ACAT 2019

Sunday 10 March 2019 - Friday 15 March 2019

Steinmatte conference center

Book of Abstracts

Contents

Welcome drinks / reception 1

DataForge: declarative approach to scientific data processing automation 1

Kotlin - new language for scientific programming 1

Using Deep Learning in Ultra-High Energy Cosmic Ray Experiments 2

Rings: an efficient library for polynomial rings 2

Decoding the nature of Dark Matter at current and future experiments 2

CutLang: a cut based HEP analysis description language and interpreter 3

Multi-threaded checksum computation for the ATLAS high-performance storage software 3

Evaluating InfluxDB and ClickHouse database technologies for improvements of the ATLAS operational monitoring data archiving 4

An Ensemble of Neural Networks for Online Electron Filtering at the ATLAS Experiment. 5

Performance evaluation of distributed file systems for the phase-II upgrade of the ATLAS experiment at CERN 5

ELisA: the ATLAS logbook facility extensions 6

Numerical multi-loop integration on heterogeneous many-core processors 6

Deep generative models for fast shower simulation in ATLAS 7

Parallelized Kalman-Filter-Based Reconstruction of Particle Tracks on Many-Core Architectures with the CMS Detector 7

Simulating Diverse HEP Workflows on Heterogeneous Architectures 8

ATLAS High Level Trigger within the multi-threaded software framework AthenaMT . . 9

Multi-threaded Event Reconstruction with JANA 9

The Central Hint and Information Processor system for automation, error detection and recovery in the ATLAS TDAQ Controls framework 9

Phase Space Integration of initial state radiation in p-p interaction 10

FPGA-accelerated machine learning inference as a solution for particle physics computing challenges . . . 10

hls4ml: deploying deep learning on FPGAs for trigger and data acquisition . . . 11

The “FELIX” detector interface for the ATLAS Trigger and Data Acquisition upgrades and its deployment in the ATLAS Inner Tracker demonstrator setup . . . 11

In-situ analysis and visualization of massively parallel computations of transitional and turbulent flows . . . 12

A hybrid deep learning approach to vertexing . . . 13

Machine Learning Techniques in the ATLAS TDAQ Network Monitoring System . . . 13

Track propagation for different detector and magnetic field setups in Acts . . . 13

The design and performance of the ATLAS Inner Detector trigger in high pile-up collisions at 13 TeV at the Large Hadron Collider . . . 14

Riemann-Theta Boltzmann Machine . . . 14

The second generation of the ATLAS Production System: expertise and future evolution 15

Boosted decision trees in the classification tasks with the Telescope Array surface detector data . . . 15

Design Pattern for Analysis Automation on Interchangeable, Distributed Resources using Luigi Analysis Workflows . . . 16

Data Sampling challenges in the ALICE O2 distributed processing system . . . 17

Recurrent GANs for particle-based simulation at the LHC . . . 17

Incorporation of Systematic Uncertainties in the Training of Multivariate Methods . . . 17

Introduction to the Full Bandwidth Amplitude Analysis Software (FALLS) . . . 18

Tagging Higgs bosons at the LHC with Interaction Networks . . . 18

Pileup mitigation with graph neural networks . . . 19

HEP.TrkX Charged Particle Tracking using Graph Neural Networks . . . 19

TAG based data analysis in BESIII offline sofware . . . 20

reFORM: designing a new symbolic manipulation toolkit . . . 20

Probing Neural Networks for the Gamma/Hadron Separation of CTA . . . 20

Particle Identification In PICO Using Semi-supervised Learning . . . 21

Controls and Configuration software for the ATLAS DAQ system in LHC Run 2 . . . 21

Global fits of BSM physics models with Gambit . . . 22

A New Visual Analytics toolkit for ATLAS metadata . . . 22

Parallel computing of Overlapping Grids Applications based on A Patch-based Communication Algorithm . . . 23

CMS Software and Offline preparation for future runs . . . 23

The NanoAOD event data format in CMS . . . 24

An innovative monitoring and maintenance model for the INFN CNAF Tier-1 data center infrastructure. . . . 24

Accelerating dark matter search in emulsion SHiP detector by Deep Learning . . . 25

Continuous Analysis in ATLAS: Running User-Defined Container Images on the Grid . . 25

pyhf: auto-differentiable binned statistical models . . . 26

Numerical calculation of high-order QED contributions to the electron anomalous magnetic moment . . . 26

A Domain Specific IDE for Parallel Application Development . . . 27

Machine Learning on sWeighted data . . . 27

Electromagnetic calorimeter reconstruction in Belle II . . . 27

Selective background Monte Carlo simulation at Belle II . . . 28

Heterogenous computing for the local reconstruction algorithms of the CMS calorimeters 28

Full Event Intepretation at Belle II . . . 29

New Features in FeynArts & Friends, and how they got used in FeynHiggs . . . 29

Pixel detector background simulation using generative adversarial networks at Belle II . 29

RooFit parallelization efforts . . . 30

The Belle II Simulation Library . . . 31

Parameter tuning of distributed storage system based on reinforcement learning . . . 31

Digitization of Cylindrical GEM Detector . . . 32

INFERNO: Inference-Aware Neural Optimisation . . . 32

Energy reconstruction of the ATLAS Tile Calorimeter under high pile-up conditions using the Wiener filter . . . 32

Trilinear Higgs boson coupling variations for di-Higgs production with full NLO QCD predictions in POWHEG . . . 33

Jupyter Notebook Support of the Belle 2 Software . . . 33

Automated and Intelligent Data Migration Strategy in High Energy Physical Storage Systems . . . 34

Double Higgs boson production in the high- and low-energy limits . . . 34

QED and electroweak radiative corrections to polarized Bhabha scattering 35

MCSANCee generator with one-loop electroweak corrections for processes with polarized e^+e^- beams 35

Production experience and performance study for HEP data production at HPC-TianheII 36

Deep Learning applied to hit classification for BESIII drift chamber 36

In-Memory Parallel Computing for Partial Wave Analysis 36

DD4hep: an integrated detector description tool for HEP 37

Core software challenges of the GPU High Level Trigger 1 of LHCb 37

Physics and computing performance of reconstruction algorithms for the GPU High Level Trigger 1 of LHCb 38

Adversarial Neural Network-based data-simulation corrections for jet-tagging at CMS . 38

Tracking performance for long living particles at LHCb 39

Machine Learning approach to boosting neutral particles identification in the LHCb calorimeter 39

Fast Data-Driven simulation of Cherenkov Detectors Using Generative Adversarial Networks. 39

Reinforced Sorting Networks for Particle Physics Analyses 40

Variational Dropout Sparsification for Particle Identification speed-up. 40

Machine Learning for Muon Identification at LHCb 40

Real-time reconstruction of long-lived particles at LHCb using FPGAs. 41

Real-time cluster finding for LHCb silicon pixel VELO detector using FPGA 41

Updates on SModelS 42

A 30 MHz software trigger and reconstruction for the LHCb upgrade 42

Deep learning for certification of the quality of the data acquired by the CMS experiment 42

Highly performant, Deep Neural Networks with sub-microsecond latency on FPGAs for trigger applications 43

Fips: An OpenGL based FITS viewer 43

Speeding HEP Analysis with ROOT Bulk I/O 44

ROOT I/O compression algorithms and their performance impact within Run 3 44

Evolution of ROOT package management 45

Migrating large codebases to C++ Modules 45

A 3D Track Finder for the Belle II CDC L1 Trigger 46

Air shower reconstruction with hexagonal convolutional neural networks 46

Machine Learning Study of Open Supernova Catalog 47

Meta-Learning for Artificial Neural Network Hyper-Parameter Optimization for CERN CMS Offline Data Certification 47

Scaling the training of semantic segmentation on MicroBooNE events to multiple GPUs 48

Electromagnetic-shower generation with Graphical GANs 48

Vectorization techniques for probability distribution function using VecCore 49

Physics inspired feature engineering with Lorentz Boost Networks 49

Primary Vertex Reconstruction with Deep Learning in FPGAs for the Phase-2 Upgrade of the Level-1 Trigger of CMS 50

Monitoring of time evolution of the trigger rates exploiting deep representation learning at the CMS experiment 50

Prong Formation in NOvA Reconstruction Pipeline 51

The computing model of the LHCb Upgrade 51

The core software framework for the LHCb Upgrade 51

Nested data structures in array and SIMD frameworks 52

Method for results comparison of various versions of simulation program by the use of experimental data 52

DeepXS: Fast approximation of MSSM electroweak cross sections at NLO 53

PhenoAI and iDarkSurvey: Learning (from) high-dimensional models 53

Deep learning for Directional Dark Matter search 54

Uncertainty reduction by gradient descent 54

Generalization of Homogeneity Tests Used in High Energy Physics Experiments 55

Further developments of FORM 55

CDF Long Term Data Preservation at the INFN CNAF Tier-1 data center: the Oracle database and software framework replica. 55

Forgotten (or not) possibility of the application of the least squares method: the analysis of the computer-based numerical results to QCD $\overline{\mathrm{MS}}$-on-shell mass relation at the four-loop level 56

Variational Autoencoders for New Physics Mining at the Large Hadron Collider 57

ConformalTracking: a geometry agnostic tracking library 57

Direct optimisation of the discovery significance when training neural networks to search for new physics in particle colliders 58

Deep Learning on HPC at NERSC 58

Vectorization of random number generation and reproducibility of concurrent particle transport simulation 59

micrOMEGAs5.0 59

Geant4 Parameter Tuning Using Professor 59

Reinforced Jet Grooming 60

Deep Learning based Algorithms in Astroparticle Physics 60

Charged Particle Tracking as a QUBO problem solved with Quantum Annealing-Inspired Optimization 61

Performance results of the GeantV prototype with complete EM physics 61

Federation of compute resources available to the German CMS community 62

An analytics driven computing model for HL-LHC 62

Submanifold Sparse Convolutional Networks for Sparse, Locally Dense Particle Image Analysis 63

DIRAC: Putting together heterogeneous resources for scientific applications 63

The ATLAS EventIndex and its evolution towards LHC Run 3 64

Using DODAS as deployment manager for smart caching of CMS data management system 65

Cross-domain Data Access System for Distributed Sites in HEP 65

The HEP.QPR Project: Quantum Pattern Recognition for Charged Particle Tracking . . . 66

Fast Deep Learning on FPGAs for the Phase-II L0 Muon Barrel Trigger of the ATLAS Experiment 67

CernVM-FS Container Image Integration 67

Track extrapolation and muon identification with GEANT4E in Belle II event reconstruction 68

Hardware Accelerated ATLAS Workloads on the WLCG 68

Belle2VR – An Interactive Virtual Reality Visualization of GEANT4 Event Histories . . . 68

Ntuples for NNLO processes 69

Applying data access policy to CernVM-FS 69

A fast and parametric digitization for triple-GEM detectors 70

Making RooFit Ready for Run 3 70

Techniques for the calculation of 4-loop QED contributions 70

STR: a Mathematica package for the method of uniqueness 71

A virtualized Tier-3g Facility Installiation in WLCG Network of CERN 71

New ROOT graphics language 72

Toward an efficient evaluation of two-loop massive scalar integralss 72

Renormalization of gauge theories at five loops 73

Towards the Increase in Granularity for the Main Hadronic ATLAS Calorimeter: Exploiting Deep Learning Methods 73

Algorithm to find an all-order in the running coupling solution to an equation of the DGLAP type 74

JupyterLab at the Saint Petersburg State University 74

GRAAL: Gem Reconstruction And Analysis Library 74

Computer algebra algorithm of simplification of tensor polynomials 75

Overall quality optimization for DQM stage in High Energy Physics experiments 75

A more Pythonic, Interoperable and Modern PyROOT 76

Generative Adversarial Networks for fast simulation: generalisation and distributed training in HPC 77

Recent developments of GRACE system 77

HepMC3 Event Record Library for Monte Carlo Event Generators 77

Machine Learning helping Monte-Carlo collider simulations 78

Boosting Performance of Data-intensive Analysis Workflows with Distributed Coordinated Caching 78

Towards a heterogeneous High Level Trigger farm for CMS 79

Using Continuous Deployment techniques to manage software change at a WLCG Tier-2 80

HEP Analyses on Dynamically Allocated Opportunistic Computing Resources 80

The Practitioner's C++17 and 20 News 81

STAR Data Production Workflow on HPC: Lessons Learned & Best Practices 81

Aligning the MATHUSLA test stand detector: Using TensorFlow 82

DIRACOS: a cross platform solution for grid tools 83

Constructing mass-decorrelated hadonic decay taggers in ATLAS 83

Identifying hadronically decaying vector bosons and top quarks in ATLAS 83

Distributed Computing Services in Belle II 84

Exploring Hybrid Deep Learning Architecture for Experimental HEP Data Analysis . . . 84

The DQM system for the SND . . . 85

Setting up technical requirements for an astroparticle data life cycle . . . 85

Excursion Set Estimation using Sequential Entropy Reduction for Efficient Searches for New Physics at the LHC . . . 86

Federated data storage evolution in HENP: data lakes and beyond . . . 87

ROOT Based Analysis: From Data to Plots, Expressive, Easy and Fast . . . 87

Weak signal extraction using matrix decomposition . . . 88

Three loop QCD corrections to heavy quark form factors . . . 88

Explore cloud solutions for ATLAS with $250,000 AWS cloud credits . . . 89

Data-driven Monte Carlo Generator for Low Energy e+e- Annihilation to Hadrons . . . 89

Next Generation of HEP CPU Benchmarks . . . 89

Computer algebra in physics research: algorithms, systems and modern developments . 90

Application on LHC High Energy Physic data analysis with IBM Quantum Computing . 90

Machine learning for an X-ray FEL . . . 91

Welcome to ACAT 2019 in Saas Fee! . . . 91

Denis Perret-Gallix (1949-2018) . . . 91

The STEM paradox: Factors affecting diversity in STEM fields . . . 91

Round table discussion . . . 92

TrackML : a tracking Machine Learning challenge . . . 92

Round table: Machine Learning . . . 93

Best Track 1 poster: Physics and computing performance of reconstruction algorithms for the GPU High Level Trigger 1 of LHCb . . . 93

Best Track 2 poster: Tracking performance for long living particles at LHCb . . . 93

Track 1 summary . . . 93

Track 2 Summary . . . 94

Track 3 Summary . . . 94

ACAT 2019 Summary . . . 94

Thank you and good bye . . . 94

Computations in Laser Particle Acceleration: heterogenous, parallel computing for physicists . . . 94

Implicit cause models 94

Keynote talk: Probabilistic Programming and Inference in Particle Physics 95

HEP.QPR Project: Quantum Pattern Recognition for Charged Particle Tracking 95

RISC-V based, silicon proven open-source processors from the PULP project 95

Making HPC friendly for scientists 96

Generative Models 96

GPU usage in theoretical calculations 96

Keynote talk: The future of HEP computing 96

Computing Architectures in the HL-LHC era 96

Decoding the nature of Dark matter at current and future experiments 96

Neuroscience and the Future of Computing 97

Keynote talk: Constraining effective field theories with machine learning 97

Beyond the Roadmap: HL-LHC HEP Software 97

Automatic Differentiation and Deep Learning 98

Higher-order QED contributions to the lepton anomalous magnetic moments 98

Machine learning in cosmology 98

Computational challenges for Generators 98

Machine Learning in Neutrino Physics 99

248

Welcome drinks / reception

Track 1: Computing Technology for Physics Research / 281

DataForge: declarative approach to scientific data processing automation

Author: Alexander Nozik[1]

[1] *INR RAS / MIPT*

Corresponding Author: altavir@gmail.com

Modern data processing (acquisition, storage and analysis) requires modern tools.
One of the problems shared by existing scientific software is "scripting" approach, when user writes an imperative script which describes the stages in which data should be processed. The main deficiency of such approach is the lack of possibility to automate the process. For example one usually needs script to manipulate or even textually generate other scripts in order to run complex tasks. Also scripted interaction usually could not be easily run in parallel or on distributed or cluster systems.

The DataForge metadata processing framework remedies this problem by presenting a declarative approach to data processing. In this approach the process described as a composition of tasks with automated task tree builder based on tree-like metadata communication. DataForge allows to write flexible simple tasks (in any programming style) and then create automatically managed task graphs of any complexity.

Poster Session / 282

Kotlin - new language for scientific programming

Author: Alexander Nozik[1]

[1] *INR RAS / MIPT*

Corresponding Author: altavir@gmail.com

One of the problems of scientific software development is lack of proper language tools to do it conveniently. Among the modern languages only few are able (have flexibility and most importantly libraries) to handle scientific tasks: C++, Python and Java. Also in some cases some niche languages like C# or Julia could be used.

The major problem of C++ is the complexity of the language and tremendous skill requirement to use it properly. Languages like Python and Julia remedy the problem of complexity, but as dynamic languages could not be used in large projects or performance critical tasks.

The Kotlin is a new language, developed by famous JetBrains company is fully compatible with Java and therefore has access to a large variety of scientific libraries. It is much less verbatim than Java which allows to write programs easier, also it addresses most of Java problems. What is more important, that Kotlin now has cross-platform support, which means that JVM, JS and native programs could be written in the same language.

In this report, we present some examples of using Kotlin language for simulation and data processing tasks in particle physics.

Plenary / 283

Using Deep Learning in Ultra-High Energy Cosmic Ray Experiments

Author: Oleg Kalashev[1]

[1] *Institute for Nuclear Research RAS*

Corresponding Author: kalashev@inr.ac.ru

The extremely low flux of ultra-high energy cosmic rays (UHECR) makes their direct observation by orbital experiments practically impossible. For this reason all current and planning UHECR experiments detect cosmic rays indirectly observing extensive air showers (EAS) initiated by cosmic ray particles in the atmosphere. Various types of shower observables are analysed in modern UHECR experiments including secondary radio signal and fluorescent light from excited nitrogen molecules. The most of data is collected by the network of surface area detectors which allows to measure horizontal EAS profile directly. The raw observables in this case are the time-resolved signals for the set of adjacent triggered detectors. To recover primary particle properties Monte Carlo shower simulation is performed. In traditional techniques the MC simulation is used to fit some synthetic observable such as shower rise time, shower front curvature and particle density normalized to a given distance from the core. In this talk we consider an alternative approach based on the deep convolutional neural network using detector signal time series as an input and trained on a large Monte-Carlo dataset. The above approach has proven its efficiency with the Monte-Carlo simulations of the Telescope Array Observatory surface detector. We will discuss in detail how we optimize network architecture for the particular task.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 284

Rings: an efficient library for polynomial rings

Author: Stanislav Poslavsky[1]

[1] *IHEP, Protvino*

Corresponding Author: stvlpos@mail.ru

The talk is devoted to the overview of *Rings* — an efficient lightweight library for commutative algebra written in Java and Scala languages. Polynomial arithmetic, GCDs, polynomial factorization and Gröbner bases are implemented with the use of modern asymptotically fast algorithms. *Rings* can be easily interacted or embedded in applications in high-energy physics and other research areas via a simple API with fully typed hierarchy of algebraic structures and algorithms for commutative algebra. The use of the Scala language brings a quite novel powerful, strongly typed functional programming model allowing to write short, expressive, and fast code for applications. At the same time *Rings* shows one of the best performances among existing software for algebraic calculations.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 285

Decoding the nature of Dark Matter at current and future experiments

Author: Alexander Belyaev[1]

[1] *University of Southampton & Rutherford Appleton Laboratory*

Corresponding Author: belyaev1967@gmail.com

Decoding the nature of Dark Matter (DM) is one of the most important problems of particle physics. DM can potentially provide unique signatures at collider and non-collider experiments. Details of these signatures which we expect to observe in the near future would allow us to delineate the properties of DM and the respective underlying theory Beyond the Standard Model (BSM). While there are many comprehensive phenomenology studies of various appealing BSM models, using "top-bottom" approach, there is no clear strategy for the reverse task of identifying underlying theory from the new signatures.

To solve this problems one should consider the comprehensive set of signatures,
database of models and use modern methods, including machine learning and artificial intelligence, which would allow us to decode the underlying theory form new signals. One of the important tools which could be helpful to solve the problem is Recently, High Energy Physics Model Database (HEPMDB) was created to make a step forward towards solving this problem. It is aimed to facilitate connection between HEP theory and experiment, to store, validate and explore BSM models and connect them to characteristic signatures. DM decoding is based on a very important complementarity of Large Hadron Collider (LHC) an DM direct and indirect detection experiments. This complementarity, modern analysis methods, comprehensive database of BSM models and their signatures are the key points for decoding of not only DM properties but the the whole underlying theory of Nature. The current status, the future prospects, strategies and tools for DM identification will be discussed.

Poster Session / 286

CutLang: a cut based HEP analysis description language and interpreter

Authors: Gokhan Unel[1]; Sezen Sekmen[2]

[1] *University of California Irvine (US)*

[2] *Kyungpook National University (KR)*

Corresponding Authors: sezen.sekmen@cern.ch, gokhan.unel@cern.ch

Nowadays, any physicist performing an analysis of the LHC data, needs to be well-versed in programming, at the level of both a system programmer and a software developer to handle the vast amounts of collision and simulation events. Even the simplest programming mistake in any of these areas can create big confusions on the analysis results. Moreover, a multitude of different analysis frameworks for similar tasks makes it difficult to communicate within the same experiment and preserve the analysis algorithms for inter-experimental discussions. The steep learning curve related to these analysis frameworks also erects a barrier between the data and the physicist who may simply wish to try out an analysis idea.
The abundance of inexpensive, powerful, easy to use computing power leads to a fundamental shift in data analysis. The development of a full-fledged text-based analysis algorithm description language that would eventually evolve into an AI-based analysis tool with which scientists can easily communicate in their own languages depends on this computing power. Such a language incorporating also logic and mathematical expressions, would eliminate all kinds of programming difficulties and errors, consequently allowing the scientist to focus on the goal, but not on the tool.
This presentation discusses the guiding principles of such an analysis description language and gives as an example the cut based analysis language and a runtime interpreter, CutLang. A number of LHC analyses of various complexities will also be shown to illustrate the advantages of a human readable analysis description language.

Poster Session / 287

Multi-threaded checksum computation for the ATLAS high-performance storage software

Authors: Fabrice Le Goff[1]; Giuseppe Avolio[1]

[1] *CERN*

Corresponding Author: giuseppe.avolio@cern.ch

ATLAS is one of the generic-purpose experiments observing hadron collisions at the LHC at CERN. Its trigger and data acquisition system (TDAQ) is responsible for selecting and transporting interesting physics events from the detector to permanent storage where the data are used for physics analysis. The transient storage of ATLAS TDAQ is the last component of the online data-flow system. It records selected events at several GB/s to non-volatile storage before transfer to the offline facilities. The transient storage is a distributed system consisting of high-performance direct-attached storage servers accounting for 480 hard drives. A distributed multi-threaded C++ application operates the hardware. Reliability and efficiency of this system are critical for the operations of ATLAS and TDAQ.
As part of the transient storage workflow, checksums of the recorded data files are calculated. The checksums are used throughout the offline data management and data distribution system to guarantee the integrity and correctness of the raw data. This paper presents the current multi-threading strategy of the transient storage software and the associated trade off between compute and storage hardware resources. We then introduce a novel multi-threaded checksum computation strategy. We discuss the key concepts of the implementation with a focus on the importance of overhead minimization. Finally the paper reports on the tests performed on the production system to demonstrate the validity of the implementation. A 30 % increase in the overall throughput performance will be demonstrated and discussed in the view of future LHC and ATLAS upgrades.

Poster Session / 288

Evaluating InfluxDB and ClickHouse database technologies for improvements of the ATLAS operational monitoring data archiving

Authors: Matei Vasile[1]; Igor Soloviev[2]; Giuseppe Avolio[3]

[1] *IFIN-HH (RO)*

[2] *University of California Irvine (US)*

[3] *CERN*

Corresponding Author: matei.vasile@cern.ch

The Trigger and Data Acquisition (TDAQ) system of the ATLAS experiment the Large Hadron Collider (LHC) at CERN currently is composed of a large number of distributed hardware and software components (about 3000 machines and more than 25000 applications) which, in a coordinated manner, provide the data-taking functionality of the overall system.
During data taking runs, a huge flow of operational data is produced in order to constantly monitor the system and allow proper detection of anomalies or misbehaviors. The Persistent Back-End for the ATLAS Information System of TDAQ (P-BEAST) is a system based on a custom-built time-series database and it is used to archive and retrieve for applications any operational monitoring

data. P-BEAST stores about 18 TB of highly compacted and compressed raw monitoring data per year acquired at 200 KHz average information update rate during ATLAS data taking periods.
Since P-BEAST has been put into production, 4 years ago, several promising database technologies for fast access to time-series and column-oriented data have become available. InfluxDB and ClickHouse were the most promising candidates for improving the performance and functionality of the current implementation of P-BEAST.
This paper presents a short description of main features of both technologies and a description of the synthetic tests ran on both database systems which try to leverage the best possible options for storage of the P-BEAST data using their respective data model capabilities. Then, the results of the performance testing that has been performed using a subset of archived ATLAS operational monitoring data are presented. Finally, a comparison of the results is presented.

Poster Session / 289

An Ensemble of Neural Networks for Online Electron Filtering at the ATLAS Experiment.

Author: Werner Spolidoro Freund[1]

[1] *Federal University of of Rio de Janeiro (BR)*

Corresponding Author: werner.spolidoro.freund@cern.ch

The ATLAS experiment implemented an ensemble of neural networks (NeuralRinger algorithm) dedicated to improve the performance of filtering events containing electrons in the high-input rate online environment of the Large Hadron Collider at CERN, Geneva.
This algorithm has been used online to select electrons with transverse energies above 15 GeV since 2017 and is extended to electrons with transverse energies below 15GeV in 2018. The ensemble employs a concept of calorimetry rings.
The training procedure and final structure of the ensemble are designed to keep flat detector response with respect to particle energy and position.
A detailed study was carried out to assess profile distortions in crucial offline quantities through the usage of statistical tests and residual analysis. These details and the online performance of this algorithm during the Run 2 data-taking will be presented.

Poster Session / 290

Performance evaluation of distributed file systems for the phase-II upgrade of the ATLAS experiment at CERN

Authors: Giuseppe Avolio[1]; Adam Abed Abud[2]; Fabrice Le Goff[1]

[1] *CERN*

[2] *Universita and INFN (IT)*

Corresponding Author: giuseppe.avolio@cern.ch

Over the next few years, the LHC will prepare for the upcoming High-Luminosity upgrade in which it is expected to deliver ten times more p-p collisions. This will create a harsher radiation environment and higher detector occupancy. In this context, the ATLAS experiment, one of the general purpose experiments at the LHC, plans substantial upgrades to the detectors and to the trigger system in order to efficiently select events. Similarly, the Data Acquisition System (DAQ) will have to redesign the data-flow architecture to

accommodate for the large increase in event and data rates.
The Phase-II DAQ design involves a large distributed storage system that buffers data read out from the detector, while a computing farm (Event Filter) analyzes and selects the most interesting events. This system will have to handle 5.2 TB/s of input data for an event rate of 1 MHz and provide access to 3 TB/s of these data to the filtering farm. A possible implementation for such a design is based on distributed file systems (DFS) which are becoming unavoidable among the big data industry. Features of DFS such as replication strategies and smart placement policies match the distributed nature and the requirements of the new data-flow system.
This paper presents an up-to-date performance evaluation of some of the DFS currently available: GlusterFS, HadoopFS and CephFS. After characterization of the future data-flow system's workload, we report on small-scale raw performance and scalability studies.
Finally, we conclude on the suitability of such systems to the tight constraints expected for the ATLAS experiment in phase-II and, in general, what the HEP community can profit from these storage technologies.

Poster Session / 291

ELisA: the ATLAS logbook facility extensions

Author: Alina Corso Radu[1]

[1] *University of California Irvine (US)*

Corresponding Author: alina.radu@cern.ch

Information concerning the operation, configuration and behaviour of the ATLAS experiment need to be reported, gathered and shared reliably with the whole ATLAS community which comprises over three thousand scientists geographically distributed all over the world. To provide such functionality, a logbook facility, Electronic Logbook for the information storage of ATLAS (ELisA), has been developed and actively used since the beginning of the LHC Run 2 period. The facility includes a user-friendly web interface to browse activity logs and to report on system operations with a configurable email notification system; a RESTful API used programmatically by other tools and services of the data acquisition infrastructure and a set of client API libraries and utilities to help user's interaction with the REST API.

Given its generic configuration capabilities, the ELisA facility has been recently deployed as a standalone logbook for other projects such as the commissioning of different sub-detectors and the offline assessment of data-quality. To ease this operation and to potentially extend ELisA usage to other projects, an extension of the database backend support is being implemented thus reducing one of the constraints (the ORACLE database) for the logbook deployment. Also, the deployment process of the logbook is being improved using containers for fast shipping and set up of all the necessary dependencies of the tool.

This contribution will present the status of the logbook facility as well as the extensions and improvements implemented to ease the logbook portability to other projects.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 292

Numerical multi-loop integration on heterogeneous many-core processors

Authors: Elise de Doncker[1]; Ahmed Almulihi[1]; Fukuko Yuasa[2]; Naohito Nakasato[3]; Hiroshi Daisaka[4]; Tadashi Ishikawa[2]

[1] *Western Michigan University*

[2] *High Energy Accelerator Research Organization (KEK)*

[3] *University of Aizu*

[4] *Hitotsubashi University*

Corresponding Authors: elise.dedoncker@wmich.edu, tadashi.ishikawa@kek.jp, ahmedhassan.almulihi@wmich.edu, fukuko.yuasa@kek.jp, nakasato@u-aizu.ac.jp

We report on multi-loop integral computations executed on a PEZY/Exascaler large-scale (immersion cooling) computing system. The programming model requires a host program written in C++ with an OpenCL kernel. However the kernel can be generated by the Goose compiler interface, which allows parallelizing loops according to compiler directives. As an advantage, the executable derived from a program instrumented with Goose pragmas can be run on multiple devices and multiple nodes without changes to the program. We use lattice rules and lattice copy (composite) rules on PEZY to approximate integrals for multi-loop self-energy diagrams with and without masses.

Poster Session / 295

Deep generative models for fast shower simulation in ATLAS

Author: Aishik Ghosh[1]

Co-author: Stefan Gadatsch [2]

[1] *Centre National de la Recherche Scientifique (FR)*

[2] *Universite de Geneve (CH)*

Corresponding Authors: aishik.ghosh@cern.ch, stefan.gadatsch@cern.ch

The extensive physics program of the ATLAS experiment at the Large Hadron Collider (LHC) relies on large scale and high fidelity simulation of the detector response to particle interactions. Current full simulation techniques using Geant4 provide accurate modeling of the underlying physics processes, but are inherently resource intensive. In light of the high-luminosity upgrade of the LHC and the need for ever larger simulated datasets to support physics analysis, the development of new faster simulation techniques is crucial. Building on the recent success of deep learning algorithms, Variational Auto-Encoders and Generative Adversarial Networks are investigated for modeling the response of the ATLAS electromagnetic calorimeter for photons in a central calorimeter region over a range of energies. The properties of synthesized showers using deep neural networks are compared to showers from a full detector simulation using Geant4. With this feasibility study we demonstrate the potential of using such algorithms for fast calorimeter simulation for the ATLAS experiment in the future, complementing current simulation techniques.

Track 2: Data Analysis - Algorithms and Tools / 296

Parallelized Kalman-Filter-Based Reconstruction of Particle Tracks on Many-Core Architectures with the CMS Detector

Authors: Mario Masciovecchio[1]; Steven R Lantz[2]; Allison Reinsvold Hall[3]; Daniel Sherman Riley[2]; Matevz Tadel[1]; Slava Krutelyov[1]; Avi Yagil[1]; Frank Wuerthwein[1]; Giuseppe Cerati[4]; Kevin McDermott[2]; Peter Wittich[2]; Peter Elmer[5]; Matti Kortelainen[4]; Boyana Norris[None]; Brian J. Gravelle[6]

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

[2] *Cornell University (US)*

[3] *Fermilab*

[4] *Fermi National Accelerator Lab. (US)*

[5] *Princeton University (US)*

[6] *University of Oregon*

Corresponding Authors: wittich@cornell.edu, frank.wuerthwein@cern.ch, kevin.mcdermott@cern.ch, peter.elmer@cern.ch, gravelle@cs.uoregon.edu, vyacheslav.krutelyov@cern.ch, giuseppe.cerati@cern.ch, ahall@fnal.gov, matti.kortelainen@cern.ch, steve.lantz@cornell.edu, daniel.riley@cornell.edu, a.yagil@cern.ch, mario.masciovecchio@cern.ch, brnorris03@gmail.com, matevz.tadel@cern.ch

In the High-Luminosity Large Hadron Collider (HL-LHC), one of the most challenging computational problems is expected to be finding and fitting charged-particle tracks during event reconstruction. The methods currently in use at the LHC are based on the Kalman filter. Such methods have shown to be robust and to provide good physics performance, both in the trigger and offline. In order to improve computational performance, we explored Kalman-filter-based methods for track finding and fitting, adapted for many-core SIMD and SIMT architectures. Our adapted Kalman-filter-based software has obtained significant parallel speedups using such processors, e.g., Intel Xeon Phi, Intel Xeon SP (Scalable Processors) and (to a limited degree) NVIDIA GPUs.
Recently, an effort has started towards the integration of our software into the CMS software framework, in view of its exploitation for the Run III of the LHC. Prior reports have shown that our software allows in fact for some significant improvements over the existing framework in terms of computational performance with comparable physics performance, even when applied to realistic detector configurations and event complexity. Here, we demonstrate that in such conditions physics performance can be further improved with respect to our prior reports, while retaining the improvements in computational performance, by making use of the knowledge of the detector and its geometry.

