


ACAT 2014

Monday, 1 September 2014 - Friday, 5 September 2014

Faculty of Civil Engineering



Book of Abstracts

Contents

The LHCb trigger and its upgrade	1
Massive Affordable Computing Using ARM Processors in High Energy Physics	1
Presentation of Redberry: a computer algebra system designed for tensor manipulation	1
A model independent search for new phenomena with the ATLAS detector in pp collisions at $\sqrt{s}=8\text{TeV}$	2
The Massive Affordable Computing Project: Prototyping of a High Data Throughput Processing Unit	3
Densities mixture unfolding for data obtained from detectors with finite resolution and limited acceptance	3
Adaptative track scheduling to optimize concurrency and vectorization in GeantV	3
GoSam-2.0: a tool for automated one-loop calculations within and beyond the Standard Model	4
Numerical multi-loop calculations with SecDec	4
Modern Particle Physics event generation with WHIZARD	5
The Matrix Element Method within CMS	5
B-tagging at High Level Trigger in CMS	6
Mathematica and Fortran programs for various analytic QCD couplings	6
GENFIT - a Generic Track-Fitting Toolkit	7
The TileCal Online Energy Estimation for Next LHC Operation Period	8
Quality Factor for the Hadronic Calorimeter in High Luminosity Conditions	8
Computations and generation of elements on the Hopf algebra of Feynman graphs	9
The Error Reporting in the ATLAS TDAQ system	9
ATLAS FTK challenge: simulation of a billion-fold hardware parallelism	10
Data-flow performance optimization on unreliable networks: the ATLAS data-acquisition case	11
Evolution of the ATLAS Trigger and Data Acquisition System	12

Intelligent operations of the Data Acquisition system of the ATLAS Experiment at the LHC	13
Performance and development for the Inner Detector Trigger algorithms at ATLAS . . .	13
EOS : Current status and latest evolutions.	14
Delphes 3: A modular framework for fast simulation of a generic collider experiment . .	15
Clustering analysis for muon tomography data elaboration in the Muon Portal project .	15
Parallel Cuba	16
Developments in the ATLAS Tracking Software ahead of LHC Run 2	16
Challenges of the ATLAS Monte Carlo Production during Run-I and Beyond	16
Testing and improving performance of the next generation ATLAS PanDA monitoring .	17
An automated framework for hierarchical reconstruction of B mesons at the Belle II exper- iment	17
Six-loop calculations of the critical exponents in the ϕ^4 theory	18
An overview of the DII-HEP OpenStack based CMS Data Analysis	18
A development of an accelerator board dedicated for multi-precision arithmetic operations and its application to Feynman loop integrals	19
From DIRAC towards an Open Source Distributed Data Processing Solution	20
HistFitter: a flexible framework for statistical data analysis	21
The Long Term Data Preservation (LTDP) project at INFN CNAF: CDF user case.	22
Statistical methods for cosmic ray composition analysis at the Telescope Array Observatory	22
The INFN CNAF Tier1 GEMSS Mass Storage System and Database facility activity	23
Integration of PanDA workload management system with Titan supercomputer at OLCF	23
A self-configuring control system for storage and computing departments at INFN CNAF Tier1.	24
Extracting Rigorous Conclusions from Model/Data Comparisons	25
Differential reduction of generalized hypergeometric functions in application to Feynman diagrams.	25
Identifying the Higgs boson with a Quantum Computer	26
Multilevel Workflow System in the ATLAS Experiment	26
Monitoring of IaaS and scientific applications on the Cloud using the Elasticsearch ecosys- tem	27
Analyzing data flows of WLCG jobs at batch job level	27

Tier 3 batch system data locality via managed caches	28
The ALICE analysis train system	28
New exclusive CHIPS-TPT algorithms for simulation of neutron-nuclear reactions	29
HERAFitter - an open source QCD fit framework	30
AliEn File Access Monitoring Service (FAMoS)	30
Gaudi Components for Concurrency: Concurrency for Existing and Future Experiments	30
The Run 2 ATLAS Analysis Event Data Model	32
Geant4 developments in reproducibility, multi-threading and physics	33
Toolbox for multiloop Feynman diagram calculations using R* operation	34
Scalable cloud without dedicated storage	34
A Neural Network z-Vertex Trigger for Belle II	35
EicRoot: yet another FairRoot framework clone	35
Towards a generic high performance geometry library for particle transport simulation .	36
Using Functional Languages and Declarative Programming to analyze ROOT data: LINQ- toROOT	37
Distributed job scheduling in MetaCentrum	37
Belle II distributed computing	38
VISPA: Direct access and execution of data analyses for collaborations	38
A new generator for the Drell-Yan process	39
New features of AMBRE	39
Data Recommender System for the Production and Distributed Analysis System «PanDA»	39
The Long-Baseline Neutrino Experiment Computing Model and its evolution	40
Automatic numerical integration methods for Feynman integrals through 3-loop	40
Pidrix: Particle Identification Matrix Factorization	41
Planning for distributed workflows: constraint based co-scheduling of computational jobs and data placement in distributed environments.	42
The Linear Collider Software Framework	42
Event simulation for colliders - A basic overview	43
High-speed zero-copy data transfer for DAQ applications	43

A novel robust and efficient algorithm for charge particle tracking in high background flux.	44
Designing and recasting LHC analyses with MadAnalysis 5	45
Going Visual	45
Native Language Integrated Queries with CppLINQ in C++	46
Modernising ROOT: Building Blocks for Vectorised Calculations	46
Recent Developments in the CVMFS Server Backend	46
Clad - Automatic Differentiation Using Cling/Clang/LLVM	47
Techniques and tools for measuring energy efficiency of scientific software applications	47
Heterogeneous High Throughput Scientific Computing with ARMv8 64-bit and Xeon Phi	48
Implementation of a multi-threaded framework for large-scale scientific applications . .	49
Simulation Upgrades for the CMS experiment	49
WLCG Tier-2 site in Prague: a little bit of history, current status and future perspectives	50
STAR Online Framework: from Metadata Collection to Event Analysis and System Control	50
ROOT 6	51
Particle-level pileup subtraction for jets and jet shapes	51
Traditional Tracking with Kalman Filter on Parallel Architectures	51
Combination of multivariate discrimination methods in the measurement of the inclusive top pair production cross section	52
High-resolution deconvolution methods for analysis of low amplitude noisy gamma-ray spectra	52
Evolution of the ATLAS Software Framework towards Concurrency	53
Deep learning neural networks - a TMVA perspective	54
Fast detector simulation and the GeantV project	54
Welcome	55
Czech Technical University	55
Prague municipality	55
Charles University	55
Nuclear Physics Institute AS CR	55
Institute of Physics AS CR	55

Czech Science and education, financing, physics	56
Logistics	56
CMS Software Documentation System	56
THE POTENTIAL OF GENOMIC MEDICINE—THE MANY ASPECTS OF COMPUTING ARE DEMANDING AND THE COMPLEXITY MUST BE MANAGED.	56
General formulation of the sector improved residue subtraction	57
The interrelation between stability of Vacuum (absence of new physics up to Plank scale) and value of running mass of top-quark: problems and perspectives.	57
Modern messaging solutions for distributed applications	58
NLO calculations for high multiplicity processes	59
High performance computing based on mobile processors	59
Memory Models in HPC: Successes and Problems	59
Cosmic ray propagation with CRPropa 3	60
A Survey on Distributed File System Technology	60
Multivariate Data Analysis in HEP. Successes, challenges and future outlook.	61
Summary for Track 1 - Computing Technology for Physics Research	61
Summary for Track 2 - Data Analysis - Algorithms and Tools	61
Summary for Track 3 - Computations in Theoretical Physics: Techniques and Methods - part I	62
ACAT 2014 summary	62
Dark matter review: recent astronomical and particle physics results	62
Future trends and challenges in scientific computing	62
Next Generation Workload Management System for Big Data on Heterogeneous Distributed Computing	63
Poster lightning session	64
Conference closing	64
Data Processing at the Pierre Auger Observatory	64
Statistical tools for the Higgs discovery - the status and future of collaborative statistical model building	64
Two-loop matching for Standard Model RGE analysis	65
Veřejná přednáška: Petabytová astronomie (in Czech)	65

The generalized BLM approach of fixing the scales in the multiloopQCD expressions for the physical quantities: the current status of investigations	65
Direct water cooling vs free air cooling	66
On calculations within the nonlinear sigma model	66
Concert logistics	66
Poster lightning: VISPA: Direct access and execution of data analyses for collaborations	66
Poster lightning: EOS : Current status and latest evolutions.	67
Poster lightning: Using Functional Languages and Declarative Programming to analyze ROOT data: LINQtoROOT	67
Poster lightning: Analyzing data flows of WLCG jobs at batch job level	67
Poster lightning: The Linear Collider Software Framework	67
Poster lightning: Designing and recasting LHC analyses with MadAnalysis 5	67
Federico Carminati	67
Denis Perret-Gallix	68
Fons Rademakers	68
Gordon Watts	68
David Fellingner	68
Dinner logistics	68
Summary for Track 3 - Computations in Theoretical Physics: Techniques and Methods - part II	68

Computing Technology for Physics Research / 1**The LHCb trigger and its upgrade****Authors:** Gerhard Raven¹; Johannes Albrecht²; Vladimir Gligorov³¹ *NIKHEF (NL)*² *Technische Universitaet Dortmund (DE)*³ *CERN***Corresponding Authors:** gerhard.raven@nikhef.nl, vladimir.gligorov@cern.ch

The current LHCb trigger system consists of a hardware level, which reduces the LHC inelastic collision rate of 30 MHz to 1 MHz, at which the entire detector is read out. In a second level, implemented in a farm of 20k parallel-processing CPUs, the event rate is reduced to about 5 kHz. We review the performance of the LHCb trigger system, focusing on the High Level Trigger, during Run I of the LHC. Special attention is given to the use of multivariate analyses in the High Level Trigger and their importance in controlling the output rate. We demonstrate that despite its excellent performance to date, the major bottleneck in LHCb's trigger efficiencies for hadronic heavy flavour decays is the hardware trigger. The LHCb experiment plans a major upgrade of the detector and DAQ system in the LHC shutdown of 2018. In this upgrade, a purely software based trigger system is being developed, which will have to process the full 30 MHz of inelastic collisions delivered by the LHC. We demonstrate that the planned architecture will be able to meet this challenge, particularly in the context of running stability and long term reproducibility of the trigger decisions. We discuss the use of disk space in the trigger farm to buffer events while performing run-by-run detector calibrations, and the way this real time calibration and subsequent full event reconstruction will allow LHCb to deploy offline quality multivariate selections from the earliest stages of the trigger system. We discuss the cost-effectiveness of such a software-based approach with respect to alternatives relying on custom electronics. We discuss the particular importance of multivariate selections in the context of a signal-dominated production environment, and report the expected efficiencies and signal yields per unit luminosity in several key physics benchmarks the LHCb upgrade.

Poster session - Board: 101 / 2**Massive Affordable Computing Using ARM Processors in High Energy Physics****Author:** Joshua Wyatt Smith¹¹ *University of Cape Town (ZA)***Corresponding Author:** joshua.wyatt.smith@cern.ch

High Performance Computing is relevant in many applications around the world, particularly high energy physics. Experiments such as ATLAS and CMS generate huge amounts of data which needs to be analyzed at server farms located on site at CERN and around the world. Apart from the initial cost of setting up an effective server farm the price to maintain them is enormous. Power consumption and cooling are some of the biggest costs. The proposed solution to reduce costs without losing performance is to utilize ARM processors found in nearly all smartphones and tablet computers. Their low power consumption and cost along with respectable processing speed makes them an ideal choice for future large scale parallel data processing centers. Benchmarks on the Cortex-A series of ARM processors including the HPL and PMBW suites will be presented. Results from the PROOF benchmarks will also be analyzed. Issues with currently available operating systems for ARM architectures will be discussed.

Computations in Theoretical Physics: Techniques and Methods / 3

Presentation of Redberry: a computer algebra system designed for tensor manipulation

Author: Stanislav Poslavsky¹

Co-author: Dmitry Bolotin²

¹ I

² IBCh RAS

Corresponding Author: stvlpos@mail.ru

Redberry is an open source computer algebra system with native support of tensorial expressions. It provides basic computer algebra tools (algebraic manipulations, substitutions, basic simplifications etc.) which are aware of specific features of indexed expressions: contractions of indices, permutational symmetries, multiple index types etc. The high energy physics package includes tools for Feynman diagrams calculation: Dirac and SU(N) algebra, Levi-Civita simplifications and tools for one-loop counterterms calculations in quantum field theory. In this presentation we give detailed overview of Redberry features: from basic manipulations with tensors to real Feynman diagrams calculation, accompanied by many examples.

Poster session - Board: 201 / 4

A model independent search for new phenomena with the ATLAS detector in pp collisions at $\sqrt{s}=8\text{TeV}$

Author: Simone Amoroso¹

¹ University of Freiburg

The data recorded by the ATLAS experiment have been thoroughly analyzed for specific signals of physics beyond the Standard Model (SM); although these searches cover a wide variety of possible event topologies, they are not exhaustive. Events produced by new interactions or new particles might still be hidden in the data. The analysis presented here extends specific searches with a model-independent approach. All event topologies involving electrons, photons, muons, jets, b-jets and missing transverse momentum are investigated in a single analysis. The SM expectation is taken from Monte Carlo simulation. For the 697 topologies with a SM expectation greater than 0.1 events three kinematic distributions sensitive to contributions from new physics are scanned for deviations from the SM prediction. A statistical search algorithm looks for the region of largest deviation between data and the SM prediction, taking into account systematic uncertainties. To quantify the compatibility of the data with the SM prediction, the distribution of p-values of the observed deviations is compared to an expectation obtained from pseudo-experiments that includes statistical and systematic uncertainties and their correlation between search channels. No significant deviation is found in data. The number and size of the observed deviations follow the Standard Model expectation obtained from the simulated pseudo-experiments

Summary:

The data recorded by the ATLAS experiment have been thoroughly analyzed for specific signals of physics beyond the Standard Model; although these searches cover a wide variety of possible event topologies, they are not exhaustive. Events produced by new interactions or new particles might still be hidden in the data. The analysis presented here extends specific searches with a model-independent approach. All event topologies involving electrons, photons, muons, jets, b-jets and missing transverse momentum are investigated in a single analysis, and about 700 final states are considered. The SM expectation is taken

from Monte Carlo simulation. No significant deviation is found in data and the number and size of the observed deviations follow the Standard Model expectation obtained from the simulated pseudo-experiments.

Computing Technology for Physics Research / 5

The Massive Affordable Computing Project: Prototyping of a High Data Throughput Processing Unit

Author: Mitchell Arij Cox¹

¹ *University of the Witwatersrand (ZA)*

Scientific experiments are becoming highly data intensive to the point where offline processing of stored data is infeasible. Data stream processing, or high data throughput computing, for future projects is required to deal with terabytes of data per second. Conventional data-centres based on typical server-grade hardware are expensive and are biased towards processing power rather than I/O bandwidth. This system imbalance can be solved with massive parallelism to increase the I/O capabilities, at the expense of excessive processing power and high energy consumption. The Massive Affordable Computing Project aims to use low-cost, ARM System on Chips to address the issue of system balance, affordability and energy efficiency. An ARM-based Processing Unit is currently in development, with a design goal of 40 Gb/s I/O throughput and significant processing power. Novel use of PCI-Express is used to address the typically limited I/O capabilities of consumer ARM System on Chips. A more detailed overview of the Project and Processing Unit will be presented along with to-date performance and I/O throughput tests.

Data Analysis - Algorithms and Tools / 6

Densities mixture unfolding for data obtained from detectors with finite resolution and limited acceptance

Author: Nikolay Gagunashvili¹

Co-author: Michael Schmelling²

¹ *University of Akureyri, Borgir, v/Nordurlod, IS-600 Akureyri, Iceland & Max-Planck-Institut für Kernphysik, P.O. Box 103980, 69029 Heidelberg, Germany*

² *Max-Planck-Institut für Kernphysik, P.O. Box 103980, 69029 Heidelberg, Germany*

Corresponding Author: nikolai@unak.is

A mixture density model-based procedure for correcting experimental data for distortions due to finite resolution and limited detector acceptance is presented. The unfolding problem is known to be an ill-posed problem that can not be solved without some a priori information about the solution such as, for example, smoothness or positivity. In the approach presented here the true distribution is estimated by a weighted sum of densities, with the variances of the densities acting as a regularization parameter responsible for the smoothness of the result. Cross-validation is used to determine the optimal value of this parameter and a bootstrap method for estimation statistical errors of unfolded distribution. Numerical examples, one of them for a steeply falling probability density function, are presented to illustrate the procedure.

Computing Technology for Physics Research / 7

Adaptative track scheduling to optimize concurrency and vectorization in GeantV

Authors: Andrei Gheata¹; Federico Carminati¹; Georgios Bitzes²; Guilherme Lima³; Johannes Christof De Fine Licht⁴; John Apostolakis¹; Laurent Duhem⁵; Marilena Bandieramonte⁶; Mihaly Novak¹; Oksana Shadura⁷; Philippe Canal⁸; Raman Sehgal⁹; Rene Brun¹; Sandro Christian Wenzel¹; Soon Yung Jun⁸; Victor Daniel Elvira⁸

¹ CERN

² University of Athens (GR)

³ FermiLab (US)

⁴ University of Copenhagen (DK)

⁵ INTEL

⁶ University of Catania and INAF

⁷ National Technical Univ. of Ukraine "Kyiv Polytechnic Institute

⁸ Fermi National Accelerator Lab. (US)

⁹ Bhabha Atomic Research Centre (IN)

Corresponding Author: andrei.gheata@cern.ch

The *GeantV* project aims to R&D new particle transport techniques to maximize parallelism on multiple levels, profiting from the use of both SIMD instructions and co-processors for the CPU-intensive calculations specific to this type of applications. In our approach, vectors of tracks belonging to multiple events and matching different locality criteria must be gathered and dispatched to algorithms having vector signatures. While the transport propagates tracks and changes their individual states, geometry locality becomes harder to maintain. The scheduling policy has to be changed to maintain efficient vectors while keeping an optimal level of concurrency. The model has complex dynamics requiring tuning the thresholds to switch between the normal regime and special modes, i.e. prioritizing events to allow flushing memory, adding new events in the transport pipeline to boost locality, dynamically adjusting the particle vector size or switching between vectorized to single track mode when vectorization causes only overhead. This work covers a comprehensive study for optimizing these parameters to make the behavior of the scheduler self-adapting, presenting the most recent results.

Computations in Theoretical Physics: Techniques and Methods / 8

GoSam-2.0: a tool for automated one-loop calculations within and beyond the Standard Model

Author: Johann Felix von Soden-Fraunhofen¹

¹ Max Planck Institute for Physics

Corresponding Author: jfsoden@mpp.mpg.de

We present the version 2.0 of the program package GoSam, which is a public program package to compute one-loop QCD and/or electroweak corrections to multi-particle processes within and beyond the Standard Model. The extended version of the Binoth-Les-Houches-Accord interface to Monte Carlo programs is also implemented. This allows a large flexibility regarding the combination of the code with various Monte Carlo programs to produce fully differential NLO results, including the possibility of parton showering and hadronisation. We illustrate the wide range of applicability of the code by showing various phenomenological results for multi-particle processes at NLO, both within and beyond the Standard Model.

Computations in Theoretical Physics: Techniques and Methods / 9**Numerical multi-loop calculations with SecDec****Author:** Gudrun Heinrich¹**Co-authors:** Johannes Schlenk²; Sophia Borowka²¹ *Max Planck Institute for Physics*² *Max Planck Institute for Physics Munich***Corresponding Author:** gudrun@mpp.mpg.de

SecDec is a program which can be used for the evaluation of parametric integrals, in particular multi-loop integrals. For a given set of propagators defining the graph, the program automatically constructs a Feynman parameter representation, extracts the singularities in the dimensional regularisation parameter and produces a Laurent series in this parameter, whose coefficients are then evaluated numerically. Threshold singularities are handled by an automated deformation of the integration contour into the complex plane. We present various new features of the program, which extend the range of applicability and increase the speed. We also present recent phenomenological examples of applications to two-loop integrals with several mass scales.

Computations in Theoretical Physics: Techniques and Methods / 10**Modern Particle Physics event generation with WHIZARD****Author:** juergen reuter¹**Co-authors:** Bijan Chokoufe²; Christian Weiss³; Fabian Bach²; Marco Sekulla³; Thorsten Ohl⁴; Wolfgang Kilian³¹ *DESY Hamburg, Germany*² *DESY*³ *University of Siegen*⁴ *University of Wuerzburg***Corresponding Author:** juergen.reuter@desy.de

We describe the multi-purpose Monte-Carlo event generator WHIZARD for the simulation of high-energy particle physics experiments. Besides the presentation of the general features of the program like SM physics, BSM physics, and QCD effects, special emphasis will be given to the support of the most accurate simulation of the collider environments at hadron colliders and especially at future linear lepton colliders. On the more technical side, the very recent code refactoring towards a completely object-oriented software package to improve maintainability, flexibility and code development will be discussed. Finally, we present ongoing work and future plans regarding higher-order corrections, more general model support including the setup to search for new physics in vector boson scattering at the LHC, as well as several lines of performance improvement.

