Scientific Software and Computing in the HL-LHC, EIC, and Future Collider Era

D. Piparo (CERN)

ACAT 2022

This Talk

- Introduction to the HL-LHC and EIC eras (next ~2 decades)
- High level overview of the challenges ahead of HE(N)P software and computing, based on current status and directions in hardware, software and practices
- Insights about how we can adapt to the future environment
- The future: FCC Era(s)

Distil trends and high level directions from current landscape, making predictions w/o focussing too much on single initiatives.

CAVEAT: *Prediction is very difficult, especially if it's about the future - N. Bohr*

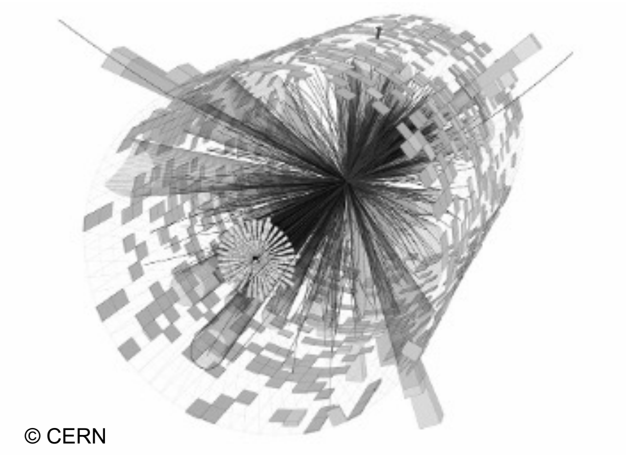
© CERN

The HL-LHC and EIC Era

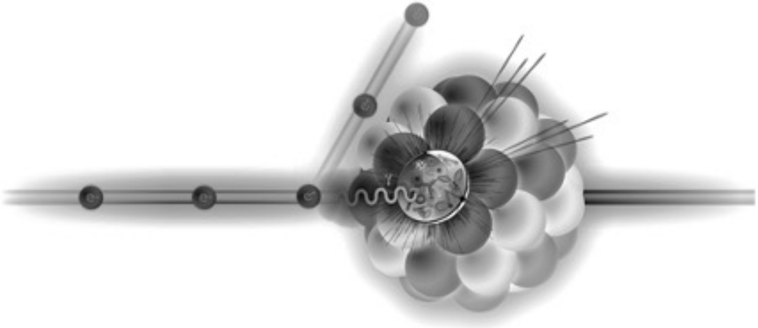

Image credit BNL

Timelines: 2030-2040

	~10% of the dataset			~90% of the dataset
	LHC			HL- LHC
	Pile-up 0-35	Pile-up 55		Pile-up 140-200
	7→ 8 → 13 TeV	13.6 TeV	HL-LHC Installation Atlas/CMS HL Upgrade	13.6-14 TeV
	190 fb^{-1} ALICE/LHCb Upgrade	450 fb^{-1}		3000 fb^{-1} 4000 fb^{-1}
	nominal lumi	2x nominal lumi		5 to 7.5x nominal lumi
	Run 1, LS1, Run 2 and LS2	Run 3	LS 3	Run 4 - 5 ...

2011 · · · 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 · · · 2040

Image:
https://www.bnl.gov/eic/

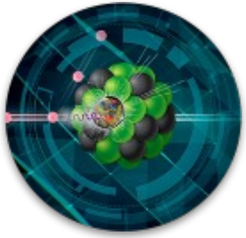

Electron-Ion Collider

2024: Team hopes to start construction

Collider operation date within the 2030s

SKA2 also starts taking data

HL-LHC and EIC will be simultaneously running!

High Luminosity Large Hadron Collider (in a nutshell!)

- Accelerator and detectors upgrade
    - ALICE and LHCb upgraded already now for Run 3
- More data: double Run 1,2,3 integrated luminosity in one year
- Mirror data (10x statistic wrt today) with (many) simulated events
- From ~50 to 140/200 parasitic collisions (pileup) per bunch crossing
- Upgraded detectors: many channels, high granularity
- 4/5x event sizes (stress on network+storage+tape infra)
- Proton physics + Heavy Ions

“Only” 78 simultaneous collisions ©CERN CMS-PHO-EVENTS-2012-006-3

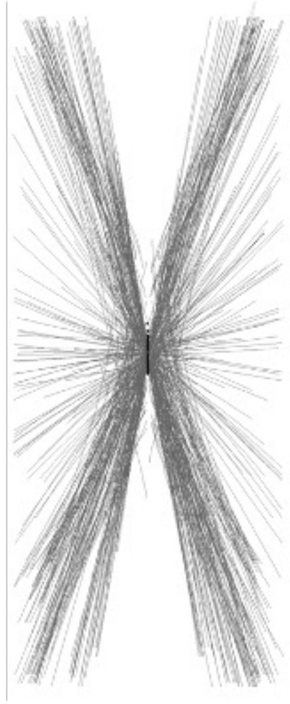

Electron-Ion Collider (in a nutshell!)