Track 1: Computing Technology for Physics Research / 297

Simulating Diverse HEP Workflows on Heterogeneous Architectures

Authors: Charles Leggett[1]; Illya Shapoval[2]; Marco Clemencic[3]; Christopher Jones[4]

[1] *Lawrence Berkeley National Lab (US)*

[2] *Lawrence Berkeley National Laboratory*

[3] *CERN*

[4] *Fermi National Accelerator Lab. (US)*

Corresponding Authors: charles.g.leggett@gmail.com, marco.clemencic@cern.ch, cdj@fnal.gov, ishapoval@lbl.gov

The next generation of HPC and HTC facilities, such as Oak Ridge's Summit, Lawrence Livermore's Sierra, and NERSC's Perlmutter, show an increasing use of GPGPUs and other accelerators in order to achieve their high FLOP counts. This trend will only grow with exascale facilities such as A21. In general, High Energy Physics computing workflows have made little use of GPUs due to the relatively small fraction of kernels that run efficiently on GPUs, and the expense of rewriting code for rapidly evolving GPU hardware. However, the computing requirements for high-luminosity LHC are enormous, and it will become essential to be able to make use of supercomputing facilities that rely heavily on GPUs and other accelerator technologies.

ATLAS has already developed an extension to AthenaMT, its multithreaded event processing framework, that enables the non-intrusive offloading of computations to external accelerator resources, and has begun investigating strategies to schedule the offloading efficiently. The same applies to LHCb, which, while sharing the same underlying framework as ATLAS (Gaudi), has considerably different workflow. CMS's framework, CMSSW, also has the ability to efficiently offload tasks to external accelerators. But before investing heavily in writing many kernels for specific offloading architectures, we need to better understand the performance metrics and throughput bounds of the workflows with various accelerator configurations. This can be done by simulating a diverse set

of workflows, using real metrics for task interdependencies and timing, as we vary fractions of offloaded tasks, latencies, data conversion speeds, memory bandwidths, and accelerator offloading parameters such as CPU/GPU ratios and speeds.

We present the results of these studies performed on multiple workflows from ATLAS, LHCb and CMS, which will be instrumental in directing effort to make HEP framework, kernels and workflows run efficiently on exascale facilities.

Poster Session / 298

ATLAS High Level Trigger within the multi-threaded software framework AthenaMT

Author: Rafal Bielski[1]

[1] *CERN*

Corresponding Author: rafal.bielski@cern.ch

Athena is the software framework used in the ATLAS experiment throughout the data processing path, from the software trigger system through offline event reconstruction to physics analysis. The shift from high-power single-core CPUs to multi-core systems in the computing market means that the throughput capabilities of the framework have become limited by the available memory per process. For Run 2 of the Large Hadron Collider (LHC), ATLAS has exploited a multi-process forking approach with the copy-on-write mechanism to reduce the memory use. To better match the increasing CPU core count and the, therefore, decreasing available memory per core, a multi-threaded framework, AthenaMT, has been designed and is now being implemented. The ATLAS High Level Trigger (HLT) system has been remodelled to fit the new framework and to rely on common solutions between online and offline software to a greater extent than in Run 2.

We present the implementation of the new HLT system within the AthenaMT framework, which is going to be used in ATLAS data-taking during Run 3 (2021 onwards) of the LHC. We also report on interfacing the new framework to the current ATLAS Trigger and Data Acquisition (TDAQ) system, which aims to bring increased flexibility whilst needing minimal modifications to the current system. In addition, we show some details of architectural choices which were made to run the HLT selection inside the ATLAS online data-flow, such as the handling of the event loop, returning of the trigger decision and handling of errors.

Track 1: Computing Technology for Physics Research / 300

Multi-threaded Event Reconstruction with JANA

Author: David Lawrence[1]

[1] *Jefferson Lab*

Corresponding Author: davidl@jlab.org

JANA2 is multi-threaded event reconstruction framework being developed for Experimental Nuclear Physics. It is an LDRD funded project that will be the successor of the original JANA framework. JANA2 is a near complete rewrite emphasizing C++ language features that have only become available since the C++11 standard. Successful and less-than-successful strategies employed in JANA and how they are being addressed in JANA2 will be presented as well as new features suited to modern and future trends in data analysis.

Track 1: Computing Technology for Physics Research / 301

The Central Hint and Information Processor system for automation, error detection and recovery in the ATLAS TDAQ Controls framework

Corresponding Author: giuseppe.avolio@cern.ch

The ATLAS experiment at the Large Hadron Collider at CERN relies on a complex and highly distributed Trigger and Data Acquisition (TDAQ) system to gather and select particle collision data obtained at unprecedented energy and rates. The TDAQ Controls system is the component that guarantees the smooth and synchronous operations of all the TDAQ components and provides the means to minimize the downtime of the system caused by runtime failures.
Given the scale and complexity of the TDAQ system and the rates of data to be analysed, the automation of the system functionality in the areas of error detection and recovery is a strong requirement. That is why in Run 2 the Central Hint and Information Processor (CHIP) service has been introduced; it can be truly considered the "brain" of the TDAQ Controls system. CHIP is an intelligent system able to supervise the ATLAS data taking, take operational decisions and handle abnormal conditions. It is based on an open-source Complex Event Processing (CEP) engine, ESPER. Currently, CHIP's knowledge base is made up of more than 300 rules organized in about 30 different contexts.
This paper will focus on the experience gained with CHIP during the whole LHC Run 2 period. Particular attention will be paid to demonstrate how the use of CHIP for automation and error recovery proved to be a valuable asset in optimizing the data taking efficiency, reducing operational mistakes, efficiently handling complex scenarios and improving the latency to react to abnormal situations. Additionally, the huge benefits brought by the CEP engine in terms of both flexibility and simplification of the knowledge base will be reported.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 302

Phase Space Integration of initial state radiation in p-p interaction

Author: Essma Redouane Salah[1]

[1] *University of M'sila Algeria*

Corresponding Author: essmars@hotmail.com

An efficient phase space integration is important for most calculations for collider experiments. We are developing a phase space integration that distribute phase space points according to the singular limit of QCD, Using the Altarelli-Parisi splitting functions as the underlying probability for a splitting, by developing and applying theoretical and computational tools

Track 1: Computing Technology for Physics Research / 303

FPGA-accelerated machine learning inference as a solution for particle physics computing challenges

Authors: Nhan Viet Tran[1]; Maurizio Pierini[2]; Philip Coleman Harris[3]; Kevin Pedro[1]; Javier Mauricio Duarte[1]; Ben Kreis[1]; Dylan Sheldon Rankin[3]; Sergo Jindariani[1]; Miaoyuan Liu[1]; Aristeidis Tsaris[4]; Jennifer Ngadiuba[2]; Burt Holzman[1]; Zhenbin Wu[5]

[1] *Fermi National Accelerator Lab. (US)*

[2] *CERN*

[3] *Massachusetts Inst. of Technology (US)*

[4] *Fermilab*

[5] *University of Illinois at Chicago (US)*

Corresponding Authors: nhan.viet.tran@cern.ch, jennifer.ngadiuba@cern.ch, dylan.sheldon.rankin@cern.ch, sergo.jindariani@cern.ch, philip.coleman.harris@cern.ch, maurizio.pierini@cern.ch, liumy871227@gmail.com, javier.mauricio.duarte@cern.ch, burt@fnal.gov, benjamin.jonah.kreis@cern.ch, aristeidistsaris@gmail.com, pedrok@cern.ch, zhenbinwu@gmail.com

Resources required for high-throughput computing in large-scale particle physics experiments face challenging demands both now and in the future. The growing exploration of machine learning algorithms in particle physics offers new solutions to simulation, reconstruction, and analysis. These new machine learning solutions often lead to increased parallelization and faster reconstructions times on dedicated hardware, here specifically Field Programmable Gate Arrays. We explore the possibility that applications of machine learning simultaneously also solve the increasing computing challenges. Employing machine learning acceleration as a web service, we demonstrate a heterogeneous compute solution for particle physics experiments that requires minimal modification to the current computing model. First results with Project Brainwave by Microsoft Azure, using the Resnet-50 image classification model as an example, demonstrate inference times of approximately 50 (10) milliseconds with our experimental physics software framework using Brainwave as a cloud (edge) service. We also adapt the image classifier, for example, physics applications using transfer learning: jet identification in the CMS experiment and event classification in the Nova neutrino experiment at Fermilab. Solutions explored here are potentially applicable sooner than may have been initially realized.

Track 2: Data Analysis - Algorithms and Tools / 304

hls4ml: deploying deep learning on FPGAs for trigger and data acquisition

Authors: Nhan Viet Tran[1]; Maurizio Pierini[2]; Philip Coleman Harris[3]; Sergo Jindariani[1]; Javier Mauricio Duarte[1]; Ben Kreis[1]; Jennifer Ngadiuba[2]; Zhenbin Wu[4]; Dylan Sheldon Rankin[3]; Sioni Paris Summers[5]; Ryan Allen Rivera[1]

[1] *Fermi National Accelerator Lab. (US)*

[2] *CERN*

[3] *Massachusetts Inst. of Technology (US)*

[4] *University of Illinois at Chicago (US)*

[5] *Imperial College Sci., Tech. & Med. (GB)*

Corresponding Authors: sioni.paris.summers@cern.ch, jennifer.ngadiuba@cern.ch, maurizio.pierini@cern.ch, benjamin.jonah.kreis@cern.ch, philip.coleman.harris@cern.ch, nhan.viet.tran@cern.ch, dylan.sheldon.rankin@cern.ch, zhenbinwu@gmail.com, javier.mauricio.duarte@cern.ch, rrivera@fnal.gov, sergo.jindariani@cern.ch

Machine learning is becoming ubiquitous across HEP. There is great potential to improve trigger and DAQ performances with it. However, the exploration of such techniques within the field in low latency/power FPGAs has just begun. We present hls4ml, a user-friendly software, based on High-Level Synthesis (HLS), designed to deploy network architectures on FPGAs. As a case study, we use hls4ml for boosted-jet tagging with deep networks at the LHC. We map out resource usage and latency versus network architectures, to identify the typical problem complexity that hls4ml could deal with. We discuss possible applications in current and future HEP experiments.

Track 1: Computing Technology for Physics Research / 305

The "FELIX" detector interface for the ATLAS Trigger and Data Acquisition upgrades and its deployment in the ATLAS Inner Tracker demonstrator setup

Corresponding Author: carlos.solans@cern.ch

The ATLAS experiment at the LHC at CERN will move to use the Front-End Link eXchange (FELIX) system in a staged approach for LHC Run 3 (2021) and LHC Run 4 (2026). FELIX will act as the interface between the data acquisition; detector control and TTC (Timing, Trigger and Control) systems; and new or updated trigger and detector front-end electronics.
FELIX functions as a router between custom serial links from front end ASICs and FPGAs to data collection and processing components via a commodity switched network. Links may aggregate many slower links or be a single high bandwidth link. FELIX also forwards the LHC bunch-crossing clock, fixed latency trigger accepts and resets received from the TTC system to front-end electronics.
The FELIX system uses commodity server technology in combination with FPGA-based PCIe I/O cards. The FELIX servers run a software routing platform serving data to network clients. Commodity servers connected to FELIX systems via the same network run innovative multi-threaded software for event fragment building, processing, buffering and forwarding.
This presentation will describe the design and status of the FELIX based readout for the Run 3 upgrade, during which a subset of the detector will be migrated. It will also show how the same concept has been successfully introduced into the demonstrator test bench of the ATLAS Pixel Inner Tracker, acting as a proof of concept towards the longer term Run 4 upgrade in which all remaining detectors will adopt a FELIX based readout.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 307

In-situ analysis and visualization of massively parallel computations of transitional and turbulent flows

Authors: Anne Cadiou[1]; Marc Buffat[2]; Lionel Le Penven[1]; Bastien Di Pierro[2]; Frédéric Alizard [2]; Christophe Pera[2]

[1] *CNRS*

[2] *Université Claude Bernard Lyon 1*

Corresponding Author: anne.cadiou@ec-lyon.fr

Massively parallel simulations generate increasing volumes of large data, whose exploitation requires large storage resources, efficient network and increasingly large post-processing facilities. In the coming era of exascale computations, there is an emerging need for new data analysis and visualization strategies.

Data manipulation, during the simulation and after, considerably slows down the analysis process, now becoming the bottleneck of high performance computing. The traditional usage consists in performing the simulations in order to write output data on disk. When dealing with three-dimensional time-dependent problems computed on thousands of cores, the volume of data generated is big and highly partitioned. As a consequence, their post-processing often requires to decrease the spatial or the time resolution in order to be performed on local platform, with less resources than on the computational machine. Another solution consists in coupling analysis with simulation, so that both are performed simultaneously.

In order to address these questions, a client-server in-situ analysis for massively parallel time-evolving computations has been developed and applied to a spectral code for the study
of turbulence and transition. It is shown to have a low impact on computational time with a reasonable increase of resource usage, while enriching data exploration. Large time sequences have been analyzed. This could not have been achieved with the traditional workflow. Moreover, computational steering has been performed with real-time adjustment of the simulation parameters, thereby getting closer to a numerical experiment process.

Track 2: Data Analysis - Algorithms and Tools / 308

A hybrid deep learning approach to vertexing

Authors: Rui Fang[1]; Henry Fredrick Schreiner[1]; Michael David Sokoloff[1]; Constantin Niko Weisser[2]; J Michael Williams[2]

[1] *University of Cincinnati (US)*

[2] *Massachusetts Inst. of Technology (US)*

Corresponding Authors: mwill@mit.edu, weisser@cern.ch, rui.fang@cern.ch, henry.fredrick.schreiner@cern.ch, michael.david.sokoloff@cern.ch

In the transition to Run 3 in 2021, LHCb will undergo a major luminosity upgrade, going from 1.1 to 5.6 expected visible Primary Vertices (PVs) per event, and will adopt a purely software trigger. This has fueled increased interest in alternative highly-parallel and GPU friendly algorithms for tracking and reconstruction. We will present a novel prototype algorithm for vertexing in the LHCb upgrade conditions.

We use a custom kernel to transform the sparse 3D space of hits and tracks into a dense 1D dataset, and then apply Deep Learning techniques to find PV locations. By training networks on our kernels using several Convolutional Neural Network layers, we have achieved better than 90% efficiency with no more than 0.2 False Positives (FPs) per event. Beyond its physics performance, this algorithm also provides a rich collection of possibilities for visualization and study of 1D convolutional networks. We will discuss the design, performance, and future potential areas of improvement and study, such as possible ways to recover the full 3D vertex information.

Track 1: Computing Technology for Physics Research / 309

Machine Learning Techniques in the ATLAS TDAQ Network Monitoring System

Authors: Oskar Wyszynski[1]; Eukeni Pozo Astigarraga[1]

[1] *CERN*

Corresponding Author: oskar.wyszynski@cern.ch

Network monitoring is of great importance for every data acquisition system (DAQ), it ensures stable and uninterrupted data flow. However, when using standard tools such as Icinga, often homogeneity of the DAQ hardware is not exploited.
We will present the application of machine learning techniques to detect anomalies among network devices as well as connection instabilities. The former exploits homogeneity of network hardware to detect device anomalies such as too high CPU or memory utilization, and consequently uncover a pre-failure state. The latter algorithm learns to distinguish between port speed instabilities caused by, e.g. failing transceiver or fiber, and speed changes due to scheduled system reboots.
All the algorithms described are implemented in the DAQ network of the ATLAS experiment.

Poster Session / 310

Track propagation for different detector and magnetic field setups in Acts

Author: Fabian Klimpel[1]

[1] *Max-Planck-Institut fur Physik (DE)*

Corresponding Author: fabian.klimpel@cern.ch

Track finding and fitting are amongst the most complex part of event reconstruction in high-energy physics, and dominates usually the computing time in high luminosity environment. A central part of track reconstruction is the transport of a given track parameterisation (i.e. the parameter estimation and associated covariances) through the detector, respecting the magnetic field setup and the traversed detector material. While a track propagation in a sparse environment (e.g. a tracking detector) can be sufficiently good approximated by considering discrete interactions at several positions, the propagation in a material dense environment (e.g. calorimeters) is better served by a continuous application of material effects. Recently, a common Tracking software project (Acts) born initially from the Common Tracking code of the ATLAS experiment has been developed in order to preserve the algorithmic concepts from the LHC start-up era and prepare them for the high luminosity era of the LHC and beyond. The software is designed in an abstract, detector independent way and prepared to allow highly parallelised execution of all involved software modules, including magnetic field access and alignment conditions. Therefore the propagation algorithm needs to be as flexible and adjustable which will be the main focus of this talk. The implemented solution for using a fourth order Runge-Kutta-Nyström integration and its extension with continuous material integration and eventual time propagation is presented, such as the navigation through different geometry setups involving different environments are shown.

Poster Session / 311

The design and performance of the ATLAS Inner Detector trigger in high pile-up collisions at 13 TeV at the Large Hadron Collider

Author: Nicola Louise Abraham[1]

[1] *University of Sussex (GB)*

Corresponding Author: nicola.abraham@cern.ch

The design and performance of the ATLAS Inner Detector (ID) trigger algorithms running online on the high level trigger (HLT) processor farm for 13 TeV LHC collision data with high pile-up are discussed.

The HLT ID tracking is a vital component in all physics signatures in the ATLAS trigger for the precise selection of the rare or interesting events necessary for physics analysis without overwhelming the offline data storage in terms of both size and rate. To cope with the high expected interaction rates in the 13 TeV LHC collisions the ID trigger was redesigned during the 2013-15 long shutdown. The performance of the ID trigger in Run 2 from 13 TeV LHC collisions exceeded expectations as the pile-up increased throughout the run periods.

The detailed efficiencies and resolutions of the trigger in a wide range of physics signatures spanning the entire Run 2 production luminosity data-taking are presented, to demonstrating that the trigger responded well under the extreme pile-up conditions. The performance of the ID trigger algorithms in ever higher pile-up collisions illustrates how the ID trigger continued to enable the ATLAS physics program and will continue to do so in the future.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 312

Riemann-Theta Boltzmann Machine

Author: Stefano Carrazza[1]

[1] *CERN*

Corresponding Author: stefano.carrazza@cern.ch

We present a novel general Boltzmann machine with continuous visible and discrete integer valued hidden states, yielding a parametric density function involving a ratio of Riemann-Theta functions. After a brief overview of the theory required to define this new ML architecture, we show how the conditional expectation of a hidden state for given visible states can be used as activation function in a feedforward neural network, thereby increasing the modeling capacity of the network. We then provide application examples for density estimation, data regression and data classification in HEP. This work is based on arXiv:1712.07581 and arXiv:1804.07768.

Track 1: Computing Technology for Physics Research / 313

The second generation of the ATLAS Production System: expertise and future evolution

Authors: Misha Borodin[1]; Alexei Klimentov[2]; Dmitry Golubkov[3]; Mikhail Titov[4]; Tadashi Maeno[2]; Fernando Harald Barreiro Megino[5]

[1] *University of Iowa (US)*

[2] *Brookhaven National Laboratory (US)*

[3] *Institute for High Energy Physics of NRC Kurchatov Institute (R*

[4] *National Research Centre Kurchatov Institute (RU)*

[5] *University of Texas at Arlington*

Corresponding Authors: mikhail.borodin@cern.ch, dmitry.v.golubkov@cern.ch, tmaeno@bnl.gov, fernando.harald.barreiro.megino@cern.ch, alexei.klimentov@cern.ch, mikhail.titov@cern.ch

ATLAS production system called ProdSys2 is used during Run2 to define and to organize workflows and to schedule, submit and execute payloads in a distributed computing infrastructure. We design ProdSys2 to manage all ATLAS workflows: data (re)processing, MC simulation, physics groups analysis objects production, High Level Trigger processing, SW release building and user analysis. It simplifies the life of ATLAS scientists by offering a web-based user interface with rich options, which implements a user-friendly environment for workflow management, such as a simple way of combining different data flows, and real-time monitoring, optimised for use with a huge amount of information to present. We present an overview of the ATLAS Production System technical implementation: job and task definitions, workflow manager and the web-based user interface. We show how it interfaces to different computing resources, for instance, HPC systems and clouds, and how the Production System interacts with the ATLAS data management system. We describe important technical design decisions, work experience acquired during the LHC Run2 and how will the Production System evolve to be used in the heterogeneous computing environment foreseen for Run3 and High-Luminosity LHC.

Poster Session / 314

Boosted decision trees in the classification tasks with the Telescope Array surface detector data

Author: Yana Zhezher[1]

Co-authors: Grigory Rubtsov [2]; Sergey Troitsky [3]; Maxim Piskunov ; Mikhail Kuznetsov

[1] *INR RAS & Lomonosov MSU*

[2] *INR RAS*

[3] *Russian Academy of Sciences (RU)*

Corresponding Authors: zhezher.yana@physics.msu.ru, sergei.troitsky@cern.ch, grisha@ms2.inr.ac.ru, maxitg@gmail.com

The Telescope Array experiment, located in Utah, USA, is aimed to the ultra-high-energy cosmic rays study with the detection of the extensive air showers (EAS). The surface detector of the Telescope Array provides multivariate data reconstructed from the waveforms of signals of the detectors which took part in a particular event. Moreover, a number of variables are composition-sensitive and may be used for the determination of the type of the primary particle. We employ the Boosted Decision Trees (BDT) technique available as the part of the ROOT::TMVA package in the mass composition study and the neutrino search with the Telescope Array surface detector. The classifier is trained with Monte-Carlo modellings: for the mass composition study proton MC set is used as background events and iron MC set is used as signal events. For the neutrino search, neutrino set is used as signal events and highly-inclined proton set is used as background events. The method was optimized to get the best separation between different primaries; the results of this approach will be presented.

Track 1: Computing Technology for Physics Research / 315

Design Pattern for Analysis Automation on Interchangeable, Distributed Resources using Luigi Analysis Workflows

Authors: Marcel Rieger[1]; Martin Erdmann[1]

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

Corresponding Authors: marcel.rieger@cern.ch, martin.erdmann@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 the luigi analysis workflow (law) Python package which is based on the open-source pipelining tool luigi, originally developed by Spotify. It establishes a generic design pattern for analyses of arbitrary scale and complexity, and shifts the focus from executing to defining the analysis logic. Law provides the building blocks to seamlessly integrate with interchangeable remote resources without, however, limiting itself to a specific choice of infrastructure. In particular, it introduces the paradigm of complete separation between analysis algorithms on the one hand, and run locations, storage locations, and software environments on the other hand.
To cope with the sophisticated demands of end-to-end HEP analyses, law supports job execution on WLCG infrastructure (ARC, gLite) as well as on local computing clusters (HTCondor, LSF), remote file access via most common protocols through the Grid File Access Library (GFAL2), and an environment sandboxing mechanism with support for Docker and Singularity containers. Moreover, the novel approach ultimately aims for analysis preservation out-of-the-box.
Law is developed open-source and entirely experiment independent. It is successfully used in ttH cross section measurements and searches for di-Higgs boson production with the CMS experiment.

Poster Session / 316

Data Sampling challenges in the ALICE O2 distributed processing system

Authors: Piotr Konopka[1]; Barthelemy Von Haller[2]

[1] *AGH University of Science and Technology (PL)*

[2] *CERN*

Corresponding Authors: piotr.jan.konopka@cern.ch, barthelemy.von.haller@cern.ch

The ALICE experiment at the CERN LHC focuses on studying the quark-gluon plasma produced by heavy-ion collisions. After the Long Shutdown 2 in 2019-2020, the ALICE Experiment will see its data input throughput increase a hundredfold, up to 3.4 TB/s. In order to cope with such a large amount of data, a new online-offline computing system, called O2, will be deployed. By reconstructing the data online, it will be possible to compress the data stream down to 90GB/s before storing it permanently.

One of the key software components of the system will be the data Quality Control (QC) that replaces the existing online Data Quality Monitoring and offline Quality Assurance. It is this framework and infrastructure, which will be responsible for all aspects related to the analysis software aimed at identifying possible issues with the data itself, and indirectly with the underlying processing done both synchronously and asynchronously. Since analyzing the full stream of data online would exceed the available computational resources, a reliable and efficient sampling will be needed. It should provide a few percent of data selected randomly in a statistically sound manner with a minimal impact on the main dataflow. Extra requirements include the option to choose messages corresponding to the same events over a group of computing nodes and the optional possibility to ensure getting a fixed amount of data at the cost of blocking the main dataflow.

In this paper we present the design of the O2 Data Sampling software. In particular, we highlight our requirements for pseudo-random number generators to be used for sampling decisions, as well as the results of the benchmark we performed to evaluate different possibilities. Finally we report on a large scale test of the O2 Data Sampling we carried out.

Track 2: Data Analysis - Algorithms and Tools / 317

Recurrent GANs for particle-based simulation at the LHC

Authors: Jean-Roch Vlimant[1]; Maurizio Pierini[2]; Jesus Arjona Martinez[1]; Maria Spiropulu[3]; Thong Nguyen[1]

[1] *California Institute of Technology (US)*

[2] *CERN*

[3] *California Institute of Technology*

Corresponding Authors: jesus.arjona.martinez@cern.ch, thong@caltech.edu, maurizio.pierini@cern.ch, jean-roch.vlimant@cern.ch, smaria@caltech.edu

Generative models, and in particular generative adversarial networks, are gaining momentum in hep as a possible way to speed up the event simulation process. Traditionally, gan models applied to hep are designed to return images. On the other hand, many applications (e.g., analyses based on particle flow) are designed to take as input lists of particles. We investigate the possibility of using recurrent GANs as a generator of particle lists. We discuss a prototype implementation, challenges and limitations in the context of specific applications.

Track 2: Data Analysis - Algorithms and Tools / 318

Incorporation of Systematic Uncertainties in the Training of Multivariate Methods

Author: Thomas Alef[1]

Co-author: Eckhard Von Torne [1]

[1] *University of Bonn (DE)*

Corresponding Authors: evt@physik.uni-bonn.de, thomas.alef@cern.ch

Multivariate analyses in particle physics often reach a precision such that its uncertainties are dominated by systematic effects. While there are known strategies to mitigate systematic effects based on adversarial neural nets, the application of Boosted Decision Trees (BDT) so far had to ignore systematics in the training.
We present a method to incorporate systematic uncertainties into a BDT, the "systematics-aware BDT" (saBDT).
We evaluate our method on open data of the ATLAS Higgs to tau tau machine learning challenge and compare our results to neural nets trained with an adversary to mitigate systematic effects.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 319

Introduction to the Full Bandwidth Amplitude Analysis Software (FALLS)

Author: Hao Cai[1]

Co-authors: Kai Zhu ; Ran Zhang

[1] *Wuhan University (CN)*

Corresponding Authors: phyzr96@gmail.com, zhuk@ihep.ac.cn, hao.cai@cern.ch

Amplitude analysis is an important tool for the research of the hadron spectrum, in which the maximum likelihood method is used to estimate the parameters of a probability density function. In each optimization step, the likelihood values of a huge number of events from both data and Monte-Carlo simulations are calculated and summed, which is the most time-consuming part of the whole optimization process. We have tried to increase their calculation speed with CUDA and finally found that it is limited by the bandwidth of GPU memory. On the Tesla V100 platform, we nearly reach the maximum limit of the bandwidth utilization, for which our analysis tool is named as Full Bandwidth Amplitude Analysis Software (FALLS). In the application to the high energy physics analysis work, the average consuming time of the whole process is reduced to several minutes while our original program also based on GPU needs several hours.

Poster Session / 320

Tagging Higgs bosons at the LHC with Interaction Networks

Authors: Jean-Roch Vlimant[1]; Javier Mauricio Duarte[2]; Maria Spiropulu[1]; Maurizio Pierini[3]; Thong Nguyen[1]; Eric Moreno[4]; Joosep Pata[1]; Avikar Periwal[1]; Olmo Cerri[1]

[1] *California Institute of Technology (US)*

[2] *Fermi National Accelerator Lab. (US)*

[3] *CERN*

[4] *California Institute of Technology*

Corresponding Authors: olmo.cerri@cern.ch, jean-roch.vlimant@cern.ch, joosep.pata@cern.ch, thong@caltech.edu, maurizio.pierini@cern.ch, javier.mauricio.duarte@cern.ch, maria.spiropulu@cern.ch, emoreno@caltech.edu, avikar.periwal@cern.ch

We show how Interaction Networks could be used for jet tagging at the Large Hadron Collider.
We take as an example the problem of identifying high-pT H->bb decays exploiting both jet substructure and secondary vertices from b quarks. We consider all tracks produced in the hadronization of the two b's and represent the jet both as a track-to-track and a track-to-vertex interaction. The representations of the two interactions are learned training two dense neural networks. The derived information is used to train a classifier of H->bb jets. Interaction networks achieve state-of-the-art discrimination performances, even when the training is prevented from learning to exploit the jet mass value as discriminating information.

Poster Session / 321

Pileup mitigation with graph neural networks

Authors: Jean-Roch Vlimant[1]; Maurizio Pierini[2]; Olmo Cerri[1]; Jesus Arjona Martinez[1]; Maria Spiropulu[3]

[1] *California Institute of Technology (US)*

[2] *CERN*

[3] *California Institute of Technology*

Corresponding Authors: jesus.arjona.martinez@cern.ch, olmo.cerri@cern.ch, smaria@caltech.edu, jean-roch.vlimant@cern.ch, maurizio.pierini@cern.ch

Mitigation of the effect of the multiple parasitic proton collisions produced during bunch crossing at the LHC is a major endeavor towards the realization of the physics program at the collider. The pileup affects many physics observable derived during the online and offline reconstruction. We propose a graph neural network machine learning model, based on the PUPPI approach, for identifying particle coming from pileup and retaining the ones from high-transverse momentum collisions. We show improvement in pileup rejection performance and energy resolution with respect to solutions currently used at the LHC.

Track 2: Data Analysis - Algorithms and Tools / 322

HEP.TrkX Charged Particle Tracking using Graph Neural Networks

Authors: Jean-Roch Vlimant[1]; Steven Andrew Farrell[2]; Alexander Zlokapa[1]; Abhishek Anand[3]; Maria Spiropulu[3]; Paolo Calafiura[2]

[1] *California Institute of Technology (US)*

[2] *Lawrence Berkeley National Lab. (US)*

[3] *California Institute of Technology*

Corresponding Authors: jean-roch.vlimant@cern.ch, smaria@caltech.edu, pcalafiura@lbl.gov, anand@college.harvard.edu, steven.farrell@cern.ch, alexander.z@cern.ch

To address the unprecedented scale of HL-LHC data, the HEP.TrkX project has been investigating a variety of machine learning approaches to particle track reconstruction. The most promising of these solutions, a graph neural network, processes the event as a graph that connects track measurements (detector hits corresponding to nodes) with candidate line segments between the hits (corresponding to edges). This architecture enables separate input features for edges and nodes, ultimately creating a hidden representation of the graph that is used to turn edges on and off, leaving only the edges that form tracks. Due to the large scale of this graph for an entire LHC event, we present new methods that allow the event graph to be scaled to a computationally reasonable size. We report the results

of the graph neural network on the TrackML dataset, detailing the effectiveness of this model on event data with large pileup. Additionally, we propose post-processing methods that further refine the result of the graph neural network, ultimately synthesizing an end-to-end machine learning solution to particle track reconstruction.

Poster Session / 323

TAG based data analysis in BESIII offline sofware

Authors: ZIYAN DENG[1]; SHENGSEN SUN[1]; XIAN XIONG[1]

[1] *Institute of High Energy Physics, Beijing, China*

Corresponding Authors: xiongxa@ihep.ac.cn, sunss@ihep.ac.cn, dengzy@ihep.ac.cn

BESIII experiment studies physics in the tau-charm energy region. Since 2009, BESIII has collected large scale data samples and many important physics results have been achieved based on these samples. Gaudi is used as BESIII offline software underlying framework, for both data production and data analysis. As data set accumulated year by year, efficiency of data analysis becomes more and more important. This presentation will report tag based analysis implemented in BESIII offline software.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 324

reFORM: designing a new symbolic manipulation toolkit

Author: Ben Ruijl[1]

[1] *Nikhef*

Corresponding Author: benruyl@gmail.com

Over the last few years manipulating expressions with millions of terms has become common in particle physics. Form is the de facto tool for manipulations of extremely large expressions, but it comes with some downsides. In this talk I will discuss an effort to modernize aspects of Form, such as the language and workflow, and the introduction of bindings to C and Python. This new tool is written in Rust and is called reFORM

Poster Session / 326

Probing Neural Networks for the Gamma/Hadron Separation of CTA

Authors: Etienne Lyard[1]; Roland Walter[1]; Vitalii Sliusar[2]; Nicolas Produit[3]

[1] *University of Geneva*

[2] *Université de Genève*

[3] *Universite de Geneve (CH)*

Corresponding Authors: vitalii.sliusar@unige.ch, etienne.lyard@unige.ch, roland.walter@unige.ch, nicolas.produit@cern.ch

The Cherenkov Telescope Array (CTA) will be the largest ground-based, gamma-ray observatory. CTA will detect the signature of gamma rays and cosmic rays hadrons and electrons interacting with the earth atmosphere. Making the best possible use of this facility requires to be able to separate events generated by gamma rays from the particle-induced background. Deep neural networks produced encouraging results, but so far there has been no evaluation of their performance for gamma/hadron separation with respect to well established approaches. In this paper we compare convolutional neural networks and a standard analysis technique, namely boosted decision trees. We compare the performance of the two techniques as applied to simulated observation data. We then looked at the Receiver Operating Characteristics (ROC) curves produced by the two approaches and discuss the similarities and differences between both. We find that neural networks outperformed classical techniques under specific conditions.

