Report on San Diego HSF CWP workshop

V. Gligorov, CNRS/LPNHE

15th May 2017

LPNHE PARIS

cnrs

Workshop layout and aims

Community White Paper (CWP) Working Groups

During October, 2016, the process of forming CWP working groups has begun. We are letting the working groups self-organize with a bit of help from the HEP experiment software/computing coordinators and the HSF startup team. The first step for each working group is to formulate a charge to describe the challenges in that area and questions which need to be answered to produce a roadmap for the CWP.

The working group charges should be world-visible for reading. If you would like to edit them, you will need to join the hsf-community-white-paper Google group and login as your google user. If you would like organize an additional working group, please post to the hsf-community-white-paper Google group.

Note: it is possible to register to all the Google groups without a Google account. In this case, send an empty email (subject and contents ignored) to the group email, adding `+subscribe` after the group name. (e.g. hsf-cwp-software-development+subscribe@googlegroups.com).

Big objective : organize a sustainable community around the big themes in HEP computing for the next decade

Key milestone : deliver the "Community White Paper" (CWP) to the NSF by this summer

Role of wkshp : build bottom-up groups of interested people around each theme, identify conveners who can push each theme forward, develop a "charge sheet" which said group should answer for the CWP

Why HSF and why the CWP?

A Software "Upgrade" for HL-LHC and 2020s HEP?

Looking forward to the next 10 years, we see a number of challenges for HEP software and computing:

- **Scale:** The HL-LHC will integrate 100 times the current data, with significantly increased data (pileup) and detector complexity.
- **Performance/cost:** Estimates of computing needs run faster than Moore's Law by factors of 3-30
- **Technology/Market evolution:** the return of heterogeneity; technology change will also make it challenging to exploit Moore's Law without software evolution.
- **Sustainability:** Most of the current software, which defines our capabilities, was designed 15-20 years ago: there are many software sustainability challenges.

"Computed hardware is a consumable. Software is what we keep, and invest in, over time." — P. Elmer

Groups/objectives

Computing Models, Facilities, and Distributed Computing WG

Google Doc

Google group: hsf-cwp-models-facilities@googlegroups.com (link to subscribe)

Plans:

Detector Simulation WG

Google Doc

Google group: hsf-cwp-simulation@googlegroups.com (link to subscribe)

Plans:

Software Trigger and Event Reconstruction WG

Google Doc

Google group: hsf-cwp-swtrig-evtreco@googlegroups.com (link to subscribe)

Visualization WG

Google Doc

Google group: hsf-cwp-visualization@googlegroups.com (link to subscribe)

Data Access and Management WG

Google Doc

Google group: hsf-cwp-data-management@googlegroups.com (link to subscribe)

Security and Access Control WG

Google Doc

Google group: None yet

Plans:

Machine Learning WG

Google Doc

Google group: hsf-cwp-machine-learning@googlegroups.com (link to subscribe)

Plans: The WG has 2 planned workshops in the coming months.

- a CWP session during the IML topical workshop at CERN, March 20 - 22, 2017
- a CWP session (TBC) during DS@HEP 2017, FNAL May 8 - 12, 2017

Conditions Database WG

Google Doc

Google group: hsf-cwp-conditionsdb@googlegroups.com (link to subscribe)

Plans:

Event Processing Frameworks WG

Google Doc

Google group: hsf-cwp-event-processing-frameworks@googlegroups.com (link to subscribe)

Plans:

hepsoftwarefoundation.org/activities/cwp.html

Groups/objectives

Physics Generators WG

Google Doc

Google group: none yet

Plans:

Math Libraries WG

Google Doc

Google group: hsf-cwp-math-libaries@googlegroups.com (link to subscribe)

Plans:

Software Development, Deployment and Validation/Verification WG

Google Doc

Google group: hsf-cwp-software-development@googlegroups.com (link to subscribe)

Plans:

Data Analysis and Interpretation WG

Google Doc

Google group: hsf-cwp-analysis@googlegroups.com (link to subscribe)

Plans:

Workflow and Resource Management WG

Google Doc

Google group: none yet

Plans:

Data and Software Preservation WG

Google Doc

Google group: none yet

Plans:

Careers, Staffing and Training WG

Google Doc

Google group: hep-sf-training-wg@googlegroups.com (link to subscribe)

Plans:

Data Acquisition Software WG

Google Doc

Status: not yet active

Google group: none yet

Plans:

Various Aspects of Technical Evolution (Software Tools, Hardware, Networking) WG

Google Doc

Status: not yet active

Google group: none yet

Plans:

Monitoring WG

Google Doc

Status: not yet active

Google group: none yet

Plans:

hepsoftwarefoundation.org/activities/cwp.html

Some observations

Some groups attracted a particularly large amount of interest

Machine Learning : very well organised by the IML and DS@HEP people, extremely broad range of people attended the discussion sessions. The charge sheet is packed and could probably consume dozens of FTEs on its own.

Some observations

Some groups attracted a particularly large amount of interest

Machine Learning : very well organised by the IML and DS@HEP people, extremely broad range of people attended the discussion sessions. The charge sheet is packed and could probably consume dozens of FTEs on its own.

Data analysis : very lively discussions. A general rejection of “real-time analysis” as actually meaning analysis by the GPD people (to the point where they wanted to define analysis as meaning something done by individual users, not in productions). Quite a few people didn’t really believe that we were serious about not reprocessing and seemed to assume this was just a gimmick.

Some observations

Some groups attracted a particularly large amount of interest

Machine Learning : very well organised by the IML and DS@HEP people, extremely broad range of people attended the discussion sessions. The charge sheet is packed and could probably consume dozens of FTEs on its own.

Data analysis : very lively discussions. A general rejection of “real-time analysis” as actually meaning analysis by the GPD people (to the point where they wanted to define analysis as meaning something done by individual users, not in productions). Quite a few people didn’t really believe that we were serious about not reprocessing and seemed to assume this was just a gimmick.

Trigger&Reco : started out as individual WGs but merged owing to the overlapping interests of the people. Will include real-time analysis as the data analysis WG explicitly refused to have it in their charge sheet. Convened, whatever that means, by David Lange and myself.