- **Understand internal structure of the proton and nuclei at the same level of the electronic structure of atoms**
- Transform RHIC into an electron-ion collider
- 10-100 fb^{-1} yearly
- Large range of c.e.m E=20-140 GeV
- Wide range of ion species
    - From protons to Uranium!
- Large detector acceptance

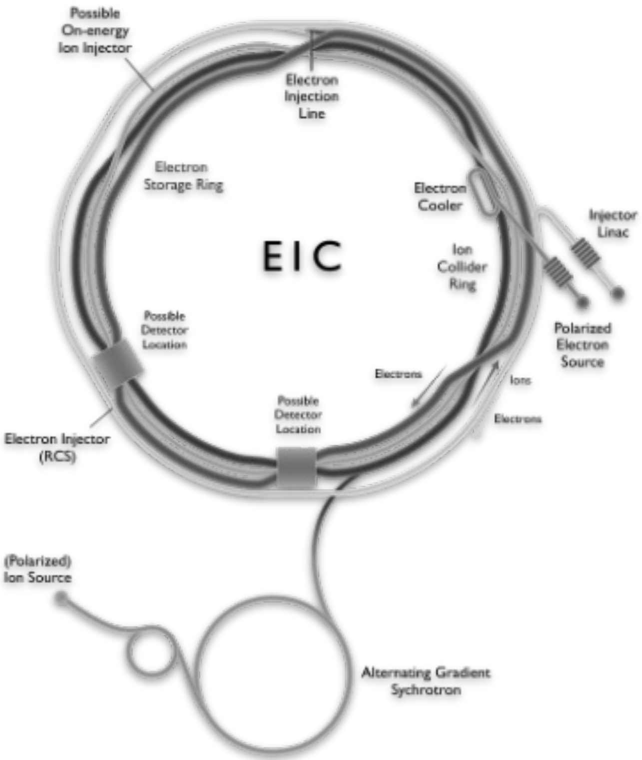

Images credit BNL

Many Differences, but Several Commonalities

- **Two discovery machines** with very different physics goals
- Five big experiments, all striving for ultimate precision, dealing with unprecedented complexity
    - … And lots of data
- Both projects will need efficient software, handling high complexity, with usable and ergonomic interfaces

All that will have to be obtained in a rapidly evolving hardware landscape

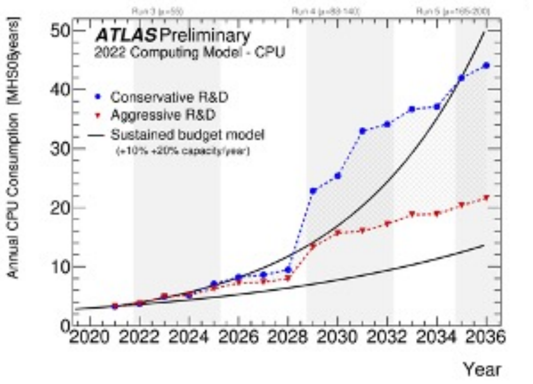

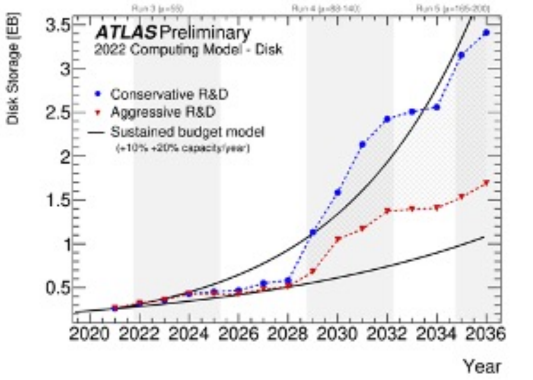

From ATLAS Software and Computing HL-LHC Roadmap


© CERN IBM 709 unloaded at Geneva airport

Hardware in the HL-LHC and EIC Era

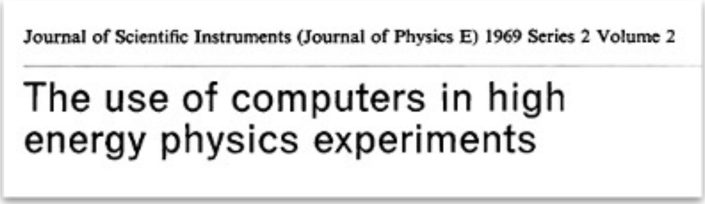
Journal of Scientific Instruments (Journal of Physics E) 1969 Series 2 Volume 2

The use of computers in high energy physics experiments

CERN-DD-DA-68-3, DOI 10.1088/0022-3735/2/1/201

Trends in Compute Hardware

- **Large uncertainty affecting all predictions**
- Hardware cost dominated by market trends, which (our) science cannot influence
- **Necessity to diversify architecture (for many years to be essentially focusing on a single one)**: CPU (x86_64, Power, ARM, RISC) and accelerators (e.g. GPUs, departing from single vendor leadership)
- Heterogeneity is hitting the mass market. E.g. for traditional reco/sim algorithms, ML training, inference.
- **Power consumption becoming more important than ever**