Track 2: Data Analysis - Algorithms and Tools / 327

Particle Identification In PICO Using Semi-supervised Learning

Authors: Brendon Matusch[None]; Gevy Cao[1]

[1] *Queen's University*

Corresponding Authors: gevy.cao@queensu.ca, brendon-m@outlook.com

PICO is a dark matter experiment using superheated bubble chamber technology. One of the main analysis challenges in PICO is to unambiguously distinguish between background events and nuclear recoil events from possible WIMP scatters. The conventional discriminator, acoustic parameter (AP), utilizes frequency analysis in Fourier space to compute the acoustic power, which is proven to be different for alpha and nuclear recoils. In a recent machine learning development, an intern collaborator demonstrated extremely powerful discriminators using semi-supervised learning. I will be presenting the results he achieved, and provide an outlook for machine learning in future analysis.

Poster Session / 329

Controls and Configuration software for the ATLAS DAQ system in LHC Run 2

Authors: Andrei Kazarov[1]; Alina Corso Radu[2]

[1] *NRC Kurchatov Institute PNPI (RU)*

[2] *University of California Irvine (US)*

Corresponding Authors: andrei.kazarov@cern.ch, alina.radu@cern.ch

The ATLAS experiment at the Large Hadron Collider (LHC) operated successfully from 2008 to 2018, which included Run 1 (2008-2013), a shutdown period and the Run 2 (2016-2018). In the course of the Run 2, the ATLAS data taking achieved an overall data taking efficiency of 97%, largely constraint by the irreducible dead-time introduced to accommodate the limitations of the detector read-out electronics. Less than 1% of the dead-time could be attributed to the central trigger and DAQ system, and out of these, a negligible fraction was due to the Controls and Configuration sub-system. The first long LHC shutdown (LS1) (2014-2015) was used to carry out a complete revision of the Controls and Configuration software, in order to suitably accommodate additional requirements that could not be seamlessly included during steady operation of the system. As well a refactorization of the software was carried out, software that had been repeatedly modified to include new features becoming less maintainable. Additionally, LS1 was the opportunity of modernizing software written at the beginning of the years 2000, thus profiting from the rapid evolution in IT technologies. This upgrade was carried out retaining the critical constraint of minimally impacting public APIs, and the operation mode of the system, in order to maximize the acceptance of the changes by the large user community. This paper summarizes and illustrates, at hand of a few selected examples, how the

work was approached and which new technologies were introduced into the ATLAS DAQ system and were used in the course of the LHC Run 2. Despite these being specific to the system, many solutions can be considered and adapted to different distributed DAQ systems. Additionally, this paper will focus on the behavior of the Controls and Configuration services through the whole Run 2 period, putting particular emphasis on robustness, reliability and performance matters.

Track 2: Data Analysis - Algorithms and Tools / 330

Global fits of BSM physics models with Gambit

Author: Andy Buckley[1]

[1] *University of Glasgow (GB)*

Corresponding Author: a.g.buckley@gmail.com

The Gambit collaboration is a new effort in the world of global BSM fitting – the combination of the largest possible set of observational data from across particle, astro, and nuclear physics to gain a synoptic view of what experimental data has to say about models of new physics. Using a newly constructed, open source code framework, Gambit have released several state-of-the-art scans of large BSM-model parameter spaces, which have revealed structures masked by the Simplified Model approach that dominates LHC collaborations' in-house data interpretations. I will present the publicly available Gambit framework for marshalling physics calculations and assembling composite likelihoods – including its use of OpenMP and MPI parallelisation, and novel scanning algorithms – as well as headline results from Gambit's programme of BSM data recasting.

Poster Session / 331

A New Visual Analytics toolkit for ATLAS metadata

Authors: Aleksandr Alekseev[1]; Alexei Klimentov[2]; Igal Milman[3]; Maria Grigoryeva[4]; Mikhail Titov[5]; Siarhei Padolski[6]; Tatiana Korchuganova[1]; Timofei Galkin[3]; Victor Pilyugin[3]

[1] *National Research Tomsk Polytechnic University (RU)*

[2] *Brookhaven National Laboratory (US)*

[3] *National Research Nuclear University "MEPhI"*

[4] *Institute for Theoretical and Experimental Physics (RU)*

[5] *National Research Centre Kurchatov Institute (RU)*

[6] *BNL*

Corresponding Authors: alexei.klimentov@cern.ch, tatiana.korchuganova@cern.ch, mikhail.titov@cern.ch, aleksandr.alekseev@cern.ch, siarhei.padolski@cern.ch, maria.grigorieva@cern.ch

The ATLAS experiment at the LHC has a complex heterogeneous distributed
computing infrastructure, which is used to process and analyse exabytes of data. Metadata are collected and stored at all stages of physics analysis and data processing. All metadata could be divided into operational metadata to be used for the quasi on-line monitoring, and archival to study the systems' behaviour over a given period of time (i.e., long-term data analysis). Ensuring the stability and efficiency of functioning of complex and large-scale systems, such as those in ATLAS computing, requires sophisticated monitoring tools, and the long-term monitoring data analysis becomes as important as the monitoring itself. Archival metadata, containing a lot of metrics (hardware and software environment descriptions, network state, application parameters, user account data, errors) accumulated for more than decade, can be successfully processed by various machine learning (ML) algorithms for classification, clustering and dimensionality reduction. However, the ML data analysis, despite the massive use, is not without shortcomings: the underlying algorithms are usually

treated as "black boxes", as there are no effective techniques for understanding their internal mechanisms, and the domain-experts involvement in the process of ML data analysis is very limited. As a result the data analysis suffers from the lack of human supervision. Moreover, sometimes the conclusions made by the algorithms with a high accuracy may have no sense regarding the real data model.
In this work we will demonstrate how the interactive data visualization can be applied to extend the routine ML data analysis methods. Visualization allows to actively use human spatial thinking to identify new tendencies and patterns found in the collected data, avoiding the necessity of struggling with the instrumental analytics tools.
The architecture and the interface prototype of visual analytics platform (VAP) for the multidimensional data analysis of ATLAS computing metadata will be presented. The general data processing and visualization methods of the VAP prototype will be implemented and tested on the slice of ATLAS jobs metadata. As a result, a web-interface will provide ATLAS jobs interactive visual clusterization and search for non-trivial behaviour and its possible reasons. Furthermore, we will demonstrate the prototype of dynamic interactive visualization, providing the possibility of the observation of changing clustering structure in different points in time.

Track 1: Computing Technology for Physics Research / 332

Parallel computing of Overlapping Grids Applications based on A Patch-based Communication Algorithm

Author: Hong Guo[1]

[1] *Institute of Applied Physics and Computational Mathematics*

Corresponding Author: guo_hong@iapcm.ac.cn

The overlapping grid technique can be used to solve partial differential equations defined on complex computational domains. However, large-scale realistic applications using overlapping grid technique under distributed memory systems are not easy. The grid points do not meet point by point and interpolation is needed. Applications with millions of grid points may consist of many blocks. A proper method that specifies the connectivity among so many blocks may be a challenge in large-scale simulations. Furthermore, a domain-decomposition method is always used in parallel computing. A proper communication schedule among a large number of subdomains may be another challenge.

In this talk, we describe a Communication Algorithm for Overlapping Grids Applications with patch-based data structure. Splitting overlapping grids into patches in order to lead to a more computationally balanced workload. The communication algorithm includes a grid mapping method and a communication schedule. Grid mapping method searches donor interpolating grid points based on patch. A tree with patch boxes as its nodes is designed in order to reduce searching cost. A ghost patch assignment strategy is used to facilitate messages pack/unpack routine. A coarse grained communication schedule is implemented in order to reduce communication latency.

According to our test results, applications based on this communication algorithm can be run efficiently on thousands of cores.

Track 1: Computing Technology for Physics Research / 333

CMS Software and Offline preparation for future runs

Author: Collaboration CMS[None]

Corresponding Authors: tommaso.boccali@cern.ch, stosi@cern.ch

The next LHC Runs, nominally RunIII and RunIV, pose problems to the offline and computing systems in CMS. RunIV in particular will needs completely different solutions, given the current estimates of LHC conditions and Trigger estimates. We want to report on the R&D process CMS has a whole has established, in order to gain insight on the needs and the possible solutions for the 2020+ CMS computing.

Track 1: Computing Technology for Physics Research / 334

The NanoAOD event data format in CMS

Author: Collaboration CMS[None]

Corresponding Author: marco.peruzzi@cern.ch

NANOAOD is an event data format that has recently been commissioned by the CMS Collaboration to serve the needs of a substantial fraction of its physics analyses. The new format is about 20 times more compact than the MINIAOD format and only includes high level physics object information. NANOAOD is easily customisable for development activities, and supports standardised routines for content validation and automated data analysis workflows. The talk will review the current status and perspectives of NANOAOD design and implementation.

Poster Session / 335

An innovative monitoring and maintenance model for the INFN CNAF Tier-1 data center infrastructure.

Author: Pier Paolo Ricci[1]

[1] *INFN CNAF*

Corresponding Author: pierpaolo.ricci@cnaf.infn.it

During the last years we have carried out a renewal of the Building Management System (BMS) software of our data center with the aim of improving the data collection capability. Considering the complex physical distribution of the technical plants and the limits of the actual building hosting our center, a system that simply monitors and collects all the necessary information and provides alarms only in case of major failures has proven to be unsatisfactory. In 2017 we suffered a major flood due to one main water pipeline failure in the public street. After this disastrous event, clearly far beyond our control, we were however forced to reconsider completely the physical site robustness of our building in addition to the current monitoring and alarm system capabilities. It was clear that in some specific cases, alerts should be triggered hours or days before the actual main problem arises in order to allow efficient human intervention and proper escalation process. This paradigm could be easily applied to almost all the infrastructure components in our site, mainly the electric power distribution and continuity systems as well as the whole cooling devices. For this reason, in parallel to a consistent increase of the sensor capillarity of our BMS data collector system, a study of a predictive maintenance approach applicability to our site has been started. Predictive maintenance techniques aims at prevent unexpected infrastructure components failures or major events with the study of the whole monitoring data collection and the creation of appropriate statistical models with the help of big data analysis and machine learning techniques. An improvement in the power distribution unit monitoring in our site and the introduction of a dedicated network of water leak sensors were the first steps for increasing the data collection information at our disposal. In addition, a high definition closed-circuit television (CCTV) system with recording capability was introduced to improve the data center remote surveillance and retrospective problem analysis. With sufficient monitoring statistical information stored in our BMS system a preliminary and exploratory predictive data analysis proof of concept could be constructed. This could lead to the model building phase and the creation of a prototype with the aim of forecasting future infrastructure main failure events and forthcoming error conditions. The general idea is, conceivably, an approach to the predictive maintenance model where it would be possible to introduce scheduled corrective actions

for the purpose of preventing potential failures in the next future and increasing the site overall reliability.

Track 2: Data Analysis - Algorithms and Tools / **336**

Accelerating dark matter search in emulsion SHiP detector by Deep Learning

Authors: Sergey Shirobokov[1]; Andrei Golutvin[2]; Andrey Ustyuzhanin[3]

[1] *Imperial College (GB)*

[2] *Imperial College London*

[3] *Yandex School of Data Analysis (RU)*

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

We investigate the problem of dark matter detection in emulsion detector. Previously we have shown, that it is very challenging but possible to use emulsion films of OPERA-like detector in SHiP experiment to separate electromagnetic showers from each other, thus hypothetically separating neutrino events from dark matter. In this study, we have investigated the possibility of usage of Target Tracker (TT) stations in OPERA-like SHiP detector to identify the energy and position of the initial particle. The idea of such search is that unlike emulsion, TT are online detectors, benefiting of zero events pile up. Thus, online observation of the excess of events with proper energy can be a signal of a dark matter.
Two different approaches were applied: classical, using Gaussian Mixtures and machine learning based on a convolutional neural network with coordinate convolution layers for energy and longitudinal position prediction. Clusterization techniques were used for transverse coordinate estimation. The obtained results are about 25% for energy resolution and about 0.8 cm for position resolution in the longitudinal direction and 1 mm in the transverse direction, without any usage of the emulsion. Obtained results are comparable to the case of multiple showers separation in the emulsion.
The obtained results will be further used to optimise the cost and parameters of the proposed SHiP emulsion detector.

Track 1: Computing Technology for Physics Research / **338**

Continuous Analysis in ATLAS: Running User-Defined Container Images on the Grid

Authors: Lukas Alexander Heinrich[1]; Alessandra Forti[2]; Paul Nilsson[3]; Tadashi Maeno[3]

[1] *New York University (US)*

[2] *University of Manchester (GB)*

[3] *Brookhaven National Laboratory (US)*

Corresponding Authors: tmaeno@bnl.gov, paul.nilsson@cern.ch, lukas.heinrich@cern.ch, alessandra.forti@cern.ch

The ATLAS software infrastructure has undergone several changes towards the adoption of Continuous Integration methodology to develop and test software. The users community can benefit from a CI environment in several ways: they can develop their custom analysis, build and test it using revision control services such as GitLab. By providing targeted official base images ATLAS enables users to also build self-contained Linux container images as part of the CI pipelines, a crucial component for analysis preservation and re-use scenarios such as reinterpretation of searches for Beyond the Standard Model physics (RECAST).

However, so far, the execution of preserved analyses was constrained to dedicated cloud infrastructure and not well-integrated into the wider WLCG computing model, where software distribution has so far relied on a combination of collaboration software distributed via the CVMFS filesystem and user software distributed ad-hoc by the workflow management system.

We describe an integration of containerized workloads into the grid infrastructure enabling users to submit self-authored or externally provided container images. To that end, the pilot process executed on the worker node has been extended to utilize the userspace container runtime singularity to execute such workloads. Further, the PanDA job configuration as well as the user-facing command line interfaces have been adapted to allow a detailed specification of the runtime environment.

Through this work a continuous grid analysis paradigm emerges, in which for each change in the revision control system, an automated pipeline of unit testing, image building and workload submission based on the freshly built image into the grid is triggered thus further streamlining physics analyses.

Poster Session / 339

pyhf: auto-differentiable binned statistical models

Authors: Lukas Alexander Heinrich[1]; Kyle Stuart Cranmer[1]; Matthew Feickert[2]; Giordon Holtsberg Stark[3]

[1] *New York University (US)*

[2] *Southern Methodist University (US)*

[3] *University of California,Santa Cruz (US)*

Corresponding Authors: kyle.cranmer@cern.ch, giordon.holtsberg.stark@cern.ch, matthew.feickert@cern.ch, lukas.heinrich@cern.ch

A large class of statistical models in high energy physics can be expressed a simultaneous measurement of binned observables. A popular framework for such binned analysis is HistFactory. So far the only implementation of the model has been within the ROOT ecosystem, limiting adoption and extensibility. We present a complete and extensible implementation of the HistFactory class of models in python, based on multi-dimensional multiarray (tensor) computations. The implementation allows for a variety of tensor backends, such as numpy, TensorFlow, PyTorch and Dask. Through the latter, likelihoods are expressed as computational graphs suitable for automatic differentiation and hardware-accelerated (e.g. GPU and TPU-based) or distributed computation. We present benchmarks showing significant performance gains compared to the existing implementation. Further the implementation introduces a new serialization format for binned likelihoods suitable for archiving on community repositories such as HepData, providing an attractive interface for reinterpretation of physics analyses.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 340

Numerical calculation of high-order QED contributions to the electron anomalous magnetic moment

Author: Sergey Volkov[None]

Corresponding Author: volkoff_sergey@mail.ru

A high-precision calculation of the electron anomalous magnetic moment requires an evaluation of QED Feynman diagrams up to five independent loops. To make this calculation practically feasible it is necessary to remove all infrared and ultraviolet divergences before integration. A procedure of removing both infrared and ultraviolet divergences in each individual Feynman diagram will be presented. The procedure is based on linear operators that are applied to the Feynman amplitudes of

ultraviolet divergent subdiagrams. The usage of linear operators allows us to avoid residual renormalizations after subtraction of divergences. This procedure leads immediately to finite Feynman parametric integrals. A method of Monte Carlo integration of these Feynman parametric integrands will be presented. The method is based on importance sampling. The probability density function is constructed for each Feynman diagram individually by using some combinatorial information from the diagram. The calculated value of the total contribution of the 5-loop QED Feynman diagrams without lepton loops to the electron anomalous magnetic moment will be presented. This result was obtained by a GPU-based computation on a supercomputer. The calculation provides double-checking of the value. The contributions of nine gauge-invariant classes of 5-loop Feynman diagrams without lepton loops will be presented for the first time. Also, the contributions of some individual 6-loop Feynman diagrams will be given for demonstration of the method.

Poster Session / 341

A Domain Specific IDE for Parallel Application Development

Author: Li Liao[1]

[1] *Institute of Applied Physics and Computational Mathematics,Beijing*

Corresponding Author: liliao@iapcm.ac.cn

The scientific computing community is suffering from a lack of good development tool that can handle well the unique problems of coding for high performance computing. It is much more difficult for domain experts to parallelize inherited serial codes written in FORTRAN which are very common in CSE research field. An automatic parallel programming IDE is developed for rapid development of large-scale numerical simulations. A domain specific graphical user interface is designed ot facilitate the algorithm design and parallel program construction based on current serial codes. A debugging approach based on the application development mechanism is built in to locate parallel errors automatically. Real applications demonstrate that is helpful on developing complex numerical applications and enabling increased software productivity greatly.

Track 2: Data Analysis - Algorithms and Tools / 342

Machine Learning on sWeighted data

Authors: Nikita Kazeev[1]; Maxim Borisyak[1]

[1] *Yandex School of Data Analysis (RU)*

Corresponding Authors: maxim.borisyak@cern.ch, nikita.kazeev@cern.ch

Analysis in high-energy physics usually deals with data samples populated from different sources. One of the most widely used ways to handle this is the sPlot technique. In this technique the results of a maximum likelihood fit are used to assign weights that can be used to disentangle signal from background. Some events are assigned negative weights, which makes it difficult to apply machine learning methods. Loss function becomes unbounded and the underlying optimization problem non-convex. In this contribution, we propose a mathematically rigorous way to apply Machine Learning methods on data with weights obtained by the sPlot. Examples of applications are also shown.

Track 2: Data Analysis - Algorithms and Tools / 343

Electromagnetic calorimeter reconstruction in Belle II

Author: Torben Ferber[1]

[1] *DESY*

Corresponding Author: torben.ferber@desy.de

The Belle II experiment at the SuperKEKB e+e- collider has completed its first-collisions run in 2018. The experiment is currently preparing for physics data taking in 2019. The electromagnetic calorimeter of the Belle II detector consists of 8,736 Thallium-doped CsI crystals with PIN-photodiode readout. Each crystal is equipped with waveform digitizers that allow the extraction of energy, time, and pulse-shape information. The talk will describe the offline reconstruction algorithm and first experience with the data taken in 2018. Further optimizations towards the high-rate data taking and high-dose background environment of Belle II will be discussed. Important steps in this process are improvements of existing regression algorithms for energy and position reconstruction, improvements of neutral and charged particle identification, and refinements to clustering itself using machine learning.

Track 2: Data Analysis - Algorithms and Tools / 344

Selective background Monte Carlo simulation at Belle II

Authors: Thomas Kuhr[None]; James Kahn[1]

[1] *Karlsruhe Institute of Technology (KIT)*

Corresponding Authors: james.kahn@kit.edu, martin.ritter@lmu.de, thomas.kuhr@lmu.de

The Belle II experiment, beginning data taking with the full detector in early 2019, is expected to produce a volume of data fifty times that of its predecessor. With this dramatic increase in data comes the opportunity for studies of rare previously inaccessible processes. The investigation of such rare processes in a high data volume environment requires a correspondingly high volume of Monte Carlo simulations to prepare analyses and gain a deep understanding of the contributing physics processes to each individual study. This presents a significant challenge in terms of computing resource requirements and calls for more intelligent methods of simulation, in particular background processes with very high rejection rates. This work presents a method of predicting in the early stages of the simulation process the likelihood of relevancy of an individual event to the target study using convolutional neural networks. The results show a robust training that is integrated natively into the existing Belle II analysis software framework.

Track 1: Computing Technology for Physics Research / 345

Heterogenous computing for the local reconstruction algorithms of the CMS calorimeters

Author: Andrea Massironi[1]

[1] *CERN*

Corresponding Author: andrea.massironi@cern.ch

The increasing LHC luminosity in Run III and, consequently, the increased number of simultaneous proton-proton collisions (pile-up) pose significant challenges for the CMS experiment. These challenges will affect not only the data taking conditions, but also the data processing environment of CMS, which requires an improvement in the online triggering system to match the required detector performance. In order to mitigate the increasing collision rates and complexity of a single event, various approaches are being investigated. Heterogenous computing resources, recently becoming

prominent and abundant, may be significantly more performant for certain types of workflows. In this work, we investigate implementations of common algorithms targeting heterogenous platforms, such as GPUs and FPGAs. The local reconstruction algorithms of the CMS calorimeters, given their granularity and intrinsic parallelizability, are among the first candidates considered for implementation in such heterogenous platforms. We will present current development status and preliminary performance results. Challenges and various obstacles related to each platform, together with the integration into CMS experiment's framework, will be further discussed.

Track 2: Data Analysis - Algorithms and Tools / 346

Full Event Intepretation at Belle II

Author: William Lawrence Sutcliffe[1]

[1] *Karlsruhe Institute of Technology (DE)*

Corresponding Author: as0015@partner.kit.edu

The Belle II experiment is an e+e- collider experiment in Japan, which begins its main physics run in early 2019. The clean environment of e+e- collisions together with the unique event topology of Belle II, in which an Y(4S) particle is produced and subsequently decays to a pair of B mesons, allows a wide range of physics measurements to be performed which are difficult or impossible at hadron colliders. A critical technique for many of these measurements is tag-side B meson reconstruction, in which one B meson in the event is reconstructed. The Full Event Interpretation is an algorithm which reconstructs tag-side B mesons at Belle II. The algorithm trains multivariate classifiers to classify O(100) unique decay channels, which allows it in turn to reconstruct O(10000) decay chains. This talk presents the algorithm and its performance relative to previous tag-side B meson reconstruction algorithms.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 347

New Features in FeynArts & Friends, and how they got used in FeynHiggs

Author: Thomas Hahn[1]

[1] *MPI f. Physik*

Corresponding Author: hahn@feynarts.de

I give an update on recent developments in FeynArts, FormCalc, and LoopTools, and show how the new features were used in making the latest version of FeynHiggs.

Poster Session / 348

Pixel detector background simulation using generative adversarial networks at Belle II

Author: Matej Srebre[1]

[1] *LMU Munich*

Corresponding Author: matej.srebre@gmail.com

The pixel vertex detector is an essential part of the Belle II experiment, allowing us to determinate the location of particle trajectories and decay vertices. The combined data from the innermost Pixel Vertex Detector (PXD), followed by the Silicon Vertex Detector (SVD), and the outermost Central Drift Chamber (CDC) are crucial in the event reconstruction phase to determine particle types, their tracks, and the decay chain. To model the effect of unwanted background noise on the track reconstruction in simulation, we add simulated or recorded background data to the simulated detector signals from the generated physics process of interest. A large batch of statistically independent samples of background noise is required to not be biased by statistical fluctuations in the background data. However, the data from the more fine-grained PXD alone is high in volume and requires a substantial amount of storage and bandwidth if we were to save all measurements. As an efficient way of producing background noise, we explore the idea of an on-demand noise generator which would produce samples statistically similar to the real background measurements and could be used in the reconstruction phase. We examine Generative Adversarial Networks (GAN) which have been extraordinarily successful in reproducing natural images and different types of datasets in various fields. In this talk, we present our progress with training GAN models on the PXD measurements of background noise, the quality of the generated samples compared to the real background data, and explore the effects of using generated noise in the reconstruction.

Track 1: Computing Technology for Physics Research / 349

RooFit parallelization efforts

Authors: Patrick Bos[1]; Wouter Verkerke[2]; Vince Croft[2]; Inti Pelupessy[3]; Carsten Daniel Burgard[2]; Jisk Attema[4]

[1] *Netherlands eScience Center / Nikhef*

[2] *Nikhef National institute for subatomic physics (NL)*

[3] *Netherlands eScience Center*

[4] *Netherlands eScience center*

Corresponding Authors: j.attema@esciencecenter.nl, carsten.burgard@cern.ch, p.bos@esciencecenter.nl, i.pelupessy@esciencecente
vincent.croft@cern.ch, verkerke@nikhef.nl

RooFit is the statistical modeling and fitting package used in many big particle physics experiments to extract physical parameters from reduced particle collision data, e.g. the Higgs boson experiments at the LHC.
RooFit aims to separate particle physics model building and fitting (the users' goals) from their technical implementation and optimization in the back-end.
In this talk, we outline our efforts to further optimize the back-end by automatically running major parts of user models in parallel on multi-core machines.
A major challenge is that RooFit allows users to define many different types of models, with different types of computational bottlenecks.
Our automatic parallelization framework must then be flexible, while still reducing run-time by at least an order of magnitude, preferably more.
We have performed extensive benchmarks and identified at least three bottlenecks that will benefit from parallelization.
To tackle these and possible future bottlenecks, we designed a parallelization layer that allows us to parallelize existing classes with minimal effort, but with high performance and retaining as much of the existing class's interface as possible.
The high-level parallelization model is a task-stealing approach.
Our multi-process approach uses socket-based communication, originally implemented using a custom built, highly performant bi-directional memory mapped pipe, while currently we are considering switching to ZeroMQ for more flexibilty.
Preliminary results show speed-ups of factor 2 to 20, depending on the exact model and parallelization strategy.

Poster Session / 350

The Belle II Simulation Library

Author: Doris Yangsoo Kim[1]

[1] *Soongsil University*

Corresponding Author: dorisykim@ssu.ac.kr

The SuperKEKB collider and the Belle II experiment have finished the second phase of their runs in 2018, which was an essential step to study the e+ e- beam collisions and prepare for the third phase of the runs. The third phase starts at the beginning of 2019, and it is planned to collected a data sample of 50/ab during the following decade.
The simulation library of the Belle II experiment is based on Geant4, which allows the full detector simulation. In this talk, we will discuss the detailed structure of the simulation library and will review the ongoing projects to optimize the usage of the Geant4 for the sake of the Belle II specific environment.

Poster Session / 351

Parameter tuning of distributed storage system based on reinforcement learning

Authors: Lu Wang[1]; Wentao Zhang[2]; Yaodong Cheng[3]

[1] *Computing Center,Institute of High Energy Physics, CAS*

[2] *Institute of High Energy Physics, Chinese Academy of Science*

[3] *IHEP*

Corresponding Authors: zhangwt@ihep.ac.cn, chyd@ihep.ac.cn, wanglu@ihep.ac.cn

HEP computing is a typical data intensive computing. Performance of distributed storage system, can largely defines the efficiency of HEP data processing and analysis. There is a large number of parameters that can be adjusted in a distributed storage system. The setting of these parameters has a great influence on the performance. At present, these parameters are either set with static values or automatically tuned by some heuristic rules defined by experienced administrators. Diversity of data access patterns and hardware capabilities, interactions between tuning actions, latency between parameter tuning and performance rewarding, and the large parameter search space determines these methods are incapable to cater the performance tuning requirements of a modern data center.
Reinforcement Learning (RL) is a branch of machine learning concerned with how an agent ought to take actions within an environment in order to maximize a certain reward. In recent years, it has many successful applications in areas of robotics and game play. Similarities between parameter tuning and these tasks encourages our idea of implementing a RL based automatic performance tuning method. Therefore, we evaluated this idea by the case of performance tuning of Lustre file system client. We used 17 different performance metrics as "state" and 8 different parameters as "tuning target", IO throughput increase as "reward". In each tuning period, a tuning agent will input the "state" to a deep neural network, and use its inference result as instructions for performance tuning, make the tuning actions on "turning target", and then read the throughput metric as "reward". After that, the "State->Tuning Action->Reward" sequence will be stored in a training database as training sample for online training of the neural network. By repeating these steps, the neural network is gradually empowered with the intelligence to make right tuning decision given an arbitrary new state input. The whole process is unsupervised, while the model can learn to adapt to new working load from its unsuccessful tuning actions.
We implemented three reinforcement learning algorithms: DQN, A2C and PPO with PyTorch (version 0.4) in this case. Experiments show that, in a small test bed with Iozone (version 3.479) workload, this method can increase the throughput by about 30% comparing to static default settings of Lustre (version 2.5.3) as the baseline. In the future, it possible to apply this method to other parameter tuning use cases in the operations of data centers.

Poster Session / 352

Digitization of Cylindrical GEM Detector

Authors: Linghui Wu[1]; Jingyi Zhao[1]; Liangliang Wang[1]

[1] *Institute of High Energy Physics, CAS*

Corresponding Authors: zhaojy@ihep.ac.cn, wulh@ihep.ac.cn, llwang@ihep.ac.cn

The inner drift chamber of the BESIII experiment is encountering an aging problem after running of several years. A Cylindrical Gas Electron Multiplier Inner Tracker (CGEM-IT) has been an important candidate for the upgrade of the inner drift chamber. In order to understand the specific detection behavior of CGEM-IT and to build a digitization model for it, a detailed simulation study with the Garfield++ program has been performed. Distributions related to the Lorentz angle, the diffusion effect, the drift time, the multiplication and the signal induction on the readouts are obtained from this simulation. Based on these studies, a preliminary digitization model is implemented in the BESIII Offline Software System and the preliminary result is reasonable.

Track 2: Data Analysis - Algorithms and Tools / 353

INFERNO: Inference-Aware Neural Optimisation

Author: Pablo de Castro[1]

Co-author: Tommaso Dorigo [1]

[1] *Universita e INFN, Padova (IT)*

Corresponding Authors: tommaso.dorigo@cern.ch, pablo.de.castro.manzano@cern.ch

Complex computer simulations are commonly required for accurate data modelling in many scientific disciplines, including experimental High Energy Physics, making statistical inference challenging due to the intractability of the likelihood evaluation for the observed data. Furthermore, sometimes one is interested on inference drawn over a subset of the generative model parameters while taking into account model uncertainty or misspecification on the remaining nuisance parameters. In this work, we show how non-linear summary statistics can be constructed by minimising inference-motivated losses via stochastic gradient descent such they provided the smallest uncertainty for the parameters of interest. As a use case, the problem of confidence interval estimation for the mixture coefficient in a multi-dimensional two-component mixture model (i.e. signal vs background) is considered, where the proposed technique clearly outperforms summary statistics based on probabilistic classification, which are a commonly used alternative but do not account for the presence of nuisance parameters.

Track 2: Data Analysis - Algorithms and Tools / 355

Energy reconstruction of the ATLAS Tile Calorimeter under high pile-up conditions using the Wiener filter

Authors: ATLAS TileCal Speakers Committee[None]; Dayane Gonçalves[1]

[1] *Universidade Federal de Juiz de Fora*

Corresponding Authors: daolivei@cern.ch, luca.fiorini@cern.ch

The ATLAS experiment records data from the proton-proton collisions produced by the Large Hadron Collider (LHC). The Tile Calorimeter is the hadronic sampling calorimeter of ATLAS in the region $|\eta| < 1.7$. It uses iron absorbers and scintillators as active material. Jointly with the other calorimeters it is designed for reconstruction of hadrons, jets, tau-particles and missing transverse energy. It also assists in muon identification. The energy deposited by the particles in the Tile Calorimeter is read out by approximately 10,000 channels. The signal provided by the readout electronics for each channel is digitized at 40 MHz and its amplitude is estimated by an optimal filtering algorithm. The increase of LHC luminosity leads to signal pile-up that deforms the signal of interest and compromises the amplitude estimation performance. This work presents the proposed algorithm for energy estimation in the Tile Calorimeter under high pile-up conditions during LHC Run 3, named Wiener Filter. The performance of the proposed method is studied under various pile-up conditions and compared with current optimal filtering method using proton-proton collision data and Monte Carlo.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 356

Trilinear Higgs boson coupling variations for di-Higgs production with full NLO QCD predictions in POWHEG

Authors: Gudrun Heinrich[1]; Stephen Philip Jones[None]; Matthias Kerner[2]; Gionata Luisoni[None]; Ludovic Michel Scyboz[3]

[1] *Max Planck Institute for Physics*

[2] *MPP Munich*

[3] *Max-Planck-Institut fur Physik (DE)*

Corresponding Authors: gionata.luisoni@cern.ch, kerner@mpp.mpg.de, ludovic.michel.scyboz@cern.ch, gudrun@mpp.mpg.de, s.jones@cern.ch

While the Higgs boson couplings to other particles are increasingly well-measured by LHC experiments, it has proven difficult to set constraints on the Higgs trilinear self-coupling λ, principally due to the very low cross-section of Higgs boson pair production. We present the results of NLO QCD corrections to Higgs pair production with full top-quark mass dependence, where the fixed-order computation up to two loops has been performed numerically, using both CPUs and GPUs. It is supplemented by parton showering within the `POWHEG` event generator framework. We use the interface between the `POWHEG-BOX-V2` program and both `Pythia8` and `Herwig7` parton showers to generate differential distributions for various values of the trilinear self-coupling λ that are still allowed by the current experimental constraints.