See dedicated slides by Gloria and Paul on the simulation and ML working groups

Will now cover some of the interesting points from hardware, facilities, and reconstruction in particular

Key points on hardware

Some of the most useful material came in the plenary talks and panel discussion on future hardware architectures. For example, Moore's law is over because there the hardware market is saturated… ok that's a slight simplification…

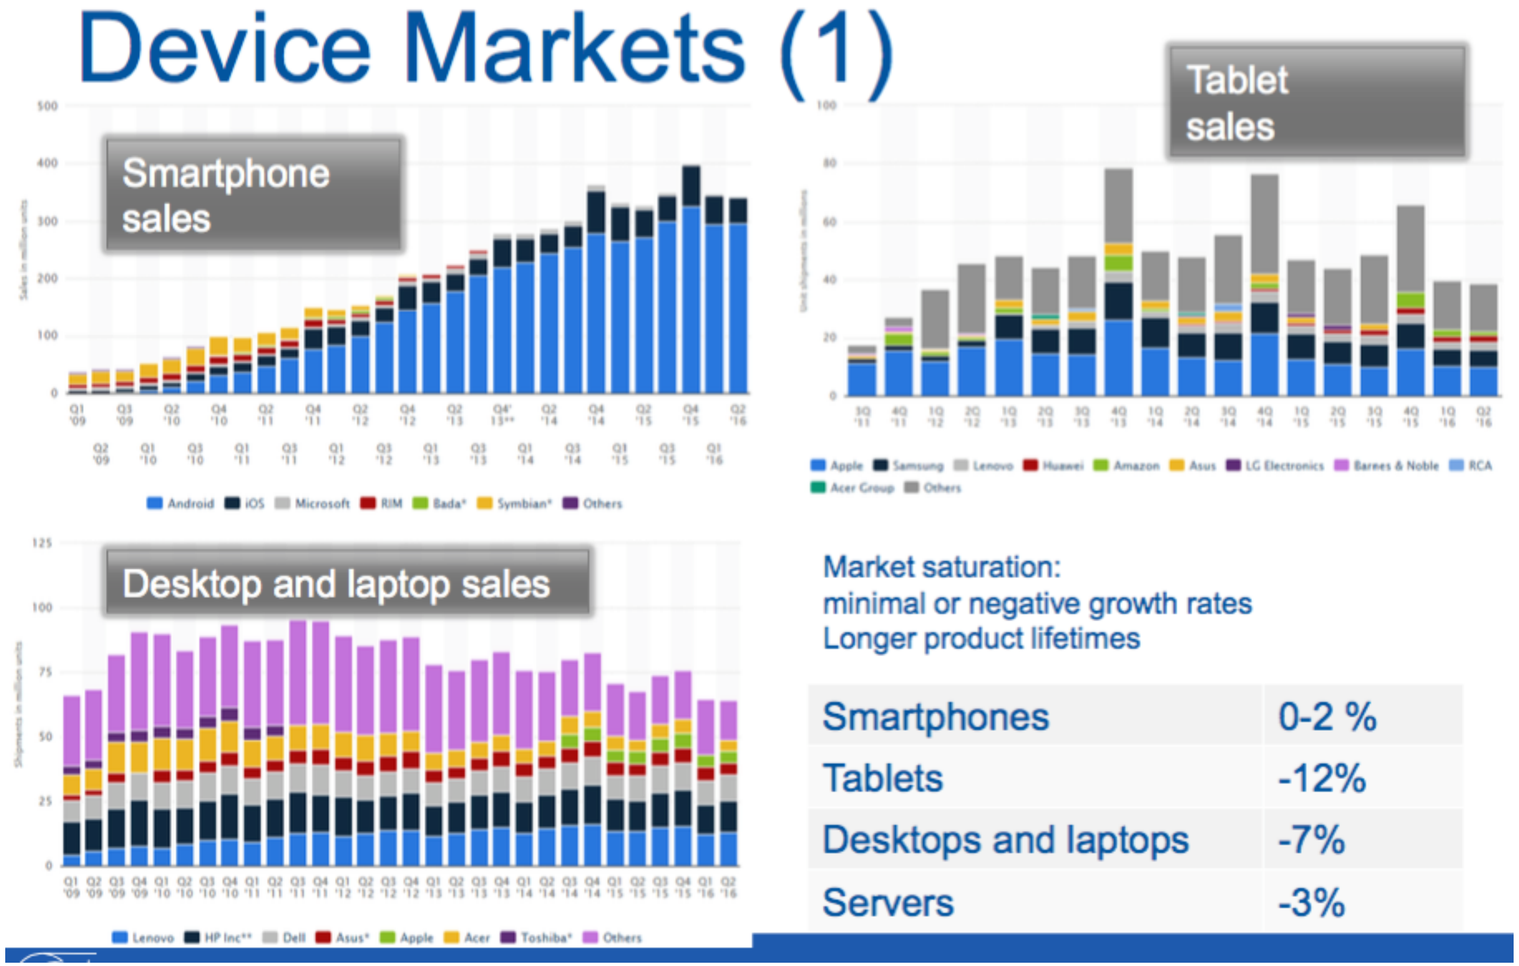

Tick-tock-tick-…-tick-…-…-…?

