




Report on the ECFA workshop

Simone Campana



The HL-LHC ECFA workshop

- Sponsored by the European Committee on Future Accelerators
- Focus on physics, accelerator and beam parameters, detector and trigger upgrades
- One session on computing with 3 talks

R&D session: Software and Computing

Conveners: Monica Pepe-Altarelli (CERN) , Pippa Wells (CERN)

16:30

Requirements and possible architecture

Speaker: Simone Campana (CERN)

 ECFA2016.pdf

16:50

Computing progress and technology options


Speaker: Helge Meinhard (CERN)

 2016-10-04-ECFAW...

17:15



Developing the roadmap for HL-LHC software

Speaker: Pere Mato Vila (CERN)

 LHC-Software-Road...


This presentation

- I deliberately decided to include all the slides presented at the workshop in the original format
 - In presenting I highlight only the key concepts but you will be able to go through the full content
- This makes it a 66 slides presentation
 - My predecessor in ATLAS computing coordination (Richard Mount) would elegantly go through it in 10 minutes
- I'll try to go through it in 30 minutes
 - And forgive me the lack of elegance

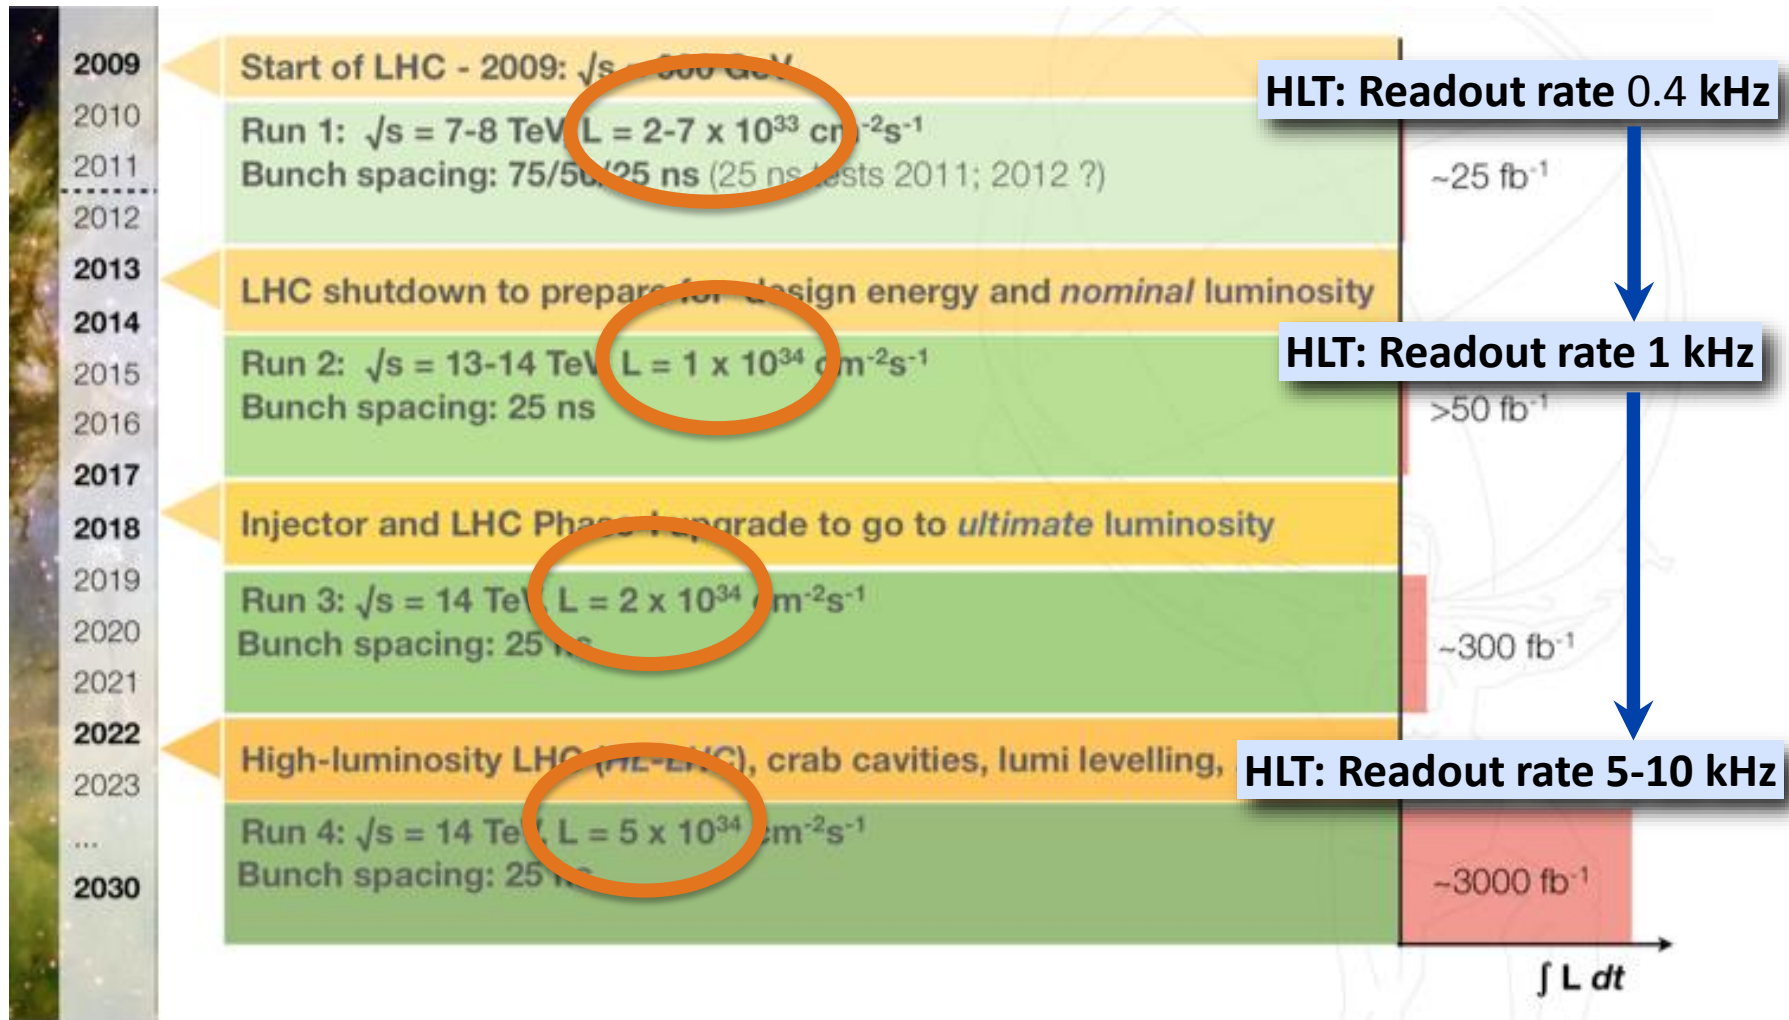


R&D session: Software and Computing Requirements and possible architecture

Simone Campana – CERN



The data rate, volume and complexity challenge



Effect of pile-up increase


The average pile-up:

- $\langle\mu\rangle=14$ in 2015
- $\langle\mu\rangle=23$ in 2016
- $\langle\mu\rangle \approx 35$ in 2017
- ...
- $\langle\mu\rangle$ up to 200 in HL-LHC (10 years)


Higher pileup means:

- Linear increase of digitization time
- Exponential increase of Reco time
- Larger events
- Lots of more memory

Reconstruction of $t\bar{t}$ at $\sqrt{s} = 13$ TeV with CMS 2016 configuration





Reconstruction of $t\bar{t}$ at $\sqrt{s} = 13$ TeV with CMS 2016 configuration



The exponential increase in reconstruction time saturates beyond Run-3 conditions ($\mu=80$)

Indicate a loss of tracking efficiency of the current detector layouts at HL-LHC

Estimates of resource needs for HL-LHC



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

CPU
 x60 from 2016

Technology at ~20%/year will bring **x6-10** in 10-11 years

=> x10 above what is realistic to expect from technology with constant cost

In this presentation...

- The resources needed for HL-LHC will be driven by ATLAS and CMS
- Alice and LHCb will face a challenge in LHC Run-3 and already evolved their computing model
- ... I will focus on ATLAS and CMS computing at HL-LHC
- I am more familiar with the ATLAS computing model and the tools to project it to the future.
- Many plots will be based on those tools and the ATLAS computing model, but the conclusions apply to both ATLAS and CMS

Input parameters, assumptions, disclaimers

Simple model based on today's computing models, but with expected HL-LHC operating parameters

ATLAS Input Parameters at HL-LHC (LOI = the ATLAS Letter of Intent for Upgrade Phase-2)

Output HLT rate: 10kHz (5 to 10 kHz in LOI)
Reco and Simul Time/Evt: from LOI
Nr Events MC / Nr Events Data = 2
Fast Simulation: 50% of MC events
LHC live seconds /year: 5.5M

CMS Input Parameters at HL-LHC



Output HLT rate 7.5 kHz
LHC live seconds /year: 6.0M
Dataset overlap factor: 1.2
Reco and Simul Time at $\mu=200$
Nr Events MC / Nr Events Data = 1.3
Analysis estimated as +60% of all other CPU usage

Simplified Computing Model with respect to 2016/2017 resource requests:



Legacy from previous years not taken into account
=> Little difference at the beginning of the Run-4 but huge
difference for Run-2 and Run-3

HL-LHC baseline resource needs

ATLAS



ATLAS





events: HLT output rate and MC needs

The output trigger rate does not determine only the amount of data per year but also the amount of Monte Carlo to be produced.

We foresees a value between 5 kHz and 10kHz.
ATLAS baseline is 10kHz, CMS is 7.5kHz

The physics case for HL-LHC will evolve in the next years.

One might assume a lower need of MC with respect to data, but generators might become more expensive seeking precision

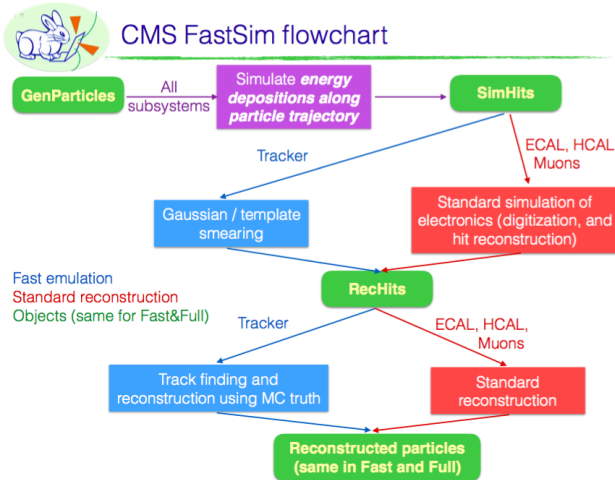



Fast Simulation and Fast Chain


G4 Fast Simulation will moderately help in HL-LHC. CPU is driven by reconstruction

Both ATLAS and CMS invested in a Fast Chain. x10 (++) faster than standard simulation

ATLAS



ATLAS Fast Chain

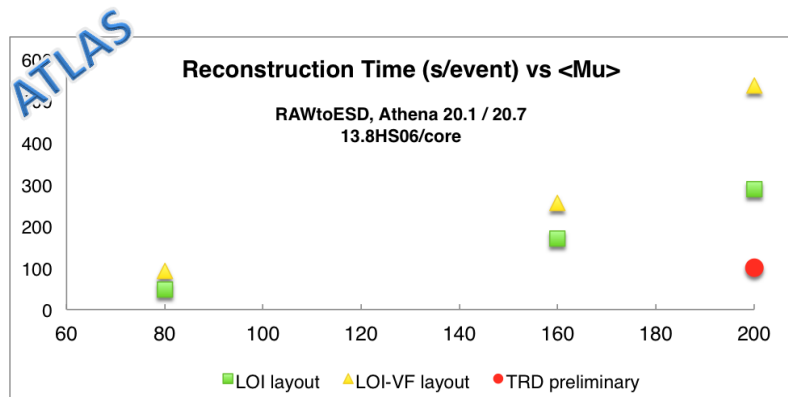


Layouts and Reconstruction


Reconstruction time dominates the CPU consumption in HL-LHC

The detector layout will play an important role, together with the optimization/tuning of algorithms. Tracking will be the main consumer

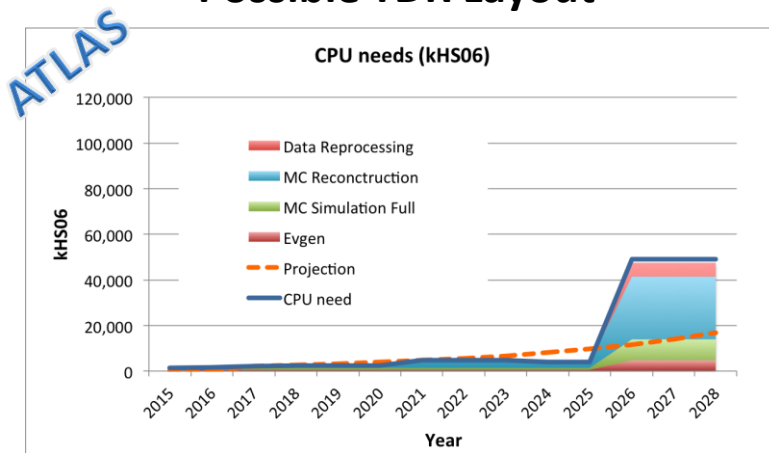
It is important to consider computing performance in designing the HL-LHC detectors. Good that this is happening



LOI Layout



Possible TDR Layout




Preliminary conclusion

- The CPU needs for HL-LHC could exceed x10 the projection of today's resources in 2026 in a pessimistic scenario
- In reality, large gains are foreseeable and we are on the right path
- Hardware trends will play a crucial role and our software will need to adapt to them
- So please listen carefully to the next two presentations

What about Storage ?