Poster Session / 357

Jupyter Notebook Support of the Belle 2 Software

Authors: Martin Ritter[1]; Nils Braun[2]; Thomas Hauth[3]; Thomas Kuhr[None]

[1] *LMU / Cluster Universe*

[2] *KIT*

[3] *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Authors: martin.ritter@lmu.de, thomas.kuhr@lmu.de, thomas.hauth@kit.edu, nils.braun@kit.edu

The Belle II experiment at the SuperKEKB e+e- collider has completed its first-collisions run in 2018. The experiment is currently preparing for physics data taking in 2019. With many scientists now preparing their analysis, the user friendliness of the Belle II software framework is of great importance.

Jupyter Notebooks allow for mixed code, documentation, and output like plots in a easy to use environment with a low entry barrier for new users. They are ideal for teaching and and exploratory development but can also be used for normal analysis. To use this potential for analysis and training we developed a Jupyter integration for the Belle II software. This integration makes sure everything is setup correctly and automatically configures rich output where appropriate.

We will give an overview of the current status of the integration and its current use in training and analysis as well as possible improvements and future plans.

Track 1: Computing Technology for Physics Research / 358

Automated and Intelligent Data Migration Strategy in High Energy Physical Storage Systems

Authors: Zhenjing Cheng[1]; Yaodong Cheng[2]; Lu Wang[3]; Qingbao Hu[2]; Haibo li[4]; Qi Xu[5]; Gang CHEN[6]

[1] *INSTITUE OF HIGH ENERGY PHYSICS*

[2] *IHEP*

[3] *Computing Center,Institute of High Energy Physics, CAS*

[4] *Institute of High Energy Physics Chinese Academy of Science*

[5] *Institute of High Energy Physics,Chinese Academy of Sciences*

[6] *INSTITUTE OF HIGH ENERGY PHYSICS*

Corresponding Authors: wanglu@ihep.ac.cn, lihaibo@ihep.ac.cn, chengzj@ihep.ac.cn, xuq@ihep.ac.cn, gang.chen@ihep.ac.cn, chyd@ihep.ac.cn, huqb@ihep.ac.cn

As a data-intensive computing application, high-energy physics requires storage and computing for large amounts of data at the PB level. Performance demands and data access imbalances in mass storage systems are increasing. Specifically, on one hand, traditional cheap disk storage systems have been unable to handle high IOPS demand services. On the other hand, a survey found that only a very small number of files have been active in storage for a period of time. Most files have never been accessed. Some enterprises and research organizations are beginning to use tiered storage architectures, such as tape, disk or solid state drives to reduce hardware purchase costs and power consumption.

As the amount of stored data grows, tiered storage requires data management software to migrate less active data to lower cost storage devices. Thus an automated data migration strategy is needed. At present, automatic data migration strategies such as LRU, CLOCK, 2Q, GDSF, LFUDA, FIFO, etc., are usually based on files' recent access mode(such as file access frequency, etc.), are mainly used to resolve data migration between memory and disk. They need to run in the operating system kernel, so the rules are relatively simple. For file access mode does not take file life cycle trend into account, some regularly accessed files are often not predicted accurately. In addition, file history access records are not considered.

Data access requests are not completely random. They are driven by the behavior of users or programs. There must be association between different files that are accessed consecutively. This paper proposes a method of file access heat prediction. Data heat trend is used as the basis for migration to a relatively low-cost storage device. Due to the limitations of traditional models, it is difficult to achieve good results in predicting at such nonlinear scenes. This paper attempts to use the deep learning algorithm model to predict the evolution trend of data access heat. This paper discussed the implementation of some initial parts of the system, in particular the trace collector and the LSTM model. Then some preliminary experiments are conducted with these parts.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 359

Double Higgs boson production in the high- and low-energy limits

Authors: Joshua Davies[1]; Matthias Steinhauser[2]; Go Mishima[1]; David Wellmann[1]

[1] *Karlsruhe Institute of Technology*

[2] *KIT*

Corresponding Authors: joshua.davies@kit.edu, matthias.steinhauser@kit.edu

In this talk, we consider some of the computational aspects encountered in recent computations of double Higgs boson production in gluon fusion. We consider the NLO virtual amplitude in the high-energy limit, and the NNLO virtual amplitude in the low-energy (or large top quark mass) limit. We discuss various optimizations which were necessary to produce our results.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 360

QED and electroweak radiative corrections to polarized Bhabha scattering

Authors: Andrey Arbuzov[1]; Serge Bondarenko[1]; Yahor Dydyshka[2]; Lidia Kalinovskaya[1]; Leonid Rumyantsev[3]; Renat Sadykov[1]

[1] *Joint Institute for Nuclear Research (RU)*

[2] *Byelorussian State University (BY)*

[3] *Joint Inst. for Nuclear Research (JINR)-Unknown-Unknown*

Corresponding Authors: lidia.kalinovskaya@cern.ch, renat.sadykov@cern.ch, leo_rumyantsev@inbox.ru, yahor.dydyshka@cern.ch, bondarenko@jinr.ru, arbuzov@theor.jinr.ru

Complete one-loop electroweak radiative corrections to polarized Bhabha scattering are presented. Higher order QED effects are evaluated in the leading logarithmic approximation. Numerical results are shown for the conditions of future circular and linear electron-positron colliders with polarized beams. Theoretical uncertainties are estimated.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 361

MCSANCee generator with one-loop electroweak corrections for processes with polarized e^+e^- beams

Authors: Andrey Arbuzov[1]; Serge Bondarenko[1]; Yahor Dydyshka[2]; Lidia Kalinovskaya[1]; Ivan Novikov[3]; Leonid Rumyantsev[3]; Renat Sadykov[1]; Vitaly Yermolchyk[3]

[1] *Joint Institute for Nuclear Research (RU)*

[2] *Byelorussian State University (BY)*

[3] *Joint Institute for Nuclear Research*

Corresponding Authors: novivanya@yandex.ru, yahor.dydyshka@cern.ch, lidia.kalinovskaya@cern.ch, vitaly.yermolchyk@cern.ch, bondarenko@jinr.ru, sadykov@cern.ch, arbuzov@theor.jinr.ru, leo12345678901@gmail.com

A new Monte Carlo event generator MCSANCee for simulation of processes at future e^+e^- colliders is presented. Complete one-loop electroweak radiative corrections and polarization of the initial beams are taken into account. The present generator includes the following processes: e^+e^- \to

e^+e^- (mu^+mu^-, tau^+tau^-, ZH, Z\gamma, \gamma\gamma). Numerical results for all of these processes are shown together with tuned comparisons with other existing codes. The plan for the further extension of the MCSANCee generators are discussed.

Poster Session / 362

Production experience and performance study for HEP data production at HPC-TianheII

Authors: Yao Zhang[1]; Ye Yuan[2]

Co-author: Ma Qiumei [3]

[1] *Institiute of High Energy Physics, China*

[2] *Institue of High Energy Physics, China*

[3] *Institute of High Energy Physics, China*

Corresponding Authors: zhangyao@ihep.ac.cn, maqm@ihep.ac.cn, yuanye@ihep.ac.cn

The mass Monte Carlo data production is the most CPU intensive process in the data analysis of for the high energy physics. The use of large scale computational resources at HPC in China is expected to increase substantially the cost-efficiency of the processing. TianheII, the second fastest HPC in China, which used to ranks first in the TOP500. We report on the technical challenges and solutions adopted to migrate offline software to TianheII, and on the experience and measured performance for mass production of COMET and BESIII experiment.

Poster Session / 363

Deep Learning applied to hit classification for BESIII drift chamber

Authors: Yao Zhang [1]; Ye Yuan[1]; Qiumei Ma[1]

[1] *Institute of High Energy Physics, China*

Corresponding Authors: maqm@ihep.ac.cn, yuanye@ihep.ac.cn, zhangyao@ihep.ac.cn

Drift chamber is the main tracking detector for high energy physics experiment like BESIII. Due to the high luminosity and high beam intensity, drift chamber is suffer from the background from the beam and electronics which represent a computing challenge to the reconstruction software. Deep learning developments in the last few years have shown tremendous improvements in the analysis of data especially for object classification. Here we present a first study of deep learning architectures applied to BESIII drift chamber real data to make the hit classification of the background and signal.

Poster Session / 364

In-Memory Parallel Computing for Partial Wave Analysis

Author: Zhanchen Wei[1]

Co-authors: Qiulan Huang [1]; Gongxing Sun [2]; Xiaoyu Liu [1]

[1] *IHEP, CAS*

[2] *INSTITUE OF HIGH ENERGY PHYSICS*

Corresponding Authors: weizc@ihep.ac.cn, liuxiaoyu@ihep.ac.cn, sungx@mail.ihep.ac.cn, huangql@ihep.ac.cn

The traditional partial wave analysis (PWA) algorithm is designed to process data serially which requires a large amount of memory that may exceed the memory capacity of one single node to store runtime data. It is quite necessary to parallelize this algorithm in a distributed data computing framework to improve its performance. Within an existing production-level Hadoop cluster, we implement PWA algorithm on the basis of Spark to process data storing on low-level storage system HDFS. But in this case, sharing data through HDFS or internal data communication mechanism of Spark is extremely inefficient. In order to solve this problem, this paper presents an in-memory parallel computing method for PWA algorithm. With this system, we can easily sharing runtime data in parallel algorithms. We can ensure complete data locality to keep compatibility with the traditional data input/output way and cache most repeated used data in memory to improve the performance, owe to the data management mechanism of Alluxio.

Poster Session / 365

DD4hep: an integrated detector description tool for HEP

Authors: Andre Sailer[1]; Frank-Dieter Gaede[2]; Marko Petric[1]; Markus Frank[1]

[1] *CERN*

[2] *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Authors: frank-dieter.gaede@cern.ch, markus.frank@cern.ch, marko.petric@cern.ch, andre.philippe.sailer@cern.ch

The detector description is an essential component in simulation, reconstruction and analysis of data resulting from particle collisions in high energy physics experiments. The main motivation behind DD4hep is to provide an integrated solution for all these stages and addresses detector description in a broad sense, including the geometry and the materials used in the device, and additional parameters describing e.g. the detection techniques, constants required for alignment and calibration, description of the readout structures and conditions data. A core part of DD4hep is DDG4 which is a powerful tool that converts arbitrary DD4hep detector geometries to Geant4 and gives access to all Geant4 action stages. It is equipped with a comprehensive plugins suite that includes handling of different IO formats, Monte Carlo truth linking and a large set of segmentation and sensitive detector classes, allowing the simulation of a wide variety of detector technologies. Another important segment of the toolkit are DDCond and DDAlign, which expose a mechanism to manage detector data simultaneously for multiple versions depending on their validity. The detector conditions data are made available to the physics algorithms through a number of transient objects grouped to collections. Such a collection represents a coherent slice of all conditions data necessary to process one or several events depending on the interval of validity of the slice, which is the intersection of the individual conditions. A multi-threaded application may hold several such collections in parallel depending on the time-stamps of the events currently processed. In this presentation, we will give an overview of the project and discuss developments in DD4hep as well as showcase recent adaptions of the framework by LHC experiments and beyond.

Track 1: Computing Technology for Physics Research / 366

Core software challenges of the GPU High Level Trigger 1 of LHCb

Corresponding Author: daniel.hugo.campora.perez@cern.ch

The LHCb detector will be upgraded in 2021, and due to the removal of the hardware-level trigger and the increase in the luminosity of the collisions, the conditions for a High Level Trigger 1 in software will become more challenging, requiring processing the full 30 MHz data-collision rate. The GPU High Level Trigger 1 is a framework that permits concurrent many-event execution targeting many-core architectures. It is designed to hide data transmission overhead with a custom memory manager and maximize GPU resource usage employing a static scheduler. We present the core infrastructure of this R&D project on many-core architectures developed in the context of the LHCb Upgrade I. We discuss the design aspects driving it, and present algorithm-specific data layout design and evaluate their impact on performance.

Poster Session / 367

Physics and computing performance of reconstruction algorithms for the GPU High Level Trigger 1 of LHCb

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: dorothea.vom.bruch@cern.ch, kmueller@physik.uzh.ch

Beginning in 2021, the upgraded LHCb experiment will use a triggerless readout system collecting data at an event rate of 30 MHz. A software-only High Level Trigger will enable unprecedented flexibility for trigger selections. During the first stage (HLT1), a sub-set of the full offline track reconstruction for charged particles is run to select particles of interest based on single or two-track selections. After this first stage, the event rate is reduced by at least a factor 30. Track reconstruction at 30 MHz represents a significant computing challenge, requiring a renovation of current algorithms and the underlying hardware. In this talk we present work based on an R&D project in the context of the LHCb Upgrade I exploring the approach of executing the full HLT1 chain on GPUs. This includes decoding the raw data, clustering of hits, pattern recognition, as well as track fitting. We will discuss the development of algorithms optimized for many-core architectures. Both the computing and physics performance of the full HLT1 chain will be presented.

Track 2: Data Analysis - Algorithms and Tools / 368

Adversarial Neural Network-based data-simulation corrections for jet-tagging at CMS

Authors: Benjamin Fischer[1]; David Josef Schmidt[1]; Dennis Noll[1]; Marcel Rieger[2]; Martin Erdmann[2]; Yannik Alexander Rath[3]

[1] *Rheinisch Westfaelische Tech. Hoch. (DE)*

[2] *Rheinisch-Westfaelische Tech. Hoch. (DE)*

[3] *RWTH Aachen University (DE)*

Corresponding Authors: david.schmidt@rwth-aachen.de, yannik.alexander.rath@cern.ch, dennis.noll@rwth-aachen.de, martin.erdmann@cern.ch, benjamin.fischer@cern.ch, marcel.rieger@cern.ch

Variable-dependent scale factors are commonly used in HEP to improve shape agreement of data and simulation. The choice of the underlying model is of great importance, but often requires a lot of manual tuning e.g. of bin sizes or fitted functions. This can be alleviated through the use of neural networks and their inherent powerful data modeling capabilities.

We present a novel and generalized method for producing scale factors using an adversarial neural network. This method is investigated in the context of the bottom-quark jet-tagging algorithms within the CMS experiment. The primary network uses the jet variables as inputs to derive the scale factor for a single jet. It is trained through the use of a second network, the adversary, which aims to differentiate between the data and rescaled simulation.

Poster Session / 369

Tracking performance for long living particles at LHCb

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: luis.miguel.garcia.martin@cern.ch, kmueller@physik.uzh.ch

The LHCb experiment is dedicated to the study of the c- and b-hadrons decays, including long living particles such as Ks and strange baryons (Lambda, Xi, etc…). These kind of particles are difficult to reconstruct from LHCb tracking systems since they escape the detection in the first tracker. A new method to evaluate the performance in terms of efficiency and throughput of the different tracking algorithms for long living particles have been developed. Special emphasis is laid on particles hitting only part of the tracking system of the new LHCb upgrade detector.

Poster Session / 370

Machine Learning approach to boosting neutral particles identification in the LHCb calorimeter

Authors: Katharina Mueller[1]; Alexey Boldyrev[2]

[1] *Universitaet Zuerich (CH)*

[2] *Moscow State University (RU)*

Corresponding Authors: alexey.boldyrev@cern.ch, kmueller@physik.uzh.ch

We present a new approach to identification of boosted neutral particles using electromagnetic calorimeters of the LHCb detector. The identification of photons and neutral pions is currently based on expected properties of the objects reconstructed in the calorimeter. This allows to distinguish single photons in the electromagnetic calorimeter from overlapping photons produced from high momentum pi0 decays. The proposed approach is based on applying machine learning techniques to primary calorimeter information, that are energies collected in individual cells around the energy cluster. The machine learning model employs extreme gradient boosting trees approach which is widely used nowadays, and separates pi0 and photon responses from "first principles". This approach allowed to significantly improve separation performance score on simulated data, reducing primary photons fake rate by factor of four. In this presentation we will present the approach, evaluate its performance obtained on real data, and discuss specific issues when transferring discriminative models from simulation to real world.

Track 2: Data Analysis - Algorithms and Tools / 371

Fast Data-Driven simulation of Cherenkov Detectors Using Generative Adversarial Networks.

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: artem.maevskiy@cern.ch, kmueller@physik.uzh.ch

The increasing luminosities of future LHC runs and next generation of collider experiments will require an unprecedented amount of simulated events to be produced. Such large scale productions are

extremely demanding in terms of computing resources. Thus new approaches to event generation and simulation of detector responses are needed. In LHCb the simulation of the RICH detector using the classical method takes a sizeable fraction of CPU time. We generate high-level reconstruction observables using a generative neural network to bypass low level details. This network is trained to reproduce the particle species likelihoods based on the track kinematic parameters and detector occupancy. The fast simulation is trained using real data samples collected by LHCb during run 2 with the help of sWeight technique. We demonstrate that this approach provides high-fidelity results along with a significant speed increase and discuss possible implication of these results. We also present an implementation of this algorithm into LHCb simulation software and validation tests.

Track 2: Data Analysis - Algorithms and Tools / 372

Reinforced Sorting Networks for Particle Physics Analyses

Authors: Martin Erdmann[1]; Benjamin Fischer[2]; Dennis Noll[2]; Yannik Alexander Rath[3]; Marcel Rieger[3]; David Josef Schmidt[2]; Marcus Wirtz[2]

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

[2] *Rheinisch Westfaelische Tech. Hoch. (DE)*

[3] *RWTH Aachen University (DE)*

Corresponding Authors: mwirtz@physik.rwth-aachen.de, benjamin.fischer@cern.ch, dennis.noll@rwth-aachen.de, marcel.rieger@cern.ch, martin.erdmann@cern.ch, david.schmidt@rwth-aachen.de, yannik.alexander.rath@cern.ch

Deep learning architectures in particle physics are often strongly dependent on the order of their input variables. We present a two-stage deep learning architecture consisting of a network for sorting input objects and a subsequent network for data analysis. The sorting network (agent) is trained through reinforcement learning using feedback from the analysis network (environment). A tree search algorithm is used to examine the large space of different possible orders.

The optimal order depends on the environment and is learned by the agent in an unsupervised approach. Thus, the 2-stage system can choose an optimal solution which is not know to the physicist in advance.

We present the new approach and its application to various classification tasks.

Track 2: Data Analysis - Algorithms and Tools / 373

Variational Dropout Sparsification for Particle Identification speed-up.

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: artem.ryzhikov@cern.ch, kmueller@physik.uzh.ch

Accurate particle identification (PID) is one of the most important aspects of the LHCb experiment. Modern machine learning techniques such as deep neural networks are efficiently applied to this problem and are integrated into the LHCb software. In this research, we discuss novel applications of neural network speed-up techniques to achieve faster PID in LHC upgrade conditions. We show that the best results are obtained using variational dropout sparsification, which provide a prediction speed increase of up to a factor five even when compared to a model with shallow networks.

Poster Session / 374

Machine Learning for Muon Identification at LHCb

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: nikita.kazeev@cern.ch, kmueller@physik.uzh.ch

Particle identification is a key ingredient of most of LHCb results. Muon identification in particular is used at every stage of the LHCb triggers. The objective of the muon identification is to distinguish muons from the rest of the particles using only information from the Muon subdetector under strict timing constraints. We use state-of-the-art gradient boosting algorithm and real data with sWeights to train such a model. In this talk we present the algorithm along with evaluation of its performance across momentum spectrum and different background sources.

Track 2: Data Analysis - Algorithms and Tools / 376

Real-time reconstruction of long-lived particles at LHCb using FPGAs.

Author: Michael J. Morello[1]

[1] *SNS and INFN-Pisa (IT)*

Corresponding Authors: giovanni.punzi@cern.ch, federico.lazzari@cern.ch, riccardo.cenci@cern.ch, andrea.di.luca@cern.ch, michael.joseph.morello@cern.ch

Finding tracks downstream of the magnet at the earliest LHCb trigger level is not part of the baseline plan of the Upgrade trigger, on account of the significant CPU time required to execute the search. Many long-lived particles, such as Ks and strange baryons, decay after the vertex track detector (VELO), so that their reconstruction efficiency is limited. We present a study of the performances of a future innovative real-time tracking system based on FPGAs, R&D developed in the context of the LHCb Upgrade Ib (LHC Run 4), dedicated to reconstructing particle downstream of the magnet in the forward tracking detector (Scintillating Fibre Tracker), that is capable of processing events at the full LHC collision rate (30 MHz).

Track 1: Computing Technology for Physics Research / 377

Real-time cluster finding for LHCb silicon pixel VELO detector using FPGA

Corresponding Author: federico.lazzari@cern.ch

The hardware trigger L0 will be removed in LHCb upgrade I, and the software High Level Trigger have to process event at full LHC collision rate (30 MHz). This is a huge task, and delegating some low-level time-consuming tasks to FPGA accelerators can be very helpful in saving computing time that can be more usefully devoted to higher level tasks. In particular, the 2-D pixel geometry of the new LHCb VELO detector makes the cluster-finding process a particularly CPU-time demanding task. We present here the first results achieved with a highly parallel clustering algorithm implemented in dedicated FPGA cards, developed in an R&D programme in the context of the LHCb Upgrade I, in view of potential future applications.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 378

Updates on SModelS

Authors: Wolfgang Waltenberger[1]; Sabine Kraml[2]; Andre Lessa[3]; Federico Ambrogi[1]; Suchita Kulkarni[1]; Humberto Reyes-Gonzales[2]; Alicia Wongel[4]; Ursula Laa[5]; Juhi Dutta[6]; Jan Heisig[7]; Philipp Neuhuber[4]; Matthias Wolf[4]

[1] *Austrian Academy of Sciences (AT)*

[2] *LPSC Grenoble*

[3] *IFGW - UNICAMP*

[4] *HEPHY Vienna*

[5] *Monash University*

[6] *Harish-Chandra Research Institute,Allahabad*

[7] *RWTH Aachen University*

Corresponding Authors: uschi.laa@gmx.at, jan.heisig@desy.de, sabine.kraml@cern.ch, andlessa@ifi.unicamp.br, suchita.kulkarni@cern.ch, federico.ambrogi@cern.ch, wolfgang.waltenberger@gmail.com, juhidutta@hri.res.in

The software framework SModelS, which has already been presented at the ACAT 2016 conference, allows for a very fast confrontation of arbitrary BSM models exhibiting a Z2 symmetry with an ever growing database of simplified models results from CMS and ATLAS. In this talk we shall present its newest features, like the extension to include searches for heavy stable charged particles (HSCPs), or the ability to combine the results from several signal regions, exploiting the simplified likelihood framework introduced by CMS. Also, the database has been greatly extended; it now comprises almost 3000 individual results from close to 100 individual analyses. Finally, we shall also discuss ongoing developments, like the use of neural networks to further speed up the software, and the extension to an even wider set of experimental signatures.

Track 1: Computing Technology for Physics Research / 379

A 30 MHz software trigger and reconstruction for the LHCb upgrade

Corresponding Author: conor.fitzpatrick@cern.ch

The first LHCb upgrade will take data at an instantaneous luminosity of 2E33cm^{-2}s^{-1} starting in 2021. Due to the high rate of beauty and charm signals LHCb has chosen as its baseline to read out the entire detector into a software trigger running on commodity x86 hardware at the LHC collision frequency of 30MHz, where a full offline-quality reconstruction will be performed. In this talk we present the challenges of triggering in the MHz signal era. We pay particular attention to the need for flexibility in the selection and reconstruction of events without sacrificing performance.

Track 1: Computing Technology for Physics Research / 380

Deep learning for certification of the quality of the data acquired by the CMS experiment

Authors: Adrian Alan Pol[1]; Virginia Azzolini[None]; Gianluca Cerminara[2]; Federico De Guio[3]; Giovanni Franzoni[2]; Cecile Germain[4]; Maurizio Pierini[2]; Tomasz Krzyzek[5]

[1] *Université Paris-Saclay (FR)*

[2] *CERN*

[3] *Texas Tech University*

[4] *Universite Paris Sud*

[5] *Jagiellonian University in Krakow*

Corresponding Authors: gianluca.cerminara@cern.ch, cecile.germain@lri.fr, giovanni.franzoni@cern.ch, adrian.pol@cern.ch, maurizio.pierini@cern.ch, virginia.azzolini@cern.ch, tomasz.krzyzek@cern.ch, federico.de.guio@cern.ch

Certifying the data recorded by the Compact Muon Solenoid (CMS) experiment at CERN which is usable for publication of physics results is a crucial and onerous task. Anomalies caused by detector malfunctioning or sub-optimal data processing are difficult to enumerate a priori and occur rarely, making it difficult to use classical supervised classification. We base out prototype towards the automation of such procedure on a semi-supervised approach using deep autoencoders. We demonstrate the ability of the model to detect anomalies with high accuracy, when compared against the outcome of the fully supervised methods. We show that the model has great interpretability of the results, ascribing the origin of the problems in the data to a specific sub-detector or physics object. Finally, we tailor the approach with a systematic method for feature filtering and address the issue of feature dependency on LHC beam intensity.

Track 1: Computing Technology for Physics Research / 381

Highly performant, Deep Neural Networks with sub-microsecond latency on FPGAs for trigger applications

Authors: Noel Aaron Nottbeck[1]; Volker Buescher[1]; Christian Schmitt[1]

[1] *Johannes Gutenberg Universitaet Mainz (DE)*

Corresponding Authors: volker.buescher@cern.ch, noel.aaron.nottbeck@cern.ch, christian.schmitt@cern.ch

Artificial neural networks are becoming a standard tool for data analysis, but their potential remains yet to be widely used for hardware-level trigger applications. Nowadays, high-end FPGAs, as they are also often used in low-level hardware triggers, offer enough performance to allow for the inclusion of networks of considerable size into these system for the first time. Nevertheless, in the trigger context, it is necessary to highly optimize the implementation of neural networks to make full use of the FPGA capabilities.

We implemented the processing data and control flow of typical NN layers, taking into account incoming data rates of up to multiple tens of MHz and sub-microsecond latency limits, but also aiming at an efficient use of the resources of the FPGA. This resulted in a highly optimized neural network implementation framework, which typically reaches 90 to 100 % computational efficiency, requires few extra FPGA resources for data flow and controlling, and achieves latencies in the order of only tens to few hundreds of nanoseconds for entire (deep) networks. The implemented layers include 2D convolutions and pooling (both with multi-channel support), as well as dense layers, all of which play a role in many physics-/detector-related applications. Significant effort was put especially into the 2D convolutional layers, to achieve a fast implementation with minimal resource usage.

A toolkit is provided which automatically creates the optimized FPGA implementation of trained deep neural network models. Results are presented, both for individual layers as well as entire networks created by the toolkit.

Poster Session / 382

Fips: An OpenGL based FITS viewer

Authors: Konstantin Malanchev[None]; Matwey Kornilov[None]

Corresponding Author: malanchev@physics.msu.ru

We present an open source GPU-accelerated cross-platform FITS 2D image viewer FIPS. Unlike other FITS viewers, FIPS uses GPU hardware via OpenGL to provide functionality such as zooming, panning and level adjustments. FIPS is the first end-to-end GPU FITS image viewer: FITS image data is fully offloaded to GPU memory as is, and then processed by OpenGL shaders.

The executables and the source is available on http://fips.space

Poster Session / 383

Speeding HEP Analysis with ROOT Bulk I/O

Authors: Brian Paul Bockelman[1]; Zhe Zhang[2]; Oksana Shadura[2]

[1] *University of Nebraska-Lincoln (US)*

[2] *University of Nebraska Lincoln (US)*

Corresponding Authors: oksana.shadura@cern.ch, zhan0915@huskers.unl.edu, brian.bockelman@cern.ch

Distinct HEP workflows have distinct I/O needs; while ROOT I/O excels at serializing complex C++ objects common to reconstruction, analysis workflows typically have simpler objects and can sustain higher event rates. To meet these workflows, we have developed a "bulk I/O" interface, allowing multiple events' data to be returned per library call. This reduces ROOT-related overheads and increases event rates – orders-of-magnitude improvements are shown in microbenchmarks.

Unfortunately, this bulk interface is difficult to use as it requires users to identify when it is applicable and they still "think" in terms of events, not arrays of data. We have integrated the bulk I/O interface into the new RDataFrame analysis framework inside ROOT. As RDataFrame's interface can provide improved type information, the framework itself can determine what data is readable via the bulk IO and automatically switch between interfaces. We demonstrate how this can improve event rates when reading analysis data formats, such as CMS's NanoAOD.

Poster Session / 384

ROOT I/O compression algorithms and their performance impact within Run 3

Authors: Oksana Shadura[1]; Brian Paul Bockelman[2]

[1] *University of Nebraska Lincoln (US)*

[2] *University of Nebraska-Lincoln (US)*

Corresponding Authors: oksana.shadura@cern.ch, brian.bockelman@cern.ch

The LHC's Run3 will push the envelope on data-intensive workflows and, at the lowest level, this data is managed using the ROOT software framework. At the beginning of Run 1, all data was compressed with the ZLIB algorithm: ROOT has since added support for multiple new algorithms (such as LZMA and LZ4), each with unique strengths. Work is continuing as industry introduces new techniques - ROOT can benefit saving disk space or reducing the I/O and bandwidth for online and offline needs of experiments by introducing better compression algorithms. In addition to alternate algorithms, we have been exploring alternate techniques to improve parallelism and apply pre-conditioners to the serialized data.

We have performed a survey of the performance of the new compression techniques. Our survey includes various use cases of data compression of ROOT files provided by different LHC experiments.

We also provide insight into solutions applied to the bottlenecks in compression algorithms for the improved ROOT performance.

Poster Session / 385

Evolution of ROOT package management

Authors: Oksana Shadura[1]; Vasil Georgiev Vasilev[2]; Brian Paul Bockelman[3]

[1] *University of Nebraska Lincoln (US)*

[2] *Princeton University (US)*

[3] *University of Nebraska-Lincoln (US)*

Corresponding Authors: oksana.shadura@cern.ch, brian.bockelman@cern.ch, vasil.georgiev.vasilev@cern.ch

ROOT is a large code base with a complex set of build-time dependencies; there is a significant difference in compilation time between the "core" of ROOT and the full-fledged deployment. We present results on a "delayed build" for internal ROOT packages and external packages. This gives the ability to offer a "lightweight" core of ROOT, later extended by building additional modules to extend the functionality of ROOT. As a part of this work, we have improved the separation of ROOT code into distinct modules and packages with minimal dependencies. This approach gives users better flexibility and the possibility to combine various build features without rebuilding from scratch.
Dependency hell is a common problem found in software and particularly in HEP software ecosystem. We would like to discuss an improvement of artifact management ("lazy-install") system as a solution to the "dependency hell" problem. HEP software stack usually consists of multiple subprojects with dependencies. The development model is often distributed, independent and noncoherent among the sub-projects. We believe that software should be designed to take advantage of other software components that are already available, or have already been designed and implemented for use elsewhere rather than "reinventing the wheel". In our contribution, we will present our approach to artifact management system of ROOT together with a set of examples and use cases.

Track 1: Computing Technology for Physics Research / 386

Migrating large codebases to C++ Modules

Authors: Oksana Shadura[1]; Vasil Georgiev Vasilev[2]; Yuka Takahashi[3]

[1] *National Academy of Sciences of Ukraine (UA)*

[2] *Princeton University (US)*

[3] *University of Tokyo*

Corresponding Authors: yuka.takahashi@cern.ch, vasil.georgiev.vasilev@cern.ch, oksana.shadura@cern.ch

ROOT has several features which interact with libraries and require implicit header inclusion. This can be triggered by reading or writing data on disk, or user actions at the prompt. Often, the headers are immutable and reparsing is redundant. C++ Modules are designed to minimize the reparsing of the same header content by providing an efficient on-disk representation of C++ Code. ROOT has released a C++ Modules-aware technology preview which intends to become the default for the next release.

In this contribution, we would like to summarize our ROOT experience migrating to C++ modules the codebase of ROOT. We outline the challenges for migration of the CMS software stack to use C++ modules, including integration of modules support in the build system while providing better functionality and correctness. We also give an insight of the continuous process of the improving

performance bottlenecks for C++ modules and also evaluate the performance benefits that experiments are expected to achieve.

Track 2: Data Analysis - Algorithms and Tools / 387

A 3D Track Finder for the Belle II CDC L1 Trigger

Authors: Sebastian Skambraks[1]; Steffen Baehr[2]; Christian Kiesling[3]; Sara McCarney[1]; Felix Meggendorfer[1]; Raynette Van Tonder[2]

[1] *Max-Planck-Institut für Physik*

[2] *Karlsruhe Institute of Technology*

[3] *Werner-Heisenberg-Institut*

Corresponding Authors: raynette.tonder@kit.edu, cmk@mppmu.mpg.de, steffen.baehr@kit.edu, fmegg@mpp.mpg.de, sskambra@mpp.mpg.de, mccarney@mpp.mpg.de

Machine learning methods are integrated into the pipelined first level track trigger of the upgraded flavor physics experiment Belle II in Tsukuba, Japan. The novel triggering techniques cope with the severe background conditions coming along with the upgrade of the instantaneous luminosity by a factor of 40 to $\mathcal{L} = 8 \times 10^{35}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Using the precise drift-time information of the central drift chamber, a neural network L1 trigger estimates the 3D track parameters of found single tracks. An extension of the present 2D Hough track finder to a 3D finder is proposed, where the single hit representations in the Hough plane are trained using Monte Carlo. This 3D finder enables an improvement of the track finding efficiency by including the stereo sense wires as input. The estimated polar track angle allows a specialization of the following neural networks to phase space sectors.