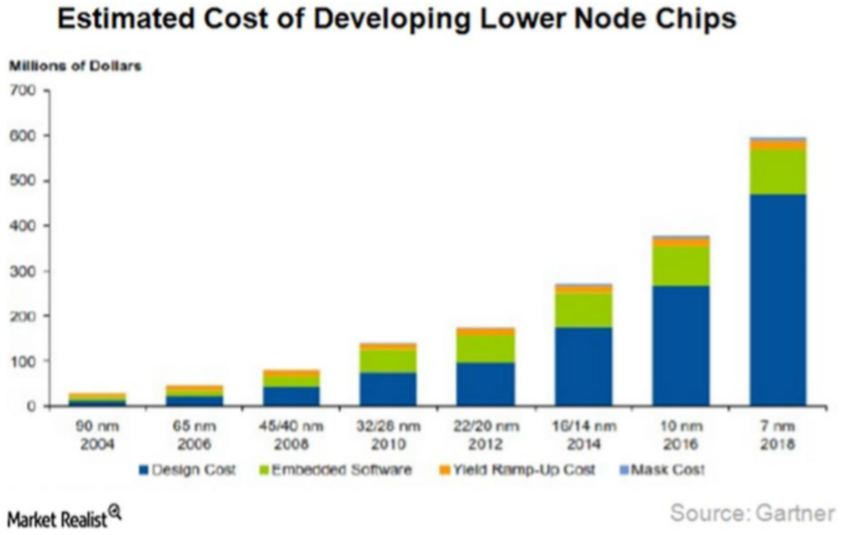

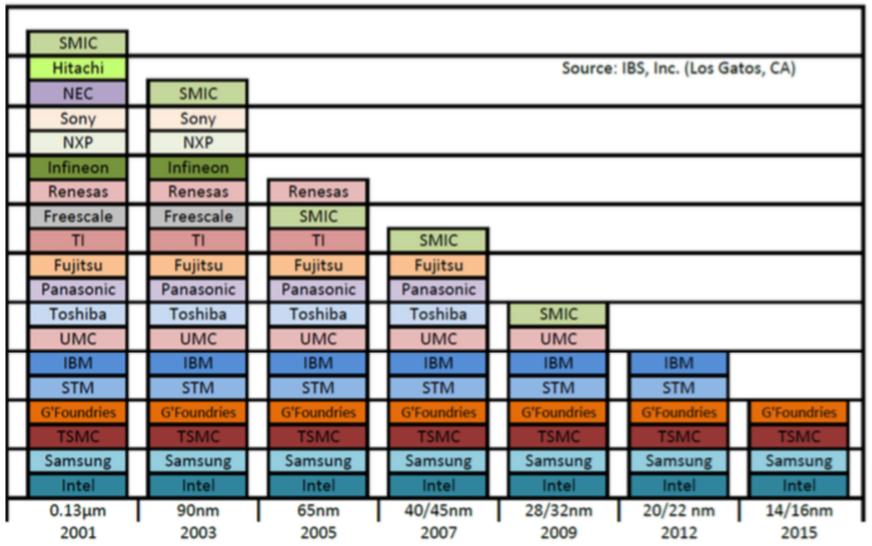

Figure 4. Dramatic Consolidation of state of the art CMOS Fabs. Source: IBS , Inc. (Los Gatos, CA).

Non-linear costs for development

- Only four companies able to fabricate 14 nm chips
- 10 nm Samsung fab costs $14 B

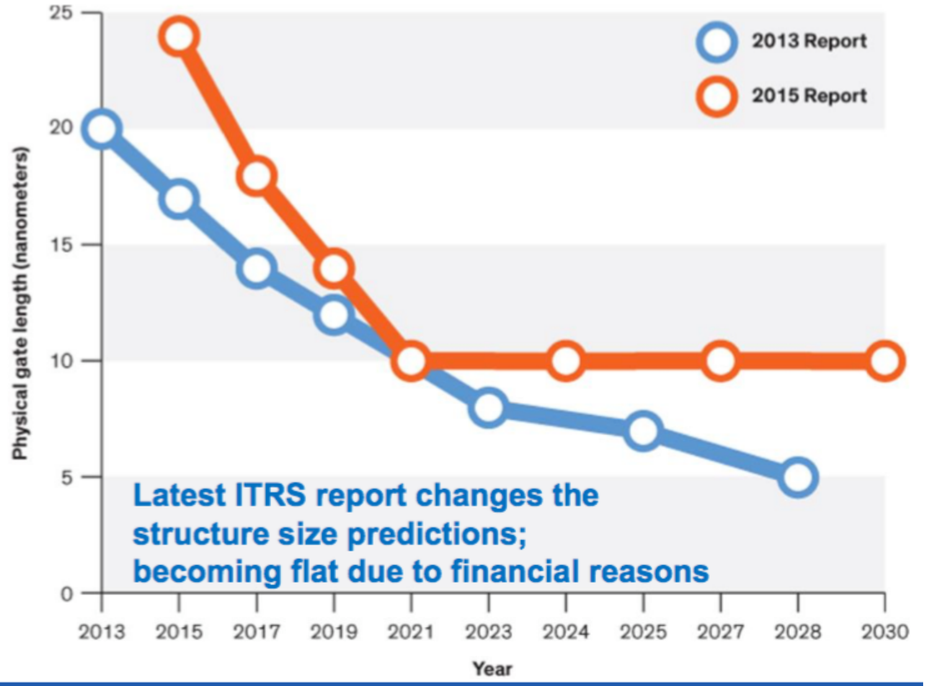

Who needs precision anyway?

Architecture fragmentation, backed up in panel discussion : focus will be more and more on specialized instruction sets for particular problems.

- New focus for graphic cards: machine learning
- Move to FP16 and even INT8 architectures, less precision → 8 bit processing !
- Google TPU Tensor Processing Unit
- New start-ups with special processor designs: e.g. KnuEdge, Nervana (just bought by Intel), krtkl, Eyeriss
- Essentially not usable as general purpose processors (online?!)
- Intel changing strategies also for their KnightsXX processors, 'forking' models (increase FP16 and decrease DP)