Data Analysis - Algorithms and Tools / 11**The Matrix Element Method within CMS****Author:** Camille Beluffi¹¹ *Universite Catholique de Louvain (UCL) (BE)***Corresponding Author:** camille.beluffi@cern.ch

The Matrix Element Method (MEM) is unique among the analysis methods used in experimental particle physics because of the direct link it establishes between theory and event reconstruction. This method was used to provide the most accurate measurement of the top mass at the Tevatron and since then it was used in the discovery of electroweak production of single top quarks. The method can in principle be used for any measurement, with a large gain compared to cut-based analysis techniques for processes involving intermediate resonances. Within CMS, this method is mainly known as a cross check to test the newly discovered boson spin (MELA), and as a way to compute the main background for the Higgs production in association with a top quark (ttH). In this contribution, the MEM is presented using the example of two ways of using it through the two CMS analysis mentioned below. The advantages and limitations of this method will also be highlighted, and the latest approved results will be presented.

Poster session - Board: 202 / 12**B-tagging at High Level Trigger in CMS****Author:** Eric Chabert¹¹ *Institut Pluridisciplinaire Hubert Curien (FR)***Corresponding Author:** eric.chabert@cern.ch

The CMS experiment has been designed with a 2-level trigger system. The Level 1 Trigger is implemented on custom-designed electronics. The High Level Trigger (HLT) is a streamlined version of the CMS offline reconstruction software running on a computer farm. Using b-tagging at trigger level will play a crucial role during the Run II data taking to ensure the Top quark, beyond the Standard Model and Higgs boson physics programme of the experiment. It will help to significantly reduce the trigger output rate which will increase due to the higher instantaneous luminosity and higher cross sections at 13 TeV. B-tagging algorithms based on the identification of tracks displaced from the primary proton-proton collision or on the reconstruction of secondary vertices have been successfully used during Run I. We will present their design and performance with an emphasis on the dedicated aspects of track and primary vertex reconstruction, as well as the improvements foreseen to meet the challenges of the Run II data taking (high track multiplicity, out-of-time pile-up)

Computations in Theoretical Physics: Techniques and Methods / 13**Mathematica and Fortran programs for various analytic QCD couplings****Author:** Gorazd Cvetič¹**Co-author:** Cesar Ayala¹¹ *Santa Maria University*

Corresponding Author: gorazd.cvetec@gmail.com

Perturbative QCD in the usual mass independent schemes gives us running coupling $a(Q^2) \equiv \alpha_s(Q^2)/\pi$ which has unphysical (Landau) singularities at low squared momenta $|Q^2| < 1 \text{ GeV}^2$ (where $Q^2 \equiv -q^2$). Such singularities do not reflect correctly the analytic (holomorphic) properties of spacelike observables $calD(Q^2)$ such as current correlators or structure function sum rules, the properties dictated by the general principles of (local) quantum field theory. Therefore, evaluating $calD(Q^2)$ in perturbative QCD in terms of the coupling $a(\kappa Q^2)$ (where $\kappa \sim 1$ is the renormalization scale parameter) cannot give us correct results at low $|Q^2|$. As an alternative, analytic (holomorphic) models of QCD have been constructed in the literature, where $A_1(Q^2)$ [the holomorphic analog of the underlying perturbative $a(Q^2)$] has the desired properties. We present our programs, written in Mathematica and in Fortran, for the evaluation of the $A_\nu(Q^2)$ coupling, a holomorphic analog of the powers $a(Q^2)^\nu$ where ν is a real power index, for various versions of analytic QCD:

(A) (Fractional) Analytic Perturbation Theory ((F)APT) model of Shirkov, Solovtsov et al. (extended by Bakulev, Mikhailov and Stefanis to noninteger ν); in this model, the discontinuity function $\rho_\nu(\sigma) \equiv \text{Im}A_\nu(-\sigma - i\epsilon)$, defined at $\sigma > 0$, is set equal to its perturbative counterpart: $\rho_\nu(\sigma) = \text{Im}a(-\sigma - i\epsilon)^\nu$ for $\sigma > 0$, and zero for $\sigma < 0$.

(B) Two-delta analytic QCD model (2 δ anQCD) of Ayala, Contreras and Cvetec; in this model, the discontinuity function $\rho_1(\sigma) \equiv \text{Im}A_1(-\sigma - i\epsilon)$ is set equal to its perturbative counterpart for high $\sigma > M_0^2$ (where $M_0 \sim 1 \text{ GeV}$), and at low positive σ the otherwise unknown behavior of $\rho_1(\sigma)$ is parametrized as a linear combination of two delta functions.

(C) The massive QCD of Shirkov, where $A_1(Q^2) = a(Q^2 + M^2)$ with $M \sim 1 \text{ GeV}$.

Summary:

We present programs, in Mathematica and in Fortran, for calculation of the general power analogs of the coupling in three different analytic (holomorphic) models of QCD.

Data Analysis - Algorithms and Tools / 14

GENFIT - a Generic Track-Fitting Toolkit

Authors: Johannes Rauch¹; Tobias Schlueter²

¹ T

² Theoretische Physik-Fakultaet fuer Physik-Ludwig-Maximilians-Uni

Corresponding Author: j.rauch@tum.de

Genfit is an experiment-independent track-fitting toolkit, which combines fitting algorithms, track representations, and measurement geometries into a modular framework. We report on a significantly improved version of Genfit, based on experience gained in the Belle II, PANDA, and FOPI experiments. Improvements concern the implementation of additional track-fitting algorithms, enhanced implementations of Kalman fitters, enhanced visualization capabilities, and additional implementations of measurement types suited for various kinds of tracking detectors. The data model has been revised, allowing for efficient track merging, smoothing, residual calculation and alignment.

Summary:

Genfit is an experiment-independent track-fitting toolkit, which combines fitting algorithms, track representations, and measurement geometries into a modular framework. We report on a significantly improved version of Genfit, based on experience gained in the Belle II, PANDA, and FOPI experiments. Improvements concern the implementation of additional track-fitting algorithms, enhanced implementations of Kalman fitters, enhanced visualization capabilities, and additional implementations of measurement types suited for various kinds of tracking detectors. The data model has been revised, allowing for efficient track merging, smoothing, residual calculation and alignment.

Poster session - Board: 203 / 15

The TileCal Online Energy Estimation for Next LHC Operation Period

Author: Bernardo Sotto-Maior Peralva on behalf of ATLAS Tile Calorimeter collaboration¹

¹ CERN

Corresponding Authors: bernardo.sotto-maior.peralva@cern.ch, henric.wilkens@cern.ch

The ATLAS Tile Calorimeter (TileCal) is the detector used in the reconstruction of hadrons, jets, muons and missing transverse energy from the proton-proton collisions at the Large Hadron Collider (LHC). It covers the central part of the ATLAS detector ($|\eta| < 1.6$). The energy deposited by the particles is read out by approximately 5,000 cells, with double readout 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, which expects a single signal with a well-defined shape. However, the LHC luminosity is expected to increase leading to signal pile-up that deforms the signal of interest. Due to limited resources, the current DSP-based hardware setup does not allow the implementation of sophisticated energy estimation methods that deal with the pile-up. Therefore, the technique to be employed for online energy estimation in TileCal for next LHC operation period must be based on fast filters such as the Matched Filter (MF) and Optimal Filter (OF).

Both the MF and OF methods envisage the use of the background second order statistics in its design, more precisely the covariance matrix. However, the identity matrix has been used to describe this quantity. Although this approximation can be valid for low luminosity LHC, it leads to biased estimators under pile-up conditions. Since most of the TileCal cell present low occupancy, the pile-up, which is often modeled by a non-Gaussian distribution, can be seen as outlier events. Consequently, the classical covariance matrix estimation does not describe correctly the second order statistics of the background for the majority of the events, as this approach is very sensitive to outliers. As a result, the MF (or OF) coefficients are miscalculated leading to a larger variance and biased energy estimator. This work evaluates the usage of a robust covariance estimator, namely the Minimum Covariance Determinant (MCD) algorithm, to be applied in the MF design. The goal of the MCD estimator is to find a number of observations whose classical covariance matrix has the lowest determinant. Hence, this procedure avoids taking into account low likelihood events to describe the background. It is worth mentioning that the background covariance matrix as well as the MF coefficients for each TileCal channel are computed offline and stored for both online and offline use.

In order to evaluate the impact of the MCD estimator on the performance of the MF, simulated data sets were used. Different average numbers of interactions per bunch crossing and bunch spacings were tested. The results show that the estimation of the background covariance matrix through MCD improves significantly the final energy resolution with respect to the identity matrix which is currently used. Particularly, for high occupancy cells, the final energy resolution is improved by more than 20%. Moreover, the use of the classical covariance matrix degrades the energy resolution for the majority of TileCal cells.

Poster session - Board: 204 / 16

Quality Factor for the Hadronic Calorimeter in High Luminosity Conditions

Author: Luiz Eduardo Balabram Filho on behalf of ATLAS Tile Calorimeter collaboration¹

¹ CERN

Corresponding Authors: seixas@lps.ufrj.br, henric.wilkens@cern.ch

The Tile Calorimeter (TileCal) is the central section of the hadronic calorimeter of ATLAS experiment and has about 10,000 electronic channels. An Optimal Filter (OF) has been used to estimate the energy

sampled by the calorimeter and applies a Quality Factor (QF) for signal acceptance. An approach using Matched Filter (MF) has also been pursued.

In order to cope with the luminosity rising foreseen for LHC operation upgrade, different algorithms have been developed. Among them, the Constrained Optimal Filter (COF) is showing good capacity in handling such luminosity rise by using a deconvolution technique, which recovers physics signals from out of time pile up. When pile up noise is low, COF switches to MF estimator for optimal performance.

Currently, the OF measure for signal acceptance is implemented through a chi-square test. At a low-luminosity scenario, such QF measure has been used as a way to describe how the acquired signal is compatible to the pulse shape pattern. However, at high-luminosity conditions, due to pile up, this QF acceptance is no longer possible when OF is employed, and the QF becomes a parameter to indicate whether the reconstructed signal suffers or not from pile up. As COF recovers the original pulse shape, the QF may be used again as signal acceptance index.

In this work, a new QF measure is introduced. It is based on divergence statistics, which measures the similarity of probability density functions. The analysis of QF measures in deconvolved pulses is performed and the chi-square measure is compared to the divergence index. Results are shown for high-luminosity Monte Carlo data.

Computations in Theoretical Physics: Techniques and Methods / 17

Computations and generation of elements on the Hopf algebra of Feynman graphs

Author: Michael Norbert Borinsky¹

¹ *Humboldt-University Berlin*

Corresponding Author: borinsky@physik.hu-berlin.de

Two programs, feyngen and feyncop, are presented. feyngen is designed to generate high loop order Feynman graphs for Yang-Mills, QED and ϕ^k theories. feyncop can compute the coproduct of these graphs on the underlying Hopf algebra of Feynman graphs.

The programs can be validated by exploiting zero dimensional field theory combinatorics and identities on the Hopf algebra which follow from the renormalizability of the theories.

Summary:

The Hopf algebra structure of Feynman graphs has been explored extensively in the last years. It proved to be valuable for the analytic computation of Feynman amplitudes by means of systematic parametric integration techniques and could lead to new non-perturbative results in the scope of Dyson-Schwinger equations. feyngen and feyncop were developed to provide input for the powerful new techniques. feyngen relies on the established nauty package by McKay for the fast generation of higher loop diagrams. feyncop can be used to calculate the coproduct of graphs on the Hopf algebra of Feynman graphs. This coproduct encodes the BPHZ algorithm necessary to evaluate the finite amplitude in a mathematically sound framework and fits well into the world of Dyson-Schwinger equations. In the framework of these equations certain identities can be obtained which were used to validate the two programs.

Poster session - Board: 102 / 19

The Error Reporting in the ATLAS TDAQ system

Author: Denis Oliveira Damazio¹

¹ *Brookhaven National Laboratory (US)*

Corresponding Authors: sergei.kolos@cern.ch, denis.oliveira.damazio@cern.ch

The ATLAS Error Reporting feature, which is used in the TDAQ environment, provides a service that

allows experts and shift crew to track and address errors relating to the data taking components and applications. This service, called the Error Reporting Service(ERS), gives software applications the opportunity to collect and send comprehensive data about errors, happening at run-time, to a place where it can be intercepted in real-time by any other system component. Other ATLAS online control and monitoring tools use the Error Reporting service as one of their main inputs to address system problems in a timely manner and to improve the quality of acquired data.

The actual destination of the error messages depends solely on the run-time environment, in which the online applications are operating. When applications send information to ERS, depending on the actual configuration the information may end up in a local file, in a database, in distributed middle-ware, which can transport it to an expert system or display it to a users, who can work around a problem. Thanks to the open framework design of ERS, new information destinations can be added at any moment without touching the reporting and receiving applications.

The ERS API is provided in three programming languages used in the ATLAS online environment: C++, Java and Python. All APIs use exceptions for error reporting but each of them exploits advanced features of a given language to simplify program writing experience. For the example, as C++ lacks language support for exceptions, a special macro have been designed to generate hierarchies of C++ exception classes at compile time. Using this approach a software developer can write a single line of code to generate a boilerplate code for a fully qualified C++ exception class declaration with arbitrary number of parameters and multiple constructors, which encapsulates all relevant static information about the given type of issues. When corresponding error occurs at run time, a program just need to create an instance of that class passing relevant values to one of the available class constructors and send this instance to ERS.

This paper presents the original design solutions exploited for the ERS implementation and describes the experience of using ERS for the first ATLAS run period, where the cross-system error reporting standardization, introduced by ERS, was one of the key points for successful launching and utilization of automated problem-solving solutions in the TDAQ online environment.

Computing Technology for Physics Research / 20

ATLAS FTK challenge: simulation of a billion-fold hardware parallelism

Author: Denis Oliveira Damazio¹

¹ *Brookhaven National Laboratory (US)*

Corresponding Authors: sasha.vanyashin@cern.ch, denis.oliveira.damazio@cern.ch

During the current LHC shutdown period the ATLAS experiment will upgrade the Trigger and Data Acquisition system to include a hardware tracker coprocessor: the Fast Tracker (FTK). The FTK accesses the 80 million of channels of the ATLAS silicon detector, identifying charged tracks and reconstructing their parameters in the entire detector at a rate of up to 100 KHz and within 100 microseconds. To achieve this performance the FTK system utilizes the computing power of a custom ASIC chip with associative memory (AM) designed to perform “pattern matching” at very high speed, and the track parameters are calculated using modern FPGAs. To control this massive system a detailed simulation has been developed with the goal of supporting the hardware design and studying the impact of such a system in the ATLAS online event selection at high LHC luminosities. The two targets, electronic design and physics performance evaluation, have different requirements: while the hardware design requires accurate emulation of a relatively small data sample, physics studies require millions of events and the efficient use of CPU is important. We present the issues related to emulating this system on a commercial CPU platform, using ATLAS computing Grid resources, and the solutions developed in order to mitigate these problems to allow the emulation to perform the studies required to support the system design, construction and installation.

Summary:

The FTK system performs the track fitting for charged tracks in p-p collisions separating the pattern matching step and the track parameters calculation in two sequential steps: the pattern matching is performed using the capability to find correlation in data using the associative memory (AM) chips, while the tracks parameters are calculated implementing a fast fit in FPGA. In both steps the high degree of parallelism and the extremely large computing available in both devices is used. This allows the reconstruction of all the tracks with transverse momentum greater than 1 GeV, in the full detector acceptance and in real time, at an event rate of up to 100 KHz within 100 microseconds per event.

The simulation of such highly parallelized system is an extremely complex task when executed using commercial computers based on CPUs. The main bottlenecks are the low bandwidth access to memory, if compared to the AM system that has I/O bandwidth of about 25 TB/s, and the lack of parallelism. In fact the AM chip uses a content addressable memory (CAM) architecture and any data inquiry is broadcasted to all memory elements simultaneously, thus data retrieval time is independent of the database size, with the possibility to perform millions of comparison per chip per second. Each incoming hit reaches all the 1 billion patterns in the whole AM system within the same clock cycle (10 ns), a very specific feature that cannot be matched by CPU based systems. Similar penalties are paid in the track fitter implementation, that in CPU cannot reach the 1 GHz rate that will be obtained in the FPGAs installed in the fit boards.

In designing and developing the FTK simulation we developed solutions to allow to run the system emulation in standard Grid working nodes, overcoming the limitation of commercial hardware in term of CPU power and memory availability. This solutions allows us to take benefit of the thousands of worker nodes available in the Grid computing facilities used by the ATLAS experiment and produce millions of physics events to complete studies on the expected performance of the system.

Poster session - Board: 103 / 21

Data-flow performance optimization on unreliable networks: the ATLAS data-acquisition case

Author: Fabrizio Salvatore¹

Co-author: Tommaso Colombo²

¹ *University of Sussex (GB)*

² *CERN and Universität Heidelberg*

Corresponding Authors: tommaso.colombo@cern.ch, p.fabrizio.salvatore@gmail.com

The ATLAS detector at CERN records proton-proton collisions delivered by the Large Hadron Collider (LHC). The ATLAS Trigger and Data-Acquisition (TDAQ) system identifies, selects, and stores interesting collision data. These are received from the detector readout electronics at an average rate of 100 kHz. The typical event data size is 1 to 2 MB. Overall, the ATLAS TDAQ can be seen as a distributed software system executed on a farm of roughly 2000 commodity PCs. The worker nodes are interconnected by an Ethernet network that at the restart of the LHC in 2015 is expected to experience a sustained throughput of several 10 GB/s.

A particular type of challenge posed by this system, and by DAQ systems in general, is the inherently bursty nature of the data traffic from the readout buffers to the worker nodes. This can cause instantaneous network congestion and therefore performance degradation. The effect is particularly pronounced for unreliable network interconnections, such as Ethernet.

In this presentation we report on the design of the data-flow software for the 2015-2018 data-taking period of the ATLAS experiment. This software will be

responsible for transporting the data across the distributed data-acquisition system. We will focus on the strategies employed to manage the network congestion and therefore minimize the data-collection latency and maximize the system performance.

We will discuss the results of systematic measurements performed on the production hardware. These results highlight the causes of network congestion and the effects on the overall system performance. Based on these results, a simulation of the distributed system communication has been developed. This enables to explore different solutions to the network congestion sources and effects, without physical intervention. These investigations will support the choice of the best data-flow control strategy for the coming data-taking period.

Poster session - Board: 104 / 22

Evolution of the ATLAS Trigger and Data Acquisition System

Author: Fabrizio Salvatore¹

¹ *University of Sussex (GB)*

Corresponding Authors: mikel.eukeni.pozo.astigarraga@cern.ch, p.fabrizio.salvatore@googlemail.com

ATLAS is a Physics experiment that collects high-energy particle collisions at the Large Hadron Collider at CERN. It uses tens of millions of electronics channels to capture the outcome of the particle bunches crossing each other every 25 ns. Since reading out and storing the complete information is not feasible (~100 TB/s), ATLAS makes use of a complex and highly distributed Trigger and Data Acquisition (TDAQ) system, in charge of selecting only interesting data and transporting those to permanent mass storage (~1 GB/s) for later analysis.

The data reduction is carried out in two stages: first, custom electronics performs an initial level of data rejection for each bunch crossing based on partial and localized information. Only data corresponding to collisions passing this stage of selection will be actually read-out from the on-detector electronics. Then, a large computer farm (~17 k cores) analyses these data in real-time and decides which ones are worth being stored for Physics analysis. A large network allows to move the data from ~1800 front-end buffers to the location where they are processed, and from there to mass storage. The overall TDAQ system is embedded in a common software framework that allows to control, configure and monitor the data taking process.

The experience gained during the first period of data taking of the ATLAS experiment (Run I, 2010-2012) has inspired a number of ideas for improvement of the TDAQ system that are being put in place during the so-called Long Shutdown 1 of the Large Hadron Collider (LHC), in 2013/14.

This paper summarizes the main changes that have been applied to the ATLAS TDAQ system and highlights the expected performance and functional improvements that will be available for the LHC Run II. Particular emphasis will be put on the evolution of the software-based data selection and of the flow of data in the system. The reasons for the modified architectural and technical choices will be explained, and details will be provided on the simulation and testing approach used to validate this system.

Poster session - Board: 105 / 23

Intelligent operations of the Data Acquisition system of the ATLAS Experiment at the LHC

Author: Fabrizio Salvatore¹

¹ University of Sussex (GB)

Corresponding Authors: giuseppe.avolio@cern.ch, p.fabrizio.salvatore@googlemail.com

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 system is composed of a large number of hardware and software components (about 3000 machines and more than 15000 concurrent processes at the end of LHC's Run 1) which in a coordinated manner provide the data-taking functionality of the overall system.

The Run Control (RC) system is

the component steering the data acquisition by starting and stopping processes and by carrying all data-taking elements through well-defined states in a coherent way (finite state machine pattern). The RC is organized as a hierarchical tree (run control tree) of run controllers following the functional de-composition into systems and sub-systems of the ATLAS detector.

Given the size and complexity of the TDAQ system, errors and failures are bound to happen and must be dealt with. The data acquisition system has to recover

from these errors promptly and effectively, possibly without the need to stop data taking operations.