Track 2: Data Analysis - Algorithms and Tools / 388

Air shower reconstruction with hexagonal convolutional neural networks

Authors: Constantin Steppa[1]; Tim Lukas Holch[2]

Co-author: Kathrin Egberts [1]

[1] *University of Potsdam*

[2] *Humboldt-University of Berlin*

Corresponding Authors: egberts@uni-potsdam.de, steppa@uni-potsdam.de, holchtim@physik.hu-berlin.de

Ground-based γ-ray astronomy relies on reconstructing primary particles' properties from the measurement of the induced air showers. Currently, template fitting is the state-of-the-art method to reconstruct air showers. CNNs represent promising means to improve on this method in both, accuracy and computational cost. Promoted by the availability of inexpensive hardware and open-source deep learning frameworks (DLFs) the applicability of CNNs for air shower reconstruction is in focus of recent and on-going studies. Thereby, the hexagonal sampling of data, which is common for Cherenkov telescopes but does not fit the input format of DLFs, poses an obstacle. It has been addressed e.g by transforming the hexagonally sampled data to an approximate representation on a rectangular grid prior to the application of CNNs. Though this procedure was shown to yield promising results, it comes at the prize of increasing computational costs. The transformation can be omitted if convolutions are directly applied on the hexagonal grid. For this purpose a Python library, called HexagDLy, was written and made publicly available. In the present study, HexagDLy was used to build CNN models for the analysis of data from the High Energy Stereoscopic System.

The performance of these models on classifying and reconstructing air-shower events will be shown and compared to alternative methods.

Track 2: Data Analysis - Algorithms and Tools / 389

Machine Learning Study of Open Supernova Catalog

Authors: Konstantin Malanchev[None]; Maria Pruzhinskaya[None]; Matwey Kornilov[None]; Emille Ishida[None]; Florian Mondon[None]; Alina Volnova[None]; Vladimir Korolev[None]

Corresponding Author: malanchev@physics.msu.ru

The next generation of astronomical surveys will revolutionize our understanding of the Universe, raising unprecedented data challenges in the process. One of them is the impossibility to rely on human scanning for the identification of unusual/unpredicted astrophysical objects. Moreover, given that most of the available data will be in the form of photometric observations, such characterization cannot rely on the existence of high resolution spectroscopic observations.

The goal of this project is to detect the anomalies in the Open Supernova Catalog with use of machine learning. We develop a pipeline where human expertise and modern machine learning techniques can complement each other. Using supernovae as a case study, our proposal is divided in two parts: a first developing a strategy and pipeline where anomalous objects are identified, and a second phase where such anomalous objects submitted to careful individual analysis. The strategy requires an initial data set for which spectroscopic is available for training purposes, but can be applied to a much larger data set for which we only have photometric observations. This project represents an effective strategy to guarantee we shall not overlook exciting new science hidden in the data we fought so hard to acquire.

Poster Session / 390

Meta-Learning for Artificial Neural Network Hyper-Parameter Optimization for CERN CMS Offline Data Certification

Authors: Mantas Stankevicius[1]; Valdas Rapsevicius[2]

[1] *Vilnius University, Institute of Data Science and Digital Technologies (LT)*

[2] *Fermi National Accelerator Lab. (US)*

Corresponding Authors: valdas.rapsevicius@cern.ch, mantas.stankevicius@cern.ch

The Compact Muon Solenoid (CMS) is one of the general-purpose detectors at the CERN Large Hadron Collider (LHC) which collects enormous amounts of physics data. Before the final physics analysis can proceed, data has to be checked for quality (certified) by passing a number of automatic (like physics objects reconstruction, histogram preparation) and manual (checking, comparison and decision making) steps. Most important yet involving and error-prone last manual step of decision making is currently under active research as a subject of future automation by applying recent advancements from computer science, specifically, machine learning (ML).

Ultimately, CMS data certification is a binary classification task where various ML techniques are being investigated for applicability. Just like in any other ML task the hyper-parameter tuning is a difficult problem, there is no golden rule and each use case is different. This study explored meta-learning applicability, it is a hyper-parameters finding technique where algorithm learns hyper-parameters from previous training experiments. Evolutionary genetic algorithm has been used to tune hyper-parameters of a neural network, like number of hidden layers, number of neurons per layer, activation functions, dropouts, training batch size and optimizer. Initially, genetic algorithm takes manu-

ally specified set of hyper-parameters and then evolves towards the near-optimal solution. Genetic stochastic operators, crossover and mutation, were applied to avoid local optimal solutions.

Study proves that by carefully seeding the initial solution the optimal is likely to be found. Proposed solution has improved AUC score of neural network used for CERN CMS data certification. Similar algorithm can be applied for other machine learning models for hyper-parameter optimization.

Poster Session / 391

Scaling the training of semantic segmentation on MicroBooNE events to multiple GPUs

Authors: Alexander Hagen[1]; Jan Fridolf Strube[None]; Eric Church[1]; Vinay Amatya[1]; Kolahal Bhattacharya[1]

[1] *PNNL*

Corresponding Authors: jan.strube@pnnl.gov, alexander.hagen@pnnl.gov, eric.church@pnnl.gov

Measurements in Liquid Argon TPC (LArTPC) neutrino detectors, such as the MicroBooNE detector at Fermilab, feature large, high fidelity event images. Deep learning techniques have been extremely successful in classification tasks of photographs, but their application to LArTPC event images is challenging, due to the large size of the events. Events in these detectors are typically two orders of magnitude larger than images found in classical challenges, like recognition of handwritten digits contained in the MNIST database or object recognition in the ImageNet database. Ideally, training would occur on many instances of the entire event data, instead of many instances of cropped regions of interest from the event data. However, such efforts lead to extremely long training cycles, which slow down the exploration of new network architectures and hyperparameter scans to improve the classification performance.
We present studies of scaling a LArTPC classification problem on multiple architectures, spanning multiple nodes. The studies are carried out on simulated events in the MicroBooNE detector. Institutional computing at Pacific Northwest National Laboratory and the Summit-dev machine at Oak Ridge National Laboratory's Leadership Computing Facility have been used. Additionally, we discuss difficulties encountered (and solutions thereof) while scaling from single to multiple GPU training for this context.

Track 2: Data Analysis - Algorithms and Tools / 392

Electromagnetic-shower generation with Graphical GANs

Authors: Vladislav Belavin[1]; Andrey Ustyuzhanin[2]

[1] *CERN*

[2] *Yandex School of Data Analysis (RU)*

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

At this moment the most convenient approach in electromagnetic shower generation is Monte-Carlo simulation produced by software packages like GEANT4. However, one of the critical problems of Monte-Carlo production is that it is extremely slow since it involves simulation of numerous subatomic interactions.

Recently, generative adversarial networks(GANs) addressed speed issue in the simulation of calorimeters response with significant speeding-up a two-three order of magnitude in comparison with the current approach. However, it is challenging to define network architecture that converges within a reasonable timeframe and define a proper figure of merit that yields realistic synthetic objects.

In this work, we propose a metric that deals successfully with the structure of the showers. The architecture of the neural network that performs nicely with shower-like objects is called graphical network. Plus the approach for the generation of electromagnetic showers with graphical neural networks fits well into a GAN-based training and produces the meaningful result. The novelty of this approach lies, firstly, in the generation of complex recursive physical process with neural network and, secondly, in significant speed-up in comparison with traditional simulation approaches.

Poster Session / 394

Vectorization techniques for probability distribution function using VecCore

Author: Oscar Roberto Chaparro Amaro[1]

Co-authors: Castro Martínez [2]; Soon Yung Jun [3]

[1] *Instituto Politécnico Nacional. Centro de Investigación en Compuutación.*

[2] *Instituto Politécnico Nacional. Centro de Investigación en Comptación.*

[3] *Fermi National Accelerator Lab. (US)*

Corresponding Authors: saga9211@hotmail.com, jesus.martinez.castro@gmail.com, syjun@fnal.gov

Probability distribution functions (PDFs) are very used in modeling random processes and physics simulations. It can be demonstrated that the relationships between PDFs are linked through functional parameters. Improving the performance of the generation of many random numbers to be used as input by the PDFs is often a very challenging task as it involves algorithms with acceptance-rejection methods. In this work we present general strategies on how to vectorize some PDFs using VecCore library. We show the results for the Exponential, Poisson and Gamma probability distributions.

Track 2: Data Analysis - Algorithms and Tools / 396

Physics inspired feature engineering with Lorentz Boost Networks

Authors: Martin Erdmann[1]; Erik Geiser[2]; Yannik Alexander Rath[3]; Marcel Rieger[1]

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

[2] *III. Physics Institute A - RWTH Aachen University*

[3] *RWTH Aachen University (DE)*

Corresponding Authors: martin.erdmann@cern.ch, yannik.alexander.rath@cern.ch, marcel.rieger@cern.ch, erik.geiser@rwth-aachen.de

A large part of the success of deep learning in computer science can be attributed to the introduction of dedicated architectures exploiting the underlying structure of a given task. As deep learning methods are adopted for high energy physics, increasing attention is thus directed towards the development of new models incorporating physical knowledge.

In this talk, we present a network architecture that utilizes our knowledge of particle combinations and directly integrates Lorentz boosting to learn relevant physical features from basic four vectors. We explore two example applications, namely the discrimination of hadronic top-quark decays from light quark and gluon jets, and the separation of top-quark pair associated Higgs boson events from a $t\bar{t}$ background. We also investigate the learned combinations and boosts to gain insights into what the network is learning.

Track 2: Data Analysis - Algorithms and Tools / 397

Primary Vertex Reconstruction with Deep Learning in FPGAs for the Phase-2 Upgrade of the Level-1 Trigger of CMS

Author: Antoni Shtipliyski[1]

[1] *Imperial College (GB)*

Corresponding Author: a.shtipliyski@cern.ch

The High-Luminosity upgrade of LHC (HL-LHC) is expected to deliver a total luminosity of 3000 fb^{-1} to the general purpose experiments. This will allow the measurement of Standard Model processes with unprecedented precision, and will significantly increase the reach of searches for new physics. Higher data rates and increased radiation levels will require substantial upgrades to the detectors and their trigger and data acquisition systems. The Phase-2 upgrade of CMS comprises a complete replacement of the silicon tracker that will for the first time provide tracking information to the Level-1 (L1) hardware trigger. The upgraded trigger is designed to reduce the event rate to 750 kHz from a collision frequency of 40 MHz by processing 50 Tbps of incoming bandwidth within 12.5 μs latency. The increased luminosity is expected to produce around 200 additional "pileup" interactions per bunch crossing. This creates a challenging environment for efficient triggering and the use of tracks in the L1 algorithms would be essential for pileup mitigation, since charged particle tracks can be matched to energy deposits in the calorimeters and tracks in the muon detectors to identify directly the particles originating from the primary interaction vertex of the hard scatter process. This talk will introduce the challenges involved in reconstructing the primary vertex at L1 and describe a Deep Learning algorithm developed for inference in an FPGA data processor. The stringent time requirements of the L1 trigger mandate that vertex reconstruction is done within O(100ns) while the computing technology poses the novel problem of balancing the trade-off between algorithm sophistication and FPGA resource usage/latency. Preliminary plans for operating the algorithm during data-taking will also be outlined.

Poster Session / 398

Monitoring of time evolution of the trigger rates exploiting deep representation learning at the CMS experiment

Authors: Adrian Alan Pol[1]; Maurizio Pierini[2]; Gianluca Cerminara[2]; Cecile Germain[3]

[1] *Université Paris-Saclay (FR)*

[2] *CERN*

[3] *Universite Paris Sud*

Corresponding Authors: cecile.germain@lri.fr, gianluca.cerminara@cern.ch, adrian.pol@cern.ch, maurizio.pierini@cern.ch

Real time monitoring of Compact Muon Solenoid (CMS) trigger system is a vital task to ensure the quality of all physics results published by the collaboration. Today, the trigger monitoring software reports on potential problems given the time evolution of the reported rates. The anomalous rates are identified given the deviation from the prediction which is calculated using a regression model generated independently for each trigger path. The CMS experiment implements a two-level trigger system. The fast hardware-based Level 1 triggers filter the data before the software-based High Level Triggers which access the full detector information. In most cases real problems present in the detector manifest themselves in abnormal trigger rates for a number of trigger paths which share a common infrastructure; whereas a few unrelated triggers misbehaving can be a result of statistical fluctuation. As such, the alarms require a considerable human interpretation. This contribution presents steps undertaken towards extending the current framework taking into account interdependence of different trigger paths. Our prototype, based on a deep autoencoder, exploits a global configuration of the trigger system and in particular its hierarchical nature.

Poster Session / 399

Prong Formation in NOvA Reconstruction Pipeline

Authors: Petr Bouř[1]; Václav Kůs[1]

[1] *FNSPE CTU in Prague*

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

Event reconstruction for NOvA experiment is a critical step preceding further data analysis. We describe the complex NOvA reconstruction pipeline (containing several unsupervised learning techniques) with focus on the specific step of so-called "prong matching". In this step, we are combining 2D prongs (projections of particle trajectories) into 3D prong objects. In order to find the best matching 2D trajectory projections, we are using Kolmogorov-Smirnov homogeneity test on energy profiles. Moreover, we present the efficiency of the current matching process and propose some improvements on homogeneity testing.

Track 1: Computing Technology for Physics Research / 400

The computing model of the LHCb Upgrade

Authors: Stefan Roiser[1]; Concezio Bozzi[2]; Marco Cattaneo[1]

[1] *CERN*

[2] *CERN and INFN Ferrara*

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

The LHCb Upgrade experiment will start operations in LHC Run 3 from 2021 onwards. Owing to the five-times higher instantaneous luminosity and higher foreseen trigger efficiency, the LHCb Upgrade will collect signal yields per unit time approximately ten times higher than that of the current experiment, with pileup increasing by a factor of six. This contribution presents the changes in the computing model and the associated offline computing resources needed for the LHCb Upgrade, that are defined by the significantly increased trigger output rate compared to the current situation, and the corresponding necessity to generate significantly larger samples of simulated events. The update of the LHCb computing model for Run 3, and beyond, is discussed with an emphasis on the optimization that has been applied to the usage of distributed computing CPU and storage resources.

Track 1: Computing Technology for Physics Research / 401

The core software framework for the LHCb Upgrade

Authors: Sebastien Ponce[1]; Concezio Bozzi[2]; Niklas Nolte[3]

[1] *CERN*

[2] *CERN and INFN Ferrara*

[3] *CERN / Technische Universitaet Dortmund (DE)*

Corresponding Authors: niklas.nolte@cern.ch, sebastien.ponce@cern.ch, bozzi@fe.infn.it

The LHCb experiment will be upgraded for data taking in Run 3 and beyond and the instantaneous luminosity will in particular increase by a factor five. The lowest level trigger of the current experiment, a hardware-based trigger that has a hard limit of 1 MHz in its event output rate, will be

removed. and replaced with a full software trigger. This new
trigger needs to sustain rates up 30 MHz for the inelastic proton-proton collisions and will thus process 5 Tb/s of data, which is over two orders of magnitude larger than the rate processed by the current LHCb experiment, all this to be achieved within the same costs of the current data processing and without compromising the physics performance. For this purpose, the Gaudi framework currently used in LHCb has been re-engineered to enable the maximally efficient usage of vector registers and of multi- and many-core architectures. In particular, a new scheduler and a re-design of the data structures were needed in order to make the most efficient usage of memory resources and speed up
access patterns.
This contribution presents these and other critical points that had to be tackled as well as the current status and an outlook of the work program that will address the challenges of the software trigger in the LHCb Upgrade.

Track 1: Computing Technology for Physics Research / 402

Nested data structures in array and SIMD frameworks

Author: Jim Pivarski[1]

Co-author: Peter Elmer [2]

[1] *Princeton University*

[2] *Princeton University (US)*

Corresponding Authors: peter.elmer@cern.ch, pivarski@princeton.edu

Nested data structures are critical for particle physics: it would be impossible to represent collision data as events containing arbitrarily many particles in a rectangular table (without padding or truncation, or without relational indirection). These data structures are usually constructed as class objects and arbitrary length sequences, such as vectors in C++ and lists in Python, and data analysis logic is expressed in imperative loops and conditionals. However, code expressed this way can thwart auto-vectorization in C++ and Numpy optimization in Python, and may be too explicit to automatically parallelize. We present an extension of the "array programming" model of APL, R, MATLAB, and Numpy, which expresses regular operations on large arrays in a concise syntax. Ordinarily, array programming only applies to flat arrays and rectangular tables, but we show that it can be extended to collections of arbitrary length lists ("jagged arrays"), nested records, polymorphic unions, and pointers. We have implemented such a library in Python called awkward-array, and we will show how it can be used to fit particle physics data into systems designed for Numpy data, such as Pandas (for analysis organization), Numba (for just-in-time compilation), Dask (for parallel processing), and CuPy (array programming on the GPU). We will also show how a proper set of primitives enables non-trivial analyses, such as combinatorial searches for particle candidates, in SIMD environments.

Poster Session / 403

Method for results comparison of various versions of simulation program by the use of experimental data

Authors: Serguei Bityukov[1]; Vladimir Ivantchenko[2]

[1] *Institute for High Energy Physics (RU)*

[2] *CERN*

Corresponding Authors: vladimir.ivantchenko@cern.ch, serguei.bitioukov@cern.ch

In this work are presented the result of the comparing two versions of GEANT4 by the use of experimental data of experiment HARP. The comparison is performed with help of a new method of statistical comparison of data sets. The method provides more information for data analysis than methods based on the chi-squared distribution.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 404

DeepXS: Fast approximation of MSSM electroweak cross sections at NLO

Authors: Sydney Otten[1]; Krzysztof Rolbiecki[2]; Sascha Caron[3]; Jong Soo Kim[None]; Roberto Ruiz De Austri[None]; Jamie Tattersall[4]

[1] *Radboud Universiteit Nijmegen*

[2] *University of Warsaw*

[3] *Nikhef National institute for subatomic physics (NL)*

[4] *RWTH Aachen*

Corresponding Authors: sascha.caron@cern.ch, krolb@fuw.edu.pl, sydney.otten@rwth-aachen.de, tattersall@physik.rwth-aachen.de, jongsoo.kim@tu-dortmund.de, rruiz@ific.uv.es

The recent years have shown an exciting development in the scientific commmunity due to the interplay between new methods from data science and artificial intelligence, increasing computational resources and physics. The fundamental object of our theories of nature is the Lagrangian whose form is determined by the symmetries found already. A famous and well-motivated extension of the SM Lagrangian is given by an additional space-time symmetry, supersymmetry. However, this extension is not only one additional theory but instead is a manifold of infinitely many theories in a parameter space with 19 effective dimensions. The quest to judge whether our models for nature are still possibly true descriptions requires the careful statistical analysis of the sea of data that is provided by experiments, e.g. at the Large Hadron Collider, in the face of the standard model of particle physics. This inspection demands a fast and accurate evaluation of cross sections at least at the next-to-leading order. However, the currently available codes take several minutes to evaluate the cross section of one parameter point. With the help of deep neural networks, expert knowledge, stacking and active learning we create a tool, DeepXS, that is seven orders of magnitude faster and only needs microseconds to calculate the cross section of supersymmetric electroweak pairs produced at the LHC with errors that are lower than the scale and PDF uncertainty. In this talk we will present how we created the AIs in DeepXS, demonstrate its performance and discuss subtleties of its validity.

Track 2: Data Analysis - Algorithms and Tools / 405

PhenoAI and iDarkSurvey: Learning (from) high-dimensional models

Authors: Jisk Attema[1]; Sascha Caron[2]; Faruk Diblen[1]; Tom Heskes[3]; Sydney Otten[4]; Krzysztof Rolbiecki[5]; Roberto Ruiz De Austri[6]; Jong Soo Kim[None]; Bob Stienen[3]

[1] *Netherlands eScience center*

[2] *Nikhef National institute for subatomic physics (NL)*

[3] *Radboud University Nijmegen*

[4] *Radboud Universiteit Nijmegen*

[5] *University of Warsaw*

[6] *Instituto de Fisica Corpuscular (ES)*

Corresponding Authors: jongsoo.kim@tu-dortmund.de, sascha.caron@cern.ch, b.stienen@science.ru.nl, j.attema@esciencecenter.nl, sydney.otten@rwth-aachen.de, t.heskes@science.ru.nl, f.diblen@esciencecenter.nl, krolb@fuw.edu.pl, roberto.ruiz.de.austri@cern.ch

Although the standard model of particle physics is successful in describing physics as we know it, it is known to be incomplete. Many models have been developed to extend the standard model, none of which have been experimentally verified. One of the main hurdles in this effort is the dimensionality of these models, yielding problems in analysing, visualising and communicating results. Because of this, most current day analyses are done using simplified models, but in this process descriptive power is lost. However, by using machine learning on simulated model points, we show that we can overcome these problems and predict both binary exclusion an continuous likelihood in any parameter space. This functionality is implemented in the PhenoAI framework, allowing non-expert users of machine learning to use trained machine learning models in their own analyses. The simulated data can be stored in our new webbased database and model visualisation tool iDarkSurvey.

Poster Session / 406

Deep learning for Directional Dark Matter search

Authors: Artem Golovatiuk[1]; Giovanni De Lellis[2]; Andrey Ustyuzhanin[3]

[1] *University of Naples (IT)*

[2] *Universita e sezione INFN di Napoli (IT)*

[3] *Yandex School of Data Analysis (RU)*

Corresponding Authors: giovanni.de.lellis@cern.ch, andrey.ustyuzhanin@cern.ch, artem.golovatiuk@cern.ch

The NEWSdm (Nuclear Emulsions for WIMP Search directional measure) is an underground Direct detection Dark Matter (DM) search experiment. The usage of recent developments in the nuclear emulsions allows probing new regions in the WIMP parameter space. The prominent feature of this experiment is a potential of recording the signal direction, which gives a chance of overcoming the "neutrino floor".

State of the art techniques lower the background contamination significantly, however, background rejection remains crucial for DM sensitivity. Deep Neural Networks were used for separation between potential DM signal and various classes of background.

In this work, we present the usage of deep 3D Convolutional Neural Networks in order to take into account the physical peculiarities of the data and achieve strong background rejection.

Track 2: Data Analysis - Algorithms and Tools / 407

Uncertainty reduction by gradient descent

Authors: Victor Estrade[1]; David Rousseau[2]; Cecile Germain[3]; Isabelle Guyon[None]

[1] *LRI*

[2] *LAL-Orsay, FR*

[3] *Universite Paris Sud*

Corresponding Authors: estrade@lri.fr, cecile.germain@lri.fr, guyon@clopinet.com, rousseau@lal.in2p3.fr

Data analysis based on forward simulation often require the use of a machine learning model for statistical inference of the parameters of interest.
Most of the time these learned model are trained to discriminate events between backgrounds and signals to produce a 1D score, which is used to select a relatively pure signal region.
The training of the model does not take into account the final objective that is to estimate the values

of the parameters of interest.
Those measurements also depends on other parameters, denoted as nuisance parameters, that will induce systematic errors on estimated values.
We propose to explore learning methods that directly minimize the measurement error (both statistical and systematic) on a realistic case coming from HEP.

Poster Session / 408

Generalization of Homogeneity Tests Used in High Energy Physics Experiments

Author: Jakub Trusina[1]

Co-authors: Jiri Franc [2]; Adam Novotny [1]

[1] *Czech Technical University (CZ)*

[2] *Czech Technical University in Prague*

Corresponding Authors: novotad2@fjfi.cvut.cz, jakub.trusina@cern.ch, jiri.franc@fjfi.cvut.cz

In High Energy Physics, tests of homogeneity are used primarily in two cases: for verification that data sample does not differ significantly from numerically produced Monte Carlo sample and for verifying separation of signal from background. Since Monte Carlo samples are usually weighted, it is necessary to modify classical homogeneity tests in order to apply them to weighted samples. In ROOT, the only homogeneity tests that allow testing weighted samples are implemented for binned data. However, after the data are binned the full information is lost. Therefore we compare these ordinary tests with modified versions of the Kolmogorov-Smirnov, Anderson-Darling and Cramér-von Mises tests that use full sample information. The proposed tests are compared by estimating a probability of type-I error which is crucial for a test's reliability.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 409

Further developments of FORM

Author: Takahiro Ueda[1]

Co-authors: Toshiaki Kaneko ; Ben Ruijl [2]; Jos Vermaseren

[1] *KEK*

[2] *Nikhef*

Corresponding Authors: tueda@st.seikei.ac.jp, t68@nikhef.nl, toshiaki.kaneko@kek.jp, benruyl@gmail.com

FORM is a symbolic manipulation system, which is especially advantageous for handling gigantic expressions with many small terms. Because FORM has been developed in tackling real problems in perturbative quantum field theory, it has some features useful in such problems, although FORM applications are not restricted to any specific research field. In this talk, we discuss recent developments of FORM and its new features.

Poster Session / 410

CDF Long Term Data Preservation at the INFN CNAF Tier-1 data center: the Oracle database and software framework replica.

Author: Pier Paolo Ricci[1]

[1] *INFN CNAF*

Corresponding Author: pierpaolo.ricci@cnaf.infn.it

The INFN CNAF Tier-1 Long Term Data Preservation (LTDP) project was established at the end of 2012 in close collaboration with Fermi National Accelerator Laboratory (FNAL) with the purpose of saving, distributing and maintaining over time the CDF Tevatron analysis framework and all the relevant scientific data produced by the experiment activity. During recent years, a complete copy of all CDF Run II raw data and user level ntuples (about 4 PB) was successfully copied from FNAL to the INFN CNAF Tier-1 tape storage. The data copy was also recovered from tape after the 2017 flood event when the storage array, containing the CDF filesystem disk area, was affected. Furthermore, there are indications that in the near future the FNAL Tevatron computing services currently running only and exclusively at FNAL (in particular the Oracle database, software methods to get calibration constants, software release availability service and specific job submission system) should be completely replicated in a second site. This operation could help ensuring the CDF physicists the possibility to complete the analyses in progress and could guarantee the future activity of the collaboration. For this reason, we decided to provide a complete copy of the Oracle database software used by CDF for the collection of not-statistical data, such as detector calibrations. This Oracle service is composed by two main DB instances, the offline one (containing info for the offline data processing and data book-keeping) and the online (data taking condition), consisting of roughly 1.9 x 1010 rows of information. Both the instances are essential, therefore we need to run the updated and replicated Oracle services at our site in order to provide information for the MC simulation, ongoing analyses and reprocessing activities. About the CDF software preservation, the releases are currently running on CERN Virtual Machine File System (CVMFS) with a server currently located at FNAL. Since the CVMFS repositories could be distributed over remote sites, we are going to install a CVMFS replica server at our site in order to make provision for the availability of the software to the community. Regarding the job submission system, the collaboration currently use the FNAL Jobsub tool and some specific CDF wrappers for emulating the specific experiment commands that were commonly exploited by CDF uses. A Jobsub head node server is being installed and tested at our site as a future job submission system that could be completely independent from FNAL services. All the software replicas running at our site, together with the Run II data already copied here, could give to the CDF collaboration community the future important possibility to continue activity in parallel with FNAL or independently at our site.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 411

Forgotten (or not) possibility of the application of the least squares method: the analysis of the computer-based numerical results to QCD $\overline{\mathrm{MS}}$-on-shell mass relation at the four-loop level

Authors: Andrei Kataev [1]; Victor Molokoedov[2]

[1] *Institute for Nuclear Research of RAS, Moscow*

[2] *Moscow Institute of Physics and Technology , Dolgoprudny*

Corresponding Authors: viktor_molokoedov@mail.ru, kataev@ms2.inr.ac.ru

The set of the four-loop numerical results for the relation between pole and running heavy quarks masses in QCD at fixed number of lighter flavors $3 \leq n_l \leq 15$, which was obtained in Ref.1 with help of the Lomonosov Supercomputer of Moscow State University, is analysed by the ordinary method of the least squares. We use its variant which allows to solve the overdetermined system of 13 linear equations and to define 4 coefficients of the polynomial of the third power in n_l in the expression for the four-loop correction to the QCD $\overline{\mathrm{MS}}$-on-shell mass relation with corresponding uncertainties of this approach. The central values of these terms are consistent with a high degree of accuracy with the results obtained in Ref.1. To demonstrate the stability of the least squares method to the number of equations we also consider the situation when the number of equations is equal to 3 at $3 \leq n_l \leq 5$. It is interesting that in both cases the central values of all unknown terms coincide with the previously obtained two unknown in analytical form values at $3 \leq n_l \leq 15$ [2,3] while the uncertainties increase no more than 10 times. Thus the least squares method allows to check for self-consistency the results of numerical computations 1 and analytical calculations [4], obtained using different methods of evaluation of Feynman diagrams and of advanced computer-oriented programs.

References.

1 P.~Marquard, A.~V.~Smirnov, V.~A.~Smirnov, M.~Steinhauser and D.~Wellmann,
"$\overline{\mathrm{MS}}$-on-shell quark mass relation up to four loops in QCD and a general SU(N) gauge group,"
Phys.\ Rev.\ D {\bf 94}, no. 7, 074025 (2016)[arXiv:1606.06754 [hep-ph]].

[2] A.~L.~Kataev and V.~S.~Molokoedov,"On the flavour dependence of the $\mathcal{O}(\alpha_s^4)$ correction to the relation between running and pole heavy quark masses," Eur.\ Phys.\ J.\ Plus {\bf 131}, no. 8, 271 (2016)
[arXiv:1511.06898 [hep-ph]].

[3] A.~L.~Kataev and V.~S.~Molokoedov,
"Multiloop contributions to the $\overline{\mathrm{MS}}$-on-shell mass relation for heavy quarks in QCD and charged leptons in QED and the asymptotic structure of the perturbative QCD series," arXiv:1807.05406 [hep-ph].

[4] R.~Lee, P.~Marquard, A.~V.~Smirnov, V.~A.~Smirnov and M.~Steinhauser,
"Four-loop corrections with two closed fermion loops to fermion self energies and the lepton anomalous magnetic moment," JHEP {\bf 1303}, 162 (2013)
[arXiv:1301.6481 [hep-ph]]

Track 2: Data Analysis - Algorithms and Tools / 412

Variational Autoencoders for New Physics Mining at the Large Hadron Collider

Authors: Jean-Roch Vlimant[1]; Maurizio Pierini[2]; Olmo Cerri[1]; Thong Nguyen[1]; Maria Spiropulu[3]

[1] *California Institute of Technology (US)*

[2] *CERN*

[3] *California Institute of Technology*

Corresponding Authors: maurizio.pierini@cern.ch, thong@caltech.edu, smaria@caltech.edu, olmo.cerri@cern.ch, jean-roch.vlimant@cern.ch

Using variational autoencoders trained on known physics processes, we develop a one-side p-value test to isolate previously unseen event topologies as outlier events. Since the autoencoder training does not depend on any specific new physics signature, the proposed procedure has a weak dependence on underlying assumptions about the nature of new physics. An event selection based on this algorithm would be complementary to classic LHC searches, typically based on model-dependent hypothesis testing. Such an algorithm would deliver a list of anomalous events, that the experimental collaborations could further scrutinize and even release as a catalog, similarly to what is typically done in other scientific domains. Repeated patterns in this dataset could motivate new scenarios for beyond-the-standard-model physics and inspire new searches, to be performed on future data with traditional supervised approaches. Running in the trigger system of the LHC experiments, such an application could identify anomalous events that would be otherwise lost, extending the scientific reach of the LHC.

Track 2: Data Analysis - Algorithms and Tools / 414

ConformalTracking: a geometry agnostic tracking library

Authors: Andre Sailer[1]; Daniel Hynds[2]; Emilia Leogrande[1]; Erica Brondolin[3]; Marko Petric[1]

[1] *CERN*

[2] *Nikhef National institute for subatomic physics (NL)*

[3] *Austrian Academy of Sciences (AT)*

Corresponding Authors: emilia.leogrande@cern.ch, erica.brondolin@cern.ch, daniel.hynds@cern.ch, andre.philippe.sailer@cern.ch, marko.petric@cern.ch

ConformalTracking is an open source library created in 2015 to serve as a detector independent solution for track reconstruction in detector development studies at CERN. Pattern recognition is one of the most CPU intensive tasks of event reconstruction at present and future experiments. Current tracking programs of the LHC experiments are mostly tightly linked to individual detector descriptions or event processing frameworks. ConformalTracking does a pattern recognition in a conformal-mapped plane, where helix trajectories of charged particles in a magnetic field are projected into straight lines, followed by a Kalman-Filter-based fit in global space. At the core of the library lies a nearest neighbour search that is optimized by means of fast KDTrees and enhanced with a cellular automaton to reconstruct the linear paths. Being based exclusively on the spatial coordinates of the hits, this algorithm is adaptable to different detector designs and beam conditions. In the detectors at CLIC and FCCee, it also profits from the low-mass silicon tracking system, which reduces complications from multiple scattering and interactions. Full-simulation studies have been performed in order to validate the algorithm and assess its performances, also in the presence of beam-induced background. In this talk, recent developments and features of the track reconstruction chain as well as results for isolated tracks and complex events with background will be discussed.

Poster Session / 415

Direct optimisation of the discovery significance when training neural networks to search for new physics in particle colliders

Authors: Adam Christopher Elwood[1]; Dirk Krücker[1]

[1] *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Authors: dirk.krucker@cern.ch, adam.elwood@desy.de

We introduce two new loss functions designed to directly optimise the statistical significance of the expected number of signal events when training neural networks to classify events as signal or background in the scenario of a search for new physics at a particle collider. The loss functions are designed to directly maximise commonly used estimates of the statistical significance, s/√(s+b), and the Asimov estimate, Z_A. We consider their use in a toy SUSY search with 30 fb^(−1) of 14 TeV data collected at the LHC. In the case that the search for the SUSY model is dominated by systematic uncertainties, it is found that the loss function based on Z_A can outperform the binary cross entropy in defining an optimal search region.