  ~100k units per year, very small market

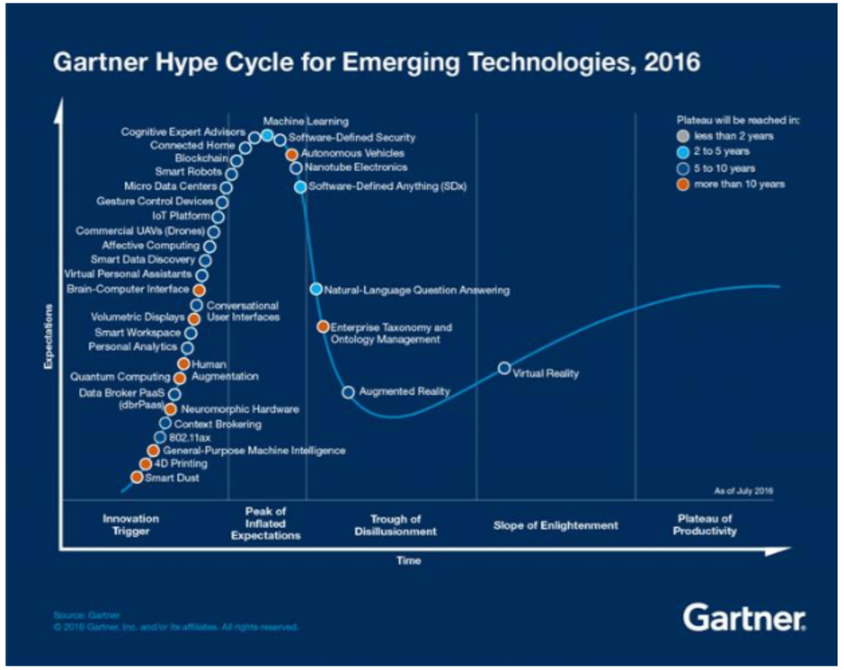

- Qualcomm plans to add neuromorphic chips into the smartphone

Disque d’or

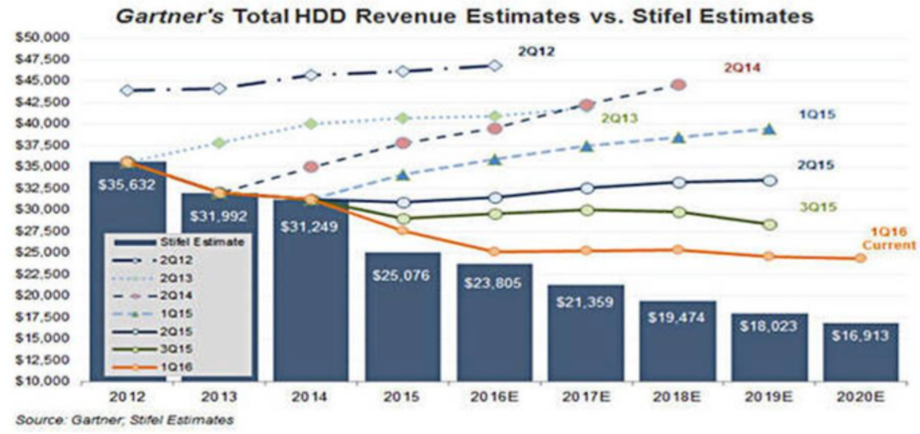

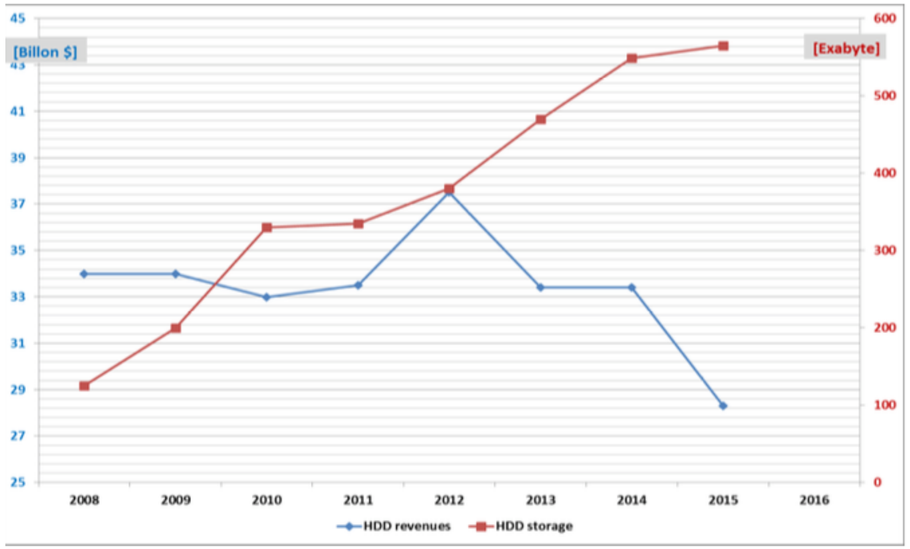

Areal Density Trends

Chart provided courtesy of the Information Storage Industry Consortium (INSIC)

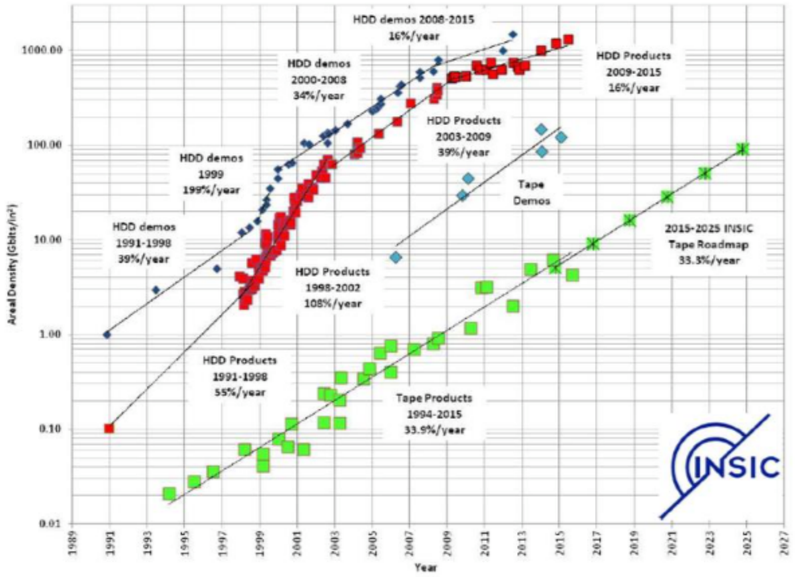

Areal density improvement dropped from ~40% to 16% per year

Overall hardware summary

- Moore's Law and Kryder's Law are slowing down
  - 18 months → >= 3 years
- Real cost/performance evolution driven by financial and market aspects rather than technology

Translation : we are a tiny market hence totally screwed by this.

