R&D: HEP Computing Infrastructure

Maria Girone, CERN

Update of the European Strategy for Particle Physics, Granada, 15 May 2019

Outline

- The need for exa-scale computing
- Moore's law evolution
- General trends and emerging technologies in the computing industry
  - Developments in CPUs
  - Exploitation of Specialized hardware
    - GPUs, TPUs, FPGAs
    - Storage Specialization
- R&D activities with industry
- Shared challenges with Data Intensive Sciences

Toward the HL-LHC and exa-scale

- HL-LHC is an **exa-scale** computing and storage challenge
  - 6x Pile-up(PU), 10x HLT rate, technology improvements of 15%/y → with a naive extrapolation of current needs, factor order 5-10 is needed

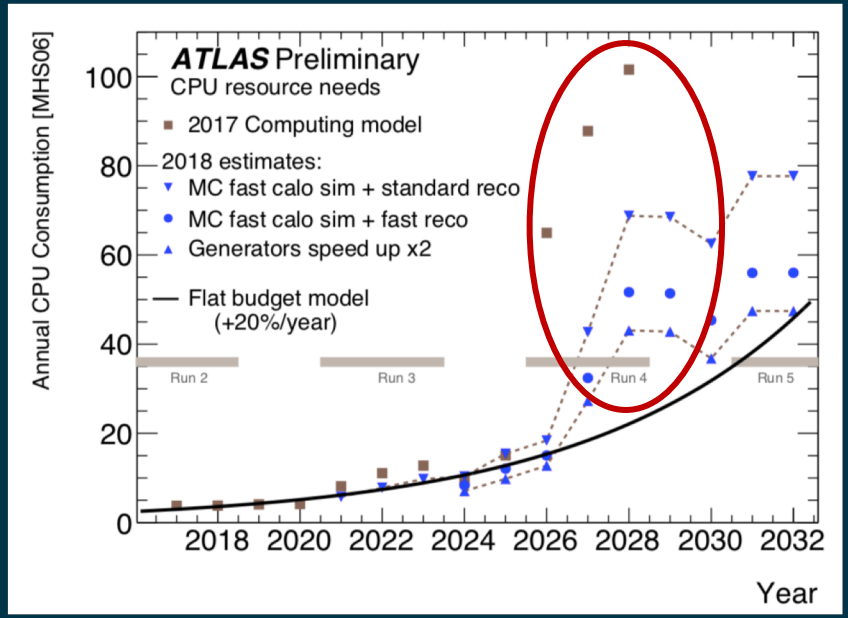

- Experiments are proposing **significant changes** in their **models**
  - Big factors of gain estimated by ATLAS with adoption of **faster simulation and faster reconstruction**
  - CMS is moving to even more **size-optimized storage formats**

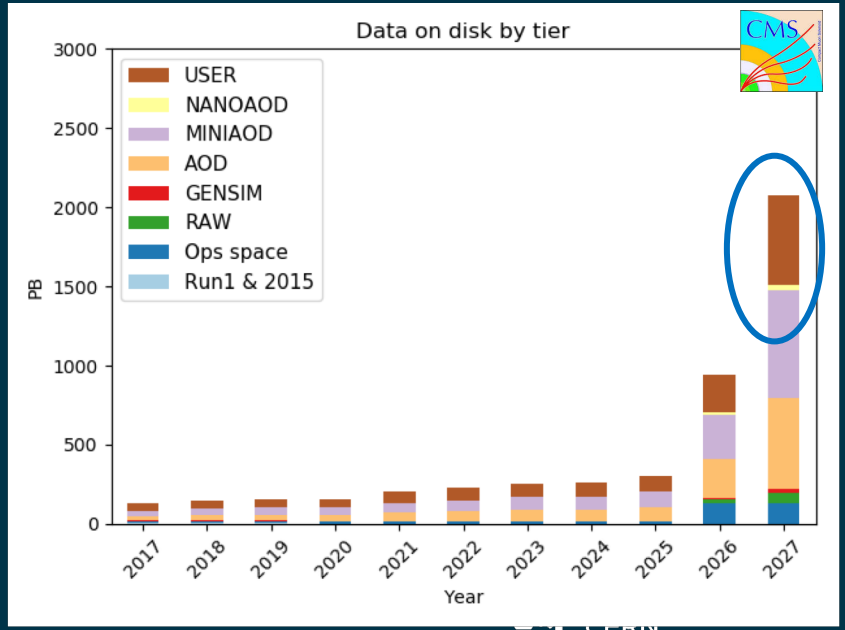

Computing Evolution

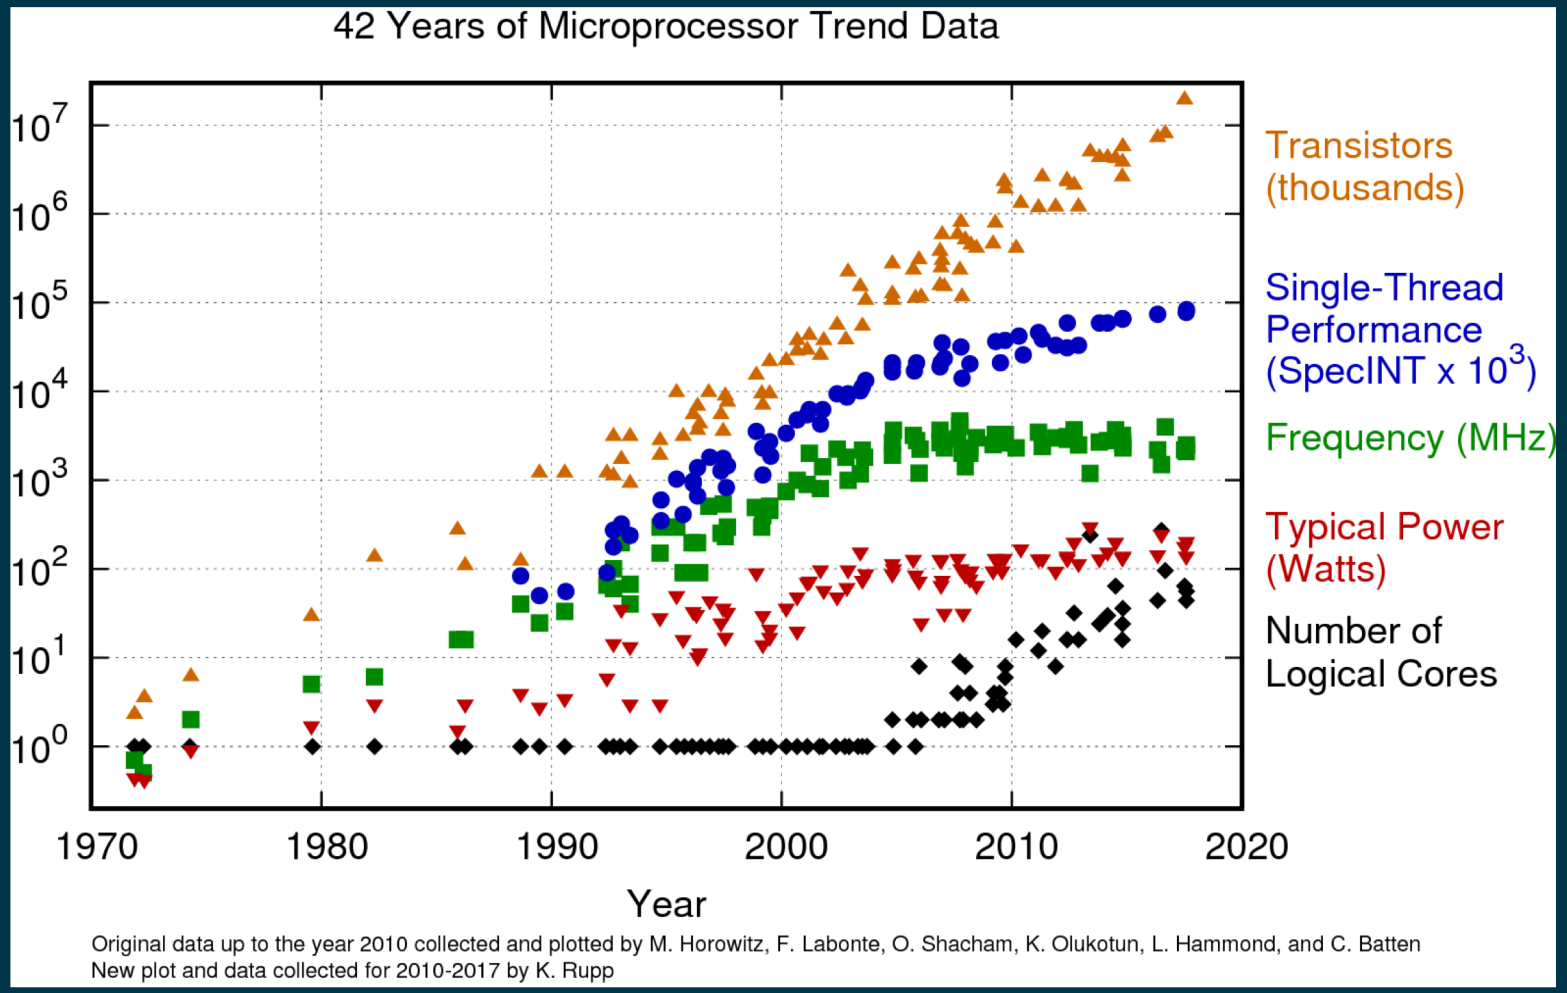

- Moore's law is barely holding on
  - CMOS technology approaches fundamental limits
- Clock speed has been largely flat since 2007
- Number of cores continues to increase

What shall we do?