In light of this crucial requirement and taking into account all the

lessons learnt during LHC's Run 1, the RC has been completely re-designed and re-implemented during the LHC Long Shutdown 1 (LS1) phase.

As a result of the new design, the RC

is assisted by the Central Hint and Information Processor (CHIP) service that can be truly considered its "brain". CHIP is an intelligent system having a global view on the TDAQ

system. It is based on a third party open source Complex Event Processing (CEP) engine, ESPER.

CHIP supervises the ATLAS data taking, takes operational decisions and handles

abnormal conditions in a remarkable efficient and reliable manner. Furthermore CHIP automates complex procedures and performs advanced recoveries.

In this paper the design,

implementation and performances of the RC/CHIP system will be described. A particular emphasis will be put on the way the RC and CHIP cooperate and on the huge benefits brought

by the CEP engine. Additionally some error recovery scenarios will be analyzed for which the intervention of human experts is now rendered unnecessary.

Poster session - Board: 106 / 24

Performance and development for the Inner Detector Trigger algorithms at ATLAS

Author: Fabrizio Salvatore¹

¹ University of Sussex (GB)

Corresponding Authors: Ondrej.penc@cern.ch, p.fabrizio.salvatore@googlemail.com

The performance of the ATLAS Inner Detector (ID) Trigger algorithms being developed for running on the ATLAS High Level Trigger (HLT) processor farm during Run 2 of the LHC are presented. During the 2013-14 LHC long shutdown modifications are being carried out to the LHC accelerator to increase both the beam energy and luminosity. These modifications will pose significant challenges for the ID Trigger

algorithms, both in terms execution time and physics performance. To meet these challenges, the ATLAS HLT software is being restructured to run as a more flexible single stage HLT, instead of two separate stages (Level2 and Event Filter) as in Run 1. This will reduce the overall data volume that needs to be requested by the HLT system, since data will no longer need to be requested for each of the two separate processing stages.

Development of the ID Trigger algorithms for Run 2, currently expected to be ready for detector commissioning near the end of 2014, is progressing well and the current efforts towards optimising the operational performance of these algorithms is discussed. The new tracking strategy employed for Run 2 will use a Fast Track Finder (FTF) algorithm to seed subsequent precision tracking, and will result in an improved track parameter resolution and faster execution times than achieved during Run 1. This will be achieved without compromising the algorithm robustness with respect to the expected increase in multiplicity of separate proton-proton interactions (pileup) per LHC bunch crossing.

The performance of the new algorithms has been evaluated using an extensive suite of profiling tools to identify those aspects where code optimisation would be most beneficial. The methods used to extract accurate timing information for each execution step are described, as well as the analysis of per-call level profiling data and the sampling of hardware counters to study the efficiency of CPU utilisation. In addition, a summary of the effective optimisation steps that have been applied to the new algorithms are discussed. The profiling infrastructure, constructed to provide prompt feedback from the optimisation, is described, including the methods used to monitor the relative performance improvements as the code evolves. This is with a view to understanding how the profiling and optimisation testing methods might be extended to other ATLAS software development.

The increased use of parallelism for HLT algorithm processing has also been explored. Possible new opportunities arising from explicit code vectorisation and the potential inclusion of co-processors to accelerate performance in key sections of the online tracking algorithms are also discussed.

Poster session - Board: 107 / 25

EOS : Current status and latest evolutions.

Author: Geoffroy Michel Adde¹

Co-authors: Andreas Joachim Peters¹; Belinda Chan Kwok Cheong¹; Dirk Duellmann¹; Elvin Alin Sindrilaru¹; Jan Iven¹; Joaquim Manuel Pereira Rocha¹; Luca Mascetti¹; Lukasz Janyst¹; Xavier Espinal Curull¹

¹ CERN

Corresponding Author: geoffray.adde@cern.ch

EOS is a distributed file system developed and used mainly at CERN. It provides low latency, high availability, strong authentication, multiple replication schemes as well as multiple access protocols and features. Deployment and operations remain simple and EOS is currently used by multiple experiments at CERN and it provides a total raw storage space of 65PB. In a first part we go through a brief overview of EOS' features and architecture and we give some operations facts. In a second part we make a focus on the new infrastructure aware file scheduler.

Data Analysis - Algorithms and Tools / 26

Delphes 3: A modular framework for fast simulation of a generic collider experiment

Author: Alexandre Jean N Mertens¹

¹ *Universite Catholique de Louvain (UCL) (BE)*

Corresponding Author: alexandre.mertens@cern.ch

Delphes is a C++ framework, performing a fast multipurpose detector response simulation. The simulation includes a tracking system, embedded into a magnetic field, calorimeters and a muon system. The framework is interfaced to standard file formats and outputs observables such as isolated leptons, missing transverse energy and collection of jets which can be used for dedicated analyses. The simulation of the detector response takes into account the effect of magnetic field, the granularity of the calorimeters and subdetector resolutions. The Delphes simulation also includes a simple event display.

Summary:

After a general description of the Delphes framework, several new features such as an emulation of the particle-flow algorithm, pile-up simulation, N-subjettiness, a simple b-tagging algorithm based on counting tracks with large impact parameter, will be discussed.

A typical use-case of the Delphes framework in the context of future collider experiments will also be presented.

Data Analysis - Algorithms and Tools / 27

Clustering analysis for muon tomography data elaboration in the Muon Portal project

Author: Marilena Bandieramonte¹

Co-authors: Alessandro Costa²; Catia Petta³; Costantino Pistagna²; Eva Sciacca²; Fabio Vitello²; Francesco Riggi³; Paola La Rocca³; Piero Massimino²; Simone Riggi²; Ugo Becciani²; Vincenzo Antonuccio De-Logu²

¹ *Dept. Of Physics and Astronomy, University of Catania and Astrophysical Observatory, Inaf Catania*

² *Astrophysical Observatory, Inaf Catania*

³ *Universita e INFN (IT)*

Clustering analysis is a set of multivariate data analysis techniques through which is possible to gather statistical data units, in order to minimize the “logical distance” within each group and to maximize the one between groups. The “logical distance” is quantified by measures of similarity/dissimilarity between defined statistical units. Clustering techniques are traditionally applied to problems like pattern recognition, image classification and color quantization. These techniques allow to infer the implicit information in the data so they are used as a data mining technique to simplify the complexity of big dataset. In this paper the authors present a novel approach to the muon tomography data analysis based on clustering algorithms. As a case study we present the Muon Portal project that aims to realize dedicated particle detector for the inspection of harbor containers to hinder the smuggling of nuclear materials. Cluster analysis successfully elaborates data in the Muon Portal project, meeting the need to make independent the tracks reconstruction and the visualization of the container’s content from the grid and the 3D-voxels. The presence of a three-dimensional grid indeed limits the automatic object identification process. The problem is relevant in scenarios where the threat to be identified has a comparable size (or even smaller) to those of the single voxel and is located in a position not aligned with the grid. Clustering techniques, working directly on points, help to detect the presence of suspicious items inside the container, acting, as it will be show, as a filter for a preliminary analysis of the data.

Computations in Theoretical Physics: Techniques and Methods / 29**Parallel Cuba****Author:** Thomas Hahn¹¹ *MPI f. Physik***Corresponding Author:** hahn@feynarts.de

The Cuba library for multidimensional numerical integration has parallelization built in since Version 3. This presentation introduces Version 4 with extended facilities for distributed computation, including support for vectorized and GPU operation.

Summary:

Features and changes for parallel and vectorized operation of the Cuba library for multidimensional numerical integration are presented.

Data Analysis - Algorithms and Tools / 30**Developments in the ATLAS Tracking Software ahead of LHC Run 2****Authors:** Andreas Salzburger¹; Massimiliano Bellomo²; Nicholas Styles³¹ *CERN*² *University of Massachusetts (US)*³ *Deutsches Elektronen-Synchrotron (DE)***Corresponding Author:** nicholas.styles@desy.de

After a hugely successful first run, the Large Hadron Collider (LHC) is currently in a shut-down period, during which essential maintenance and upgrades are being performed on the accelerator. The ATLAS experiment, one of the four large LHC experiments has also used this period for consolidation and further developments of the detector and of its software framework, ahead of the new challenges that will be brought by the increased centre-of-mass energy and instantaneous luminosity in the next run period. This is of particular relevance for the ATLAS Tracking software, responsible for reconstructing the trajectory of charged particles through the detector, which faces a steep increase in CPU consumption due to the additional combinatorics of the high-multiplicity environment. The steps taken to mitigate this increase and stay within the available computing resources while maintaining the excellent performance of the tracking software in terms of the information provided to the physics analyses will be presented. Particular focus will be given to changes to the Event Data Model, replacement of the maths library, and adoption of a new persistent output format. The resulting CPU profiling results will be discussed, as well as the performance of the algorithms for physics processes under the expected conditions for the next LHC run.

Poster session - Board: 108 / 31**Challenges of the ATLAS Monte Carlo Production during Run-I and Beyond****Authors:** Claire Gwenlan¹; James Edward Ferrando²; Jiahang Zhong¹; John Derek Chapman³; Jose Enrique Garcia Navarro⁴; Sascha Mehlhase⁵; Vakho Tsulaia⁶; Wolfgang Ehrenfeld⁷

¹ *University of Oxford (GB)*

² *University of Glasgow (GB)*

³ *University of Cambridge (GB)*

⁴ *Instituto de Fisica Corpuscular (ES)*

⁵ *Ludwig-Maximilians-Univ. Muenchen (DE)*

⁶ *Lawrence Berkeley National Lab. (US)*

⁷ *Universitaet Bonn (DE)*

Corresponding Authors: c.gwenlan1@physics.ox.ac.uk, jose.enrique.garcia@cern.ch

In this presentation we will review the ATLAS Monte Carlo production setup including the different production steps involved in full and fast detector simulation. A report on the Monte Carlo production campaigns during Run-I and Long Shutdown 1 (LS1) will be presented, including details on various performance aspects. Important improvements in the work flow and software will be highlighted.

Besides standard Monte Carlo production for data analyses at 7 and 8 TeV, the production accommodates for various specialized activities. These ranges from extended Monte Carlo validation, Geant4 validation, pileup simulation using zero bias data and production for various upgrade studies. The challenges of these activities will be discussed.

Poster session - Board: 109 / 32

Testing and improving performance of the next generation ATLAS PanDA monitoring

Authors: Alexei Klimentov¹; Jaroslava Schovancova¹; Kaushik De²; Maxim Potekhin¹; Peter Love³; Sergey Belov⁴; Tadashi Maeno¹; Torre Wenaus¹

¹ *Brookhaven National Laboratory (US)*

² *University of Texas at Arlington (US)*

³ *Lancaster University (GB)*

⁴ *Joint Inst. for Nuclear Research (RU)*

Corresponding Author: jaroslava.schovancova@cern.ch

The PanDA Workload Management System (WMS) has been the basis for distributed production and analysis of the ATLAS experiment at the Large Hadron Collider since early 2008. Since the start of data taking of LHC Run I, PanDA usage has ramped up to over 1 exa-byte of processed data in 2013, and 1.5M peak completed jobs per day in 2014. The PanDA monitor is one of the core component of the PanDA WMS. It offers a set of views convenient to the PanDA users and site operators to follow the progress of submitted workloads, monitor activity and operational behavior, and drill down to the details of job execution when necessary, e.g. to diagnose problems. The monitor is undergoing significant extensions in the context of the BigPanDA project that is generalizing PanDA for the use of exascale science communities. The next generation monitor, BigPanDAmon is designed as a modular application, separating the data visualization and data access layers, following the MVC design pattern. While the front-end benefits from HTML5 and jQuery and its plugins, the back-end part serves data using a RESTful API. User community feedback is very important for the ATLAS PanDA monitor evolution. The BigPanDAmon development team consists of several developers and external contributors distributed world-wide. As new features are introduced rapidly, it is important to make sure all the developers can be productive, and the production front-end remains stable without preventable downtimes at the same time. The changes into the code are continuously integrated and deployed. In this talk we describe challenges for the quality assurance process of the ATLAS BigPanDAmon package. We describe steps taken and lessons learned to improve performance of the user interface of the ATLAS BigPanDAmon.

Data Analysis - Algorithms and Tools / 34**An automated framework for hierarchical reconstruction of B mesons at the Belle II experiment****Authors:** Christian Pulvermacher¹; Thomas Keck¹**Co-authors:** Martin Heck²; Michael Feindt¹; Thomas Kuhr¹¹ *KIT*² *KIT, Karlsruhe***Corresponding Author:** cpulvermacher@ekp.uni-karlsruhe.de

Belle II is an experiment being built at the e^+e^- SuperKEKB B factory, and will record decays of a large number of $B\bar{B}$ pairs.

This pairwise production of B mesons allows analysts to use one correctly reconstructed B meson to deduce the four-momentum and flavour of the other (signal-side) B meson, without reconstructing any of its daughter particles.

It also permits, in conjunction with a signal-side selection, to account for all tracks and calorimeter signals in an event and thus get a cleaner sample, e.g. for decays containing neutrinos.

I will present a software framework for Belle II that reconstructs B mesons in many decay modes with minimal user intervention.

It does so by reconstructing particles in user-supplied decay channels, and then in turn using these reconstructed particles in higher-level decays.

This hierarchical reconstruction allows to cover a relatively high fraction of all B decays by specifying a limited number of particle decays.

Multivariate classification methods are used to achieve a high signal-to-background ratio in each individual channel.

The entire reconstruction, including the application of precuts and classifier trainings, is automated to a high degree and will allow users to easily add new channels or to retrain on analysis-specific Monte Carlo samples.

Computations in Theoretical Physics: Techniques and Methods / 35**Six-loop calculations of the critical exponents in the ϕ^4 theory****Authors:** Dmitrii Batkovich¹; Mikhail Kompaniets¹¹ *SPbSU*

We present results of the six loop renormalization group calculations of the critical exponents in $O(n)$ -symmetric ϕ^4 theory in the framework of ϵ -expansion (minimal subtraction scheme). Technical details of this calculations are discussed. Obtained results are compared with experimental data and with results of other theoretical approaches like $1/n$ expansion, renormalization group in fixed space dimension, high temperature expansion and Monte-Carlo simulations.

Computing Technology for Physics Research / 36**An overview of the DII-HEP OpenStack based CMS Data Analysis****Author:** Tomas Linden¹

Co-authors: Carl Oscar Kraemer ¹; John White White ¹; Lirim Osmani ²; Matti Kortelainen ¹; Miika Komu ³; Paula Eerola ⁴; Salman Toor ¹; Sasu Tarkoma ²

¹ *Helsinki Institute of Physics (FI)*

² *University of Helsinki*

³ *Ericsson Research*

⁴ *University of Helsinki (FI)*

Corresponding Authors: tomas.linden@helsinki.fi, matti.kortelainen@cern.ch

An OpenStack based private cloud with the Gluster File System has been built and used with both CMS analysis and Monte Carlo simulation jobs in the Datacenter Indirection Infrastructure for Secure High Energy Physics (DII-HEP) project. On the cloud we run the ARC middleware that allows running CMS applications without changes on the job submission side. Our test results indicate that the adopted approach provides a scalable and resilient solution for managing resources without compromising on performance and high availability.

To manage the virtual machines (VM) dynamically in an elastic fashion, we are testing the EMI authorization service (Argus) and the Execution Environment Service (Argus-EES). An OpenStack plugin has been developed for Argus-EES.

The Host Identity Protocol (HIP) has been designed for mobile networks and it provides a secure method for IP multihoming. HIP separates the end-point identifier and locator role for IP address which increases the network availability for the applications. Our solution leverages HIP for traffic management.

This presentation gives an update on the status of the work and our lessons learned in creating an OpenStack based cloud for HEP.

Computations in Theoretical Physics: Techniques and Methods / 37

A development of an accelerator board dedicated for multi-precision arithmetic operations and its application to Feynman loop integrals

Author: Shinji Motoki¹

Co-authors: Atsushi Kawai ²; Fukuko Yuasa ¹; Hiroshi Daisaka ³; Jun'ichiro Makino ⁴; Naohito Nakasato ⁵; Tadashi Ishikawa ¹; Toshiyuki Fukushige ²

¹ *High Energy Accelerator Research Organization*

² *K & F Computing Research*

³ *Hitotsubashi University*

⁴ *RIKEN Advanced Institute for Computational Science*

⁵ *University of Aizu*

Corresponding Author: smotoki@post.kek.jp

With the discovery of Higgs particle in CERN Large Hadron Collider, precision measurements are expected in future International Linear Collider to explore new physics beyond Standard Model. In tandem with experiments, an accurate theoretical prediction is required. To meet such a request, higher order correction in perturbative quantum field theory becomes more and more important and the method to evaluate multi-loop integrals precisely should be provided. We have been developing DCM (Direct Computation Method) for loop integrals and it is a fully numerical method in the combination of multi-dimensional integration and the extrapolation technique. In multi-dimensional integration we encounter cases where the double precision operation is not enough due to loss of significances or trailing digits. For example, diagrams with infrared divergence or several greatly different masses in loops fit the cases. We present a hardware accelerator, GRAPE9-MPX, which

is based on GRAPE technique originally developed for gravitational N-body simulations. GRAPE-9 consists of FPGA boards in which a lot of processor elements are implemented. Each processor element has the dedicated logic for quadruple, hexuple and octuple precision arithmetic operation, and enables us high-speed computations in an SIMD way. We describe the design of GRAPE9-MPX and show some performance results taking the multi-precision computation of multi-loop integrals as examples.

Poster session - Board: 110 / 38

From DIRAC towards an Open Source Distributed Data Processing Solution

Authors: Andrei Tsaregorodtsev¹; Ricardo Graciani Diaz²

Co-authors: Jiri Chudoba ; Tomasz Szepieniec³

¹ *Marseille*

² *University of Barcelona (ES)*

³ *Unknown-Unknown-Unknown*

Corresponding Authors: graciani@ecm.ub.edu, andrei.tsaregorodtsev@cern.ch

The Open DISData Initiative is focusing on today's challenges of e-Science in a collaborative effort shared among different scientific communities, relevant technology providers and major e-Infrastructure providers. The target will be to evolve from existing partial solutions towards a common platform for distributed computing able to integrate already existing grid, cloud and other local computing and storage resources. This common platform will be guided by the needs of the scientist and their research. By joining this effort, and using this common platform to implement their own solutions, scientists will ensure at the same time robustness, interoperability and reusability as well as an important economy of scales. Sustainability will be achieved by selling customized solutions based on the common platform and its support to interested scientific and industrial clients.

In order to achieve this target we propose to build from existing solutions and to work in two directions addressing in parallel the big science and the long tail of science challenges. The first challenge refers to a relatively small number of well-organized large communities with very large data access and processing requirements. The second one refers to a large number of small loosely organized communities with an almost infinite variety of different applications and use cases.

Although started from the DIRAC technology, great care has been taken to cover all other relevant areas like: Workload and Data management (dCache.org and ARC) and advanced user interfaces including Portals and Identity management (InSilicoLab, SCI-BUS and Catania Science Gateway). Major e-Infrastructure projects and providers like EUDAT, EGI, NDGF, NeCTAR, OSG or EDGI are strongly supporting this Initiative. At the same time, some IT related private companies like Bull, Dell and ETL are willing to participate in different areas of the Initiative contributing with their industrial and marketing experience.

The long-term goal of the Open DISData Initiative is to build a self-sustained Collaboration, keeping the common platform up to date with new requirements and technologies, and offering a high quality

but affordable support model, with continuous security training and audit, and following industrial quality standards and procedures where appropriate.

Summary:

The sustainability of relevant software developments for scientific computing is a challenge. Many Research projects will have lifetimes of tens of years, ensuring appropriate support and further development of the necessary software is not a negligible effort that it is often not even considered. This initiative will try to promote mechanism to make “used” software sustainable.

Data Analysis - Algorithms and Tools / 39

HistFitter: a flexible framework for statistical data analysis

Authors: Alex Koutsman¹; David Cote²; Geert-Jan Besjes³; Jeanette Lorenz⁴; Max Baak⁵

¹ TRIUMF

² University of Texas

³ Radboud Universiteit Nijmegen

⁴ Ludwig-Maximilians-Universitat Munchen

⁵ CERN

Corresponding Authors: jeanette.miriam.lorenz@cern.ch, geert-jan.besjes@cern.ch

We present a software framework for statistical data analysis, called *HistFitter*, that has been used extensively in the ATLAS Collaboration to analyze data of proton-proton collisions produced by the Large Hadron Collider at CERN. Most notably, HistFitter has become a de-facto standard in searches for supersymmetric particles since 2012, with some usage for Exotic and Higgs boson physics.

HistFitter coherently combines several statistics tools in a programmable and flexible framework that is capable of bookkeeping hundreds of data models under study using thousands of generated input histograms. The key innovations of HistFitter are to weave the concepts of control, validation and signal regions into its very fabric, and to treat them with rigorous methods, while providing multiple tools to visualize and interpret the results through a simple configuration interface, as will become clear throughout this presentation.