Clouds everywhere

Most of the slides from CERN about infrastructure were about clouds. A personal comment, also driven by the way AFS is being phased out : there seems to be a desire to cut costs by shifting the analysis model to laptops+cloud, which will hurt the poorest groups which don't have local clusters the most.

An interesting tension between the analysis WG which seems to want to treat analysis as "what individuals do" and the drive from facilities side towards production style work for anything but the most cursory data inspection.

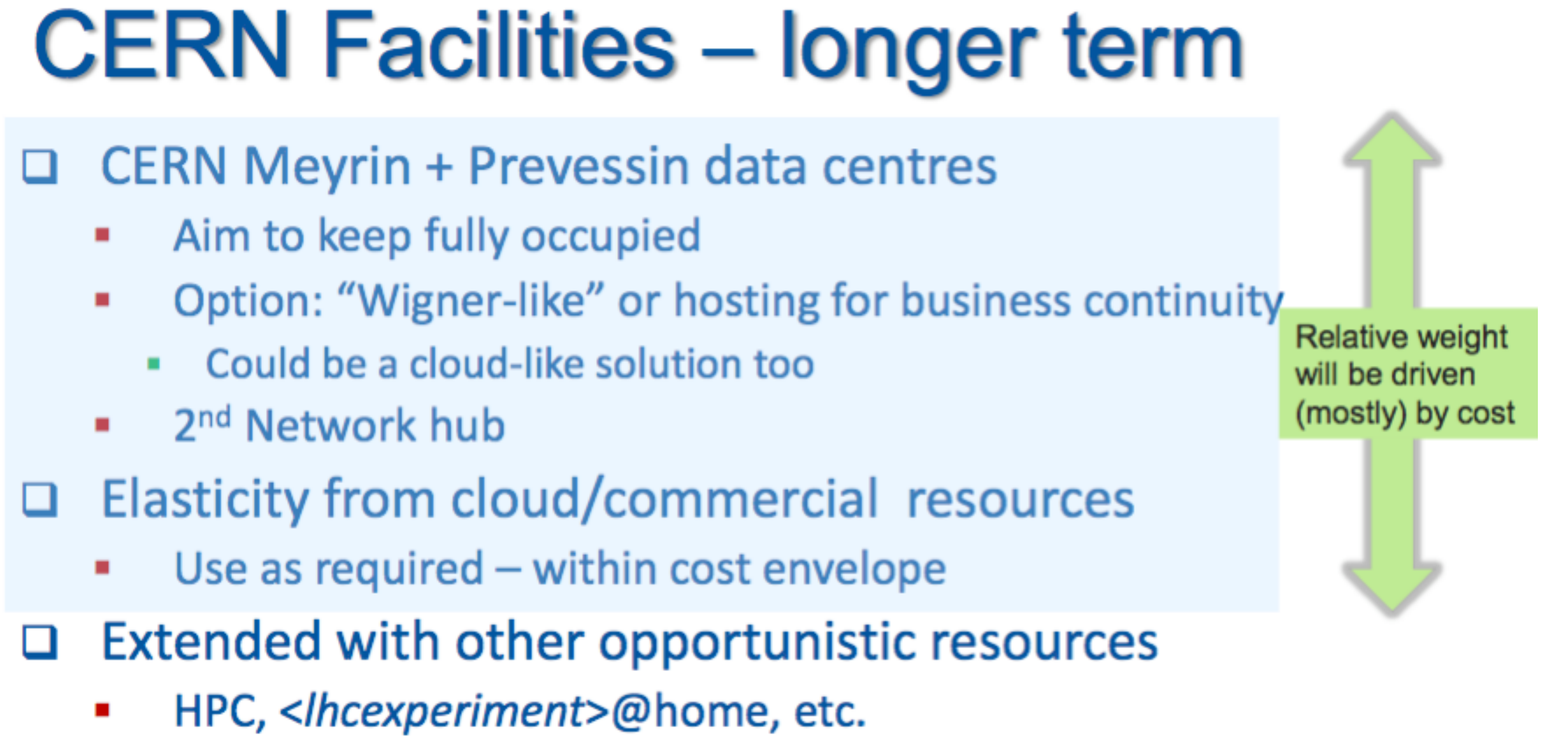

Options for Future Computing

- **Example: LHC Run 4 (2026 and beyond) will start the exabyte era for HEP!**
  - **How will the data be processed and analyzed?**

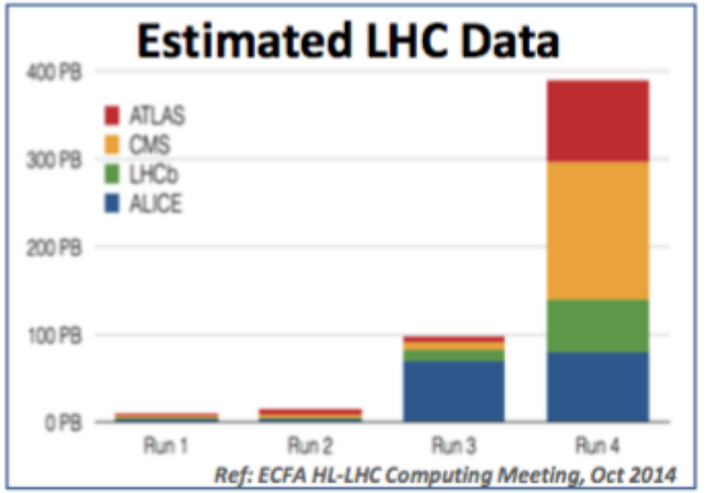

☐ **Buy facilities**
- ✓ **Pro: Own it! No impediment to running at full capacity when needed**
- ✗ **Con: Must invest for peak utilization even if not used**

☐ **Use services from other providers**
- ✓ **Pro: Others make capital investments**
- ✗ **Con: Will usage be available/affordable when needed?**

- **Current computing model provides examples of infrastructure not owned by HEP**
  - **Like ESNet, NERSC, commercial clouds; not necessary for HEP to purchase all hardware**

☑ **Hybrid model**
- ✓ **Own baseline resources that will be used at full capacity**
  - **Reliable cycles available for reconstruction, MC generation, etc.**
- ✓ **Use service providers for peak cycles when needed**
  - **Conference analysis season, special collaboration needs, etc.**
- ? **Community and agencies exploring this approach, *but* future cost model uncertain**

- **To achieve P5 global vision, all partners need to bring in their available resources!**

Options for Future Computing

A lot of chat about peak vs. on demand use, capital investment, and so on.

