




Software Update


Graeme Stewart, for EP-SFT and HSF

2020-06-03




Workshops and Generic Activities

Google Summer of Code & Season of Docs



- GSoC
 - Another successful year for the HSF as an umbrella organisation for GSoC
 - 36 slots from Google were assigned (of 37 possible; +3 from 2019) and have been taken up by projects
 - We are now in the coding period
- Season of Docs
 - Last year successful documentation project for ROOT was our first test
 - This year 3 proposals
 - AllPix2; ROOT; Rucio

CERN-HSF GSoC participation



LHCC HL-LHC Review

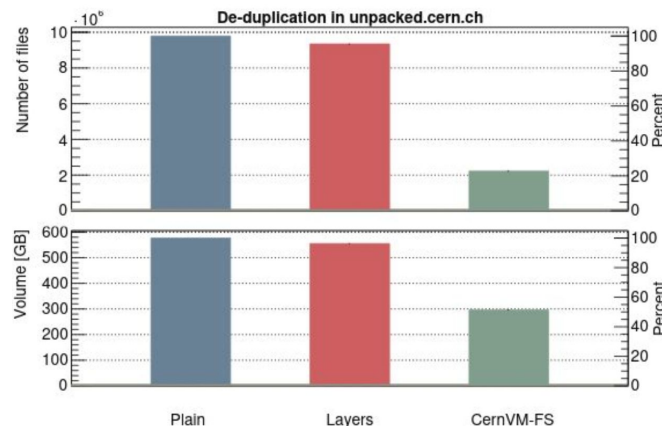
- Significant effort put into the HSF document on Common Tools and Community Software
 - <https://zenodo.org/record/3779250>
 - The document is an excellent survey on how things have advanced in the key areas since the Community White Paper
 - Eagerly awaiting the official feedback from the review panel
- Supplementary documents
 - ROOT: <https://arxiv.org/abs/2004.07675>
 - Generators Working Group: <https://arxiv.org/abs/2004.13687>
 - *The HSF is very happy that these teams took the time to provide their focused input for the software development path towards HL-LHC*

Software Deployment pre-GDB I

- Remote only pre-GDB, examining the challenges and opportunities of software deployment for the coming decade
 - Looking at many perspectives: experiment librarian, platform providers, infrastructure developers to site administrators
- Containers here to stay as the default abstraction layer between the software and the hardware
 - Encapsulation of user environments as a way to support analysis (supported by ATLAS; under investigation in CMS and LHCb)
 - But scaling issues when starting massive numbers of container images from the default

Software Deployment pre-GDB II

- CVMFS offers much better deduplication and deployment scaling of only the pieces needed by the jobs
 - x10 to x50 better image distribution efficiency
 - unpacked.cern.ch, singularity.opensciencegrid.org
 - Target making user experience as seamless as *docker push*, hide the conversion
- Software deployment to HPCs still an ongoing challenge (N-challenges across N sites...)
 - Fat containers can solve the problem in the short term
 - Shrink-wrapped CVMFS
 - Newer linux kernels offer unprivileged mounts from inside the container
- Kubernetes (k8s) interesting for sites, industry standard container orchestrator
 - Ongoing development to better support filesystem based containers and CVMFS





WLCG-HSF Workshop

- Planned face-to-face workshop in Lund from 11-15 May was unable to take place
- Decided not to lose the slot, but organise a virtual workshop, taking place over 3 days, 2 hours a day
 - 16-18h CERN time - fairly good for Europe and US; hard for Asia
 - Strong feeling that virtual meetings demand higher concentration and cannot run for as long as normal workshops
- **New Architectures, Portability, and Sustainability theme**
 - Would have been the all-day Tuesday plenary in Lund
 - Monday: Application Software
 - Tuesday: Processing Frameworks
 - Wednesday: Validation and Accounting

A Success!