Summary:

HistFitter is a software framework for statistical data analysis used widely throughout the ATLAS Collaboration at the LHC. HistFitter interfaces with the statistical tools HistFactory and RooStats to construct parametric models and to perform statistical tests of the data, and extends these tools in four key areas:

1. *Programmable framework:* HistFitter puts tools from several sources together in a coherent and programmable framework, capable of performing a complete statistical analysis of pre-formatted input data samples.
2. *Bookkeeping:* HistFitter can perform statistical tests and scan over parameter values of hundreds of signal hypotheses in an organized way from a single user-defined configuration file.
3. *Analysis strategy:* HistFitter has built-in the concepts of control, signal and validation regions, which are used to constrain, extrapolate and validate data-model predictions across analysis regions. HistFitter also introduces a rigorous treatment of validation regions that is new in high-energy physics.
4. *Presentation and interpretation:* the HistFitter framework keeps track of data models before and after fits to the measured data, and includes a collection of methods to determine the statistical significance using various statistical tests and to produce tables and plots expressing these results with publication-quality style.

Poster session - Board: 111 / 40

The Long Term Data Preservation (LTDP) project at INFN CNAF: CDF user case.

Authors: Daniele Gregori¹; Luca dell'Agnello²; Pier Paolo Ricci³; Silvia Amerio⁴; michele pezzi⁵

¹ *Istituto Nazionale di Fisica Nucleare (INFN)*

² *INFN-CNAF*

³ *INFN CNAF*

⁴ *University of Padova & INFN*

⁵ *Infn-cnaf*

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

In the last years the problem of digital preservation of valuable scientific data has significantly become one of the most important point to consider inwards scientific collaborations. In particular the long term preservation of almost all experimental data, raw and all related derived formats including calibration information, is one of the emerging requirements within the High Energy Physics (HEP) community for experiments that has already concluded the data taking phase. The DPHEP group (Data Preservation in HEP) coordinates the local teams within the whole collaboration and the different Tiers (computing centers). The INFN CNAF Tier1 is one of the reference site for data storage and computing in the LHC community but it also offers resources to many other HEP and non-HEP collaborations. In particular the CDF experiment has used the INFN CNAF Tier1 resources for many years and after the end of data taking in 2011, it is now facing the challenge to both preserve the large amount of data produced during several years and to retain the ability to access and reuse the whole amount of it in the future. According to this task the CDF Italian collaboration, together with the INFN CNAF computing center, has developed and is now implementing a long term future data preservation project in collaboration with the FNAL computing sector. The project comprises the copy of all CDF raw data and user level ntuples (about 4 PB) at the INFN CNAF site and the setup of a framework which will allow to access and analyze the data in the long term future. Therefore a big portion of the 4 PB of data (raw data and analysis-level ntuples) are currently being copied from FNAL to the INFN CNAF tape library backend and the system, which subsequently allows the data access, is being set up. In addition to this data access system, a data analysis framework is being developed in order to run the complete CDF analysis chain in the long term future, from raw data reprocessing to analysis-level ntuples production. In this contribution we first illustrate the difficulties and the technical solutions adopted to copy, store and maintain CDF data at the INFN CNAF Tier1 computing center. In addition we describe how we are exploiting virtualization techniques for the purpose of building the long term future analysis framework, and we also briefly illustrate the validation tests and techniques under development in order to check data integrity and software operation efficiency over time.

Computations in Theoretical Physics: Techniques and Methods / 41

Statistical methods for cosmic ray composition analysis at the Telescope Array Observatory

Authors: Grigory Rubtsov¹; Sergey Troitsky¹

¹ *INR RAS*

Corresponding Author: grisha@ms2.inr.ac.ru

The Telescope Array (TA) surface detector (SD) stations record the temporal development of the signal from the extensive air shower front which carries information about the type of the primary particle. We develop methods for studies of the primary mass composition of ultra-high-energy cosmic rays based on multivariate analysis (MVA). We propose to convert each observable into percentile rank with respect to Monte-Carlo. These ranks demonstrate stronger composition sensitivity

than raw data values. We report the results of the technique on TA primary mass composition and primary mass and neutrino search.

Poster session - Board: 112 / 42

The INFN CNAF Tier1 GEMSS Mass Storage System and Database facility activity

Authors: Alessandro Cavalli¹; Daniele Gregori²; Luca dell'Agnello¹; Pier Paolo Ricci³; Vincenzo Vagnoni⁴; Vladimir Sapunenko⁵

¹ *INFN-CNAF*

² *Istituto Nazionale di Fisica Nucleare (INFN)*

³ *INFN CNAF*

⁴ *Universita e INFN (IT)*

⁵ *INFN*

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

The consolidation of Mass Storage services at the INFN CNAF Tier1 Storage department that has occurred during the last 5 years resulted in a reliable, high performance and moderately easy-to-manage facility that can provide virtually all the storage archive, backup and database software services to several different use-cases. At present the INFN CNAF Tier1 GEMSS Mass Storage System installation, based upon an integration between the IBM GPFS parallel filesystem and the Tivoli Storage Manager (TSM) tape management software, is one of the biggest and most dependable hierarchical storage sites in Europe. It provides storage resources for about 12% of entirely LHC data, as well as for other non-LHC experiments, which could access the data using the standard SRM Grid services provided by the Storage Resource Manager (StoRM) software or, alternately, with access methods based on Xrootd and Http/Webdav, in particular for specific user cases currently under development. Besides these services, an Oracle Database facility is in production, running databases for storing and accessing relational data objects and for providing database services to the currently active use-cases, with a proven effective level of redundancy and availability. This installation takes advantage of several Oracle technologies, like Real Application Cluster (RAC), Automatic Storage Manager and the Enterprise Manager central management tools, making it possible to investigate recent use-cases, together with the other technologies available on the Oracle Database for performance optimization, ease of management and downtime reduction.

The objective of the present work is an illustration of the state-of-the-art of the INFN CNAF Tier1 Storage department software services and a site report of the success stories, results obtained occurred during the last period of activity and some brief descriptions of future projects. Particular attention is paid to the description of the administration, monitoring and problem-tracking tools that play primary roles in managing the whole framework in a complete, and relatively easy to learn, mode.

Poster session - Board: 113 / 43

Integration of PanDA workload management system with Titan supercomputer at OLCF

Author: Sergey Panitkin¹

¹ *Brookhaven National Laboratory (US)*

Corresponding Author: panitkin@bnl.gov

Experiments at the Large Hadron Collider (LHC) face unprecedented computing challenges. Heterogeneous resources are distributed worldwide, thousands of physicists analyzing the data need

remote access to hundreds of computing sites, the volume of processed data is beyond the exabyte scale, and data processing requires more than billions of hours of computing usage per year. The PanDA (Production and Distributed Analysis) workload management system (WMS) was developed to meet the scale and complexity of LHC distributed computing for the ATLAS experiment. While PanDA currently uses more than 100,000 cores at well over 100 Grid sites with a peak performance of 0.3 petaFLOPS, next LHC data taking run will require more resources than Grid computing can possibly provide.

The Worldwide LHC Computing Grid (WLCG) infrastructure will be sufficient for the planned analysis and data processing, but it will be insufficient for Monte Carlo (MC) production and any extra activities.

Additional computing and storage resources are therefore required.

To alleviate these challenges, ATLAS is engaged in an ambitious program to expand the current computing model to include additional resources such as the opportunistic use of supercomputers.

In turn this activity drives evolution of the PanDA WMS.

We will describe a project aimed at integration of PanDA WMS with

Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF).

Current approach utilizes modified PanDA pilot framework for job submission to

Titan's batch queue and local data management, with light-weight MPI wrappers to run single threaded workloads in parallel on Titan's multi-core worker nodes.

It also gives PanDA new capability to collect, in real time, information about unused worker nodes on Titan, which allows

to precisely define the size and duration of jobs submitted to Titan according to available free resources.

This capability can reduce job wait time and improve Titan's utilization efficiency.

This implementation was tested with Monte Carlo simulation jobs and is suitable for deployment with many other supercomputing platforms.

Poster session - Board: 114 / 44

A self-configuring control system for storage and computing departments at INFN CNAF Tier1.

Author: Daniele Gregori¹

Co-authors: Andrea Prosperini²; Pierpaolo Ricci²; Stefano dal pra²

¹ *Istituto Nazionale di Fisica Nucleare (INFN)*

² *INFN*

Corresponding Author: gregori@bo.infn.it

The storage and farming departments at the INFN CNAF Tier1 manage approximately thousands of computing nodes and several hundreds of servers that provides access to the disk and tape storage. In particular, the storage server machines should provide the following services: an efficient access to about 15 petabytes of disk space with different cluster of GPFS file system, the data transfers between LHC Tiers sites (Tier0, Tier1 and Tier2) via GridFTP cluster and Xrootd protocol and finally the writing and reading data operations on magnetic tape backend.

One of the most important and essential point in order to get a reliable service is a control system that can warn if problems arise and which is able to perform automatic recovery operations in case of service interruptions or major failures.

Moreover, during daily operations the configurations can change, i.e. if the GPFS cluster nodes roles can be modified and therefore the obsolete nodes must be removed from the control system production, and the new servers should be added to the ones that are already present.

The manual management of all these changes is an operation that can be somewhat difficult in case of several changes, it can also take a long time and is easily subject to human error or misconfiguration. For these reasons we have developed a control system with the feature of self-configure itself if any change occurs. Currently, this system has been in production for about a year at the INFN CNAF Tier1 with good results and hardly any major drawback.

There are three major key points in this system. The first is a software configurator service (e.g. Quattor or Puppet) for the servers machines that we want to monitor with the control system; this service must ensure the presence of appropriate sensors and custom scripts on the nodes to check

and should be able to install and update software packages on them. The second key element is a database containing information, according to a suitable format, on all the machines in production and able to provide for each of them the principal information such as the type of hardware, the network switch to which the machine is connected, if the machine is real (physical) or virtual, the possible hypervisor to which it belongs and so on. The last key point is a control system software (in our implementation we choose the Nagios software), capable of assessing the status of the servers and services, and that can attempt to restore the working state, restart or inhibit software services and send suitable alarm messages to the site administrators.

The integration of these three elements was made by appropriate scripts and custom implementation that allow the self-configuration of the system according to a decisional logic and the whole combination of all the above-mentioned components will be deeply discussed in this paper.

Plenary / 45

Extracting Rigorous Conclusions from Model/Data Comparisons

Author: Scott Pratt¹

¹ *Michigan State University*

Corresponding Author: prattsc@msu.edu

Many fields of science have developed multi-scale multi-component models to address large-scale heterogenous data sets. Constraining model parameters is made difficult by the inherent numerical cost of running such models and by the intertwining dependencies between parameters and observables. I will describe how the MADAI Collaboration has developed a suite of statistical tools based on the strategy of model emulators to meet these challenges. The tools have been applied to problems in galaxy formation and in relativistic heavy ion collisions, and have been formulated so that they can be transferred or expanded for numerous other problems. The tools assist with the distillation of data, the creation of model emulators and the exploration of parameter space via Markov Chain Monte Carlo. I will focus on the application to relativistic heavy ion collisions, where these methods are providing the means to reach the field's first rigorous quantitative conclusions.

Computations in Theoretical Physics: Techniques and Methods / 48

Differential reduction of generalized hypergeometric functions in application to Feynman diagrams.

Author: Vladimir Bytev¹

¹ *JINR*

Corresponding Author: bvv@jinr.ru

The differential reduction algorithm which allow one to express generalized hypergeometric functions with arbitrary values of parameters in terms of functions with fixed values of parameters differing from the original ones by integers is discussed in a context of evaluation of Feynman diagrams.

It is shown that the criterion of reducibility of multiloop Feynman integrals can be reformulated in terms of the criterion of reducibility of hypergeometric functions.

The HYPERDIRE - Mathematica based program, for differential reduction of hypergeometric functions of one and two variables with non-exceptional values of parameters to a set of basic functions is presented.

Data Analysis - Algorithms and Tools / 49

Identifying the Higgs boson with a Quantum Computer

Authors: Alexander Mott¹; Maria Spiropulu¹

¹ *California Institute of Technology*

A novel technique to identify events with a Higgs boson decaying to two photons and reject background events using neural networks trained on a quantum annealer is presented. We use a training sample composed of simulated Higgs signal events produced through gluon fusion and decaying to two photons and one composed of simulated background events with Standard Model two-photon final states. We design a problem such that minimizing the error of a neural network classifier is mapped to a quantum binary optimization problem (QUBO). This problem is encoded on the quantum annealer, which is designed to employ quantum adiabatic evolution to find the optimal configuration of qubits to solve the optimization problem. This is also the configuration of the network that minimizes the classification error. With the current hardware we are able to encode a problem with up to 30 correlated input variables and obtain solutions that have high efficiency to accept signal and reject background.

Poster session - Board: 115 / 50

Multilevel Workflow System in the ATLAS Experiment

Author: Alexandre Vaniachine¹

Co-authors: Alexei Klimentov²; Dmitry Golubkov³; Jose Enrique Garcia Navarro⁴; Kaushik De⁵; Misha Borodin⁶; Tadashi Maeno²

¹ *ANL*

² *Brookhaven National Laboratory (US)*

³ *Institute for High Energy Physics (RU)*

⁴ *Unknown*

⁵ *University of Texas at Arlington (US)*

⁶ *Moscow State Engineering Physics Institute (RU)*

Corresponding Author: vaniachine@anl.gov

The ATLAS experiment is scaling up Big Data processing for the next LHC run using a multilevel workflow system comprised of many layers. In Big Data processing ATLAS deals with datasets, not individual files. Similarly a task (comprised of many jobs) has become a unit of the ATLAS workflow in distributed computing, with about 0.8M tasks processed per year. In order to manage the diversity of LHC physics (exceeding 35K physics samples per year), the individual data processing tasks are organized into workflows. For example, the Monte Carlo workflow is composed of many steps: generate or configure hard-processes, hadronize signal and minimum-bias (pileup) events, simulate energy deposition in the ATLAS detector, digitize electronics response, simulate triggers, reconstruct data, convert the reconstructed data into ROOT ntuples for physics analysis, etc. Outputs are merged and/or filtered as necessary to optimize the chain.

The bi-level workflow manager - ProdSys2 - generates actual workflow tasks and their jobs are executed across more than a hundred distributed computing sites by PanDA –the ATLAS job-level workload management system. On the outer level, the Database Engine for Tasks (DEFT) empowers production managers with templated workflow definitions. On the next level, the Job Execution and Definition Interface (JEDI) is integrated with PanDA to provide dynamic job definition tailored to the sites capabilities.

We report on scaling up the production system to accommodate a growing number of requirements from main ATLAS areas: Trigger, Physics and Data Preparation.

Summary:

We report on scaling up the ATLAS production system to accommodate a growing number of requirements for the next LHC run.

Poster session - Board: 116 / 52

Monitoring of IaaS and scientific applications on the Cloud using the Elasticsearch ecosystem

Authors: Andrea Guarise¹; Dario Berzano²; Sara Vallero³; Stefano Bagnasco⁴; Stefano Lusso⁵

¹ *Istituto Nazionale Fisica Nucleare (IT)*

² *CERN*

³ *Universita e INFN (IT)*

⁴ *I.N.F.N. TORINO*

⁵ *INFN-TO*

Corresponding Authors: svallero@to.infn.it, bagnasco@to.infn.it

The private Cloud at the Torino INFN computing centre offers IaaS services to different scientific computing applications. The infrastructure is managed with the OpenNebula cloud controller. The main stakeholders of the facility are a grid Tier-2 site for the ALICE collaboration at LHC, an interactive analysis facility for the same experiment and a grid Tier2 site for the BESIII collaboration, plus a number of other smaller tenants that will increase in the near future. Besides keeping track of the usage, the automation of dynamic allocation of resources to tenants requires detailed monitoring and accounting of the resource usage.

As a first investigation towards this, we set up a monitoring system to inspect the site activities both in terms of IaaS and applications running on the hosted virtual instances. For this purpose we used the Elasticsearch, Logstash and Kibana stack. In the current implementation, the heterogeneous accounting information is fed to different MySQL databases and sent to Elasticsearch via a custom Logstash plugin. For the IaaS metering, we developed sensors for the OpenNebula API. The IaaS level information gathered through the API is sent to the MySQL database through an ad-hoc developed RestFUL WEB service, which is also used for accounting purposes. Concerning the application level, we used the Root plugin TProofMonSenderSQL to collect accounting data from the interactive analysis facility. The BESIII virtual instances used to be monitored with Zabbix, as a proof of concept we also retrieve the information contained in the Zabbix database. Each of these three cases is indexed separately in Elasticsearch. We are now starting to consider dismissing the intermediate level provided by the SQL DB and evaluating a NoSQL option as a unique central database for all the monitoring information. We setup a set of Kibana dashboards with pre-defined queries in order to monitor the relevant information in each case. In this way we have achieved a uniform monitoring interface for both the IaaS and the scientific applications mostly leveraging off-the-shelf tools.

Poster session - Board: 117 / 53

Analyzing data flows of WLCG jobs at batch job level

Author: Eileen Kuhn¹

Co-authors: Andreas Petzold¹; Christopher Jung¹; Manuel Giffels¹; Max Fischer¹

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

Corresponding Authors: christopher.jung@kit.edu, eileen.kuhn@cern.ch

With the introduction of federated data access to the workflows of WLCG, it is becoming increasingly important for data centers to understand specific data flows regarding storage element accesses, firewall configurations, as well as the scheduling of batch jobs themselves. As existing batch system monitoring and related system monitoring tools do not support measurements at batch job level, a new tool has been developed and put into operation at the GridKa Tier1 center for monitoring continuous data streams and characteristics of WLCG jobs and pilots. Long term measurements and data collection are in progress. These measurements already have been proven to be useful analyzing misbehaviors and various issues. Therefore we aim for an automated, realtime approach for anomaly detection. As a requirement, prototypes for standard workflows have to be examined. Based on measurements of several months, different features of HEP jobs are evaluated regarding their effectiveness for data mining approaches to identify these common workflows.

The presentation will introduce the actual measurement approach and statistics as well as the general concept and first results classifying different HEP job workflows derived from the measurements at GridKa.

Poster session - Board: 118 / 54

Tier 3 batch system data locality via managed caches

Author: Max Fischer¹

Co-authors: Christopher Jung¹; Eileen Kuhn¹; Gunter Quast¹; Manuel Giffels¹

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

Corresponding Authors: christopher.jung@kit.edu, max.fischer@cern.ch

Modern data processing increasingly relies on data locality for performance and scalability, whereas the common HEP approaches aim for uniform resource pools with minimal locality, recently even across site boundaries.

To combine advantages of both, the High Performance Data Analysis (HPDA) Tier 3 concept opportunistically establishes data locality via coordinated caches.

In accordance with HEP Tier 3 activities, the design incorporates two major assumptions:

1. Only a fraction of data is accessed regularly and thus the deciding factor for overall throughput.
2. Data access may fallback to non-local, making permanent local data availability an inefficient resource usage strategy.

Based on this, the HPDA design generically extends available storage hierarchies into the batch system.

Using the batch system itself for scheduling file locality, an array of independent caches on the worker nodes is dynamically populated with high-profile data.

Cache state information is exposed to the batch system both for managing caches and scheduling jobs.

As a result, users directly work with a regular, adequately sized storage system.

However, their automated batch processes are presented with local replications of data whenever possible.

We highlight the potential and limitations of currently available technologies in light of HEP Tier 3 activities, showcase the current design and implementation of the HPDA data locality, and present first experiences with our prototype.

The ALICE analysis train system

Author: Markus Bernhard Zimmermann¹

¹ *CERN and Westfaelische Wilhelms-Universitaet Muenster (DE)*

Corresponding Author: markus.bernhard.zimmermann@cern.ch

In order to cope with the large recorded data volumes (around 10 PBs per year) at the LHC, the analysis within ALICE is done by hundreds of analysis users on a GRID system. High throughput and short turn-around times are achieved by a centralized system called the 'LEGO'trains. This system combines analysis of different users to so-called analysis trains which are then executed within the same GRID jobs reducing the number of times the data needs to be read from the storage systems. To prevent that a single failing analysis jeopardizes the results of all users within a train, an automated testing procedure has been developed. Each analysis is tested separately for functionality and performance before it is allowed to be submitted to the GRID. The analysis train system steers the job management, the merging of the output and sends notifications to the users. Clear advantages of such a centralized system are improved performance, good usability for users and the means of bookkeeping important for the reproducibility of the results.

The train system builds upon the already existing ALICE tools, i.e. the analysis framework as well as the GRID submission and monitoring infrastructure. The entry point to the train system is a web interface which allows to configure the analysis and the desired datasets as well as to test and submit the train. While the analysis configuration is done directly by the users, datasets and train submission are controlled by a smaller group of operators. The analysis train system is operational since early 2012 and has quickly gained popularity with a continuously increasing trend. Throughout 2013, about 4800 trains have been submitted consuming about 2600 CPU years while analyzing 75 PB of data. This constitutes about 57% of the resources consumed for analysis in ALICE in 2013.

Within the GRID environment which by its nature has changing availability of resources, it has been very challenging to achieve a fast turn around time. Various measures have been implemented, e.g. to obtain a speedy merging process and to avoid that a few problematic GRID jobs stall the completion of a train.

The talk will introduce the analysis train system which has become very important for the daily analysis within ALICE. Further, the talk will focus on bottlenecks which have been identified and addressed by dedicated improvements. Finally, the lessons learned when setting up an organized analysis system for a user group which is in the hundreds will be discussed.

