

# Common Software Updates

---

Graeme Stewart (CERN EP-SFT) and Liz Sexton-Kennedy (FNAL)

LHCC Referees, 2022-03-08



# Activities and Planning

- Much activity at the start of 2022 devoted to review of 2021 accomplishments and plans of work
- We shall give some of the highlights here, but for details please see:
  - HSF Planning
    - Coordination Meeting on [20 January](#)
    - [2022 Planning Page](#)
  - SFT Group Meetings
    - 24 January: [SPI, Key4hep, HSF](#)
    - 31 January: [CVMFS](#)
    - 7 February: [Geant4](#)
    - 28 February: [ROOT](#)

# GSoC



- It's Google Summer of Code time again!
- [CERN/HSF](#) has just been accepted again as an organisation
- Project proposals this year
  - 18 projects
  - 25 organisations
  - 39 proposals
- Note that this year Google are allowing again longer projects of 350 hours, which generally works better for us

# Generators


- Working on a Snowmass input document
  - [Topical meeting](#) this week to discuss that
- Establishing better links with the nuclear (EIC) and neutrino community
- Continue to work on highlight topics
  - Negative weights at NLO and NNLO
    - Lots of recent papers to discuss
  - Porting to GPU and better use of modern CPUs (with IT-SC-RD, Madgraph and Sherpa teams)
  - Optimising generator use by the experiments (e.g., [arXiv:2112.09588](#))
    - Would like to review this in the context of new generator releases with features that improve physics and performance

# Simulation



- [Geant4 11.0](#) released December 2021
- Selected Highlights (Kernel, Geometry & Infrastructure)
  - New Tasking system based on PTL or Intel/TBB as parallelism scheme
  - Enhanced CMake configuration; C++17 now standard
  - New release [VecGeom v1.1.18](#)
    - New navigation algorithm with BVH acceleration (for CPU and GPU)
    - New option to run in single-precision
- Selected Highlights (Physics) in release 11.0
  - Quantum entanglement of  $e^+$  annihilation
  - Updated electromagnetic integral approach
  - Substantially simplified fluctuation model sampling algorithm for faster computation
  - Extended interface of the Fritiof (FTF) string model
  - Introduced six light hypernuclei and their corresponding antiparticles
    - Extended the interface of hadronic models for projectile hypernuclei and anti-hypernuclei (via Pythia8)
- Courses @ CERN on Geant4
  - [Beginners Course](#), 9-13 May 2022 (fully Virtual with on hands-on);
- New calorimeter test-beam benchmarks in the [GeantVal](#) physics validation portal


# Simulation R&D and HSF Activity



- New extended example in Geant4.11 demonstrating the incorporation of inference libraries (ONNX, LWTNN) into Geant4 for Machine-Learning based fast simulation
  - The current example is uses a VAE network, trained on some generalised calorimeter geometries, conditioned on the energy and angle of the incoming particle
- LPCC Fast Sim Workshop was held [22-23 November](#)
  - Now to be held annually, based on feedback from participants and experiments
- Completion of the new AdePT example demonstrating the interfacing to Geant4 through the fast simulation hooks
  - This examples allows one to run Geant4 simulation with the GPU based transport (AdePT) enabled in specific parts of the geometry (calorimeter) for the EM physics
  - Energy depositions are collected in selected volumes and transferred back to the host
  - Full validation for CMS geometry, benchmarking and optimisation is ongoing
- HSF WG
  - Use of FPGAs for simulation [looked at last year](#)
  - This year
    - Non-Geant4 simulation ([MARS, 28 March](#))
    - Geometry topics (links to reconstruction geometry to be examined)
    - ML for simulation

# Reconstruction and Software Triggers

- Discussions on Patatrack and ACTS tomorrow
  - Joint [meeting](#) with the Compute Accelerator Forum
- 4D Reconstruction Algorithms, [16 March](#)
  - For detectors with timing information
- [Open Data Detector](#)
  - Development of TrackML detector
  - To be a testbed for experiment independent datasets for algorithm development
- Machine learning topics to be addressed with special ‘learning to discover’ (Institut Pascal Workshop) discussion



# ROOT



- ROOT v6.26/00 has been released with several significant new features to make analysis simpler and faster
  - ROOT's super-efficient systematic variations with a wonderfully simple interface
  - Distributed version of RDataFrame, the PROOF successor
- All experiments agreed to stay for the time being with the ROOT 6.24 release series for their production workflows, for the start of Run 3
  - Significantly simplifies the project's life w.r.t. long-term maintenance!
- Plan of work will be [presented on Wednesday](#) to the “grand public” (i.e. experiments and physicists), and already includes input by experiments
- [ROOT Users Workshop](#) (virtual, organized by Fermilab) on May 9-11
  - ROOT roadmap
  - New ROOT features
  - Updated RNTuple performance studies (result: faster than any other format, for HEP data)
  - User feedback



# Analysis and PyHEP WG

- Analysis
  - Analysis working group paper on Metadata is now on [arXiv:2203.00463](https://arxiv.org/abs/2203.00463)
  - Continue to look at workflow management tools
    - Very popular topic in [December](#)
  - Implementation of systematic uncertainty tools
  - Interesting meetings on Analysis Facilities at the GDB in [December](#) and [January](#)
- PyHEP
  - Restarted topical meetings, with a look at [Boost-Histogram](#) and Hist
  - Planning the 2022 PyHEP Workshop for September - primarily virtual event
  - LHCC input document is [arXiv:2202.02194](https://arxiv.org/abs/2202.02194)