In my personal view, this ignores or sidesteps the tradeoffs between disk space and CPU usage which will become more relevant as we go on into the HL-LHC era.

Partly I mean that the less disk you have, the more you have to use CPU. But also what I mean is that e.g. we could easily burn all our CPU 10x over generating MC all the time for analysis, but we wouldn't have the disk space to store it even if we did.

However as we move towards analysis trains and true real-time reproducible analyses, we will in principle be able to generate all the MC we want, run it through the analysis, and then discard it (or keep a heavily reduced analysis format for reproducibility). So should never be idling anywhere.

This is also a bigger problem for LHCb because I can see more and more precision analyses asking for MC closure tests before approval, which in the case of things like Charm will mean infinite samples and CPU time even with "fast" options. This is a real difference with ATLAS/CMS and should be communicated better to outside world.

Reconstruction

My favourite talk was on the FCC : insane stuff

Completely absurd without 4D tracking, likely silly even with this.

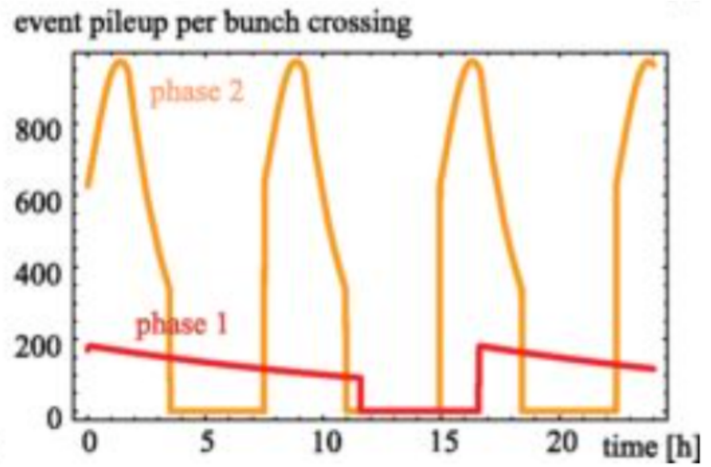

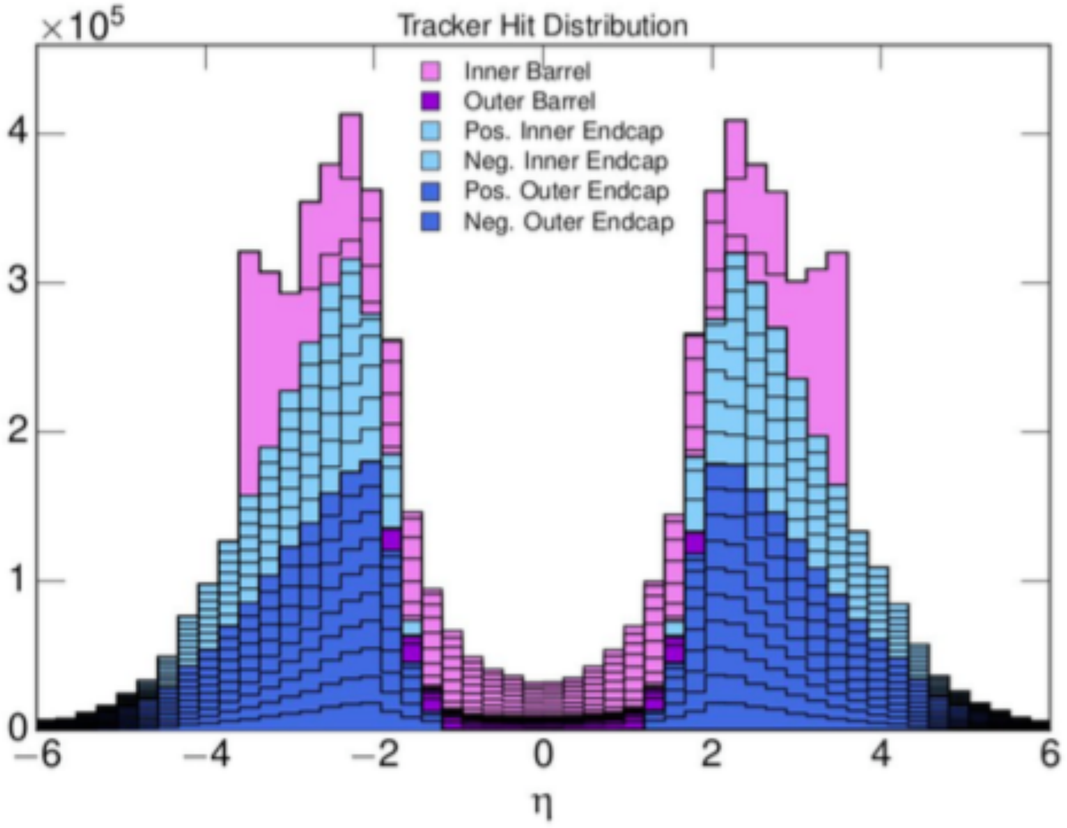

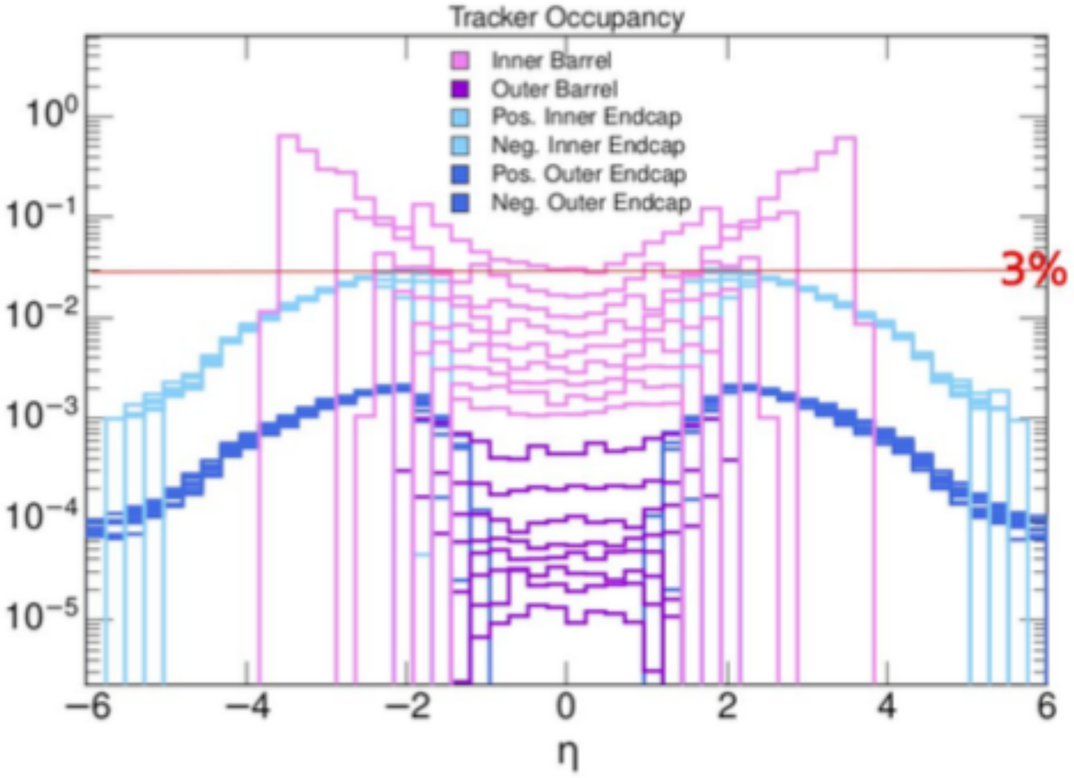

AIDA etc.

- *MarlinTrk* is a generic tracking toolkit
  - based on the **LCIO-EDM** and **DD4Hep** geometry description
  - used by ILC, CLICdp and CEPC
- *IMarlinTrk* interface separates pattern recognition from concrete track fitters
- *DDRec* provides surfaces with material properties -automatically extracted from the detailed DD4hep detector geometry

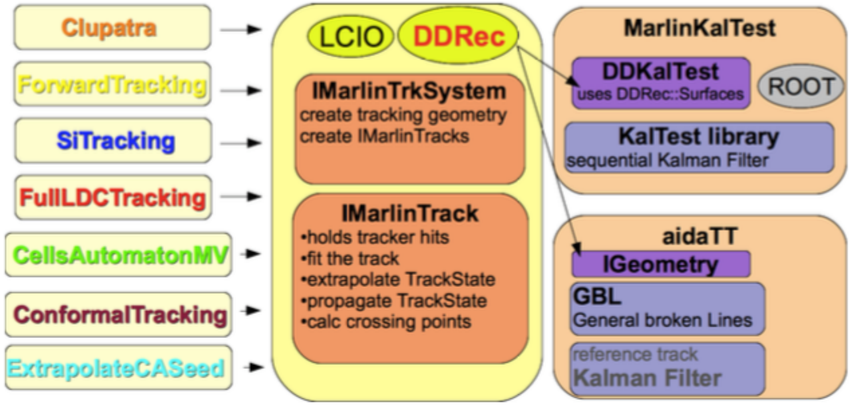

- plan to evolve into a truly generic HEP tracking toolkit
- need to investigate synergies with *ACTS* ?

Lots of chat about generic tracking frameworks. It is not obvious to me that this is really what will help us with our real data problems, although I can see the conceptual interest from ILC/CLIC/FCC etc. people.

On the other hand, implementing our baseline tracking in this, combined with DELPHES, could make it a lot easier to perform standalone proof-of-concept studies and collaborate with non-LHCb people, which would be great.