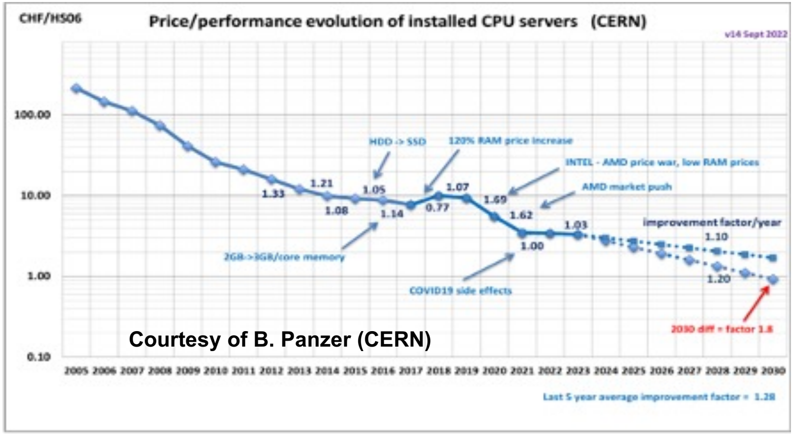

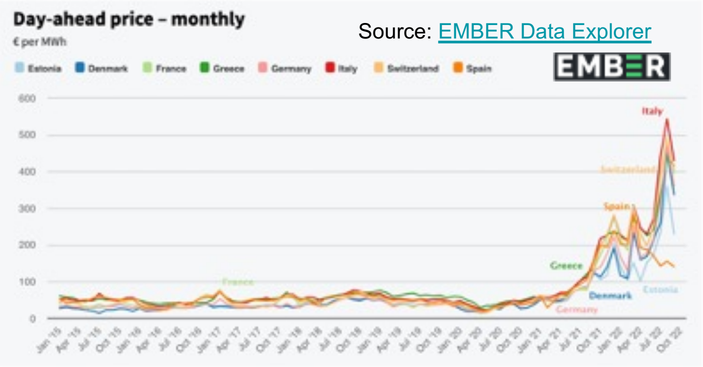

50 years of CPUs

Some easy take-home messages:

- **We hit the power wall 20y ago**
- Luckily, technology steadily provides to us mechanisms to increase runtime performance
- **Data+task parallelism *have to* be at the heart of algorithms and frameworks**
- Computing capacity *must also come* from devices other than CPUs
- **New architectures entered the game!**

***Perfectly* exponential**

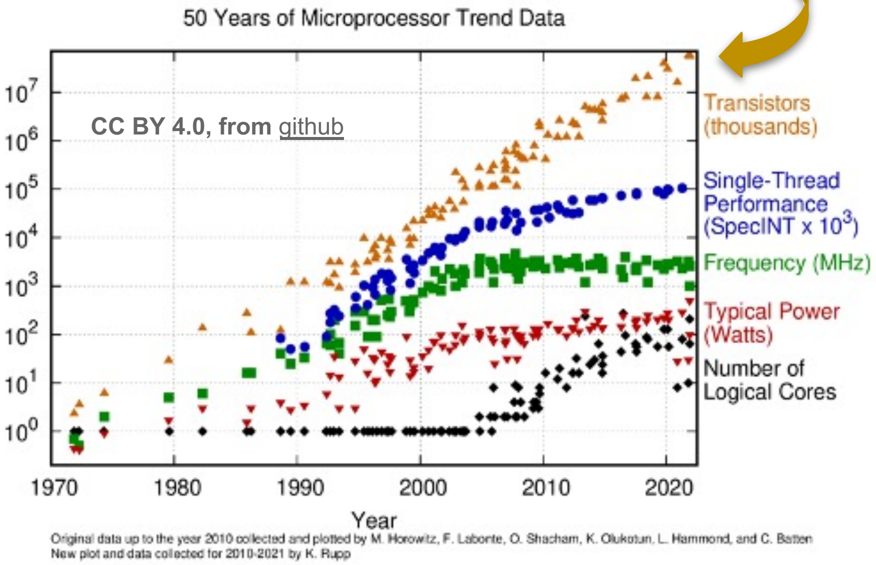

High Performance Computing

- **HPCs more and more visible in the scientific computing infrastructure**
  - Massive national and supranational investments
- Several Exascale machines will be available during the next decades
- HPCs used by all LHC experiments to some extent
  - **Does a future scenario w/o HPCs for HENP computing even exist?**
- HPCs for High Throughput Computing come at a price, e.g.
  - Data access (access, bandwidth, caches ...)
  - Environment less open than Grid one (OS, access policies, ...)
  - Node configuration (low RAM/Disk, ...)
  - Relationships between providers and experiments are decades long
  - **Top performance comes with exotic architectures**

Solving these challenges also requires investing in people

Some trends in Storage

- Tape: libraries+medium. Hard to show clear cost trend. **LTO media technology seems healthy, still one single vendor for libraries, IBM.**
- Disk: purchasing power increase somewhat stopped
- **Storage management technologies can come to rescue**, e.g. Erasure Coding (one shot improvement)
- EC: more logical space with the same physical space
  - Used in production by ALICE for its O2 100 PB (!) disk buffer at CERN
- **Sync'n'Share very popular, also for scientific data**, e.g. mirror locally output of Grid jobs for fast analysis
  - Plays well with small data formats!

