

# HSF workshop notes, May 2016, LAL

## Contents

[Timetable](#)

[Project support in HSF \(Mon am\)](#)

[HSF status \(Mon pm\)](#)

[Community white paper - Peter Elmer](#)

[News from projects \(Mon pm\)](#)

[Future conditions DB - Andrea Formica](#)

[AIDA2020 WP3 - F. Gaede](#)

[DIANA-HEP - P. Elmer](#)

[SW&C Knowledge Base - T. Wenaus](#)

[WikiToLearn - R. Iaconelli](#)

[Learning from other communities \(Mon pm\)](#)

[Bioconductor - W. Huber](#)

[Netherlands eScience Center](#)

[Depsy - J. Priem](#)

[Machine learning \(Tue am\)](#)

[Impact of Deep Learning for HEP SW&C - A. Farbin](#)

[OpenLab ML and Data Analytics Workshop - M. Girone](#)

[Recent Developments in ROOT/TMVA - L. Moneta](#)

[Data Analysis and Reproducibility Tools for HEP - A. Ustyuzhanin](#)

[Keras/Theano - M. Paganini](#)

[OpenData in CMS - K. Lassila-Perini](#)

[A Common Tracking Software - A. Salzburger](#)

[Inter-experiment ML Group and HSF - S. Gleyzer](#)

[RAMP: ML Hackathon on Anomaly Detection \(Tue pm\)](#)

[Packaging \(Tue pm\)](#)

[Spack Presentation by Patrick Gartung. > 30 participants](#)

[Geant4 technical form \(Tue pm\)](#)

[Software performance \(Wed am\)](#)

[SW Performance @ALICE - D. Rohr](#)

[ATLAS Observations - G. Stewart](#)

[CMS - D. Lange](#)

[Astrophysics Experience - O. Iffrig](#)

[ROOT Experience and Challenges - P. Mato](#)

[Art/LArSoft - M. Paterno](#)

[GeantV - J. Apostolakis](#)

[Discussion panel](#)

[Next steps & wrap-up \(Wed pm\)](#)

[Logo](#)

[Meeting notes](#)

[GeantV review](#)

[Journal proposal "Computing and Software for Data-intensive Physics"](#)

[HSF in StackExchange?](#)

[HSF as legal entity](#)

[HSF governance](#)

[HSF Center](#)

[HSF \(Human\) resources](#)

[HSF communication](#)

[Community white paper & road map](#)

[Outcomes, conclusions, next steps](#)

[Actions out of the workshop](#)

## Timetable

### Project support in HSF (Mon am)

HSF is about SW: project support an important goal

- Good practices and appropriate tools can help sharing SW
- Plenty of things to consider, often forgotten or postponed

A list of best practices have been drafted

- On GitHub as everything done by HSF (documents repository)
- Feedback is more than welcome!
- License is one of the critical topic: HSF Technical Note available to help
- Importance of documenting contribution workflow to get contributors

Recommendations are not HEP specific which is as it should be

- Ensure there aren't hidden dependencies like assumptions on afs or cvmfs presence

A template is available for new projects: `hsf_create_project.py`

- Build a CMake project
- Configure build, tests, packaging, install: some technologies chosen among the possible ones
- Created files with detailed documentation
- Genesis of the starter template was a real project stripped down to its essentials

Test resources: plenty of resources available (for free) for continuous integration

- In addition CERN proposed access to TechLab for specific, fancy HW. FNAL ready to do the same...

Define a set of standards to be met to be accepted as an HSF packages?

- Checklist is a good start
- Should not be exclusive or prescriptive, point is to give guidelines

Oriented towards standalone projects. Another aspect is a project that needs to build against a number of other large pieces, and adds incremental functionality. Define standard ways of interfacing to existing projects.

- ROOT, for example.
- If project is cmake based, have a set of tools to interface to any other cmake based package
- A point of the packaging WG is to define a canonical stack to enable projects following guidelines to work together
- It's an objective of the spack code sprint tomorrow
- Set things up so that "if you use ROOT, just uncomment this line" etc

Should include requirement for documentation of needed env variables, other environment setup

**Please give feedback on the draft best practices guideline document linked in Benedikt's slides. Would be a very good output of the workshop to have a document representing the collective view of this group.**

### HSF status (Mon pm)

HSF objectives remembered: share expertise, raise awareness around projects, make easier for people to start new projects...

HSF concrete work organized in 6 WGs: information exchange, training, SW packaging, SW licensing, SW projects, dev tools and services

Communication: web, mailing list/fora, knowledge base, technical notes

- Topical forums: SW Technology (former Concurrency Forum), Reconstruction Algorithms, Machine Learning
- KB: emphasize easy adding/editing of contents, cross-linking
- TN: anything useful that you want to share, already 4 available

Training: focus on sharing training material

- Important focus on WikiToLearn, a wiki-based platform to make this easier. WikiToLearn objective: make possible for students/trainees to improve the material. Contents is what you contribute!