- 221 people registered
- Slides were posted in advance for review
- Workshop notebook was available in advance
 - A sort of Live Notes++
- Attendance peaked at 175 Monday, 150 Tuesday, 110 Wednesday
 - We had a clash with LHCOPN/LHCONE meeting on Wednesday :-(
- As this was the first event of this type we hosted we put effort into
 - Post-workshop survey
 - Identifying outcomes and follow-ups

Monday - Application Software

- Code Portability
 - Increasingly large number of possible non-CPU devices available
 - Clear that the community cannot support N codes for N platforms
 - Industry knows this too, hence proliferation of toolkits
- How to assess the best?
 - This is an orthogonal question to redesigning code for at least one parallel architecture
- DOE HEP-CCE Project
 - Portable Parallelization Strategies
 - Assess metrics for toolkits on real HEP examples

	OpenMP Offload	Kokkos	dpc++ / SYCL	HIP	CUDA	Alpaka	
NVidia GPU			<i>Intel/codeplay</i>				Supported
AMD GPU		<i>prototype</i>	<i>via hipSYCL</i>				Under Development
Intel GPU						<i>very early development</i>	3rd Party
CPU							Not Supported
Fortran							
FPGA						<i>possibly via SYCL</i>	

Monday - Application Software



GPU ray tracing of Track ML geometry using VecGeom

- Heterogeneous Architectures and Detector Simulation
 - Simulation a very significant part of HEP computing; common engine in Geant4
 - GeantV ([arXiv:2005.00949](https://arxiv.org/abs/2005.00949)) taught valuable lessons about how to optimise
 - Optimisation from SIMD far less than hoped for - data preparation costs to use vector registers is high (see Pere's talk, [November 2019 LHCC Referees](#))
 - Modernising and reducing code size can bring up to x2
 - Dedicated libraries to do pieces of HEP specific code can be reused (VecGeom)
 - Ideas for the future: ray tracing on GPUs?
- TensorFlow as a Compute Engine
 - Using highly optimised libraries with built in GPU support
 - Quite a few HEP projects: TensorFlowAnalysis, zfit, pyhf
 - Engine is designed for different purpose to ours, impedance matching can be awkward

Tuesday - Application Frameworks

- Heterogeneous Experimental Frameworks
 - Goal is *optimal* use of heterogeneous resources
 - Easier on owned resources (HLT) as opposed to HPCs or other sites
 - Separate process spaces (ALICE-FAIR approach in O2)
 - Great code separation, dynamically balance CPU resource use at process level
 - Accelerator only approach
 - Ideal for R&D projects; optimal performance for the target devices
 - Can leave CPUs idle when they could do useful work
 - Hybrid approach
 - Asynchronous execution, so most complex for framework
 - But maybe the biggest prize
 - Smart underlying schedulers (TBB, HPX) help maximise CPU usage
 - CMSSW is a good example of implementing this
 - In all cases there are hurdles for the experiment developers and some steep learning curves

Survey Highlights

- Slides in advance and notebook were liked
- Live talks strongly supported over pre-recorded talks
- Format was about what people wanted (duration, timeslot)
- We should be much more disciplined with time keeping
 - More important for virtual meetings
- Topics covered were supported
- People want to see conclusions and outcomes
 - We are in the process of identifying these now
- Large support for continuing with virtual meetings in the future
 - But people want face-to-face meetings as well

We will probably organise a follow-up virtual event in September

Projects and Working Group Updates

Geant4



- [Plan of work for 2020](#) for Geant4 available since last March
 - Details in backup...
- Defined release schedule for 2020
 - Next public release 10.7, on December 4th; Beta release on June 26th
 - New [patch release 10.6.p01](#) deployed last February; [patch release 10.6.p02](#) in May
- New activities started, as part of the [Geant4 Task Force for R&Ds](#)
- Courses ([beginners](#) & [advanced](#))
 - CERN technical training - beginners held last January; advanced postponed to September
- [Virtual] Geant4 Collaboration meeting, on 21-25 September

HSF Detector Simulation

- Topical meetings
 - Detector geometry issues
 - Simulation requirements from non-LHC experiments
- Review of successful GPU projects in the simulation domain
 - Medical physics MPEXS (sadly closed source)
 - Optical photon transport with NVIDIA OptiX engine (used for JUNO)
 - Neutron transport for reactors
- Next meeting will be for lightning talks for simulation on GPUs R&D projects
 - Link in to Geant4 R&D task force and to all other ideas in the community