No AOD on disk



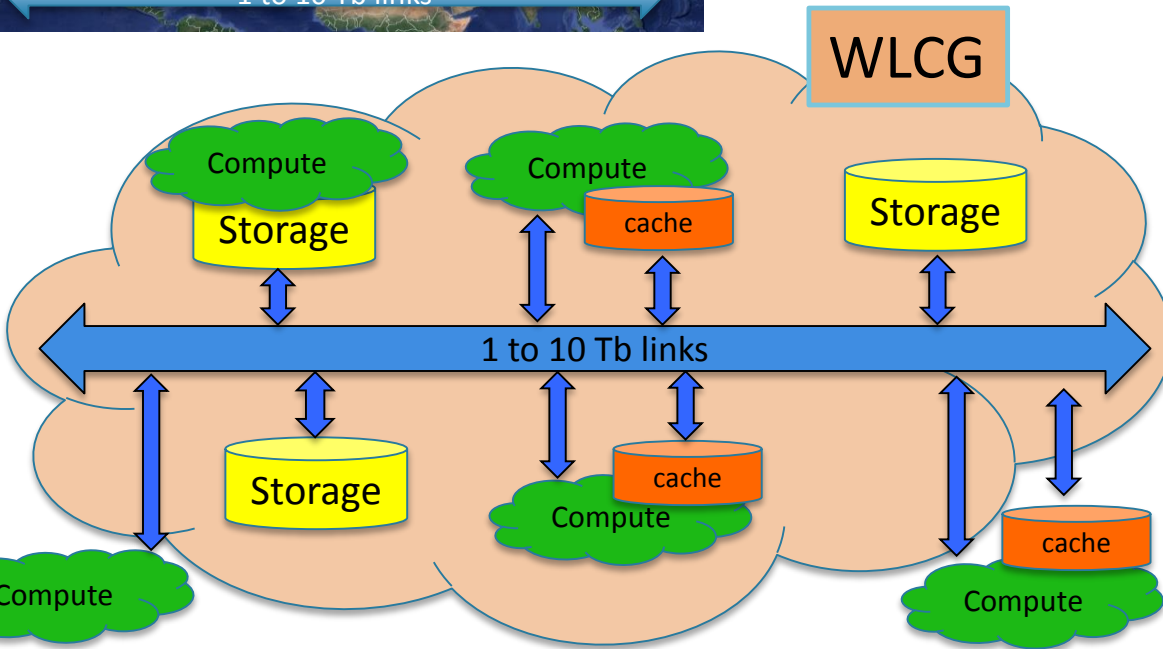

Storage is really the hard part. Even in an optimistic scenario, we are still far from solving the problem

AODs and their derived formats are the main consumers.

With no AOD on disk you get x4 above the resource projection (left plot)

The remaining gain must come from re-thinking of distributed data management, distributed storage and data access. A network driven data model allows to reduce the amount of storage, particularly for disk. Tape today costs at least 4 times less than disk.

Computing infrastructure in HL-LHC



A data cloud for science

Storage and Compute loosely coupled but connected through a fast network

Heterogeneous Computing facilities (Grid/Cloud/HPC/ ...) both in and outside the cloud

Different centers with different capabilities, for different use cases

Data Management: Challenges and Opportunities

- “Funny how tape never seems like the cheap option when you have to pay for it”. One could say the same about network
- A fast WAN does not imply fast data access. The infrastructure and the I/O layers need to be optimized from end to end
- Multilevel caching should be built **IN** the infrastructure rather than **ON** top of it
- A unique opportunity to define and implement a common data management and data access layer
- Today WLCG is a data Grid. Tomorrow we will have a data cloud
The challenge is always the data

Conclusions

- We identified a concrete set of steps in preparation for computing at HL-LHC
- To keep cost of computing under control in 2026 we need to invest effort from now. Data will be the challenge.
- The effort spans many areas: online, offline software, distributed computing, physics, infrastructure and facilities. The detector layout will play a crucial role
- It is important to consider cost of computing when choices are made
- We are on schedule to define a computing model for HL-LHC in the next three years



Computing Progress and Technology Options

ECFA Workshop on HL-LHC, Aix-les-Bains, 03 – 06 October 2016

Helge Meinhard / CERN

Presenting material prepared by Bernd Panzer-Steindel / CERN

Outline


- Semiconductor market
- Device market
- Processors
- Hard Disk
- Solid-State Disks
- Memory
- Tapes
- Server