Poster session - Board: 205 / 56

New exclusive CHIPS-TPT algorithms for simulation of neutron-nuclear reactions

Author: Dmitry SAVIN¹

Co-author: Mikhail KOSOV¹

¹ *VNIIA, Moscow*

The CHIPS-TPT physics library is being developed for simulation of neutron-nuclear reactions on the new exclusive level. The exclusive modeling conserves energy, momentum and quantum numbers in each neutron-nuclear interaction. The CHIPS-TPT algorithms are based on the exclusive CHIPS library, which is compatible with Geant4. Special CHIPS-TPT physics lists in the Geant4 format are provided, which help to use the CHIPS-TPT libraries in the Geant4 simulation. The calculation time for an exclusive CHIPS-TPT simulation is comparable to the time of the corresponding Geant4-HP simulation. In addition to the reduction of the deposited energy fluctuations, which is a consequence of the energy conservation, the CHIPS-TPT libraries provide a possibility of simulation of the secondary particles correlation, e.g. secondary gammas, and of the Doppler broadening of gamma lines in the spectrum, which can be measured by germanium detectors.

Data Analysis - Algorithms and Tools / 57**HERAFitter - an open source QCD fit framework****Authors:** Ringaile Placakyte¹; Voica Ana Maria Radescu¹**Co-author:** Andrey Sapronov²¹ *Deutsches Elektronen-Synchrotron (DE)*² *Joint Inst. for Nuclear Research (RU)***Corresponding Authors:** andrey.sapronov@cern.ch, voica.ana.maria.radescu@cern.ch

We present the HERAFitter project, a unique platform for QCD analyses of hadron-induced processes in the context of multi-process and multi-experiment setting. Based on the factorisable nature of the hadronic cross sections into universal parton distribution functions (PDFs) and process dependent partonic scattering cross sections, HERAFitter allows determination of the PDFs from various hard scattering measurements.

The project successfully encapsulates a wide variety of QCD tools to facilitate investigations of the experimental data and theoretical calculations. HERAFitter is the first open source platform which is optimal for benchmarking studies. It allows for direct comparisons of various theoretical approaches under the same settings, a variety of different methodologies in treating of the experimental and model uncertainties. The growth of HERAFitter benefits from its flexible modular structure driven by QCD advances.

Poster session - Board: 119 / 60**AliEn File Access Monitoring Service (FAMoS)****Authors:** Armenuhi Abramyan¹; Narine Manukyan¹**Co-authors:** Ara Grigoryan¹; Costin Grigoras²; Latchezar Betev²; Miguel Martinez Pedreira²; Pablo Saiz²; Predrag Buncic²¹ *A.I. Alikhanyan National Science Laboratory (Yerevan Physics Institute) Foundation*² *CERN***Corresponding Authors:** nmanukya@mail.yerphi.am, aabramya@mail.yerphi.am

FAMoS leverages the information stored in the central AliEn file catalogue, which describes every file in a Unix-like directory structure, as well as metadata on file location and replicas. In addition, it uses the access information provided by a set of API servers, which are used by all Grid clients to access the catalogue. The main functions of FAMoS are to sort the file accesses by logical groups, access time, user and storage element. The collected data can identify rarely used groups of files, as well as those with high popularity over different time periods. This can be further used to optimize file distribution and replication factors, thus increasing the data processing efficiency. This paper will describe in detail the FAMoS structure and user interface and will present the results obtained in one year of operation of the service.

Computing Technology for Physics Research / 61

Gaudi Components for Concurrency: Concurrency for Existing and Future Experiments

Authors: Daniel Funke¹; Illya Shapoval²

Co-authors: Benedikt Hegner³; Danilo Piparo³; Marco Clemencic³; Pere Mato Vila³

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

² *CERN, KIPT, UNIFE, INFN-FE*

³ *CERN*

Corresponding Author: daniel.funke@cern.ch

HEP experiments produce enormous data sets at an ever-growing rate. To cope with the challenge posed by these data sets, experiments' software needs to embrace all capabilities modern CPUs offer. With decreasing $\text{memory}/\text{core}$ ratio, the one-process-per-core approach of recent years becomes less feasible. Instead, multi-threading with fine-grained parallelism needs to be exploited to benefit from memory sharing among threads.

Gaudi is an experiment-independent data processing framework, used for instance by the ATLAS and LHCb experiments at CERN's Large Hadron Collider. It has originally been designed with only sequential processing in mind. In a recent effort, the framework has been extended to allow for multi-threaded processing. This includes components for concurrent scheduling of several algorithms – either processing the same or multiple events, thread-safe data store access and resource management.

In the sequential case, the relationships between algorithms are encoded implicitly in their pre-determined execution order. For parallel processing, these relationships need to be expressed explicitly, in order for the scheduler to be able to exploit maximum parallelism while respecting dependencies between algorithms. Therefore, means to express and automatically track these dependencies need to be provided by the framework.

The experiments using Gaudi have built a substantial code base, thus a minimally intrusive approach and a clear migration path for the adoption of multi-threading is required for the extended framework to succeed.

In this paper, we present components introduced to express and track dependencies of algorithms to deduce a precedence-constrained directed acyclic graph, which serves as basis for our algorithmically sophisticated scheduling approach for tasks with dynamic priorities.

We introduce an incremental migration path for existing experiments towards parallel processing and highlight the benefits of explicit dependencies even in the sequential case, such as sanity checks and sequence optimization by graph analysis.

Summary:

Non-intrusiveness is one of the primary goals of the project, in order to ease the effort of introducing multi-threading to the experiments' large existing code bases.

The concurrent components therefore present an option to the user and do not interfere with existing sequential production code.

This allows for an incremental revision of the existing implementation and adoption of the components offered.

In order to process data concurrently the user must

1. explicitly declare data in- and outputs to each algorithm,

2. declare which reusable modules – called Tools in the Gaudi context – are used by an algorithm, and
3. revise thread-unsafe code such as use of caches, back-channel communication, race conditions in the update of shared data, ...

The first two points can be achieved by using the newly introduced data and tool handles, respectively. To ease their use, the interfaces of the handles were designed according to components familiar to Gaudi developers.

The third point requires a case by case analysis of the operations performed by the algorithm.

Having explicitly stated the algorithms' in- and outputs allows the automated deduction of the data flow among them.

In Gaudi, the execution path is not only influenced by data dependencies but also by control flow constructs.

These are expressed by grouping algorithms in sequences.

Each algorithm produces a binary decision, which causes the further execution or abortion of its sequence.

By composing several sequences, a complex control flow can be modeled.

Including both data dependencies and control flow constructs in a unified graph, gives a coherent picture of the relationships between algorithms.

The information required to construct the control flow portion of the graph can be extracted from the sequence hierarchy automatically and requires no further user intervention.

The unified data- and control-flow graph is a precedence-constrained directed acyclic graph.

The scheduler identifies algorithms with fulfilled data dependencies in the graph, which are candidates for execution.

An algorithm's execution priority is determined by its connectedness and membership to the critical path;

it can become zero if it is no longer required to execute by the control flow.

Candidates with non-zero priority are submitted by the scheduler for execution in the order of their priorities.

The graph structure also serves a secondary purpose:

it allows the analysis of a given configuration statically for unfulfillable data dependencies and superfluous control flow constructs.

Furthermore, the length of the critical path and the maximum number of concurrently executable algorithms

can be determined to aid in resource planning.

The gradual migration strategy and the benefits of static configuration analysis

present a strong argument for the adoption of the components and concepts developed in the concurrent Gaudi project.

The adoption presents an opportunity for general code improvement as well as the possibility to explore the use of accelerators and co-processors in the future.

Data Analysis - Algorithms and Tools / 62

The Run 2 ATLAS Analysis Event Data Model

Authors: Andy Buckley¹; Attila Krasznahorkay²; Dag Gillberg²; Edward Moyses³; Karsten Koeneke⁴; Marcin Nowak⁵; Markus Elsing²; Peter Van Gemmeren⁶; Scott Snyder⁵; Till Eifert⁷

¹ University of Glasgow (GB)

² CERN

³ University of Massachusetts (US)

⁴ Albert-Ludwigs-Universitaet Freiburg (DE)

⁵ Brookhaven National Laboratory (US)

⁶ *Argonne National Laboratory (US)*

⁷ *SLAC National Accelerator Laboratory (US)*

Corresponding Authors: marcin.nowak@cern.ch, attila.krasznahorkay@cern.ch

During the LHC's first Long Shutdown (LS1) ATLAS set out to establish a new analysis model, based on the experience gained during Run 1. A key component of this is a new Event Data Model (EDM), called the xAOD. This format, which is now in production, provides the following features:

- A separation of the EDM into interface classes that the user code directly interacts with, and data storage classes that hold the payload data. The user sees an Array of Structs (AoS) interface, while the data is stored in a Struct of Arrays (SoA) format in memory, thus making it possible to efficiently auto-vectorise reconstruction code.
- A simple way of augmenting and reducing the information saved for different data objects. This makes it possible to easily decorate objects with new properties during data analysis, and to remove properties that the analysis does not need.
- A persistent file format that can be explored directly with ROOT, either with or without loading any additional libraries. This allows fast interactive navigation without additional overheads, while maintaining the possibility of using the interface EDM to its full potential. Both compiled C++ or interactive Python code can be used after loading a minimal set of libraries.

The presentation will describe the design of the xAOD data format, showing the first results on reconstruction and analysis performance.

Summary:

A presentation of the ATLAS Analysis Event Data Model (xAOD) developed for LHC's Run 2.

Data Analysis - Algorithms and Tools / 63

Geant4 developments in reproducibility, multi-threading and physics

Author: John Apostolakis¹

Co-authors: Alberto Ribon¹; Andrea Dotti²; Dennis Herbert Wright²; Gabriele Cosmo¹; Gunter Folger¹; Makoto Asai²; Marc Verderi³; Marek Gayer¹; Michael Kelsey²; Tatiana Nikitina⁴; Tatsumi Koi²; Vladimir Ivantchenko¹; Vladimir Uzhinskiy⁵; Witold Pokorski¹

¹ *CERN*

² *SLAC National Accelerator Laboratory (US)*

³ *Ecole Polytechnique (FR)*

⁴ *Gangneung-Wonju National University (KR)*

⁵ *Joint Inst. for Nuclear Research (RU)*

Corresponding Author: john.apostolakis@cern.ch

The Geant4 toolkit is used in the production detector simulations of most recent High Energy Physics experiments, and diverse applications in medical physics, radiation estimation for satellite electronics and other fields. We report on key improvements relevant to HEP applications that were provided in the most recent releases: 9.6 (Dec 2012) and 10.0 (Dec 2013).

'Strong' reproducibility of events is a requirement of LHC experiments: the same results (hit data) must be obtained when restarting a job from an intermediate event. A significant extension of testing and several corrections in releases 9.6 and 10.0 achieve this for most applications.

Multithreading(MT) is included in release 10.0 for parallelism at the event level. It can harness the computing power of multicore machines scalably, with a small increase in memory footprint per additional thread. We report on performance and memory size measurements for release 10.0 on multi-core CPUs and on Xeon Phi (™). Strong reproducibility is required for repeatability of results when using multithreading (MT).

Release 10.0 contains the first version of ‘USolids’, a common library of shape primitives being developed to replace both Geant4 and Root implementations. The USolid implementations for many solids are an option at installation. Field propagation is extended to particles with dipole moments coupling to a magnetic field gradient. There is a new option to turn on gravity, which used to require a user’s own code.

Geant4 9.6 includes the new INCL 5.1 cascade, reengineered and written in C++ for reactions of pions and nuclei (H-alpha) up to 3 GeV. Release 10.0 introduces tracking of long-lived meta-stable nuclides (isomers), and isomer production in the de-excitation models and radioactive decay. A new improved neutron capture model and revised neutron cross sections below 20 MeV are used in production physics lists (rel. 10.0). These improve simulations that cannot afford the higher CPU cost of the detailed NeutronHP package. Results for tungsten calorimeters, are brought close to those of NeutronHP.

Many other improvements and refinements were made in other physics models. The Bertini-inspired cascade was extended to provide gamma- and electro-nuclear interactions and the capture at rest of negative hadrons and muons. The phase-space generation for multi-body states and the two-body final state angular distributions were improved. Development of the FTF string model includes the extension to handle nucleus-nucleus collisions and improved diffraction dissociation of protons, pions and kaons. In addition the LEP and HEP parameterised physics models (inherited from GHEISHA and refined) and the CHIPS models have been retired and removed from release 10.0.

Poster session - Board: 301 / 64

Toolbox for multiloop Feynman diagram calculations using R^* operation

Authors: Batkovich Dmitry¹; Loran Adzhemyan¹; Mikhail Kompaniets¹; Yury Kirienco¹

¹ *St. Petersburg State University (RU)*

Corresponding Authors: batya239@gmail.com, mikhail.kompaniets@cern.ch

We present the set of tools for computations on Feynman diagrams. Various package modules implement:

- graph manipulation, serialization, symmetries and automorphisms
- calculators, which are used to calculate integrals by particular methods (analytical or numerical).
- UV-counterterms calculation using IR-rearrangement and R^* operation (minimal subtraction scheme)

The following calculators are available out of the box: reduction to master integrals (using LiteRed IPB and DRR rules), sector decomposition and the Gegenbauer polynomial x-space technique. These set of calculators can be extended by creating your own Feynman diagrams calculators using API provided. Library implemented in python (≥ 2.6 compatibility) and uses GiNAC as computer algebra engine.

Poster session - Board: 120 / 65

Scalable cloud without dedicated storage

Authors: Andrey Zarochentsev¹; Dmitry Batkovich¹; Mikhail Kompaniets¹

¹ *St. Petersburg State University (RU)*

Corresponding Author: batya239@gmail.com

We present a prototype of a scalable computing cloud. It is intended to be deployed on the basis of a cluster without a separate dedicated storage. The dedicated storage is replaced by the distributed software storage. In addition, all cluster nodes are used both as computing nodes and as storage nodes. This increases utilization of the cluster resources as well as improves the fault tolerance and performance of distributed storage. Another advantage of this solution is high scalability with relatively low initial and maintenance cost. The solution is built on the basis of the open source components like CloudStack, Ceph, etc.

Data Analysis - Algorithms and Tools / 66

A Neural Network z -Vertex Trigger for Belle II

Authors: Alois Knoll¹; Christian Kiesling²; Fernando Abudinen²; Jochen Schieck³; Martin Heck⁴; Michael Feindt⁴; Rudolf Frühwirth³; Sara Neuhaus¹; Sebastian Skambraks¹; Stephan Paul¹; Yang Chen¹

¹ *TU München*

² *MPI für Physik, München*

³ *Inst. f. Hochenergiephysik, Wien*

⁴ *KIT, Karlsruhe*

Corresponding Authors: sneuhaus@mpp.mpg.de, rudolf.fruehwirth@oeaw.ac.at

The Belle II experiment, the successor of the Belle experiment, will go into operation at the upgraded KEKB collider (SuperKEKB) in 2016. SuperKEKB is designed to deliver an instantaneous luminosity $\mathcal{L} = 8 \times 10^{35} \text{cm}^{-2}\text{s}^{-1}$, a factor of 40 larger than the previous KEKB world record. The Belle II experiment will therefore have to cope with a much larger machine background than its predecessor Belle, in particular from events outside of the interaction region. We present the concept of a track trigger, based on a neural network approach, that is able to suppress a large fraction of this background by reconstructing the z (longitudinal) position of the event vertex within the latency of the first level trigger.

The trigger uses the topological and drift time information of the hits from the Central Drift Chamber (CDC) of Belle II within narrow cones in polar and azimuthal angle as well as in transverse momentum (sectors), and estimates the z -vertex without explicit track reconstruction. The preprocessing for the track trigger is based on the track information provided by the standard CDC trigger. It takes input from the 2D track finder, adds information from the stereo wires of the CDC, and finds the appropriate sectors in the CDC for each track in a given event. Within each sector, the z -vertex of the associated track is estimated by a specialized neural network, with the wire hits from the CDC as input and a continuous output corresponding to the scaled z -vertex.

The neural algorithm will be implemented in programmable hardware. To this end a Virtex 7 FPGA board will be used, which provides at present the most promising solution for a fully parallelized implementation of neural networks or alternative multivariate methods. A high speed interface for external memory will be integrated into the platform, to be able to store the $\mathcal{O}(10^9)$ parameters required.

The contribution presents the results of our feasibility studies and discusses the details of the envisaged hardware solution.

Poster session - Board: 121 / 67

EicRoot: yet another FairRoot framework clone

Author: Alexander Kiselev¹

¹ BNL

Corresponding Author: ayk@bnl.gov

The long-term upgrade plan for the RHIC facility at BNL foresees the addition of a high-energy polarized electron beam to the existing hadron machine thus converting RHIC into an Electron-Ion Collider (eRHIC) with luminosities exceeding $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. GEANT simulation framework for this future project (EicRoot) is based on FairRoot and its derivatives (PandaRoot, CbmRoot, FopiRoot). General layout of the EicRoot framework, as well as its distinguishing features (user-friendly tools for momentum/energy resolution studies of basic tracker and EM calorimeter setups, a powerful CAD-to-ROOT geometry converter, etc) will be presented.

Computing Technology for Physics Research / 69

Towards a generic high performance geometry library for particle transport simulation

Authors: Andrei Gheata¹; Federico Carminati¹; Georgios Bitzes²; Guilherme Lima³; Johannes Christof De Fine Licht⁴; John Apostolakis¹; Laurent Duhem⁵; Marilena Bandieramonte⁶; Mihaly Novak¹; Oksana Shadura⁷; Philippe Canal⁸; Raman Sehgal⁹; Rene Brun¹; Sandro Christian Wenzel¹; Soon Yung Jun⁸; Victor Daniel Elvira⁸

¹ CERN

² University of Athens (GR)

³ FermiLab (US)

⁴ University of Copenhagen (DK)

⁵ Intel

⁶ Universita e INFN (IT)

⁷ National Technical Univ. of Ukraine "Kyiv Polytechnic Institute

⁸ Fermi National Accelerator Lab. (US)

⁹ Bhabha Atomic Research Centre (IN)

Corresponding Author: sandro.wenzel@cern.ch

Thread-parallelization and single-instruction multiple data (SIMD) “vectorization” of software components in HEP computing is becoming a necessity to fully benefit from current and future computing hardware. In this context, the Geant-Vector/GPU simulation prototypes aim to reengineer current software for the simulation of the passage of particles through detectors in order to increase the event throughput in simulations of particle collisions at higher energy. As one of the core modules of detector simulation, the geometry library (treating tasks such as distance calculations to detector elements) plays a central role and vectorizing its algorithms will be one of the cornerstones towards achieving good CPU performance.

Here, we report on the progress made in vectorizing various algorithms, as well as in applying other new C++ template based optimizations of existing code available in the Geant4, ROOT or USolids geometry libraries. Next to a presentation of performance improvements achieved so far, we will focus on a discussion of our software development approach that aims to bridge the gap between the need to provide optimized code for all use cases of the library (e.g., single particle and many-particle APIs) and a required support for different architectures (CPU and GPU) on the one hand, and the necessity to keep the code base small, manageable and maintainable on the other hand. We report on a generic and templated C++ geometry library as a continuation of the USolids project and we argue that the experience gained with these developments will be beneficial to other parts of the simulation software, such as for the optimization of the physics library, and possibly to other parts of the experiment software stack, such as reconstruction and analysis.

Poster session - Board: 122 / 70

Using Functional Languages and Declarative Programming to analyze ROOT data: LINQtoROOT

Author: Gordon Watts¹

¹ University of Washington (US)

Corresponding Author: gwatts@uw.edu

Modern high energy physics analysis is complex. It typically requires multiple passes over different datasets, and is often held together with a series of scripts and programs. For example, one has to first reweight the jet energy spectrum in Monte Carlo to match data before plots of any other jet related variable can be made. This requires a pass over the Monte Carlo and the Data to derive the reweighting, and then another pass over the Monte Carlo to plot the variables the analyzer is really interested in. With most modern ROOT based tools this requires separate analysis loops for each pass, and script files to glue to the results of the two analysis loops together. A framework has been developed that uses the functional and declarative features of the C# language and its Language Integrated Query (LINQ) extensions to declare the analysis. The framework uses language tools to convert the analysis into C++ and runs ROOT or PROOF as a backend to get the results. This gives the analyzer the full power of an object-oriented programming language to put together the analysis and at the same time the speed of C++ for the analysis loop. The tool allows one to incorporate C++ algorithms written for ROOT by others. A byproduct of the design is the ability to cache results between runs, dramatically reducing the cost of adding one-more-plot and also to keep a complete record associated with each plot for data preservation reasons. The code is mature enough to have been used in ATLAS analyses. The package is open source and available on the open source site CodePlex.

Summary:

A tool that converts functional queries over ROOT data in the C# language into C++ that runs in the ROOT or PROOF environment is described. This open source translation software uses declarative programming to describe the tasks an analyzer wants accomplished. A few lines can generate a large number of plots from several different datasets, and the results can be combined, all in one program. A byproduct of the technique means that previously generated plots can be cached to avoid having to re-run later. This could also have data preservation implications.

