marcin Slodkowski¹; Daniel Kikola¹; Joanna Porter-Sobieraj²; Patryk Marcinkowski¹; Patryk Gawryszewski¹

¹ *Faculty of Physics, Warsaw University of Technology (PL)*

² *Faculty of Mathematics and Information Science, Warsaw University of Technology (PL)*

Corresponding Authors: kikola@if.pw.edu.pl, piro.plock@gmail.com, j.porter@mini.pw.edu.pl, slodkow@if.pw.edu.pl, patgaw09318@gmail.com

This work is focused on the influence of energy deposited by jets in the medium on the behavior of bulk nuclear matter.

In the heavy ion reactions jets are widely used as probes in the study of Quark-Gluon-Plasma (QGP). Modeling using relativistic hydrodynamics with jets perturbation is employed to extract the properties of the QGP. In order to observe a modification of the collective characteristic of the matter, we use our (3+1) relativistic hydrodynamic code and jet energy loss algorithm implemented on the Graphics Cards (GPU).

The program uses 7th order WENO algorithm and the Cartesian coordinate system to provide high spatial resolution and high accuracy in hydrodynamic simulations required to analyze the propagation of jets in the nuclear matter. We present how the propagation of jets in the medium

could affect the measurements of the properties of a strongly interacting nuclear matter.

Track 1: Computing Technology for Physics Research / 109

Supercomputers, Clouds and Grids powered by BigPanDA for Brain studies

Authors: Alexandre Beche¹; Kaushik De²; Fabien Delalondre¹; Alexei Klimentov³; Felix Schürmann¹; Ruslan Mashinistov⁴

¹ *Ecole polytechnique fédérale de Lausanne (CH)*

² *University of Texas at Arlington (US)*

³ *Brookhaven National Laboratory (US)*

⁴ *Russian Academy of Sciences (RU)*

Corresponding Authors: ruslan.mashinistov@cern.ch, fabien.delalondre@epfl.ch, alexei.klimentov@cern.ch, felix.schuermann@epfl.ch, alexandre.beche@epfl.ch, kaushik@uta.edu

The PanDA WMS - Production and Distributed Analysis Workload Management System - has been developed and used by the ATLAS experiment at the LHC (Large Hadron Collider) for all data processing and analysis challenges. BigPanDA is an extension of the PanDA WMS to run ATLAS and non-ATLAS applications on Leadership Class Facilities and supercomputers, as well as traditional grid and cloud resources. The success of the BigPanDA project has drawn attention from other compute intensive sciences such as biology. In 2017, a pilot project was started between BigPanDA and the Blue Brain Project (BBP) of the Ecole Polytechnique Federal de Lausanne (EPFL) located in Lausanne, Switzerland. This proof of concept project is aimed at demonstrating the efficient application of the BigPanDA system to support the complex scientific workflow of the BBP which relies on using a mix of desktop, cluster and supercomputers to reconstruct and simulate accurate models of brain tissue.

In the first phase, the goal of this joint project is to support the execution of BBP software on a variety of distributed computing systems powered by PanDA. The targeted systems for demonstration include: Intel x86-NVIDIA GPU based BBP clusters located in Geneva (47 TFlops) and Lugano (81 TFlops), BBP IBM BlueGene/Q supercomputer (0.78 PFlops and 65 TB of DRAM memory) located in Lugano, the Titan Supercomputer with peak theoretical performance 27 PFlops operated by the Oak Ridge Leadership Computing Facility (OLCF), and Cloud based resources such as Amazon Cloud.

To hide execution complexity and simplify manual tasks by end-users, we developed a web interface to submit, control and monitor user tasks and seamlessly integrated it with the BigPanDA WMS system. The project demonstrated that the software tools and methods for processing large volumes of experimental data, which have been developed initially for experiments at the LHC accelerator, can be successfully applied to other scientific fields.

Track 1: Computing Technology for Physics Research / 110

Virtualization of the PNNL High Energy Physics Computing Infrastructure

Authors: Kevin Fox¹; David Cowley¹; Malachi Schram¹; Evan Felix¹; James Czebotar¹; Smith Gary¹

¹ *PNNL*

Corresponding Authors: gary.smith@pnnl.gov, evan.felix@pnnl.gov, kevin.fox@pnnl.gov, malachi.schram@pnnl.gov, david.cowley@pnnl.gov, james.czebotar@pnnl.gov

At PNNL, we are using cutting edge technologies and techniques to enable the Physics communities we support to produce excellent science. This includes Hardware Virtualization using an on premise OpenStack private Cloud, a Kubernetes and Docker based Container system, and Ceph, the leading Software Defined Storage solution. In this presentation we will discuss how we leverage these technologies, along with industry established Grid software such as DIRAC, gridftp, and condor to provide resources to the collaboration.

Poster Session / 111

Real-time alignment and reconstruction: performance and recent developments at the LHCb experiment

Authors: Agnieszka Dziurda¹; Francesco Polci²; Lucia Grillo³

Co-author: Michael David Sokoloff⁴

¹ CERN

² Centre National de la Recherche Scientifique (FR)

³ Universita & INFN, Milano-Bicocca (IT)

⁴ University of Cincinnati (US)

Corresponding Authors: francesco.polci@cern.ch, agnieszka.dziurda@cern.ch, lucia.grillo@cern.ch

The LHCb detector is a single-arm forward spectrometer, which has been designed for the efficient reconstruction decays of c- and b-hadrons.

LHCb has introduced a novel real-time detector alignment and calibration strategy for LHC Run II. Data collected at the start of the fill are processed in a few minutes and used to update the alignment, while the calibration constants are evaluated for each run. This procedure permits to obtain the same quality of the processed events in the trigger system as in the offline reconstruction. In addition, the larger timing budget available allows to process the events using the best performing reconstruction in the trigger, which fully includes the particle identification selection criteria. This approach greatly increases the efficiency, in particular for the selection of charm and strange hadron decays.

In this talk the strategy and performance are discussed followed by presentation of the recent developments implemented for the 2017 run of data taking. The topic is discussed in terms of operational performance and reconstruction quality.

Track 2: Data Analysis - Algorithms and Tools / 113

Application of deep learning to the analysis for B -> K* Gamma in Belle II

Authors: Jan Strube¹; Malachi Schram²; Samuel Cunliffe¹; Vishnu Abhinav¹; Charles Siegel¹

¹ PNNL

² PNNL

Corresponding Authors: charles.siegel@pnnl.gov, samuel.cunliffe@pnnl.gov, malachi.schram@pnnl.gov, abhinav.vishnu@pnnl.gov, jan.strube@pnnl.gov

The Belle II experiment is expected to start taking data in early 2018. Precision measurements of rare decays are a key part of the Belle II physics program and machine learning algorithms have played an important role in the measurement of small signals in high energy physics over the past several

years. The authors report on the application of deep learning to the analysis of the B to K* Gamma analysis. We report on the implementation using the Machine Learning Toolkit for Extreme Scale (MaTEX) and the deployment on an HPC system.

Poster Session / 114

Opportunistic data locality for HEP analysis workflows

Authors: Günter Quast¹; Manuel Giffels¹; Max Fischer²; Matthias Jochen Schnepf¹; Christoph Heidecker¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

² *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

Corresponding Authors: christoph.heidecker@cern.ch, matthias.jochen.schnepf@cern.ch, max.fischer@cern.ch, manuel.giffels@cern.ch, gunter.quast@cern.ch

The heavily increasing amount of data delivered by current experiments in high energy physics challenge both end users and providers of computing resources. The boosted data rates and the complexity of analyses require huge datasets being processed. Here, short turnaround cycles are absolutely required for an efficient processing rate of analyses. This puts new limits to the provisioning of resources and infrastructure since already existing approaches are difficult to adapt to HEP requirements and workflows.

The CMS group at the KIT has developed a prototype enabling data locality for HEP analysis processing via coordinated caches. This concept successfully solves key issues of data analyses for HEP:

- Caching reduces the limiting factor of data transfers by joining local high performance devices with large background storages.
- Throughput optimization is reached by selecting and allocating critical data within user workflows
- Transparent integration into the batch system including the usage of container technology solves compatibility issues

Since this prototype has sped up user analyses by several factors, but is limited in scope, our focus is to extend this setup to serve a wider range of analyses and a larger amount of resources. Since it is a static setup under own control of hard- and software, new developments focus not only on extending the setup, but also make it flexible for volatile resources like cloud computing. Usually, data storages and computing farms are deployed by different providers, which leads to data delocalization and a strong influence of the interconnection transfer rates. Here, a caching solution combines both systems into a highly performant setup and enables fast processing of throughput dependent analysis workflows.

Track 1: Computing Technology for Physics Research / 115

The management of heterogeneous resources in Belle II

Authors: Malachi Schram¹; Vikas Bansal¹; Antonio Ledesma¹

¹ *PNNL*

Corresponding Authors: malachi.schram@cern.ch, malachi.schram@gmail.com, vikas.bansal@pnnl.gov, malachi.schram@pnnl.gov, antonio.ledesma@pnnl.gov

The Belle II experiment at the SuperKEKB collider in Tsukuba, Japan, will start taking physics data in early 2018 and aims to accumulate 50/ab, or approximately 50 times more data than the Belle

experiment.

The collaboration expects it will manage and process approximately 190 PB of data.

Computing at this scale requires efficient and coordinated use of the compute grids in North America, Asia and Europe and will take advantage of high-speed global networks.

We present the general Belle II computing model, the distributed data management system and the results of recent network data transfer stress tests.

Additionally, we present how U.S. Belle II is using virtualization techniques to augment computing resources by leveraging Leadership Class Facilities (LCFs).

Poster Session / 116

Applications of A-Common-Tracking-Software (ACTS) in the Future Circular Collider Design Study

Author: Valentin Volk1¹

¹ *University of Innsbruck (AT)*

Corresponding Author: valentin.volk1@cern.ch

The ACTS project aims to decouple the experiment-agnostic parts of the well-established ATLAS tracking software into a standalone package. As the first user, the Future Circular Collider (FCC) Design Study based its track reconstruction software on ACTS. In this presentation we describe the usecases and performance of ACTS in the dense tracking environment of the FCC proton-proton (FCC-hh) collider. An optimized tracking geometry description in ACTS is obtained by means of automated conversion from DD4hep, a generic toolkit for HEP detector description. The internal data model and extrapolation engine are the basis for the description and calculation of particle tracks. Reconstruction of the track parameters is done using the ACTS implementation of the Kalman Filter. The high pileup rates foreseen at the FCC-hh collider are prohibitive for algorithms that scale unfavorably with the number of tracks per event. Assuming sufficient time resolution of the track detectors, it is conceivable to mitigate the performance degradation due to pileup by using precise time measurements as a fourth hit coordinate. ACTS is used to study the potential of 4D Tracking for FCC-hh.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 117

Direct numerical computation and its application to the higher-order radiative corrections

Authors: Kiyoshi Kato¹; Elise de Doncker²; Tadashi Ishikawa³; Fukuko YUASA³

¹ *Kogakuin University*

² *Western Michigan University*

³ *KEK*

Corresponding Authors: elise.dedoncker@wmich.edu, kato@cc.kogakuin.ac.jp, ishika@suchix.kek.jp, fukuko.yuasa@kek.jp

The direct computation method(DCM) is developed to calculate the multi-loop amplitude for general masses and external momenta. The ultraviolet divergence is under control in dimensional regularization.

We discuss the following topics in this presentation.

From the last report in ACAT2016, we have extended the applicability of DCM to several scalar multi-loop integrals.

Also it will be shown that some 2-loop self-energy function can be calculated in the series of epsilon though DCM.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 118**Loopedia, a Database for Loop Integrals****Authors:** Thomas Hahn¹; Viktor Papara¹**Co-authors:** Sophia Borowka²; Gudrun Heinrich¹; Stephen Jones¹; Andreas von Manteuffel³; Christian Bogner⁴; Matthias Kerner¹; Erik Panzer⁵¹ *MPIf. Physik*² *CERN*³ *MSU Department of Physics and Astronomy*⁴ *HU Berlin*⁵ *University of Oxford***Corresponding Authors:** erikpanzer@gmail.com, kerner@mpp.mpg.de, hahn@mpp.mpg.de, papara@mpp.mpg.de, bogner@math.hu-berlin.de, sophia.borowka@cern.ch, gudrun@mpp.mpg.de, manteuffel@pa.msu.edu, sjones@mpp.mpg.de

Loopedia is a new database for bibliographic (and other) information on loop integrals. Its bibliometry is orthogonal to that of SPIRES or arXiv in the sense that it admits searching for graph-theoretical objects, e.g. a graph's topology. We hope it will in time be able to answer the query "Find all papers pertaining to graph X ."

Poster Session / 119**CODE-RADE a user centric software delivery system for science****Authors:** Sean Murray¹; bruce becker²¹ *University of Cape Town (ZA)*² *Meraka Institute***Corresponding Authors:** sean.murray@cern.ch, bbecker@csir.co.za

CODE-RADE is a platform for user-driven, continuous integration and delivery of research applications in a distributed environment. Starting with 6 hypotheses describing the problem at hand, we put forward technical and social solutions to these. Combining widely-used and thoroughly-tested tools, we show how it is possible to manage the dependencies and configurations of a wide range of scientific applications, in an almost fully-automated way, via constant integration tools harnessing docker instances and volume storage for the building and storage of the final applications, and delivery into cvmfs.

Due to the complexity and number both of scientific packages as well as computing platforms, delivering these applications to end users has always been a significant challenge through the grid era, and remains so in the cloud era.

The CODE-RADE platform is a means for developing trust between public computing and data infrastructures on the one hand and various developer and scientific communities on the other hand. Predefined integration tests are specified for any new application, allowing the system to be user-driven. This greatly accelerates time-to-production for scientific applications, while reducing the workload for administrators of HPC, grid and cloud installations together with the people maintaining the software. Specific examples will be given for the HPC facility in Cape Town and the distributed grid resources with in South Africa. Finally, we will give some insight into how this platform could be extended to address issues of reproducibility and collaboration in scientific research in Africa.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 120**A development of an accelerator board dedicated for multi-precision arithmetic operations and its application to Feynman loop integrals II****Authors:** Hiroshi DAISAKA¹; Naohito NAKASATO²; Tadashi ISHIKAWA³; Fukuko YUASA³; Keigo NITADORI⁴¹ *Hitotsubashi University*² *University of Aizu*³ *KEK*⁴ *RIKEN/AICS***Corresponding Author:** daisaka@phys.science.hit-u.ac.jp

Evaluation of a wide variety of Feynman diagrams with multi-loop integrals and physical parameters, and its comparison with high energy experiments are expected to investigate new physics beyond the Standard Model. We have been developing a direct computation method (DCM) of multi-loop integrals of Feynman diagrams. One of the features of our method is that we adopt double exponential (DE) rule for numerical integrations which enables us to evaluate the loop integral with boundary singularities. Other feature is that in order to accelerate the numerical integrations with multi-precision calculations, we develop an accelerator system with Field Programmable Gate Array (FPGA) boards on which processing elements (PE) with dedicated logic for quadruple/hexuple/octuple precision arithmetic operations are implemented. We presented in ACAT 2014 performance results of the small system consists of 4 FPGA boards, and its usability by performing numerical integration of two-loop planar box diagrams. Here we present details of implementation of the dedicated logic on FPGA, our development environment designed for easy use of the system, and the current system consists of 64 FPGA boards. We also present numerical results of higher-loop diagrams performed on our system.