- Summary
- References

General Market

Few companies dominating the markets

Worldwide Semiconductor Revenues
Year-to-Year Percent Change


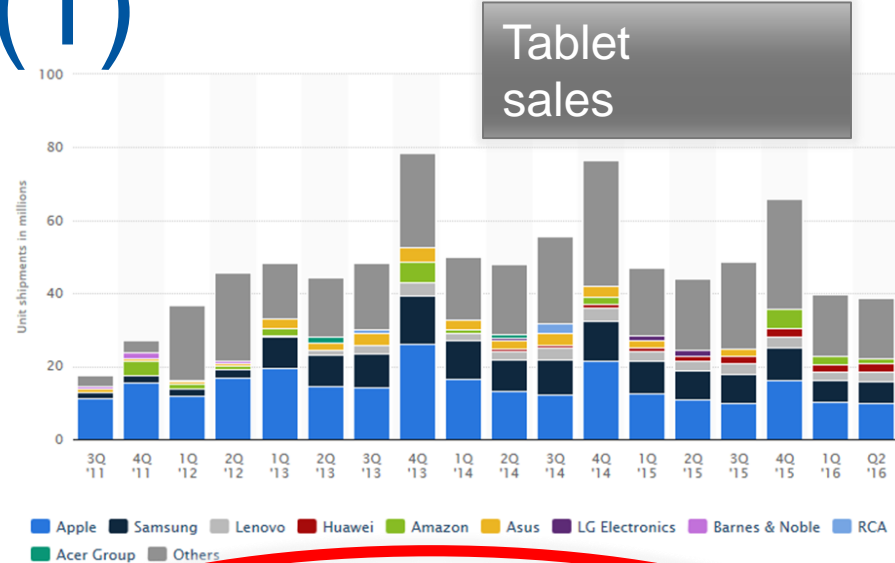
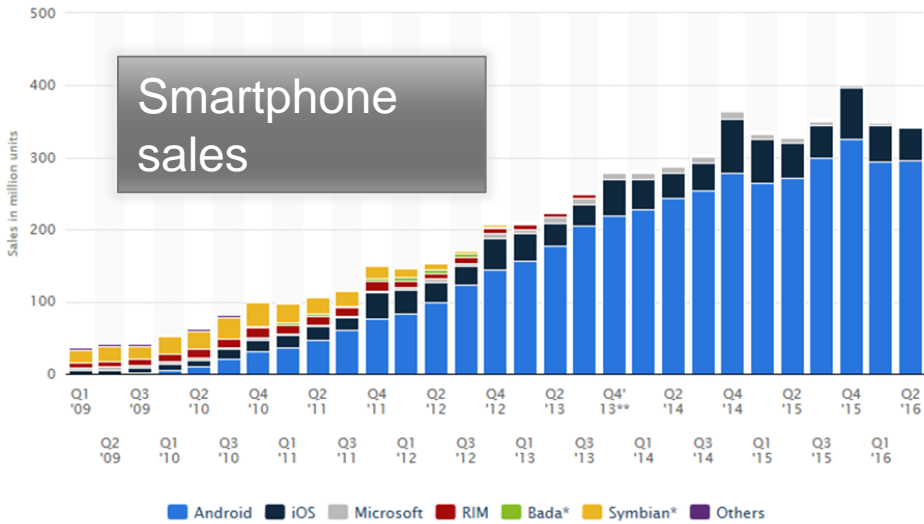


Source: WSTS

Small or negative growth rates

Server CPUs	Intel (99%)
FPGA	Xilinx (49%), Intel (38%)
GPU	Intel (72%), Nvidia (14%), AMD (14%)
Hard disks	Western Digital (44%), Seagate (40%), Toshiba
Tape drives	HP, IBM, Oracle
Tape media	Fujifilm, Sony
NAND	Samsung (45%), Toshiba, Western Digital, Intel
DRAM	Samsung (47%), Hynix, Micron/Intel

Device Markets (1)




**Market saturation:
no or negative growth rates**

Smartphones	0%
Tablets	-12%
Desktops and laptops	-7%
Servers	-3%

Device Markets (2)

New mobile subscriptions Q1 2016






63 million new mobile subscriptions globally in Q1 2016

Top 5 countries by net additions Q1 2016

- 1 India +21 million
- 2 Myanmar +5 million
- 3 Indonesia +5 million
- 4 USA +3 million
- 5 Pakistan +3 million


5 BILLION subscribers

The number of mobile subscriptions exceeds the population in many countries. This is largely due to inactive subscriptions, multiple device ownership or optimization of subscriptions for different types of calls. This means the number of subscribers is lower than the number of subscriptions. Today there are around 5 billion subscribers compared to 7.4 billion subscriptions.




Saturation:
7.3 B phone subscriptions world-wide – more than the population

Replacement bump expected in 2018

Processors (1)

Estimated Cost of Developing Lower Node Chips



Market Realist[®]

Source: Gartner


Source: IBS, Inc. (Los Gatos, CA)

Year	2001	2003	2005	2007	2009	2012	2015
SMIC							
Hitachi							
NEC							
Sony							
NXP							
Infineon							
Renesas							
Freescall							
TI							
Fujitsu							
Panasonic							
Toshiba							
UMC							
IBM							
STM							
G'Foundries							
TSMC							
Samsung							
Intel							

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

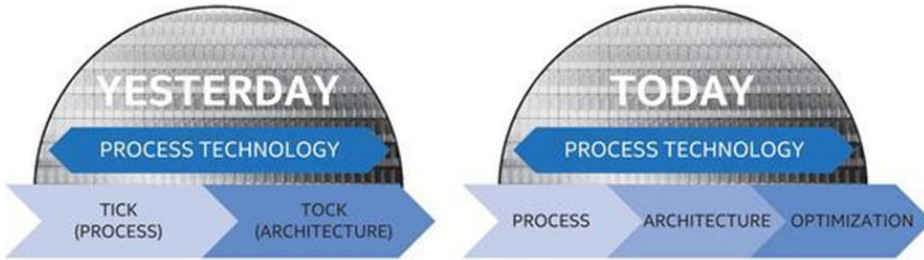
Non-linear costs for development

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




Latest ITRS report changes the structure size predictions; becoming flat due to financial reasons

Processors (2)



Intel moved from 2-year cycle to 3 years or more




Feature	Intel	TSMC	Samsung
Gate length (nm)	24	33	30
Min contacted gate pitch (nm)	70	90	78
Fin height under gate (nm)	42	37	37
Fin pitch (nm)	43	45	49
Min metal pitch (nm)	52	70	67

• Intel transistors are smaller than TSMC or Samsung


Incubation Time

- Strained Silicon
 - 1992->2003
- HKMG
 - 1996->2007
- Raised S/D
 - 1993->2009
- MultiGates
 - 1997->2011



~ 12-15 years

Decrease of feature size goes along with new material technologies
R&D → production needs 12-15 years



7nm structures need new technologies: nanowires and non-silicon material

Accelerators: GPU (1)

Embedded market shares
(CPU+GPU): Intel 72%,
Nvidia 16%, AMD 12%

Discrete GPU cards: Nvidia 77%,
AMD 23%

Desktop and notebook shipments declining




Figure 2: While PC shipments have returned to predictable patterns, graphics shipments have been erratic and defy any seasonal attributes

Focus: high-end Gamer (DP artificially reduced)

Professional workstation cards and HPC:
small niche, ~5 million cards per year
(compared to 350 million total GPUs)