Computing Technology for Physics Research / 71

Distributed job scheduling in MetaCentrum

Authors: Miroslav Ruda¹; Šimon Tóth¹

¹ CESNET

Corresponding Author: simont@mail.muni.cz

MetaCentrum, Czech national grid, provides access to various resources across Czech Republic. In this talk, we will describe unique features of job scheduling system used in MetaCentrum. System is based on heavily modified Torque batch system, which is improved to support requirements of such large installation. We will describe distributed setup of several standalone servers, which can work as independent servers, while preserving global scheduling via cooperating schedulers, as well as extensions supporting scheduling of GPU jobs, support for encapsulation of jobs into virtual machines (started on-demand) or even virtual clusters (hidden in on-demand prepared private virtual network).

Computing Technology for Physics Research / 72**Belle II distributed computing****Author:** Pavel Krokovny¹**Co-author:** Takanori Hara²¹ *Budker Institute of Nuclear Physics (RU)*² *on behalf of the Belle II computing group***Corresponding Author:** pavel.krokovny@cern.ch

The existence of large matter-antimatter asymmetry (CP violation) in the b -quark system as predicted in the Kobayashi-Maskawa theory was established by the B -Factory experiments. However, this cannot explain the magnitude of the matter-antimatter asymmetry of the universe we live in today. This indicates undiscovered new physics exists. The Belle II experiment, the next generation of the B -Factory, is expected to reveal the new physics by accumulating 50 times more data ($\sim 50\text{ab}^{-1}$) than Belle by 2022. The Belle II computing system has to handle an amount of beam data eventually corresponding to several tens of PetaByte per year under an operation of the SuperKEKB accelerator with a designed instantaneous luminosity. Under this situation, it cannot be expected that one site, KEK, will be able to provide all computing resources for the whole Belle II collaboration including the resources not only for the raw data processing but also for the MC production and physics analysis done by users. In order to solve this problem, Belle II employed the distributed computing system based on DIRAC, which provides us the interoperability of heterogeneous computing systems such as grids with different middleware, clouds and the local computing clusters. Since the last year, we performed the MC mass production campaign to confirm the feasibility and find out the possible bottleneck of our computing system. In parallel, we also started the data transfer challenge through the transpacific and transatlantic networks. This presentation describes the highlights of the Belle II computing and the current status. We will also present the experience of the latest MC production campaign in 2014.

Poster session - Board: 123 / 73**VISPA: Direct access and execution of data analyses for collaborations****Author:** Christian Glaser¹**Co-authors:** Daniel Van Asseldonk²; Gero Müller¹; Marcel Rieger³; Martin Erdmann²; Martin Urban¹; Robert Fischer²; Thorben Quast¹¹ *RWTH Aachen University*² *Rheinisch-Westfaelische Tech. Hoch. (DE)*³ *R***Corresponding Author:** glaser@physik.rwth-aachen.de

The VISPA web framework opens a new way of collaborative work. All relevant software, data and computing resources are supplied on a common remote infrastructure. Access is provided through a web GUI, which has all functionality needed for working conditions comparable to a personal computer. The analyses of colleagues can be reviewed and executed with just one click. Furthermore,

code can be modified and extended –given the necessary permissions –either directly via shared files or through a repository. VISPA can be extended to fit the specific needs of an experiment. A GUI interface to the analysis framework “Offline” of the Pierre Auger collaboration is already in use.

Computations in Theoretical Physics: Techniques and Methods / 74

A new generator for the Drell-Yan process

Author: Yahor Dydyska¹

Co-author: Vitaly Yermolchik¹

¹ *Byelorussian State University (BY)*

Corresponding Authors: vitaly.yermolchik@cern.ch, yahor.dydyshka@cern.ch

We present the Monte Carlo event generator *LePaProGen* for lepton pair production at hadron colliders. *LePaProGen* focuses on the description of higher-order electroweak radiative corrections. The generator is implementing a new algorithm for the selection of the optimal variables for phase-space parametrization.

Computations in Theoretical Physics: Techniques and Methods / 75

New features of AMBRE

Author: Tord Riemann¹

Co-authors: Ievgen Dubovyk¹; Janusz Gluza²; Michal Ochman¹

¹ *DESY*

² *University of Silesia*

Corresponding Author: tord.riemann@desy.de

AMBRE is a Mathematica package for the evaluation of Feynman integrals with the aid of their Mellin-Barnes representations.

The new packages for

(i) a proper and efficient Mellin-Barnes presentation of non-planar topologies

and

(II) an automatic derivation of a subsequent representation by multiple sums are presented.

The latter one allows a linking of packages for automatic summation.

Poster session - Board: 124 / 76

Data Recommender System for the Production and Distributed Analysis System «PanDA»

Author: Mikhail Titov¹

Co-authors: Alexei Klimentov²; Gergely Záruba¹; Kaushik De¹

¹ *University of Texas at Arlington (US)*

² *Brookhaven National Laboratory (US)*

Corresponding Author: mikhail.titov@cern.ch

The Production and Distributed Analysis system (PanDA) is a distributed computing workload management system for processing user analysis, group analysis, and managed production jobs on the grid. The main goal of the recommender system for PanDA is to utilize user activity to build a corresponding model of user interests that can be considered in how data needs to be distributed. As an implicit outcome, the recommender system provides a quantitative assessment of users' potential interest in new data. Furthermore, relying on information about computer centers that are in users' activity zones, it provides an estimated list of computing centers as possible candidates for data storage. As an explicit outcome, the system recommends data collections to users by estimating/predicting the likelihood of user interest in such data.

The proposed recommender system is based on data mining techniques and combines two basic approaches: content-based filtering and collaborative filtering. Each approach has its own advantages, while their combination helps to increase the accuracy of the system. Content-based filtering is focused on creating user profiles based on data features and group of features including corresponding weights that show the significance of features for the user. Collaborative filtering can reveal the similarity between users and between data collections, thus such similarity measure may indicate how "close" objects are, i.e., how close pairs of users' preferences are to each other; or how close data-sets are to those that individual users have used previously.

Information about processed jobs taken from a PanDA database (in this study focusing on data coming from the ATLAS experiment) provides the recommender system with corresponding objects: users and input data (items in terms of recommender systems), and relations between them. This is the minimum required information to build the user-item matrix (i.e. utility matrix, where each element is an implicit rating of item per user). The herein proposed recommender system is not intrusive, i.e., it does not change any part of PanDA system but it can be used as an added-value service to increase efficiency of data management for PanDA.

Poster session - Board: 125 / 77

The Long-Baseline Neutrino Experiment Computing Model and its evolution

Author: Maxim Potekhin¹

¹ *Brookhaven National Laboratory*

Corresponding Author: maxim.potekhin@cern.ch

The Long-Baseline Neutrino Experiment (LBNE) will provide a unique, world-leading program for the exploration of key questions at the forefront of particle physics and astrophysics. Chief among its potential discoveries is that of matter-antimatter symmetry violation in neutrino flavor mixing. To achieve its ambitious physics objectives as a world-class facility, LBNE has been conceived around three central components: an intense, wide-band neutrino beam; a fine-grained near neutrino detector just downstream of the neutrino source; a massive liquid argon time-projection chamber (LArTPC) deployed as a far neutrino detector deep underground, 1300 km downstream. Every stage in the life-cycle of the experiment, from R&D to operations to data analysis, requires the use of sophisticated "physics tools" software as well as robust and efficient software and computing infrastructure to support the work of the many members of LBNE Collaboration, which include more than five hundred scientists in the US and a few countries abroad. In this talk we describe the organization and planning of the LBNE Software and Computing effort, discuss challenges encountered so far and present its evolving Computing Model.

Computations in Theoretical Physics: Techniques and Methods / 78

Automatic numerical integration methods for Feynman integrals through 3-loop

Author: Elise de Doncker¹

Co-authors: Fukuko Yuasa²; Kiyoshi Kato³; Omofalakanmi (Fola) Olagbemi¹; Tadashi Ishikawa²

¹ *Western Michigan University*

² *High Energy Accelerator Research Organization (KEK), Japan*

³ *Kogakuin University*

Corresponding Author: elise.dedoncker@wmich.edu

The paper will include numerical integration results for Feynman loop diagrams through 3-loop such as those covered by (Laporta, 2000). While Laporta generated solutions by solving systems of difference equations, the current methods are based on automatic adaptive integration, using iterated integration with programs from the QuadPack package, or multivariate techniques from the ParInt package. The QuadPack programs have been parallelized with OpenMP for multicore systems. In particular, the Dqags algorithm allows handling boundary singularities of fairly general types. ParInt is a package for multivariate integration layered over MPI (Message Passing Interface), which runs on clusters and incorporates advanced parallel/distributed techniques such as load balancing among processes that may be distributed over a network of nodes. We will give results for 3-loop self-energy diagrams without IR (infra-red) or UV (ultra-violet) singularities, and 2-loop self-energy diagrams with UV terms. The latter can be treated with automatic numerical integration allowing for boundary singularities, and numerical extrapolation. These cases include 2-loop self-energy diagrams with three, four and five internal lines.

Data Analysis - Algorithms and Tools / 79

Pidrix: Particle Identification Matrix Factorization

Author: Evan Sangaline¹

¹ *Michigan State University*

Corresponding Author: esangaline@gmail.com

Probabilistically identifying particles and extracting particle yields are fundamentally important tasks required in a wide range of nuclear and high energy physics analyses. Quantities such as ionization energy loss, time of flight, and Čerenkov angle can be measured in order to help distinguish between different particle species, but distinguishing becomes difficult when there is no clear separation between the measurements for each type of particle. The standard approach in this situation is to model the measurement distributions in some way and then perform fits to the recorded data in order to extract yields and constrain the model parameters. This carries the risk that even very small disagreements between the model and the true distributions can result in significant biases in both the extracted yields and their estimated uncertainties. We propose a new approach to particle identification that does not require the modeling of measurement distributions. It instead relies on the independence of measurement errors between different detectors to pose the problem as one of matrix factorization which is then solved using iterative update equations. This allows for the unsupervised determination of both the measurement distribution and yield of each particle species.

Summary:

An unsupervised learning technique for particle identification, called pidrix, is introduced.

Computing Technology for Physics Research / 80

Planning for distributed workflows: constraint based co-scheduling of computational jobs and data placement in distributed environments.

Authors: Dzmitry Makatun¹; Hana Rudova²; Jerome LAURET³; Michal Sumbera¹

¹ Nuclear Physics Institute (CZ)

² Masaryk University (CZ)

³ BROOKHAVEN NATIONAL LABORATORY

Corresponding Author: d.i.makatun@gmail.com

When running data intensive applications on distributed computational resources long I/O overheads may be observed when access to remotely stored data is performed. Latencies and bandwidth can become the major limiting factors for the overall computation performance and reduce the application's CPU/WallTime ratio due to excessive IO wait. For this reason further optimization of data management may imply increasing availability of data "closer" to the computational task and by then, reducing the overheads due to data access over long distance on the Grid. This ideal is in high demand in data intensive computational fields such as the ones in the HENP communities.

In previous collaborative work of BNL and NPI/ASCR, we addressed the problem of efficient data transferring in a Grid environment and cache management. The transfer considered an optimization of moving data at N sites while the data may be located at M locations. However, the co-scheduling of data placement and processing was not (yet) approached as the hard problem was decomposed into simpler tasks.

Leveraging the knowledge of our previous research, we propose a constraint programming based planner that schedules computational jobs, data placement (transfers) in a distributed environment in order to optimize resource utilization and reduce the overall processing completion time. The optimization is achieved by ensuring that none of the resources (network links, data storages and CPUs) are over-saturated at any moment of time and either (a) that the data is pre-placed at the site where the job runs or (b) that the jobs are scheduled where the data is already present. Such an approach would eliminate the idle CPU cycles and would have wide application in the community.

In this talk, we will present the theoretical model behind our planner. We will further present the results of simulation based on input data extracted from log files of batch and data-management systems of experiment's STAR computation facility.

Poster session - Board: 126 / 81

The Linear Collider Software Framework

Author: Andre Sailer¹

Co-authors: Akiya Miyamoto²; Frank-Dieter Gaede³; Norman Anthony Graf⁴

¹ CERN

² High Energy Accelerator Research Organization (JP)

³ Deutsches Elektronen-Synchrotron (DE)

⁴ *SLAC National Accelerator Laboratory (US)*

Corresponding Author: andre.philippe.sailer@cern.ch

For the future experiments at linear electron-positron colliders (ILC or CLIC), detailed physics and detector optimisation studies are taking place in the CLICdp, ILD, and SiD groups. The physics performance of different detector geometries and technologies has to be estimated realistically. These assessments require sophisticated and flexible full detector simulation and reconstruction software. At the heart of the linear collider detectors lies the particle flow reconstruction requiring the combination of fine-grained calorimeters and advanced clustering software.

The similarities between the different detector concepts allow for the use of common software tools. All the concepts share an event data and persistency format which enables the sharing of files and applications across the concepts. Particle flow clustering, vertexing and flavour tagging is already provided by stand alone packages via lightweight interfaces. In the near future the geometry information for all detector layouts will be provided by a unique source for the simulation and reconstruction programs, providing further re-use of software between the collaborations. In addition a track reconstruction package is currently under development.

The sharing and development of flexible software tools not only saves precious time and resources. Using common tools for different detectors also helps to uncover bugs or inefficiencies that would be harder to spot without multiple users. The concept of generic software tools and some of the programs themselves can be beneficial to experiments beyond the linear collider community.

Plenary / 82

Event simulation for colliders - A basic overview

Author: Christian Reuschle¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: christian.reuschle@kit.edu

In this article we will discuss the basic calculational concepts to simulate particle physics events at high energy colliders. We will mainly focus on the physics in hadron colliders and particularly on the simulation of the perturbative parts, where we will in turn focus on the next-to-leading order QCD corrections.

Poster session - Board: 127 / 83

High-speed zero-copy data transfer for DAQ applications

Author: Flavio Pisani¹

Co-authors: Daniel Hugo Campora Perez ²; Niko Neufeld ²

¹ *Universita e INFN, Roma I (IT)*

² CERN

Corresponding Authors: niko.neufeld@cern.ch, flavio.pisani@cern.ch

The LHCb Data Acquisition (DAQ) will be upgraded in 2020 to a trigger-free readout. In order to achieve this goal we will need to connect 500 nodes with a total network capacity of 40 Tb/s. To get such an high network capacity we are testing zero-copy technology in order to maximise the theoretical link throughput without adding excessive CPU and memory bandwidth overhead, leaving free resources for data processing. More CPU power available means less machines needed for accomplishing the same task resulting in less power, space and money used for the same result.

We had developed two test applications one using non zero-copy protocols (TCP/UDP) and the other using the OFED libverbs API, which can provide low level access and high throughput. The libverbs API offers a good level of flexibility allowing the application to be compatible with different RDMA solutions, like Infiniband and Internet Wide Area RDMA Protocol (iWARP), and it provides us the possibility to perform tests on different technologies using the same application for a more comprehensive evaluation of different implementations of an RDMA protocol over different network technologies.

We will present throughput, CPU and memory overhead measures comparing Infiniband and 40 GbE solutions using RDMA, those measures will be present for several network configurations to test the scalability of the system.

The comparison between zero-copy and non zero-copy results will be presented to evaluate the impact of high speed Ethernet communication (40 Gb/s now and 100 Gb/s then) on the host machine in terms of CPU and memory usage. These results are relevant to wide range of high-speed, low-cost PC based data-acquisition.

Data Analysis - Algorithms and Tools / 84

A novel robust and efficient algorithm for charge particle tracking in high background flux.

Authors: Alessio Del Dotto¹; Cristiano Fanelli²; Evaristo Cisbani³

¹ *Università di Roma Tre and INFN, Roma, Italy*

² *INFN Sezione di Roma, Università di Roma 'La Sapienza', Roma, Italy*

³ *INFN, Sezione di Roma, gruppo Sanità and Istituto Superiore di Sanità, I-00161 Rome, Italy*

Corresponding Author: fanellic@roma1.infn.it

A new tracker based on the GEM technology is under development for the upcoming experiments in Hall A at Jefferson Lab, where a longitudinally polarized electron beam of 11 GeV, combined with innovative polarized targets, will provide luminosity up to $10^{39}/(\text{s cm}^2)$ opening exciting opportunities to investigate unexplored aspects of the inner structure of the nucleon and the dynamics of its constituents.

At this luminosity, the expected background flux, mainly due to low energy (~ 1 MeV) photons, is up to 500 MHz/cm^2 which generates about 200 kHz/cm^2 hits in each tracker chamber.

In such a context, an efficient, computational time effective and precise tracks reconstruction is mandatory.

We propose a novel algorithm based on a Hopfield neural network (NN) combined to filter techniques.

A preliminary clustering of the GEM hits exploits all spatial and timing information of the acquired signals coming from the GEM strips, to maximally reduce the data to be processed.

The NN, within a mean field theory framework, provides a robust association of the GEM hits, (then drastically reducing the potential hit combinations),

whereas a Kalman filter associated to a Rauch-Tung-Striebel smoother, is used for final accurate reconstruction.

Results of the first tests on simulated and real data will be presented as well as a description of the method and of its original aspects.

Poster session - Board: 206 / 85

Designing and recasting LHC analyses with MadAnalysis 5

Authors: Benjamin Fuks¹; Eric Conte²

¹ CERN

² GRPHE

Corresponding Author: eric.conte@uha.fr

The LHC experiments are currently pushing limits on new physics to a further and further extent. The interpretation of the results in the framework of any theory however relies on our ability to accurately simulate both signal and background processes. This task is in general achieved by matching matrix-element generator predictions to parton showering, and further employing hadronization and fast detector simulation algorithms. Phenomenological analyses can in this way be performed at several levels of the simulation chain, i.e., at the parton-level, after hadronization or after detector simulation.

This talk focuses on MadAnalysis 5, a unique analysis package dedicated to phenomenological investigations to be achieved at any step of the simulation chain. Within this framework, users are invited, through a user-friendly Python interpreter, to perform physics analyses in a very simple manner. An associated C++ code is then automatically generated, compiled and executed. Very recently, the expert mode of MadAnalysis 5 has been extended so that the notion of signal/control regions is now handled and additional observables are included. In addition, the program features an interface to several fast detector simulation packages, one of them being an optimized tune of the Delphes 3 package. As a consequence, it is now possible to easily recast existing CMS or ATLAS analyses within the MadAnalysis 5 framework. Finally, the new release of the program is more platform-independent and benefits from the graphical components of GnuPlot, Matplotlib and ROOT.

Poster session - Board: 207 / 86

Going Visual

Author: Vasil Georgiev Vasilev¹

¹ CERN

Corresponding Author: vasil.georgiev.vasilev@cern.ch

Since the silicon era, programming languages thrive: assembler, macro assembler, Fortran, C, C++, LINQ. A common characteristic between the generations is the level of abstraction. While assembly languages didn't provide abstractions, macro assemblers, Fortran, C and C++ promised to improve the deficiencies of the abstractions of the older ones. The increasing popularity of the domain specific languages showed that a single textual (ascii) language cannot supply convenient concepts, necessary for multidisciplinary frameworks.

In many situations, the offered details by C++ language and ROOT become a burden for users in their every day job. One approach to hide some of the details is to provide a multistage interface layers. This allows rich graphical user interface (GUI) to be build on top of them, turning framework's GUI into a domain specific programming language. In this paper we present a few modern technologies, which helped to reduce the complexity of producing simulation and analysis algorithms

in various domains. Mashup technologies such as Yahoo Pipes, Presto Wires, OpenWires and visual programming languages such as LabVIEW, KNIME, VESPA solved a wide variety of problems. We argue for a mixed approach of using visual and textual algorithm design and implementation. We outline a methodology of common steps typically taken in data analysis. The work discusses the advantages and disadvantages going visual at every step of the data analysis. We give insights and scenarios, where going visual in the software development is more optimal in field of high energy physics.

Computing Technology for Physics Research / 87

Native Language Integrated Queries with CppLINQ in C++

Author: Vasil Georgiev Vasilev¹

¹ CERN

Corresponding Author: vasil.georgiev.vasilev@cern.ch

Programming language evolution brought to us the domain-specific languages (DSL). They proved to be very useful for expressing specific concepts, turning into a vital ingredient even for generic-purpose frameworks. Supporting declarative DSLs (such as SQL) into imperative languages (such as C++) can happen in the manner of language integrated query (LINQ).

We approach to integrate a similar to LINQ programming language, native to C++. We review its usability in the context of the high energy physics. We present examples using CppLINQ for many common workflows done by the end-users doing data analysis and simultaion. We discuss evidences how this DSL technology can simplify massively parallel grid system such as PROOF.

Computing Technology for Physics Research / 88

Modernising ROOT: Building Blocks for Vectorised Calculations

Authors: Danilo Piparo¹; Lorenzo Moneta¹; Sandro Christian Wenzel¹

¹ CERN

Corresponding Authors: sandro.wenzel@cern.ch, danilo.piparo@cern.ch

The evolution of the capabilities offered by modern processors in the field of vectorised calculus has been steady in the recent years. Vectorisation is indeed of capital importance to increase the throughput of scientific computations (e.g. for Biology, Theory, High Energy and Solid State Physics), especially in presence of the well known cpu clock frequency stagnation. On the other hand, the expression of parallelism through vectorisation is not straightforward, especially in popular languages like C++.