Poster Session / 121**Performance studies of GooFit on GPUs versus RooFit on CPUs while estimating the global statistical significance of a new physical signal****Author:** Alexis Pompili¹¹ *Universita e INFN, Bari (IT)***Corresponding Author:** alexis.pompili@cern.ch

Graphical Processing Units (GPUs) represent one of the most sophisticated and versatile parallel computing architectures available that are nowadays entering the High Energy Physics field. GooFit is an open source tool interfacing ROOT/RooFit to the CUDA platform on nVidia GPUs (it also supports OpenMP). Specifically it acts as an interface between the MINUIT minimization algorithm and a parallel processor which allows a Probability Density Function (PDF) to be evaluated in parallel.

In order to test the computing capabilities of GPUs with respect to traditional CPU cores, a high-statistics pseudo-experiment technique has been implemented both in ROOT/RooFit and GooFit frameworks with the purpose of estimating the local statistical significance of the structure observed by CMS close to the kinematical threshold of the J/ψ ϕ invariant mass in the B^+ to J/ψ ϕ K^+ decay. As already shown in ACAT2016, the optimized GooFit application running on GPUs provides striking speed-up performances with respect to the RooFit application parallelised on multiple CPU workers through the PROOF-Lite tool.

Now the described technique has been extended to situations when, dealing with an unexpected signal, a global significance must be estimated. The LEE is taken into account by means of a scanning technique in order to consider - within the same background-only fluctuation and everywhere in the relevant mass spectrum - any fluctuating peaking behavior with respect to the background model. The execution time of the fitting procedure for each MC toy considerably increases, thus the RooFit-based approach is not only time-expensive but gets unreliable and the use of GooFit as a reliable tool is mandatory to carry out this p-value estimation method.

Track 1: Computing Technology for Physics Research / 122

Design and implementation of data cache and access system across remote sites

Author: Yaodong CHENG¹

¹ *IHEP, Beijing*

Corresponding Authors: yaodong.cheng@cern.ch, ycheng@cern.ch

Distributed computing system is widely used in high energy physics such as WLCG. Computing job is usually scheduled to the site where the input data was pre-staged in using file transfer system. It will lead to some problems including low CPU utility for some small sites lack of storage capacity. Furthermore, It is not flexible in dynamic cloud computing environment. Virtual machines will be created in different cloud platforms on demand. VM need access data immediately after it is created. Cloud platforms may be located in different places for example commercial clouds such as EC2, or private clouds such as CERNCloud and IHEPCloud. It is not possible to stage in data to all cloud platforms before the VM is created. So we designed and implemented a remote data access system based on streaming and cache mechanism. The goal of the system is to export data in one site to remote sites like the behavior of NFS exporting data from one host to others. The system is called LEAF, which means it is one extension of one site storage system. LEAF system is composed of three components, including storage gateway, cache daemon and client module. Storage gateway is deployed in main site which exports specified data repositories. Data repositories is a list of local directory or file system space such as HDFS or EOS. Cache daemon is deployed in remote site which receives requests from client module and then get data from storage gateway in main site. Data is transferred using high performance HTTP connection supported by tornado web framework. Client module is implemented as a file system based on FUSE. Most of file system semantics are handled in local cache daemon, so LEAF system can achieve better performance than directly mounted file system at remote site. The testbed was deployed at two sites in about 2000KM distance. The test result showed LEAF has about 5 times faster than traditional file system such as EOS. The paper will describe the architecture, key technologies, implementation, use cases and performance evaluation of LEAF system.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 123

High-speed evaluation of loop integrals using lattice rules

Authors: Elise de Doncker¹; Ahmed Almulihi¹; Fukuko Yuasa²

¹ *Western Michigan University*

² *High Energy Accelerator Research Organization (KEK)*

Corresponding Authors: elise.dedoncker@wmich.edu, fukuko.yuasa@kek.jp

Using lattice generators, we implement lattice rules in CUDA for a many-core computation on GPUs. We discuss a high-speed evaluation of loop integrals, based on lattice rules combined with a suitable transformation.

The theoretical background of the method and its capabilities will be outlined. Extensive results have been obtained for various rules and integral dimensions, and for classes of diagrams including 2-loop box and 3-loop self-energy diagrams with massive internal lines. Its application for extrapolation with respect to the dimensional regularization parameter is also tested. The current status of the project will be presented.

Poster Session / 124

Weakly Supervised Classifiers in High Energy Physics

Authors: Francesco Rubbo¹; Ben Nachman²; Lucio Dery³; Ariel Gustavo Schwartzman⁴; Eric Metodiev⁵; Patrick Komiske⁶; Matthew Schwartz^{None}

¹ *SLAC National Accelerator Laboratory (US)*

² *Lawrence Berkeley National Lab. (US)*

³ *Stanford*

⁴ *SLAC National Accelerator Laboratory (US)*

⁵ *MIT*

⁶ *Massachusetts Institute of Technology*

Corresponding Authors: ldery@stanford.edu, ariel.gustavo.schwartzman@cern.ch, benjamin.philip.nachman@cern.ch, metodiev@mit.edu, schwartz@physics.harvard.edu, francesco.rubbo@cern.ch, pkomiske@mit.edu

As machine learning algorithms become increasingly sophisticated to exploit subtle features of the data, they often become more dependent on simulations. This paper presents a new approach called weakly supervised classification in which class proportions are the only input into the machine learning algorithm. Using one of the most challenging binary classification tasks in high energy physics - quark versus gluon tagging - we show that weakly supervised classification can match the performance of fully supervised algorithms. Furthermore, by design, the new algorithm is insensitive to any mis-modeling of discriminating features in the data by the simulation. Weakly supervised classification is a general procedure that can be applied to a wide variety of learning problems to boost performance and robustness when detailed simulations are not reliable or not available.

Track 1: Computing Technology for Physics Research / 125

Parallel computing of SNIper based on Intel TBB

Authors: Jiaheng Zou¹; Tao Lin¹

Co-authors: Teng LI²; Weidong Li¹; Xing-Tao Huang³; Ziyang Deng¹; Guofu Cao¹; Zhengyun You⁴

¹ *IHEP*

² *Shandong University, China*

³ *Shandong University (CN)*

⁴ *Sun Yat-sen University*

Corresponding Authors: youzhy5@mail.sysu.edu.cn, huangxt@sdu.edu.cn, dengzy@ihep.ac.cn, lintao@ihep.ac.cn, zoujh@ihep.ac.cn, caogf@ihep.ac.cn, liteng_shiyang@163.com, weidong.li@ihep.ac.cn

SNIper is a general purpose software framework for high energy physics experiments. During the development, we pay more attention to the requirements of neutrino and cosmic ray experiments.

Now SNI_{PER} has been successfully adopted by JUNO (Jiangmen Underground Neutrino Observatory) and LHAASO (Large High Altitude Air Shower Observatory). It has an important effect on the research and design of these projects.

The detector scale and data amount are very large for both JUNO and LHAASO. A great number of computing resources are necessary. Instead of increasing clock speed, increasing cores number became the trend in recent years in the CPU industry. It is natural to implement parallel computing to improve our data processing efficiency with multi-core CPUs. Intel TBB, as a powerful high level library, emancipates us from trivial and complex details of raw threads.

For we have taken parallel computing into account at the beginning of SNI_{PER}'s design, it is possible to achieve multi-threads in a non-invasive way. More than one SNI_{PER} TopTask instances can coexist without interfering between each other. We can simply map each SNI_{PER} TopTask to a TBB task. Events are transmitted from TBB tasks to SNI_{PER} TopTasks one by one, and all SNI_{PER} TopTasks can be executed concurrently. In this way we implemented an event level parallel computing wrapper for SNI_{PER}. A bright characteristic is that the SNI_{PER} kernel module and Intel TBB are absolutely de-coupled. Threads features are transparent to most users, except a few conventions such as global variables. This approach will significantly reduce the costs of migration from serial to parallel computing. However, it is more complicated for critical resources handling, such as disk I/O and memory management. We will also present some attempts for these services, by which they are also non-invasive to the SNI_{PER} kernel module.

Poster Session / 126

Improving Science Yield for NASA Swift with Automated Planning Technologies.

Author: Aaron Tohuvavohu¹

¹ Penn State University

Corresponding Author: aaron.kholodenko@gmail.com

The Swift Gamma-Ray Burst Explorer is a uniquely capable mission, with three on-board instruments and rapid slewing capabilities. It often serves as a fast-response space observatory for everything from gravitational-wave counterpart searches to cometary science. Swift averages 125 different observations per day, and is consistently over-subscribed, responding to about one-hundred Target of Opportunity (ToO) requests per month from the general astrophysics community, as well as co-pointing and follow-up agreements with many other observatories. Since launch in 2004, the demands put on the spacecraft have grown consistently in terms of number and type of targets as well as schedule complexity. To facilitate this growth, various scheduling tools and helper technologies have been built by the Swift team to continue improving the scientific yield of the Swift mission. In this study, we examine various approaches to the automation of observation schedules for the Swift spacecraft, comparing the efficiency and quality of these tools to each other and to the output of well-trained human Science Planners. Because of the computational complexity of the scheduling task, no automation tool has been able to produce a plan of equal or higher quality than that produced by a well-trained human, given the necessary time constraints. We detail here several approaches towards achieving this goal of surpassing human quality schedules using classical optimization and algorithmic techniques, as well as machine learning and recurrent neural network (RNN) methods. We then quantify the increased scientific yield and benefit to the wider astrophysics community that would result from the further development and adoption of these technologies.

Poster Session / 127

a loosely coupled scalable cloud infrastructure

Author: Tao Cui¹

Co-author: Yaodong Cheng ²

¹ *IHEP(Institute of High Energy Physics, CAS,China)*

² *IHEP*

Corresponding Authors: cuit@ihep.ac.cn, chyd@ihep.ac.cn

Large-scale virtual computing system requires a loosely coupled virtual resource management platform that provides the flexibility to add or subtract physical resources and the Convenience to upgrade the platform and so on. Openstack provides large-scale virtualization solution such as “Cells” and “Tricircle/ Trio2o “. But because of the complexity, it’s difficult to be deployed and maintain for small cloud computing teams. We discuss a loosely coupled cloud cluster infrastructure. It is based on the database and plug-in to achieve information collection, virtual machine scheduling and controlling and uses a simplified method to achieve unified management of network and images. It can flexibly expand or reduce cluster members, deploys different versions of Openstack and even heterogeneous cloud platform. We will show and analysis this infrastructure. The test bed works well in IHEP.

Poster Session / 128

Speeding up experiments software with VecCore, a portable SIMD library

Authors: Guilherme Amadio¹; Philippe Canal²; Sandro Christian Wenzel¹

Co-author: Danilo Piparo ¹

¹ *CERN*

² *Fermi National Accelerator Lab. (US)*

Corresponding Authors: sandro.wenzel@cern.ch, pcanal@fnal.gov, daniло.piparo@cern.ch, amadio@cern.ch

Portable and efficient vectorization is a significant challenge in large software projects such as Geant, ROOT, and experiment frameworks. Nevertheless, taking advantage of the expression of parallelism through vectorization is required by the future evolution of the landscape of particle physics, which will be characterized by a drastic increase in the amount of data produced.

In order to bridge the widening gap between data processing and analysis needs, and available computing resources, the particle physics scientific software stack needs to be upgraded to fully exploit SIMD. While libraries exist that wrap SIMD intrinsics in a convenient way, they don’t always support every available architecture, or perform well only in a subset of them. This situation needs an improvement.

VecCore provides a solution. It features a simple API to express SIMD-enabled algorithms that can be dispatched to one or more backends, such as CUDA, or other widely adopted SIMD libraries such as Vc or UME::SIMD. In this talk we discuss the programming model associated to VecCore, the most relevant details of its implementation, and some use cases in HEP software packages such as ROOT and GeantV. Outlooks on possible usage in experiments’ software are also highlighted.

Performance figures from benchmarks on NVidia GPUs, and on Intel Xeon and Xeon Phi processors are discussed that demonstrate nearly optimal gains from SIMD parallelism.

Poster Session / 129

Optimizing ROOT I/O for Analysis**Authors:** Brian Paul Bockelman¹; Zhe Zhang^{None}**Co-author:** Jim Pivarski²¹ *University of Nebraska-Lincoln (US)*² *Princeton University***Corresponding Authors:** pivarski@fnal.gov, brian.bockelman@cern.ch, zhan0915@huskers.unl.edu

The ROOT I/O (RIO) subsystem is foundational to most HEP experiments - it provides a file format, a set of APIs/semantics, and a reference implementation in C++. It is often found at the base of an experiment's framework and is used to serialize the experiment's data; in the case of an LHC experiment, this may be hundreds of petabytes of files! Individual physicists will further use RIO to perform their end-stage analysis, reading from intermediate files they generate from experiment data.

RIO is thus incredibly flexible: it must serve as a file format for archival (optimized for space) and for working data (optimized for read speed). To date, most of the technical work has focused on improving the former use case. We present work designed to help improve RIO for analysis. We analyze the real-world impact of LZ4 to decrease decompression times (and the corresponding cost in disk space). We introduce new APIs that read RIO data in bulk, removing the per-event overhead of a C++ function call. We compare the performance with the existing RIO APIs for simple structure data and show how this can be complimentary with efforts to improve the parallelism of the RIO stack.

Track 2: Data Analysis - Algorithms and Tools / 130**Machine Learning for Antihydrogen Detection at ALPHA**

Authors: Andrea Capra¹; Mostafa Ahmadi²; Bruno Ximenez Rodrigues Alves³; Marcelo Baquero Ruiz⁴; William Alan Bertsche⁵; Eoin Butler⁶; Celeste Carruth⁴; Claudio Lenz Cesar⁷; Michael Charlton⁸; Snir Cohen⁹; Robert Andrew Collister¹; Stefan Eriksson¹⁰; Andrew Evans¹¹; Nathan Andrew Evetts¹; Joel Fajans⁴; Tim Friesen³; Makoto C. Fujiwara¹; David Russell Gill¹; Andrea Gutierrez¹; Jeffrey Scott Hangst³; Walter Hardy¹; Mike Hayden¹²; Aled Isaac¹⁰; Akira Ishida¹³; Mark Andrew Johnson⁵; Steven Armstrong Jones¹⁰; Svante Jonsell¹⁴; Leonid Kurchaninov¹; Niels Madsen¹⁰; Melissa Rae Mathers¹⁵; Daniel Thomas Maxwell¹⁰; Joseph Mc Kenna¹; Scott Menary¹; Juan Mario Michan¹; Takamasa Momose¹⁶; Justine Joyce Munich¹²; Paul Nolan¹⁷; Konstantin Olchanski¹; Art Olin¹; Petteri Pusa²; Chris Rasmussen³; Francis Robicheaux¹⁸; Rodrigo Lage Sacramento⁷; Muhammed Sameed¹⁰; Elazar Sarid⁹; Daniel De Miranda Silveira⁷; Simone Stracka¹⁹; Chukman So²⁰; Graham Stutter³; Tim Tharp³; James Edward Thompson¹⁵; Robert Thompson¹¹; Dirk Peter Van Der Werf¹⁰; Jonathan Syrkin Wurtele⁴

¹ *TRIUMF (CA)*² *University of Liverpool (GB)*³ *Aarhus University (DK)*⁴ *University of California Berkeley (US)*⁵ *University of Manchester (GB)*⁶ *Imperial College (GB)*⁷ *Federal University of of Rio de Janeiro (BR)*⁸ *Swansea University*⁹ *At.En. Comm., Nucl. Res. Cen. Negev (IL)*¹⁰ *Swansea University (GB)*¹¹ *University of Calgary Dep. of Phys. and Astronomy (CA)*¹² *Simon Fraser University (CA)*