Software trigger & reco WG chapter

Software trigger and event reconstruction

https://docs.google.com/document/d/1QRO8RA488fwfSg5CSjmvm16-pZpGApSA01666g_mS_0/edit

Software trigger & reco WG chapter

Challenges posed by Future Facilities

Goals:

1. Describe expected evolution in accelerator operating conditions. For CMS/Atlas, its the move towards higher pileup (200, perhaps limited by experiment and/or software algorithm performance (physics and technical), similarly for LHCb Phase II upgrade
2. Describe the overabundance of signal problem faced by experiments, e.g. with reference to Kaon/Charm physics at LHCb or B-physics at CMS/ATLAS

Landscape of Experimental apparatus

Goals:

1. Describe expected experimental hardware changes on 5-15 year timescale. Focus on new experimental features that trigger the need for R&D. Eg, move towards silicon trackers, CMS high-granularity calorimeter, tracking triggers, high-precision timing detectors, desire to fully exploit the LHC's cross-section for light-quark physics which necessitates a reconstruction of ultra-low-PT objects

https://docs.google.com/document/d/1QRO8RA488fwfSg5CSjmvm16-pZpGApSA01666g_mS_0/edit

Software trigger & reco WG chapter

Challenges from Event rates and real-time processing

Goals:

1. Describe evolution in trigger systems expected on 5-15 year timescale. Examples are
   - Increase in event rates. Eg, CMS/Atlas up to 1MHz L1 rate (HLT input rate), and 7.5 kHz output rate (event reconstruction input rate)
   - Move towards quasi-realtime processing of ~all events (examples at LHCb, CMS, etc)
2. Describe implications for trigger, reconstruction, monitoring, calibration and alignment

Challenges from Evolutions in Computing technology

Goals:

1. Summarize challenges faced by current algorithms in making efficient use of today's and future computing resources. Eg, vectorization, manycore, memory locality, I/O...
2. Describe the challenges posed by data structure optimization for both processing and readback (eg, analysis). Include in this the increasing emphasis on online data

https://docs.google.com/document/d/1QRO8RA488fwfSg5CSjmvm16-pZpGApSA01666g_mS_0/edit