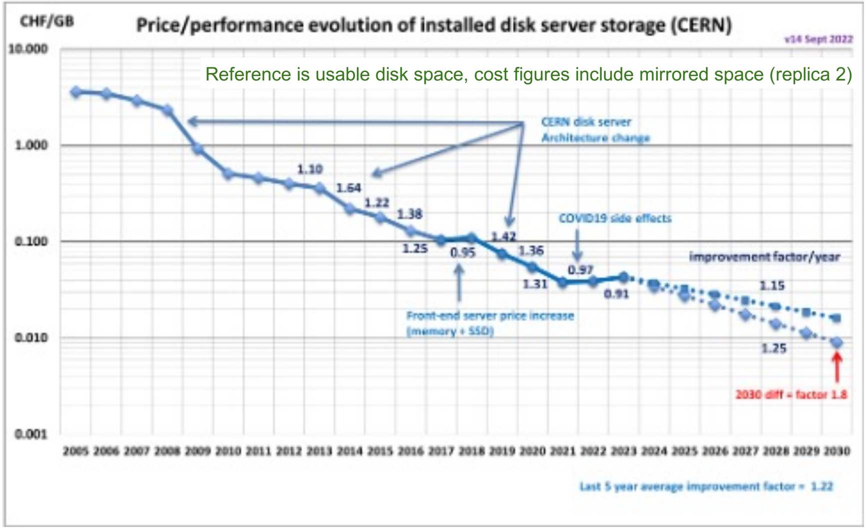

Network

- **The Third Pillar of computing models with storage and compute**
  - E.g. usable to balance the need of other resources
- An Exabyte challenge ahead of us in the next decade
- Network situation seems healthy *now*, however…
- **Not just a matter of total bandwidth, but also when the bandwidth is available**
- HE(N)P will have to share network with other sciences
  - E.g. SKA2

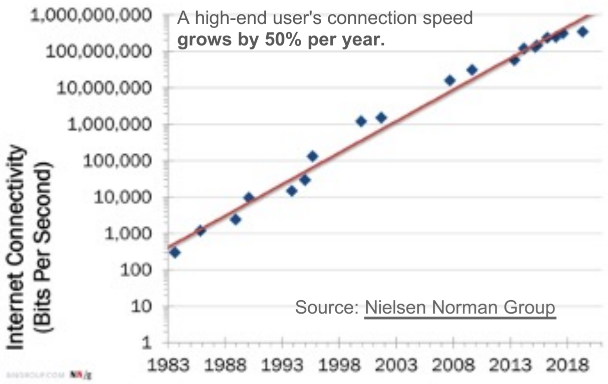

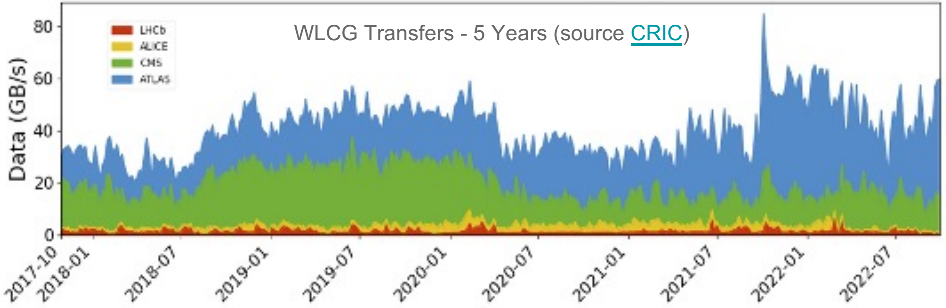

Quantum Computing

- An emerging field of physics and engineering with the potential to revolutionise science and society
- Some fields of interest:
  - Quantum sensors
  - Communication and networking
  - Computing and simulation
- **Currently, quantum computers cannot outperform traditional computers for many tasks**
  - As digital cameras did not outperform film cameras in their early days (nor telephony could outperform telegraphy)
- Difficult to plan on QC for the next decade of HE(N)P computing, but **it is farsighted to explore potential future applications**

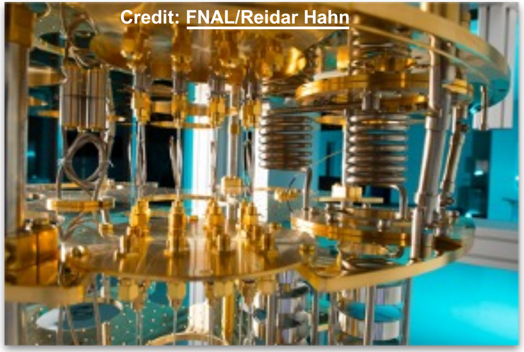
Credit: FNAL/Reidar Hahn

How Can We Adapt?