¹³ *University of Tokyo (JP)*

¹⁴ *Stockholm University (SE)*

¹⁵ *York University (CA)*

¹⁶ *University of British Columbia*

¹⁷ *University of Liverpool*

¹⁸ *Purdue University (US)*

¹⁹ *Universita & INFN, Pisa (IT)*

²⁰ *University of Calgary (CA)*

Corresponding Authors: simone.stracka@cern.ch, pjn@ns.ph.liv.ac.uk, tim.tharp@cern.ch, mike.hayden@cern.ch, joseph.mckenna@cern.ch, jonathan.syrkin.wurtele@cern.ch, arthur.olin@cern.ch, menary@yorku.ca, mostafa.ahmadi@cern.ch, bruno.ximenez@cern.ch, chukman.so@cern.ch, petteri.pusa@cern.ch, celeste.carruth@cern.ch, aisaac@cern.ch, claudio.cesar@cern.ch, justine.joyce.munich@cern.ch, robert.ian.thompson@cern.ch, tim.friesen@cern.ch, jeffrey.hangst@cern.ch, esarid@cern.ch, snir.cohen@cern.ch, daniel.thomas.maxwell@cern.ch, m.charlton@swansea.ac.uk, niels.madsen@cern.ch, juan.mario.michan@cern.ch, mark.andrew.johnson@cern.ch, stefan.johannes.eriksson@cern.ch, william.bertsche@cern.ch, m.sameed@cern.ch, robert.andrew.collister@cern.ch, momose@chem.ubc.ca, melissa.rae.mathers@cern.ch, andrea.gutierrez@cern.ch, nathan.andrew.evetts@cern.ch, olchansk@triumf.ca, james.edward.thompson@cern.ch, eoin.butler@cern.ch, dirk.peter.van.der.werf@svante.jonsell@cern.ch, leonid.kurchaninov@cern.ch, david.russell.gill@cern.ch, chris.rasmussen@cern.ch, makoto.fujiwara@cern.ch, daniel.miranda@cern.ch, andrew.evans@cern.ch, andrea.capra@cern.ch, steven.armstrong.jones@cern.ch, graham.stutter@cern.ch, jfajans@berkeley.edu, akira.ishida@cern.ch, francis.robicheaux@cern.ch, walter.hardy@cern.ch, marcelo.baquero.ruiz@cern.ch, rodrigo.lage.sacramento@cern.ch

The ALPHA experiment at CERN is designed to produce, trap and study antihydrogen, which is the antimatter counterpart of the hydrogen atom. Since hydrogen is one of the best studied physical system, both theoretically and experimentally, experiments on antihydrogen permit a precise direct comparison between matter and antimatter. Our basic technique consists of driving an antihydrogen resonance which will cause the antiatom to leave our trap and annihilate. This resonant frequency can be compared with its corresponding value in hydrogen. The antihydrogen annihilation location, called the vertex, is determined by reconstructing the trajectories of the annihilation products and by finding the point where they pass closest to each other. The main background to antihydrogen detection is due to cosmic rays. When an experimental cycle extends for several minutes, while the number of trapped antihydrogen remains fixed, background rejection can become challenging. The use of “cuts-based” analysis is often not sufficient to reach the target statistical significance. Machine learning methods have been employed in ALPHA for several years, leading to a dramatic reduction of the background contamination. Thanks to these techniques, the ALPHA collaboration observed for the first time a transition between Zeeman levels of the antihydrogen ground state 1, placed the most stringent upper limit to the antihydrogen electric charge 2, and performed the first laser spectroscopy experiment 3. These results will be presented along with the optimization of the analysis methods employed in these measurements.

1 C. Amole *et al.*, *Nature* **483**, 439-443 (2012)

2 M. Ahmadi *et al.*, *Nature* **529**, 373-376 (2016)

3 M. Ahmadi *et al.*, *Nature* **541**, 506-510 (2017)

Track 3: Computations in Theoretical Physics: Techniques and Methods / 131

Calculation of One Loop Integrals with Package-X and Collier-Link

Author: Hiren Patel¹

¹ *UMass Amherst*

Corresponding Author: hhpatel@umass.edu

Package-X is a Mathematica package for analytically computing and symbolically manipulating dimensionally regulated one-loop Feynman integrals, and CollierLink is an upcoming interface to the

COLLIER library. In this talk, I will review new features in the upcoming release of Package-X: calculation of cut discontinuities, and command-line readiness. Additionally, features of CollierLink will be discussed: numerical evaluation of Passarino-Veltman functions using the COLLIER library from Mathematica, and automatic code generation and compilation of functions for rapid evaluation of loop integrals. Time permitting, I will emphasize the importance of user friendliness of public packages.

Track 2: Data Analysis - Algorithms and Tools / 132

Novel functional and distributed approaches to data analysis available in ROOT

Authors: Enric Tejedor Saavedra¹; Danilo Piparo¹; Philippe Canal²; Enrico Guiraud³; Xavier Valls Pla⁴; Gerardo Ganis¹; Pere Mato Vila¹; Guilherme Amadio¹; Lorenzo Moneta¹

Co-author: Jakob Blomer ¹

¹ CERN

² Fermi National Accelerator Lab. (US)

³ CERN, University of Oldenburg (DE)

⁴ University Jaume I (ES)

Corresponding Authors: jakob.blomer@cern.ch, xavier.valls.pla@cern.ch, danilo.piparo@cern.ch, enric.tejedor.saavedra@cern.ch, enrico.guiraud@cern.ch, gerardo.ganis@cern.ch, pcanal@fnal.gov, pere.mato@cern.ch, lorenzo.moneta@cern.ch, amadio@cern.ch

The bright future of particle physics at the Energy and Intensity frontiers poses exciting challenges to the scientific software community. The traditional strategies for processing and analysing data are evolving in order to cope with the ever increasing complexity and size of the datasets.

The traditional strategies for processing and analysing data are evolving in order to (i) offer higher-level programming model approaches and (ii) exploit parallelism to cope with the ever increasing complexity and size of the datasets.

This contribution describes how the ROOT framework, a cornerstone of software stacks dedicated to particle physics, is preparing to provide adequate solutions for the analysis of large amount of scientific data on parallel architectures.

The functional approach to parallel data analysis provided with the ROOT TDataFrame interface is then characterised. The design choices behind this new interface are described also comparing with other widely adopted tools such as Pandas and Apache Spark. Commonalities and differences with ReactiveX and Ranges v3 are highlighted.

The programming model is illustrated highlighting the reduction of boilerplate code, composability of the actions and data transformations as well as the capabilities of dealing with different data sources such as ROOT, json, csv or databases.

Details are given about how the functional approach allows transparent implicit parallelisation of the chain of operations specified by the user.

The progress done in the field of distributed analysis is examined. In particular, the power of the integration of ROOT with Apache Spark via the PyROOT interface is shown.

In addition, the building blocks for the expression of parallelism in ROOT are briefly characterised together with the structural changes applied in the building and testing infrastructure which were necessary to put them in production.

All new ROOT features are accompanied by scaling and performance measurements of real life use cases on highly parallel and distributed architectures.

Track 1: Computing Technology for Physics Research / 133**Increasing Parallelism in the ROOT I/O subsystem**

Authors: Brian Paul Bockelman¹; Zhe Zhang^{None}; Philippe Canal²; Enric Tejedor Saavedra³; Danilo Piparo³; Guilherme Amadio³

¹ *University of Nebraska-Lincoln (US)*

² *Fermi National Accelerator Lab. (US)*

³ *CERN*

Corresponding Authors: amadio@cern.ch, danilo.piparo@cern.ch, zhan0915@huskers.unl.edu, brian.bockelman@cern.ch, enric.tejedor.saavedra@cern.ch, pcanal@fnal.gov

When dealing with the processing of large amount of data, the rate at which the reading and writing can take place is a critical factor. High Energy Physics data processing relying on ROOT based persistification is no exception. The recent parallelisation of LHC experiments' software frameworks and the analysis of the ever increasing amount of collision data collected by experiments further emphasised this issue underlying the need of increasing the implicit parallelism expressed within the ROOT I/O.

In this contribution we highlight the improvements of the ROOT I/O subsystem which targeted a satisfactory scaling behaviour in a multithreaded context. The effect of parallelism on the individual steps which are chained by ROOT to read and write data, namely (de)compression, (de)serialisation, access to storage backend, are discussed.

Details relevant for the programming model associated to these innovations are characterised as well as description of the design choices adopted such as the streamlining of the asynchronous operations via a task based approach relying on the same engine exploited by experiments to guarantee parallel execution, the Intel TBB library. Measurements of the benefit of the aforementioned advancements are discussed through real life examples coming from the set of CMS production workflows on traditional server platforms and highly parallel architectures such as Xeon Phi.

Track 2: Data Analysis - Algorithms and Tools / 134**Speeding up prediction performance of the boosting decision trees-based learning models.**

Authors: Egor Khairullin¹; Andrey Ustyuzhanin²

¹ *Moscow Institute of Physics and Technology, Yandex School of Data Analysis (RU)*

² *Yandex School of Data Analysis (RU)*

Corresponding Authors: mikari@yandex-team.ru, andrey.ustyuzhanin@cern.ch

The result of many machine learning algorithms are computational complex models. And further growth in the quality of the such models usually leads to a deterioration in the applying times. However, such high quality models are desirable to be used in the conditions of limited resources (memory or cpu time).

This article discusses how to trade the quality of the model for the speed of its applying a novel boosted trees algorithm called Catboost. The idea is to combine two approaches: training fewer trees and uniting trees into huge cubes. The proposed method allows for pareto-optimal reduction of the computational complexity of the decision tree model with regard to the quality of the model. In the considered example number of lookups was decreased from 5000 to only 6 (speedup factor of 1000) while AUC score of the model was reduced by less than per mil.

Track 2: Data Analysis - Algorithms and Tools / 135**Learning to Remove Pileup at the LHC with Jet Images****Authors:** Eric Metodiev¹; Patrick Komiske¹; Ben Nachman²; Matthew Schwartz³¹ *Massachusetts Institute of Technology*² *Lawrence Berkeley National Lab. (US)*³ *Harvard University***Corresponding Authors:** benjamin.philip.nachman@cern.ch, schwartz@physics.harvard.edu, pkomiske@mit.edu, metodiev@mit.edu

Many simultaneous proton-proton collisions occur in each bunch crossing at the Large Hadron Collider (LHC). However, most of the time only one of these collisions is interesting and the rest are a source of noise (pileup). Several recent pileup mitigation techniques are able to significantly reduce the impact of pileup on a wide set of interesting observables. Using state-of-the-art machine learning techniques, we develop a new method for pileup mitigation based on the jet images framework. We demonstrate that our algorithm outperforms existing methods on a wide range of simple and complex jet observables up to pileup levels of 140 collisions per bunch crossing. We also investigate what aspects of the event our algorithms are utilizing and also test the robustness of the trained pileup mitigation algorithm.

Track 2: Data Analysis - Algorithms and Tools / 136**Analysis Preservation and Systematic Reinterpretation within the ATLAS Experiment****Authors:** Lukas Alexander Heinrich¹; Kyle Stuart Cranmer¹¹ *New York University (US)***Corresponding Authors:** kyle.cranmer@cern.ch, lukas.heinrich@cern.ch

The LHC data analysis software used in order to derive and publish experimental results is an important asset that is necessary to preserve in order to fully exploit the scientific potential of a given measurement. Among others, important use cases of analysis preservation are the reproducibility of the original results and the reusability of the analysis procedure in the context of new scientific studies. A prominent use-case for the latter is the systematic reinterpretation of searches for new Physics in terms of signal models that not studied in the original publication (RECAST).

This paper presents the usage of the graph-based workflow description language yadage to drive the reinterpretation of preserved HEP analyses. The analysis software for individual states in the analysis is preserved using Docker containers, while the workflow structure is preserved using plain JSON documents. This allows the re-execution of complex analysis workflows on industry standard container-based distributed computing clusters (Kubernetes via OpenStack Magnum)

We present re-interpretations of ATLAS analyses based on both the original ATLAS analysis code and third-party re-implementations such as CheckMATE and integrations with other analysis preservation efforts such as the CERN Analysis Preservation Portal.

Poster Session / 137**Belle II Conditions Database Interface****Authors:** Martin Ritter^{None}; Thomas Kuhr^{None}; Christian Pulvermacher¹

¹ KEK

Corresponding Authors: martin.ritter@lmu.de, thomas.kuhr@lmu.de, christian.pulvermacher@kek.jp

The Belle II experiment at the SuperKEKB e^+e^- accelerator is preparing for taking first collision data next year. For the success of the experiment it is essential to have information about varying conditions available in the simulation, reconstruction, and analysis code.

The online and offline software has to be able to obtain conditions data from the Belle II Conditions Database in a robust and reliable way under very different situations. As such the client side interface to the Belle II Conditions Database has been designed with a variety of access mechanisms which allow the software to be used with and without internet connection. Different methods to access the payload information are implemented to allow for high level of customization per site and to simplify testing of new payloads locally. The framework obtains objects from the back-end database only when needed and caches them during their validity. The mechanism also enables transparent handling of validity ranges which are smaller than the finest granularity supported by the database.

The user API to the conditions data was designed to make the life for developers as easy as possible. Two classes, one for single objects and one for arrays of objects, provide type-safe access. Their interface resembles that of the classes for the access to event-level data with which the developers are already familiar. Changes to the conditions data are usually transparent to the client code but users can actively check whether an object has changed or register callback functions to be called whenever a conditions data object is updated. In addition a command line user interface has been developed to simplify inspection and modification of the database contents.

The talk will present the design of the conditions database interface in the Belle II software, show examples of its application, and report about usage experiences in large-scale Monte Carlo productions and calibration exercises.

Track 1: Computing Technology for Physics Research / 138

FELIX: the new detector readout system for the ATLAS experiment

Corresponding Author: kevin.thomas.bauer@cern.ch

Starting during the upcoming major LHC shutdown from 2019-2021, the ATLAS experiment at CERN will move to the the Front-End Link eXchange (FELIX) system as the interface between the data acquisition system and the trigger and detector front-end electronics. FELIX will function as a router between custom serial links and a commodity switch network, which will use industry standard technologies to communicate with data collection and processing components. The FELIX system is being developed using commercial-off-the-shelf server PC technology in combination with a FPGA-based PCIe Gen3 I/O card hosting GigaBit Transceiver links and with Timing, Trigger and Control connectivity provided by an FMC-based mezzanine card. FELIX functions will be implemented with dedicated firmware for the Xilinx FPGA (Virtex 7 and Kintex UltraScale) installed on the I/O card alongside an interrupt-driven Linux kernel driver and user-space software. On the network side, FELIX is able to connect to both Ethernet or Infiniband network architectures. This presentation will describe the FELIX system design as well as reporting on results of the ongoing development program.

Poster Session / 139

Multi-core supports in JUNO distributed computing

Author: xiaomei zhang¹

¹ *IHEP, Beijing*

Corresponding Author: zhangxm@mail.ihep.ac.cn

The IHEP distributed computing system has been built on DIRAC to integrate heterogeneous resources from collaboration institutes and commercial resource providers for data processing of IHEP experiments, and began to support JUNO in 2015. The Jiangmen Underground Neutrino Observatory (JUNO) is a multipurpose neutrino experiment located in southern China to start in 2019. The study on applying parallel computing in JUNO software is on-going to fasten JUNO data processes, and fully use capability of multi-core and many-core CPUs. Therefore, it is necessary for the IHEP distributed computing system to explore the way to support single-core and multi-core jobs in a consistent way. A series of changes on job descriptions, scheduling and accounting will be considered and discussed. The pilot-based scheduling with mixture of single-core and multi-core jobs is the most complicated part. In the report, two ways of scheduling and their efficiency has been studied, one way using separated pilots for single-core and multi-core jobs, and the other using dynamic partitionable common pilots for both jobs. Their advantages and disadvantages will be discussed.

