Summary of HSF Workshop 2017 in Annecy

Benedikt Hegner

EP/SFT Group Meeting, 17.7.2017

- Context of the CWP and the Workshop
- The HSF Workshop
  - Goal
  - The Working Groups
  - Plenaries
- Whitepaper Status & Next Steps

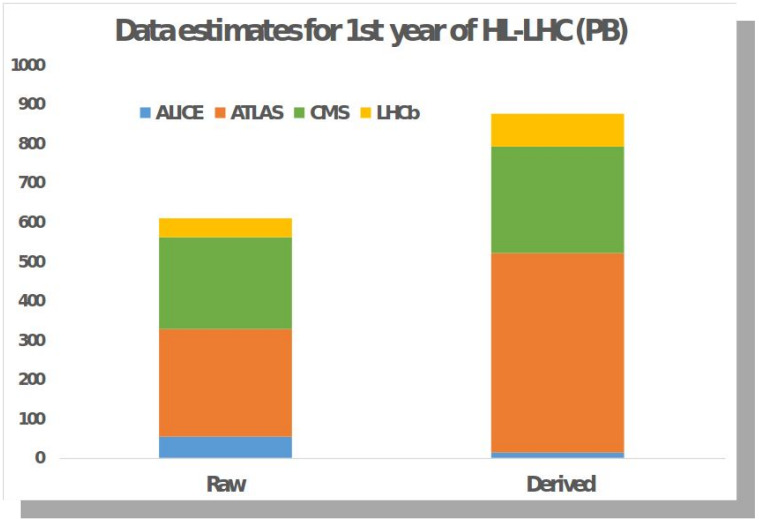

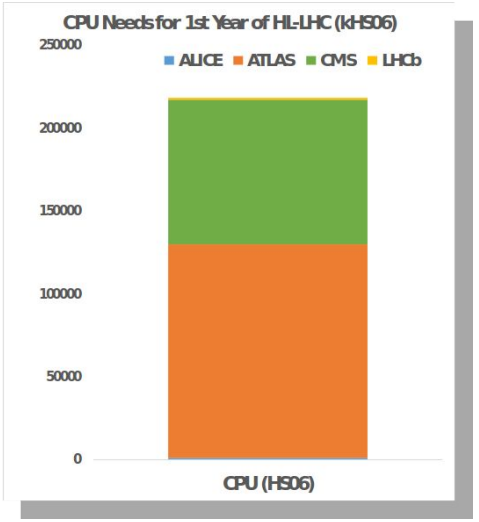

Data:

- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU:

- x60 from 2016

Hardware technology at ~20%/year will bring x6-10 in 10-11 years (assuming a flat budget)

Missing factor of 10 in CPU has to come from SW innovation

See presentation by Ian Bird:
http://indico.cern.ch/event/570249/contributions/2404394/attachments/1400414/2136977/LHC-Future-Challenges-CWP.pdf

- A **Community White Paper (CWP)** should describe a global vision for software and computing for the HL-LHC era and HEP in the 2020s

- The CWP will identify and prioritize the software research and development investments required:
  - to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
  - to enable new approaches to computing and software that could radically extend the physics reach of the detectors
  - to ensure the long term sustainability of the software through the lifetime of the HL-LHC

- The HSF is engaging the HEP community to produce the CWP via a community process
  - Initiated as an HL-LHC planning process
  - Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider, ...)
  - The resulting roadmap will be used for the HL-LHC computing TDR and other strategic plans

CWP Kick-Off Workshop in San Diego (23.-36. January)

- ~110 participants, mainly US + CERN
  - Unfortunately very few Europeans from outside CERN
- 2.5 days of parallel topical WG meetings
  - Agenda : http://indico.cern.ch/event/570249/timetable/#all
  - From infrastructure to reconstruction and analysis, through simulation, data management…
  - Notes from (almost) all WG discussions in the WG Google Docs, summary slides in the agenda

This was the beginning for many other topical workshops

Last month: (almost) concluding HSF Workshop at LAPP/Annecy (26.-30. June)

90 participants:

- US: 48 (8 FNAL)
- CERN: 20 (7 EP/SFT)
- France: 14 (7 LAPP)
- Italy: 3
- UK: 2
- Germany: 2
- Switzerland: 1

Organization:

- Indico Agenda
- Monday: Introduction+Status
- Tuesday: Parallel Sessions of WGs
- Wednesday: Plenaries+WGs
- Thursday: Plenaries+WGs
- Friday: Closeout