SW Packaging: see this morning and tomorrow's hackathon

- Work so far summarized in a TN

Licensing: guidelines summarized in a TN

- Second version planned based on your feedback

SW projects: see this morning too

HSF is also about fostering collaboration around SW: some successes in the last year (Next-generation Conditions DB, Track reconstruction, Gaudi

- Some satellite projects like DIANA (4-year NSF funded), AIDA2020) taking care of being coordinated with HSF

Food for thoughts during the workshop: to be discussed on Wednesday

- HSF need a logo: vote for your preferred logo!
- Project reviews: recently asked by GeantV as a first candidate. Review as an open process but report and conclusions by the review committee. Other interested project should contact HSF.
- HSF StackExchange? A complement/alternative to forums. Overlap with other tools like RootTalk?
- HSF Journal: publish S&C papers in peer-reviewed journals. May be good for larger visibility across communities, may help with career recognition. Springer ready to collaborate.
- HSF need more dedicated resources/people: required to achieve our goals, needs to be funded. Endorsement from ICFA or similar bodies?
- HSF as a legal entity: currently not a foundation! May help with IPR/copyright handling, facilitate collaboration with industry... Also cons... discuss during the workshop, try to agree on a plan on Wednesday

Discussion

- Conda missing from packaging tools report? Didn't have a Conda advocate in the group, though it was presented. Document was in circulation for comment for weeks, could have had input at any time. Did try to do some Conda follow-up and didn't find volunteers. Can still welcome a revision that includes Conda properly. Could do it in this workshop.
- In seeking resources by 'HSF resources' we do not mean HSF owned resources. We mean people identified and assigned with a fraction of their time to work through HSF projects. E.g. someone working on re-engineering tracking, work on that through the HSF to make the effort common as much as possible.

## Community white paper - Peter Elmer

HSF demonstrated some initial collaborative activities but to address the challenges ahead of us (e.g. HL-LHC) we need more and dedicated resources

- 2 projects recently funded, DIANA (NSF) and AIDA2020 WP3 (EU): not started as part of HSF but they are examples of projects which will contribute to the overall HSF effort
- Would help to put together a community roadmap describing HEP S&C challenges and directions (similar to P5 for HEP experiments in the US)
- Community White Paper (CWP): R&D to prepare TDR, common framework that could help...

- Could also provide a better context for applying for resources, engage computing people...
- An NSF Software Institute plan is one example of a project proposal which could benefit from a clear community roadmap
- CWP process: the proposal is for a series of HSF-branded workshops, starting this Fall and converging on a final workshop next Spring.
- Need participation of funding agencies, coordination bodies...

## News from projects (Mon pm)

### Future conditions DB - Andrea Formica

Conditions data are used at different stages of our workflows and are also refined during these stages (online processing, prompt reconstruction, bulk processing, ...)

CMS developed during LS1 a new schema trying to simplify what was done in Run 1: this schema is in production in Run 2

- ATLAS expressed interest: matching its own ideas of useful metadata after Run 1 experience
- Condition data as Blobs: storage system not aware about the internal structure. Metadata in the tag table provides the necessary information to deserialize the object.
- CMS chose Boost library for the serialization: several others choices possible (Google Protobuf, HDF5, ROOT or ascii formats like JSON and XML)
- 

Architecture requirement: decouple clients from backend using a REST API, support many different kinds of backends including several relational platforms (Oracle, PostgreSQL,...), but eventually also NoSQL or file system. All business code to manage the conditions data is inside a server. The clients can be implemented in several languages because the communication is done via HTTP using JSON for the messages exchanged with the server.

Current prototype implemented in Java (based on JEE, Spring): easier integration into Frontier, the idea being to profit at most of the experience gained by the Frontier development group in terms of caching etc....

Client implementation for the moment on going. For python easy via automatic generation of client API based on swagger documentation.

### AIDA2020 WP3 - F. Gaede

AIDA2020: detector R&D, not specifically SW

- WP3: SW related WP, promote commonalities, develop SW specific to projects only if needed
- 7 tasks around SW challenges
- Strong links with current and future experiments: milestones to ensure that what is done in WP3 will be useful

DD4HEP: 1 of the HSF incubator projects providing a generic detector description toolkit

USolids/VecGeom: generic shape library, introduced in G4 10.x

- Integration with SIMD instructions

Alignment: developed a fully automatized fast alignment procedure for LHCb VELO, in production

EDM Toolkit: PODIO project, also in HSF incubator, efficient I/O with PODs, currently being evaluated in the context of FCC/Icio

Framework extensions: improving parallel scheduling in frameworks, currently focused on Gaudi, later plan to use it in Marlin and PandoraSDK

DDG4: interface between DD4HEP and G4

Advanced tracking tools: follow-up of aidaTT from AIDA project

- Parallelized tracking algorithms
- Integration of LHC tracking algorithms into aidaTT

- Connected to the Common Tracking Forum started recently

Advance Particle Flow Algorithms (Pandora PFA toolkit)

- Framework independent
- LC, HL-LHC, LAr neutrino...