Track 2: Data Analysis - Algorithms and Tools / 140

Identification of Hadronically Decaying W Bosons and Top Quarks using Multivariate Techniques at ATLAS

Authors: Salvador Marti I Garcia¹; Christophe Clement²; David Miller³; Tatsumi Nitta⁴

¹ *IFIC-Valencia (UV/EG-CSIC)*

² *Stockholm University (SE)*

³ *University of Chicago (US)*

⁴ *Waseda University (JP)*

Corresponding Authors: tatsumi.nitta@cern.ch, david.w.miller@cern.ch, martis@ific.uv.es, christophe.clement@cern.ch

By colliding protons and examining the particle emitted from the collisions, the Large Hadron Collider aims to study the interactions of quarks and gluons at the highest energies accessible in a controlled experimental way. In such collisions, the types of interactions that occur may extend beyond those encompassed by the Standard Model of particle physics. Such interactions typically occur at energy scales much higher than the rest mass of the incoming or outgoing particles. Because of this, reconstructing highly relativistic particles emitted from these collisions is becoming increasingly important. In particular, the ability to identify the originating particle which decays hadronically using distinguishing features of the radiation pattern of the jet plays a central role in searches. This is typically done by the use of a single physically motivated observable constructed from the constituents of the jet. In this work, multiple complementary observables are combined using boosted decision trees and neural networks to increase the ability to distinguish W bosons and top quarks from light quark jets in the ATLAS experiment.

Poster Session / 141

Hydra: A framework for data analysis in massively parallel platforms

Authors: Antonio Augusto Alves Junior¹; Antonio Augusto Alves Junior²; Michael David Sokoloff¹

¹ *University of Cincinnati (US)*

² *Univercity of Cincinnati*

Corresponding Authors: aalvesju@cern.ch, antonio.augusto.alves.junior@cern.ch, michael.david.sokoloff@cern.ch

Hydra is a templated header-only, C++11-compliant library for data analysis on massively parallel platforms targeting, but not limited to, the field High Energy Physics research.

Hydra supports the description of particle decays via the generation of phase-space Monte Carlo, generic function evaluation, data fitting, multidimensional adaptive numerical integration and histogramming.

Hydra is open source and the code is hosted in GitHub.

The library deploys a series of techniques in order to achieve optimal performance in both computing and management of memory resources.

The overall design exploits heavily C++ variadic templates in order to implement static polymorphism, kernel fusion and coalesced memory access patterns, avoiding completely the usage of function pointers, virtual methods calls and other known potential performance degrading constructs.

Hydra is developed on top of the Thrust library and runs on Linux systems and can deploy transparently NVidia CUDA-enabled GPUs as well as multicore CPUs and accelerators.

This contribution summarizes the main features of Hydra. A basic description of the user interface and some examples of applications are provided, along with measurements of performance in a variety of environments.

Track 1: Computing Technology for Physics Research / 142

Enabling High Availability Service with oVirt Virtualization and CephFS

Authors: Michael Poat¹; Jerome LAURET¹

¹ *Brookhaven National Laboratory*

Corresponding Authors: jlauret@bnl.gov, mpoat@bnl.gov

The online computing environment at STAR has generated demand for high availability of services (HAS) and a resilient uptime guarantee. Such services include databases, web-servers, and storage systems that user and sub-systems tend to rely on for their critical workflows. Standard deployment of services on bare metal creates a problem if the fundamental hardware fails or loses connectivity. Additionally, the configuration of redundant fail-over nodes (a secondary Web service for example) requires constant syncing of configuration and content and sometimes manual interaction for switching hardware (DNS name comes to mind). Beyond those uses, and within any computing environment, over-provisioned systems with unused CPU and memory resources could be put to use with container or virtualization technologies. How to achieve HAS using OpenSource packages and our experience will be the objective of our presentation.

We will focus on two tools: oVirt and Ceph. For Ceph, we have presented in past conferences our testing experience, performance improvement attempts as well as the status of our production system. Growing in popularity as a distributed storage system, it appeared natural to leverage our experience for this project. oVirt is an OpenSource virtualization management application that enables central management of hardware nodes, storage, and network resources used to deploy and monitor virtual machines. oVirt supports the deployment of a virtual environment for your data center leveraging automatic provisioning, live migration, and the ability to easily scale the number of hypervisors. oVirt enables the use of multiple storage technologies where you can store virtual

machines, images, and templates within one or multiple storage systems. STAR's recent efforts focused on deploying a CephFS POSIX compliant distributed storage system, we would enable the ability to couple our Ceph storage with the oVirt virtualization management system.

When designing an intricate system for hosting critical services, it is a requirement to circumvent single points of failure. This work will involve the testing and viability of such an approach along with test cases for high availability, live migration, and service on demand.

Track 1: Computing Technology for Physics Research / 143

Invited parallel: A Machine Learning tool for fast simulation in GeantV

Authors: Sofia Vallecorsa¹; Federico Carminati¹; Andrei Gheata¹; Gulrukh Khattak¹; Patricia Mendez Lorenzo¹

¹ CERN

Corresponding Authors: gul.rukh.khattak@cern.ch, sofia.vallecorsa@cern.ch, andrei.gheata@cern.ch, federico.carminati@cern.ch, patricia.mendez@cern.ch

The GeantV project introduces fine grained parallelism, vectorisation, efficient memory management and NUMA awareness in physics simulations. It is being developed to improve accuracy, while preserving, at the same time, portability through different architectures (Xeon Phi, GPU). This approach brings important performance benefits on modern architectures and a good scalability through a large number of threads.

Within the GeantV framework we have started developing a machine learning based tool for fast simulation. Machine learning techniques have been used in different applications by the HEP community, however the idea of using them to replace detector simulation is still rather new. Our plan is to provide, in GeantV, a fully configurable tool to train a neural network to reproduce the detector response and replace standard Monte Carlo simulation. This represents a completely generic approach in the sense that such a network could be designed and trained to simulate any kind of detector and, eventually, the whole data processing chain in order to get, directly in one step, the final reconstructed quantities. Such development is intended to address the ever increasing need for simulated events expected for LHC experiments and their upgrades, such as the High Luminosity LHC. We will present results of the first tests we run on several machine learning and deep learning models, including computer vision techniques, to simulate particle showers in calorimeters.

Poster Session / 144

Upgrade of the YARR DAQ system for the ATLAS Phase-II Pixel detector readout chip

Author: Nikola Lazar Whallon¹

Co-authors: Timon Heim²; Hideyuki Oide³; Karolos Potamianos²

¹ University of Washington (US)

² Lawrence Berkeley National Lab. (US)

³ INFN-Genova

Corresponding Authors: karolos.potamianos@cern.ch, timon.heim@cern.ch, alokin@uw.edu, hideyuki.oide@cern.ch

The Yet Another Rapid Readout (YARR) system is a DAQ system designed for the readout of the current generation ATLAS Pixel FE-I4 chip, which has a readout bandwidth of 160 Mb/s, and the latest readout chip currently under design by the RD53 collaboration which has a much higher bandwidth

up to 5 Gb/s and is part of the development of new Pixel detector technology to be implemented in High-Luminosity Large Hadron Collider experiments.

YARR utilises a commercial-off-the-shelf PCIe FPGA card as a reconfigurable I/O interface, which acts as a simple gateway to pipe all data from Pixel modules via the high speed PCIe connection into the host system's memory.

All data processing is done on a software level in the host CPU(s), utilising a data-driven, multi-threaded, parallel processing paradigm.

This processing is designed to support a scalable, modular distribution to multiple CPUs with an asynchronous, message-oriented system control.

It is also designed to allow for a flexible configuration, enabling the system to adapt to various hardware boards and use-case scenarios.

As such, the software is designed to cover a large range of operational environments, from prototyping in the laboratory, to full scale implementation in the experiment -

this is one of the core design goals, as it conserves manpower and builds a larger user-base compared to more specialised systems.

YARR is also able to interface directly with software emulators of front-end chips which frees the software development from the necessity of readout hardware and is useful for unit and coverage tests.

The overall concept and data flow of YARR will be outlined, as well as a demonstration of the system's DAQ and calibration capabilities and performance results of the PCIe transfer rate.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 146

N-jettiness subtraction

Author: Xiaohui Liu¹

¹ *Beijing Normal University*

Corresponding Author: xiaohuiliu@gmail.com

In this talk, I will review the current status of N-jettiness subtraction scheme and its application to V_j production at the LHC.

Track 2: Data Analysis - Algorithms and Tools / 147

GooFit 2.0

Authors: Michael David Sokoloff¹; Henry Fredrick Schreiner¹

¹ *University of Cincinnati (US)*

Corresponding Authors: michael.david.sokoloff@cern.ch, henry.fredrick.schreiner@cern.ch

The GooFit package provides physicists a simple, familiar syntax for manipulating probability density functions and performing fits, but is highly optimized for data analysis on NVIDIA GPUs and multithreaded CPU backends. GooFit is being updated to version 2.0, bringing a host of new features. A completely revamped and redesigned build system makes GooFit easier to install, develop with, and run on virtually any system. Unit testing, continuous integration, and advanced logging options are improving the stability and reliability of the system. Developing new PDFs now uses standard CUDA terminology and provides a lower barrier for new users. The system now has built-in support for multiple graphics cards or nodes using MPI, and is being tested on a wide range of different systems.

GooFit also has significant improvements in performance on some GPU architectures due to optimized memory access. Support for time-dependent four body amplitude analyses has also been added.

Poster Session / 148**Deep Learning Method for Inferring Cause of Data Anomalies****Author:** Maxim Borisyak¹**Co-authors:** Denis Derkach¹; Andrey Ustyuzhanin¹; Olga Koval²; Fedor Ratnikov¹¹ *Yandex School of Data Analysis (RU)*² *Yandex***Corresponding Authors:** fedor.ratnikov@gmail.com, olga-koval@yandex-team.ru, maxim.borisyak@cern.ch, denis.derkach@cern.ch, andrey.ustyuzhanin@cern.ch

Daily operation of a large-scale experiment is a resource consuming task, particularly from perspectives of routine data quality monitoring. Typically, data comes from different channels (subdetectors or other subsystems) and the global quality of data depends on the performance of each channel. In this work, we consider the problem of prediction which channel has been affected by anomalies in the detector behaviour.

We introduce a generic deep learning model and prove, that, under reasonable assumptions, the model learns to identify 'channels' affected by an anomaly. Such model could be used for data quality manager cross-check and assistance and identifying good channels in anomalous data samples. The main novelty of the method is that the model does not require ground truth labels for each channel, only global flag is used. This effectively distinguishes the model from classical classification methods.

Evaluation of the method on data collected by the CERN CMS experiment is presented.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 150**Numerical techniques for 2- and 3-loop integrals****Author:** Ayres Freitas¹¹ *University of Pittsburgh***Corresponding Author:** afreitas@pitt.edu

There are currently no analytical techniques for computing multi-loop integrals with an arbitrary number of different massive propagators. On the other hand, the challenge for numerical methods is to ensure sufficient accuracy and precision of the results. This contribution will give a brief overview on numerical approaches based on differential equations, dispersion relations and Mellin-Barnes relations. Some recent developments and applications will be highlighted.

Poster Session / 151**Making Containers Lazy With Docker and CernVM-FS****Author:** Nikola Hardi¹**Co-authors:** Jakob Blomer¹; Gerardo Ganis¹; Radu Popescu¹¹ *CERN***Corresponding Authors:** nikola.hardi@cern.ch, gerardo.ganis@cern.ch, radu.popescu@cern.ch, jakob.blomer@cern.ch

Containerisation technology is becoming more and more popular because it provides an efficient way to improve deployment flexibility by packaging up code into software micro-environments. Yet, containerisation has limitations and one of the main ones is the fact that entire container images need to be transferred before they can be used. Container images can be seen as software stacks and High-Energy Physics has long solved the distribution problem for large software stacks with CernVM-FS. CernVM-FS provides a global, shared software area, where clients only load the small subset of binaries that are accessed for any given compute job.

In this paper, we propose a solution to the problem of efficient image distribution using CernVM-FS for storage and transport of container images. We chose to implement our solution for the Docker platform, due to its popularity and widespread use. To minimise the impact on existing workflows, our implementation comes as a Docker plugin, meaning that users will continue to pull, run, modify, and store Docker images using standard Docker tools.

We introduced the concept of a “thin” image, whose contents are served on demand from CernVM-FS repositories. Such behavior closely reassembles the lazy evaluation strategy in programming language theory. Our measurements confirm that the time before a task starts executing only depends on the size of the files actually used, minimizing the cold start-up time in all cases.

Track 1: Computing Technology for Physics Research / 152

A quantitative review of data formats for HEP analyses

Author: Jakob Blomer¹

¹ CERN

Corresponding Author: jakob.blomer@cern.ch

The analysis of High-Energy Physics (HEP) data sets often take place outside the realm of experiment frameworks and central computing workflows, using carefully selected “n-tuples” or Analysis Object Data (AOD) as a data source. Such n-tuples or AODs may comprise data from tens of millions of events and grow to hundred gigabytes or a few terabytes in size. They are typically small enough to be processed by an institute’s cluster or even by a single workstation. N-tuples and AODs are often stored in the ROOT file format, in an array of serialized C++ objects in columnar storage layout. In recent years, several new data formats emerged from the data analytics industry. We provide a quantitative comparison of ROOT and other popular data formats, such as Apache Parquet, Apache Avro, Google Protobuf, and HDF5. We compare speed, read patterns, and usage aspects for the use case of a typical LHC end-user n-tuple analysis. The performance characteristics of the relatively simple n-tuple data layout also provides a basis for understanding performance of more complex and nested data layouts. From the benchmarks, we derive performance tuning suggestions both for the use of the data formats and for the ROOT (de-)serialization code.