Adaptation

An adaptation is a feature that arose and was favored by natural selection for its current function. ***Adaptations help an organism survive*** and/or reproduce in its current environment.

Source: evolution.berkley.edu

More *survival* than *convenience*…

United We Have a Stronger Voice

- **Experiments should continue to talk to each other and adopt common solutions (tools or approaches)** ***whenever it makes sense***
  - On the Grid, we were (almost) the sole users, in the new world (HPC and beyond) our relative weight is smaller!
- From the HEP perspective, HPC landscape in Europe looks fragmented, e.g. wrt the U.S.
  - Very different funding schemes, several national entities vs one country
- **Expanding sites with nearby HPCs (or hosting sites at HPCs) works. However...**
- **Is it possible to federate HPC/Cloud resources at European level for the benefit of HEP?**
  - CERN seems to be the natural "federator"
  - Facilitate grants and solution of technical hurdles
  - Lightweight process to join the "CERN HEP Compute pool"
- Examples of working models exist, e.g. OSG or HepCloud

Il quarto stato, G. Pellizza da Volpedo, from Wikipedia

Accelerators

- Heterogeneous platforms (CPU + accelerators) became common for scientific computation, also driven by ML training and inference
- **GPUs used already in Run 3 by three LHC experiments online**
  - LHCb's Allen GPU aided online processing @ 30 MHz, CMS HLT offloading to GPU's 30% of the online sequence, ALICE online successfully used 250 cpu nodes equipped with 8 GPUs with the power of 1800 cpu nodes. ATLAS also heavily engaged (see article)
- Adaptability of applications and workflows might become a necessary condition to use HPCs
  - Scientific excellence alone might not be a sufficient condition anymore
  - Having the two we will have prime access to these resources
- **Heterogeneity paradigm shift: three pillars**
  - Efficient offloading support in a parallel data processing framework
  - Re-invention of algorithms for the extreme accelerator architectures
  - Performance portability (one codebase, many accelerator types)

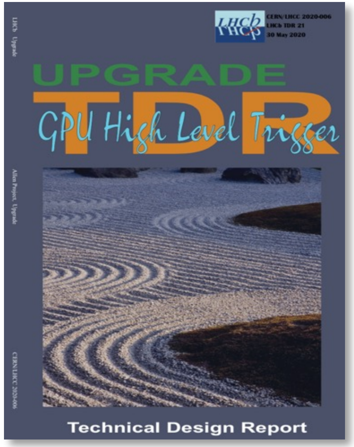

2 AMD Milan + 2 Nvidia T4
@CMS HLT © **CERN**

Data Processing Frameworks

- Nexus of data processing
- HEP: support task parallelism expressed via multithreading paradigm
  - E.g. Gaudi, CMSSW. One process, from RAW data to quality monitoring histos and analysis datasets, algorithms in C++ or for accelerators.
- There are notable exceptions (BASF-II through multiprocessing)
  - Possible to think to independent processes communicating with standardised messages in some context ?
- **Framework development requires constant effort and expertise**
  - **Architects are highly specialised profiles**
  - **Few experts, high costs for training**
  - **Important to attract new effort**

https://gaudi-framework.readthedocs.io

https://software.belle2.org/development/sphinx/index.html

Many Types of Accelerators, A Unified Code Base

- Compilers transform code to binaries for different CPU archs seamlessly
- Often (always?) different languages are needed to support accelerators
- **Performance portability is required (N implementations=too expensive)**
- Different paths to achieve performance portability can be adopted:
  - Experiment specific portability layer or wrappers
  - 3rd party portability library, e.g. Alpaka, Kokkos or OneAPI
- We converged on the TBB for multithreading at LHC
- Can we also converge one portability layer?
- **Perf. portability does not solve the issue of "physics results portability"**
  - Validation of physics results on N architectures is a delicate issue to be addressed carefully

kokkos.org

github.com/alpaka-group

www.oneapi.io

Generation and Simulation

- Precision physics (N(N)LO, fine-grained detector description) calls for increased accuracy and needs more computing power
- Mitigation: constant profiling and code improvements for CPUs, usage of accelerators, either through traditional approaches or ML
- **Improve *the* common tool for simulation, Geant 4**
- Experiment specific fast MC chain: from gen to analysis data formats
- **VERY fast MC chain: from gen to analysis formats w/o intermediate steps through ML**
    - Hard to learn >1k variables per event (full "AOD" info)
    - Is it feasible to learn a few 100's for small end user analysis data formats?