Track 2: Data Analysis - Algorithms and Tools / 416

Deep Learning on HPC at NERSC

Authors: Wahid Bhimji[1]; Steven Andrew Farrell[1]; Mustafa Mustafa[2]

[1] *Lawrence Berkeley National Lab. (US)*

[2] *Lawrence Berkeley National Laboratory*

Corresponding Authors: wahid.bhimji@cern.ch, steven.farrell@cern.ch, mmustafa@lbl.gov

We present recent work in deep learning for particle physics and cosmology at NERSC, the US Dept. of Energy mission HPC centre. We will describe activity in new methods and applications; distributed training across HPC resources; and plans for accelerated hardware for deep learning in NERSC-9 (Perlmutter) and beyond.
Some of the HEP methods and applications showcased include conditional Generative Adversarial Networks on large full-detector HEP images and high-resolution dark-matter cosmology simulations; bayesian inference via probabilistic programming for LHC analyses; alternative representations of HEP data for NN training (such as GraphNNs); and architecture search approaches. We also describe computational developments and infrastructure for training these models at large scale on NERSC supercomputers through productive interfaces.

Track 1: Computing Technology for Physics Research / **417**

Vectorization of random number generation and reproducibility of concurrent particle transport simulation

Authors: Soon Yung Jun[1]; John Apostolakis[2]; Andrei Gheata[2]; Philippe Canal[1]; Lorenzo Moneta[2]

[1] *Fermi National Accelerator Lab. (US)*

[2] *CERN*

Corresponding Authors: lorenzo.moneta@cern.ch, pcanal@fnal.gov, john.apostolakis@cern.ch, andrei.gheata@cern.ch, syjun@fnal.gov

Efficient random number generation with high quality statistical properties and exact reproducibility of Monte Carlo simulation are important requirements in many areas of computational science. VecRNG is a package providing pseudo-random number generation (pRNG) in the context of a new library VecMath. This library bundles up several general-purpose mathematical utilities, data structures and algorithms having both SIMD and SIMT(GPUs) support based on VecCore. Several state-of-the-art RNG algorithms are implemented as kernels supporting parallel generation of random numbers in both scalar, vector and Cuda workflows. In this report, we will present design considerations, implementation details and computing performance of parallel pRNG engines on both CPU and GPU. Reproducibility of propagating multiple particles in parallel in HEP event simulation is demonstrated, using GeantV based examples, for both sequential and fine-grain track-level concurrent simulation workflows. Strategies for efficient uses of vectorized pRNG and non-overlapping streams of random number sequences in concurrent computing environments will be also discussed.

Track 3: Computations in Theoretical Physics: Techniques and Methods / **418**

micrOMEGAs5.0

Author: Alexander Pukhov[1]

[1] *Skobeltsyn institute of Nuclear Physics*

Corresponding Author: pukhov@lapp.in2p3.fr

micrOMEGAs is a package for the calculation of the relic density of Dark Matter and of different observables related with Dark Matter searches. The talk will present the general structure of the package and several recent developments including freeze-in relic abundance calculation, interface with different packages that compute collider observables, and recent improvements in direct detection signals.

Poster Session / 419

Geant4 Parameter Tuning Using Professor

Authors: Victor Daniel Elvira[1]; Laura Fields[2]; Krzysztof Genser[3]; Robert William Hatcher[4]; Vladimir Ivantchenko[5]; Soon Yung Jun[1]; Michael Kelsey[6]; Tatsumi Koi[6]; Gabriel Perdue[4]; Alberto Ribon[5]; Vladimir Uzhinsky[7]; Dennis Herbert Wright[None]; Julia Yarba[1]

[1] *Fermi National Accelerator Lab. (US)*

[2] *Fermi National Accelerator Laboratory*

[3] *Fermi National Accelerator Laboratory (US)*

[4] *Fermilab*

[5] *CERN*

[6] *SLAC National Accelerator Laboratory (US)*

[7] *Joint Institute for Nuclear Research (RU)*

Corresponding Authors: tkoi@slac.stanford.edu, uzhinsky@cern.ch, kelsey@slac.stanford.edu, alberto.ribon@cern.ch, perdue@fnal.gov, julia.yarba@cern.ch, dennis.herbert.wright@cern.ch, ljf26@fnal.gov, genser@fnal.gov, syjun@fnal.gov, vladimir.ivantchenko@cern.ch, rhatcher@fnal.gov, daniel@fnal.gov

The Geant4 toolkit is used extensively in high energy physics to simulate the passage of particles through matter and to estimate effects such as detector responses, efficiencies and smearing. Geant4 uses many underlying models to predict particle interaction kinematics, and uncertainty in these models leads to uncertainty in the interpretation of experiment measurements. The Geant4 collaboration recently made some parameters in physics models accessible for uncertainty studies. We present a study of the impact of varying parameters in three Geant4 hadronic physics models on agreement with thin target data sets and describe fits to these data sets using the Professor model tuning framework. We find that varying parameters produces substantially better agreement with some data sets, but that more degrees of freedom are required for full agreement. This work is a first step towards a common framework for propagating uncertainties on Geant4 models to high energy physics measurements, and we outline future work required to complete that goal.

Track 2: Data Analysis - Algorithms and Tools / 420

Reinforced Jet Grooming

Authors: Frederic Alexandre Dreyer[1]; Stefano Carrazza[2]

[1] *Oxford*

[2] *CERN*

Corresponding Authors: stefano.carrazza@cern.ch, frederic.dreyer@physics.ox.ac.uk

We introduce a novel implementation of a reinforcement learning algorithm which is adapted to the problem of jet grooming, a crucial component of jet physics at hadron colliders. We show that the grooming policies trained using a Deep Q-Network model outperform state-of-the-art tools used at the LHC such as Recursive Soft Drop, allowing for improved resolution of the mass of boosted objects. The algorithm learns how to optimally remove soft wide-angle radiation, allowing for a modular jet grooming tool that can be applied in a wide range of contexts.

Track 2: Data Analysis - Algorithms and Tools / 421

Deep Learning based Algorithms in Astroparticle Physics

Authors: Jonas Glombitza[1]; Marcus Wirtz[1]; Martin Erdmann[1]; Michael Dohmen[None]; Vieweg Maximilian[1]

[1] *RWTH Aachen*

In recent years, the astroparticle physics community has successfully adapted supervised learning algorithms for a wide range of tasks, including event reconstruction in cosmic ray observatories1, photon identification at Cherenkov telescopes[2], and the extraction of gravitational wave signals from time traces[3]. In addition, first unsupervised learning approaches of generative models at observatories for cosmic rays showed promising results[4]. Besides simulation acceleration, here, the refinement of physics simulations was investigated by training a refiner network to make simulated time traces to look like data traces. This may have groundbreaking outcomes on machine learning algorithms and shows the potential to explore unsupervised learning for physics research.
In this presentation we summarize the latest developments in machine learning in
the context of astroparticle physics and discuss the far-reaching scope of future applications.

1 DOI: 10.1016/j.astropartphys.2017.10.006
[2] DOI: 10.1016/j.astropartphys.2018.10.003
[3] DOI: 10.1016/j.physletb.2017.12.053
[4] DOI: 10.1007/s41781-018-0008-x

Track 2: Data Analysis - Algorithms and Tools / **422**

Charged Particle Tracking as a QUBO problem solved with Quantum Annealing-Inspired Optimization

Authors: Alexander Zlokapa[1]; Jean-Roch Vlimant[1]; Abhishek Anand[2]; Daniel Lidar[3]; Maria Spiropulu[2]; Joshua Job[3]

[1] *California Institute of Technology (US)*

[2] *California Institute of Technology*

[3] *USC*

Corresponding Authors: jjob@usc.edu, smaria@caltech.edu, alexander.z@cern.ch, jean-roch.vlimant@cern.ch, anand@college.harvard.edu, lidar@usc.edu

With the upgrade of the LHC to high luminosity, an increased rate of collisions will place a higher computational burden on track reconstruction algorithms. Typical algorithms such as the Kalman Filter and Hough-like Transformation scale worse than quadratically. However, the energy function of a traditional method for tracking, the geometric Denby-Peterson (Hopfield) network method, can be described as a quadratic unconstrained binary optimization (QUBO) problem. Quantum annealers have shown shown promise in their ability to solve QUBO problems despite being NP-hard. We present a novel approach for track reconstruction by applying a quantum annealing-inspired algorithm to the Denby-Peterson method. We propose additional techniques to divide an LHC event into disjoint subgraphs in order to allow the problem to be embeddable on existing quantum annealing hardware, using multiple anneals to fit tracks to a single event. To accommodate this dimension reduction, we use Bayesian methods and further algorithms to pre- and post-process the data. Results on the TrackML dataset are presented, demonstrating the successful application of quantum annealing-inspired algorithms to the track reconstruction problem.

Track 1: Computing Technology for Physics Research / **423**

Performance results of the GeantV prototype with complete EM physics

Author: Andrei Gheata[1]

[1] *CERN*

Corresponding Author: andrei.gheata@cern.ch

Improving the computing performance of particle transport simulation is an important goal to address the challenges of HEP experiments in the coming decades (i.e. HL-LHC), as well as the needs of other fields (i.e. medical imaging and radiotherapy).
The GeantV prototype includes a new transport engine, based on track level parallelization by grouping a large number of tracks in flight into "baskets", with improved use of caches and vectorisation. The main goal is to investigate what performance increase this new approach can deliver compared to the performance of the Geant4 toolkit.
We have implemented a prototype of the transport engine and auxiliary components, including a work scheduler, vectorized code for geometry, transport and magnetic field handling, as well as a complete set of vectorized electromagnetic physics models compatible with a recent Geant4 version. Based on this prototype, computing performance benchmarking and software optimization will allow us to determine the performance gains achievable in realistic conditions.
An analysis of the current performance results will be presented, both for a complete LHC detector and a simplified sampling calorimeter setup. The individual sources of the observed performance gain will be discussed together with the experienced limitations, and an estimate of the final performance will be given.

Poster Session / 424

Federation of compute resources available to the German CMS community

Authors: Ralf Florian Von Cube[1]; Manuel Giffels[2]; Christoph Heidecker[2]; Gunter Quast[2]; Martin Benedikt Sauter[2]; Matthias Jochen Schnepf[2]

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

[2] *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Authors: gunter.quast@cern.ch, martin.sauter@student.kit.edu, matthias.jochen.schnepf@cern.ch, manuel.giffels@cern.ch, christoph.heidecker@cern.ch, ralf.florian.von.cube@cern.ch

The German CMS community (DCMS) as a whole can benefit from the various compute resources, available to its different institutes. While Grid-enabled and National Analysis Facility resources are usually shared within the community, local and recently enabled opportunistic resources like HPC centers and cloud resources are not. Furthermore, there is no shared submission infrastructure available.

In this contribution we present a concept, how connecting pools using HTCondors's flocking and routing mechanisms allows for a transparent and more efficient usage of all resources available to the DCMS community. In addition to the statically provisioned resources, also dynamically allocated resources from external cloud providers can be integrated. However, the usage of such dynamically allocated resources gives rise to additional complexity. Constraints on access policies of the resources, as well as workflow necessities have to be taken care of.

To maintain a well-defined and reliable execution environment on each resource, virtualization and containerization technologies such as virtual machines, Docker, and Singularity, are used.

We give an overview about the concepts and first experiences on how to provide DCMS resources dynamically and transparently to the community.

Track 1: Computing Technology for Physics Research / 425

An analytics driven computing model for HL-LHC

Authors: David Lange[1]; Frank Wurthwein[2]; Nick Smith[3]; Tommaso Boccali[4]

[1] *Princeton University (US)*

[2] *UCSD*

[3] *Fermi National Accelerator Lab. (US)*

[4] *INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P*

Corresponding Authors: tboccali@cern.ch, fkw@ucsd.edu, nick.smith@cern.ch, david.lange@cern.ch

The HL-LHC program has seen numerous extrapolations of its needed computing resources that each indicate the need for substantial changes if the desired HL-LHC physics program is to be supported within the current level of computing resource budgets. Drivers include detector upgrades, large increases in event complexity (leading to increased processing time and analysis data size) and trigger rates needed (5-10 fold increases) for the HL-LHC program. In this presentation, we discuss the newly developed modeling techniques in use for improving the accuracy of CMS computing resource needs for HL-LHC. Our emphasis is on monitoring-data driven techniques for model construction, parameter determination, and importantly, model extrapolations. Additionally we focus on uncertainty quantification as a critical component for understanding and properly interpreting our results.

Track 2: Data Analysis - Algorithms and Tools / 426

Submanifold Sparse Convolutional Networks for Sparse, Locally Dense Particle Image Analysis

Authors: Laura Domine[1]; Kazuhiro Terao[2]

[1] *Stanford University/SLAC*

[2] *SLAC National Accelerator Laboratory*

Corresponding Authors: kterao@fnal.gov, ldomine@stanford.edu

From a breakthrough revolution, Deep Learning (DL) has grown to become a de-facto standard technique in the fields of artificial intelligence and computer vision. In particular Convolutional Neural Networks (CNNs) are shown to be a powerful DL technique to extract physics features from images: They were successfully applied to the data reconstruction and analysis of Liquid Argon Time Projection Chambers (LArTPC), a class of particle imaging detectors which records the trajectory of charged particles in either 2D or 3D volumetric data with a breathtaking resolution (~3mm/pixel). The CNNs apply a chain of matrix multiplications and additions, and can be massively parallelized on many-core systems such as GPUs when applied on image data analysis. Yet a unique feature of LArTPC data challenges traditional CNN algorithms: it is locally dense (no gap in a particle trajectory) but generally sparse. A typical 2D LArTPC image has less than 1% of pixels occupied with non-zero value. This makes standard CNNs with dense matrix operations very inefficient. Submanifold sparse convolutional networks (SSCN) have been proposed to address exactly this class of sparsity challenges by keeping the same level of sparsity throughout the network. We demonstrate their strong performance on some of our data reconstruction tasks which include 3D semantic segmentation for particle identification at the pixel-level. They outperform a standard, dense CNN in an accuracy metric with substantially less computations. SSCN can address the problem of computing resource scalability for 3D DL-based data reconstruction chain R&D for LArTPC detectors.

Poster Session / 427

DIRAC: Putting together heterogeneous resources for scientific applications

Authors: Andrei Tsaregorodtsev[1]; Luisa Arrabito[None]; Johan Bregeon[None]; Vanessa HAMAR[2]; Sorina POP[3]; Fabio Hernandez[4]

[1] *Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France*

[2] *CC-IN2P3*

[3] *CNRS*

[4] *IN2P3/CNRS Computing Centre*

Corresponding Authors: arrabito@in2p3.fr, fabio@in2p3.fr, atsareg@in2p3.fr, sorina.pop@creatis.insa-lyon.fr, johan.bregeon@lupm.in2p3.fr, vanessa.hamar@cc.in2p3.fr

Scientific user communities are widely using computing and storage resources provided by large grid infrastructures. More and more capacity is provided by these infrastructures in a form of cloud resources. Cloud resources are much more flexible for usage but provide completely different access interfaces. Furthermore, grid infrastructure users are often getting access to extra computing resources like HPC centers, home lab clusters, etc. These resources are not part of any formal infrastructure and do not have uniform access mechanisms. Therefore, providing access to all these resources in a transparent way with a unique interface becomes a serious issue. In this contribution we describe the way in which the DIRAC Interware solution can be used to address this issue. The solution will be illustrated by the example usage of the DIRAC4EGI and FG-DIRAC services provided by the EGI and France-Grilles grid infrastructure projects.

Track 1: Computing Technology for Physics Research / **428**

The ATLAS EventIndex and its evolution towards LHC Run 3

Authors: Dario Barberis[1]; Evgeny Alexandrov[2]; Igor Aleksandrov[3]; Zbigniew Baranowski[4]; Gancho Dimitrov[4]; Alvaro Fernandez Casani[5]; Elizabeth Gallas[6]; Carlos Garcia Montoro[5]; Santiago González de la Hoz[7]; Julius Hrivnac[8]; Alexander Iakovlev[9]; Andrei Kazymov[2]; Mikhail Mineev[2]; Fedor Prokoshin[10]; Grigori Rybkin[8]; Javier Sanchez[5]; Jose Salt[5]; Petya Vasileva[11]; Miguel Villaplana[12]

[1] *Università e INFN Genova (IT)*

[2] *Joint Institute for Nuclear Research (RU)*

[3] *Joint Institute for Nuclear Research (JINR)*

[4] *CERN*

[5] *Univ. of Valencia and CSIC (ES)*

[6] *University of Oxford (GB)*

[7] *IFIC-Valencia*

[8] *Universite de Paris-Sud 11 (FR)*

[9] *Joint Inst. for Nuclear Research (JINR)*

[10] *Federico Santa Maria Technical University (CL)*

[11] *PU*

[12] *Università degli Studi e INFN Milano (IT)*

Corresponding Authors: miguel.villaplana.perez@cern.ch, dario.barberis@cern.ch, evgeny.alexandrov@cern.ch, alvaro.fernandez.casani@cern.ch, petya.tsv.vasileva@gmail.com, igor.alexandrov@cern.ch, mikhail.mineev@cern.ch, julius.hrivnac@cern.ch, javier.sanchez@cern.ch, elizabeth.gallas@physics.ox.ac.uk, jose.salt@cern.ch, sgonzale@ific.uv.es, gancho.dimitrov@cern.ch, andrei.kazymov@cern.ch, carlos.garcia.montoro@cern.ch, zbigniew.baranowski@cern.ch, fedor.prokoshin@cern.ch, alexander.iakovlev@cern.ch, grigori.rybkine@cern.ch

The ATLAS experiment produced so far hundreds of petabytes of data and expects to have one order of magnitude more in the future. This data are spread among hundreds of computing Grid sites around the world. The EventIndex is the complete catalogue of all ATLAS events, real and simulated, keeping the references to all permanent files that contain a given event in any processing stage. It provides the means to select and access event data in the ATLAS distributed storage system, and

provides support for completeness and consistency checks and trigger and offline selection overlap studies. The EventIndex employs various data handling technologies like Hadoop and Oracle databases, and is integrated with other systems of the ATLAS distributed computing infrastructure, including those for data, metadata, and production management. The project is in operation since the start of LHC Run 2 in 2015, and is in permanent development in order to fit the production and analysis demands and follow technology evolutions. The main data store in Hadoop, based on MapFiles and HBase, has worked well during Run 2 but new solutions are explored for the future. This paper reports on the current system performance and on the studies of a new data storage prototype that can carry the EventIndex through Run 3.

Track 1: Computing Technology for Physics Research / 429

Using DODAS as deployment manager for smart caching of CMS data management system

Authors: Daniele Spiga[1]; Tommaso Boccali[2]; Daniele Bonacorsi[3]; Giacinto Donvito[4]; Davide Salomoni[5]; Luciano Gaido[6]; Diego Ciangottini[7]; Valentin Y Kuznetsov[8]; Marica Antonacci[9]; Mirco Tracolli[10]; Doina Cristina Duma[9]

[1] *Universita e INFN, Perugia (IT)*

[2] *INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P*

[3] *University of Bologna*

[4] *INFN-Bari*

[5] *Universita e INFN, Bologna (IT)*

[6] *Universita e INFN (IT)*

[7] *INFN, Perugia (IT)*

[8] *Cornell University (US)*

[9] *INFN*

[10] *INFN Perugia*

Corresponding Authors: giacinto.donvito@ba.infn.it, daniele.bonacorsi@bo.infn.it, mirco.tracolli@cern.ch, marica.antonacci@ba.infn.it, cristina.aiftimiei@cnaf.infn.it, davide.salomoni@cnaf.infn.it, luciano.gaido@to.infn.it, daniele.spiga@cern.ch, diego.ciangottini@cern.ch, valentin.kuznetsov@cern.ch, tommaso.boccali@cern.ch

DODAS stands for Dynamic On Demand Analysis Service and is a Platform as a Service toolkit built around several EOSC-hub services designed to instantiate and configure on-demand container-based clusters over public or private Cloud resources. It automates the whole workflow from service provisioning to the configuration and setup of software applications. Therefore, such solution allows to use "any cloud provider", with almost zero effort. In this talk, we demonstrate how DODAS can be adopted as deployment manager to set up and manage compute resources and services, required to develop an AI solution for smart data caching. The smart caching layer may reduce the operational cost and increase flexibility with respect to regular centrally managed storage of the current CMS computing model. The cache space should be dynamically populated with the most requested data. In addition, clustering such caching systems will allow to operate them as Content Delivery System between data providers and end-users. Moreover a geographically distributed caching layer will be functional also to a data-lake based model, where many satellite computing centers might appear and disappear dynamically. In this context, our strategy is to develop a flexible and automated AI environment for smart management of the content of such clustered cache system. In this contribution we will describe the identified computational phases required for the AI environment implementation, as well as the related DODAS integration. Therefore we will start with the overview of the architecture for the pre-processing step, based on Spark, which has the role to prepare data for a Machine Learning technique. A focus will be given on the automation implemented through DODAS. Then, we will show how to train an AI-based smart cache and how we implemented a training facility managed through DODAS. Finally we provide a overview of the inference system, based on the CMS-TensorFlow as a Service and also deployed as a DODAS service.

Track 1: Computing Technology for Physics Research / 430

Cross-domain Data Access System for Distributed Sites in HEP

Authors: Qi Xu[1]; Zhenjing Cheng[2]; Yaodong Cheng[3]; Gang CHEN[4]

[1] *Institute of High Energy Physics, Chinese Academy of Sciences*

[2] *INSTITUE OF HIGH ENERGY PHYSICS*

[3] *IHEP*

[4] *INSTITUTE OF HIGH ENERGY PHYSICS*

Corresponding Authors: xuq@ihep.ac.cn, gang.chen@ihep.ac.cn, chengzj@ihep.ac.cn, chyd@ihep.ac.cn

A large amount of data is produced by large scale scientific facilities in high energy physics (HEP) field. And distributed computing technologies has been widely used to process these data. In traditional computing model such as grid computing, computing job is usually scheduled to the sites where the input data was pre-staged in. This model will lead to some problems includ-ing low CPU utilization, inflexibility, and difficulty in highly dynamic cloud environment. The paper proposed a cross-domain data access system (CDAS), which presents one same file system view at local and the remote sites, supporting directly data access on demand. Then the computing job can run everywhere no need to know where data is located. For the moment the system has been implemented including these functionalities such as native access for remote data, quick response, data transmission and management on demand based on HTTP, data block hash and store, uniform file view and so on. The test results showed the performance was much better than tradi-tional file system on high-latency WAN.

Track 2: Data Analysis - Algorithms and Tools / 433

The HEP.QPR Project: Quantum Pattern Recognition for Charged Particle Tracking

Authors: Frederic Bapst[1]; Wahid Bhimji[2]; Paolo Calafiura[2]; Steven Andrew Farrell[2]; Heather Gray[3]; Wim Lavrijsen[2]; Lucy Linder[1]; Illya Shapoval[4]

[1] *School of Engineering and Architecture of Fribourg (Switzerland)*

[2] *Lawrence Berkeley National Lab. (US)*

[3] *LBNL*

[4] *Lawrence Berkeley National Laboratory*

Corresponding Authors: pcalafiura@lbl.gov, wahid.bhimji@cern.ch, frederic.bapst@hefr.ch, lucylinder@hefr.ch, steven.farrell@cern.ch, heather.gray@cern.ch, ishapoval@lbl.gov, wim.lavrijsen@cern.ch

Universal Quantum Computing may still be a few years away, but we have entered the Noisy Intermediate-Scale Quantum era which ranges from D-Wave commercial Quantum Annealers to a wide selection of gate-based quantum processor prototypes. These provide us with the opportunity to evaluate the potential of quantum computing for HEP applications.
We will present early results from the DOE HEP.QPR project (https://hep-qpr.lbl.gov) on the impact of Quantum Computing on charged particle tracking. Due to the increase in data rates and event complexities, tracking has become one of the most pressing HEP computational problems.
In HEP.QPR, we are studying the potential of the Quantum Associative Memory (QuAM) and Quantum Annealing (QA) algorithms. QuAM provides in principle an exponential increase in storage capacity compared to the classical associative memory algorithm used, e.g., for LHC data triggering. We will present a prototype implementation of the QuAM protocols and analyze the topological limitations for running QuAM on the IBM Q chips - a family of superconducting gate-based quantum processors. We will review several difficulties integrating the end-to-end quantum pattern recognition into a real-time production workflow and discuss possible techniques for mitigation.

We will also present some promising results we achieved expressing the LHC track finding problem as a Quadratic Unconstrained Binary Optimization (QUBO) that can be solved using a D-Wave Quantum Annealer. We generated QUBOs that encode the pattern recognition problem at the LHC on the TrackML dataset, and we solved them using D-Wave qbsolv hybrid optimizer. Those early experiments achieved a performance exceeding 99% for purity, efficiency, and for the TrackML score at low track densities. We plan to extend the performance of such algorithms at higher track densities by improving seeding algorithms, geographical partitioning, and our QUBO models. We will also evaluate if the hybrid classical/quantum annealing approach used by qbsolv provides performance improvements compared to purely classical QUBO solvers.

Track 1: Computing Technology for Physics Research / 434

Fast Deep Learning on FPGAs for the Phase-II L0 Muon Barrel Trigger of the ATLAS Experiment

Author: Simone Francescato[1]

[1] *Sapienza Universita e INFN, Roma I (IT)*

Corresponding Author: simone.francescato@cern.ch

The Level-0 Muon Trigger system of the ATLAS experiment will undergo a full upgrade for HL-LHC to stand the challenging performances requested with the increasing instantaneous luminosity. The upgraded trigger system foresees to send RPC raw hit data to the off-detector trigger processors, where the trigger algorithms run on new generation of Field-Programmable Gate Arrays (FPGAs). The FPGA represents an optimal solution in this context, because of its flexibility, wide availability of logical resources and high processing speed. Studies and simulations of different trigger algorithms have been performed, and novel low precision deep neural network architectures (based on ternary dense and convnet networks) optimized to run on FPGAs and to cope with sparse data are presented. Both physics performances in terms of efficiency and fake rates, and FPGA logic resource occupancy and timing obtained with the developed algorithms are presented.

Poster Session / 435

CernVM-FS Container Image Integration

Author: Simone Mosciatti[1]

Co-authors: Jakob Blomer [2]; Gerardo Ganis [2]; Radu Popescu [2]

[1] *Politecnico di Milano (IT)*

[2] *CERN*

Corresponding Authors: simone.mosciatti@cern.ch, jakob.blomer@cern.ch, gerardo.ganis@cern.ch, radu.popescu@cern.ch

Linux containers have gained widespread use in High Energy Physics, be it for services using container engines such as containerd/kubernetes, for production jobs using container engines such as Singularity or Shifter, or for development workflows using Docker as a local container engine. Thus the efficient distribution of the container images, whose size usually ranges from a few hundred megabytes to a few tens of gigabytes, is becoming a pressing concern. Because container images show similar characteristics than scientific application stacks, unpacking the images in CernVM-FS can remedy the distribution issues provided that the container engine at hand is able to use such unpacked images from CernVM-FS.
In this contribution, we'll report on recent advances in the integration of Singularity, Docker, and containerd with CernVM-FS. We show improvements in the publishing of container images from a Docker registry that rely on new means of directly ingesting images tarballs. We'll also show a repository file system structure for storing container images that are optimized for storing both

container engines using flat root file systems (Singularity) as well as container engines using layers (containerd, Docker). To evaluate the benefits of our approach, we show concrete use cases and figures for production and development images from LHC experiment stored in the recently created “unpacked.cern.ch” repository.

Poster Session / 436

Track extrapolation and muon identification with GEANT4E in Belle II event reconstruction

Author: Leo Piilonen[1]

[1] *Virginia Tech*

Corresponding Author: piilonen@vt.edu

I describe the charged-track extrapolation and muon-identification modules in the Belle II data-analysis code framework (basf2). These modules use GEANT4E to extrapolate reconstructed charged tracks outward from the Belle II Central Drift Chamber into the outer particle-identification detectors, the electromagnetic calorimeter, and the K-long and muon detector (KLM). These modules propagate the position, momentum, 6-dimensional covariance matrix, and time of flight through the detailed Belle II detector geometry to permit comparison of the extrapolated track with the hits detected in the outer detectors. In the KLM, a Kalman filter is employed to adjust the extrapolation based on the matching measurements. Downward-moving cosmic rays in the upper half of the detector are back-propagated upward. Several modifications were made to permit GEANT4E to interoperate with GEANT4 within a single basf2 job for event generation+simulation+reconstruction and to expand the number of particle species that are extrapolated.

Poster Session / 437

Hardware Accelerated ATLAS Workloads on the WLCG

Authors: Lukas Alexander Heinrich[1]; Manuel Guth[2]; Alessandra Forti[3]

[1] *New York University (US)*

[2] *Albert Ludwigs Universitaet Freiburg (DE)*

[3] *University of Manchester (GB)*

Corresponding Authors: alessandra.forti@cern.ch, lukas.heinrich@cern.ch, manuel.guth@cern.ch

In recent years the usage of machine learning techniques within data-intensive sciences in general and high-energy physics in particular has rapidly increased, in part due to the availability of large datasets on which such algorithms can be trained as well as suitable hardware, such as graphics or tensor processing units which greatly accelerate the training and execution of such algorithms. Within the HEP domain, the development of these techniques has so far relied on resources external to the primary computing infrastructure of the WLCG. In this paper we present an integration of hardware-accelerated workloads into the Grid through the declaration of dedicated queues with access to hardware accelerators and the use of linux container images holding a modern data science software stack. A frequent use-case of in the development of machine learning algorithms is the optimization of neural networks through the tuning of their hyper parameters. For this often a large range of network variations must be trained and compared, which for some optimization schemes can be performed in parallel – a workload well suited for grid computing. An example of such a hyper-parameter scan on Grid resources for the case of Flavor Tagging within ATLAS is presented.

Track 2: Data Analysis - Algorithms and Tools / 438

Belle2VR – An Interactive Virtual Reality Visualization of GEANT4 Event Histories

Author: Leo Piilonen[1]

[1] *Virginia Tech*

Corresponding Author: piilonen@vt.edu

I describe a novel interactive virtual reality visualization of the Belle II detector at KEK and the animation therein of GEANT4-simulated event histories. Belle2VR runs on Oculus and Vive headsets (as well as in a web browser and on 2D computer screens, in the absence of a headset). A user with some particle-physics knowledge manipulates a gamepad or hand controller(s) to interact with and interrogate the detailed GEANT4 event history over time, to adjust the visibility and transparency of the detector subsystems, to translate freely in 3D, to zoom in or out, and to control the event-history timeline (scrub forward or backward, speed up or slow down). A non-expert uses the app - during public outreach events, for example - to explore the world of subatomic physics via electron-positron collision events in the Belle II experiment at the SuperKEKB colliding-beam facility at KEK in Japan. Multiple simultaneous users, wearing untethered locomotive VR backpacks and headsets, walk about a room containing the virtual model of the Belle II detector and each others' avatars as they observe and control the simulated event history. Developed at Virginia Tech by an interdisciplinary team of researchers in physics, education, and virtual environments, the simulation is intended to be integrated into the undergraduate physics curriculum. I describe the app, including visualization features and design decisions, and illustrate how a user interacts with its features to expose the underlying physics in each electron-positron collision event.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 439

Ntuples for NNLO processes

Author: Daniel Maitre[None]

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

In this contribution I will discuss the practicalities of storing events from a NNLO calculation on disk with the view of "replaying" the simulation for a different analysis and under different conditions, such as a different PDF fit or a different scale setting.

Poster Session / 441

Applying data access policy to CernVM-FS

Author: Minh Duc Nguyen[1]

Co-author: Alexander Kryukov [2]

[1] *Lomonosov Moscow State University Skobeltsyn Institute of Nuclear Physics*

[2] *M.V. Lomonosov Moscow State University (RU)*

Corresponding Authors: alexander.kryukov@cern.ch, nguyendmitri@gmail.com

CernVM-FS is a solution to scalable, reliable and low-maintenance software distribution that is widely used in various High Energy Physics collaborations. The information that can be distributed by CernVM-FS is not limited to software but any other data. By default, the whole CernVM-FS repository containing all subdirectories and files is available to all users in read-only mode after mounting.

This behaviour is not acceptable when CernVM-FS is used to distribute data that require different access levels. In this paper, we introduce an approach to apply data access policy to CernVM-FS so that a user can list and read only certain subdirectories and files according to his access level. All other subdirectories and files are hidden.

Poster Session / 442

A fast and parametric digitization for triple-GEM detectors

Authors: Lia Lavezzi[1]; Riccardo Farinelli[2]

[1] *Universita e INFN Torino (IT)*

[2] *Universita e INFN, Ferrara (IT)*

Corresponding Authors: rfarinelli@fe.infn.it, lia.lavezzi@cern.ch

Triple-GEM detectors are a well known technology in high energy physics. In order to have a complete understanding of their behavior, in parallel with on-beam testing, a Monte Carlo code has to be developed to simulate their response to the passage of particles. The software must take into account all the physical processes involved from the primary ionization up to the signal formation, *e.g.* the avalanche multiplication and the effect of the diffusion on the electrons. In the case of gas detectors, existing software such as GARFIELD already perform a very detailed simulation but are CPU time consuming. A description of a reliable but faster simulation is presented here: it uses a parametric description of the variables of interest obtained by suitable preliminary GARFIELD simulations and tuned to the test beam data. It can reproduce the real values of the charge measured by the strip, needed to reconstruct the position with the Charge Centroid method. In addition, particular attention was put to the simulation of the timing information, which permits to apply also the micro-Time Projection Chamber position reconstruction, for the first time on a triple-GEM. A comparison between simulation and experimental values of some sentinel variables in different conditions of magnetic field, high voltage settings and incident angle will be shown.