- We need to better exploit the hardware on the market
  - Faster software will significantly reduce the the resource gap
    - Vectorization, parallelization, CPU specific instruction sets and optimized libraries [Next talk]

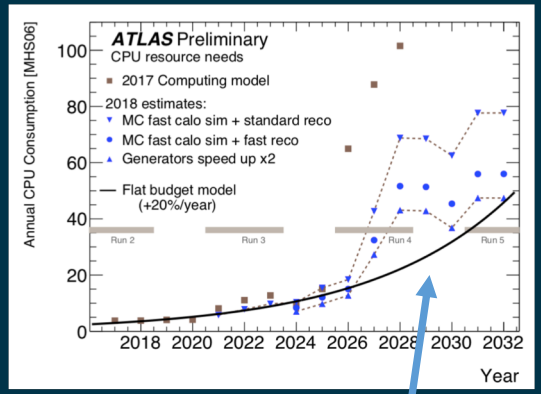

- There is specialized accelerated hardware
  - GPUs, TPUs, FPGAs (new languages and programming paradigms)
- New techniques and methods
  - Machine Learning
- Disruptive technologies:

  Neuromorphic, Quantum Computing

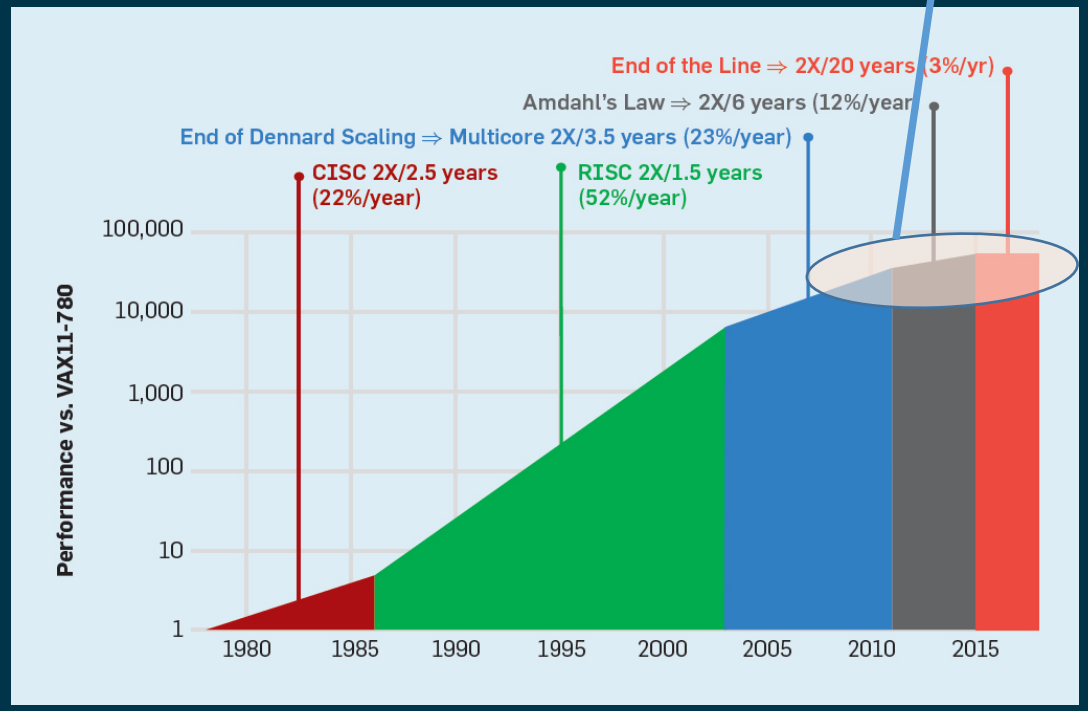

https://cacm.acm.org/magazines/2019/2/234352-a-new-golden-age-for-computer-architecture/fulltext

Improvements are flattening

- **A very optimistic 15% growth is used for calculations!!**

CPU Improvements

- General purpose processors are currently the primary workhorse of HEP computing
  - **x86** architecture drives our software and computing
  - **Our community doesn't drive architecture and design**

- **Intel** dominates the market
- **AMD** offers smaller feature size and higher core count (Zen2)
- **IBM** is an early adopter of technology in the Power series
  - Next generation PCI and NVLink
- **ARM**: we ported our software
  - Watching evolution (performance, power, cost)

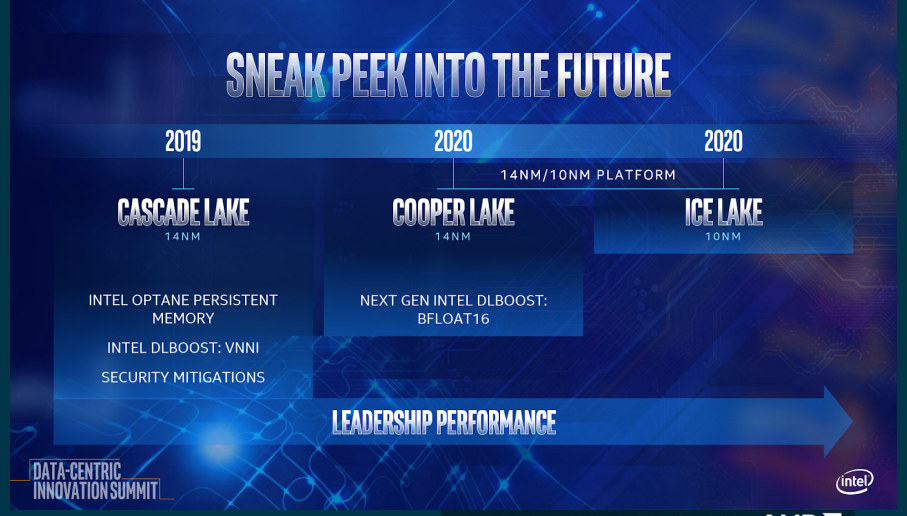

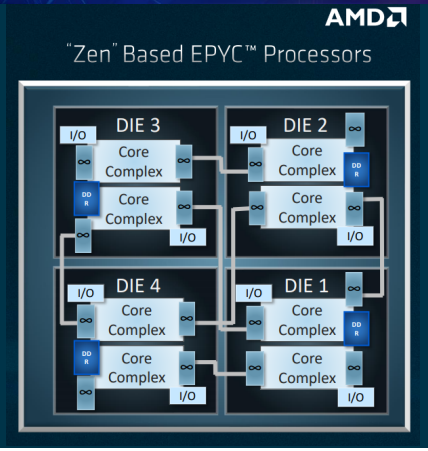

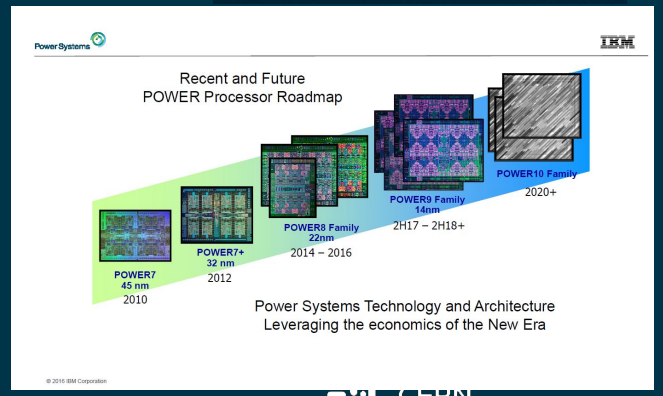

GPU Improvements

- Extremely parallel architectures are widely used in industry and other sciences
  - Can provide better performance/price and energy efficiency
  - GPU performance is on a an 18-month doubling cycle

- R&D effort in the LHC experiments to offload parts of the **reconstruction workflows** to GPUs
  - For both high level trigger and offline (on CUDA)

- ML frameworks profit from GPU acceleration and substantially reduce turnaround times
  - Industry has invested heavily in optimizing machine learning libraries on GPUs (training and inference)
  - Increasing R&D effort in HEP

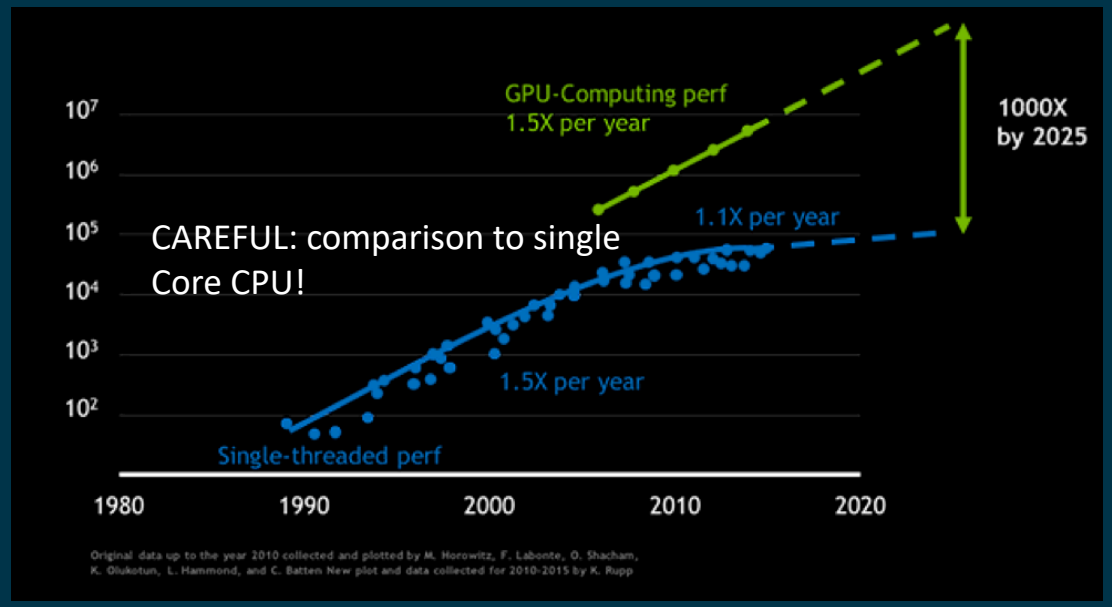

Industry is producing low-power GPUs
- Can be used to upgrade existing machines
- Latest **NVIDIA** GPUs (i.e. T4) use 25% of the power at significantly lower cost
- **AMD** GPUs are tested; excellent performance/price
  - Software is catching up

Workflow Specific Accelerators

- Industry has been investing in alternative architectures
  - **FPGAs** can be programmed to execute a calculation at the logic gate level
    - Good for inference, pattern recognition and data transformation with good power to performance ratio
    - Low level programming
  - **TPUs** are custom ASICs designed for ML
    - Excellent memory bandwidth
    - Less expensive than GPUs with a much higher performance per watt