## DIANA-HEP - P. Elmer

Data intensive ANALisys for HEP: collaborative efforts around analysis tools to make them a sustainable infrastructure in our community

- 4 year project (2015-2019)
- Not a new SW
- Not only about SW development but SW eco-system, career recognition, training...
- ROOT as one core component: performance, ROOT I/O
- Several PIs with a strong experience in complementary fields for addressing these challenges, involved in various HEP experiments or computer sciences (e.g. machine learning)

## SW&C Knowledge Base - T. Wenaus

<http://hepsoftware.org>: the last generation (hopefully the last one, works nicely!)

- Javascript app in the browser
- No support for access restricted links yet (experiment-specific links may come)
- Software projects organized by categories and can belong to one or several
- Also allow descriptions and links to institutions, experiments...
- Links to related projects
- Also descriptions/links to sciences
- Add your software, publicize SW&C KB!

See slides: many examples in them

## WikiToLearn - R. Iaconelli

Collaborative textbooks: promote collaboration around training materials

- Centered on students/trainees
- Currently 160 contributors, among them 40 core editors

Text at the heart of the system (external links or pdf are accepted but not encouraged)

Every user can make a book from a piece of training material

Support for draft pages (basically personal version of the material)

Tracking and notifications of modifications

Result of collaboration: sharing of effort

Ability to run snippets of code in the browser or in a container (AKHET, desktop as a container with Webdav file access)

- Currently resources provided by 3 GARR sites

Features to come

- Offline editor
- Collaborative editor à la Google Docs

# Learning from other communities (Mon pm)

## Bioconductor - W. Huber

A use case illustrating the challenge in biomed: leukemia, a disease with an heterogeneity of causes posing a real challenge for treatment, drug research...

- Also heterogeneity of data sources/types
- Complex correlations
- SW required to extract value from data

Bioconductor: collaborative and open-source SW

- Interoperable components, rapid development, code reuse
- Based on R
- Computational reproducibility
- Low barrier of entry for users, training
- 10K users: world largest bioinformatics project
- Driven by scientists

Contributor community increased over time

Several modes of interaction developed, including web site, mailing list, video-conference...

Importance of documentation: not only manual pages but user “vignettes” (narrative overview), workflows, citable papers with peer review

Increased use of GitHub for package development to allow early and open peer-review

Lessons learnt

- Trade-off between standardization and being open and inclusive
- Initial resistance of contributors afraid of losing their credits
- Need a core team
- Career of people involved have benefited

## Netherlands eScience Center

An organisation bridging scientific communities and computing infrastructures

- A core team of experts on optimized data handling, big data analytics, efficient computing
- Covering all scientific fields
- Reuse community solution in other contexts: open-source, open access. Build upon expertise gained in different contexts
- Funding done by NeSC: allow to retain expertise and to do the “technology transfer”
- Very similar to HSF in a way... but cross-discipline

One example: analysis with ROOT for Xenon1T (dark matter).

- Interconnection with Panda/mongoDB. Learned about the difficulty to glue different pieces! For example, problem with a bug in ROOT6 (memory leak with Python 3) not yet fixed upstream...
- packaging/dependencies with Conda: Conda recipes for ROOT6 available on GitHub (nlesc project): available in Anaconda cloud (NLeSC/Packages/root-numpy)

## Depsy - J. Priem

Depsy: NSF funded project

- Promote credits for SW contributions
- Retrieve statistics about SW packages, projects reusing them, contributors...
- Currently focused on Python and R: using the central repositories for packages like PyPi

SW citation often only informal and doesn't permit identification and credit of contributors  
Depsy takes into account indirect contributions (transitive credits): contribution to a project heavily reused  
Aggregated impact of people across many projects

#### Discussion

- Pere: how much challenging will it be to do the same thing for C/C++ packages? Jason: yes, more challenging as there is no central, well-recognized, package repo.

## Machine learning (Tue am)

### Impact of Deep Learning for HEP SW&C - A. Farbin

Motivations, potential: see slides

- Very successful use in the context of LArTPC

GPUs critical for performance

Need to provision a large amount of resources: no longer embarrassingly //

DL in reconstruction: will require distribution of large datasets for training DNN

- Will require available of public dataset: also required to communicate with DL experts

Proposal of a R&D project consisting to build a HEP framework on top of DL

- Use Google TensorFlow framework to the DL part (open-source)
- Success would mean HEP science is closer to other sciences: would benefit from more experts and offer a greater potential for collaboration with Data Scientists (and industry)
- Need to start now if we want to be ready in 10 years
- DL is not necessary easy... can also complicate things...