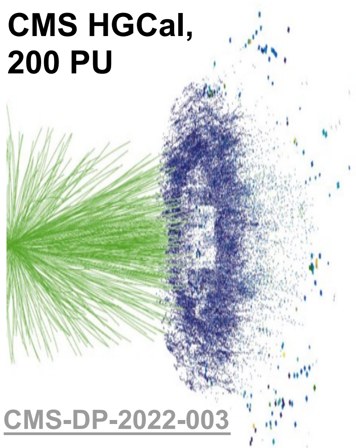

CMS-DP-2022-003

Geant4 - A simulation toolkit

Reconstruction

- Deal effectively with high occupancy
  - 35 PU → 50 PU → 140 PU → 200 PU
- ~~Port~~ Re-invent our algorithms for accelerators
  - Simple data layout, parallelizable
  - Intermediate data products stay on the device
  - Rediscover existing code and write it anew from scratch
  - **Exciting times for physicists-developers!**
- Complement traditional approaches with ML when possible, keeping compute capacity needs in check
  - Flexibility to run inference on CPU and accelerators efficiently

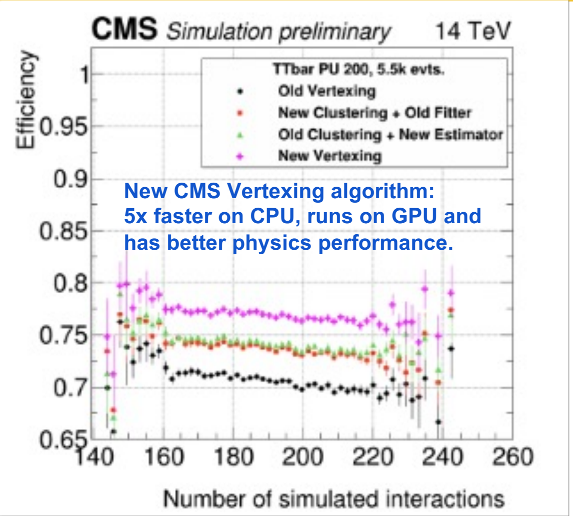

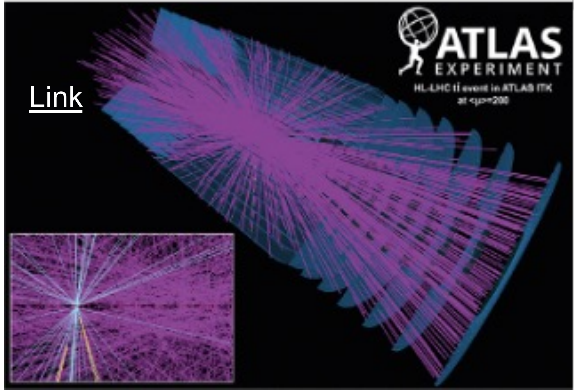

The Role of Machine Learning

- Not new (already used at LEP)
- **Inference engines: hard to integrate them into our sw stacks**
    - Imagined for stand-alone usage, with some exceptions
    - **Reliably transform models into CPU/GPU code?**
- Closely follow advancements in ML research
- **Pave the way to paradigm shifts e.g. completely replacing current practices ("end-to-end" ML approaches...)**
- Address at our labs/institutes the needs of model training
    - Can be heavy, both computationally and in terms of IO
- Computing: serve data for training fast, make accelerators easily discoverable.
    - Possibly use also remote accelerators through network?
    - Challenges ahead: scheduling of workloads, load on the network

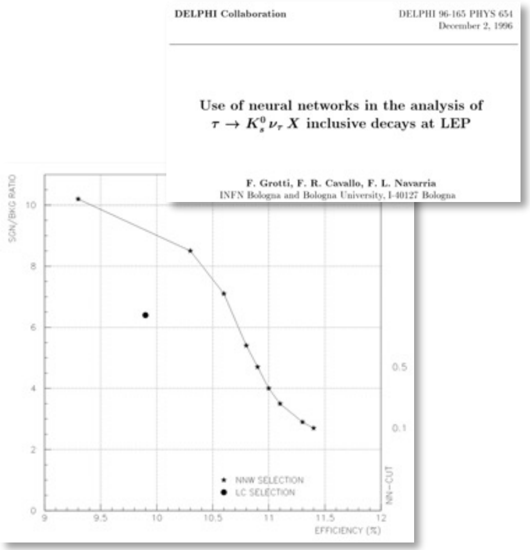
DELPHI Collaboration

DELPHI 96-165 PHYS 654
December 2, 1996

Use of neural networks in the analysis of
$\tau \to K^0_s \nu_\tau X$ inclusive decays at LEP

F. Grotti, F. R. Cavallo, F. L. Navarria
INFN Bologna and Bologna University, I-40127 Bologna


Image Credit: FNAL/CERN

End User Analysis

- **Future analysis will have to adapt to available hardware**
- **Need for small data formats easily readable efficiently, potentially by more than one tool**
- Further evolve our existing analysis facilities to ease the task?
    - AFs do exist: lxplus+lxbatch at CERN, ALICE's AF at Wigner, ALICE's AF at GSI, LPC Cluster at FNAL, SWAN at CERN, Tier-2's, group clusters at universities...
    - Through data centre design obtain even more efficient delivery of data, data caching, true interactivity, dedicated nodes with many cores, fast storage, etc.