Track 1: Computing Technology for Physics Research / 153

DAQExpert - An expert system to increase CMS data-taking efficiency

Authors: Jean-Marc Olivier Andre¹; Ulf Behrens²; James Gordon Branson³; Olivier Chaze⁴; Sergio Cittolin³; Cristian Contescu¹; Georgiana Lavinia Darlea⁵; Christian Deldicque⁴; Zeynep Demiragli⁶; Marc Dobson⁴; Nicolas Doualot¹; Samim Erhan⁷; Jonathan Fulcher⁴; Dominique Gigi⁴; Maciej Szymon Gladki⁸; Frank Glege⁴; Guillermo Gomez Ceballos Retuerto⁵; Jeroen Hegeman⁴; Andre Holzner³; Mindaugas Janulis⁹; Lorenzo Masetti⁴; Frans Meijers⁴; Emilio Meschi⁴; Remi Mommsen¹; Srecko Morovic⁸; Vivian O’Dell^{None}; Luciano Orsini⁴; Christoph Paus⁵; Petia Petrova⁴; Andrea Petrucci¹⁰; Marco Pieri³; Dinyar Rabady⁴; Attila Racz⁴; Thomas Reis⁴; Hannes Sakulin⁴; Christoph Schwick⁴; Dainius Simelevicius⁸; Michail Vougioulas⁸; Petr Zejdl¹; Philipp Maximilian Brummer⁴

¹ Fermi National Accelerator Lab. (US)

² *Deutsches Elektronen-Synchrotron (DE)*

³ *Univ. of California San Diego (US)*

⁴ *CERN*

⁵ *Massachusetts Inst. of Technology (US)*

⁶ *Massachusetts Institute of Technology, Cambridge, Massachusetts, USA*

⁷ *University of California Los Angeles (US)*

⁸ *CERN, Geneva, Switzerland*

⁹ *Vilnius University (LT)*

¹⁰ *Rice University (US)*

Corresponding Authors: dominique.gigi@cern.ch, hannes.sakulin@cern.ch, sergio.cittolin@cern.ch, ulf.behrens@desy.de, marco.pieri@cern.ch, srecko.morovic@cern.ch, dinyar.rabady@cern.ch, odell@fnal.gov, thomas.reis@cern.ch, philipp.maximilian.brummer@cern.ch, andrea.petrucchi@cern.ch, branson@ucsd.edu, maciej.szymon.gladki@cern.ch, guillermo.gomez.ceballos@cern.ch, lorenzo.masetti@cern.ch, christian.deldicque@cern.ch, attila.racz@cern.ch, lavinia.darlea@cern.ch, michail.vougioukas@cern.ch, andre.holzner@cern.ch, jonathan.fulcher@cern.ch, nicolas.doualot@cern.ch, olivier.chaze@cern.ch, petr.zejdl@cern.ch, jm.andre@cern.ch, paus@mit.edu, frank.glege@cern.ch, emilio.meschi@cern.ch, remigius.mommsen@cern.ch, luciano.orsini@cern.ch, christoph.schwick@cern.ch, cristian.contescu@cern.ch, jeroen.hegeman@cern.ch, samim.erhan@cern.ch, marc.dobson@cern.ch, zeynep.demiragli@cern.ch, frans.meijers@cern.ch, petia.petrova@cern.ch, mindaugas.janulis@cern.ch, dainius.simelevicius@cern.ch

The efficiency of the Data Acquisition (DAQ) in the new DAQ system of the Compact Muon Solenoid (CMS) experiment for LHC Run-2 is constantly being improved. A significant factor on the data taking efficiency is the experience of the DAQ operator. One of the main responsibilities of DAQ operator is to carry out the proper recovery procedure in case of failure in data-taking. At the start of Run-2, understanding the problem and finding the right remedy could take a considerable amount of time, sometimes up to minutes. This was caused by the need to manually diagnose the error condition and to find the right recovery procedure out of an extended list which changed frequently over time. Operators heavily relied on the support of on-call experts, also outside working hours. Wrong decisions due to time pressure sometimes lead to an additional overhead in recovery time.

To increase the efficiency of CMS data-taking we developed a new expert system, the DAQExpert which provides shifters with optimal recovery suggestions instantly when the failure occurs. This tool significantly improves the response time of operators and the success rate of recovery procedures. Our goal is to cover all known failure conditions and to eventually trigger the recovery without human intervention wherever possible. This paper covers how we achieved two goals - making CMS more efficient and building a generic solution that can be used in other projects as well. More specifically we discuss how we: determine the optimal recovery suggestion, inject expert knowledge with minimum overhead, facilitate post-mortem analysis and reduce the amount of calls to on-call experts without deterioration of CMS efficiency. DAQExpert is a web application analyzing frequently updating monitoring data from all DAQ components and identifying problems based on expert knowledge expressed in small, independent logic-modules written in Java. Its results are presented in real-time in the control room via a web-based GUI and a sound-system in a form of short description of the current failure, and steps to recover. Additional features include SMS and e-mail notifications and statistical analysis based on reasoning output persisted in a relational database.

Track 2: Data Analysis - Algorithms and Tools / 155

Modeling detector digitization and read-out with adversarial networks

Authors: Maxim Borisyak¹; Andy Nelson²; Andrey Ustyuzhanin¹; Chase Owen Shimmin³

¹ *Yandex School of Data Analysis (RU)*

² *University of California Irvine (US)*

³ *Yale University (US)*

Corresponding Authors: chase.shimmin@cern.ch, andrey.ustyuzhanin@cern.ch, maxim.borisyak@cern.ch, andrew.james.nelson@cern.ch

Tools such as GEANT can simulate volumetric energy deposition of particles down to a certain energy and length scales.

However, fine-grained effects such as material imperfections, low-energy charge diffusion, noise, and read-out can be difficult to model exactly and may lead to systematic differences between the simulation and the physical detector.

In this work, we introduce a method based on Generative Adversarial Networks (GANs) that learns and corrects for these systematic effects.

The network transforms a simplistic GEANT simulation of a sensor into a realistic model that matches data from a physical sensor.

We also consider an extension of the GAN model based on Cycle-GAN, that allows for the introduction of explicit constraints on the network based on physical assumptions.

As a test case, we consider the simulation of cosmic ray interactions within mobile phone cameras, designed for use in the Cosmic RAYs Found In Smartphones (CRAYFIS) experiment. On the dataset from the experiment we demonstrate viability of the proposed approach and show its competitive advantages over known so far methods.

Track 2: Data Analysis - Algorithms and Tools / 158

Parallelized Kalman-Filter-Based Reconstruction of Particle Tracks on Many-Core Architectures

Authors: Giuseppe Cerati¹; Peter Elmer²; Slava Krutelyov³; Steven R Lantz⁴; Matthieu Lefebvre²; Mario Masciovecchio³; Kevin McDermott⁴; Daniel Sherman Riley⁴; Matevz Tadel³; Peter Wittich⁴; Frank Wuerthwein⁵; Avi Yagil³

¹ *Fermi National Accelerator Lab. (US)*

² *Princeton University (US)*

³ *Univ. of California San Diego (US)*

⁴ *Cornell University (US)*

⁵ *UCSD*

Corresponding Authors: ml15@princeton.edu, vyacheslav.krutelyov@cern.ch, peter.elmer@cern.ch, giuseppe.cerati@cern.ch, mario.masciovecchio@cern.ch, a.yagil@cern.ch, steve.lantz@cornell.edu, matevz.tadel@cern.ch, daniel.riley@cornell.edu, wittich@cornell.edu, fkw@fnal.gov, kevin.mcdermott@cern.ch

Faced with physical and energy density limitations on clock speed, contemporary microprocessor designers have increasingly turned to on-chip parallelism for performance gains. Examples include the Intel Xeon Phi, GPGPUs, and similar technologies. Algorithms should accordingly be designed with ample amounts of fine-grained parallelism if they are to realize the full performance of the hardware. This requirement can be challenging for algorithms that are naturally expressed as a sequence of small-matrix operations, such as the Kalman filter methods widely in use in high-energy physics experiments. In the High-Luminosity Large Hadron Collider (HL-LHC), for example, one of the dominant computational problems is expected to be finding and fitting charged-particle tracks during event reconstruction; today, the most common track-finding methods are those based on the Kalman filter. Experience at the LHC, both in the trigger and offline, has shown that these methods are robust and provide high physics performance. Previously we reported the significant parallel speedups that resulted from our efforts to adapt Kalman-filter-based tracking to many-core architectures such as Intel Xeon Phi. Here we report on how effectively those techniques can be applied to more realistic detector configurations and event complexity.

Track 1: Computing Technology for Physics Research / 159

Massive Parallel Lattice QCD Computing on FPGA Accelerator with Data-Flow Programming

Author: Thomas Janson^{None}

Co-author: Udo Kebschull¹

¹ *Johann-Wolfgang-Goethe Univ. (DE)*

Corresponding Authors: janson@iri.uni-frankfurt.de, uk@rz.uni-frankfurt.de

In this talk, we explore the data-flow programming approach for massive parallel computing on FPGA accelerator, where an algorithm is described as a data-flow graph and programmed with MaxJ from Maxeler Technologies. Such a directed graph consists of a small set of nodes and arcs. All nodes are fully pipelined and data moves along the arcs through the nodes. We have shown that we can implement complex algorithms like the Wilson Dirac operator from Lattice QCD. Our implementation collects all nearest neighbour terms on the four dimensional lattice to perform all arithmetic operations simultaneously.

Poster Session / 162

Present and future of PanDA WMS integration with Titan super-computer at OLCF

Authors: Alessio Angius¹; Danila Oleynik²; Matteo Turilli¹; Sergey Panitkin³; Shantenu Jha⁴

Co-authors: Alexei Klimentov⁵; Fernando Harald Barreiro Megino⁶; Jack Wells⁷; Kaushik De⁸; Paul Nilsson⁵; Pavlo Svirin⁹; Ruslan Mashinistov¹⁰; Sarp Oral¹¹; Siarhei Padolski¹²; Tadashi Maeno⁵; Torre Wenaus⁵

¹ *Rutgers University*

² *Joint Institute for Nuclear Research (RU)*

³ *Brookhaven National Lab*

⁴ *Rutgers University (US)*

⁵ *Brookhaven National Laboratory (US)*

⁶ *University of Texas at Arlington*

⁷ *Oak Ridge National Laboratory*

⁸ *University of Texas at Arlington (US)*

⁹ *National Academy of Sciences of Ukraine (UA)*

¹⁰ *Russian Academy of Sciences (RU)*

¹¹ *Oak Ridge National Lab*

¹² *BNL*

Corresponding Authors: ruslan.mashinistov@cern.ch, pavlo.svirin@cern.ch, tmaeno@bnl.gov, panitkin@bnl.gov, wenaus@gmail.com, matteo.turilli@gmail.com, kaushik@uta.edu, danila.oleynik@cern.ch, wellsjc@ornl.gov, alessio.angius@gmail.com, alexei.klimentov@cern.ch, fernando.harald.barreiro.megino@cern.ch, oralhs@ornl.gov, shantenu.jha@rutgers.edu, paul.nilsson@cern.ch, siarhei.padolski@cern.ch

Modern physics experiments collect peta-scale volumes of data and utilize vast, geographically distributed computing infrastructure that serves thousands of scientists around the world.

Requirements for rapid, near real time data processing, fast analysis cycles and need to run massive detector simulations to support data analysis pose special premium on efficient use of available computational resources.

A sophisticated Workload Management System (WMS) is needed to coordinate the distribution and processing of data and jobs in such environment.

In this talk we will discuss PanDA WMS developed by the ATLAS experiment at the LHC.

Even though PanDA was originally designed for workload management in Grid environment, it was successfully extended to include cloud resources and supercomputers.

In particular we'll described current state of PanDA integration with Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF).

Our approach utilizes a modified PanDA pilot framework for job submission to Titan's batch queues and for data transfers to and from OLCF .

The system employs lightweight MPI wrappers to run in parallel multiple, independent, single node payloads on Titan's multi-core worker nodes. It also gives PanDA a new capability to collect, in real time, information about unused worker nodes on Titan, which allows to precisely define the size and duration of jobs submitted to Titan according to available free resources. The initial implementation of this system already allowed to collect in 2016 more than 70M core hours of otherwise left unused resources on Titan and execute tens of millions of PanDA jobs. Based on experience gained on Titan the PanDA development team is exploring designs of next generation components and services for workload management on HPC, Cloud and Grid resources. In this talk we'll give an overview of these new components and discuss their properties and benefits.

Poster Session / 164

Gradient reversal for MC/real data calibration

Authors: Ryzhikov Artem^{None}; Andrey Ustyuzhanin¹

¹ *Yandex School of Data Analysis (RU)*

Corresponding Authors: andrey.ustyuzhanin@cern.ch, artemryzhikov@gmail.com

In the research, a new approach for finding rare events in high-energy physics was tested. As an example of physics channel the decay of $\tau \rightarrow 3 \mu$ is taken that has been published on Kaggle within LHCb-supported challenge. The training sample consists of simulated signal and real background, so the challenge is to train classifier in such way that it picks up signal/background differences and doesn't overfit to simulation-specific features. The approach suggested is based on cross-domain adaptation using neural networks with gradient reversal. The network architecture is a dense multi-branch structure. One branch is responsible for signal/background discrimination, the second branch helps to avoid overfitting on Monte-Carlo training dataset. The tests showed that this architecture is a robust a mechanism for choosing tradeoff between discrimination power and overfitting, moreover, it also improves the quality of the baseline prediction. Thus, this approach allowed us to train deep learning models without reducing the quality, which allow us to distinguish physical parameters, but do not allow us to distinguish simulated events from real ones. The third network branch helps to eliminate the correlation between classifier predictions and reconstructed mass of the decay, thereby making such approach highly viable for great variety of physics searches.

Track 2: Data Analysis - Algorithms and Tools / 165

Event Reconstruction on NOvA and DUNE with Deep Convolutional Networks

Author: Fernanda Psihas¹

¹ *Indiana University*

Corresponding Author: psihas@fnal.gov

Deep Convolutional Neural Networks (CNNs) have been widely used in the field of computer vision to solve complex problems in image recognition and analysis. Recently, we have applied a Deep Convolutional Visual Network (CVN), to identify neutrino events in the NOvA experiment. NOvA is a long baseline neutrino experiment whose main goal is the measurement of neutrino oscillations. It relies on the accurate identification and reconstruction of the neutrino flavor in observed interactions. In 2016 the NOvA collaboration observed oscillations in the $\nu_\mu \rightarrow \nu_e$ channel, a first HEP result employing CNNs. In this presentation, I will describe the core concepts of CNNs, how we apply them to identify NOvA events, and the details of their implementation in the Caffé framework.

I will also discuss the recent developments in the application of CNNs for particle tagging and other ongoing work in neutrino event reconstruction in NOvA and DUNE.

Track 1: Computing Technology for Physics Research / 166

Real-time heterogeneous stream processing with NaNet in the NA62 experiment.

Authors: Alessandro Lonardo¹; Piero Vicini¹; Andrea Biagioni²; Luca Pontisso¹; Paolo Cretaro³; Gianluca Lamanna⁴; Mattia Barbanera⁵; Michele Martinelli⁶; Ottorino Frezza⁷; Roberto Ammendola⁸

Co-authors: Roberto Piandani⁵; Marco Sozzi⁹; Pier Stanislao Paolucci⁷; Simula Francesco⁷; Paolo Valente¹; Elena Pastorelli⁷; Stefano Di Lorenzo¹⁰; Francesca Lo Cicero⁷

¹ *Sapienza Università e INFN, Roma I (IT)*

² *Università e INFN, Roma I (IT)*

³ *INFN - National Institute for Nuclear Physics*

⁴ *INFN e Laboratori Nazionali di Frascati (IT)*

⁵ *INFN Sezione di Pisa, Università' e Scuola Normale Superiore, P*

⁶ *INFN*

⁷ *INFN Roma 1*

⁸ *Università e INFN Roma Tor Vergata (IT)*

⁹ *Università di Pisa & INFN (IT)*

¹⁰ *Università & INFN, Pisa (IT)*

Corresponding Authors: andrea.biagioni@cern.ch, gianluca.lamanna@cern.ch, luca.pontisso@cern.ch, francesco.simula@roma1.infn.it, roberto.piandani@cern.ch, alessandro.lonardo@cern.ch, pier.paolucci@roma1.infn.it, stefano.di.lorenzo@cern.ch, marco.sozzi@cern.ch, paolo.cretaro@roma1.infn.it, paolo.valente@roma1.infn.it, roberto.ammendola@roma2.infn.it, francesca.locicero@roma1.infn.it, michele.martinelli@roma1.infn.it, piero.vicini@cern.ch, mattia.barbanera@cern.ch, ottorino.frezza@roma1.infn.it, elena.pastorelli@roma1.infn.it

The use of GPUs to implement general purpose computational tasks, known as GPGPU since fifteen years ago, has reached maturity. Applications take advantage of the parallel architectures of these devices in many different domains.

Over the last few years several works have demonstrated the effectiveness of the integration of GPU-based systems in the high level trigger of various HEP experiments. On the other hand the use of GPUs in the DAQ and low level trigger systems, characterized by stringent real-time constraints, poses several challenges.

In order to achieve such a goal we devised NaNet, a FPGA-based PCI-Express Network Interface Card design capable of: i) direct (zero-copy) data transfer with CPU and GPU (GPUDirect); ii) processing of incoming and outgoing data streams; iii) support for multiple link technologies (1/10/40GbE and custom ones).