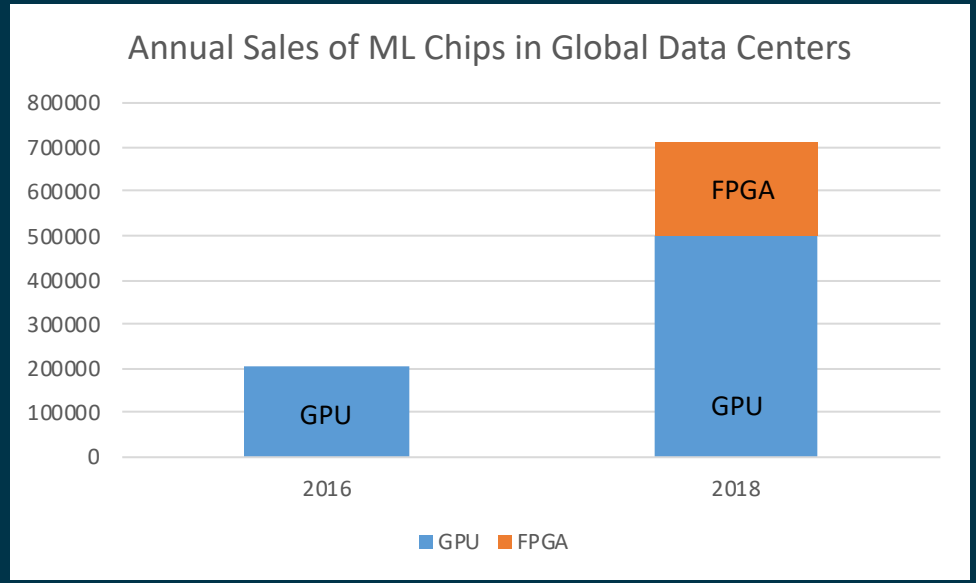

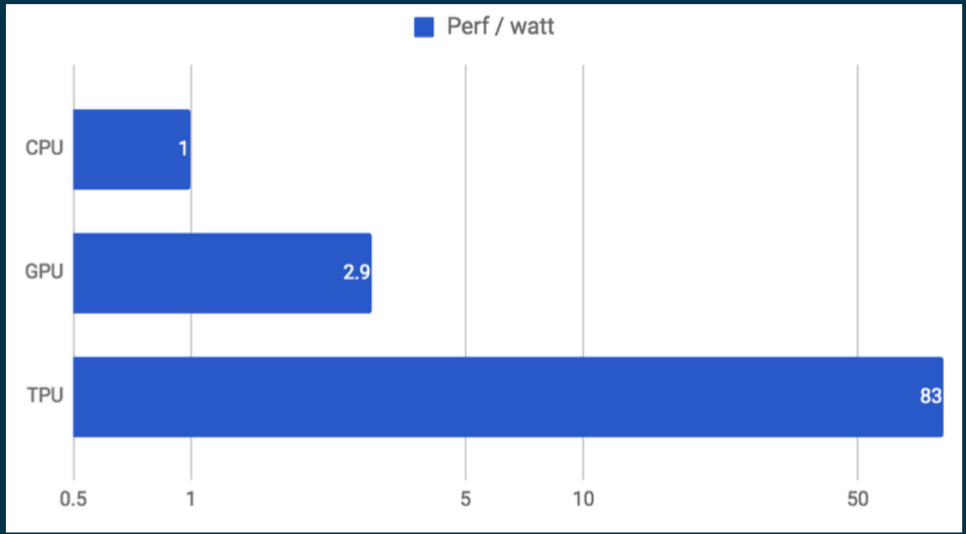

Industry and HEP Applications

Technology	Industry Applications	HEP Experiments R&D (see backup slides)
GPUs 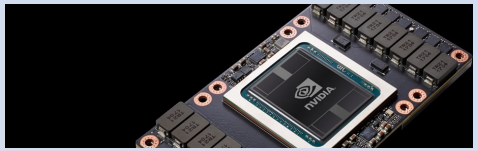	Machine learning, AI, block chain, signal processing, voice recognition, etc.	Online pattern recognition, object reconstruction, fast simulation, data quality monitoring, computing resource optimization, machine learning training and inference
TPUs 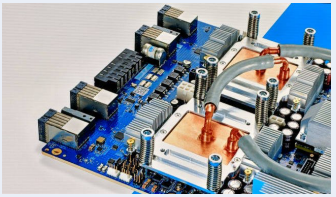	Machine learning inference and training, data processing, matrix multiplication	Machine Learning inference and training
FPGAs 	Pattern recognition, inference application, low latency real-time applications	Triggering, machine learning inference, image recognition
ARM 	Mobile applications, embedded systems, low power highly parallelized many core systems	Alternative low power general purpose architectures

An example of R&D on reconstruction on GPU
CMS - Patatrack

- Demonstrated advantage of heterogeneous reconstruction (on GPU) from RAW to Pixel Vertices at the CMS HLT
  - 1 order of magnitude both in speed-up and energy efficiency wrt full Xeon
  - Running within the CMS software framework
- Parallelization of more algorithms to run in production during Run 3 and Run 4
- R&D now includes ECAL/HCAL local energy reconstruction

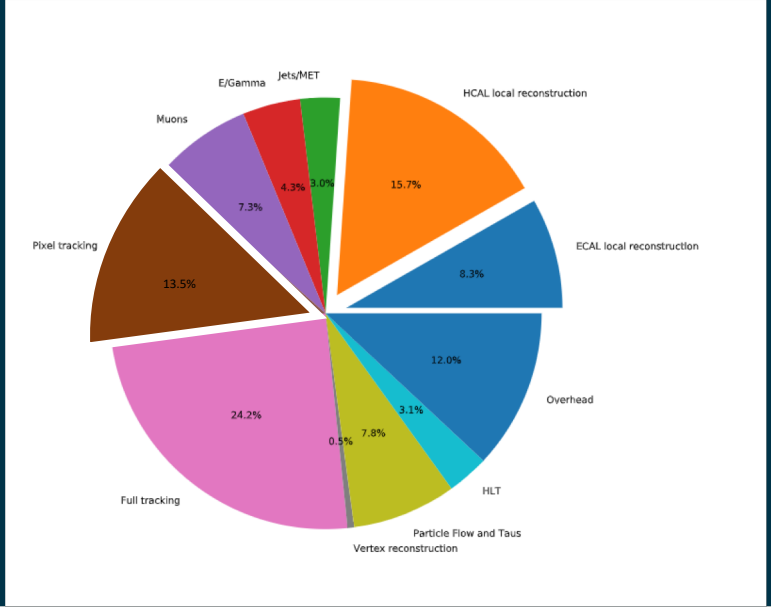

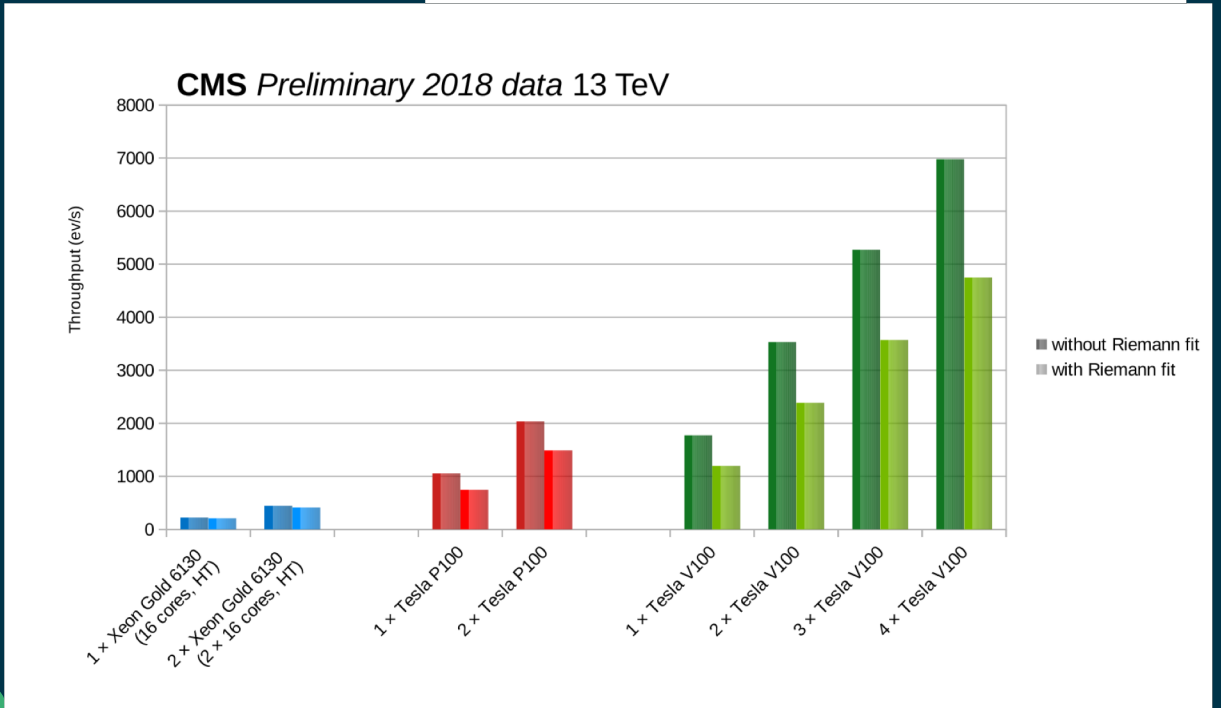

More info: A. Bocci, F. Pantaleo

Fast simulation with GANs

- Classical electromagnetic shower simulation can be very demanding
- Demonstrated three orders of magnitude better performance by running inference on a GAN wrt to classical approach
- Quite accurate physics performance

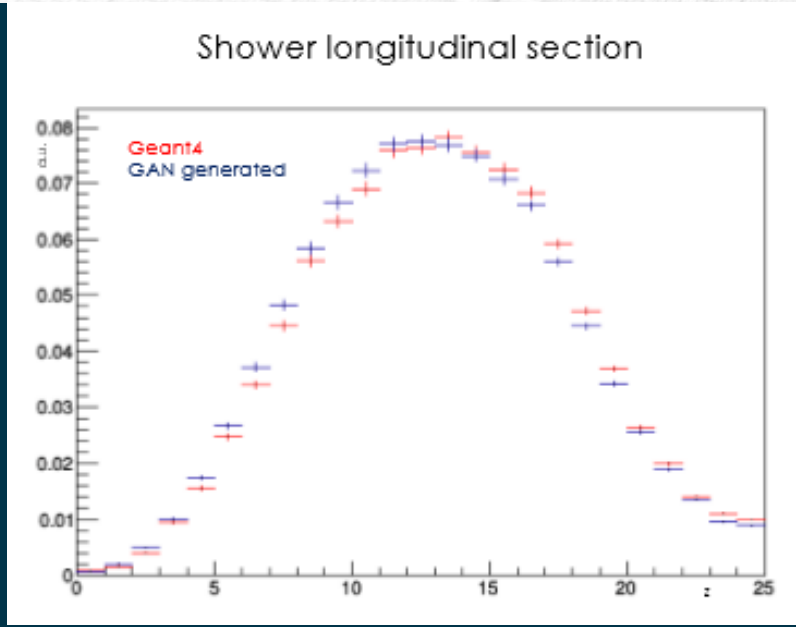

More info: F. Carminati, G. Khattak, S.Vallecorsa

FPGA-accelerated Deep Learning for HEP

- Challenge: preserve physics discoveries at future large-scale experiments with limited computing resources
- Possible solutions: (ex, LHC detectors)
  1. **recast physics problems into a deep learning problem** ➔ DL proved high efficiency for physics of interest with large noise rejection
  2. **improve event filtering** to reduce data rates to manageable levels ➔ pure FPGA implementation for DL algorithms running @ 40MHz with O(μs) latencies
  3. improve current CPU-based algorithms with **heterogenous FPGA+CPU systems**