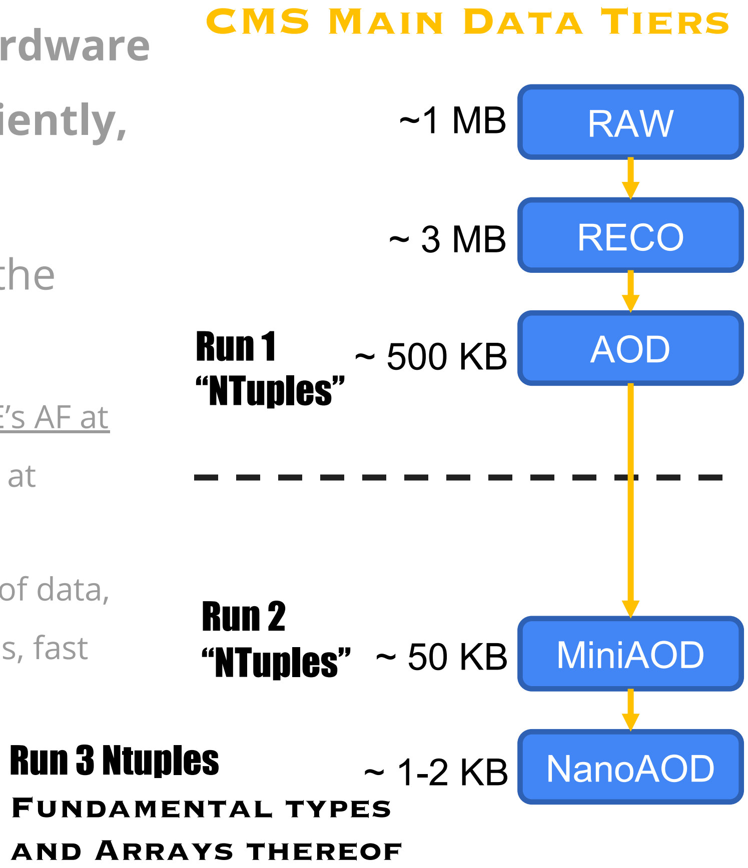

Imagining Future Storage

- Overcome disk-tape duality and reason in terms of cold and warm storage:
  - Tape vault for the data you won't reprocess soon but want to keep?
  - Stateless data caches?
  - Very fast storage for analysis at selected locations?
- One or two copies of RAW data on tape?
  - One for ALICE, and, selectively for CMS (parked dataset)
  - Not an easy one: lots of implications to discuss
  - What about one copy of some form of "reduced RAW"?

- Federate storage as we can federate compute resources?
  - Multiple storage pools behind a single entry point, with different quality of service
  - Sync'n'share federations do exist already

Is a Revolution always better than an Evolution?

Probably not. Our challenges can benefit from an evolutionary approach:

- Improve reduced analysis data formats guided by analysis community
- Optimise detector specific fast and the Geant 4 based sim
- Strive to increase code performance
- Agile physics validation and constant profiling
- Re-invent existing algorithms for accelerators
- Bring successful R&Ds to production


La Liberté guidant le peuple, E. Delacroix, from Wikipedia

The above

1) **Yielded and will yield substantial reduction of computing resource needs**
2) **They need constant support for maintenance and operation tasks by FA's**

See also K. Bloom, Snowmass 2021

Effort and Expertise

- **Continual improvement of existing software/tools as well as R&D is needed**
  - We must invest in both
- Training efforts: need to get people engaged
  - One single motivated individual can make the difference!
- Allow for a healthy diversity of ideas and influx of new members
- **Career paths** need to be provided for software/computing professionals, much like we do now for engineers.
- **Invest in young talent, work to create the most rewarding and welcoming work environment**

The FCC Era(s?)

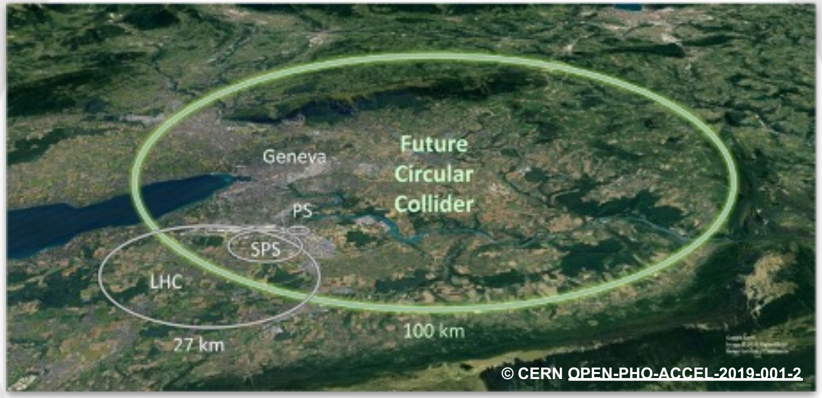
Geneva
Future
Circular
Collider
PS
SPS
LHC
27 km
100 km
© CERN OPEN-PHO-ACCEL-2019-001-2


cottonbro on Pexels

How far in the Future?