Potential of DL to find "unexpected things" still to be assessed in our context

### OpenLab ML and Data Analytics Workshop - M. Girone

@CERN, April 29, <https://indico.cern.ch/event/514434/timetable/#20160429.detailed>

A lot of investment in industry: interesting contributions by several companies. Many good reasons to collaborate but must also take into account the different culture regarding collaboration. OpenLab and its NDAs infrastructure can help.

Event classification: interest by LHCb and ALICE to meet the Run3 challenges, CMS also has plans

Object identification: great potential, raised by all experiments. Could benefit from industry experience with image recognition (self-driving cars...) if we could formulate the problem in a similar way...

Anomaly detection: a growing use case in the industry (e.g. security), may be relevant to HEP

Data analysis: all experiments with plan to have events almost ready for analysis when they leave the detector (online reconstruction for ALICE and LHCb). Potential for streamlined analysis, want to look at tools like Spark/Hadoop, OpenLab to help setting up a testbed

A lot of tools available to optimize data access and analysis, produced by industry but open-source

Event classification and triggering probably the most challenging use case: nothing really matching this in the industry but several frameworks can help

### Recent Developments in ROOT/TMVA - L. Moneta

TMVA future discussed last September. Core requirements identified.

- Maintain a set of core algorithms for HEP
- Be able to interface with Python and R to allow the use of modern ML tools (done, PyMVA and RMVA)
- Workflows integrated with ML frameworks

Several actions decided, most of them already done or in progress: see slides

New method to compute Feature Importance based on contribution to the classifier

- Independent of the actual classifier used

New DL classes recently added supporting recent developments in the field: currently being optimized by TMVA developers

Cross validation also recently added

- Support for // execution (Spark, multi-processes) being added

Work in progress in TMVA

- Better separation of classification and regression and improved regression
- Improved perf and memory usage
- Better support for parallelization
- Integration with ROOT-Books (Jupyter notebooks): Service for Web-based Analysis (SWAN). No SW installation required
- Use general ROOT I/O instead of being limited to XML/JSON, import training output from other frameworks, import data from non ROOT formats (e.g. HDF5)

5 students will work this summer funded by Google (Google Summer of Code program)

Welcoming more contributions to TMVA development

- Join the developer team or issue pull request on github

V. Innocente: CMS has faced performance/mem footprint issues with TMVA DL classifiers in the context of reconstruction.

We need to find solutions, e.g. reduced precision (FP16) which proves to be enough in many case.

Amir Farben: like TVMA interface but skeptical about trying to integrate anything into TMVA. Results is that it makes everything more complex, need to reinvent parallelisation optimisations...

- M. Schulz:

## Data Analysis and Reproducibility Tools for HEP - A. Ustyuzhanin

YANDEX: 2 companies related to data science (YANDEX Data Factory) and research education (YANDEX School of Data Analytics, non profit)

- School of Data Analytics member of LHCb since 2011
- Several contributions to LHCb: topological trigger, data storage optimisation...

Several developments targeting reproducible research (all on GitHub, open-source, Apache 2.0)

- Reproducible Experiment Platform: Python based, Jupyter friendly, interface with scikit-learn

## Keras (Theano, TensorFlow) - M. Paganini

Keras: Python library interfacing with tensor manipulation frameworks like TensorFlow (TF) and Theano

- Keras provides building blocks for users to construct deep learning models at a high level of abstraction
- Theano and TF are high level wrappers around C++: efficiently call C++ from high level language
- Similar in many ways:
  - Computation described as a directed acyclic graph (DAG) with tensors flowing along the edges
  - Theano older, more stable API, faster for CNNs + interfaces with CuDNN (NVIDIA optimized kernels for NN operations)
  - TF backed by Google + easier for distributed/cluster computing
- Theano: dynamic C code generation, stability optimization, execution speed optimisation, support for GPUs
- TensorFlow: meant for large scale deployment on heterogeneous systems/HW, support synchronous and asynchronous training, evolved from DistBelief

In ATLAS:

- Keras-trained RNNs for flavor tagging already in Athena ([RNNIPTag](#)) based on lightweight NN class ([lwtnn](#))
- Ongoing work to try to integrate TensorFlow with Athena

## OpenData in CMS - K. Lassila-Perini

Challenge: knowledge preservation



- We are good at keeping/saving “immediate” metadata related to data taking conditions
- Not as good to save the context allowing to analyze properly the data: things trivial at the time of the experiment but easily lost

OpenData/preservation effort forces to address the “context metadata” preservation challenge. Beneficial for HEP community.

Lesson learned: much better to start the open data effort at the same time as data analysis is done but competing for human and computing resources...

- Also making data public is not enough to make the data easy to understand/use... Requires a specific effort (LHCb has one FTE dedicated to maintain a guide describing all the trigger lines...)