ex, data flow @ ATLAS, CMS

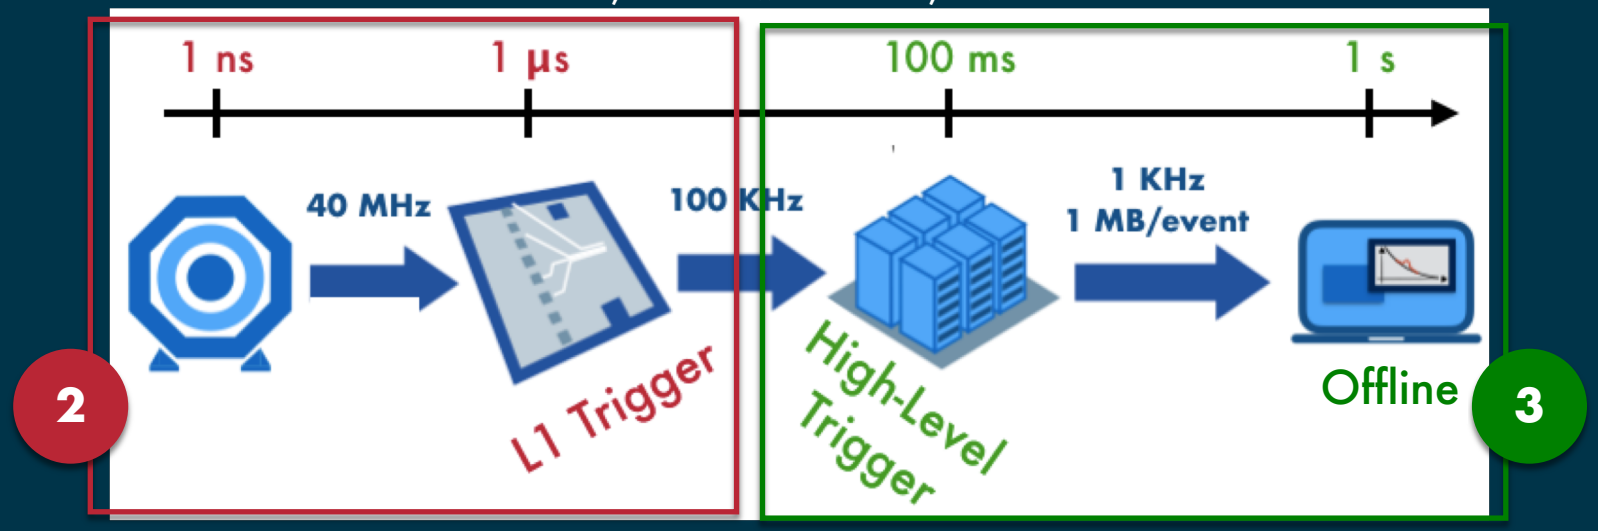

Jennifer.Ngadiuba@cern.ch

Storage Evolution

Storage remains a challenge as data volumes increase

- Similar to processing, industry is specializing

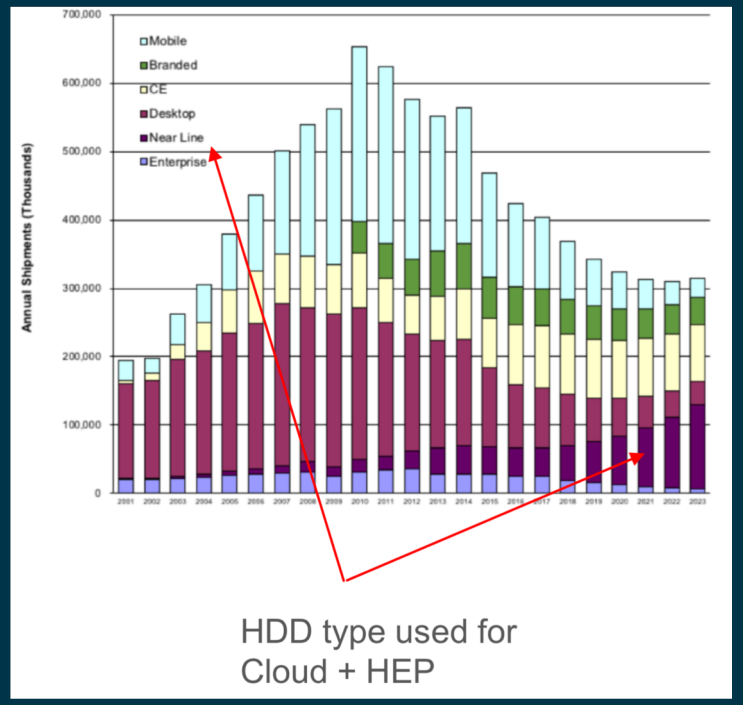

- For disk and directly accessible memory
  - Technology improvements **for spinning disk** continue for capacity but slower improvement for bandwidth and IOPS (Shingled disk, Helium drives, MAMR and HAMR)
    - Market is shrinking
  - High capacity, fast and affordable **SSD**-flash based are taking over
  - Smaller but extremely fast solutions like 3D Xpoint NVRAM blur the boundary between memory and storage, for a price

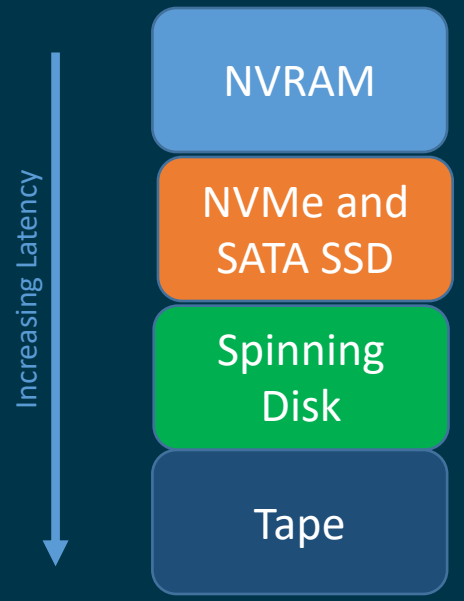

Archival Storage

- **Tape** remains the preferred (cheap) solution of HEP experiments for long-term storage
  - As data volumes increase, LHC experiments are looking at larger roles for tape for data with very structured access
    - Tape roadmap extends another decade and a factor of 20 in capacity

- Disruptive technologies on the far horizon for archiving
  - **DNA Storage** relying on the density and stability of DNA molecules to encode large volumes of data

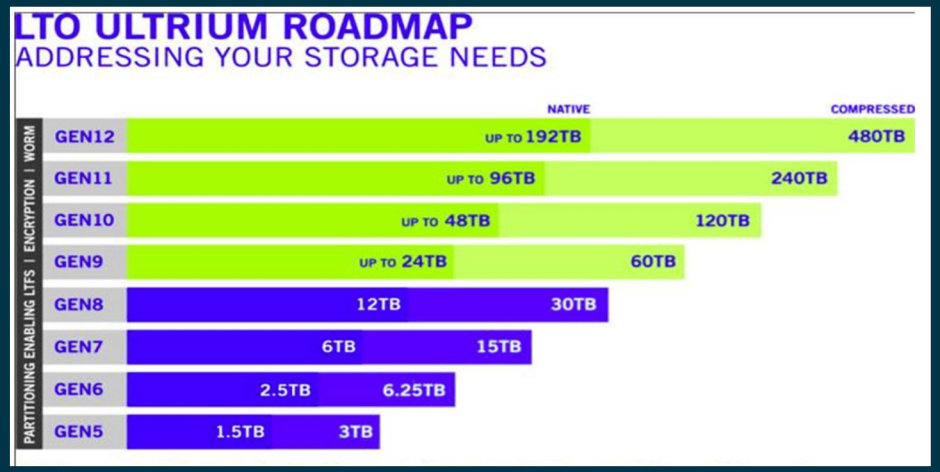

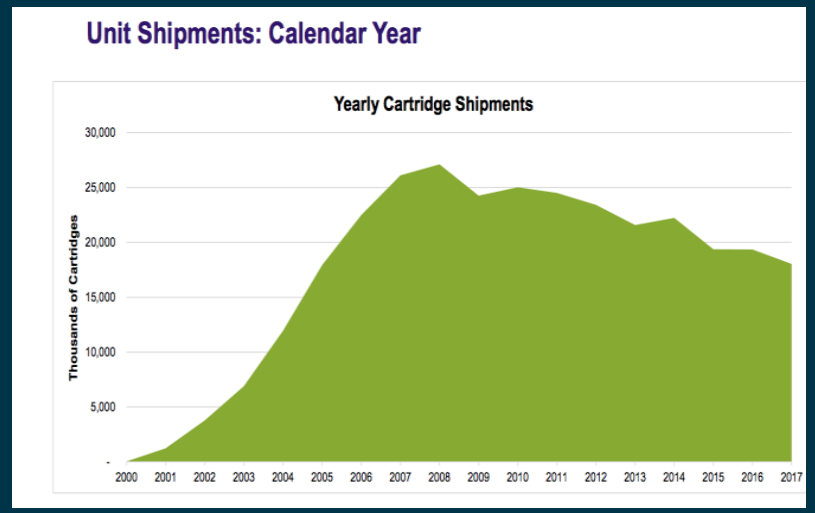

ID: 53, 59, 79, 162

Investments at HPC Sites

- **HPC sites** will **grow** by a factor 20 on the time scale of HL-LHC
  - Large investments in USA, Europe and Asia
  - Mostly available to sciences through **allocations**

- **High performance accelerated hardware** will provide the **bulk of** the processing capacity (~10% expected from CPUs) . Implies:
  - R&D on software and techniques to better exploit processing architectures, low latency networking, edge services, ..
  - R&D in interoperability across HPC (and Cloud) sites, via consistent services (containers , CVMFS, ..)

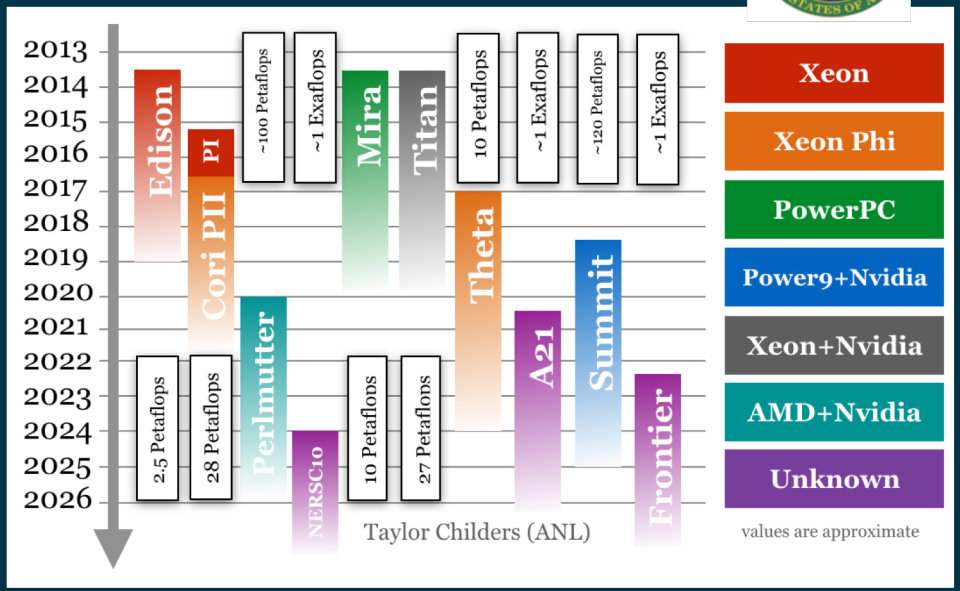

Name	Supplier & Model	Compute power (Peak) [GFLOP/s]	Energy consumption under Linpack (kw)
Piz Daint gpu	Cray XC50 (Hybrid)	25'326'000	2'272
Piz Daint mc	Cray XC40 (Multicore)	2'193'000	734
Installed 2017			