- Above schedule: FCC-ee starts in ~30 years from now
- Need to strike a balance between full HL-LHC exploitation and FCC delivery
- Different experiences in the past
  - LEP→LHC: 10 years (same tunnel)
  - *SppS*→LEP: -2 years.
- **Key Factor: human and financial resources**

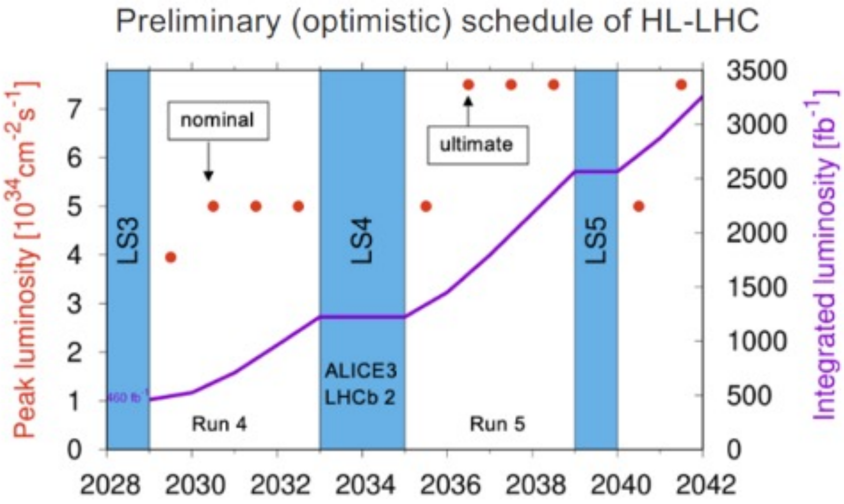

See also: F. Simon, G.Wilkinson 5th FCC Physics Workshop

Status of the Computing and Software

- **Computing resources are not unsurmountable!**
- LHC-scale computing is nearly sufficient (within some uncertainty factor) for simulation needed for the Z-pole run of a FCC-ee detector. Storage similar to full HL-LHC.
- **Current approach based on edm/key/dd-4hep: seems a solid basis** for the feasibility study
- Leveraging commonalities across future-collider community broadens engagement
- **Continuous need to emphasize the bridge from core development to integration and support of FCC detector geometry and algorithms**
- Encourage documentation (from the core outward..), good examples, clear recipes
- Be welcoming of software contributions and encourage developers to **take the last step of integrating their work with the collaboration software**

See also D. Lange: FCC Physics Workshop 2022

Wrap-up

Conclusions

- Two discovery machines will start operations in the 2030's: EIC and HL-LHC
- Different approaches, but several commonalities. Both will have to operate in a rapidly evolving hardware landscape
- Actions can be taken to adapt, e.g.
  - Keep storage usage in check, striving for the smallest data formats
  - Efficient usage of accelerators and HPCs: paradigm shift ahead of us!
  - Well monitored and managed network usage (HENP will not be alone using the infrastructure)
- All that can benefit from an evolutionary approach
  - Operations, integration of successful R&Ds
  - Without forgetting bold, revolutionary new techniques of which we may take advantage in the future

Backup

Programming Languages in HEP

- Hard to summarise in one slide: a full ~~talk~~ lecture might be needed!
- Broad categories of tasks in HEP and current language (approximate) mix:
  - data processing (algorithms and framework): C++/ (abstractions of) CUDA-HIP
  - data analysis: C++/Python
  - computing tools/services: Python/Go
- Golang: emerging trend for computing tools and services
  - Compiled, statically linked, portable binaries, plays well with containerisation and k8s deployment, simpler management of dependencies and packages
- **Hard to predict *exactly* what we'll use in the next two decades: we have to keep our eyes open!**

More References:

- HL-LHC timeline: https://hilumilhc.web.cern.ch/sites/default/files/HL-LHC_Janvier2022.pdf
- EIC Dates: https://www.bnl.gov/newsroom/news.php?a=219454
- EIC Project director https://indico.cern.ch/event/934666/contributions/4154187/attachments/2168360/3660239/EIC%20Project%20Overview%20Epiphany%20Yeck%20sm.pdf
- Lancium https://indico.cern.ch/event/1160140/contributions/5008195/attachments/2503004/4300041/2022-Sustainabilitiy-Lancium-and-HEP.pdf
- K. Bloom @ Snowmass Future Computing in HEP: https://indico.fnal.gov/event/22303/contributions/245364/attachments/157980/206998/bloom_resource_gaps.pdf
- FCC Physics Workshop: https://indico.cern.ch/event/1066234/timetable/#b-449470-next-steps
    - Closing remarks: https://indico.cern.ch/event/1066234/contributions/4708145/attachments/2390022/4085522/PED_wkshp_closing_2022.pdf
- Fermilab quantum initiative https://quantum.fnal.gov/
- CERN Quantum initiative https://quantum.cern/