Sustainability of software and jobs

Software as infrastructure

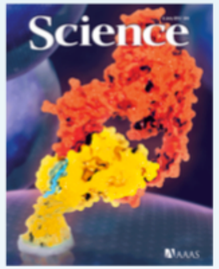

- Software (including services) essential for the bulk of science
  - About half the papers in recent issues of Science were software-intensive projects
  - Research becoming dependent upon advances in software
  - Wide range of software types: system, applications, modeling, gateways, analysis, algorithms, middleware, libraries

- Software is not a one-time effort, it must be sustained
  - Development, production, and **maintenance** are people intensive
  - Software life-times are long vs hardware
  - Software has under-appreciated value

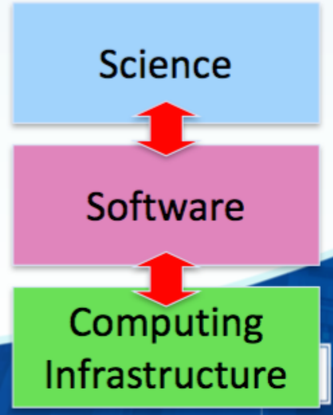

Can't do this justice, see whole slides at the workshop page

Sustainability of software and jobs

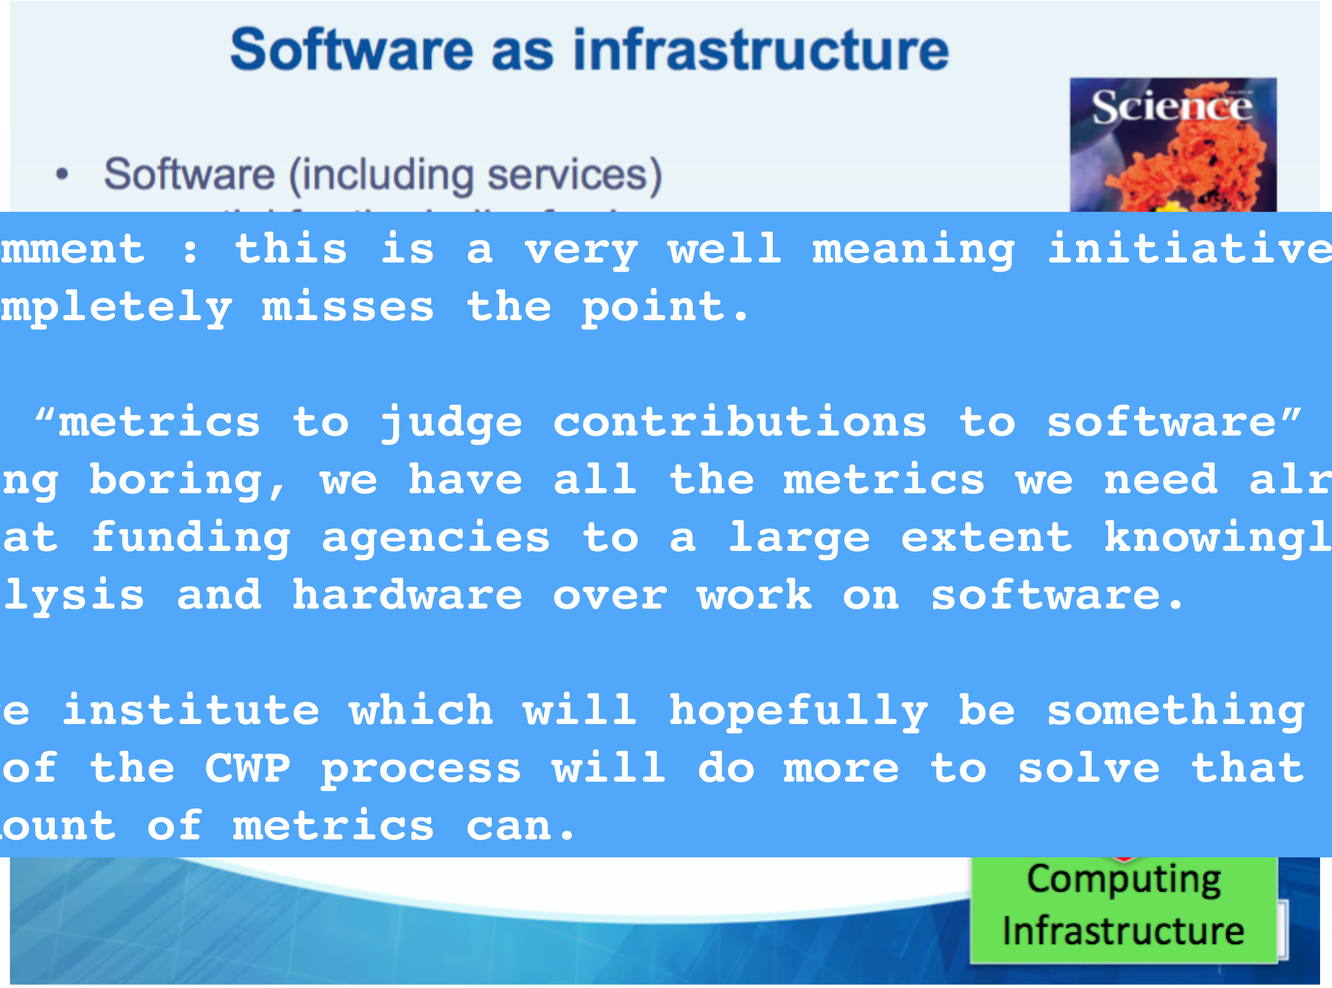

Personal comment : this is a very well meaning initiative but I think it completely misses the point.

Focus is on "metrics to judge contributions to software" but at the risk of being boring, we have all the metrics we need already. The issue is that funding agencies to a large extent knowingly favour work on analysis and hardware over work on software.

The software institute which will hopefully be something concrete coming out of the CWP process will do more to solve that problem than any amount of metrics can.

Can't do this justice, see whole slides at the workshop page

Please contribute to CWP and attend

https://docs.google.com/document/d/1QRO8RA488fwfSg5CSjmvm16-pZpGApSA01666q_mS_0/edit
https://indico.cern.ch/event/613093/