The validity of our approach has been tested in the context of the NA62 CERN experiment, harvesting the computing power of last generation nVIDIA Pascal GPUs and of the FPGA hosted by NaNet to build in real-time refined physics-related primitives for the RICH detector (i.e. the Cerenkov rings parameters) that enable the building of more stringent conditions for data selection in the low level trigger.

Track 2: Data Analysis - Algorithms and Tools / 168

Automated proton track identification in MicroBooNE using gradient boosted decision trees

Author: Katherine Woodruff¹

¹ *New Mexico State University*

Corresponding Author: kwoodruf@nmsu.edu

MicroBooNE is a liquid argon time projection chamber (LArTPC) neutrino experiment that is currently running in the Booster Neutrino Beam at Fermilab. LArTPC technology allows for high-resolution, three-dimensional representations of neutrino interactions. A wide variety of software tools for automated reconstruction and selection of particle tracks in LArTPCs are actively being developed. Short, isolated proton tracks, the signal for low-momentum-transfer neutral current (NC) elastic events, are easily hidden in a large cosmic background. Detecting these low-energy tracks will allow us to probe interesting regions of the proton's spin structure. An effective method for selecting NC elastic events is to combine a highly efficient track reconstruction algorithm to find all candidate tracks with highly accurate particle identification using a machine learning algorithm. We present our work on particle track classification using gradient tree boosting software (XGBoost) and the performance on simulated neutrino data.

Track 2: Data Analysis - Algorithms and Tools / 169

ROOT at the center of the future analysis ecosystem

Author: Axel Naumann¹

¹ *CERN*

Corresponding Author: axel.naumann@cern.ch

ROOT <https://root.cern> is evolving along several new paths. At the same time it is reconsidering existing parts. This presentation will try to predict where ROOT will be in three years from now: the main themes of development and where we are already now, the big open questions as well as some of the questions that we didn't even ask yet. The oral presentation will cover the new graphics and GUI system, the new personality of ROOT's programming interfaces, new analysis approaches, as well as a whole chapter on concurrency.

Poster Session / 171

Alignment and Calibration Framework for the Belle II detector

Author: Jakub Kandra¹

Co-author: Tadeas Bilka²

¹ *Charles University in Prague*

² *Charles University, Prague*

Corresponding Authors: jakub.kandra@karlov.mff.cuni.cz, bilka@ipnp.mff.cuni.cz

The Belle II experiment is approaching its first physics run in 2018. Its full capability to operate at the precision frontier will need not only excellent performance of the SuperKEKB accelerator and the detector, but also advanced calibration methods combined with data quality monitoring.

To deliver data in a form suitable for analysis as soon as possible, an automated Calibration Framework (CAF) has been developed. The CAF integrates various calibration algorithms and their input

collection methods for event-level data. It allows execution of the calibration workflow using different backends from local machines to a computing cluster, resolution of dependencies among algorithms, management of the produced calibration constants, and database access across possible iterations.

One of the main algorithms fully integrated in the framework uses Millepede II to solve a large minimization problem emerging in the track-based alignment and calibration of the pixel and strip detector, the central drift chamber, and the muon system. Advanced fitting tools are used to properly describe the detector material and field and include measurements of different sub-detectors into a single global fit required for Millepede.

This talk will present the design of the calibration framework, the integration of the Millepede calibration, and its current performance.

Poster Session / 172

Parallelisation and Vectorisation of ROOT Fitting classes

Authors: Xavier Valls Pla¹; Lorenzo Moneta²

¹ *University Jaume I (ES)*

² *CERN*

Corresponding Authors: xavier.valls.pla@cern.ch, lorenzo.moneta@cern.ch

In order to take full advantage of new computer architectures and to satisfy the requirement of minimizing the CPU usage with increasing amount of data to analysis, parallelisation and vectorisation have been introduced in the ROOT mathematical and statistical libraries.

We report first on the improvements obtained in the function evaluation, used for data modelling, by adding the support for SIMD vectorisation using the convenient API provided by the VecCore package, which has been recently integrated in ROOT. We then present how the evaluation of the likelihood and the least square functions used for fitting ROOT histograms, graphs and trees have been parallelized using different paradigms.

We describe in detail how the vectorisation and parallelisation has been introduced and how the support for different SIMD's backed libraries and for different parallelisation strategies, such as those based on multi-threads or multi-process, has been included. Furthermore, we present some new generic classes supporting a task based parallelisation model, which have been introduced in ROOT, and how these new classes can be used also for other complex computational tasks.

Track 2: Data Analysis - Algorithms and Tools / 173

How easily can neural networks learn relativity?

Authors: Kartik Chitturi¹; Peter Onyisi¹

¹ *University of Texas (US)*

Corresponding Authors: kartik.prabhat.chitturi@cern.ch, ponyisi@utexas.edu

We study the ability of different deep neural network architectures to learn various relativistic invariants and other commonly-used variables, such as the transverse momentum of a system of particles, from the four-vectors of objects in an event. This information can help guide the optimal design of networks for solving regression problems, such as trying to infer the masses of unstable particles produced in a collision.

Track 2: Data Analysis - Algorithms and Tools / 176**A Toolkit to Study Sensitivity of the Geant4 Predictions to the Variations of the Physics Model Parameters**

Authors: Krzysztof Genser¹; Robert William Hatcher²; Michael Kelsey³; Gabriel Perdue²; Hans-Joachim Wenzel⁴; Dennis Herbert Wright^{None}; Julia Yarba⁴; Laura Fields⁵

¹ *Fermi National Accelerator Laboratory (US)*

² *Fermilab*

³ *SLAC National Accelerator Laboratory (US)*

⁴ *Fermi National Accelerator Lab. (US)*

⁵ *Fermi National Accelerator Laboratory*

Corresponding Authors: ljf26@fnal.gov, syjun@fnal.gov, genser@fnal.gov, kelsey@slac.stanford.edu, julia.yarba@cern.ch, perdue@fnal.gov, rhatcher@fnal.gov, dennis.herbert.wright@cern.ch, wenzel@fnal.gov

Geant4 is the leading detector simulation toolkit used in high energy physics to design detectors and to optimize calibration and reconstruction software. It employs a set of carefully validated physics models to simulate interactions of particles with matter across a wide range of interaction energies. These models, especially the hadronic ones, rely largely on directly measured cross-sections and phenomenological predictions with physically motivated parameters estimated by theoretical calculation or measurement. Because these models are tuned to cover a very wide range of possible simulation tasks, they may not always be optimized for a given process or a given material.

This raises several critical questions, e.g.:

How sensitive Geant4 predictions are to variations of the model parameters ?

What uncertainties are associated with a particular tune of one or another Geant4 physics model, or a group of models ?

How to consistently derive guidance for Geant4 model development and improvement from a wide range of available experimental data ?

We have designed and implemented a comprehensive, modular, user-friendly software toolkit to study and address such questions from the users community. It allows to easily modify parameters of one or several Geant4 physics models involved in the simulation, and to perform collective analysis of multiple variants of the resulting physics observables of interest and their statistical comparison against a variety of corresponding experimental data.

Based on modern event-processing infrastructure software, the toolkit offers a variety of attractive features, e.g. flexible run-time configurable workflow, comprehensive bookkeeping, easy to expand collection of analytical components.

Design, implementation technology, and key functionalities of the toolkit will be presented and illustrated with results obtained with Geant4 key hadronic models.

Keywords: Geant4 model parameters variation, systematic uncertainty in detector simulation

Track 3: Computations in Theoretical Physics: Techniques and Methods / 177**NNPDF: Neural Networks for precision PDF determinations**

Authors: Alberto Guffanti¹; Stefano Carrazza²

¹ *Università degli Studi di Torino*

² *CERN*

Corresponding Authors: stefano.carrazza@cern.ch, alberto.guffanti@unito.it

Parton Distribution Functions (PDFs) are a crucial ingredient for accurate and reliable theoretical predictions for precision phenomenology at the LHC.

The NNPDF approach to the extraction of Parton Distribution Functions relies on Monte Carlo techniques and Artificial Neural Networks to provide an unbiased determination of parton densities with a reliable determination of their uncertainties.

I will discuss the NNPDF methodology in general, the latest NNPDF global fit (NNPDF3.1) and then present ideas to improve the training methodology used in the NNPDF fits.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 181

High-precision calculation of the 4-loop contribution to the electron $g-2$ in QED

Author: Stefano Laporta¹

¹ *Dipartimento di Fisica, Universita di Bologna*

Corresponding Author: stefano.laporta@bo.infn.it

In this talk I will describe the results of evaluation up to 1100 digits of precision of the mass-independent contribution of the 891 4-loop Feynman diagrams contributing to the electron $g-2$ in QED.

I will show the analytical expressions fitted to the high-precision values, which contain polylogarithms of sixth-root of unity and one-dimensional integrals of products of complete elliptic integrals.

I will discuss also some technical aspects of my program SYS used to perform all the calculations.

Poster Session / 183

A federated Xrootd cache

Authors: Alja Mrak Tadel¹; Matevz Tadel¹; Frank Wuerthwein¹; Ben Steer²; Terrence Martin³; Edgar Fajardo Hernandez¹

¹ *Univ. of California San Diego (US)*

² *Monash University*

³ *UCSD*

Corresponding Authors: bsste2@student.monash.edu, frank.wuerthwein@cern.ch, tmartin@physics.ucsd.edu, matevz.tadel@cern.ch, alja.mrak.tadel@cern.ch, edgar.mauricio.fajardo.hernandez@cern.ch

With the shift in the LHC experiments from the computing tiered model where data was prefetched and stored at the computing site towards a bring data on the fly, model came an opportunity. Since data is now distributed to computing jobs using XrootD federation of data, a clear opportunity for caching arose.

In this document, we present the experience of installing and using a Federated Xrootd Cache (A Xrootd Cache consistent of several independent nodes). There is some fine tuning towards and scaling tests performed to make it fit for the CMS Analysis case.

Finally, we show how this federated cache can be expanded into a federation of caches in which the caches can be distributed among computing centers.

Track 2: Data Analysis - Algorithms and Tools / 184

Exploring end-to-end deep learning solutions for event classification at CMS

Authors: Michael Andrews¹; Manfred Paulini¹

Co-author: Sergei Gleyzer²

¹ *Carnegie-Mellon University (US)*

² *University of Florida (US)*

Corresponding Authors: sergei.gleyzer@cern.ch, manfred.paulini@cern.ch, michael.andrews@cern.ch

An essential part of new physics searches at the Large Hadron Collider at CERN involves event classification, or distinguishing signal decays from potentially many background sources. Traditional techniques have relied on reconstructing particle candidates and their physical attributes from raw sensor data. However, such reconstructed data are the result of a potentially lossy process of forcing raw data into progressively more physically intuitive kinematic quantities. However, powerful image-based machine learning algorithms have emerged that are able to directly digest raw data and output a prediction, so-called end-to-end deep learning classifiers. We explore the use of such algorithms to perform physics classification using raw sensory data from the CMS detector. As proof of concept, we classify photon versus electron identification using data from the CMS electromagnetic calorimeter. We show that for single particle shower images, we are able to exploit higher-order features in the shower to improve discrimination versus traditional shower shape variables. Furthermore, for full event classification, we show that these techniques are able to exploit correlations between different showers in the event to achieve strong discrimination.

Poster Session / 185

Networking requirements for Tier 2 sites

Author: Roger Jones¹

¹ *Lancaster University (GB)*

Corresponding Author: roger.jones@cern.ch

The LHC and other experiments are evolving their computing models to cope with the changing data volumes and rate, changing technologies in distributed computing and changing funding landscapes. The UK is reviewing the consequent network bandwidth provision required to meet the new models, there will be increasing consolidation of storage into fewer sites and increase use of caching and data streaming to exploit sites dominated by CPU capacity. While broad-brush arguments can be used to set the general scale of the bandwidths required, more detailed modelling and calculation is useful when convincing the providers of networking. An attempt to build a more detailed model based on ATLAS requirements in the UK will be presented.

Track 1: Computing Technology for Physics Research / 186

A container model for resource provision at a WLCG Tier-2

Authors: Gareth Douglas Roy¹; Gordon Stewart²; David Crooks¹; Samuel Cadellin Skipsey^{None}; David Britton¹

¹ *University of Glasgow (GB)*

² *University of Glasgow*

Corresponding Authors: david.crooks@cern.ch, gordon.stewart@glasgow.ac.uk, gareth.douglas.roy@cern.ch, david.britton@cern.ch, samuel.cadellin.skipsey@cern.ch

Containers are more and more becoming prevalent in Industry as the standard method of software deployment. They have many benefits for shipping software by encapsulating dependencies and turning complex software deployments into single portable units. Similar to Virtual Machines, but with a lower overall resource requirement, greater flexibility and more transparency they are a compelling choice for software deployment. The use of containers is becoming attractive to WLCG experiments as a means to encapsulate their payloads, ensure that userland environments are consistent and to segregate running jobs from one another to improve isolation. Technologies such as Docker and Singularity are already being used and tested by larger WLCG experiments along with CERN IT.

Our purpose here is to explore the use of containers at a medium to large WLCG Tier-2 as a method of reducing the manpower required to run such a site. By looking at the requirements of WLCG payloads (such as the availability of CVMFS, Trust Anchors or VOMS information) a model of a contained compute platform will be developed and presented. It is hoped that novel ways of interaction with experiment frameworks will be apparent along with the ability to leverage new technologies such as Docker-Swarm, Kubernetes or CoreOS to allow compute resources to be turned up quickly and effectively. Along with providing the compute it is hoped that readily available monitoring solutions can be bundled to provide a complete toolbox for local System Administrators to provide resources quickly and securely.

Track 2: Data Analysis - Algorithms and Tools / 187

The HEP.TrkX Project: deep neural networks for HL-LHC online and offline tracking

Authors: Dustin James Anderson¹; Paolo Calafiura²; Giuseppe Cerati³; Steven Andrew Farrell²; Lindsey Gray³; Jim Kowalkowski⁴; Mayur Mudigonda⁵; Prabhat⁵; Panagiotis Spentzouris⁴; Maria Spiropulu¹; Aristeidis Tsaris⁴; Jean-Roch Vlimant¹; Stephan Zheng⁶

¹ *California Institute of Technology (US)*

² *Lawrence Berkeley National Lab. (US)*

³ *Fermi National Accelerator Lab. (US)*

⁴ *Fermilab*

⁵ *Lawrence Berkeley National Laboratory*

⁶ *California Institute of Technology*

Corresponding Authors: maria.spiropulu@cern.ch, jean-roch.vlimant@cern.ch, aristeidistsaris@gmail.com, lindsey.gray@cern.ch, jbk@fnal.gov, dustin.james.anderson@cern.ch, prabhat@lbl.gov, spentz@fnal.gov, giuseppe.cerati@cern.ch, steven.farrell@cern.ch, pcalafiura@lbl.gov

Charged particle reconstruction in dense environments, such as the detectors of the High Luminosity Large Hadron Collider (HL-LHC) is a challenging pattern recognition problem. Traditional tracking algorithms, such as the combinatorial Kalman Filter, have been used with great success in HEP experiments for years. However, these state-of-the-art techniques are inherently sequential and scale quadratically or worse with increased detector occupancy. The HEP.TrkX project is a pilot project with the aim to identify and develop cross-experiment solutions based on machine learning algorithms for track reconstruction. Machine learning algorithms bring a lot of potential to this problem thanks to their capability to model complex non-linear data dependencies, to learn effective representations of high-dimensional data through training, and to parallelize easily on high-throughput architectures such as FPGAs or GPUs. In this talk we will discuss the evolution and performance of our recurrent (LSTM) and convolutional neural networks moving from basic 2D models to more complex models and the challenges of scaling up to realistic dimensionality/sparsity.