CSCS Centro Svizzero di Calcolo Scientifico Swiss National Supercomputing Centre

Engaging and working together with Super Computing Centers is essential for HEP. Opportunities for training in software engineering

PRACE EuroHPC Joint Undertaking BDEC2 BIG DATA AND EXTREME-SCALE COMPUTING DEEP Projects ETP 4 HPC EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

Maria.Girone@cern.ch

CERN openlab

Commercial Cloud Investments

- Commercial clouds have been investigated by experiments and sites
  - Challenges of scale and data access are similar to HPC
  - Elasticity to absorb peaks of experiments' activities well demonstrated

- **Costs** (compute, storage, data egress) are evolving to be more competitive
  - Watch the (positive) trend of the market

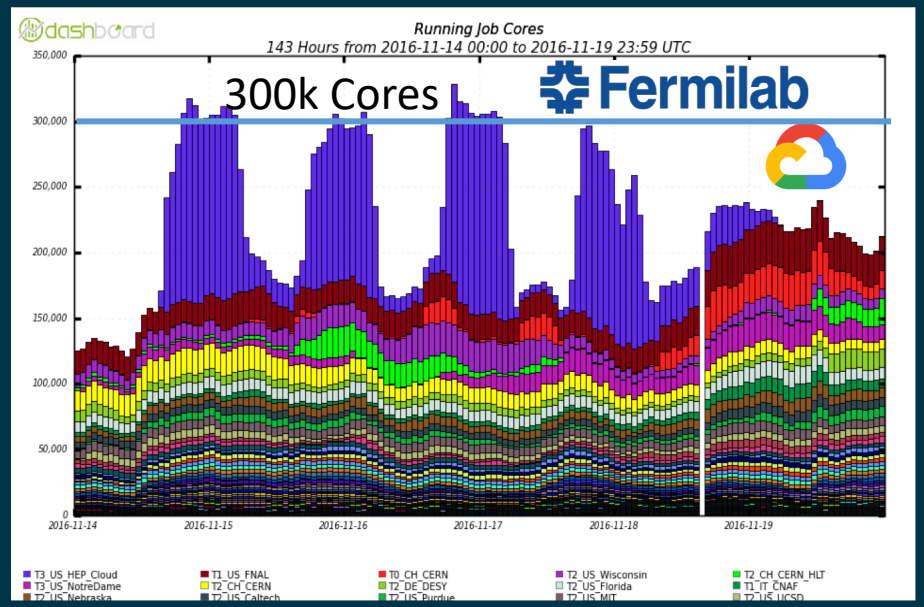

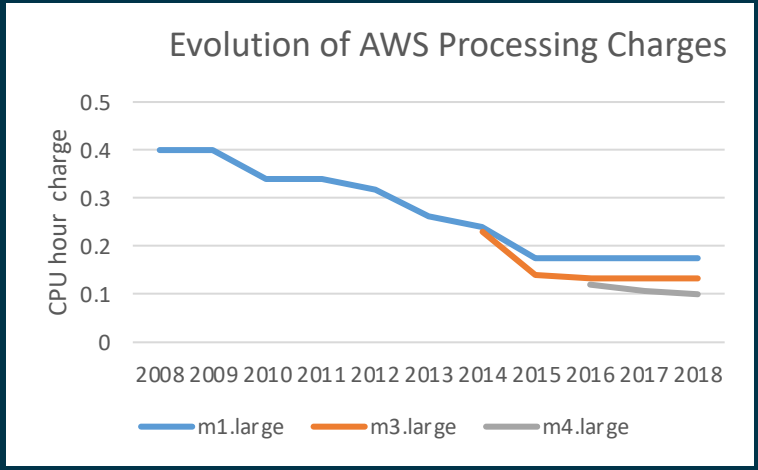

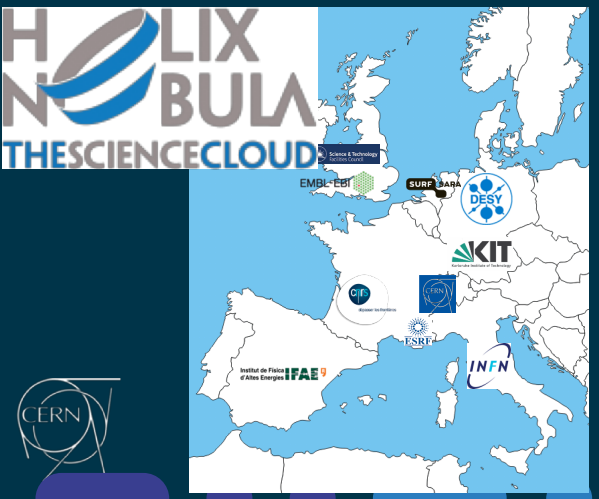

EU supported initiatives within EOSC for cloud procurement and adoption

EOSC
The European Open Science Cloud

OCRE Open Clouds for Research Environments

Maria.Girone@cern.ch

CERN openlab

Networking, Data and Workflow Management

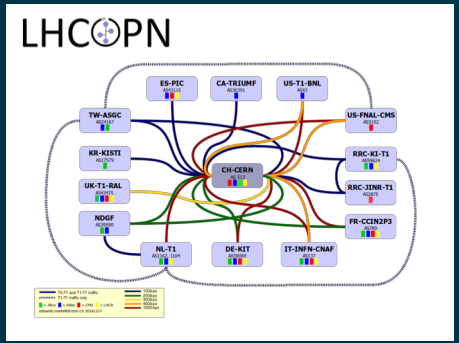

- Increasing size of dedicated resources, commercial clouds, and HPC farms will require **Tb/s** networks
  - The original LHCOPN was 10Gb/s, many links have since been upgraded to 100Gb/s
- Currently using ~1M CPU cores continuously
  - HPC and clouds will make bursts much larger

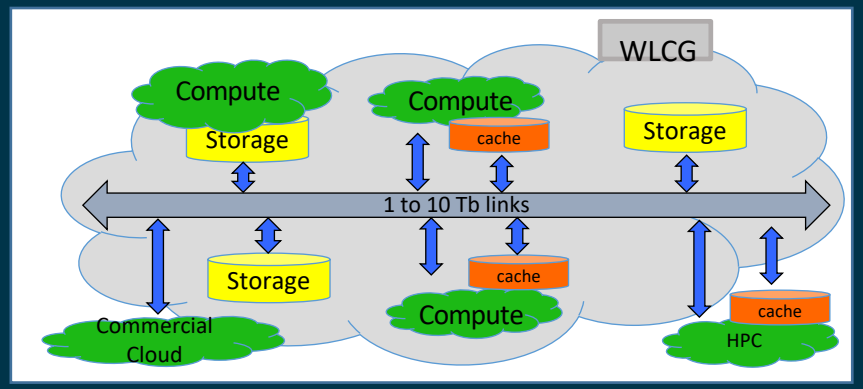

- Ongoing R&D in Data Organization Management and Access **(DOMA)**, experiment **data & workflow management systems** to deliver data and scheduling processing resources at unprecedented scales
  - Rucio, Dirac, … examples of multi-experiment projects with high potential to evolve and embrace the Run3&Run4 challenges

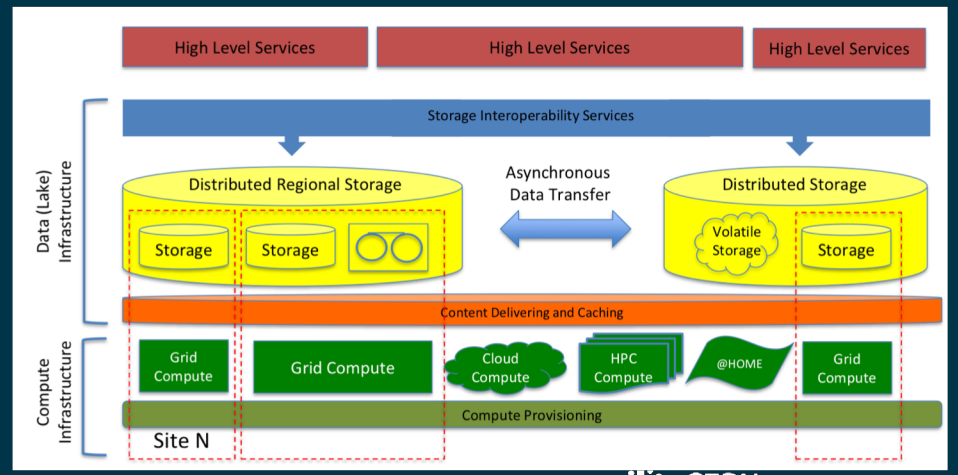

Disruptive Technologies

Neuromorphic Computing

- Very large scale systems to mimic neuro-biological architectures present in the nervous system

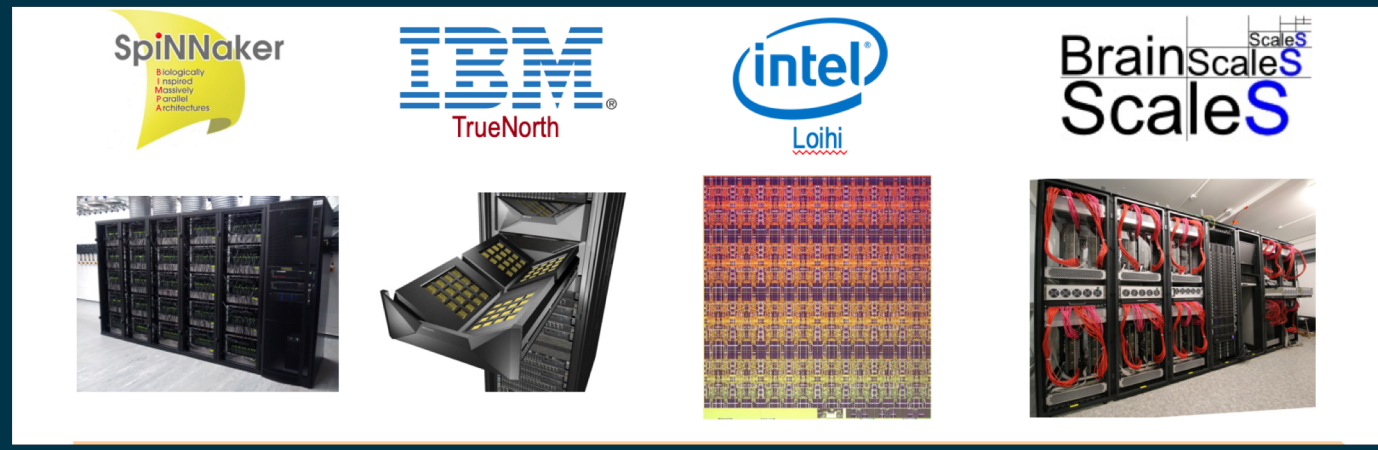

Quantum Computing

- Computing systems relying on quantum properties/behaviors to perform calculations
  - Quantum Annealing like D-Wave systems
  - Quantum Gate models like IBM, Rigetti, Intel, Google

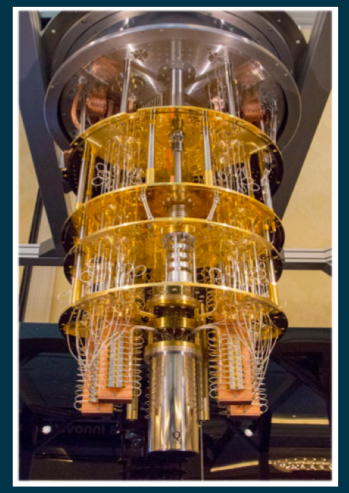

Advances in QC

ID: 59, 128

- Technology is advancing rapidly, yet unclear what the impact will be on the time scale of the HL-LHC
  - investments in quantum programming and application development

Quantum Chip	Qubits	Announced	Qubit Archetype	Computing Model
D-Wave XX & 2000Q	~5000	02/2019	Superconducting **flux** qubits	Quantum annealing
	2048	01/2017		
IBM 20Q and 50Q	20	11/2017	Superconducting **transmon** qubits	Quantum circuits
	50	11/2017 (tests)		
Rigetti 19Q	19	12/2017	Superconducting **transmon** qubits	Quantum circuits
Intel Tangle Lake	49	01/2018 (tests)	Superconducting qubits	Quantum circuits
Google Bristlecone	72	03/2018 (tests)	Superconducting **transmon** qubits	Quantum circuits
UC Berkeley QNL	4 (8)	2017	Superconducting **transmon** qubits	Quantum circuits
	64	2022 ?		