Reconstruction and Software Triggers

- Meetings
 - Projects and progress for HL-LHC tracking meeting
 - Held jointly with IRIS-HEP
 - Multi-threading and code optimisation in ATLAS and CMS
- Significant input for HL-LHC document
 - Can help become a place for people to discover different projects
 - Ideas to improve HSF webpages with list of projects from community input
- HSF will host a repository to be used for cross-experiment algorithm development and validation
 - First topic: Graph Neural Network tracking approaches

ROOT



- New production release v6.20/00, February 2020
 - Zstd compression, always better than zlib
 - Redclaration improvements, very useful for notebooks
 - Help system improvements (`.help TTree::branch` opens the correct documentation)
 - RooFit fixes and speed-ups
- Follow on patch releases in March and April
- Major developments (see Axel's talk, [February 2020 LHCC Referees](#)):
 - RNTuple replacement for for TTrees (smaller and faster)
 - PyROOT new version for next release
- Plus, HL-LHC document, <https://arxiv.org/abs/2004.07675>

HSF Analysis WG

- Pre-CHEP Workshop held in Adelaide
 - Analysis Systems: From Future Facilities to Final Plots
 - Cross-over with DOMA
 - Extremely well attended, almost 200 participants
 - Novel interactive session (brain writing) to help engage all attendees
- Follow-ups
 - More interaction with DOMA experts
 - Analysis description languages workshop (planned - delayed due to COVID-19)
 - “Exclusive” analyses that don’t fit into the model

PyHEP WG

- Second PyHEP Workshop last October
 - <https://indico.cern.ch/e/PyHEP2019>
 - Visualisation keynote, Python and accelerators, analysis tools, histogramming, fitting
- Next workshop planned for July 13-17
 - Was to be co-located with SciPy-2020 in Austin TX
 - SciPy has become virtual, due to COVID-19
 - [PyHEP 2020](#) virtual event planned now
 - Double time-slot format to accommodate Asia, Europe and Americas
 - Remarkable level of interest
 - *536 people have registered*
 - Planning for keynote talks as well as hands-on tutorials and “notebook-talks”



Frameworks WG

- New Frameworks Working Group established
 - Group has been very active holding meetings reviewing the state of current experiment frameworks
 - Organisation of frameworks session in the last workshop
- Discussions on strategy for the use of threading libraries
 - TBB and any eventual replacement
 - Underlines the role of HSF for providing a coherent view of the whole software stack used in HEP

Generators

- Full version of HL-LHC paper submitted to Computing and Software for Big Science [[arXiv:2004.13687](#)]
 - Understand most effective choices that can be made
 - Improvements in generation efficiency
 - Particularly negative weights: recent papers [[arXiv2005.09375](#), [arXiv:2002.12716](#)]
 - Technical improvements: thread safety to accelerator use
- Work on GPU version of MadGraph ongoing

Infrastructures and Tools

- New LCG releases
 - LCG_96b - ROOT 6.18
 - LCG_97 - ROOT 6.20
 - Including nightly builds for latest ROOT patches
- Consolidation of container builds
- Deep investigation of use of [Spack](#) for build orchestration (LLNL tool for building science software stacks)
 - Links to EP R&D project on Turnkey Software Stacks
 - Future accelerator software experts involved: CLIC, ILC, FCC, CEPC
 - Presented also at pre-GDB meeting

Training

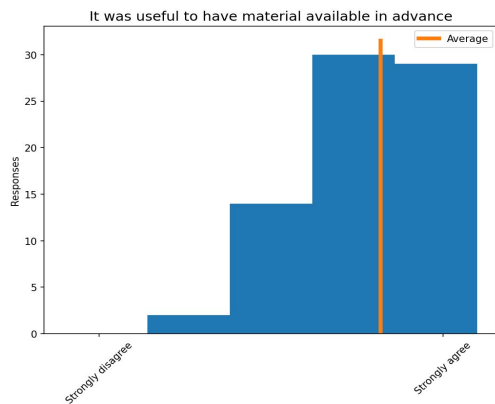
- Software Carpentry @CERN held in November 2019
 - Mixed general software skills, taught by Carpentries instructors with more specific HEP analysis techniques
 - Subsequent Software Carprenties event postponed due to COVID-19
- Analysis Preservation Training
- Virtual Pipelines Training event happening right now...
- Advanced training in Alpaka was postponed from March (COVID-19)
 - Virtual training event in preparation, planning for July
- Template for development of HEP training materials is in development
 - Range of materials for different levels of learning
 - <https://hepsoftwarefoundation.org/training/curriculum>