Poster Session / 188**Federating distributed storage for clouds in ATLAS****Authors:** Rolf Seuster¹; Frank Berghaus¹¹ *University of Victoria (CA)***Corresponding Authors:** frank.berghaus@cern.ch, rolf.seuster@cern.ch

Input data for applications that run in cloud computing centres can be stored at distant repositories, often with multiple copies of the popular data stored at many sites. Locating and retrieving the remote data can be challenging, and we believe that federating the storage can address this problem. A federation would locate the closest copy of the data currently on the basis of GeoIP information. Currently we are using the DynaFed data federation software solution developed by CERN IT. DynaFed supports several industry standards for connection protocols like Amazon's S3, Microsofts Azure, as well as WebDav and HTTP. Protocol dependent authentication is hidden from the user by using their X509 certificate. We have setup an instance of DynaFed and integrated it into the ATLAS Data Distribution Management system. We report on the challenges faced during the installation and integration. We have tested ATLAS analysis jobs submitted by the PanDA production system and we report on our first experiences with its operation.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 189**Numerical Methods and the 4-point 2-loop Higgs amplitudes****Author:** Stephen Jones¹¹ *MPI, Munich***Corresponding Author:** sjones@mpp.mpg.de

We introduce pySecDec, a toolbox for the numerical evaluation of multi-scale integrals, and highlight some of the new features. The use of numerical methods for the computation of multi-loop amplitudes is described, with a particular focus on Sector Decomposition combined with Quasi-Monte-Carlo integration and importance sampling. The use of these techniques in the computation of multi-loop Higgs amplitudes are reviewed and the potential to use similar methods for further multi-loop computations is discussed.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 190**A Java library to perform S-expansions of Lie algebras****Authors:** Carlos Inostroza¹; Kondrashuk Igor²; Felip Nadal³; Nelson Merino⁴¹ *Universidad de Concepcion*² *Univerdidad del Bío-Bío*³ *Instituto de Física Corpuscular (IFIC)*⁴ *Universite Paris Diderot***Corresponding Authors:** c.inostroza@gmail.com, igor.kondrashuk@gmail.com, nemerino@gmail.com, felip.nadal@gmail.com

The contraction method is a procedure that allows to establish non-trivial relations between Lie algebras and has had successful applications in both mathematics and theoretical physics. This work deals with generalizations of the contraction procedure with a main focus in the so called

S-expansion method as it includes most of the other generalized contractions. Basically, the S-expansion combines a Lie algebra G with a finite abelian semigroup S in order to define new S-expanded algebras. After giving a description of the main ingredients used in this paper, we present a Java library that automatizes the S-expansion procedure. With this computational tool we are able to represent Lie algebras and semigroups, so we can perform S-expansions of Lie algebras using arbitrary semigroups. We explain how the library methods has been constructed and how they work; then we give a set of example programs aimed to solve different problems. They are presented so that any user can easily modify them to perform his own calculations, without being necessarily an expert in Java. Finally, some comments about further developments and possible new applications are made.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 191

Algorithm to find S-related Lie algebras

Authors: Carlos Inostroza¹; Kondrashuk Igor²; Nelson Merino³; Felip Nadal⁴

¹ *Universidad de Concepcion*

² *Univerdidad del Bío-Bío*

³ *Universite Paris Diderot*

⁴ *Instituto de Física Corpuscular (IFIC)*

Corresponding Authors: felip.nadal@gmail.com, c.inostroza@gmail.com, nemerino@gmail.com, igor.kondrashuk@gmail.com

S-expansion of Lie algebras is a procedure that contains the Inonu-Wigner contraction and most of its generalizations. Based on a recent work where we have presented a Java library to perform S-expansions of Lie algebras [arXiv:1703.04036] we provide an extension allowing to solve different problems. In particular, in this work we complement our library of [arXiv:1703.04036] with new methods allowing not only to answer whether two given algebras can be S-related, but also to determine and classify new expanded algebras satisfying certain conditions of interest in different problems. We also give and explain some program examples which are presented in such a way that a user, not necessarily an expert in Java, can easily modify them to perform his own calculations.

Plenary / 192

New developments in FORM

Corresponding Author: tueda@nikhef.nl

Symbolic computation is an indispensable tool for theoretical particle physics, especially in the context of perturbative quantum field theory. In this talk, I will review FORM, one of computer algebra systems widely used in higher-order calculations, its design principles and advantages. The newly released version 4.2 will also be discussed.

Plenary / 193

Deep Learning and Particle Physics

Corresponding Author: kyle.cranmer@cern.ch

Plenary / 194

Deep Learning usage by large experiments

Corresponding Author: benjamin.philip.nachman@cern.ch

Modern machine learning (ML) has introduced a new and powerful toolkit to High Energy Physics. While only a small number of these techniques are currently used in practice, research and development centered around modern ML has exploded over the last year(s). I will highlight recent advances with a focus on jet physics to be concrete. Themselves defined by unsupervised learning algorithms, jets are a prime benchmark for state-of-the-art ML applications and innovations. For example, I will show how deep learning has been applied to jets for classification, regression, and generation. These tools hold immense potential, but incorporating domain-specific knowledge is necessary for optimal performance. In addition, studying what the machines are learning is critical for robustness and may even help us learn new physics!

Plenary / 195

ML in theoretical physics: from PDFs to MC tools

Corresponding Author: stefano.carrazza@cern.ch

We start the discussion by summarizing recent and consolidated applications of ML in TH-HEP. We then focus our discussion on recent studies about parton distribution functions determination and related tools based on machine learning algorithms and strategies. We conclude by showing future theoretical applications of ML to Monte Carlo codes.

Plenary / 196

Statistics, Data Mining, and Machine Learning in Astronomy

Corresponding Author: jakevdp@cs.washington.edu

Plenary / 197

Meta: Toward generative C++

Corresponding Author: hsutter@microsoft.com

Can we evolve the C++ language itself to make C++ programming both more powerful and simpler, and if so, how? The only way to accomplish both of those goals at the same time is by adding abstractions that let programmers directly express their intent—to elevate comments and documentation to testable code, and elevate coding patterns and idioms into compiler-checkable declarations.

This talk covers one such experimental feature I'm currently working on, provisionally called "meta-classes," which aims to provide a way to generatively write C++ types more simply and flexibly, and includes design motivation and how this can affect C++ programming in the future in many domains.

Plenary / 198

New results on the g-2 calculation

Corresponding Author: stefano.laporta@bo.infn.it

Plenary / 199

FPGA Datacenters – The New Supercomputer

Corresponding Author: andrew.putnam@microsoft.com

The emergence of Cloud Computing has resulted in an explosive growth of computing power, where even moderately-sized datacenters rival the world's most powerful supercomputers in raw compute capacity.

Microsoft's Catapult project has augmented its datacenters with FPGAs (Field Programmable Gate Arrays), which not only expand the compute capacity and efficiency for scientific computing, but also allows for the creation of customized accelerated networks for communicating between compute nodes.

In this talk, I will describe the architecture of Microsoft's Catapult system, the results that we have achieved with compute offload, network acceleration, and even machine learning, and look toward areas where FPGAs can usher in unprecedented performance and efficiency on scientific computing workloads.

Plenary / 200

Generative models for fast simulation

Corresponding Author: sofia.vallecora@cern.ch

Machine Learning techniques have been used in different applications by the HEP community: in this talk, we discuss the case of detector simulation. The need for simulated events, expected in the future for LHC experiments and their High Luminosity upgrades, is increasing dramatically and requires new fast simulation solutions. We will present results of several studies on the application of computer vision techniques to the simulation of detectors, such as calorimeters. We will also describe a new R&D activity within the GeantV project, aimed at providing a configurable tool capable of training a neural network to reproduce the detector response and replace standard Monte Carlo simulation. This represents a generic approach in the sense that such a network could be designed and trained to simulate any kind of detector and, eventually, the whole data processing chain in order to get, directly in one step, the final reconstructed quantities, in just a small fraction of time. We will present the first three-dimensional images of energy showers in a high granularity calorimeter, obtained using Generative Adversarial Networks.

Plenary / 201

Physics and Computing Challenges in Lattice QCD

Corresponding Author: ask@fnal.gov

In this talk, I will give a quick overview of physics results and computational methods in lattice QCD. Then I will outline some of the physics challenges, especially those of interest to particle physicists. Last, I will speculate on how machine-learning ideas could be applied to accelerate lattice-QCD algorithms.

Plenary / 202

HEP Software Development in the next decade; the views of the HSF community

Corresponding Author: sexton@fnal.gov

We've known for a while now that projections of computing needs for the experiments running in 10 years from now are unaffordable. Over the past year the HSF has convened a series of workshops aiming to find consensus on the needs, and produce proposals for research and development to address this challenge. At this time many of the software related drafts are far enough along to give a clear picture of what will result from this process. This talk will synthesize and report on some of the key elements to come out of this community work.

Plenary / 203

Development and optimization of the MCFM generator

Corresponding Author: giele@fnal.gov

Plenary / 204

HPC for HEP

Corresponding Author: djbard@lbl.gov

High Performance Computing (HPC) has been an integral part of HEP computing for decades, but the use of supercomputers has typically been limited to running cycle-hungry simulations for theory and experiment. Today's supercomputers offer spectacular compute power but are not always simple to use - supercomputers have a highly specialized architecture that means that code that runs well on a laptop or a small compute cluster will rarely scale up efficiently.

In this talk I will discuss initiatives developed at NERSC (the primary computing center for the DoE Office of Science) that are designed to enable scientists to work productively with supercomputers. In particular, I will describe how we are taking advantage of container technology (via the Shifter project) to solve the problems of portability and scalability on supercomputers at NERSC.

Plenary / 205

Deep Learning on the Intel Nervana platform

Corresponding Author: ravi.panchumarthy@intel.com

This presentation will share details about the Intel Nervana Deep Learning Platform and how a data scientist can use it to develop solutions for deep learning problems. The Intel Nervana DL Platform is a full-stack platform including hardware and software tools that enable data scientists to build high-accuracy deep learning solutions quickly and cost effectively than with alternative approaches. The Nervana DL platform is available in two deployment options: the hosted Intel Nervana Cloud and the on-premise deep learning appliance.

Plenary / 206

D-wave Quantum Computing

Corresponding Author: cpwilliams@dwavesys.com

Plenary / 207

Games and loop integrals

Corresponding Author: benruyl@gmail.com

Project HEPGame was created to apply methods from AI that have been successful for games, such as MCTS for Go, to solve problems in High Energy Physics. In this talk I will describe how MCTS helped us simplify large expressions. Additionally, I will describe how we managed to compute four loop (and some five loop) integrals in an automated way. I close with some interesting challenges for AI in theoretical high energy physics.

Plenary / 208

Survey of how DL is used in a large experiment

Corresponding Author: benjamin.philip.nachman@cern.ch

Plenary / 209

Machine Learning in HEP

Corresponding Author: sergei.gleyzer@cern.ch

Plenary / 210

Round Table on Machine Learning

Corresponding Author: pushpalatha.bhat@cern.ch

The round table will be animated by the following panelists

Kyle Cranmer
Wahid Bhijmi
Michela Paganini
Andrey Ustyuzhanin
Sergei Gleyzer

Plenary / 211

Deep Learning in Particle physics

Corresponding Author: daniel.whiteson@cern.ch

Plenary / 212

Data, Software, and Knowledge Preservation of Scientific Results

Corresponding Author: mikeh@omega.hep.nd.edu

Simply preserving the data from a scientific experiment is rarely sufficient to enable the re-use or re-analysis of the data. Instead, a more complete set of knowledge describing how the results were obtained, including analysis software and workflows, computation environments, and other documentation may be required. This talk explores the challenges in preserving the various knowledge products and how emerging technologies such as linux containers may provide solutions that simultaneously enable preservation and computational portability. In particular, the work of the Data and Software Preservation for Open Science (DASPOS) project and the building of the CERN Analysis Preservation Portal will be highlighted.

Plenary / 213

Track reconstruction algorithms in high pile up environments

Corresponding Author: heather.gray@cern.ch

The reconstruction of particle trajectories in the tracking detectors is one of the most complex parts in analysing the data at hadron colliders. Maximum luminosity is typically achieved at the cost of a large number of simultaneous proton-proton interactions between beam crossing. The large number of particles produced in such interactions introduces challenges both in terms of maintaining excellent algorithmic performance as well as meeting computing constraints. I will review the development of track reconstruction algorithms at hadron colliders and highlight how they have evolved to cope with pile up. I will also discuss some novel ideas for how track reconstruction algorithms may look in the future.

Plenary / 214

Why Cultural Diversity Matters: Becoming Culturally competent

Author: George Langford¹

¹ *Syracuse University*

Corresponding Author: glangfor@syr.edu

Research has shown that diversity enhances creativity. It encourages the search for novel information and perspectives leading to better decision making and problem solving, and leads to unfettered discoveries and breakthrough innovations. Even simply being exposed to diversity can change the way you think.

Professional development opportunities are needed to train faculty and staff to improve their ability to navigate issues of inclusion, equity, and cultural awareness, all of which improve research outcomes. Faculty and staff need training to address diversity factors directly whenever they mentor and teach non-traditional students. The training will help them avoid miscommunication, avoid privileging dominant cultural norms, and misappropriating expectations due to differing value orientations and work styles. Faculty must move away from the idea that science is neutral to cultural diversity factors and embrace the idea that cultural diversity matters.

Plenary / 215

Round Table on Diversity

Corresponding Authors: maria.girone@cern.ch, jlauret@bnl.gov

Our panel will cover the topics of “How to create/hire diversity into teams and the competitive advantage of diverse teams”.

We would like to collect questions you may have in advance so panelists have time to prepare comprehensive answers. We will collect them until Wednesday 23rd, noon. The form for this is at https://docs.google.com/forms/d/1J3_42C-qMNt3Pev3LaK2Z9-gZgm4hE5zU3GMbx1-5p0

The panelists will be

- Jennifer Barnes - Vice President, Operations & Communications at OneEnergy Renewables OneEnergy Inc.
- Paul Chiames - Chief Human Resource Officer, Stanford Linear Accelerator Center
- Tom Gallant - Co-chair of Lambda Alliance at Lawrence Berkeley Lab
- George Langford - Professor, Biology and Dean Emeritus of the College of Arts and Sciences, Syracuse University

Their short bio is attached below.

Plenary / 216

Summary of Track 1

Corresponding Author: schsu@uw.edu

Plenary / 217

Summary of Track 2

Corresponding Author: sergei.gleyzer@cern.ch

Plenary / 218

Summary of Track 3

Corresponding Author: afreitas@physik.unizh.ch

Plenary / 219

Summary of the Conference

Corresponding Author: pushpalatha.bhat@cern.ch

Plenary / 220

Machine Learning at NVIDIA

Corresponding Authors: vxtomgibbs@gmail.com, tgibbs@nvidia.com

Track 3: Computations in Theoretical Physics: Techniques and Methods / 221

Go-HEP: writing concurrent software with ease and Go

Author: Sebastien Binet¹

¹ *IN2P3/LPC*

Corresponding Author: sebastien.binet@cern.ch

High Energy and Nuclear Physics (HENP) libraries are now required to be more and more multi-thread-safe, if not multi-thread-friendly and multi-threaded.

This is usually done using the new constructs and library components offered by the C++11 and C++14 standards.

These components are however quite low-level (threads, mutexes, locks, ...) and hard to use and compose, or easy to misuse.

However, Go – a somewhat new language – provides a set of better building blocks for tackling concurrency: goroutines and channels.

This language is now used by the whole cloud industry: docker/moby, rkt, kubernetes are obvious flagships for Go.