## A Common Tracking Software - A. Salzburger

Code optimizations for Run1/2 can allow to meet the CPU challenge (ATLAS and CMS already achieved a 5x since beginning of Run1) in trigger/tracker: not the case for HL-LHC

- Need to start a R&D but many barriers
- Small community, but many challenges ahead (HL-LHC, FCC)

ML currently used in tracking mainly for classification: just opening to pattern recognition with ML Tacking challenge

LHC detector SW has really been stress-tested: idea of starting an experiment-agnostic toolkit based on the experience gained (ACTS)

- At the same time, often used on very old idea: such a toolkit may help to test new ideas
- Data model of tracking was designed to serve both tracking experts and end-users. Lead to many compromises. Turned out not to be needed / not being used as anticipated.
- Core tracking software put into standalone package ACTS now - on CERN's gitlab. Decoupled from Gaudi/Athena, but with required adaptors to it for ease of integration.
- Doesn't want to replace/duplicate framework features: persistency, geometry... plugins for integration with frameworks. One the same line, do not duplicate G4.
- ACTS will be used by ATLAS at the core piece of its Run3 tracking
- Proof of concept by FCC connecting ACTS with DD4HEP
- ACTS example for the tracking machine learning challenge
- Hope to demonstrate GaudiHive integration in ATLAS TIM workshop (June)
- License a question still to be tackled; question to the audience for help

Tracking ML challenge: can be an important step but need to agree on what we expect from it. CMS and ATLAS have different expectations from the LVL1 trigger, need to build kind of an hybrid detector description.

## Inter-experiment ML Group and HSF - S. Gleyzer

IML founded mi-2015: community effort to modernize ML tools used in HEP

- <http://iml.cern.ch/>
- By the community and for the community
- Monthly open meeting with work on-going between meetings: continuous effort, many topics on the list of future meetings
- All LHC experiments and some non LHC (e.g. DUNE)
- Discussing various ideas, including new emerging ones
- Creating an HEP team around the topic
- Promotes common ground for testing/evaluating tools
- Contributing to providing interfaces to interesting ML tools in our SW ecosystem

Recent new activity: connection between MEM (Matrix Element Methods) and ML

- MEM experts working together with ML experts

IML working closely with CERN software group to ensure that ML packages provide good perms and are long-term supported

Tutorials: critical to attract new users

- some of them may become contributors later

IML also wants to promote/increase collaboration with ML experts

- Ensure HEP experts have knowledge about the ML tools
- Allow ML experts to understand HEP
- Good synergy with DS@LHC (DataScience@LHC) workshops
- Wants to contribute back to ML community HEP developments of general interest

IML evolved so much since its inception that an update of SW&C KB is needed! Will do it...

- IML and HSF are sharing many objectives
- Proposal made that IMF is used at the HSF forum for ML discussion: great idea, as long as we keep in mind that IMF has participants outside HSF

## RAMP: ML Hackathon on Anomaly Detection (Tue pm)

RAMP: the ML challenge idea with a group of people in the same room during 1 day

- Emphasizing collaboration rather than competition
- Objective: get codes that can be reused in real life: code accessible, not just the results
- Several RAMPs run in last year with a good success

## Packaging (Tue pm)

### Spack Presentation by Patrick Gartung. > 30 participants

Dependencies defined as a DAG: checked at installation time

Spack add-on repo created for HEP: on HEP-SF GitHub

## Geant4 technical form (Tue pm)

About 10 people in the room, more on video. Reports on G4 10.2 status, performance, em and hadronic physics developments, review of outstanding requirements being worked, detailed physics developments.

## Software performance (Wed am)

### SW Performance @ALICE - D. Rohr

Current situation: fast online reconstruction + offline reconstruction

- Most compute-intensive task is tracking
- GPU-based tracker: found the balanced config was 1 GPU per CPU
- Same code used for GPU and CPU: macros used for specific things, 90+% common code. Can use CUDA, OpenCL...
- Had to accept different floating rounding leading to a different result for a very small fraction of the clusters

Managed to get tracker CPU time increasing linearly with the number of tracks: confident it will work for Run3

Run 3 challenge: currently <2 Khz readout rate, moving to continuous readout at 50 Khz interaction rate

- Much more data: more efficient compression needed, based on reconstruction. Thus the need for full online reconstruction.
- Reconstruction requires precise calibration, thus online calibration
- New framework developed for Run 3: O2. Common development with FAIR (ALFA/FairROOT)

- Better exploit of GPU concurrency: currently the number of tracks matches the internal GPU cores but will not be the case in the future: need to process several events concurrently in the same GPU. First tests promising: improved throughput. Current setup should be able to sustain 40 Khz.

Future directions

- Remove multiple copy operations between CPU/GPU memory for the same event

## ATLAS Observations - G. Stewart

ATLAS Performance Monitoring Board (PMB) in charge of monitoring information from job instrumentation and to report/understand in significant change in performance, memory footprint...