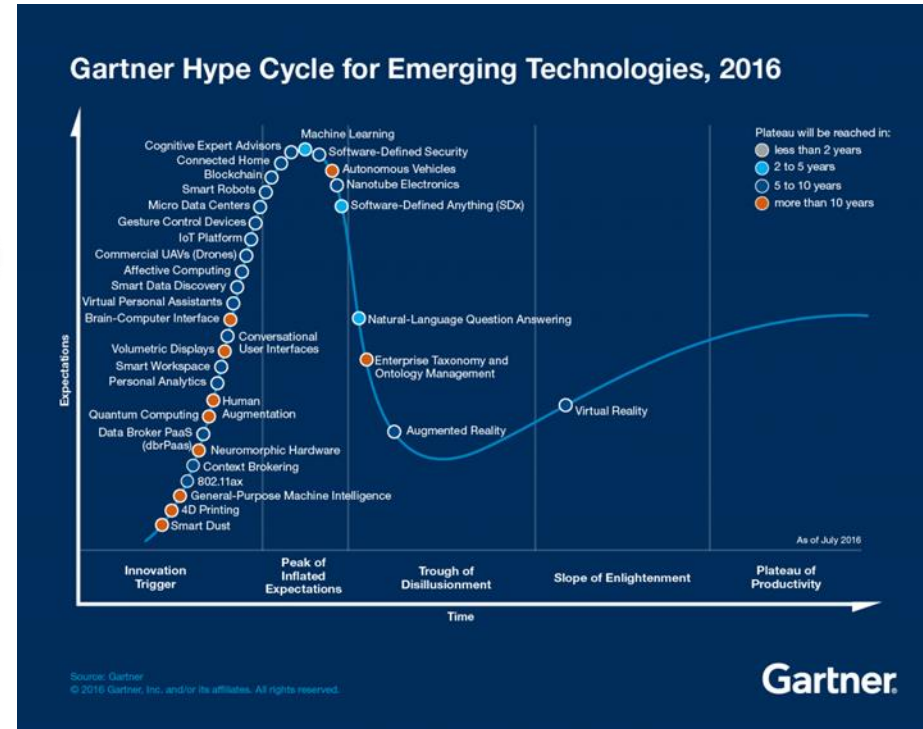


Figure 3: Add-in board shipments over time


Accelerators: GPU (2)

- New focus for graphic cards : machine learning
- Move to FP16 and INT8 architectures, less precision
→ 8 bit processing !
- Google TPU Tensor Processing Unit
- New start-ups with special processor designs: KNUEDGE, Nervana (just bought by Intel), krtkl, Eyeriss
- Essentially not usable as general purpose processors
- Intel changing strategies also for their Knights xxx processors

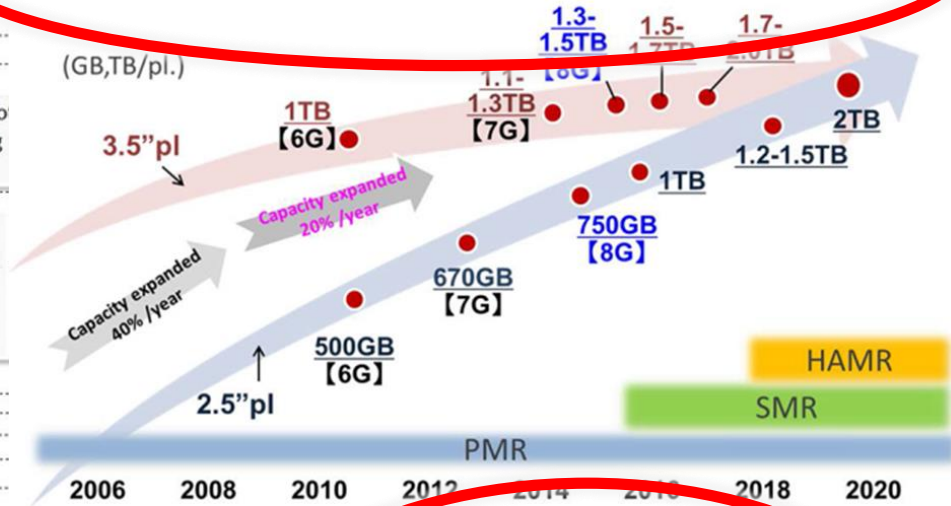


Hard Disks (1)

ASTC Technology Roadmap




[Road map for storage density increase] (SDK forecast)



Combining bit density (30% annual growth rate) and volume density (number of platters, helium) → 100 TB in 2025 conceivable

PMR limit at 1 TbPSI
SMR adds ~25%, market small
HAMR should provide 5 TbPSI

HAMR delayed, production in 2018



Hard Disks (2)


Continuous decrease in revenues
Forecast changes every year

Gartner's Total HDD Revenue Estimates vs. Stifel Estimates




Areal Density Trends

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



©2016 Information Storage Industry Consortium - All Rights Reserved


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



Hard Disks (3)



Shipments of HDDs by Seagate, Western Digital and Toshiba

All numbers are in millions of units, data by Seagate and WD. Shipments of Toshiba are estimates.



- HDD sales decreasing, related to PC sales decline
- Pressure from SSDs in the notebook area and in the enterprise performance drives (FC, 15krpm)
- Stable sales for capacity cloud drives

Solid-State Disks vs. Hard Disks




- 14 times more HDD capacity than SSD
- Price per TB decreasing about the same way
- Difference SSD/HDD costs per TB ~5-10 will slowly decrease
- Fab investment of \$100...200 B necessary to achieve HDD EB deliveries

Memory: DRAM


Memory technology trend

- GDDR6 with over 14Gbps, beyond 10Gbps GDDR5
- LP5, 20% more power-efficient than LP4X



Limited future improvements on performance and energy efficiency


Figure 1: DRAM Spot Price Trend



DRAM Technology Review

TECHINSIGHTS

DRAM Process Node Roadmap (Manufacturers)



Source: DRAMeXchange

2 Chinese companies will enter the DRAM market in 2017

Further price decay possible

New Memory Technologies

- 3d xpoint: new technology from Intel and Micron, presumably a variant of Phase Change Memory
- Specs are changing:
Announcement 2015: 1000x faster, 1000x endurance, 10x denser than NAND
IDF 2016: 10x faster, 3x endurance, 4x denser than NAND
- Will enter the high end server market in Q1 2017
- Memristors: developed since 2008; HPE now collaborating with SanDisk (ReRAM)
- Spin torque MRAM in larger production units available (Everquest + Globalfoundries)
- Low density and high price
- Tantalum memory, Rice University
- RRAM or ReRAM , various new categories being developed: Oxide RAM (OxRAM), Conductive-Bridge RAM (CBRAM) or Self-Rectifying Cells (SRC)
- But... NAND fab investments are high, extended technology lifetime with 3D, hard to replace in the short term


Magnetic Tapes (1)

LTO Roadmap



	LTO-3	LTO-4	LTO-5	LTO-6	LTO-7	LTO-8	LTO-9	LTO-10
Shipment Year	2005	2007	2010	2013	2015	TBD	TBD	TBD
Native Capacity	400GB	800GB	1.5TB	2.5TB	6.0TB	Up to 12.8TB	Up to 25TB	Up to 50TB
Compressed Capacity	800GB	1.6TB	3.0TB	6.25TB	15TB	Up to 32TB	Up to 62.5TB	Up to 125TB
Native Transfer Rate	80 MB/s	120 MB/s	140 MB/s	160 MB/s	300 MB/s	Up to 472 MB/s	Up to 708 MB/s	Up to 1100 MB/s
Compressed Transfer Rate	160 MB/s	240 MB/s	280 MB/s	400 MB/s	750 MB/s	Up to 1180 MB/s	Up to 1770 MB/s	Up to 2750 MB/s