This is the reason why the ROOT toolkit recently added to its rich collection of mathematical tools, components to ease vectorisation like the VC and VDT libraries. In this paper we present the potential of these novel components with documented examples coming from areas like the simulation of passage of particles through matter, random number generation, reconstruction of HEP events and statistical calculations. We also show their integration and synergy with the existing mathematical routines and classes traditionally offered by ROOT.

Computing Technology for Physics Research / 89

Recent Developments in the CVMFS Server Backend

Authors: Jakob Blomer¹; Rene Meusel¹

Co-authors: Gerardo Ganis¹; Predrag Buncic¹; Seppo Heikkila¹

¹ CERN

Corresponding Author: rene.meusel@cern.ch

The CernVM-File System (CVMFS) is a snapshotting read-only file system designed to deliver software to grid worker nodes over HTTP in a fast, scalable and reliable way. In recent years it became the de-facto standard method to distribute HEP experiment software in the WLCG and starts to be adopted by other grid computing communities outside HEP.

This paper focusses on the recent developments of the CVMFS Server, the central publishing point of new file system snapshots. Using a union file system, the CVMFS Server allows for direct manipulation of a (normally read-only) CVMFS volume with copy-on-write semantics. Eventually the collected changeset is transformed into a new CVMFS snapshot, constituting a transactional feedback loop.

The generated repository data is pushed into a content addressable storage requiring only a RESTful interface and gets distributed through a hierarchy of caches to individual grid worker nodes. Besides practically all POSIX-compliant file systems, lately CVMFS allows to use highly-scalable key-value storage systems through the Amazon S3 API.

Additionally we describe recent features, such as file chunking, repository garbage collection, fast replication and file system history that enable CVMFS for a wider range of use cases.

Data Analysis - Algorithms and Tools / 90

Clad - Automatic Differentiation Using Cling/Clang/LLVM

Author: Vasil Georgiev Vasilev¹

Co-author: Lorenzo Moneta¹

¹ CERN

Corresponding Author: vasil.georgiev.vasilev@cern.ch

Differentiation is ubiquitous in high energy physics, for instance for minimization algorithms in fitting and statistical analysis, detector alignment and calibration, theory. Automatic differentiation (AD) avoids well-known limitations in round-offs and speed, which symbolic and numerical differentiation suffer from, by transforming the source code of functions.

We will present how AD can be used to compute the gradient of multi-variate functions and functor objects. We will explain approaches to implement an AD tool. We will show how LLVM, Clang and Cling (ROOT's C++11 interpreter) simplifies creation of such a tool. We describe how the tool could be integrated within any framework. We will demonstrate a simple proof-of-concept prototype, called clad, which is able to generate n-th order derivatives of C++ functions and other language constructs. We also demonstrate how clad can offload laborious computations from the CPU using OpenCL.

Computing Technology for Physics Research / 91

Techniques and tools for measuring energy efficiency of scientific software applications

Authors: David Abdurachmanov¹; Filip Nybäck Filip²; Giulio Eulisse³; Goncalo Marques Pestana⁴; Jukka Nurminen²; Peter Elmer⁵; Robert Knight⁶; Tapio Petteri Niemi⁷; Zhonghong Ou²

¹ *Vilnius University (LT)*

² *Aalto University*

³ *Fermi National Accelerator Lab. (US)*

⁴ *H*

⁵ *Princeton University (US)*

⁶ *Princeton University*

⁷ *Helsinki Institute of Physics (FI)*

Corresponding Authors: goncalo.marques.pestana@cern.ch, giulio.eulisse@cern.ch, david.abdurachmanov@cern.ch

As both High Performance Computing (HPC) and High Throughput Computing (HTC) are sensitive to the rise of energy costs, energy-efficiency has become a primary concern in scientific fields such as High Energy Physics (HEP). There has been a growing interest in utilizing low power architectures, such as ARM processors, to replace traditional Intel x86 architectures. Nevertheless, even though such solutions have been successfully used in mobile applications with low I/O and memory demands, it is still unclear if they are suitable and more energy-efficient in the scientific computing environment. Furthermore, there is still lack of tools to derive and compare power consumption for these types of workloads, and eventually to support software optimizations for energy efficiency.

To that end, we have performed several physical and software-based measurements of workloads from CERN running on ARM and Intel architectures, to compare their power consumption and performance. We leverage several profiling tools to extract different aspects of the experiments, including hardware usage and software characteristics. We report the results of these measurements and the experience gained in developing a set of measurement techniques and profiling tools to accurately assess the power consumption for scientific workloads.

Computing Technology for Physics Research / 92

Heterogeneous High Throughput Scientific Computing with ARMv8 64-bit and Xeon Phi

Authors: Brian Paul Bockelman¹; David Abdurachmanov²; Giulio Eulisse³; Peter Elmer⁴; Robert Knight⁵

¹ *University of Nebraska (US)*

² *Vilnius University (LT)*

³ *Fermi National Accelerator Lab. (US)*

⁴ *Princeton University (US)*

⁵ *Princeton University*

Corresponding Authors: david.abdurachmanov@cern.ch, peter.elmer@cern.ch

Electrical power requirements will be a constraint on the future growth of Distributed High Throughput Computing (DHTC) techniques as used in High Energy Physics. Performance-per-watt is a critical metric for the evaluation of computer architecture for cost-efficient computing. Additionally, the future performance growth comes from

heterogeneous, many-core, and high computing density platforms with specialized processors. In this paper, we examine the Intel Xeon Phi Many Integrated Cores (MIC) co-processor and Applied Micro's XGene ARMv8 64-bit low-power server system-on-a-chip (SoC) solutions for scientific computing applications.

We report our experience on software porting, performance and energy efficiency and evaluate the potential for use of such technologies in the context of the distributed computing systems such as the Worldwide LHC Computing Grid (WLCG).

Computing Technology for Physics Research / 93

Implementation of a multi-threaded framework for large-scale scientific applications

Author: David Lange¹

Co-authors: Christopher Jones²; Elizabeth Sexton-Kennedy²; W.David Dagenhart²

¹ *Lawrence Livermore Nat. Laboratory (US)*

² *Fermi National Accelerator Lab. (US)*

Corresponding Authors: sexton@fnal.gov, david.lange@cern.ch

The CMS experiment has recently completed the development of a multi-threaded capable application framework. In this presentation, we will discuss the design, implementation and application of this framework to production applications in CMS. For the 2015 LHC run, this functionality is particularly critical for both our online and offline production applications, which depend on faster turn-around times and a reduced memory footprint relative than before. These applications are complex codes, each including a large number of physics-driven algorithms. While the framework is capable of running a mix of thread-safe and "legacy" modules, algorithms running in our production applications need to be thread-safe for optimal use of this multi-threaded framework at a large scale. Towards this end, we discuss the types of changes which were necessary for our algorithms to achieve good performance of our multi-threaded applications in a full-scale application.

Data Analysis - Algorithms and Tools / 94

Simulation Upgrades for the CMS experiment

Author: David Lange¹

Co-authors: Ianna Osborne²; Mike Hildreth³; Vladimir Ivantchenko⁴

¹ *Lawrence Livermore Nat. Laboratory (US)*

² *Fermi National Accelerator Lab. (US)*

³ *University of Notre Dame (US)*

⁴ *CERN*

Corresponding Author: david.lange@cern.ch

Over the past several years, the CMS experiment has made significant changes to its detector simulation application. The geometry has been generalized to include modifications being made to the CMS detector for 2015 operations, as well as model improvements to the simulation geometry of the current CMS detector and the implementation of a number of approved and possible future detector configurations. These include both completely new tracker and calorimetry systems. We have

completed the transition to Geant4 version 10, we have made significant progress in reducing the CPU resources required to run our Geant4 simulation. These have been achieved through both technical improvements and through numerical techniques. Substantial speed improvements have been achieved without changing the physics validation benchmarks that the experiment uses to validate our simulation application for use in production. In this presentation, we will discuss the methods that we implemented and the corresponding demonstrated performance improvements deployed for our 2015 simulation application.

Computing Technology for Physics Research / 95

WLCG Tier-2 site in Prague: a little bit of history, current status and future perspectives

Authors: Dagmar Adamova¹; Jiri Chudoba²

Co-authors: Jan Svec²; Lukas Fiala²; Marek Elias²; Milos Lokajicek²; Tomas Kouba³

¹ *NPI AS CR Prague/Rez*

² *Institute of Physics AS CR Prague*

³ *CERN/Institute of Physics AS CR Prague*

Corresponding Author: dagmar.adamova@cern.ch

The High energy physics is one of the research areas where the accomplishment of scientific results is inconceivable without a complex distributed computing infrastructure. This includes also the experiments at the Large Hadron Collider (LHC) at CERN where the production and analysis environment is provided by the Worldwide LHC Computing Grid (WLCG). A very important part of this system is represented by sites classified as Tier-2s: they deliver a half of the computing and disk storage capacity of the whole WLCG. In this contribution we present an overview of the Tier-2 site praguec2 in Prague, the largest site in the Czech republic providing computing and storage services for particle physics experiments. A brief history flashback, current status report and future perspectives of the site will be presented.

Computing Technology for Physics Research / 96

STAR Online Framework: from Metadata Collection to Event Analysis and System Control

Authors: Dmitry Arkhipkin¹; Jerome LAURET²

¹ *Brookhaven National Laboratory*

² *BROOKHAVEN NATIONAL LABORATORY*

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

In preparation for the new era of RHIC running (RHIC-II upgrades and possibly, the eRHIC era), the STAR experiment is expanding its modular Message Interface and Reliable Architecture framework (MIRA). MIRA allowed STAR to integrate meta-data collection, monitoring, and online QA components in a very agile and efficient manner using a messaging infrastructure approach. In this paper, we will briefly summarize our past achievements, provide an overview of the recent development activities focused on messaging patterns and describe our experience with the complex event processor (CEP) recently integrated into the MIRA framework. CEP was used in the recent RHIC Run 14 which provided practical use cases. Finally, we will present our requirements and expectations for the planned expansion of our systems, which will allow our framework to acquire features typically associated with Detector Control Systems. Special attention will be given to aspects related

to latency, scalability and interoperability within heterogeneous set of services, various data and metadata acquisition components coexisting in STAR online domain.

Data Analysis - Algorithms and Tools / 97

ROOT 6

Author: Axel Naumann¹

¹ CERN

Corresponding Author: axel.naumann@cern.ch

The recently published ROOT 6 is the first major ROOT release since nine years. It opens a whole world of possibilities with full C++ support at the prompt and a built-in just-in-time compiler, while staying almost completely backward compatible. The ROOT team has started to make use of these new features, offering for instance an improved new implementation of TFormula, fast and type-safe access to TTrees, and access to native, runtime-compiled functions through its new interpreter cling. Many other, unrelated new features will be presented, too, for instance a better engine for formula layout, Latex export, support for transparency and shading, and much improved graphics interactivity.

Poster session - Board: 208 / 99

Particle-level pileup subtraction for jets and jet shapes

Author: Peter Berta¹

¹ Charles University (CZ)

Corresponding Author: peter.bera@cern.ch

The ability to correct jets and jet shapes for the contributions of multiple uncorrelated proton-proton interactions (pileup) largely determines the ability to identify highly boosted hadronic decays of W, Z, and Higgs bosons, or top quarks. We present a new method that operates at the level of the jet constituents and provides both performance improvement and simplification compared to existing methods. Comparisons of the new method with existing methods along with predictions of the impact of pileup on jet observables during the LHC Run II will be presented. We will also discuss methods that may remove pileup contribution from the whole event.

Poster session - Board: 209 / 100

Traditional Tracking with Kalman Filter on Parallel Architectures

Authors: Avi Yagil¹; Daniel Sherman Riley²; Frank Wuerthwein¹; Giuseppe Cerati¹; Ian Macneill¹; Kevin McDermott²; Matevz Tadel¹; Peter Elmer³; Peter Wittich²; Steven Lantz⁴

¹ Univ. of California San Diego (US)

² Cornell University (US)

³ Princeton University (US)

⁴ Cornell University

Corresponding Authors: david.abdurachmanov@cern.ch, peter.elmer@cern.ch

Power density constraints are limiting the performance improvements of modern CPUs. To address this we have seen the introduction of lower-power, multi-core processors, but the future will be even more exciting. In order to stay within the power density limits but still obtain Moore's Law performance/price gains, it will be necessary to parallelize algorithms to exploit larger numbers of lightweight cores and specialized functions like large vector units. Example technologies today include Intel's Xeon Phi and GPGPUs.

Track finding and fitting is one of the most computationally challenging problems for event reconstruction in particle physics. At the High Luminosity LHC, for example, this will be by far the dominant problem. The need for greater parallelism has driven investigations of very different track finding techniques including Cellular Automata or returning to Hough Transform techniques originating in the days of bubble chambers. The most common track finding techniques in use today are however those based on the Kalman Filter. Significant experience has been accumulated with these techniques on real tracking detector systems, both in the trigger and offline. They are known to provide high physics performance, are robust and are exactly those being used today for the design of the tracking system for HL-LHC. We report the results of our investigations into the potential and limitations of these algorithms on the new parallel hardware.

Data Analysis - Algorithms and Tools / 101

Combination of multivariate discrimination methods in the measurement of the inclusive top pair production cross section

Author: Jiri Franc¹

¹ *Czech Technical University in Prague*

Corresponding Author: jiri.franc@fifi.cvut.cz

The application of multivariate analysis techniques in experimental high energy physics have been accepted as one of the fundamental tools in the discrimination phase, when signal is rare and background dominates. The purpose of this study is to present new approaches to the variable selection based on phi-divergences, together with various statistical tests, and the combination of new applied MVA methods together with familiar ROOT TMVA methods in the real data analysis. The results and quality of separation of the Generalized Linear Models (GLM), Gaussian Mixture Models (GMM), Neural Networks with Switching Units (NNSU), TMVA Boosted Decision Trees, and Multi-layer Perceptron (MLP) in the measurement of the inclusive top pair production cross section employing $D0$ Tevatron full RunII data ($9.7 fb^{-1}$) will be presented. Possibilities of improvement in discrimination will be discussed.

Data Analysis - Algorithms and Tools / 102

High-resolution deconvolution methods for analysis of low amplitude noisy gamma-ray spectra

Author: Vladislav Matoušek¹

Co-author: Miroslav Morhac ¹

¹ *Institute of Physics, Slovak Academy of Sciences*

Corresponding Author: vladislav.matousek@savba.sk

The deconvolution methods are very efficient and widely used tools to improve the resolution in the spectrometric data. They are of great importance mainly in the tasks connected with decomposition of low amplitude overlapped peaks (multiplets) in the presence of noise.

In the talk we will present a set of deconvolution algorithms and a study of their decomposition capabilities from the resolution point of view. We have proposed improvements in the efficiency of the iterative deconvolution methods by introducing further modifications into deconvolution process, e.g. noise suppression operations during iterations and improved blind deconvolution methods. We will illustrate their suitability for processing of noisy spectrometric data. It will be shown, that using the new developed algorithms we are able to improve the resolution in spectrometric data. The methods are able better detect hidden peaks in the noisy gamma-ray spectra and decompose the overlapped peaks by concentrating the peak areas into a few channels. The efficiency of the above mentioned algorithms and their comparison will be presented also.

Computing Technology for Physics Research / 103

Evolution of the ATLAS Software Framework towards Concurrency

Authors: Benjamin Michael Wynne¹; Charles Leggett²; Graeme Andrew Stewart³; Roger Jones⁴

¹ *University of Edinburgh (GB)*

² *Lawrence Berkeley National Lab. (US)*

³ *University of Glasgow (GB)*

⁴ *Lancaster University (GB)*

Corresponding Authors: roger.jones@cern.ch, graeme.andrew.stewart@cern.ch

The ATLAS experiment has successfully used its Gaudi/Athena software framework for data taking and analysis during the first LHC run, with billions of events successfully processed. However, the design of Gaudi/Athena dates from early 2000 and the software and the physics code has been written using a single threaded, serial design. This programming model has increasing difficulty in exploiting the potential of current CPUs, which offer their best performance only through taking full advantage of multiple cores and wide vector registers. Future CPU evolution will intensify this trend, with core counts increasing and memory per core falling. Maximising performance per watt will be a key metric, so all of these cores must be used as efficiently as possible.

In order to address the deficiencies of the current framework, ATLAS has embarked upon two projects: first, a practical demonstration of the use of multi-threading in our reconstruction software, using the GaudiHive framework; second, an exercise to gather requirements for an updated framework, going back to the first principles of how event processing occurs.

In this paper we report on both these aspects of our work. For the hive based demonstrators, we discuss what changes were necessary in order to allow the serially designed ATLAS code to run, both to the framework and to the tools and algorithms used. We report on what general lessons were learned about the code patterns that had been employed in the software and which patterns were identified as

particularly problematic for multi-threading. These lessons were fed into our considerations of a new framework and we present preliminary conclusions on this work. In particular we identify areas where the framework can be simplified in order to aid the implementation of a concurrent event processing scheme. Finally, we discuss the practical difficulties involved in migrating a large established code base to a multi-threaded framework and how this can be achieved for LHC Run 3.

Poster session - Board: 210 / 104

Deep learning neural networks - a TMVA perspective

Author: Eckhard Von Torne¹

Co-authors: Andreas Hoecker²; Helge Voss³; Jan Therhaag⁴; Joerg Stelzer²; Peter Speckmayer

¹ *Universitaet Bonn (DE)*

² *CERN*

³ *Max-Planck-Gesellschaft (DE)*

⁴ *Univ. Bonn, Physikalisches Institut-Unknown-Unknown*

Corresponding Author: evt@physik.uni-bonn.de

Deep learning neural networks are feed-forward networks with several hidden layers. Due to their complex architecture, such networks have been successfully applied in several difficult non-HEP applications such as face recognition. Recently the application of such networks has been explored in the context of particle physics. We discuss the construction and training of such neural nets within the Toolkit for MultiVariate Analysis (TMVA) and present recent improvements of TMVA in that field.

Plenary / 105

Fast detector simulation and the GeantV project

Author: Andrei Gheata¹

¹ *CERN*

Corresponding Author: andrei.gheata@cern.ch

Particle transport Monte Carlo simulation has a fundamental role in High Energy and Nuclear Physics (HENP) experiments. It enables an experiment's designers to predict its measurement potential, and to disentangle detector effects from the physics signal.

High-energy physics detector simulation is increasingly relied upon due to the increasing complexity of the experimental setups, which scales with the number of sub-detectors and analysed channels. The first LHC run and the corresponding arrival of the GRID era boosted the production of such simulations to an unprecedented scale, with each experiment simulating billions of events using full detailed simulation.

This revealed both the power of the state of the art physics embedded in current detailed simulation models, as well as important shortcoming of throughput with respect to the increasing demands of simulated data samples.

The talk will review the ongoing efforts of the community to develop fast detector simulation applications, which tend to cluster into frameworks. This trend justifies the R&D of more generic solutions to either improve the performance of the traditional simulation tools by integrating fast simulation components, or make use of modern computing techniques allowing to increase the throughput. The second part will describe how this is being addressed by the GeantV framework, the current status, lessons and challenges faced by the project.

Opening / 106

Welcome

Author: Denis Perret-Gallix¹

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

Corresponding Author: denis.perret-gallix@cern.ch

Opening / 107

Czech Technical University

Author: Vojtech Petracek¹

¹ *Czech Technical University (CZ)*

Corresponding Author: vojtech.petracek@cern.ch

Opening / 108

Prague municipality

Opening / 109

Charles University

Opening / 110

Nuclear Physics Institute AS CR

Author: Petr Lukas^{None}

Opening / 111

Institute of Pysics AS CR

Author: Jan Ridky^{None}

Opening / 112

Czech Science and education, financing, physics

Author: Jiri Chyla¹

¹ *Acad. of Sciences of the Czech Rep. (CZ)*

Corresponding Author: chyla@fzu.cz

Plenary / 113

Logistics

Author: Tomas Davidek¹

¹ *Charles University (CZ)*

Corresponding Author: tomas.davidek@cern.ch

Poster session - Board: 128 / 114

CMS Software Documentation System

Author: Ali Mehmet Altundag¹

¹ *Cukurova University (TR)*

Corresponding Author: ali.mehmet.altundag@cern.ch

CMS Software is huge software development project with a large amount of source code. In large scale and complex projects, it is important to have a software documentation system as possible. The core of the documentation should be version-based and available online with the source code. CMS uses Doxygen and Twiki as main tools to provide automated and non-automated documentation. Both of them are heavily cross-linked to prevent duplication of information. Doxygen software documentation tool is used to generate documentation with UML graphs. This note describes the design principles, the basic functionalities and the technical implementations of the CMSSW documentation.

Plenary / 115

THE POTENTIAL OF GENOMIC MEDICINE—THE MANY ASPECTS OF COMPUTING ARE DEMANDING AND THE COMPLEXITY MUST BE MANAGED.

Author: Janis Landry-Lane¹

¹ IBM

Corresponding Author: martina_kocourkova@cz.ibm.com

Janis will speak about Building an architected genomics pipeline platform that will extend to support the analytics, data management, data provenance, long-term retention, and especially the issues as we genomics becomes part of the clinical information for patients.