Poster Session / 443

Making RooFit Ready for Run 3

Author: Stephan Hageboeck[1]

Co-author: Lorenzo Moneta [2]

[1] *University of Bonn (DE)*

[2] *CERN*

Corresponding Authors: stephan.hageboeck@cern.ch, lorenzo.moneta@cern.ch

RooFit and RooStats, the toolkits for statistical modelling in ROOT, are used in most searches and measurements at the Large Hadron Collider. The data to be collected in Run 3 will enable measurements with higher precision and models with larger complexity, but also require faster data processing.

In this talk, first results on vectorising and multi-threading likelihood fits in RooFit will be presented. These improvements will enable the LHC experiments to process larger datasets without having to compromise with respect to model complexity.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 444

Techniques for the calculation of 4-loop QED contributions

Author: Stefano Laporta[1]

[1] *University of Padova*

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

We describe a little in detail some techniques used in the calculation of
4-loop QED contributions to some quantities like g-2,
slope of the Dirac form factor, renormalization constants;
in particular, some different approaches to the parallelization of some
parts of the calculations. Some recent results will be also presented.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 445

STR: a Mathematica package for the method of uniqueness

Author: Michelangelo Preti[1]

[1] *Ecole Normale Supérieure - CNRS*

Corresponding Author: michelangelo.preti@gmail.com

The Mathematica package STR (Star-Triangle Relations) is a recently developed tool designed to solve Feynman diagrams by means of the method of uniqueness in any (Euclidean) spacetime dimension D. The method of uniqueness is a powerful technique to solve multi-loop Feynman integrals in theories with conformal symmetry imposing some relations between D and the powers of propagators. In our algorithm we include both identities for scalar and Yukawa type integrals. The package is equipped with a graphical environment in which is possible to draw the desired diagram with the mouse input and a set of tools to modify and compute it. Throughout the use of a graphic interface, the package should be easily accessible to users with little or no previous experience on diagrams computation.

Poster Session / 447

A virtualized Tier-3g Facility Installiation in WLCG Network of CERN

Author: Ilkay Turk Cakir[1]

Co-authors: Agah Alici [2]; Hasan Saygın [2]; Sinan Kuday [3]

[1] *Giresun University*

[2] *Istanbul Aydin University*

[3] *Istanbul Aydin University (TR)*

Corresponding Authors: sinan.kuday@cern.ch, ilkay.turk.cakir@cern.ch

A Tier-3g Facility within the computing resources of Istanbul Aydin
University has been planned and installed in collaboration with TR-ULAKBIM national
Tier-2 center. The facility is intended to provide an upgraded data analysis
infrastructure to CERN researchers considering the recent nation-wide projects of ATLAS and
CMS experiments. The fundamental design of Tier-3g has been detailed in this work with an emphasis
on technical implementations of the following parts: Virtualization

of all nodes, VOMS usage for reaching fast experimental data in the WLCG network, batch cluster / multicore computing with HTCONDOR and PROOF systems, usage of grid proxies to access code libraries in AFS and CVMFS, dynamic disc space allocation and remote system mounting of EOS. We also present the performance test results that was obtained during a typical simulation of official analysis tools.

Track 1: Computing Technology for Physics Research / 448

New ROOT graphics language

Authors: Serguei Linev[1]; Iliana Betsou[2]; Bertrand Bellenot[3]; Olivier Couet[3]

[1] *GSI Darmstadt*

[2] *National Technical Univ. of Athens (GR)*

[3] *CERN*

Corresponding Authors: bertrand.bellenot@cern.ch, olivier.couet@cern.ch, s.linev@gsi.de, iliana.betsou@cern.ch

For two decades, ROOT brought its own graphics system abstraction based on a graphics model inspired by the popular graphics systems available at that time. (X11, OpenGL, Cocoa …)

With the emergence of modern C++ and recent graphics systems based on client/server models, it was time to redefined completely ROOT graphics.

This has been been done in the context of ROOT 7 which provides the new Graphics library using modern C++ serving JavaScript-based clients over the web.

This new approach re-think the High Energy Physics graphics language targeting the production of plots designed for usability with new graphics style and optimal defaults

Track 3: Computations in Theoretical Physics: Techniques and Methods / 449

Toward an efficient evaluation of two-loop massive scalar integralss

Authors: Jean-Philippe Guillet[1]; Eric Pilon[1]; Yoshimitsu Shimizu[2]; Sadek Zidi[3]

[1] *LAPTh CNRS*

[2] *KEK*

[3] *Jijel University, Algeria*

Corresponding Authors: guillet@lapth.cnrs.fr, pilon@lapth.cnrs.fr, zidi@lapth.cnrs.fr

A key ingredient in an automated evaluation of two-loop multileg processes is a fast and numerically stable evaluation of scalar Feynman integrals. In this respect, the calculation of two-loop three- and four-point functions in the general complex mass case so far relies on multidimensional numerical integration through sector decomposition whereby a reliable result has a high computing cost, whereas the derivation of a fully analytic result remains beyond reach. It would therefore be useful to perform part of the Feynman parameter integrations analytically in a systematic way to let only a reduced number of integrations to be performed numerically. Such a working program has been initiated for the calculation of massive two-loop N-point functions

using analytically computed building blocks. This approach is based on the implementation of two-loop scalar N-point functions in four dimensions ${}^{(2)}I_N^4$ as double integrals in the form:
$ ${}^{(2)}I_N^4 \sim \sum \int_0^1 d\rho \int_0^1 d\xi \, P(\rho,\xi) \; {}^{(1)}\widetilde{I}_{N+1}^4(\rho,\xi)$ wherethebuildingblocks^{(1)}\widetilde{I}{N+1}^{4}(\rho,\xi)involvedintheint
loop(N+1)−pointFeynman−typeintegrals, andwhereP(\rho,\xi)areweightingfunctions.The ^{(1)}\widetilde{I}{N+1}^{4}(\r
0 \leq z_{j} \leq 1, j = 1,\cdots,N+1; \sum_{j=1}^{N+1} z_{j}=1\}atworkfortheone−loop(N+1)−pointfunction, butanotherdom
loopN−pointfunctionconsidered.Thegeneralisationconcernsalsotheunderlyingkinematics, which, besidesexternalm
loop(N+1)−particleprocessesatcolliders.Theonlytworemainingintegrationsover\rho,\xitobeperformednumericallyre
loopintegrals.

As a first step in this direction, the method developped has been successfully applied to the usual one-loop four-point function for arbitrary masses and kinematics as a "proof of concept", showing its ability to circumvent the subtleties of the various analytic continuations in the kinematical variables in a systematic way, in a series ofthree articles. The target work, namely its practical implementation to compute the building blocks ^{(1)}\widetilde{I}_{N+1}^{4}(\rho,\xi)$ istobeelaboratedandpresented inafutureseriesofarticles.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 451

Renormalization of gauge theories at five loops

Author: York Schröder[1]

[1] *UBB Chillán*

Corresponding Author: yorkschroder@gmail.com

I briefly review the recently finished 5-loop renormalization program of QCD, and explain the status and prospects of the computer-algebraic techniques involved.

Track 2: Data Analysis - Algorithms and Tools / 452

Towards the Increase in Granularity for the Main Hadronic ATLAS Calorimeter: Exploiting Deep Learning Methods

Corresponding Author: philipp.do.nascimento.gaspar@cern.ch

An extensive upgrade programme has been developed for LHC and its experiments, which is crucial to allow the complete exploitation of the extremely high-luminosity collision data. The programme is staggered in two phases, so that the main interventions are foreseen in Phase II.
For this second phase, the main hadronic calorimeter of ATLAS (TileCal) will redesign its readout electronics but the optical signal pathway will be kept unchanged.
However, there is a technical possibility to increase the detector granularity, without changing its mechanical structure, by modifying only the calorimeter readout. During the high luminosity regime, particle jets with high transverse momentum tend to deposit its energy in the last layers of TileCal. Therefore, dividing the actual calorimeter cells into new subregions will improve momentum reconstruction, mass, transverse energy and angular position of those jets, allowing future analysis benefit from a finer-grained granularity detector.
The light emitted by the calorimeter tiles is collected by a set of WLS fibers grouped in a bundle, one per calorimeter cell, coupled by a light mixer to a Photomultiplier Tube (PMT).
Aiming at extracting additional information on the spatial distribution of the energy deposited within each cell, the original PMT is substituted by a Multi-Anode Photomultiplier
Tube (MA-PMT) with 64 photosensors distributed in a grid of 8 x 8 pixels. This makes it possible to increase the detector granularity by means of an algorithm whose purpose is to match the image pattern formed in the grid of pixels to a topological subregion within a given cell. Calibration data

are used for algorithm development, which is costly to produce, demanding time and man power. Therefore, the Generative Adversarial Network (GAN) is used to simulate the interaction of particles in a calorimeter cell, and, thereafter, leveraging the amount of statistics for the final classification model development.
Using a variant of the GAN model based on deep layers (DCGAN), a substantial increase in the number of images was obtained. As a consequence, a supervised deep learning approach based on Convolutional Neural Network (CNN) could be developed for mapping the signal image information onto two regions of the 8 x 8 grid. During the development
stage, the synthetic images produced with the generative model were used to train the CNN and its performance was evaluated in real calibration data. The preliminary results show an accuracy of more than 95% in both splits. This is encouraging for a possible solution of a such important step in calorimeter upgrade in the ATLAS experiment.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 454

Algorithm to find an all-order in the running coupling solution to an equation of the DGLAP type

Authors: Gustavo Álvarez[None]; Igor Kondrashuk[1]

[1] *UBB*

Corresponding Authors: igor.kondrashuk@gmail.com, gustavo_0408@hotmail.com

We propose an algorithm to find a solution to an integro-differential equation of the DGLAP type for all the orders in the running coupling α with splitting functions given at a fixed order in α. Complex analysis is significantly used in the construction of the algorithm, we found a way to calculate the involved integrals over contours in the complex planes in more simple way than by any of the methods known at present. Then, we write a code in Mathematica based on the proposed algorithm. We apply these algorithm and code to the DGLAP equation for singlet parton distributions of QCD and compare our solution with the results which may be obtained by using the existing numerical or symbolical software tools, for example, we compare it in this talk with the results obtained for singlet parton distribution functions by using QCDNUM.

Poster Session / 455

JupyterLab at the Saint Petersburg State University

Authors: Andrey Erokhin[1]; Andrey Zarochentsev[1]

[1] *St Petersburg State University (RU)*

Corresponding Authors: andrey.erokhin@cern.ch, andrey.zarochencev@cern.ch

The poster focuses on our experience in usage and extending of JupyterLab
in combination with EOS and CVMFS for HEP analysis within a local university group.
We started with a copy of CERN SWAN environment, after that our project evolved independently.
A major difference is that we switched from classic Jupyter Notebook
to JupyterLab, because our users are more insterested in text editor plus
terminal workflow rather than in Notebook workflow.
Like in SWAN, we are using CVMFS to load Jupyter kernels
and other software, and EOS to store user data.

Poster Session / 456

GRAAL: Gem Reconstruction And Analysis Library

Authors: Riccardo Farinelli[1]; Lia Lavezzi[2]; Gianluigi Cibinetto[3]

[1] *Universita e INFN, Ferrara (IT)*

[2] *Universita e INFN Torino (IT)*

[3] *INFN Ferrara*

Corresponding Authors: cibinett@fe.infn.it, lia.lavezzi@cern.ch, rfarinelli@fe.infn.it

MPGD are the new frontier in gas trackers. Among this kind of devices, the GEM chambers are widely used. The experimental signals acquired with the detector must obviously be reconstructed and analysed. In this contribution, a new offline software to perform reconstruction, alignment and analysis on the data collected with APV-25 and TIGER ASICs will be presented. GRAAL (Gem Reconstruction And Analysis Library) is able to measure the performance of a MPGD detector with a strip segmented anode (presently). The code is divided in three parts: reconstruction, where the hit s are digitized and clusterized; tracking, where an procedure fits the points from the tracking system and uses that information to align the chamber with rotations and shifts; analysis, where the performance is evaluated (e.g. efficiency, spatial resolution,etc.). The user must set the geometry of the setup and then the program returns automatically the analysis results, taking care of different conditions of gas mixture, electric field, magnetic field, geometries, strip orientation, dead strip, misalignment and many others.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 457

Computer algebra algorithm of simplification of tensor polynomials

Author: Alexander Kryukov[1]

Co-author: Grigory Shpiz [1]

[1] *SINP MSU*

Corresponding Authors: kryukov@theory.sinp.msu.ru, gregory.shpiz@gmail.com

Tensor calculations are an important case in many natural sciences like mathematics and physics. To simplify such expressions, computer algebra is widely used. There are a number of approaches for solving this problem, namely, the component calculations, the calculations when tensor is considered as an abstract symbol with indices possessing some symmetry properties, and finally a pure abstract calculus like exterior algebra. In this paper we describe a method for reducing tensor expressions to the canonical form using group algebra method. The focus is on taking into account symmetry properties with respect to various permutations of indices inside the tensor, symmetries associated with renaming summation indices, as well as linear relationships between tensors, such as Bianchi identities. We define the canonical representation for polynomial (multiplicative) tensor expressions, which is the result of averaging the tensor expression over the orbits of certain finite group (stabilizer). We present a sketch of the algorithm which reduces tensor expression to the canonical form. The approach is illustrated by examples containing Riemann curvature tensors.

Poster Session / 458

Overall quality optimization for DQM stage in High Energy Physics experiments

Authors: Nectarios Benekos[1]; Manuel Jesus Parra Royon[2]; Jose Benitez[3]

[1] *University of Peloponnese (GR)*

[2] *Universidad de Granada (ES)*

[3] *DECSAI, University of Granada, Spain*

Corresponding Authors: manuel.jesus.parra.royon@cern.ch, j.m.benitez@decsai.ugr.es, nectarios.benekos@cern.ch

Data Quality Monitoring (DQM) is a very significant component of all high-energy physics (HEP) experiments. Data recorded by Data Acquisition (DAQ) sensors and devices are sampled to perform live monitoring of the status of each detector during data collection. This gives to the system and scientists the ability to identify problems with extremely low latency, minimizing the amount of data that would otherwise be unsuitable for physical analysis. In the offline DQM environment, the system is used to review the results of the final reconstruction of data on a continuous baseline, serving as the certified database used in all scientific analysis. DQM performs a large set of operations on the data, such as Fast Fourier Transform (FFT), Clustering algorithms, Region of Interest (RoI) or Classification problems, for instance. All those operations suppose an intensive processing workflow with massive volumes of data to be performed on High Performance Computing (HPC) platforms.
This workflow involves the use of HPC resources and time, both depending on the number of events of the experiment, the complexity of the tasks or the data quality levels, among others. The use of Machine Learning (ML) techniques in the DQM stage can significantly improve the general performance in the execution of these workflows, achieving a better performance in the management and distribution of the processes that run on HPC.
In this context,the main objective of our work is the application of ML algorithms in order to improve the use of HPC resources and increase the overall quality levels of the data processed in DQM. As an effective solution a Multi-Objective Evolutionary Algorithm (MOEA) has been proposed to accomplish this optimization, considering objectives, resources and constraints. MOEA has been used due to its good balance between computational cost and quality solutions reached.
Finally a tool for decision making has been deployed allowing the scientists to decide more efficiently on DQM process.

Track 1: Computing Technology for Physics Research / 459

A more Pythonic, Interoperable and Modern PyROOT

Author: Stefan Wunsch[1]

Co-author: Enric Tejedor Saavedra [2]

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

[2] *CERN*

Corresponding Authors: enric.tejedor.saavedra@cern.ch, stefan.wunsch@cern.ch

PyROOT is the name of ROOT's Python bindings, which allow to access all the ROOT functionality implemented in C++ from Python. Thanks to the ROOT type system and the Cling C++ interpreter, PyROOT creates Python proxies for C++ entities on the fly, thus avoiding to generate static bindings beforehand.

PyROOT is in the process of being enhanced and modernised to meet the demands of the HEP Python community. In particular, the ongoing work in PyROOT comprises three areas: first, making PyROOT more pythonic, by adding the so-called "pythonisations" to make it simpler to access C++ from Python; second, improving the interoperability of PyROOT with the Python data science ecosystem tools (for instance, NumPy); third, redesigning PyROOT on top of the Cppyy library, in order to benefit from the modern C++ features supported by the latter.

Track 1: Computing Technology for Physics Research / 460

Generative Adversarial Networks for fast simulation: generalisation and distributed training in HPC

Authors: Sofia Vallecorsa[1]; Jean-Roch Vlimant[2]; Vladimir Loncar[3]; Thong Nguyen[2]; Gulrukh Khattak[1]; Maurizio Pierini[1]; Federico Carminati[1]; Felice Pantaleo[1]

[1] *CERN*

[2] *California Institute of Technology (US)*

[3] *University of Belgrade (RS)*

Corresponding Authors: felice.pantaleo@cern.ch, thong@caltech.edu, maurizio.pierini@cern.ch, jean-roch.vlimant@cern.ch, gul.rukh.khattak@cern.ch, federico.carminati@cern.ch, sofia.vallecorsa@cern.ch, vladimir.loncar@cern.ch

Deep Learning techniques have are being studied for 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 describe an R&D activity within CERN openlab, 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 in just a small fraction of time. We will present the first application of three-dimensional convolutional Generative Adversarial Networks to the simulation of high granularity electromagnetic calorimeters. We will describe detailed validation studies comparing our results to Geant4 Monte Carlo simulation, showing, in particular, the very good agreement we obtain for high level physics quantities (such as energy shower shapes) and detailed calorimeter response (single cell response). Finally we will show how this tool can easily be generalized to describe a larger class of calorimeters, opening the way to a generic machine learning based fast simulation approach. To achieve generalization we will leverage advanced optimization algorithms (using Bayesian and/or Genetic approach) and apply state of the art data parallel strategies to distribute the training process across multiple nodes in HPC and Cloud environment. Performance of the parallelization of GAN training on HPC clusters will also be discussed in details.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 461

Recent developments of GRACE system

Author: Yoshimasa Kurihara[1]

[1] *KEK*

Corresponding Author: yoshimasa.kurihara@kek.jp

The Grace system is an automatic system to calculate cross sections based on the standard model and MSSM including one-loop corrections.
I would like to report recent progress of the GRACE system including optimization of generated codes.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 462

HepMC3 Event Record Library for Monte Carlo Event Generators

Authors: Andrii Verbytskyi[1]; Andy Buckley[2]; David Grellscheid[3]; Dima Konstantinov[4]; James William Monk[5]; Leif Lönnblad[6]; Tomasz Przedzinski[7]; Witold Pokorski[8]

[1] *Max-Planck-Institut fur Physik (DE)*

[2] *University of Glasgow (GB)*

[3] *IPPP Durham*

[4] *IHEP*

[5] *Niels Bohr Institute, University of Copenhagen*

[6] *Lund University (SE)*

[7] *Jagiellonian University (PL)*

[8] *CERN*

Corresponding Authors: d.grellscheid@gmail.com, dkonst@mail.cern.ch, andrii.verbytskyi@cern.ch, tomasz.przedzinski@cern.ch, a.g.buckley@gmail.com, witold.pokorski@cern.ch, leif.lonnblad@thep.lu.se, james.william.monk@cern.ch

We present the HepMC3 library designed to perform manipulations with event records of High Energy Physics Monte Carlo Event Generators (MCEGs). The library is a natural successor of HepMC and HepMC2 libraries used in the present and in the past. HepMC3 supports all functionality of previous versions and significantly extends them.

In comparison to the previous versions, the default event record has been simplified, while an option to add arbitrary information to the event record has been implemented. Particles and vertices are stored separately in an ordered graph structure, reflecting the evolution of a physics event and enabling usage of sophisticated algorithms for event record analysis.

The I/O functionality of the library has been extended to support common input and output formats of HEP MCEGs, including formats used in Fortran HEP MCEGs, formats used in HepMC2 library and ROOT. The functionality of the library allows user to implement customized input or output format.

The library is already supported by popular modern MCEGs (e.g. Sherpa and Pythia8) and can replace the older HepMC versions in many others.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 463

Machine Learning helping Monte-Carlo collider simulations

Author: Valentin Hirschi[1]

[1] *ETHZ*

Corresponding Author: valentin.hirschi@gmail.com

The high-energy community recently witnessed the first attempts at leveraging machine (deep) learning techniques for improving the efficiency of the numerical Monte-Carlo integrations that lie at the core of most high-energy physics simulations.
The first part of my talk will characterise the various type of integrations necessary in these simulations as well as the type of improvements that could significantly impact their efficiency.
The second part will focus on reviewing the objectives and achievements of the first attempts at applying modern machine learning techniques in this context.

Poster Session / 464

Boosting Performance of Data-intensive Analysis Workflows with Distributed Coordinated Caching

Authors: Gunter Quast[1]; Manuel Giffels[1]; Max Fischer[2]; Eileen Kuehn[1]; Matthias Jochen Schnepf[1]; Ralf Florian Von Cube[3]; Martin Benedikt Sauter[1]; Christoph Heidecker[1]

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

[2] *GSI - Helmholtzzentrum fur Schwerionenforschung GmbH (DE)*

[3] *Rheinisch-Westfaelische Tech. Hoch. (DE)*

Corresponding Authors: ralf.florian.von.cube@cern.ch, christoph.heidecker@cern.ch, matthias.jochen.schnepf@cern.ch, manuel.giffels@cern.ch, gunter.quast@cern.ch, max.fischer@cern.ch, martin.sauter@student.kit.edu, eileen.kuehn@kit.edu

Data-intensive end-user analyses in High Energy Physics requires high data throughput to reach short turnaround cycles.
This leads to enormous challenges for storage and network infrastructure, especially when facing the tremendously increasing amount of data to be processed during High-Luminosity LHC runs.
Including opportunistic resources with volatile storage systems into the traditional HEP computing facilities makes this situation more complex.

Bringing data close to the computing units is a very promising approach to solve throughput limitations and improve the overall performance.
We focus on coordinated distributed caching, where we coordinate the placement of critical data on distributed caches and match work-flows to the most suitable host in terms of cached files.
The coordination of data allows to efficiently use limited cache volume by reducing redundant data storage on distributed caches.
In addition, workflow coordination optimizes overall processing efficiency by improving data access for data-intensive analysis workflows.

The NaviX coordination service developed at KIT realizes this concept by connecting an XRootD cache proxy server infrastructure with an HTCondor batch system.
The usage of distributed caches on opportunistic resources was tested to enable efficient processing of data-intensive workflows there.
In addition, after successfully running a prototype system, we are building a Throughput-Optimized Analysis-System (TOPAS), where about 600 CPU cores are directly connected to a distributed 1PB cache and 11 NVME SSD 1TB caches.
Our system with coordinate distributed caches enables fast analysis of large amounts of data as required for future HEP experiments.

In this contribution, we provide an overview of the concept and the experience gained in coordinated distributed caching.

Track 1: Computing Technology for Physics Research / 465

Towards a heterogeneous High Level Trigger farm for CMS

Author: Andrea Bocci[1]

Co-authors: Marco Rovere [1]; Felice Pantaleo [1]; Vincenzo Innocente [1]; Matti Kortelainen [2]

[1] *CERN*

[2] *Fermi National Accelerator Lab. (US)*

Corresponding Authors: matti.kortelainen@cern.ch, vincenzo.innocente@cern.ch, felice.pantaleo@cern.ch, andrea.bocci@cern.ch, marco.rovere@cern.ch

The CMS experiment has been designed with a two-level trigger system: the Level 1 Trigger, implemented on custom-designed electronics, and the High Level Trigger (HLT), a streamlined version of the CMS offline reconstruction software running on a computer farm. A software trigger system requires a trade-off between the complexity of the algorithms running on the available computing resources, the sustainable output rate, and the selection efficiency.

During its "Phase 2" the LHC will reach a luminosity of $7\times10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ with a pileup of 200 collisions. To fully exploit the higher luminosity, the CMS experiment will increase the full readout rate from 100 kHz to 750 kHz. The higher luminosity, pileup and input rate present an unprecedented challenge to the HLT, that will require a processing power larger than today by at least a factor 20. This exceeds by far the expected increase in processing power for conventional CPUs, demanding an alternative approach.

Industry and HPC have been successfully using heterogeneous computing platforms, that can achieve higher throughput and better energy efficiency by matching each job to the most appropriate architecture.

The reliable use of a heterogeneous platform at the HLT since the beginning of Phase 2 requires the careful assessment of its performance and characteristics, which can only be attained by running a prototype in production already during Run 3. The integration of heterogeneous computing in the CMS reconstruction software depends upon improvements to its framework and scheduling, together with a tailoring of the reconstruction algorithms to the different architectures.

This R&D work began in 2017, and by the end of 2018 produced a demonstrator working in realistic conditions. This presentation will describe the results of the development and the characteristics of the system, along with its future perspectives.

Poster Session / 466

Using Continuous Deployment techniques to manage software change at a WLCG Tier-2

Authors: Gareth Douglas Roy[1]; Samuel Cadellin Skipsey[None]; Gordon Stewart[2]; Emanuele Simili[2]; David Britton[None]

[1] *University of Glasgow (GB)*

[2] *University of Glasgow*

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

In software development Continuous Integration (CI), the practice of bringing together multiple developers' code modifications into a single repository, and Continuous Delivery (CD), the practice of automatedly creating and testing releases are well known. CI/CD pipelines are available in many automation tools (such as GitLab) and act to enhance and speed up software development.

Continuous Deployment, the next step which is the practice of pushing releases to the production environment through automation is not as well established due to business or legal requirements. In the WLCG we have no such limitations and it is therefore a good candidate for Continuous Deployment of Grid resources.

In this paper we develop the work presented previously on containerised worker node environments and show how the use of Continuous Deployment techniques (and tooling) in conjunction with CI/CD can enhance the management of a WLCG Tier-2 resource, reducing downtime due to code changes and middleware updates.

Poster Session / 467

HEP Analyses on Dynamically Allocated Opportunistic Computing Resources

Author: Matthias Jochen Schnepf[1]

Co-authors: Ralf Florian Von Cube [2]; Max Fischer [3]; Manuel Giffels [1]; Christoph Heidecker [1]; Andreas Heiss [1]; Eileen Kuehn [1]; Andreas Petzold [1]; Gunter Quast [1]; Martin Benedikt Sauter [1]

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

[2] *Rheinisch-Westfaelische Tech. Hoch. (DE)*

[3] *GSI - Helmholtzzentrum fur Schwerionenforschung GmbH (DE)*

Corresponding Authors: martin.sauter@student.kit.edu, eileen.kuehn@kit.edu, ralf.florian.von.cube@cern.ch, andreas.petzold@cern.ch, matthias.jochen.schnepf@cern.ch, manuel.giffels@cern.ch, max.fischer@cern.ch, christoph.heidecker@cern.ch, gunter.quast@cern.ch, andreas.heiss@kit.edu

The ever growing amount of HEP data to be analyzed in the future requires as of today the allocation of additional, potentially only temporary available non-HEP dedicated resources. These so-called opportunistic resources are also well-suited to cover the typical unpredictable peak demands for computing resources in end-user analyses. However, their temporary availability requires a dynamic allocation, integration, and management, while their heterogeneity results in the challenge to allocate always the best matching resources in order to maintain a high resource utilization.

For this purpose, we are developing the multi-agent resource manager TARDIS (Transparent Adaptive Resource Dynamic Integration System) which reacts on the current utilization of the integrated resources. A feedback loop implemented by COBalD (Opportunistic Balancing Daemon) ensures the further allocation of well-used resources while reducing the amount of insufficiently used ones. TARDIS is able to allocate and manage resources from various resource providers such as HPC centers or commercial and public clouds while ensuring a dynamic allocation and efficient utilization of these heterogeneous opportunistic resources.

Furthermore, opportunistic resources are integrated into one overlay batch system which provides a single point of entry for all users. In order to provide the dedicated HEP software environment virtualization and container technologies are used.

In this contribution, we will give an overview of our developments of TARDIS/COBalD as well as the current status of the integration of opportunistic resources for HEP user analyses.

Track 1: Computing Technology for Physics Research / 468

The Practitioner's C++17 and 20 News

Author: Axel Naumann[1]

[1] *CERN*

Corresponding Author: axel.naumann@cern.ch

When it comes to number-crunching, C++ is at the core of HENP's software. But while C++17 is old news, many of us did not get to use it yet. And why would we? This presentation introduces some of the main reasons to move to C++17 - focusing on performant, readable code and robust interfaces.

Where C++17 has many new features that help, C++20 might come as "your next C++11", a major step forward for C++: it will most likely introduce concepts, contracts and ranges; fairly likely the "spaceship operator", coroutines and networking. Some of these will change the way we want to write code. As today's compilers are already implementing many of tomorrow's features, now is a good time to see where C++ is heading, and to learn how this affects our usage of C++.

Track 1: Computing Technology for Physics Research / 469

STAR Data Production Workflow on HPC: Lessons Learned & Best Practices

Authors: Michael Poat[1]; Jerome LAURET[1]

Co-authors: Jefferson Porter ; Jan Balewski

[1] *Brookhaven National Laboratory*

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

The Solenoidal Tracker at RHIC (STAR) is a multi-national supported experiment located at Brookhaven National Lab. The raw physics data captured from the detector is on the order of tens of PBytes per data acquisition campaign, which makes STAR fit well within the definition of a big data science experiment. The production of the data has typically run on standard nodes or on standard Grid computing environments. All embedding simulations (complex workflow mixing real and simulated events) have been run a standard Linux resources at NERSC aka PDSF. However, HPC resources such as Cori have become available for STAR's data production as well as embedding, and STAR has been the very first experiment to show feasibility of running a sustainable data production campaign on this computing resource.

The use of Docker containers with Shifter is required to run on HPC @ NERSC – this approach encapsulates the environment in which a standard STAR workflow runs. From the deployment of a tailored Scientific Linux environment (requiring many of its own libraries and special configurations required to run) to the deployment of third-party software and the STAR specific software stack, it has become impractical to rely on a set of containers containing each specific software release. To this extent, solutions based on CVMFS for the deployment of software and services have been employed in HENP, but one needs to make careful scalability considerations when using a resource like Cori, such as not allowing all software to be deployed in containers or bare node. Additionally, CVMFS clients are not compatible on Cori nodes and one needs to rely on an indirect NFS mount scheme. In our contribution, we will discuss our strategies from the past and our current solution based on CVMFS. Furthermore, running on HPC is not a simple task as each aspect of the workflow must be enabled to scale, run efficiently, and the workflow needs to fit within the boundaries of the provided queue system (SLURM in this case). Lastly, we will also discuss what we have learned to be the best method for grouping jobs to maximize a single 48 core HPC node within a specific time frame and maximize our workflow efficiency.
We hope both aspects will serve the community well as well as those following the same path.

Poster Session / 470

Aligning the MATHUSLA test stand detector: Using TensorFlow

Author: Gordon Watts[1]

[1] *University of Washington (US)*

Corresponding Author: gwatts@uw.edu

MATHUSLA has been proposed as a second detector that sits over 100m from an LHC interaction point, on the surface, to look for ultra long-lived particles. A test stand was constructed with 2 layers of scintillator paddles and 3 layers of RPC's, on loan from the DZERO and Argo-YBJ. Downward and upward going tracks from cosmics and muons from the interaction point have been reconstructed. To align the detector we have used 10K tracks and doing a simultaneous fit to find the location of the detectors. We used TensorFlow to drive the fit and align the detector. All aspects of this work will be discussed, from the technical details of how the alignment equations were converted to matrices for the TensorFlow kernel, the alignment equations expressed in TensforFlow language, and, finally, some high-level comments on the approach.

Poster Session / 471

DIRACOS: a cross platform solution for grid tools

Authors: Ben Couturier[1]; Christophe Haen[1]; Marko Petric[1]

[1] *CERN*

Corresponding Authors: marko.petric@cern.ch, christophe.denis.haen@cern.ch, ben.couturier@cern.ch

All grid middleware require external packages to interact with computing elements, storage sites… In the case of the DIRAC middleware this was historically divided into two bundles, one called externals containing Python and standard binary libraries and the other called the LCGBundle containig libraries form the grid world (gfal, arc, etc). The externals were provided for several platforms whereas the LCGBundle was released only for SLC6 and CC7. Such a setup was difficult to test and hindered agile development of DIRAC. Therefore we developed DIRACOS to produce a coherent environment for grid interaction. Additionally it solves the binary incompatibility we reached by using a python version newer than the native system one (SLC6). It is spawned form a single list of required packages from where we use SRPMs to pull all dependencies down to the level of glibc. With such an approach we can provide the same packages for our clients, servers, and several platforms. It is an extendible setup with an agile development cycle in mind. The core build functionality of DIRACOS is based on Fedora Mock. DIRACOS also introduces its own grammar, and to handle specific cases it also allows patching (some SRPM require tweeking, which the user can do by providing a diff) as well as routines for pre/post/instead actions of compilation. With this approach DIRAC was able to provide a single bundle for clients and servers, that is reliable, flexible, easy to test and relatively small (250 MB). It allows for a smooth transition from SLC6 to CC7 and provides a clear roadmap for possible extension of DIRAC to a myriad of platforms.