■ TAPE: source NSIC 2013





- Enterprise drives:
Oracle 2017: 8.5 TB → 12 TB
IBM 2018: 10 TB → 16 TB
- Technology in the lab:
Fujifilm 154 TB, Sony 185 TB,
IBM 220 TB

- Good improvements of price/capacity


Euro/GB Tape price evolution (street prices)



Magnetic Tapes (2)





Unit Shipments: Calendar Year



- More NAND than LTO shipped
- Steady decrease of tapes shipped and revenues
- Will Oracle and/or IBM sell or drop these products?


Servers (1)

- Server market is saturated: flat revenues and unit shipments
- High profit market
- Single vendor: Intel, 99% market share
- Several initiatives to change that:
 - OpenPower (IBM): consortium with many members
 - But revenues still going down, little impact so far
 - Announcement of POWER9 might help
 - ARM server:
 - AppliedMicro, Qualcomm, Cavium: new high end products
 - Announcements for 2H2017 (third ARMv8 Wave 2017-2018),
 - First two waves had little impact
 - Fujitsu ARM-powered supercomputer
 - Add large vector instructions to the ARM design
 - Aimed for 2020, now ~2022



Servers (2)


Preliminary extrapolation of CPU and disk server costs (based on CERN procurements)



Pessimistic and reasonable improvement extrapolations

Influence of changing software and hardware architecture requirements to be taken into account (programs, data model, data centre, ...)

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



Summary (1)

- Device markets (smartphones, tablets, PCs, notebooks, servers, HPC) saturated – negative growth
 - Replacement market
- Moore's Law in trouble, financial issues
 - Not clear how this effects price/performance evolution
 - So far okay for CPU and disk servers
- Technology improvements still continuing, but require high CAPEX needed
 - End-product price tag evolution more complicated
- Market dominance of few companies increases, competition diminishing

Summary (2)

- Technology unlikely to solve the HL-LHC computing problem
 - Not much more to be expected than minor contributions

References

<http://electroiq.com/blog/2016/05/global-semiconductor-sales-increase-slightly-in-march/>
<http://www.statista.com/statistics/266219/global-smartphone-sales-since-1st-quarter-2009-by-operating-system/>
<http://www.statista.com/statistics/263393/global-pc-shipments-since-1st-quarter-2009-by-vendor/>
<http://www.statista.com/statistics/276651/global-media-tablet-shipments-since-3rd-quarter-2011-by-vendor/>
<https://www.ericsson.com/res/docs/2016/ericsson-mobility-report-2016.pdf>
<http://www.nasdaq.com/article/the-evolution-of-smartphone-markets-where-growth-is-going-cm619105>
http://www.potomac institute.org/steps/images/PDF/Articles/FritzeSTEPS_2016Issue3.pdf
<http://www.pcper.com/news/Processors/Intel-officially-ends-era-tick-tock-processor-production>
<http://semimd.com/chipworks/>
<http://www.extremetech.com/extreme/223022-the-myths-of-moores-law>
<http://www.forbes.com/sites/gartnergroup/2016/08/29/track-three-trends-in-the-2016-gartner-hype-cycle-for-emerging-technologies/#59fc4d787286>
http://jonpeddie.com/publications/market_watch
<http://www.anandtech.com/show/10613/discrete-desktop-gpu-market-trends-q2-2016-amd-grabs-market-share-but-nvidia-remains-on-top>
<http://www.computerworld.com/article/3041947/data-storage/how-these-technologies-will-blow-the-lid-off-data-storage.html>
<http://www.computerworld.com/article/2852233/want-a-100tb-disk-drive-youll-have-to-wait-til-2025.html>
<http://www.anandtech.com/show/9866/hard-disk-drives-with-hamr-technology-set-to-arrive-in-2018>
http://www.theregister.co.uk/2016/05/31/hdd_revenues_to_plummet_as_ssd_penetration_rises/
<http://www.anandtech.com/show/10315/market-views-hdd-shipments-down-q1-2016>
http://www.trendfocus.com/ssd-cq116_update/
<http://www.forbes.com/sites/tomcoughlin/2016/02/03/flash-memory-areal-densities-exceed-those-of-hard-drives/#5dc24d2b4026>
<http://www.anandtech.com/show/10589/hot-chips-2016-memory-vendors-discuss-ideas-for-future-memory-tech-ddr5-cheap-hbm-more>
<http://asia.nikkei.com/Business/Trends/NAND-flash-memory-prices-likely-to-climb-again>
<http://www.techinsights.com/techinsights/about-techinsights/articles/deep-dive-into-the-intel-micron-3D-32L-FG-NAND/>
<http://amigobulls.com/articles/micron-technology-inc-stock-is-the-next-big-idea-for-2016>
<http://wccftech.com/micron-compete-samsung-16-nm-dram/>
<http://storageconference.us/2016/Slides/BobFontana.pdf>
<http://searchsolidstatestorage.techtarget.com/feature/New-memory-technologies-generate-attention-as-successors-to-NAND-flash>
<http://www.itjungle.com/tfh/tfh061316-story05.html>
<http://www.forbes.com/sites/tomcoughlin/2016/01/15/digital-storage-projections-for-2016-part-2/2/#35b1916a3aa8>
http://www.lto.org/wp-content/uploads/2016/03/LTO_Media-Shipment-Report_3.22.16.pdf
<http://semiaccurate.com/2016/09/12/intels-xpoint-pretty-much-broken/>



Developing the Roadmap for HL-LHC Software

ECFA High Luminosity LHC Experiments Workshop - 2016
October 3-6 2016, Aix-Les-Bains