- Every change has to be justified: many related to bugs
- A justified change generally has to be compensated by another reduction

LS1 lesson: too many things to do (implement/test) in //

Linear algebra: CLHEP too slow, replaced by Eigen after a detailed evaluation

- SIMD implementation available
- Painful migration but done implementing an abstract interface to help future migration if needed

Magnetic field access was a big CPU consumer: big impact on simulation, a lot of improvements in this area

- All together: 5x improvement in reconstruction performance

xAOD EDM: an important step to make analysis more efficient

- Extensible format, ROOT readable

Emphasis on code quality: critical to understand the code to improve it!

- Use several compilers, gcc plugin to detect some ATLAS coding standard violation
- Coverity

Run3: framework evolution, no radical change

- More concurrency, parallelism
- Gaudi work in conjunction with LHCb and FCC

Offline code reviews can be a good occasion to make progress, identify problems... and ensure that developers document their design and implementation!

## CMS - D. Lange

As others, reconstruction is the primary performance target: target, particle flow based object identification, high granularity calorimeter

- Already a lot of improvements done since the beginning of data taking to improve tracking and reconstruction time
- Need more to meet the HL-LHC challenge: computing resources are not going to match the increased time required by current implementation/approach

A lot of recent work in several different area: igprof has been an important tool to identify hotspots

- Threading optimization: Intel VTune used successfully

## Astrophysics Experience - O. Iffrig

Challenge of a fluid dynamics simulation: large number of points  $O(1000)$  in the 3 directions, 10 double precision variables

- One step depends of the previous one: synchronisation required at each step
- At least 2 steps at the same time in memory

Parallelization was a requirement: based on MPI, each process has its own data, data exchanges at their border

- Bootlenecks: communications, border data duplication, load balancing
- Currently works up to 10k processors: not enough
- On the current HPC machine used, a run lasts 10 days of WC time (not counting data transfer, maintenance, etc): limited number of simulations per campaign ( $\mathcal{O}(4)$ )

Future directions explored:

- Multiple threads per MPI tasks (openMP, pthreads...): benefit from shared memory

- Use GPUs: vectorize as much as possible, optimize data exchange

Analysis is another challenge: many algorithms I/O bound

- Try to read data once
- Parallel algorithms

Challenge of evolving a large code-base ( $O(100K)$  lines)... still work in progress

- Between 10 and 100 developers

## ROOT Experience and Challenges - P. Mato

20 year old: reengineering/rewriting required in several areas, collaboration with the community required

- Aim to get much of the changes for performances transparent for the users
- Profiling is one particular area where collaboration is beneficial: see the igprof experience

Parallelisation: multithreading and multi-processing

- PROOF remains the most used approach for perf improvement
- New MultiProc package
- Threading: new ThreadProc, same interface as multi-proc, output merging properly handled

Multi-node, on-demand analysis: SWAN

I/O perf also deserves work: exploring several approaches, including new serialization formats

Plan to exploit JIT capability of LLVM/CLANG for perf optimization at run time

Exploring functional chains à la Spark: user specifies what, ROOT decides how, providing room for optimizations

Need to explore multiple approaches to meet the challenges: not a single solution for all the use cases

## Art/LArSoft - M. Paterno

Recent experience in MC G4 following an identified problem with memory footprint: problem turned out to be more complex than initially thought, a team of experts via different profiles assembled

- People from experiments + G4 + experts in C++ profiling
- Was used as a good occasion for mentoring
- Importance of validation tests: experiment tests are not enough. Changes should not impact physics output. Was critical to convince users to use the new code.

Lesson learned

- Clear out cached data as soon as they are no longer needed in member functions
- Don't cache data into algorithmic objects
- Avoid complex objects for small lookups: high memory cost
- Concrete changes in a short time motivate people

Would be better to catch design problems earlier: main way to achieve it is collaboration during development, like for the development of an analysis

- Peer reviews

## GeantV - J. Apostolakis

Simulation represents 50% of LHC computing: GV wants to improve perf by a 2.5 to 5x

- Redesign code with fine-grained granularity at the core
- Scalar & vector interfaces (efficient use of SIMD instructions)
- Concentrate first on geometry: 30-40% of the simulation CPU

Every component/class has a test and a benchmark both for scalar and vector interfaces

- Benchmark results compared with Geant4, ROOT, USolids

Importance of I/O

Basketization is another critical part

- Includes a lock-free scheduler

Several developments could be backported to G4.

# Discussion panel

Graeme: hard limit to perf improvement set by HW. Changes in HW have a much higher latency than SW

- Ultimate metrics is event throughput

Olivier: relying on HPC, the HW question is different. Resources provided by national facilities. No control of it, GPUs will be part of next generation of machines.

D. Rohr (D.R.): GPUs are here, no discussion about using it. The question is the right balance between GPU and CPU: may change over time, also a chance to do it incrementally.

Liz: emphasize Oliver's point of view. Same in the US: HEP encourage to join the HPC community and the HPC community has a roadmap that includes GPUs as a significant part of the next machine generation. Need to deal with HW heterogeneity increases the pressure on the build/packaging system.

Amir: need to work together as a community around the R&D about these issues and come with a common framework that could be used as the basis for the computing infrastructure in 10 years. HSF could be the good place for this effort.