Track 2: Data Analysis - Algorithms and Tools / 472

Constructing mass-decorrelated hadonic decay taggers in ATLAS

Corresponding Author: andreas.sogaard@cern.ch

A large number of physics processes as seen by ATLAS at the LHC manifest as collimated, hadronic sprays of particles known as 'jets'. Jets originating from the hadronic decay of a massive particle are commonly used in searches for both measurements of the Standard Model and searches for new physics. The ATLAS experiment has employed machine learning discriminants to the challenging task of identifying the origin of a given jet, but such multivariate classifiers exhibit strong non-linear correlations with the invariant mass of the jet, complicating many analyses which wish to make use of the mass spectrum. Adversarially trained neural networks (ANN) are presented as a way to construct mass-decorrelated jet classifiers by jointly training two networks in a domain-adversarial fashion. The use of neural networks further allows this method to benefit from high-performance computing platforms for fast development. A comprehensive study of different mass-decorrelation techniques is performed in ATLAS simulated datasets, comparing ANNs to designed decorrelated taggers (DDT), fixed-efficiency k-NN regression, convolved substructure (CSS), and adaptive boosting for uniform efficiency (uBoost). Performance is evaluated using metrics for background jet rejection and mass-decorrelation.

Poster Session / 473

Identifying hadronically decaying vector bosons and top quarks in ATLAS

Author: Steven Schramm[1]

[1] *Universite de Geneve (CH)*

Corresponding Authors: oliver.majersky@cern.ch, steven.schramm@cern.ch

Hadronic decays of vector bosons and top quarks are increasingly important to the ATLAS physics program, both in measurements of the standard model and searches for new physics. At high energies, these decays are collimated into a single overlapping region of energy deposits in the detector, referred to as a jet. However, vector boson and top quarks are hidden under an enormous background of other processes producing jets. The ATLAS experiment has employed boosted decision trees and deep neural networks to the challenging task of identifying hadronically decaying vector boson and top quarks and rejecting other jet backgrounds. These discriminants are becoming increasingly complex and using more advanced machine learning techniques. The methods currently used to tag these objects are described. In order to improve the tagger performance on the signal efficiency and background rejection, new in-situ techniques are applied, thus directly evaluating the agreement between data and simulation after applying an arbitrarily complex classifier. The precision obtained by applying the in-situ techniques is presented.

Poster Session / 474

Distributed Computing Services in Belle II

Authors: Benedikt Hegner[1]; Carlos Fernando Gamboa[1]; Hironori Ito[1]; Maxim Potekhin[1]; Paul James Laycock[1]; Ruslan Mashinistov[1]; Siarhei Padolski[1]

[1] *Brookhaven National Laboratory (US)*

Corresponding Authors: ruslan.mashinistov@cern.ch, maxim.potekhin@cern.ch, cgamboa@bnl.gov, benedikt.hegner@cern.ch, siarhei.padolski@cern.ch, hito@bnl.gov, paul.james.laycock@cern.ch

In 2019 Belle II will start the planned physics runs with the entire detector installed. Compared to current collider experiments at the LHC, where all critical services are provided by the CERN as host lab and only storage and CPU resources are provided externally, Belle II and KEK chose a different, more distributed strategy. In particular, it provides easier access to existing expertise and resources at the participating institutions.

Many of the services are hosted outside the host lab. DESY runs the suite of collaborative tools for issue tracking and code management, and the Brookhaven National Lab runs the data management and conditions infrastructure. Proper orchestration of these services and sites is critical. Thus choosing this service distribution model increases both the pressure for and provides opportunities for using community-wide or industry-standard solutions. Better standardization allows to eventually implement fallback instances for the most critical services, which is very hard in the setups chosen by previous experiments.

We will present our experience with setting up and running computing services of Belle II at BNL, the challenges this system and the cross-lab handshakes pose, and where and why we base the work on widely accepted tools like e.g. RUCIO. In addition, we will give an outlook for where we see future potential for more community-wide solutions on tackling conditions and other important experiment services

Poster Session / 475

Exploring Hybrid Deep Learning Architecture for Experimental HEP Data Analysis

Author: Yavar Taheri Yeganeh[1]

[1] *Shahid Beheshti University*

Corresponding Author: y.taheriyeganeh@mail.sbu.ac.ir

Deep learning has shown a promising future in physics' data analysis and is anticipated to revolutionize LHC discoveries.
Designing an optimal algorithm may seem to be the most challenging task in machine learning progress especially in HEP due to the high dimensionality and extreme complexity of the data.
Physical knowledge can be employed in designing and modifying of the algorithm's modules as well as constructing high-level features, however, few researchers suggested that it may be sup-optimal (especially latter case).
Hybrid architectures aims to achieve a complicated target based on fusion of different modules such as Convolutional Neural Network. This topic has been exploited in several computer vision researches, consequently, it can also be considered for Jet physics.
Hybrid Deep Learning Architecture concerned with taking full advantage of expertise on the particular environment of the task.

Poster Session / 476

The DQM system for the SND

Author: Konstantin Pugachev[1]

[1] *Budker Institute of Nuclear Physics (RU)*

Corresponding Author: sekrasoft@gmail.com

The SND detector has been operating at the VEPP-2000 collider (BINP, Russia) for several years unveiling amazing knowledge. Being a scientific facility it experiences constant improvements. One of the improvements worth mentioning is the DQM system for the SND detector.

First, information is collected automatically by DQM scripts and then could be corrected/confirmed by the detector operators and experts. It's available in the SND information system via web interface. The DQM system takes into account multiple parameters (numbers, histograms, user opinions) to form a data quality decision (good, bad, etc). It supports several parameters sets to be customized for different users: e.g. a general set of parameters for operators, specific detector subsystem parameters sets for experts, extended sets for further offline data processing, etc.

The input of the DQM system is hundreds of numeric parameters and histograms for each experiment run (lasts for about 30m.-2h.) produced by the SND software. The DQM system stores its data (configuration and cached scripts results) in an MySQL database. Scripts are executed when a user requests a run information and there is no cached results. A DQM script (a ROOT macro) could be written by detector subsystem expert. It takes run-related values and yields some decisions using simple C++ interface. A number or a histogram is accompanied by its reference value in the web interface for non-expert to make a reasonable decision.

The system put into production. It's consistent as an infrastructure however there are some scripts to implement and the previous experiment runs to leave opinions about. So the next move is physicists'.

Poster Session / 477

Setting up technical requirements for an astroparticle data life cycle

Authors: Victoria Tokareva[1]; Andreas Haungs[None]; Dmitriy Kostunin[2]

[1] *Karlsruhe Institute of Technology (KIT)*
[2] *KIT*

Corresponding Authors: andreas.haungs@kit.edu, victoria.tokareva@kit.edu, dmitriy.kostunin@kit.edu

Increasing data rates opens up new opportunities for astroparticle physics by improving the precision of data analysis and by deploying advanced analysis techniques that demand relatively large data volumes, e.g. deep learning . One of the ways to increase statistics is to combine data from different experimental setups for joint analysis. Moreover, such data integration provides us with an opportunity to carry out multi-messenger astroparticle studies and to search for hidden patterns within the data.

A data life cycle (DLC), namely the data processing pipeline, that is clearly defined and maximally automated at each step, from receiving data to obtaining final results of the analysis, allows us to facilitate and speed up data processing and to make the calculations more reliable and reproducible at each stage. The German-Russian Astroparticle Data Life Cycle Initiative (GRADLCI) is aimed to develop such DLC for combined analysis of data from the KASCADE-Grande (Karlsruhe, Germany) and Tunka-133 (Tunka valley, Russia) experiments. The important features of astroparticle DLC include scalability for handling large amounts of data, heterogeneous data integration, and exploiting parallel and distributed computing at every possible stage of the data processing. This demands special technical requirements necessary to perform the analysis in the hardware and software environment of a computing cluster.

In the talk we discuss the plans on DLC organization that are being implemented in the GRADLCI. This include accelerating the KASCADE-Grande database by the use of end-to-end indexing; developing a database for TAIGA/Tunka-133; organizing fast search within distributed data with a help of a proxy server and metadata database; developing access infrastructure and interfaces for scientists and general public to interact with data; performing a joint data analysis; configuring DLC for safe handling client analysis requests on the server side; and maximum automation of distributed run-wise simulations. The talk addresses the choice of a distributed data storage system, and of virtualization tools.

Track 2: Data Analysis - Algorithms and Tools / 479

Excursion Set Estimation using Sequential Entropy Reduction for Efficient Searches for New Physics at the LHC

Authors: Lukas Alexander Heinrich[1]; Gilles Louppe[1]; Kyle Stuart Cranmer[1]

[1] *New York University (US)*

Corresponding Authors: kyle.cranmer@cern.ch, g.louppe@cern.ch, lukas.heinrich@cern.ch

A common goal in the search for new physics is the determination of sets of New Physics models, typically parametrized by a number of parameters such as masses or couplings, that are either compatible with the observed data or excluded by it, where the determination into which category a given model belong requires expensive computation of the expected signal. This problem may be abstracted into the generalized problem of finding excursion sets (or, equivalently, iso-surfaces) of scalar, multivariate functions in n dimensions.

We present an iterative algorithm for choosing points within the problem domain for which the functions are evaluated in order to estimate such sets at a significantly lower computational cost. The algorithm implements a Bayesian Optimization procedure, in which a information-based acquisition function seeks to maximally reduce the uncertainty on a excursion set. Further extension of the basic algorithm to the simultaneous estimation of excursion sets of multiple functions as well as batched selection of multiple points is presented.

Finally, a python package, excursion1, is presented, which implements the algorithm and performance benchmarks are presented comparing this active-learning approach to other strategies commonly used in the high energy physics context, such as random sampling and grid searches.

1 https://github.com/diana-hep/excursion

Track 1: Computing Technology for Physics Research / 481

Federated data storage evolution in HENP: data lakes and beyond

Authors: Xavier Espinal[1]; Andrey Kirianov[2]; Jaroslava Schovancova[1]; Andrey Zarochentsev[3]

[1] *CERN*

[2] *Petersburg Nuclear Physics Institut (RU)*

[3] *St Petersburg State University (RU)*

Corresponding Authors: xavier.espinal@cern.ch, andrey.zarochencev@cern.ch, jaroslava.schovancova@cern.ch, andrey.kiryanov@cern.ch

Storage have been identified as the main challenge for the future distributed computing infrastructures: Particle Physics (HL-LHC, DUNE, Belle-II), Astrophysics and Cosmology (SKA, LSST). In particular, the High Luminosity LHC (HL-LHC) will begin operations in the year of 2026 with expected data volumes to increase by at least an order of magnitude as compared with the present systems. Extrapolating from existing trends in disk and tape pricing, and assuming flat infrastructure budgets, the implications for data handling for end-user analysis are significant. HENP experiments need to manage data across a variety of mediums based on the types of data and its uses: from tapes (cold storage) to disks and solid state drives (hot storage) to caches (including world wide access data in clouds and "data lakes"). DataLake R&D project aims at exploring an evolution of distributed storage while bearing in mind very high demands of HL-LHC era. Its primary objective is to optimize hardware usage and operational costs of a storage system deployed across distributed centers connected by fat networks and operated as a single service. Such storage would host a large fraction of the data and optimize the cost, eliminating inefficiencies due to fragmentation. In this talk we will highlight current status of the project, its achievements, interconnection with other research activities in this field like WLCG-DOMA and ATLAS-Google DataOcean, and future plans.

Track 1: Computing Technology for Physics Research / 482

ROOT Based Analysis: From Data to Plots, Expressive, Easy and Fast

Authors: Danilo Piparo[1]; Enric Tejedor Saavedra[1]; Enrico Guiraud[2]; Guilherme Amadio[1]; Javier Cervantes Villanueva[1]; Stefan Wunsch[3]

[1] *CERN*

[2] *CERN, University of Oldenburg (DE)*

[3] *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Authors: lorenzo.moneta@cern.ch, amadio@cern.ch, javier.cervantes.villanueva@cern.ch, danilo.piparo@cern.ch, stefan.wunsch@cern.ch, enrico.guiraud@cern.ch, enric.tejedor.saavedra@cern.ch

During the past two years ROOT's analysis tools underwent a major renovation, embracing a declarative approach.

This contribution explores the most recent developments of the implementation of such approach, some real-life examples from LHC experiments as well as present and future R&D lines.

After an introduction of the tool offering access to declarative analysis, RDataFrame, the newly introduced syntax for the treatment of collections

is described together with examples concerning the analysis of Open Datasets.
The tooling for visualising and studying computation graphs built with
RDataFrame in the form of diagrams is then presented.

Example real-life analyses based on RDataFrame from collider and non-collider
experiments are then discussed from the programming model and performance
perspective.

Finally, the status of existing R&D lines as well as future direction is
discussed, most notably the integration of RDataFrame with big data technologies
to distribute interactive calculations on massive computing resources.

Track 2: Data Analysis - Algorithms and Tools / 483

Weak signal extraction using matrix decomposition

Author: steven prohira[1]

[1] *The Ohio State University*

Corresponding Author: prohira.1@osu.edu

In radio-based physics experiments, sensitive analysis techniques are often required to extract signals at or below the level of noise. For a recent experiment at the SLAC National Accelerator Laboratory to test a radar-based detection scheme for high energy neutrino cascades, such a sensitive analysis was employed to dig down into a spurious background and extract a signal. This analysis employed singular-value decomposition (SVD) to decompose the data into a basis of patterns constructed from the data itself. Expansion of data in a decomposition basis allows for the extraction, or filtration, of patterns which may be unavailable to other analysis techniques. In this talk we briefly present the results of this analysis in the context of experiment T-576 at SLAC, and detail the analysis method which was used to extract a hint of a radar signal at a significance of 2.3σ.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 484

Three loop QCD corrections to heavy quark form factors

Author: Narayan Rana[1]

Co-authors: Jakob Ablinger ; Johannes Bluemlein [2]; Peter Marquard [2]; Carsten Schneider [3]

[1] *INFN Milan*

[2] *DESY*

[3] *RISC, JKU*

Corresponding Authors: carsten.schneider@risc.jku.at, peter.marquard@desy.de, johannes.bluemlein@desy.de, narayan.rana@mi.infn.it

We present an algorithm which allows to solve analytically linear systems of differential equations which factorize to first order. The solution is given in terms of iterated integrals over an alphabet where its structure is implied by the coefficient matrix of the differential equations. These systems appear in a large variety of higher order calculations in perturbative Quantum Field Theories. We apply this method to calculate the master integrals of the three-loop massive form factors for different currents, as an illustration, and present the results for all the form factors in detail. Here the solution space emerging is given by the cyclotomic harmonic polylogarithms and their associated special constants. No special basis representation of the master integrals is needed. The algorithm can be applied as well to more general cases factorizing at first order, which are based on more general alphabets, iterated integrals and associated constants.

Poster Session / 485

Explore cloud solutions for ATLAS with $250,000 AWS cloud credits

Authors: Yongsheng Gao[1]; Joshua Moss[1]; Harinder Singh Bawa[1]; Kathryn Grimm[1]

[1] *California State University (US)*

Corresponding Authors: yongsheng.gao@cern.ch, harinder.singh.bawa@cern.ch, kathryn.grimm@cern.ch, joshua.moss@cern.ch

Maintaining the huge computing grid facilities for LHC experiments and replacing their hardware every few years has been very expensive. The California State University (CSU) ATLAS group just received $250,000 AWS cloud credit from the CSU Chancellor's Office to build the first virtual US ATLAS Tier 3 to explore cloud solutions for ATLAS. We will use this award to set up full ATLAS computing environments on the cloud for ATLAS physics analysis frame works, MC generation, simulation and production. We will also develop policies for ATLAS members to submit jobs to the cloud and develop an economic model focused especially on the cost effectiveness of cloud solutions for ATLAS through extensive real user experience. The results will help ATLAS computing and physics communities decide future directions with incoming LHC upgrades.

Track 3: Computations in Theoretical Physics: Techniques and Methods / 486

Data-driven Monte Carlo Generator for Low Energy e+e- Annihilation to Hadrons

Authors: Simon Eidelman[1]; Alexandr Korobov[2]

[1] *Budker Institute of Nuclear Physics and Novosibirsk State University*

[2] *BINP,NSU*

Corresponding Author: korobovfech@gmail.com

The Monte Carlo generator to simulate events of single-photon annihilation to hadrons at center-of-mass energies below 2.5 GeV is described. The generator is based on existing data on cross sections of various exclusive channels of e+e- annihilation obtained in various e+e- experiments by the scan and ISR methods. It is extensively used in the software packages for analysis of experiments at Novosibirsk colliders VEPP-2000 and VEPP-4 aimed at high-precision measurements of hadronic cross sections for calculations of the hadronic vacuum polarization for the muon anomaly problem.

Poster Session / 487

Next Generation of HEP CPU Benchmarks

Author: Domenico Giordano[1]

[1] *CERN*

Corresponding Author: domenico.giordano@cern.ch

HEPSPEC-06 is a decade old suite used to benchmark CPU resources for WLCG.
Its adoption spans from the hardware vendors, to the site managers, funding agencies and software experts.
It is stable, reproducible, accurate, however it is reaching the end of its life.
Initial hints of lack of correlations with the HEP applications have been collected.
Looking for suitable alternatives the HEPiX Benchmarking Working Group has evaluated SPEC CPU 2017 and a number of fast benchmarks.
The studies done so far do not show major advantage in adopting SPEC CPU 2017 respect to HS06.

A suite based on the workloads that HEP experiments run can be an alternative to industrial standard benchmarks.
The adoption by the experiments of modern software development techniques simplifies the ability to package, distribute and maintain a field specific benchmark suite.
The HEPiX Benchmarking Working Group is actively working to make this possible.

This report summarises the progress of the HEPiX Benchmarking Working Group in building a benchmarking suite based on HEP workloads.
Comparisons of results with SPEC CPU 2017 and HS06 will be discussed.

Plenary / 488

Computer algebra in physics research: algorithms, systems and modern developments

Author: Stanislav Poslavskii[1]

[1] *Institute for High Energy Physics of NRC Kurchatov Institute (R*

Corresponding Author: stanislav.poslavskii@cern.ch

Computer algebra is one of a key tools in modern physics research. In this talk I will give an overview of the main mathematical and programming concepts that lie in the basis of modern computer algebra tools and how they are applied for solving modern theoretical physics and some engineering problems. I will also give a sketch overview of modern computer algebra software, including general purpose systems and dedicated tools, how they compare by functionality and performance, and what are the modern trends in the development and programming of computer algebra software.

Poster Session / 489

Application on LHC High Energy Physic data analysis with IBM Quantum Computing

Authors: Sau Lan Wu[1]; Jay Chan[1]; Wen Guan[2]; Shaojun Sun[1]; Alex Zeng Wang[1]; Chen Zhou[1]; Federico Carminati[3]; Panagiotis Barkoutsos [4]; Ivano Tavernelli[5]; Stefan Woerner[6]; Christa Zoufal[7]

[1] *University of Wisconsin Madison (US)*

[2] *University of Wisconsin (US)*

[3] *CERN*

[4] *IBM Research Zurich*

[5] *IBM Zurich Research Laboratory*

[6] *IBM Research - Zurich*

[7] *IBM Research Laboratory Zurich*

Corresponding Authors: sau.lan.wu@cern.ch, ita@zurich.ibm.com, wen.guan@cern.ch, federico.carminati@cern.ch, bpa@zurich.ibm.com, chen-hsun.chan@cern.ch, ouf@zurich.ibm.com, shaojun.sun@cern.ch, chen.zhou@cern.ch, wor@zurich.ibm.com, awang89@wisc.edu

We will present our experiences and preliminary studies on LHC high energy physics data analysis with quantum simulators and IBM quantum computer hardware using IBM Qiskit. The performance is compared with the results using a classical machine learning method applied to a physics process in Higgs-coupling-to–two-top-quarks as an example. This work is a collaboration between University of Wisconsin-Madison, CERN openlab and IBM.

Plenary / **490**

Machine learning for an X-ray FEL

Author: Daniel Ratner[1]

[1] *Stanford University*

Corresponding Author: dratner@stanford.edu

X-ray Free Electron Lasers (XFELs) are among the most complex accelerator projects in the world today. With large parameter spaces, sensitive dependence on beam quality, huge data rates, and challenging machine protection, there are diverse opportunities to apply machine learning (ML) to XFEL operation. This talk will summarize promising ML methods and highlight recent examples of successful applications at the Linac Coherent Light Source (LCLS).

Plenary / **491**

Welcome to ACAT 2019 in Saas Fee!

Corresponding Author: federico.carminati@cern.ch

Plenary / **492**

Denis Perret-Gallix (1949-2018)

Corresponding Author: werlen@lapth.cnrs.fr

Plenary / **493**

The STEM paradox: Factors affecting diversity in STEM fields

Corresponding Author: cherneyi@merrimack.edu

Women obtain more than half of U.S. undergraduate degrees in biology, chemistry, and mathematics, yet they earn less than 20% of computer science, engineering, and physics undergraduate degrees

(NSF, 2014). Why are women represented in some STEM fields more than others? The STEM Paradox and the Gender Equality Paradox show that countries with greater gender equality have a lower percentage of female STEM graduates. This phenomenon as well as other factors explaining gender disparities in STEM participation will be discussed.

Plenary / 494

Round table discussion

Plenary / 495

TrackML : a tracking Machine Learning challenge

Authors: Andreas Salzburger[1]; Andrey Ustyuzhanin[2]; Cecile Germain[3]; David Rousseau[4]; Edward Moyse[5]; Heather Gray[6]; Isabelle Guyon[None]; Jean-Roch Vlimant[7]; Laurent Roger Igor Basara[8]; Mikhail Hushchyn[2]; Moritz Kiehn[9]; Paolo Calafiura[10]; Sabrina Amrouche[11]; Steven Andrew Farrell[12]; Tobias Golling[9]; Victor Estrade[13]; Vincenzo Innocente[1]; Vladimir Gligorov[14]; Yetkin Yilmaz[15]

[1] *CERN*

[2] *Yandex School of Data Analysis (RU)*

[3] *Universite Paris Sud*

[4] *LAL-Orsay, FR*

[5] *University of Massachusetts (US)*

[6] *LBNL*

[7] *California Institute of Technology (US)*

[8] *Universita degli Studi di Trento è INFN (IT)*

[9] *Universite de Geneve (CH)*

[10] *Lawrence Berkeley National Lab. (US)*

[11] *Université de Geneve (CH)*

[12] *University of California-Unknown-Unknown*

[13] *LRI*

[14] *University of Glasgow*

[15] *LAL-Orsay (FR)*

Corresponding Authors: rousseau@lal.in2p3.fr, vincenzo.innocente@cern.ch, guyon@clopinet.com, edward.moyse@cern.ch, v.gligorov@physics.gla.ac.uk, msmk@cern.ch, andrey.ustyuzhanin@cern.ch, pcalafiura@lbl.gov, cecile.germain@lri.fr, salzburg@cern.ch, estrade@lri.fr, mikhail.hushchyn@cern.ch, tobias.golling@unige.ch, laurent.basara@cern.ch, heather.gray@cern.ch, c.amrouche@cern.ch, yetkin.yilmaz@cern.ch, steven.andrew.farrell@cern.ch, jean-roch.vlimant@cern.ch

The HL-LHC will see ATLAS and CMS see proton bunch collisions reaching track multiplicity up to 10.000 charged tracks per event. Algorithms need to be developed to harness the increased combinatorial complexity. To engage the Computer Science community to contribute new ideas, we organize a Tracking Machine Learning challenge (TrackML). Participants are provided events with 100k 3D points, and are asked to group the points into tracks; they are also given a 100GB training dataset including the ground truth. The challenge is run in two phases. The first "Accuracy" phase has run on Kaggle platform from May to August 2018; algorithms were judged judged only on a score related the fraction of correctly assigned hits. The second "Throughput" phase runs Sep 2018 to March 2019 on Codalab, will require code submission; algorithms are there ranked by combining accuracy and speed. The first phase has seen 653 participants, with top performers with innovative approaches. The second phase will finish at the time of ACAT. The talk will report on the first lessons from the challenge.

Plenary / 496

Round table: Machine Learning

Plenary / 497

Best Track 1 poster: Physics and computing performance of reconstruction algorithms for the GPU High Level Trigger 1 of LHCb

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: dorothea.vom.bruch@cern.ch, kmueller@physik.uzh.ch

Beginning in 2021, the upgraded LHCb experiment will use a triggerless readout system collecting data at an event rate of 30 MHz. A software-only High Level Trigger will enable unprecedented flexibility for trigger selections. During the first stage (HLT1), a sub-set of the full offline track reconstruction for charged particles is run to select particles of interest based on single or two-track selections. After this first stage, the event rate is reduced by at least a factor 30. Track reconstruction at 30 MHz represents a significant computing challenge, requiring a renovation of current algorithms and the underlying hardware. In this talk we present work based on an R&D project in the context of the LHCb Upgrade I exploring the approach of executing the full HLT1 chain on GPUs. This includes decoding the raw data, clustering of hits, pattern recognition, as well as track fitting. We will discuss the development of algorithms optimized for many-core architectures. Both the computing and physics performance of the full HLT1 chain will be presented.

Plenary / 498

Best Track 2 poster: Tracking performance for long living particles at LHCb

Author: Katharina Mueller[1]

[1] *Universitaet Zuerich (CH)*

Corresponding Authors: luis.miguel.garcia.martin@cern.ch, kmueller@physik.uzh.ch

The LHCb experiment is dedicated to the study of the c- and b-hadrons decays, including long living particles such as Ks and strange baryons (Lambda, Xi, etc…). These kind of particles are difficult to reconstruct from LHCb tracking systems since they escape the detection in the first tracker. A new method to evaluate the performance in terms of efficiency and throughput of the different tracking algorithms for long living particles have been developed. Special emphasis is laid on particles hitting only part of the tracking system of the new LHCb upgrade detector.

Plenary / 499

Track 1 summary

Corresponding Author: patricia.mendez@cern.ch

Plenary / 500

Track 2 Summary

Corresponding Author: andrey.ustyuzhanin@cern.ch

Plenary / 501

Track 3 Summary

Author: Ben Ruijl[1]

[1] *Nikhef*

Corresponding Author: benruyl@gmail.com

Plenary / 502

ACAT 2019 Summary

Corresponding Author: gwatts@uw.edu

Plenary / 503

Thank you and good bye

Corresponding Author: federico.carminati@cern.ch

Plenary / 504

Computations in Laser Particle Acceleration: heterogenous, parallel computing for physicists

Author: Bussmann Michael[1]

[1] *HZDR, Dresden*

Corresponding Author: m.bussmann@hzdr.de

Plenary / 505

Implicit cause models

Corresponding Author: rajeshr@cims.nyu.edu

Plenary / **506**

Keynote talk: Probabilistic Programming and Inference in Particle Physics

Author: Atilim Gunes Baydin[1]

[1] *University of Oxford*

Corresponding Author: gunes@robots.ox.ac.uk

We present a novel framework that enables efficient probabilistic inference in large-scale scientific models by allowing the execution of existing domain-specific simulators as probabilistic programs, resulting in highly interpretable posterior inference. Our framework is general purpose and scalable, and is based on a cross-platform probabilistic execution protocol through which an inference engine can control simulators in a language-agnostic way. We demonstrate the technique in particle physics, on a scientifically accurate simulation of the τ (tau) lepton decay, which is a key ingredient in establishing the properties of the Higgs boson. High-energy physics has a rich set of simulators based on quantum field theory and the interaction of particles in matter. We show how to use probabilistic programming to perform Bayesian inference in these existing simulator codebases directly, in particular conditioning on observable outputs from a simulated particle detector to directly produce an interpretable posterior distribution over decay pathways. Inference efficiency is achieved via inference compilation where a deep recurrent neural network is trained to parameterize proposal distributions and control the stochastic simulator in a sequential importance sampling scheme, at a fraction of the computational cost of Markov chain Monte Carlo sampling.

Plenary / **507**

HEP.QPR Project: Quantum Pattern Recognition for Charged Particle Tracking

Corresponding Author: heather.gray@cern.ch

Plenary / **508**

RISC-V based, silicon proven open-source processors from the PULP project

Author: Frank K. Gürkaynak[1]

[1] *ETH*

Corresponding Author: kgf@ee.ethz.ch

Since 2013, ETH Zürich and University of Bologna have been working on the PULP project to develop energy efficient computing architectures suitable for a wide range of applications starting from the IoT domain where computations have to be done in a few milliWatts, all the way to the HPC domain where the goal is to extract the maximum number of calculations within a given power budget. For this project, we have adopted an open source approach. Our main computation cores are based on the open RISC-V ISA, and we have developed highly optimized 32bit and 64bit RISC-V cores. Together with a rich set of peripherals, we have released a series of open source computing platforms from single-core microcontroller, to multi-cluster systems with tens of cores. So far we have designed and tested nearly 30 ASICs as part of the PULP project and our open source offering has been used by many companies including Google, IBM and NXP. In this talk, I will give an overview of the PULP project and show what we are currently working on.

Plenary / 509

Making HPC friendly for scientists

Corresponding Author: alam@cscs.ch

Plenary / 510

Generative Models

Corresponding Author: danilor@google.com

Plenary / 511

GPU usage in theoretical calculations

Corresponding Author: s.jones@cern.ch

Plenary / 512

Keynote talk: The future of HEP computing

Author: Rene Brun[1]

[1] *CERN*

Corresponding Author: rene.brun@cern.ch

Plenary / 513

Computing Architectures in the HL-LHC era

Corresponding Author: pcalafiura@lbl.gov

Plenary / 514

Decoding the nature of Dark matter at current and future experiments

Author: Alexander Belyaev[1]

[1] *Southampton University*

Corresponding Author: belyaev1967@gmail.com

Plenary / **515**

Neuroscience and the Future of Computing

Corresponding Author: felix.schuermann@epfl.ch

Modern electronic general-purpose computing has been on an unparalleled path of exponential acceleration for more than 7 decades. From the 1970 onwards, this trend was driven by the success of integrated circuits based on silicon technology. The exponential growth has become a self-fulfilling (and economically driven) prophecy commonly referred to as Moore's Law. The end of Moore's law has been augured many times before, but now the economic equation fueling Moore's law is increasingly broken leading to actual technology delays. Ground rule scaling of the underlying technology is expected to saturate in less than 10 years. If computational performance needs to keep increasing beyond this horizon, alternative sources of advancements will have to be found. We will have to rely much more than before on software innovations, specialized chips and ultimately new computing paradigms. This talk will cover these challenges and will discuss which role neuroscience may play in the search for novel computing paradigms, in particular neuromorphic computing.

Plenary / **516**

Keynote talk: Constraining effective field theories with machine learning

Corresponding Author: johann.brehmer@nyu.edu

An important part of the LHC legacy will be precise limits on indirect effects of new physics, framed for instance in terms of an effective field theory. These measurements often involve many theory parameters and observables, which makes them challenging for traditional analysis methods. We discuss the underlying problem of "likelihood-free" inference and present powerful new analysis techniques that combine physics insights, statistical methods, and the power of machine learning. We have developed MadMiner, a new Python package that makes it straightforward to apply these techniques. In example LHC problems we show that the new approach lets us put stronger constraints on theory parameters than established methods, demonstrating its potential to improve the new physics reach of the LHC legacy measurements. While we present techniques optimized for particle physics, the likelihood-free inference formulation is much more general, and these ideas are part of a broader movement that is changing scientific inference in fields as diverse as cosmology, genetics, and epidemiology.

Plenary / **517**

Beyond the Roadmap: HL-LHC HEP Software

Author: Gordon Watts[1]

[1] *University of Seattle*

Corresponding Author: gwatts@uw.edu

Abstract:
The HEP software ecosystem faces new challenges in 2020 with the approach of the High Luminosity LHC (HL-LHC) and the turn-on of a number of large new experiments. Current software development is organized around the experiments: No other field has attained this level of self-organization and collaboration in software development.

During 2017 the community produced a roadmap for the software R&D needed to address the software challenges of the 2020's, with a focus on the HL-LHC. Members of the community that produced over 20 papers and strategic roadmaps included individual researchers, the US and European labs, LHC-based experiments, non-collider experiments, and members of industry. The field's organization was apparent during this process.

HEP must now build on its past successes to address the challenges in front of it. Both technical solutions and new models of collaboration will be required. Funding agencies are ready to make new investments, such as the recent US NSF funded "Institute in Software in High Energy Physics (IRIS-HEP)" and recent calls for funding in Europe. In this talk, I will discuss the process by which we have gotten to this point and possible way forward.

Plenary / **518**

Automatic Differentiation and Deep Learning

Corresponding Author: soumith@gmail.com

Plenary / **519**

Higher-order QED contributions to the lepton anomalous magnetic moments

Corresponding Author: nio@riken.jp

The anomalous magnetic moment of the electron a_e and that of the muon a_μ occupy the special positions for precision tests of the Standard Model of elementary particles. Both have been precisely measured, 0.24 ppb for a_e and 0.5 ppm for a_μ, and new experiments of both a_e and a_μ are on-going aiming to reduce the uncertainties. Theoretical calculations of a_e and a_μ starting from the Lagrangian of the Standard model can also be achieved to the same precision of the future measurements. However, to do so, we need to carry out the five-loop QED calculation without any approximation. I will overview the computation method invented by T. Kinoshita in 1960's that enables us to numerically calculate the entire five-loop QED contribution to the lepton anomalous magnetic moment. I also discuss the current status of the precision tests of the lepton anomalies and the fine-structure constant α.

Plenary / **520**

Machine learning in cosmology

Author: Francois Lanusse[1]

[1] *Lawrence Berkeley National Laboratory*

Corresponding Author: flanusse@andrew.cmu.edu

Plenary / **521**

Computational challenges for Generators

Corresponding Author: a.g.buckley@gmail.com

Plenary / **522**

Machine Learning in Neutrino Physics

Author: Kazuhiro Terao[1]

[1] *SLAC*

Corresponding Author: kterao@stanford.edu