Pere Mato/CERN

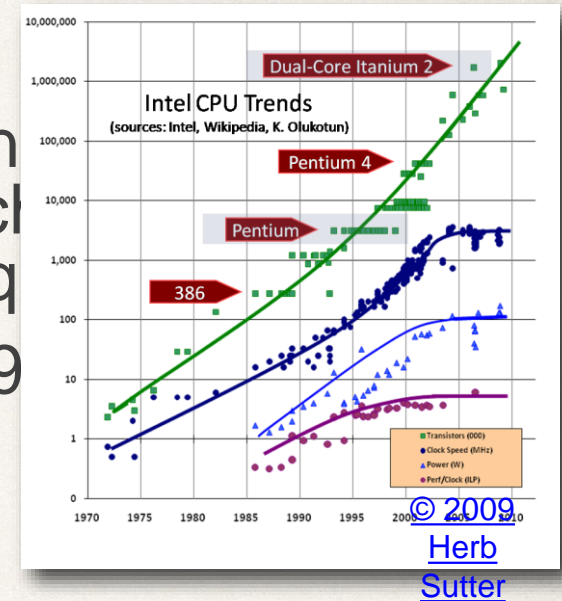
October 3-6 2016

Coping with HL-LHC Needs

- ❖ To cope with the expected enormous computing demands for HL-LHC we have two solutions:
 - ❖ Invest on more computing: more hardware, more centers, ...
 - ❖ Invest on better software
 - ❖ or a combination of both
- ❖ What is better software?
 - ❖ Better algorithms
 - ❖ Better adapted to the current and future hardware architectures
 - ❖ Better optimisations
 - ❖ Better quality
 - ❖ Better sustainability

CPU Technology Trends

- ❖ Until ~2004 we have had an easy life in HEP software and computing
- ❖ Year after year up to 2x increase in capacity thanks to the #transistor/clock (Moore's law) and higher clock frequency
- ❖ The same program that in year 1999 needing 10 seconds, would need 1 second in 2002
- ❖ The "easy life" is now over
 - ❖ The available transistors are used for adding new CPU cores while keeping the clock frequency basically constant thus limiting the power consumption
 - ❖ We need to **introduce parallelism** into applications to fully exploit the continuing exponential CPU throughput gains



Technical Challenges

- ❖ Big chunks of the LHC software is more than 20 years old, and some parts **require re-engineering and modernization**
 - ❖ Need to exploit modern hardware (many-core, GPU, etc.) to boost performance
 - ❖ Modernize implementations (C++11/14 constructs, use more modern and performant libraries, etc.)
- ❖ Many **algorithms will need to be re-designed** to be run in parallel but integrating them to run in a single application is highly non-trivial
 - ❖ It will require new levels of expertise that need to be acquired by the community
- ❖ Changes in the code of running experiments must be done gradually whilst **preserving the correctness of the physics output**

Paradigm Shift

- ❖ Most of the scientific software and algorithms was designed for sequential processor in use for many decades and **will require significant re-engineering**
- ❖ Migrating sequential applications to multi-threaded is highly non-trivial
 - ❖ Difficult to **develop**: we not only need to code what needs to be done but also how this is done in parallel
 - ❖ Difficult to **debug**: nasty data race conditions will be difficult to reproduce, and so to fix
 - ❖ Difficult to **maintain**: latent threading bugs may take years to be visible
- ❖ The community needs to **develop expertise in concurrent programming**
 - ❖ Similarly to the OOP migration, training will be eagerly needed

Collaboration Challenges

- ❖ LHC experiments **cannot afford to undertake this software [r]evolution independently**
 - ❖ There are many common software packages that would require common efforts, thus coordination
 - ❖ General wish to increase the level of commonality and re-use
- ❖ Require the collaboration of the whole HEP community to ensure evolution and sustainability
 - ❖ Show a common and coherent roadmap to funding agencies
 - ❖ Establish structures to facilitate contributions to the HEP software stack
- ❖ The adoption of a **collective response** will help to meet the challenges using available expertise and resources and within the required timescale

Prospects for HEP Software


- ❖ Potential gains can be made by exploiting features of today's CPUs' micro architecture
 - ❖ by making use of vector registers, instruction pipelining, multiple instructions per cycle
 - ❖ by improving data and code locality and making use of hardware threading
- ❖ New architectures to **off-load large computations to accelerators** (GPGPUs, Xeon Phi™) or the new integrated architectures with heterogenous processors (AMD)
 - ❖ specific memory models will force explicit memory programming
 - ❖ new programming languages (Cuda, OpenCL, etc.)

Prospects for HEP Software (2)

- ❖ Today multi-core architectures employing $O(10)$ cores are well exploited using a multi-process model (1 job/core)
- ❖ However this performance will **not scale** to future generations of many-core architectures employing $O(100)$ cores **due to memory issues**
 - ❖ there are technical issues related to connecting many cores to shared memory that will reduce the amount of memory available to each core
 - ❖ whereas the memory footprint of HEP code is increasing due to increasing event complexity as the energy and luminosity of the LHC is increased
 - ❖ in addition, we may see new architectures with non-uniform memory access

Addressing the Challenges

- ❖ **HEP Software Foundation (HSF)** as the umbrella for addressing these challenges together!
 - ❖ Collection of ideas and proposals in 2014 and **startup-team** formed
 - ❖ Kick-off workshop Jan 2015 at SLAC established concrete activities
 - ❖ Workshop in May 2016 at LAL to review progress and setting directions
- ❖ In addition, the HSF aims at
 - ❖ Support **career development** for software and computing specialists
 - ❖ Provide a framework for **attracting effort and support** to S&C common projects
 - ❖ Provide a **structure to set priorities** and goals for the work
 - ❖ Facilitate wider connections; it should be open enough to form the basis for **collaboration with other sciences**



Current Status and Activities

❖ Sharing expertise

- ❖ Schools, trainings and courses (not always easy to find)
- ❖ Adopting [wikiToLearn](#) as a platform for training material
- ❖ [HEP S&C Knowledge Base](#)
 - ❖ Database of software packages, categories, experiments, organisations, languages, meetings, workshops, etc.
- ❖ HSF Technical notes
- ❖ Pursuing a journal on "SW&C for Big Science"
- ❖ Topical fora and working groups in the HSF web

❖ New hardware architectures and technologies

- ❖ [Concurrency forum](#), evolved into a general software technology forum
- ❖ Usage of resources provided on best-effort basis by e.g. CERN's TechLab / Openlab
- ❖ Porting to new architectures efforts within the LHC experiments

Current Status and Activities II

- ❖ Software performance
 - ❖ Simulation: parallelisation of Geant4; GeantV R&D activity
 - ❖ HSF is organising "software community meeting" to review the progress made in simulation R&D
 - ❖ Reconstruction: HSF common tracking SW forum + IML forum
 - ❖ I/O: parallel ROOT I/O, key-value-store evaluations
 - ❖ Mathematics: MetaLibm, parallelisation of fitting, etc.
 - ❖ Ad-hoc improvements and parallelisation in various SW projects
 - ❖ Performance tools (e.g. igprof, FOM tools)