Goal was to progress with the CWP preparations

Outcome of each WG should be a roadmap to HL-LHC with objectives for 1, 3 and 5 years:

- 1 year prototypes and initial studies
- 3 year studies to give input into LHC experiment TDRs
- 5 year real projects to deliver software for high luminosity

Leading principles to the working groups:

- catalyse common projects
- promote commonality
- attract new effort
- set priorities

Links between the groups should be made for a coherent approach

e.g., training and machine learning are really cross cutting themes

Working Groups active during the workshop

- SW Trigger and Reconstruction
- Machine Learning
- Data Access, Organization and Management
- Software Development, Deployment and Verification/Validation
- Data Analysis and Interpretation
- Conditions Database
- Data and Software Preservation
- Event Processing Frameworks
- Physics Generators
- Workflow and Resource Management
- Visualization
- Computing Models, Facilities and Distributed Computing
- Careers, Staffing and Training

Full list of all working groups and their working documents:
http://hepsoftwarefoundation.org/cwp/cwp-working-groups.html

When High Performance Computing Meets Astronomy - a Concrete Case (Pierre Aubert et al.)

Gamma-Ray Astronomy with a new *Cherenkov Telescope Array (CTA)*

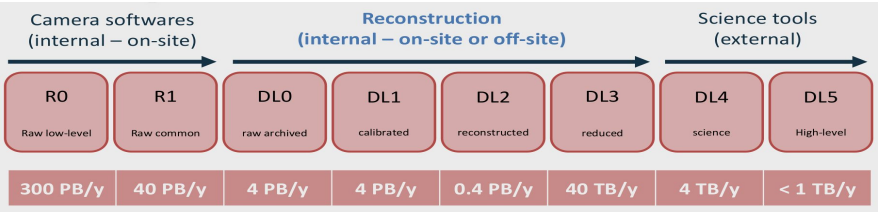

Vectorization of existing algorithms

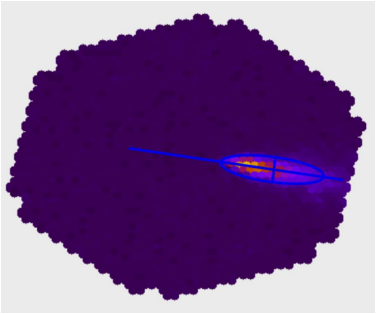

Data format ×40

Vectorization of the reduction

	Speed (cy/el)	Speed up
Initial data format	2125.5	1
Optimized data format	53.1365	**40.0**
+ Intrinsic Vectorized SSE4	6.39931	332
+ Intrinsic Vectorized AVX	**2.98499**	**712**

Fruitful cooperation of HPC and Physics Experts

Combination of methods for faster compression

Raw data compression and speed

Test File
- 475 MB
- Up to 99% ADC values

	Compression ratio	Time	File size (MB)
LZMA (7Z)	4.84	7 min 48s	98
Advanced Polynomial Compression	3.74	3.7 s	127
Advanced Polynomial Compression + LZMA	4.84	24.646 s	98

Same Compression ratio

19 × faster

The Elements That Are Driving Innovation (Tom Gibbs, NVidia)

Contrasting classical and artifical intelligence approaches and performance

HPC	AI
+40 years of Algorithms based on first principles theory Proven statistical models for accurate results in multiple science domains	New methods to improve predictive accuracy, insight into new phenomena and response time with previously unmanageable data sets

Training a Deep Neural Network DNN and running it on dedicated resources can give both better results and computational performance

Transformation HPC + AI couple simulation with live data in real time detection/control system	Augmentation HPC + AI combined to improve simulation time to science > orders of magnitude	Modulation HPC + AI combined to reduce the number of runs needed for a parameter sweep
Experimental/simulated data is used to train a NN that is used to for detection/control of an experiment or clinical delivery system in real time. The NN is improved continuously as new simulated / live data is acquired	Experimental/simulated data is used to train a NN that is used to replace all or significant runtime portions of a conventional simulation. The NN is improved continuously as new simulated / live data is acquired	Experimental/simulated data used to train a NN which steers simulation/experiment within/btwn runs The steering NN can be trained continuously as new simulated / live data is acquired

Concrete example from e.g. LIGO was a x1000 speedup in data analysis

Designing for Data Access (Gaurav Kaul, Intel)

Problem of *Memory Wall*

- Memory capacity and bandwidth increasing slower than number of cores
- So far resulted in ever increasing cache infrastructures (and energy consumption)

Various emerging technologies

- Embedded DRAM, 3D memory stack, …
- Execute arithmetic and address ops nearer to memory subsystem and exploit parallelism in memory access (NDP)

Despite their quality, slides worth a read!