Pere: first step with GPU is to demonstrate the gain and for this, we need to identify the areas to concentrate on.

D.R.: easier to rewrite/adapt a specialized application (like tracking) than frameworks like ROOT and GEANT

Vincenzo: is there still a place for commodity HW? Pressed to join the HPC community: at the same time, less opportunity for commodity with automatic power-off capabilities in new HW... ARM vs. HPC: may achieve the same throughput without the big pressure on parallelization.

Graeme: no doubt we'll have to rewrite a significant part of our code but it needs to match a large part of the architecture phase space. Need to rely on compiler and build systems to help with this.

J. Apostolakis (J.A.): impossible to say what will be the dominating infrastructure 5 years from now, need to remain flexible and be able to support multiple infrastructures at a low cost.

Olivier: current approach is to put the implementation details for each architecture in libraries and hide them from the users. No clear standard that could simplify the problem: still need to adapt to each architecture.

D.R.: ALICE took the opportunity of the requirement to rewrite the tracker to redesign it with GPU in mind. But originally not written for GPU.

Amir: expect to see different GPU HW for gaming/deep learning (16-bit precision for higher perf) and HPC (double precision). AMD being out of the HW business, it is in a better position to change the SW landscape around these new HW (well aware of the need to support many HW/programming environments).

Pere: need to prevent users specifying low-level things. If the user gives a high level description, it is easier for the system to do internal optimization.

Amir: need to look at what big players did to offload things like compression to GPUs. Highly connected with I/O and efficient I/Os from GPUs.

Pere: we need to assess the exact impact on the overall workflow of offloading one particular part to specialized HW. And compare it with the effort to support them.

Amir: HSF could have the role of identifying the needed R&D and convince funding agencies to support it. We need topical workshops or topical sessions in regular meetings. Need to be proactive.

M. Sokoloff: if current NSF proposal moves forward, we'll need to produce the Community White Paper in the next 15 months: The HSF is the natural organisation to do it.

Jeff: we need a metric for HSF work: maximize efficiency of overall people contributions.

Liz: ultimate goal remains to get additional people, not to play a zero-sum game.

Dario: importance of trainings to meet the challenge by increasing existing people expertise.

## Next steps & wrap-up (Wed pm)

### Logo

- Thank you to all contributors for the high quality entries!
- Thanks particularly and congratulations to Joschka Lingemann, CERN Fellow working with EP-SFT, who produced the winner preferred by a strong plurality (17 votes in 37 responses). A bottle of wine will be forthcoming for Joschka :-)



### Meeting notes

- Michel, in particular, has kept comprehensive [meeting notes](#), thank you!
- Next HSF Newsletter (next week?) will provide a summary of the workshop

- Clear message from the [doodle](#) is preferred date is week of Oct 24-28, probably early that week. Location not yet firmly decided, probably CERN.
- If this doesn't work for you, and you very much want to attend, let us know

## Journal proposal “Computing and Software for Data-intensive Physics”

- [Christian Caron, Springer, presentation](#)

Proposal to have a journal refereed, abstracted, indexed about HEP computing

- Authoritative and central reference archive
- Help with career paths
- Do not restrict to HEP strictly: open to Data-Intensive Physics. Fields organized as large collaborations around large-scale experiments
- Scope should cover all aspects of computing: from infrastructure to data analysis...
- Continuous publishing: no paper/volume

Several open questions

- Exact size/focus
- Business model: hybrid (OA access per article, based on authors choice/constraints) vs. pure OA. Pure OA requires sponsorship.
- Article types: regular articles, reviews, advanced tutorials, no letter, no proceedings
- Editorial structure

Discussion

- Allow publishing software (with DOI) or only articles about software?

- May be useful to recognise merit of individuals, not so much internally in HEP but rather for people who will leave the field
- Should be wider than HEP: we have more and more collaborations with other communities, we have common issues. At the same time must be somewhat restricted to people with similar (scale) problems.

## HSF in StackExchange?

- Should we set up `hepswcomp.stackexchange.com`?
- If yes, who will initiate it and take care of it?
- Need an initial set of 40 questions
- Andrea will try to bootstrap it

## Open questions

- Public vs. competition - one experiment will launch a private Q&A site with stackexchange look&feel, internal to only collaboration members

## HSF as legal entity

- Proposal from Elizabeth Sexton-Kennedy and Fons Rademakers to seed HSF as a legal entity with an anchor project -- much as Apache Software Foundation did with httpd -- namely ROOT
  - [Slides from Liz](#)
- Do this in parallel with -- and without interference with -- HSF's ongoing activities
- Would imply adding a Foundation Board, otherwise should not necessarily drive governance changes
- Main motivation is to have an entity that can take ownership of SW contributions in the community: clear ownership required to allow clear licensing and collaboration with other fields.
  - Has to be neutral
  - Need a "core product" like Apache Server for ASF: ROOT?