Current Status and Activities III

- ❖ Supporting developers and participating projects
 - ❖ Providing best practices to facilitate integration into HEP eco-system
 - ❖ Project templates for bootstrapping new projects
 - ❖ Development services
 - ❖ Help in selecting the proper SW license
- ❖ Quite some activity in HSF, even though participation in the startup-team is on volunteer/best-effort level
- ❖ Need to allocate soon some dedicated resources to keep momentum and ensure continuity

HSF Working Groups

Working Group	Objectives	Forum - Mailing list
Communication and information exchange	Address communication issues and building the knowledge base Technical notes	hep-sf-tech-forum
Training	Organization of training and education, learning from similar initiatives	hep-sf-training-wg
Software Packaging	Package building and deployment, runtime and virtual environments	hep-sf-packaging-wg
Software Licensing	Recommendation for HSF licence(s)	hep-sf-tech-forum
Software Projects	Define incubator and other project membership or association levels. Easy-start project templates	hep-sf-tech-forum
Development tools and services	Access to build, test, integration services and development tools	hep-sf-tech-forum

HSF Topical Fora

- ❖ Software Technology Forum
 - ❖ Technical issues to embrace new technology in our software
 - ❖ Ongoing activity
- ❖ Reconstruction Algorithms Forum
 - ❖ All matters of event reconstruction and pattern recognition software
 - ❖ Several in-person meetings, “Connecting the Dots” workshop
- ❖ Machine Learning Forum
 - ❖ ML discussions and code development in the context of HEP
 - ❖ Development of relevant tools, methodology and applications

Cross-experiments Collaborations

- ❖ Experiment frameworks
 - ❖ Gaudi (ATLAS, LHCb, FCC)
 - ❖ FAIRRoot (FAIR, ALICE)
 - ❖ ART (CMS, Neutrino programme)
- ❖ Common Conditions Data Project
 - ❖ Discussion/cooperation between ATLAS, Belle II, CMS and LHCb
- ❖ Common Software Build and Packaging Tool efforts
 - ❖ Working group of HSF comparing HEP and non-HEP solutions
 - ❖ Starting point was LCG's Librarians and Integrators Meeting
- ❖ Cooperation on Reconstruction Software
 - ❖ "Connecting the Dots" tracking workshop extended by HSF session about common tracking implementations
- ❖ AIDA2020 (EU funded)
 - ❖ DD4hep for detector description (LCD, FCC, potentially LHCb)
 - ❖ PODIO data model library (FCC, LCD, potentially LHCb)
- ❖ DIANA (Data Intensive ANALysis) (NSF funded)
 - ❖ 4-year project on on analysis software, including ROOT and its ecosystem

Examples of cross-experiment collaborations, with involvement or moderation of the HSF - going beyond LHC

Defining Longer-term Strategy

- ❖ HL-LHC computing requires a major ‘software upgrade’
- ❖ A **Community White Paper** (CWP) on the overall strategy and roadmap for software and computing **has been proposed**
 - ❖ The scope should not be restricted only to HL-LHC
 - ❖ However, it can be used to identify research required to prepare the LHC experiment’s TDRs in advance of HL-LHC
 - ❖ Some early software components could be built, tested and used by experiments in LHC Run3
- ❖ Organised by the HEP Software Foundation (HSF)
- ❖ Paper to be delivered by Summer 2017, important ingredient for
 - ❖ Future funding opportunities, including a US opportunity for an NSF-funded Software Institute
 - ❖ WLCG computing roadmap for HL-LHC

CWP

- ❖ The CWP should **identify and prioritise** 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 experiments
 - ❖ to ensure the **long term sustainability** of the software through the lifetime of the HL-LHC
- ❖ We need to **engage the HEP community** in this process through a **series of workshops**
 - ❖ Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider so far)

CWP: Getting Organised

- ❖ First organisational discussion took place two weeks ago
 - ❖ Reviewed draft charge, initial **working group organisation** and asked attendees to encourage people in their communities to join up and participate
 - ❖ Working groups will self-organise, the “**do-ocracy**” determining the proactive people who emerge as conveners
 - ❖ Created a single CWP mailing list
 - ❖ <https://groups.google.com/forum/#!forum/hsf-community-white-paper>
 - ❖ **Please subscribe** if you want to participate or follow progress
 - ❖ A GoogleDoc page will be setup for each WG to start planning and writing

CWP: Getting Organised II

- ❖ The next step: Sun Oct 9 during the pre-CHEP WLCG meeting
 - ❖ The afternoon is an HSF session, to be devoted mainly to CWP getting organised
 - ❖ Flesh out the charges, the initial ideas, plans for the WGs
 - ❖ Ideally with early volunteers in at least some WGs having brought some initial written ideas
 - ❖ Only a subset of the interested community will be present, asked for Vidyo
- ❖ The real launch: a **workshop at UCSD San Diego Jan 23-26**
 - ❖ Start real writing after a few months post-CHEP gestation in the WGs
 - ❖ Discussions on more controversial topics, reach consensus
 - ❖ Detailed plans and responsibilities for delivering white paper by summer 2017

CWP: Working Groups

Working Group	Challenges and Comments
Computing models, facilities, technology evolution	range of possible models, costing
Physics generators	better models, better precision, code optimisations
Detector simulation	full and fast simulations, hi-pileup environments
Triggering	algorithms, GPUs and/or FPGAs
Event reconstruction	new approaches to event reconstruction
Data access and management	scaling to the exabyte level
Workflow and resource management	millions of jobs in heterogenous systems
Data analysis and interpretation	efficient use of many-core, modern techniques
Software development, deployment and validation/verification	improved modularity and quality, easy
Data and software preservation	
Visualization	
Careers, staffing, training	

This list will evolve. Additional working groups could be formed if it makes sense (e.g. on specific technology issues)

re
each
per 65

Summary

- ❖ We need as a community to **invest on better software** to cope with the high demands of the HL-LHC
- ❖ Existing software needs to be re-engineered, and a lot of new software needs to be developed using new ways: **paradigm shift**
 - ❖ The community needs to develop expertise in concurrent programming
- ❖ Initiated the **HEP Software Foundation** (HSF) as the umbrella for addressing these challenges together!
 - ❖ Need to put dedicated resources soon to keep momentum
- ❖ Working on a **Community White Paper** (CWP) to define the strategy and roadmap for long-term software and computing
 - ❖ Input for funding opportunities and WLCG roadmap for HL-LHC
 - ❖ Call for participation to the defined WGs

Executive Summary

- The LHCC and the Funding Agencies asked to start a process to define and address the cost of computing for HL-LHC
- We (WLCG) started this process and what I just summarized are the first steps forward
- New Alice and LHCb computing TDRs target Run-3
- The plan is to produce ATLAS and CMS TDRs on the timescale of 2020
- MHO: innovation toward HL-LHC needs to be an adiabatic process involving SW, computing and infrastructure. And needs to start now