Designing for Data Access (Gaurav Kaul, Intel)

NVME revolution: a new layer in the storage hierarchy

- 3D XPoint: combining a latency close to DRAM and a capacity close to NAND

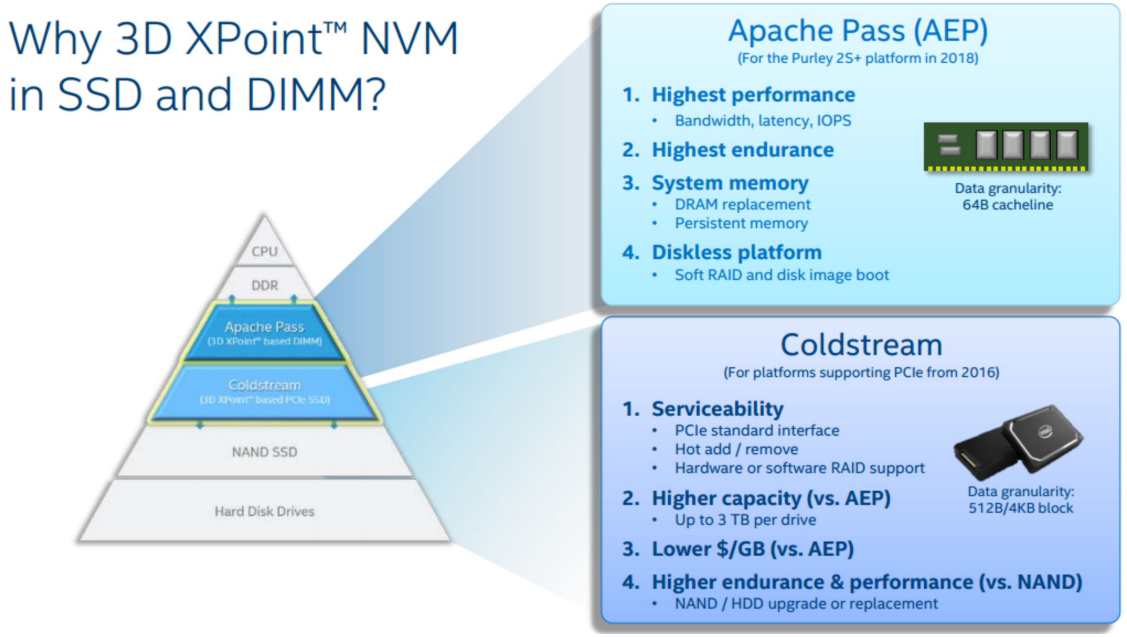

- Several usage modes: memory, storage, direct control by application (open-source DAOS library)
- Project of 1U JBOD based on XPoint with a 1 PB capacity

In total more than 200 pages of text and figures produced so far

- Ranging from bare LaTeX-Template up to almost finished document
- Vastly different levels of detail and length
  - From ~20 up to 59 pages for “content-complete” drafts
- Most documents describe the status-quo and the problems nicely
- The quality on the 1,3 and 5 year goals differs though
  - Some of the working groups laid out visionary plans to try new approaches
  - Others conservatively extrapolate existing approaches

Please have a look at the existing documents and give feedback to the authors

You are part of the community and your voice counts!

My personal opinion

We should make sure no chances for innovation are forgotten

The need for innovation was one of the reasons why the CWP process was started

Finish the documents with the aim of

- Keeping the reports short so that people can actually read it, aim for 30-50 pages including executive summary
- Including an overall roadmap to have some prioritisation across the WGs
- Making results available on ArXiv

Proposed timeline:

- Finalized WG chapters by end of July
- Finalized CWP by end of August

What was still unclear during the workshop

- How to incorporate feedback by LHC experiments? Do we need any form of sign-off by them?
- Who are actually the authors of the CWP? Authors and supporters?

Please note:
Based on the CWP there will be an independent TDR for the WLCG end of the year with explicit sign-off by the LHC experiments.