Credit: I. Shapoval

Examples:

1. Quantum SVM for Higgs boson searches (U of Wisconsin)
   - Higgs coupling to top quark, Higgs decay to muons. Double Higgs production, Dark Matter
2. Quantum SVM for boosted Higgs (U of Washington)
   - Identify jets originating from Higgs boson decay in order to study Higgs coupling to Standard Model particles
3. Classical optimizations
   - Particle trajectory reconstruction (Intel)
   - Grid workload optimization (Google)

Quantum Flagship in EU and National Quantum Initiative in USA

ID: 162

Example of Collaboration with Industry

CERN openlab

Unique approach to foster collaboration with industry and other sciences, an example with high potential for supporting the R&D and innovation

- **Evaluate** state-of-the-art technologies in a challenging environment and improve them
- **Test** in a research environment today technologies that will be used in many business sectors tomorrow
- **Train** the next generation of engineers & researchers
- **Collaborate** and exchange ideas to create knowledge and innovation

PARTNERS	CONTRIBUTORS	ASSOCIATES	RESEARCH
intel	EXTREME PARTNER NETWORK	COMTRADE	INFN
ORACLE	rackspace	open systems	Fermilab
SIEMENS	E4 COMPUTER ENGINEERING		TU/e
HUAWEI	IBM		King's College London
Micron	be studys		Newcastle University
Google			Simpulse
			IVI Innovation Value Institute
			EMBL-EBI
			European Society of Preventive Medicine

Created at the start of LHC, today includes collaborations with other experiments (i.e. SKA, DUNE) and other sciences

- **Satellite image** analysis for **UNOSAT:/UNITAR**: ML algorithms for information on refugee camps
- **BioDynamo:** a platform for simulations of biological tissue dynamics, such as brain development
- **CERN Livinglab**: new initiative for an open ML platform for large-scale systems biology studies

CERN openlab R&D areas

High-bandwidth fabrics, specialized d processors for data acquisition

Online and offline reconstruction algorithms on GPU, fast simulation data quality monitoring, physics data reduction, large-scale multi-disciplinary platforms, anomaly detection

Data analyses, classical optimizations

Simulation and reconstruction algorithms optimization

Cloud federations, containers, scalability

Software Defined Networks, Security

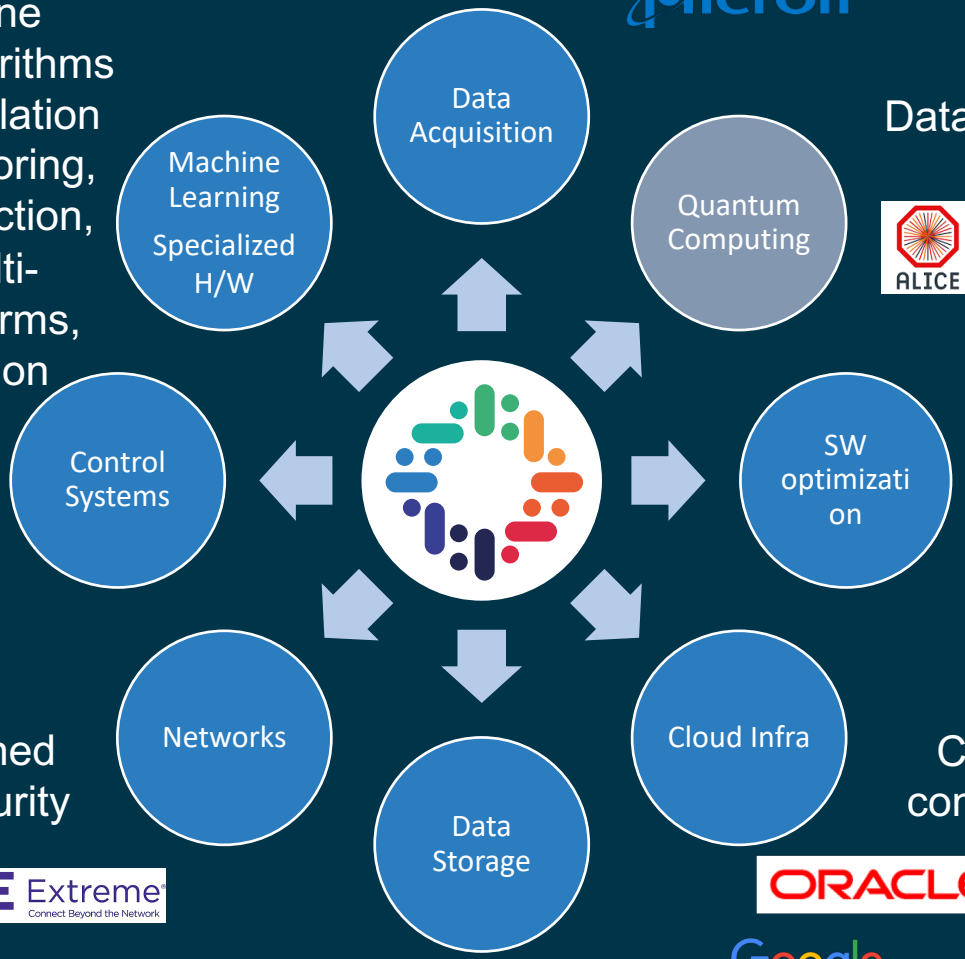

Maria.Girone@cern.ch

Outlook

- Traditional computing improvements are slowing down
- Industry investments are in heterogenous hardware and specialized architectures
  - Driven by AI and big data
- Specialized hardware can significantly change the time to physics and reshape the building blocks of WLCG

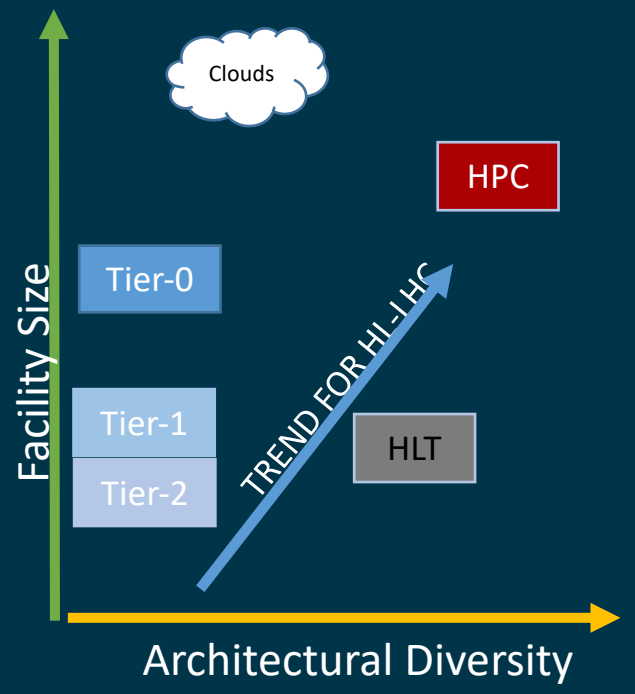

- HEP needs to progress on how to exploit heterogenous hardware
- R&D is needed to understand potential benefits of Super Computing centers
- Data Challenges at the LHC upgrades rate would help validating the scale of the proposed solutions

Open Questions

- How much can we gain by moving to specialized hardware and techniques, and can we afford it?
  - Only R&D will show. Significant investments are needed
- How do we expand the pool of computing resources for HEP to include HPC, commercial clouds, heterogenous specialized facilities and future disruptive technologies?
  - Submissions 59 (INFN Input on S&C), 128 (QC), 162 (Industry collaboration and R&D)
- How do we shift the field toward modern data science tools and techniques?
  - Submissions 5 (European Data Science Institute) and 150 (Tools for Particle Physics)

Many thanks!

J. Blomer, T. Boccali, P, Calafiura, S. Campana, F. Carminati, D. Duellmann, I. Fisk, O. Gutsche, M. Klute, V. Innocente, V. Khristenko, J. Ngadiuba, M. Lassnig, A. di Meglio, F. Pantaleo, M. Pierini, G. Stewart, A. Sciaba, F. Stagni, S. Vallecorsa, J.R. Vlimant

BACKUP SLIDES

Top 10 Supercomputers from top500.org

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,397,824	143,500.0	200,794.9	9,783
2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
5	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	387,872	21,230.0	27,154.3	2,384
6	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , Cray Inc. DOE/NNSA/LANL/SNL United States	979,072	20,158.7	41,461.2	7,578
7	**AI Bridging Cloud Infrastructure (ABCI)** - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , Fujitsu National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	32,576.6	1,649
8	**SuperMUC-NG** - ThinkSystem SD530, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path , Lenovo Leibniz Rechenzentrum Germany	305,856	19,476.6	26,873.9	
9	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x , Cray Inc. DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
10	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom , IBM DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890

CERN openlab

- CVMFS is the most efficient technology on the market for the global distribution of massive application software stacks
  - >1B files under management, >150 production sites
  - Example of LHC software R&D that evolved into a HEP de-facto standard and is disseminated beyond our field (e.g. astrophysics, bio sciences, ...)
- HEP driven development that gives our community a **competitive edge**

- In a future computing landscape with increased resource sharing among scientific communities, we should **invest in integration efforts**, particularly wrt.
  - *Container ecosystems (docker, kubernetes, singularity)*
  - *HPC centers / supercomputers*
  - *non-HEP content-distribution networks*.