Summary


- Many activities ongoing for R&D in software in EP-SFT and all experiments
- HSF continues to enhance communication between experiments
 - Working groups are active
 - Forum for exchange of ideas
- Preparation of HL-LHC computing review documents was a good opportunity to assess status after the CWP
- Adaption to virtual working environments is an ongoing process
 - Second virtual event in September

Backup


Workshop Survey Highlights I



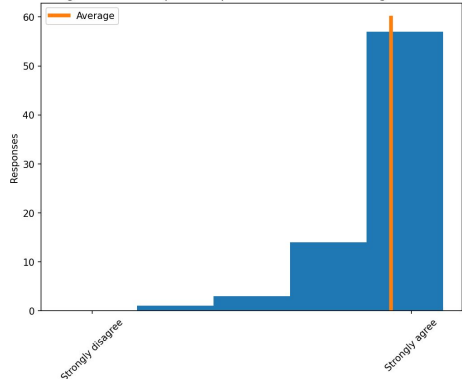
I reviewed the pre-workshop material before the conference




I posted questions in the notebook before the workshop



It was good to have speakers present slides live during the workshop



I would prefer that all of the talks are pre-recorded instead of presented live




- Slides in advance and notebook were liked
- Live talks supported over pre-recorded talks


Workshop Survey Highlights II

- Format was about what people wanted
- We should be much more disciplined with time keeping
- Decent numbers of people came to whole event and areas outside their immediate interests


Organising a virtual workshop in 2 hour blocks was




Given the live presentation format, the planned 50% of time for discussion was



Sticking strictly to the scheduled timetable is important (even cutting off speakers if they badly overrun)

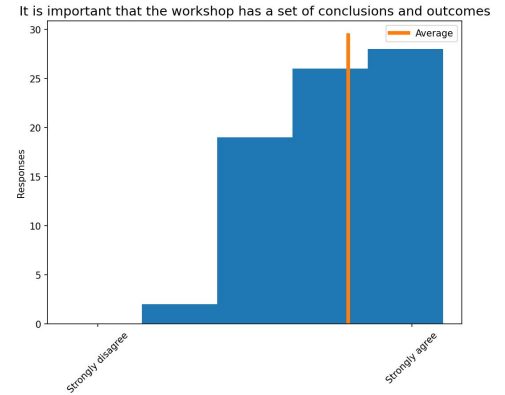







For the workshop




Workshop Survey Highlights III

The sessions that were organised were very useful...



Assuming easier travel in the future, I would like to see
An even mix of virtual meetings and face-to-face events



- Topics were supported
- People want to see conclusions and outcomes
- Large support for continuing with virtual meetings in the future
 - But people want face-to-face meetings as well

Geant4



GEANT4
A SIMULATION TOOLKIT

- Plan of work for 2020 for Geant4:
 - New alternative specialised transport for e+ and gamma
 - data base for ionisation cross-sections of K-, L-, and M- shells by heavy ion collisions
 - linear gamma polarization options into all gamma models
 - new ion energy fluctuation model
 - updates and tuning to Goudsmit-Saunderson model
 - new model for three gamma annihilation
 - improvements to diffuse elastic model
 - extension of High-Precision model to higher energies
 - extension of Bertini intra-cascade model for light-target collisions
 - hadronic shower improvements for FTF and QGS models
 - new and improved biasing fast-simulation techniques
 - interface with VecGeom navigator
 - templated equation of motion and field steppers
 - tasking parallelism
 - revision of production thresholds