- Need to find sponsors, including funding agencies

Vincenzo: counter-example of Geant4, with a different form of organization, namely the G4 collaboration

Samir: if we want to increase collaboration with others, not clear that they all want to give up their IPR.

Andy: supportive of HSF being able of taking ownership of IPR but skeptical that existing big projects with huge IPR are good gunieapigs.

Markus: any Foundation would be established under a national law; would US institutes/agencies agree to give IP to a foundation under EU/CH law, or viceversa?

## HSF governance

- So long as HSF remains an entirely unresourced volunteer activity, there is no motivation or justification to go to a heavier governance
- If/when the HSF becomes 'resourced' somehow, governance will need to be addressed

## HSF Center

- A suggestion from SLAC (Andy Hanushevsky): invite institutes to establish an 'HSF Center' as an anchor for their material (not person power) contributions to HSF
- E.g. if institute X supports a Jenkins service for the use of HSF & associated projects, the contribution comes as part of their HSF Center
- Would allow both the HSF and the institute to clearly identify contributions
  - For the institute, a means of clearly identifying participation and contribution to HEP S&C common projects via HSF support

- For the HSF, concretely identifies contributing institutes and their contributions, and can help incite more
- Maintain a list of services we would like Centers to pick up

## HSF (Human) resources

- What can/should be said, concluded, done?
- Make a wish/priority list for dedicated effort?

To be defined by CWP & roadmap

## HSF communication

- Do people feel informed enough about what we do?
- Should we continue like now or do something differently?

## Community white paper & road map

- Broad community white paper that can be referenced, used, distilled, adapted for particular uses.
  - Such as the the white paper output of the NSF SI2-S2I2 project conceptualization phase
  - Such as WLCG's HL-LHC directed computing planning?
- Seems there is consensus that the HSF is the right place to do this
- An investment towards gaining new resources through successful proposals
- Look soon at what has to be investigated, what R&D is needed, do we need task forces to be formed, ...

## Discussion

- John: LHCC asked WLCG to write a TDR about HL-LHC computing. Same timescale, covering similar issues
- Amir: collect the list of problems/challenges already existing

- Peter: need more than a list of problems and possible solutions, funding agencies/governance bodies would like to see a roadmap.
- R&D vs. production SW: roadmap should define the steps
- M. Sokoloff: one important goal of CWP is to answer requests from funding agencies that experiments develop common solutions to common problems. At least common elements that can be picked/assembled by the different experiments.
- Kickoff meeting in the US this Fall: “co-locate” with CHEP? Peter: should be driven/branded by HSF.

## Outcomes, conclusions, next steps

- New activity around SW performance? Web site area? Should happen inside the SW technology forum but give more visibility to it in the “HSF space”.
- Project template, packaging
  - The way the Mon morning discussion developed, towards managing a software stack and the issues taken up in the Packaging WG, suggests that we’re on the right track with the Packaging WG’s investigations into solutions
  - Great interest in Spack but recognition in the packaging discussion that there’s no silver bullet
    - Bigger interest by smaller projects/experiments not having the resources of LHC experiments
  - Discussions here affirm the usefulness of the document describing and charting the features, pros, cons of the alternative
  - Input here towards a V2 of the document
    - Explicit request to add conda to it.
  - Project template considered a good idea to show/implement best practices
    - Should be maintained actively

- Feature request for putting a ROOT example in there
- Common github repo created for next-gen conditions DB
  - <https://github.com/HEP-SF/PhysCondDB>
  - Group: <https://github.com/orgs/HEP-SF/teams/conditiondata>
- HSF status, directions, scope
  - White paper
  - Once we have HSF road map, consider seeking ICFA (& ECFA) endorsement
    - ECFA workshop coming up this fall, involve HSF to have a role in this in long range computing planning?
  - No new things have popped up that ought to be added to HSF scope?
    - (Distributed software/systems is out there, one day...)
  - ML is the hot item of the day, and there we have the IML here today, and whatever relationship with the IML develops in the future. IML will use SW KB. Could also be part of HEP StackExchange.
- Next F2F: CHEP
  - Report status, discuss hot topics, look for ways to leverage CHEP content
  - Investigate possibility of a plenary slot
  - A couple of hours? No approaches made yet to CHEP to schedule it
  - **Investigate the possibility of the CWP kickoff meeting at the beginning of the week after CHEP (or before...)**
- Next workshop like this: ~ a year from now
  - Should it be structured more or less like this one? This format seemed to work well.
- Smaller topical workshop(s) according to need

# Actions out of the workshop

- Investigate the possibility of a community white paper meeting at the beginning of the week after CHEP, or before CHEP, in the Bay Area (SLAC willing to host)
- Andrea Valassi will follow up on trying out the idea of [hepswcomp.stackexchange.com](https://hepswcomp.stackexchange.com); first will be to come up with 40 typical questions
- Liz will pursue questions raised at the workshop about making the foundation a legal entity.