# Frameworks and Tools

- Frameworks
  - Started to look at [C++ features](#)
    - First of a series of topics on C++20, also to look at modules, reflection, ... with tools group
  - Efficient I/O for parallel data processing
  - Common topics with other groups
    - Scheduling accelerator workloads (esp. for reconstruction)
    - Conditions database interactions
- Tools and Infrastructure
  - Continue to look at Spack (becoming the packaging tool of choice in HEP)
  - C++20 features and training topics
    - Joint interest with Frameworks and Training

# CernVM File System and Software Infrastructure



- SPI
  - Several Layers for ATLAS LCG\_101\_ATLAS\_[7,8,9,10,11,12,13,14]
  - Added all packages used by ATLAS to aarch64/ARM builds
  - Added nightly builds for Centos Stream 9
  - Preparing for LCG\_102 (ROOT 6.26/00, Geant4 11.0)
  - Continued work on the Spack built stack
- CVMFS
  - Major release 2.9 came out in November 2021: publishing from ephemeral containers, S3 backend improvements, faster unpacking of container images, bug fixes and quality-of-life improvements (see [release notes](#))
    - Patch release 2.9.1 is due next week
  - Senior fellowship position for R&D on CernVM-FS funded by jumptrading.
    - Focusing on client performance with very large file sets
    - Fellow will start in Aug 2022
  - The goal of the CernVM team in 2022 is laying the foundations that enable addressing scalability limits expected on the timescale of HL-LHC (see [CernVM program of work for 2022](#))

# Training

- Software Carpentry Training
  - [December 2021](#) edition was held (with The Carpentries, ROOT and IRIS-HEP)
  - [March 2021](#) edition to come
  - Plan to host this 3-4 times a year
  - New: will hold Carpentry Trainer Event for our community (accreditation)
- Matplotlib Training close to first run
- C++ Training
  - Regular material development meetings
  - Split course into 2: *Essentials* and *Advanced*
  - Next training is [Essentials, 15-17 March](#)
    - Two new teachers from EP-SFT and IT-SC-RD

# Workshops

- Detector Simulation on GPU Community Meeting
  - Discuss and review progress on the GPU detector simulation projects
  - AdePT and Celeritas will both participate
  - Virtual format: 4 European afternoons, 3-6 May
- ROOT Users Workshop
  - ROOT roadmap and new ROOT features
  - Updated RNTuple performance studies
  - User feedback
- Analysis Ecosystems Workshop 2
  - Continuing the quest for reduced ‘time-to-insight’, 5 years after the Amsterdam workshop
  - Looking at the latest developments in ROOT, data science tools, analysis facilities
  - Discussion on next steps and missing pieces for end-to-end analysis - to produce a written report
  - In person (!) event at IJCLab, 23-25 May
- PyHEP 2022
  - Mostly virtual, anticipated for early September 2022
- ACAT2022
  - Our big HEP software conference this year with significant SFT involvement
  - In person, Bari 24-28 October

# New Activities

- Analysis Facilities Forum
  - Community platform for those interested in contributing to the development of analysis facilities for use by HEP experiments
    - Very active area in projects like IRIS-HEP, GDB discussions and experiments
  - Linking elements into an integrated data, software and computational resource
  - Kick-off meeting 25 March
  - Should aim to produce a report on a 12-18 month timescale
- Conditions Databases
  - Early HSF activity
  - Renewed interest from experts in ATLAS, CMS, Belle II, DUNE
  - Work to defining an API embodying best practice

# Summary



- HSF Activities starting up again for 2022
- New areas of interest are identified and the HSF can provide a home
- Looking forward to interesting focused workshops again

# Backup

---



# Simulation: 11.0 Release Details



- Highlights (Kernel, Geometry & Infrastructure)
  - New Tasking system based on PTL or Intel/TBB as default parallelism scheme
  - Enhanced CMake configuration system supporting modular builds & Jenkins pipelines
  - Migration to C++17 ISO Standard, now set as default
  - New example enabling VecGeom navigation in Geant4 ( **$\beta$**  version)
  - New interface allowing for custom tracking managers
  - New release [VecGeom v1.1.18](#)
    - New navigation algorithm with BVH acceleration (for CPU and GPU)
    - New ability to run in single-precision on-demand
    - New CUDA manager
  - New prototype web site based on Jekyll

# Simulation: 11.0 Release Details



- Highlights (Physics) in release 11.0
  - Quantum entanglement of  $e^+$  annihilation
  - Updated electromagnetic integral approach
  - Substantially simplified fluctuation model sampling algorithm for faster computation
  - Updated Livermore models to use EPICS2017 data
  - New class store for optical material properties, allowing also predefined properties usage
  - Extended interface of the Fritiof (FTF) string model to get the number of projectile/target spectator nucleons, number of nucleon-nucleon collisions and impact parameter
  - Introduced six light hypernuclei and their corresponding antiparticles; extended the interface of hadronic models for projectile hypernuclei and anti-hypernuclei
  - Added implementation of hyper-cluster emission in INCLXX
  - General code clean-up in both EM and hadronic modules; removed deprecated processes and models