But to be able to perform any meaningful physics analysis, one needs a set of basic libraries (matrix operations, linear algebra, plotting, I/O, ...)

We present Go-HEP, a set of packages to easily write concurrent software to interface with legacy HENP C++ physics libraries.

Go-HEP provides packages (in pure Go, no C++ required) to:

- read ROOT files (and their content: TH1x, TH2x, TTrees, TGraph)

- read/write YODA and numpy files

- read/write HepMC files
- read/write SLHA files
- read/write Les Houches files
- read/write SIO and LCIO files
- fill, save and load 1D and 2D histograms, scatters and profiles
- create PDF/LaTeX, PNG, JPEG, SVG and interactive plots

But Go-HEP also provides more physics-oriented packages:

- fmom: 4-vectors
- fads: a fast detector simulation toolkit, a reimplementaion of C++ Delphes in (concurrent) Go,
- fastjet: a reimplementaion of C++ FastJet, in concurrent Go,
- fit: a MINUIT-like minimization package,
- heppdt: HEP particle data tables,
- pawgo: a Physics Analysis Workstation in Go (REPL, plots, I/O)

We will first re-introduce – quickly – the concurrent programming building blocks of Go, its great development environment and tools (easy refactoring, quick deployment, fast edit/compile/run cycle, etc...).

We then describe two packages (rootio and fads) enabling physics analyses.

Finally, the performances (CPU, VMem) of two applications built with Go-HEP will be compared to their C++ counterparts: Delphes and Rivet.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 222

On the notion ”significance of the difference”

Author: Serguei Bitjukov¹

¹ *Institute for High Energy Physics (RU)*

Corresponding Author: serguei.bitjukov@cern.ch

The concept of “the significance of the difference” is widely used in data analysis in high-energy physics. We consider the significance Type C 1. This significance is used in works [2,3,4]. It is shown that when comparing two independent samples obtained from one general population, the distribution of estimates of the “significance of the difference” for mean values, obtained for these samples, is close to the standard normal distribution.

1 Bitjukov S., Krasnikov N., Nikitenko A., Smirnova V., Two approaches to Combining Significances, PoS (ACAT08) 118.

2 Aubert B., et al (BABAR Collaboration), Search for CP Violation in Neutral D Meson Cabibbo-suppressed Three-body Decays, Phys.Rev. D78 (2008) 051102.

3 Bediaga I., Bigi I.I., Gomes A., Guerrer G., Miranda J., dos Reis A.C., On a CP Anisotropy Measurement in the Dalitz Plot, Phys.Rev. D80 (2009) 096006.

4 Bitjukov S., Krasnikov N., Nikitenko A., Smirnova V., On the distinguishability of histograms, Eur.Phys. J. Plus (2013) 128:143.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 223

Complete one-loop electroweak radiative corrections to polarized Bhabha scattering in SANC

Author: Renat Sadykov¹

¹ *Joint Institute for Nuclear Research (RU)*

Corresponding Author: sadykov@cern.ch

We continue to study Bhabha scattering at the one-loop level in SANC system. It is our first step to development EW library for radiative corrections for processes $e+e \rightarrow f\bar{f}$ with longitudinal polarization for future colliders. Higher-order leading logarithmic QED corrections are taken into account by means of structure function approach. Comparison with the existing results for unpolarized Bhabha scattering cross-sections at LEP1 and LEP2 were performed.

Also we present numerical results for ILC and CLIC setup with various magnitude of polarization.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 224

THESEUS event generator for collisions at energies available at NICA and FAIR

Author: Sergey Merts^{None}

Corresponding Author: sergey.merts@gmail.com

In this work, we present results obtained with the new Three-fluid Hydrodynamics-based Event Simulator Extended by UrQMD final State interactions (THESEUS) and apply it to the description of heavy-ion collisions in the NICA/FAIR energy range. It presents the 3FH output in terms of a set of observed particles and the afterburner can be run starting from this output by means of the UrQMD model. Thus THESEUS as a new tool allows to discuss the multifaceted physics challenges at FAIR and NICA energies. The new simulation program has the unique feature that it can describe a hadron-to-quark matter transition of first order which proceeds in the baryon stopping regime that is not accessible to previous simulation programs designed for higher energies, like QGSM or PHSD.

The simulation proceeds in two steps: first, a numerical solution of the 3-fluid hydrodynamics is computed with the corresponding code. Based on the output of the hydrodynamic part, a Monte Carlo procedure is used to sample the ensemble of hadron distributions and the UrQMD code is engaged to calculate final state hadronic rescatterings.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 225

Transfer Learning in Astronomy: A New Machine Learning Paradigm

Author: Ricardo Vilalta¹

¹ *University of Houston*

Corresponding Author: r.vilalta.us@ieee.org

The widespread dissemination of machine learning tools in science, particularly in astronomy, has revealed the limitation of working with simple single-task scenarios in which any task in need of a predictive model is looked in isolation, and ignores the existence of other similar tasks. In contrast, a new generation of techniques is emerging where predictive models can take advantage of previous experience to leverage information from similar tasks. The new emerging area is referred to as "Transfer Learning". In this paper we briefly describe the motivation behind the use of transfer learning techniques, and explain how such techniques can be used to solve popular problems in astronomy. As an example, we show how a prevalent problem in astronomy is to estimate the class of an object (e.g., supernova Ia) using a generation of photometric light-curve datasets where data abounds, but class labels are scarce; such analysis can benefit from spectroscopic data where class

labels are known with high confidence, but data is of small size. Transfer learning provides a robust and practical solution to leverage information from one domain to improve the accuracy of a model built on a different domain. In the example above, transfer learning would look to overcome the difficulty in the compatibility of models between spectroscopic data and photometric data, since data properties such as size, class priors, and underlying distributions, are all expected to be significantly different.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 226

Round table discussion: Analytical vs. numerical methods for NNLO+ computations for LHC

Corresponding Authors: tueda@nikhef.nl, kato@cc.kogakuin.ac.jp, giele@fnal.gov, sjones@mpp.mpg.de

In recent years, we have seen an explosion of new results at the NNLO level and beyond for LHC processes. These advances have been achieved through both analytical and numerical techniques, depending on the process and the group that performed the calculation.

This panel discussion will address such as how much the minimization of computer running time is desirable and if the possibility to incorporate new results into an event generator is always important (which tends to favor analytical techniques). On the other hand, there may be relevant processes for which it is not conceivable to compute them without numerical methods.

The event will start with brief statements by each of the panel members, followed by an open discussion.

Track 1: Computing Technology for Physics Research / 227

Round-Table: Using heterogeneous resources for HEP computing

Author: Ian Fisk¹

¹ *Simons Foundation*

Corresponding Authors: gareth.roy@glasgow.ac.uk, david.lange@cern.ch, tgibbs@nvidia.com, fisk@cern.ch

Track 2: Data Analysis - Algorithms and Tools / 228

Background Suppression with the Belle II Neural Network Trigger

Corresponding Author: sebastian.skambraks@tum.de

Track 2: Data Analysis - Algorithms and Tools / 229

The ATLAS Electron and Photon Trigger

Corresponding Author: samuel.david.jones@cern.ch

Track 2: Data Analysis - Algorithms and Tools / 230

Boosted Decision Trees in the Level 1 Muon Endcap Trigger at CMS

Corresponding Author: carnesandnoble@gmail.com

Track 2: Data Analysis - Algorithms and Tools / 231

Tomographic Reconstruction of LArTPC Events using Wire-cell

Corresponding Author: chao@bnl.gov

Track 2: Data Analysis - Algorithms and Tools / 232

Muon reconstruction of JUNO experiment with convolutional neural networks

Corresponding Author: lu.wang@cern.ch

Poster Session / 233

Implementing the Belle II Conditions Database using Industry-Standard Tools

Author: Lynn Wood¹

Co-authors: Todd Elsethagen²; Kevin Fox³; Jeter Hall²; Bibi Raju²; Malachi Schram²; Eric Stephan²

¹ *Pacific Northwest National Laboratory, USA*

² *Pacific Northwest National Laboratory*

³ *PNNL*

Corresponding Authors: bibi.raju@pnnl.gov, eric.stephan@pnnl.gov, todd.elsethagen@pnnl.gov, lynn.wood@pnnl.gov, kevin.fox@pnnl.gov, jeter.hall@pnnl.gov, malachi.schram@gmail.com

The Belle II Experiment at KEK is preparing for first collisions in early 2018. Processing the large amounts of data that will be produced will require conditions data to be readily available to systems worldwide in a fast and efficient manner that is straightforward to both the user and maintainer. The Belle II Conditions Database was designed to make maintenance as easy as possible. To this end, a HTTP REST service was developed with industry-standard tools such as Swagger for the API interface development, Payara for the Java EE application server, and the Hazelcast in-memory data grid for support of scalable caching as well as transparent distribution of the service across multiple sites. This talk will present the design of the conditions database environment at Belle II, as well as go into detail about the actual implementation, capabilities, and performance.

Track 1: Computing Technology for Physics Research / 234

Belle II Conditions Database Overview

Authors: Martin Ritter^{None}, Lynn Wood¹

Co-authors: Thomas Kuhr²; Marko Bracko²; Christian Pulvermacher³; Todd Elsethagen⁴; Kevin Fox⁵; Jeter Hall⁴; Bibi Raju⁴; Malachi Schram⁴; Eric Stephan⁴

¹ Pacific Northwest National Laboratory, USA

² Jozef Stefan Institute (SI)

³ KIT

⁴ Pacific Northwest National Laboratory

⁵ PNNL

Corresponding Authors: lynn.wood@pnnl.gov, cpulvermacher@ekp.uni-karlsruhe.de, todd.elsethagen@pnnl.gov, jeter.hall@pnnl.gov, thomas.kuhr@lmu.de, kevin.fox@pnnl.gov, eric.stephan@pnnl.gov, malachi.schram@pnnl.gov, martin.ritter@lmu.de, bibi.raju@pnnl.gov, marko.bracko@ijs.si

The Belle II experiment at KEK is preparing for taking first collision data in early 2018. For the success of the experiment it is essential to have information about varying conditions available to systems worldwide in a fast and efficient manner that is straightforward to both the user and maintainer. The Belle II Conditions Database was designed to make maintenance as easy as possible. To this end, a HTTP REST service was developed with industry-standard tools such as Swagger for the API interface development, Payara for the Java EE application server, and the Hazelcast in-memory data grid for support of scalable caching as well as transparent distribution of the service across multiple sites.

On the client side, the online and offline software has to be able to obtain conditions data from the Belle II Conditions Database in a robust and reliable way under very different situations. As such the client side interface to the Belle II Conditions Database has been designed with a variety of access mechanisms which allow the software to be used with and without internet connection. Different methods to access the payload information are implemented to allow for high level of customization per site and to simplify testing of new payloads locally.

Changes to the conditions data are usually handled transparently but users can actively check whether an object has changed or register callback functions to be called whenever a conditions data object is updated. In addition a command line user interface has been developed to simplify inspection and modification of the database contents.

This talk will give an overview the design of the conditions database environment at Belle II on the server and client side and report about usage experiences and performance in large-scale Monte Carlo productions.

Poster Session / 235

Machine-Learning techniques for electro-magnetic showers identification in OPERA datasets

Authors: Andrey Ustyuzhanin¹; Sergey Shirobokov¹; Vladislav Belavin²; Artem Filatov³

¹ Yandex School of Data Analysis (RU)

² CERN

³ YSDA

Corresponding Authors: andrey.ustyuzhanin@cern.ch, vladislav.belavin@cern.ch, sergey.shirobokov@cern.ch

We investigate different approaches to the recognition of electromagnetic showers in the data which was collected by the international collaboration OPERA. The experiment initially was designed to detect neutrino oscillations, but the data collected can also be used for the development of the machine learning techniques for electromagnetic shower detection in photo emulsion films. Such showers

may be used as signals of Dark Matter interaction. Due to the design of the detector and exposure time, emulsion films contain few million of traces of cosmic rays and around 1000 signal tracks attributed to single shower. We propose three different algorithms for the shower identification. All the algorithms achieve higher performance than baseline and can completely clean the detector volume from the background tracks saving about a half of the signal tracks.

Plenary / 236

Final remarks

Corresponding Author: gwatts@uw.edu

Poster Session / 237

Mastering Opportunistic Computing Resources for HEP

Authors: Matthias Jochen Schnepf¹; Max Fischer²; Manuel Giffels¹; Andreas Heiss¹; Andreas Petzold¹; Christoph Heidecker¹; Eileen Kuehn¹; Gunter Quast¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

² *GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE)*

Corresponding Authors: matthias.jochen.schnepf@cern.ch, eileen.kuehn@kit.edu, andreas.petzold@cern.ch, andreas.heiss@kit.edu, christoph.heidecker@cern.ch, gunter.quast@cern.ch, max.fischer@cern.ch, manuel.giffels@cern.ch

As results of the excellent LHC performance in 2016, more data than expected has been recorded leading to a higher demand for computing resources. It is already foreseeable that for the current and upcoming run periods a flat computing budget and the expected technology advance will not be sufficient to meet the future requirements. This results in a growing gap between supplied and demanded resources. Physics is likely to be limited by the available computing and storage resources.

One option to reduce the emerging lack of computing resources is the utilization of opportunistic resources such as local university clusters, public and commercial cloud providers, HPC centers and volunteer computing. However, to use opportunistic resources additional challenges have to be tackled.

The traditional HEP Grid computing approach leads to a complex software framework that has special dependencies in operation system and software requirements, which currently prevents HEP from using these additional resources. To overcome these obstacles the concept of pilot jobs in combination with virtualization and/or container technology is the way to go. In this case the resource providers only needs to operate the “Infrastructure as a Service”, whereas HEP manages its complex software environment and the on-demand resource allocation. This approach allows us to utilize additional resources in a dynamically fashion on different kind of opportunistic resource providers.

Another challenge that has to be addressed is that not all workflows are suitable for opportunistic resources. For the HEP workflows the deciding factor is mainly the external network usage. To identify suitable workflows that can be outsourced to external resource providers, we propose an online clustering of workflows to identify those with low external network usage. This class of workflows can than be transparently outsourced to opportunistic resources dependent on the local site utilization.

Our approach to master opportunistic resources for the HEP community in Karlsruhe is currently evaluated and refined. Since the general approach is not tailored to HEP, it can be easily adapted by other communities as well.

Plenary / 238

Software Citations

Corresponding Author: d.katz@ieee.org

Plenary / 239

Welcome from the Department of Physics Chair

Corresponding Author: blayneh@uw.edu

Plenary / 240

Welcome To ACAT 20171

Corresponding Author: gwatts@uw.edu

241

ACAT Informal Bar Crawl

Three bars near Alder Hall know we might be coming... Join them, and discuss all things ACAT, Seattle, and relax... Go early, go late - they are open!

- College Inn
- Shultzy's
- Big Time

Poster Session / 242

Parllized JUNO simulation software based on SNIPER

Author: Tao Lin¹

¹ IHEP

Corresponding Author: lintao@ihep.ac.cn

Poster Session / 243

Paralleization and vectorization of ROOT fitting classes

Author: Xavier Valls Pla¹

¹ *University Jaume I (ES)*

Corresponding Author: xavier.valls.pla@cern.ch

Poster Session / 244

High Statistics and GPU Accerated Data Analysis in ICECUBE

Author: Philipp Eller¹

¹ *Penn State University*

Corresponding Author: pde3@psu.edu

Poster Session / 245

Speeding up Software with VECCORE, a portable SIMD library

Author: Guilherme Amadio¹

¹ *CERN*

Corresponding Author: amadio@cern.ch

246

Banquet

Poster Session / 247

Awards!

Corresponding Author: gwatts@uw.edu