High-performance best practices in computing and storage solutions are required to process the data produced by Next Generation Sequencing that is doubling every five months. Of the four phases in a sequencing project: a. Experimental design and sample collection, b. Sequencing c. Data management and d. Downstream analysis, IBM has optimized the data management aspect. It takes a highly optimized HPC platform to keep pace with the genomic data analysis. The algorithms are typically I/O intensive. Additionally, the genomics pipeline workflow must be optimized in order to fully utilize the available resources. The data must typically be archived so that it stored cost effectively and, in the cases of clinical genomics, for many years. IBM, working with our life cycle management and hierarchical storage management will address the requirement for long-term data archive.

The primary goal of a sequencing project is to use the data for extensive downstream analytics with corresponding phenotypic information, image analysis, published scientific discovery, and the use of other internal and external data sources, allowing researchers to obtain insights. IBM will address the work being done to integrate genomics data into translational platforms, with the goal of personalized medicine.

Computations in Theoretical Physics: Techniques and Methods / 116

General formulation of the sector improved residue subtraction

Author: David Heymes¹

¹ RWTH Aachen

Corresponding Author: dheymes@physik.rwth-aachen.de

The main theoretical tool to provide precise predictions for scattering cross sections of strongly interacting particles is perturbative QCD. Starting at next-to-leading order (NLO) the calculation suffers from unphysical IR-divergences that cancel in the final result. At NLO there exist general subtraction algorithms to treat these divergences during a calculation. Since the LHC demands for more precise theoretical predictions, general subtraction methods at next-to-next-to-leading order (NNLO) are needed.

This talk is about the four-dimensional formulation of the sector improved residue subtraction. I explain how the subtraction scheme STRIPPER can be extended to arbitrary multiplicities. Therefore, it furnishes a general framework for the calculation of NNLO cross sections in perturbative QCD.

Plenary / 117

The interrelation between stability of Vacuum (absence of new physics up to Plank scale) and value of running mass of top-quark: problems and perspectives.

Author: Mikhail Kalmykov¹

¹ *II. Institut für Theoretische Physik, Universität Hamburg*

Corresponding Author: kalmykov.mikhail@gmail.com

After the discovery of the Higgs boson - the last important building block of the Standard Model (SM) required by its renormalizability - and the still missing direct detection of new physics beyond SM at the LHC, the self-consistency of the SM has attracted a lot of notice. One of the approaches to determine the scale at which SM may break down is based on the renormalization group (RG) analysis of the SM running couplings, specifically of the Higgs self-coupling and the question whether it stays positive up to the Planck scale which would imply the vacuum to remain stable.

Recently, the detailed renormalization group analysis of the stability of vacuum was done by a few groups with common conclusion that the current values of Higgs boson and top-quark masses imply that our vacuum is metastable. It was observed, that higher order radiative corrections to RG equations and matching conditions as well as the numerical value of top-quark mass play an important role in this analysis. The small variation of the top-quark mass gives rise to stability of vacuum so that new physics does not appear up to Planck scale, there is no Landau problem for Higgs self-coupling in this case, does not exist the hierarchy problem in Standard Model, and even more, the Higgs of Standard Model can serve as inflaton.

One of the manifestations of stable-vacuum scenario is the large EW contribution to the value of running mass of top-quark so that the sum of QCD and EW contribution are almost perfectly canceled. We treat this effect as indication of importance of two-loop electroweak radiative corrections which play the same role as three- and/or four-loop QCD corrections to any physical processes.

Plenary / 118

Modern messaging solutions for distributed applications

Author: Luca Magnoni¹

¹ *CERN*

Corresponding Author: luca.magnoni@cern.ch

Modern software applications rarely live in isolation and nowadays it is common practice to rely on services or consume information provided by remote entities. In such a distributed architecture, integration is key.

Messaging, for more than a decade, is the reference solution to tackle challenges of a distributed nature, such as network unreliability, strong-coupling of producers and consumers and the heterogeneity of applications. Thanks to a strong community and a common effort towards standards and consolidation, message brokers are today the transport layer building blocks in many projects and services, both within the physics community and outside. Moreover, in recent years, a new generation of messaging services has appeared, with a focus on low-latency and high-performance use cases, pushing the boundaries of messaging applications.

This talk will present messaging solutions for distributed applications going through an overview of the main concepts, technologies and services.

Plenary / 119

NLO calculations for high multiplicity processes

Author: Daniel Maitre¹

¹ *University of Durham (GB)*

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

In this contribution I will review the state-of-the-art for NLO calculations and expose the current computational challenges faced in computing these predictions.

Plenary / 120

High performance computing based on mobile processors

Author: Filippo Mantovani¹

¹ *Barcelona Supercomputing Center*

In the late 1990s, (mostly) economic reasons led to the adoption of commodity desktop processors in high-performance computing. This transformation has been so effective that in 2014 the TOP500 list is still dominated by x86-based computers.

More recently, around 2005-2008, always for economic/market reasons commodity GPUs became interesting and powerful enough devices to be used as coprocessors in high-performance computing. Heterogeneous computing based on CPU+coprocessors has evolved till today where the first entry of TOP500 is an heterogeneous system.

The story tells us that a “technological circle” moves innovation from HPC to commodity market and back to HPC in the moment in which commodity market makes devices enough cost/compute effective.

In 2013, the largest commodity market in computing is not PCs or servers, but mobile computing, comprising smartphones and tablets, most of which are built with ARM-based SoCs. This leads to the guess that, once mobile SoCs deliver sufficient performance, mobile SoCs can help reduce the cost of HPC.

In view of the experiences within the Mont-Blanc project at the Barcelona Supercomputing Center, this talk will describe possibilities and challenges involved in developing a high-performance computing platform from low cost and energy efficient mobile processors and commodity components.

Plenary / 121

Memory Models in HPC: Successes and Problems

Author: Antun Balaz¹

¹ *Institute of Physics Belgrade*

Performance and scaling of many algorithms in scientific and distributed computing crucially depend on the memory model, organization and speed of access to the memory hierarchy. We will review parallel computing strategies and memory models in modern HPC, discuss MPI and OpenMP parallelization paradigms, as well as hybrid programming approach and memory mismatch problem. We will comment on the emerging hybrid programming hierarchy due to introduction of GPGPUs, FPGAs and other types of accelerators.

Computations in Theoretical Physics: Techniques and Methods / 122

Cosmic ray propagation with CRPropa 3

Author: David Walz¹

Co-authors: Arjen van Vliet²; Carmelo Evoli²; Daniel Kuempel¹; Gero Mueller¹; Guenter Sigl²; Karl-Heinz Kampert³; Martin Erdmann¹; Rafael Alves Batista²; Tobias Winchen³

¹ *RWTH Aachen*

² *University of Hamburg*

³ *University of Wuppertal*

Corresponding Author: walz@physik.rwth-aachen.de

Solving the question of the origin of ultra-high energy cosmic rays (UHECR) requires the development of detailed simulation tools in order to interpret the experimental data and draw conclusions on the UHECR universe. CRPropa is a public Monte Carlo code for the galactic and extragalactic propagation of cosmic ray nuclei above $\sim 10^{17}$ eV, as well as their photon and neutrino secondaries. In this contribution the new algorithms and features of CRPropa 3, the next major release, are presented.

CRPropa 3 introduces time-dependent scenarios to include cosmic evolution in the presence of cosmic ray deflections in magnetic fields.

The usage of high resolution magnetic fields is facilitated by shared memory parallelism, modulated fields and fields with heterogeneous resolution.

Galactic propagation is enabled through the implementation of galactic magnetic field models, as well as an efficient forward propagation technique through transformation matrices.

To make use of the large Python ecosystem in astrophysics, CRPropa 3 can be steered and extended in Python.

Plenary / 123

A Survey on Distributed File System Technology

Author: Jakob Blomer¹

¹ *CERN*

Corresponding Author: jakob.blomer@cern.ch

Distributed file systems provide a fundamental abstraction to location-transparent, permanent storage. They allow distributed processes to co-operate on hierarchically organized data beyond the

life-time of each individual process. The great power of the file system interface lies in the fact that applications do not need to be modified in order to use distributed storage. On the other hand, the general and simple file system interface makes it notoriously difficult for a distributed file system to perform well under a variety of different workloads. This has led to today's landscape with a number of popular distributed file systems, each tailored to a specific use case.

This contribution provides a survey of distributed file systems and key ideas of their internal mechanics. Early implementations merely execute file system calls on a remote server, which limits scalability and resilience to failures. Such limitations have been greatly reduced by modern techniques such as distributed hash tables, content-addressable storage, distributed consensus algorithms, or erasure codes. In the light of upcoming scientific data volumes at the exabyte scale, two trends are emerging. First, the previously monolithic design of distributed file systems is decomposed into services that independently provide a hierarchical namespace, data access, and distributed coordination. Secondly, the segregation of storage and computing resources yields to a storage architecture in which every compute node also participates in providing persistent storage.

Plenary / 124

Multivariate Data Analysis in HEP. Successes, challenges and future outlook.

Author: Helge Voss¹

¹ *Max-Planck-Gesellschaft (DE)*

Corresponding Author: helge.voss@cern.ch

Extensive use of multivariate techniques has allowed the HEP experiments to improve the information content extracted from their data. This affected both the event reconstruction from the detector response as well as the selection process along a given physics signature. While in many aspects at the forefront of technology, modern statistical analysis tools have only slowly moved from the world of computer science to everyday physics analysis. This presentation reviews multivariate techniques used in HEP, discusses their strengths and challenges, and provides insight into techniques developed elsewhere and their possible usefulness for HEP.

Summary / 125

Summary for Track 1 - Computing Technology for Physics Research

Authors: Axel Naumann¹; Clara Gaspar¹

¹ *CERN*

Corresponding Authors: clara.gaspar@cern.ch, axel.naumann@cern.ch

Summary / 126

Summary for Track 2 - Data Analysis - Algorithms and Tools

Authors: Alina Gabriela Grigoras¹; Martin Spousta²

¹ *CERN*

² *Charles University*

Corresponding Authors: martin.spousta@cern.ch, alina.gabriela.grigoras@cern.ch

Summary / 127

Summary for Track 3 - Computations in Theoretical Physics: Techniques and Methods - part I

Author: Radja Boughezal¹

¹ *Argonne National Laboratory*

Corresponding Author: rboughezal@anl.gov

Summary / 128

ACAT 2014 summary

Author: Jerome LAURET¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Corresponding Author: jlauret@bnl.gov

Plenary / 131

Dark matter review: recent astronomical and particle physics results

Author: Michael Prouza¹

¹ *Institute of Physics Prague*

Corresponding Author: prouza@fzu.cz

The dark matter is undoubtedly one of the greatest enigmas of the modern physics. Its another mysterious companion in the energy budget of the universe, the dark energy, remains fully unexplained, but the dark matter recently started to reveal its secrets. During last years we have marked several important observations in the area of a direct detection based on the particle physics methods, as well as in the area of astronomical indirect observations of the dark matter. Using this new knowledge we are starting to pin down the properties of the dark matter and we are able to exclude some of the former favourite particle candidates for this role. In my talk I will review these most recent advances and observations, I will mention the methods and approaches used in the data analysis, and I will present the current conclusions that we are able to draw concerning the nature and the origin of the dark matter.

Plenary / 133

Future trends and challenges in scientific computing

Author: Branislav Jansik¹¹ I

Corresponding Author: branislav.jansik@vsb.cz

The future and limits of computing in physics are largely dependent on the ever changing capabilities of the contemporary supercomputers. In this contribution, we focus on current and near future computer architectures, their limits and capabilities, in context of scientific computations. Further we discuss trends in high performance computing and their impact on the physics codes. The accelerated, multi-node, many-core architectures offer outstanding computational performance. However, the performance is coming at a cost of encountering unique challenges in efficient utilization of the contemporary and future computational resources.

Plenary / 136

Next Generation Workload Management System for Big Data on Heterogeneous Distributed Computing

Author: Alexei Klimentov¹

Co-authors: Alexandre Vaniachine ²; Danila Oleynik ³; Jack Wells ⁴; Jeff Porter ⁵; Kaushik De ⁶; Kenneth Read ⁷; Paul Nilsson ⁶; Predrag Buncic ⁸; Richard Philip Mount ⁹; Sergey Panitkin ¹; Shantenu Jha ¹⁰; Tadashi Maeno ¹; Torre Wenaus ¹

¹ Brookhaven National Laboratory (US)² ATLAS³ Joint Inst. for Nuclear Research (RU)⁴ ORNL⁵ Lawrence Berkeley National Lab. (US)⁶ University of Texas at Arlington (US)⁷ Oak Ridge National Laboratory - (US)⁸ CERN⁹ SLAC National Accelerator Laboratory (US)¹⁰ Rutgers University

Corresponding Author: alexei.klimentov@cern.ch

The Large Hadron Collider (LHC), operating at the international CERN Laboratory in Geneva, Switzerland, is leading Big Data driven scientific explorations. Experiments at the LHC explore the fundamental nature of matter and the basic forces that shape our universe, and were recently credited for the discovery of a Higgs boson. ATLAS, one of the largest collaborations ever assembled in the sciences, is at the forefront of research at the LHC. To address an unprecedented multi-petabyte data processing challenge, the ATLAS experiment is relying on a heterogeneous distributed computational infrastructure. The ATLAS experiment uses PanDA (Production and Data Analysis) Workload Management System for managing the workflow for all data processing on hundreds of data centers. Through PanDA, ATLAS physicists see a single computing facility that enables rapid scientific breakthroughs for the experiment, even though the data centers are physically scattered all over the world. The scale is demonstrated by the following numbers: PanDA manages O(102) sites, O(105) cores, O(108) jobs per year, O(103) users and ATLAS Data Volume is O(1017) bytes. In 2013 we started an ambitious program to expand PanDA to all available computing resources, including

opportunistic use of commercial and academic clouds and Leadership Computing Facilities (LCF). The project titled 'Next Generation Workload Management and Analysis System for Big Data' (BigPanDA) is funded by DOE ASCR and DOE HEP. Extending PanDA to clouds and LCF presents new challenges in managing heterogeneity and supporting workflow. The BigPanDA project is under way to setup and tailor PanDA at Oak Ridge Leadership Computing Facility (OLCF), and at National Research Center "Kurchatov Institute" together with ALICE Distributed Computing and ORNL computing professionals. Our approach for integration of the HPC platforms at OLCF and elsewhere is to reuse, as much as possible, existing components of the PanDA system. The next generation of PanDA will allow many data-intensive sciences employing a variety of computing platforms to benefit from ATLAS' and ALICE experience and proven tools in highly scalable processing. We will present our current accomplishments with running PanDA WMS at OLCF and other supercomputers and demonstrate our ability to use PanDA as a portal independent of the computing facilities infrastructure for High Energy and Nuclear Physics as well as other data-intensive science applications.

137

Poster lightning session

Conference closing / 138

Conference closing

Author: Lukas Fiala¹

¹ *Institute of Physics AS CR*

Corresponding Author: fialal@fzu.cz

Computations in Theoretical Physics: Techniques and Methods / 139

Data Processing at the Pierre Auger Observatory

Author: Jakub Vicha on behalf of Pierre Auger Collaboration¹

¹ *Institute of Physics AS CR*

Corresponding Author: vicha@fzu.cz

Cosmic rays of ultra-high energy (above 10^{18} eV) are very rare events and still of unknown origin. They provide a unique opportunity e.g. to study hadronic interactions at CMS energies more than one order of magnitude higher than it is achievable at LHC. The existence of the most energetic events (around 10^{20} eV) is theoretically very hard to explain mostly because of the opacity of the Universe at these energies. The Pierre Auger Observatory combines surface and fluorescence detection techniques and it measures these particles already for ten years with the largest exposure ever. The computing tools developed for processing of Monte Carlo and measured data will be presented. A brief overview of selected scientific results will be also given.

Plenary / 140

Statistical tools for the Higgs discovery - the status and future of collaborative statistical model building

Author: Wouter Verkerke¹

¹ NIKHEF (NL)

Corresponding Author: verkerke@nikhef.nl

The discovery of the Higgs boson by the ATLAS and CMS is the result of an elaborate statistical analysis of many signal and control samples for which a set of common tools has been used that were specially developed for the LHC. The key feature of this tool design has been a logical and practical separation between model building, the formulation of the likelihood function, and the statistical inference procedures, which invariably take the likelihood function as experimental input. By allowing the likelihood functions to be expressed in a uniform language (RooFit), components of full analysis model of the Higgs could be formulated by independent teams of physicists, each focusing on a particular Higgs decay mode, and be assembled in a full model in a relatively short time frame. The ability to persist physics likelihood models of arbitrary complexity in ROOT files has further contributed to the exchange of ideas and analysis components, with physicists being able to perform each others full statistical analysis with literally a few lines of code. The RooStats suite of analysis tools that perform the statistical tests on these models, (construction of confidence intervals, upper limits etc), also benefits from this model uniformity: the statistical problem to be solved can be fully specified by a (persisted) RooFit model and a declaration of the parameters of interest, providing a compact uniform interface to a variety of calculation methods: Bayesian, Frequentist or Likelihood-based. I will present an overview of the design and practical successes of RooFit/RooStats tool suite, and their prospects for future use in particle physics.

Computations in Theoretical Physics: Techniques and Methods / 141

Two-loop matching for Standard Model RGE analysis

Author: Andrey Pikelner¹

¹ Joint Inst. for Nuclear Research (RU)

Corresponding Author: andrey.pikelner@cern.ch

After three-loop beta functions for all Standard Model(SM) couplings in \overline{MS} scheme become available, the last missing part for three-loop RGE analysis of SM become two-loop threshold corrections for low-energy input.

In this work we present full two-loop Electro-Weak corrections for top and bottom mass, it's Yukawa couplings, Higgs mass and it's self-coupling, and masses of W and Z bosons. Relations between pole masses and running couplings are implemented in form of computer code and made publicly available.

142

Veřejná přednáška: Petabytová astronomie (in Czech)

Corresponding Author: grygar@fzu.cz

Computations in Theoretical Physics: Techniques and Methods / 143

The generalized BLM approach of fixing the scales in the multi-loop QCD expressions for the physical quantities: the current status of investigations

Author: Andrei Kataev¹

¹ *Institute for Nuclear Research of the Academy of Sciences of Russia*

Corresponding Author: kataev@ms2.inr.ac.ru

The brief review of the modern developments of the studied by Grunberg, Kataev, Chyla, Mikhailov, Brodsky, Wu, Mojaza and others generalized Brodsky-Lepage-McKenzie approach, for fixing the scales in multiloop corrections to the renormalization-group invariant observable quantities in QCD are presented. The special attention is paid to the procedure of splitting the perturbative QCD coefficients, evaluated in the \overline{MS} -like schemes, to the parts, which respect conformal symmetry, and to the ones, which are violating it within perturbation theory. The procedures and of the resummations of the latter ones are discussed. The ideas of their computer implementations are mentioned in brief.

Plenary / 144

Direct water cooling vs free air cooling

Computations in Theoretical Physics: Techniques and Methods / 145

On calculations within the nonlinear sigma model

Author: Karol Kampf¹

¹ *Charles University (CZ)*

Corresponding Author: karol.kampf@cern.ch

The chiral $SU(N)$ nonlinear sigma model represents one of the simplest cases of effective field theory. For the last several decades it has played an extremely important role not only for the low energy phenomenology but also in many other areas of theoretical physics. In this talk we will focus on the tree-level scattering amplitudes of the n -Goldstone bosons. It will be shown that it can be reconstructed using the BCFW-like recursion relations. This method, which does not rely on the Lagrangian description, is much more efficient than the standard Feynman diagram techniques.

Plenary / 146

Concert logistics

Author: Milos Lokajicek^{None}

Corresponding Author: lokajicek@fzu.cz

Summary / 147

Poster lightning: VISPA: Direct access and execution of data analyses for collaborations

Corresponding Author: glaser@physik.rwth-aachen.de

Summary / 148

Poster lightning: EOS : Current status and latest evolutions.

Corresponding Author: geoffray.adde@cern.ch

Summary / 149

Poster lightning: Using Functional Languages and Declarative Programming to analyze ROOT data: LINQtoROOT

Corresponding Author: gwatts@uw.edu

Summary / 150

Poster lightning: Analyzing data flows of WLCG jobs at batch job level

Corresponding Author: christopher.jung@cern.ch

Summary / 151

Poster lightning: The Linear Collider Software Framework

Corresponding Author: andre.philippe.sailer@cern.ch

Summary / 152

Poster lightning: Designing and recasting LHC analyses with MadAnalysis 5

Corresponding Author: eric.conte@ires.in2p3.fr

Expanding software collaboration beyond HEP: pros, cons, dos and donts / 153

Federico Carminati

Corresponding Author: federico.carminati@cern.ch

Expanding software collaboration beyond HEP: pros, cons, dos and donts / 154

Denis Perret-Gallix

Corresponding Author: denis.perret-gallix@cern.ch

Expanding software collaboration beyond HEP: pros, cons, dos and donts / 155

Fons Rademakers

Corresponding Author: fons.rademakers@cern.ch

Expanding software collaboration beyond HEP: pros, cons, dos and donts / 156

Gordon Watts

Corresponding Author: gwatts@uw.edu

Expanding software collaboration beyond HEP: pros, cons, dos and donts / 157

David Fellingner

Plenary / 158

Dinner logistics

Summary / 159

Summary for Track 3 - Computations in Theoretical Physics: Techniques and Methods - part II

Author: Grigory Rubtsov¹

¹ *INR RAS*

Corresponding Author: grisha@ms2.inr.ac.ru