Jakob.Blomer@cern.ch

Rucio - Scientific Data Management

- Originally developed by ATLAS, now a growing community of large-scale sciences
  - Adopted by **Xenon**, **AMS**, **CMS**, **DUNE**, **IceCube**, ongoing with **SKA**, **LIGO**, **Belle-II,** and many more
  - Fully open source and **community-driven** development model
- A complete and generic scientific data management service
  - Data can be scientific observations, measurements, objects, events, images, and many more
  - Facilities can be distributed at multiple locations belonging to different administrative domains
  - Designed with more than a decade of operational experience in large-scale data management!
- Smart data in a distributed heterogeneous world
  - Creation, location, transfer, and deletion of data
  - Orchestration according to experiment policies
  - Dynamic optimisation based on usage
  - Monitoring, analytics, reporting
  - Corruption detection and recovery
  - Commercial cloud storage integration
  - Interface with workflow management

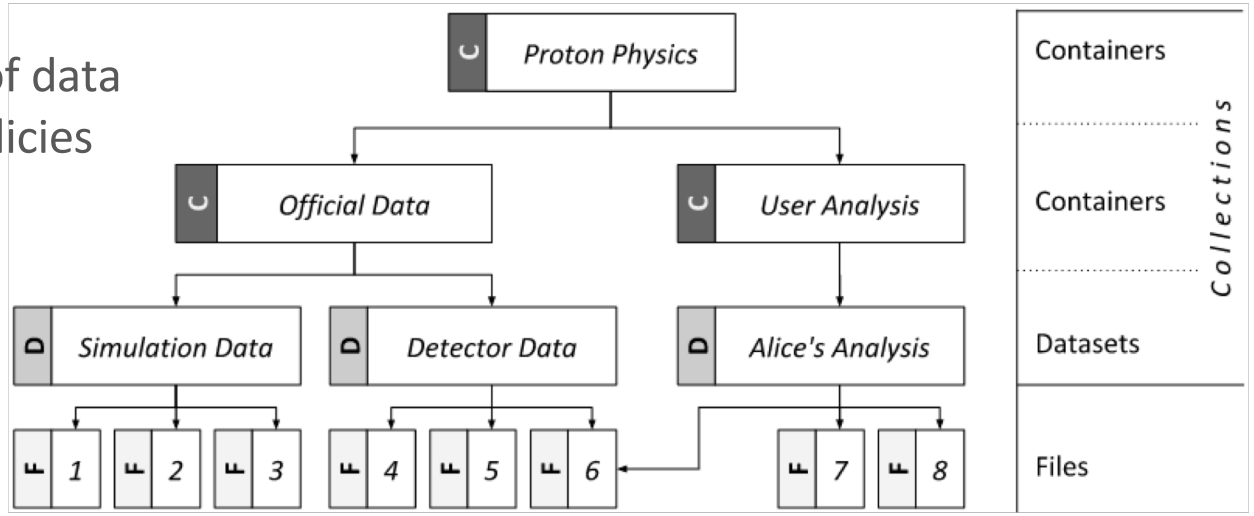

Multi-experiment data management

- Shared use of the global research infrastructures will become the norm, especially with sciences at the scale of HL-LHC, DUNE, and SKA
  - Competing requests on very limited sets of storage and network, data centres will be multi-experiment
  - **Compute** is well-covered, e.g., via scheduling, common interfaces, and requirement languages
  - But **Data** historically has been experiment-specific and needs a solution

- Objective: *Ensure fair use of available data resources across multiple experiments*
  - Allocate **storage** and **network** based on science needs, not statically based on administrative domains
  - Orchestrate data policies within these claims and needs across experiments
  - Dynamically support compute workflows with adaptive data allocations
  - Unify reporting, monitoring and analytics to data centres and administration

Allows more efficient use of the available resources while giving the sciences tangible schedules

Summary of main developments

Development	Production level	Caveats/notes
New WebApp, based on ExtJS6	Now (v6r22)	
OAuth2 support	~Now (v6r22) (with extension)	Will move to use native DIRAC HTTPs, when deployed
DIRACOS	~Now	Can be installed with dirac-install since v6r21, will be default from v7r0
ProductionSystem	v7r0 (2019)	
Pilot3 logging	2019 (?)	
DIPs → HTTPs	2020	• M2Crypto (v7r0) • Patched Tornado
Python 3	2020	DIRAC dependencies
DIRAC→Rucio bridge	? [t b det here]	

Other developments

Development	Production level	Caveats/notes
Gitlab-CI for “matrix style” integration tests	2019	CERN summer student work
JobParameters on ES	v7r0 (2019)	...needs ES (for who wants it)
Multi-VO RSS	v7r1 (?)	
MySQL 8, ES 7	?	
DIRAC for HPC edge nodes	?	
Package+run with Docker+orchestrators	?	ATM not the biggest priority, but mostly missing personpower

Algorithms and Frameworks

Acceleration of algorithms with GPUs

Online: LHCb

- Big changes during LS2
- Event selection based on two software-based HLT levels
- HLT-1 reduces 5TB/s input to 130GB/s:
  - Track reconstruction, muon-id, two-tracks vertex/mass reconstruction
  - GPUs can be used to accelerate the entire HLT-1 from RAW data
- HLT-2 reduces 130GB/s to 10GB/s:
  - Full offline quality reconstruction, alignment and calibration

More info: D. vom Bruch, V. Gligorov, D. Campora Perez

Online & frameworks: CMS - Patatrack

- Demonstrated advantage of heterogeneous reconstruction (on GPU) from RAW to Pixel Vertices at the CMS HLT
  - 1 order of magnitude both in speed-up and energy efficiency wrt full Xeon
  - Running within the CMS software framework
  - Benchmarks executed at Flatiron institute
- Parallelization of more algorithms to run in production during Run 3 and Run 4
- Performance portability
- Definition of a composable farm approach with remote offload

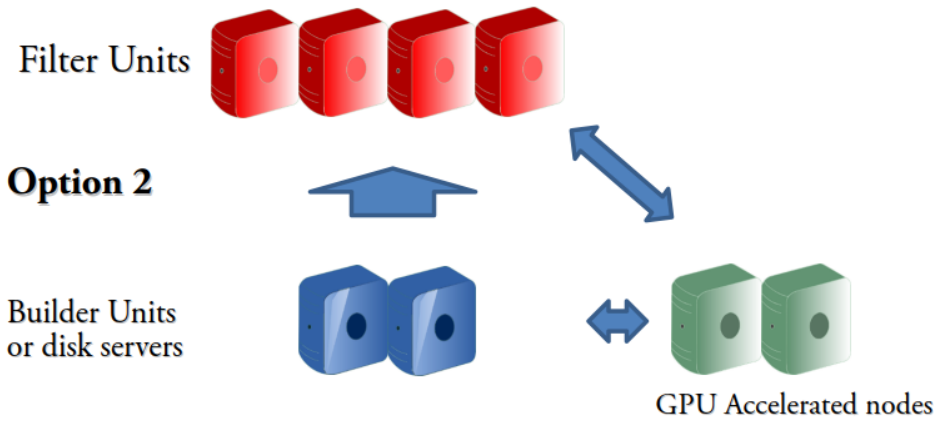

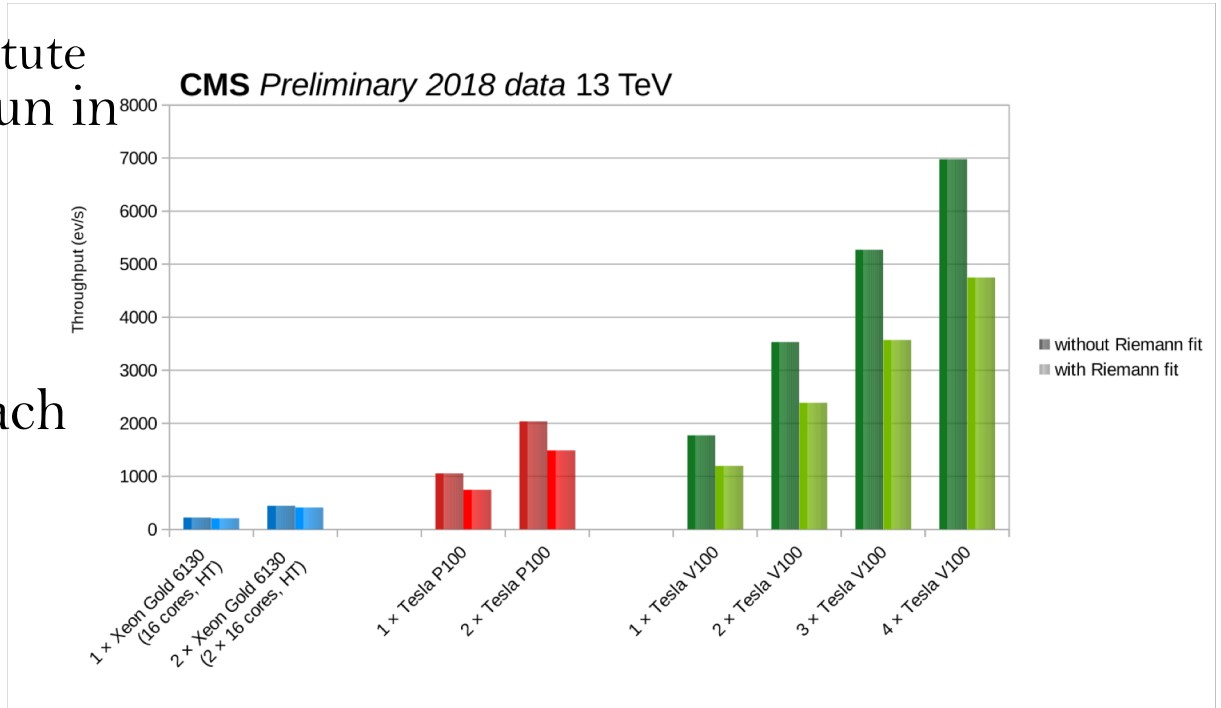

More info: A. Bocci, F. Pantaleo

Ecal/Hcal Reconstruction on GPU

- Heterogenous Execution for CMSSW
  - Concentrating on HCAL / ECAL Local Energy Reconstruction

Current Calorimeters take 20-25% RECO time
And both use the same algorithm -> fast NNLS

Table 2.1: Time spent into the various HLT reconstruction steps

Step	Real-Time	Percentage
ECAL local reconstruction	38.9 ms	8.25%
HCAL local reconstruction	73.9 ms	15.67%
Jets/MET	14 ms	2.97%
E/Gamma	20.4 ms	4.33%
Muons	34.2 ms	7.25%
Pixel tracking	65.7 ms	13.93%
Full tracking	114.2 ms	24.22%
Vertex reconstruction	2.3 ms	0.49%
Particle Flow and Taus	36.8 ms	7.8%
HLT	14.7 ms	3.12%
Overhead	56.4 ms	11.96%
Total	471.5 ms	100%

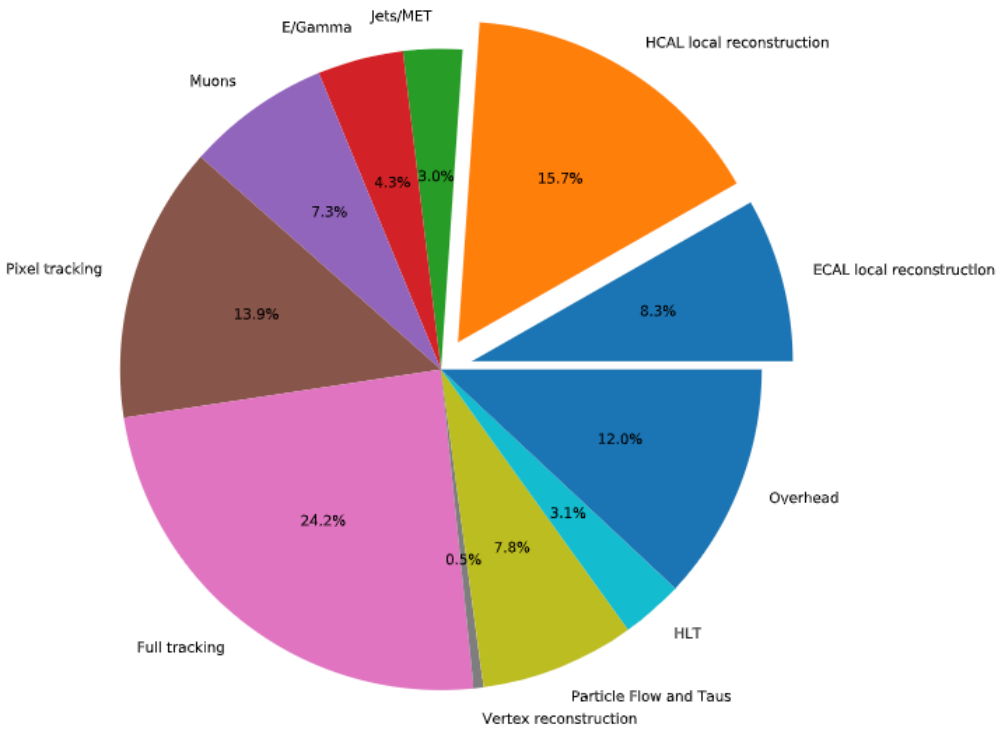

Machine Learning

FPGA-accelerated Deep Learning for HEP

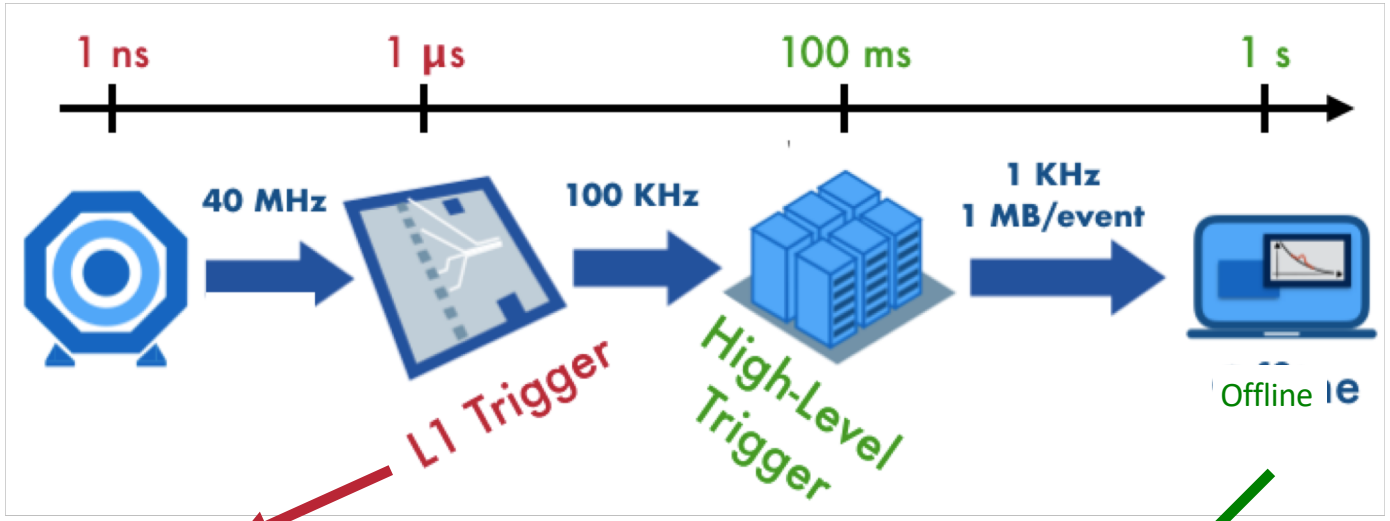

HLS4ML: Fast inference of deep neural networks in FPGAs for particle physics
(JINST, 13, 2018)

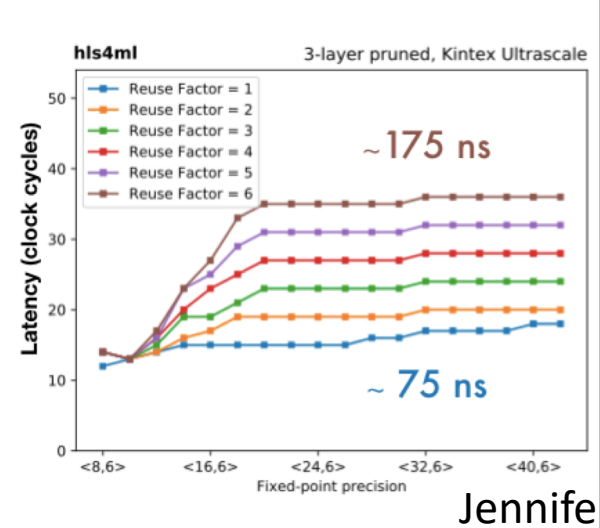

SONIC: FPGA-accelerated machine learning inference as a service for particle physics computing
(arxiv.1904.08986)

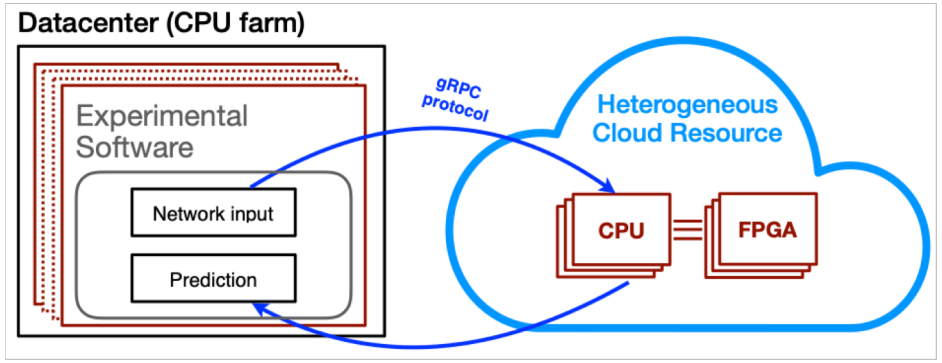

Jennifer.Ngadiuba@cern.ch

Deep Learning for fast detector simulation

Detailed simulation of subatomic particles interactions is essential

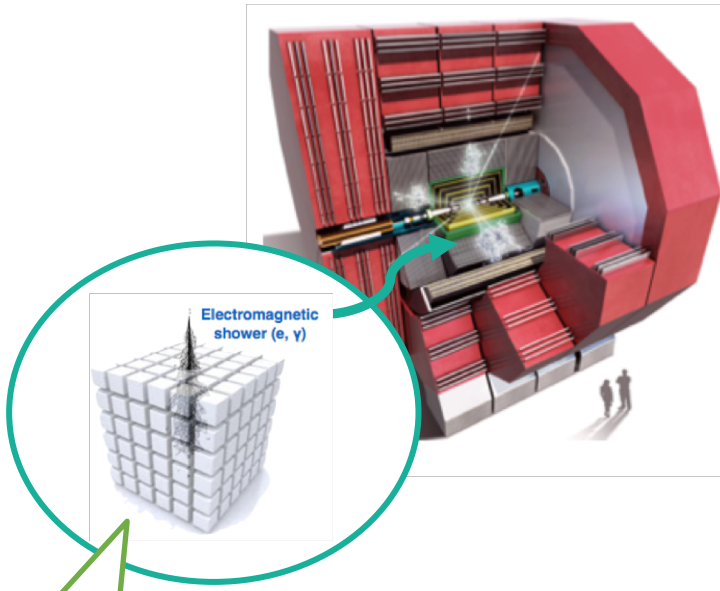

Monte Carlo approach is not fast enough for HL-LHC needs

- 3D convolutional GAN generate realistic detector output
- >2000x faster than Monte Carlo on a Intel Xeon processor

Training takes ~1 day on NVIDIA P100

- Use parallel approach to distribute training across multiple nodes

Detector output as 3D images

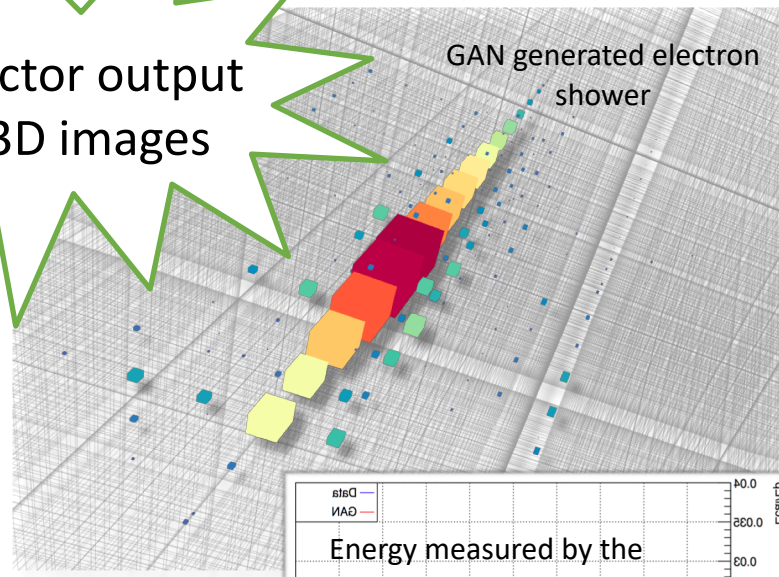

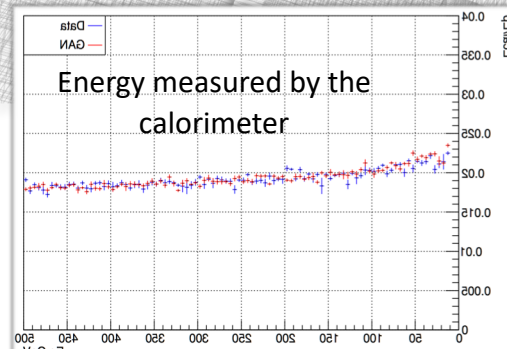

Use Techlab GTX-1080Ti to do basic tests

To generalise we need systems with larger memory and clusters

Sofia.vallecorsa@cern.ch

Distributed training

Different MPI-based data parallel training tests

- Successfully run on different HPC centers (Oakridge – TACC)
- Deployment on cloud environment using docker + kubeflow+kubernetes
  - Openstack
  - Private cloud providers (through HNSciCloud)

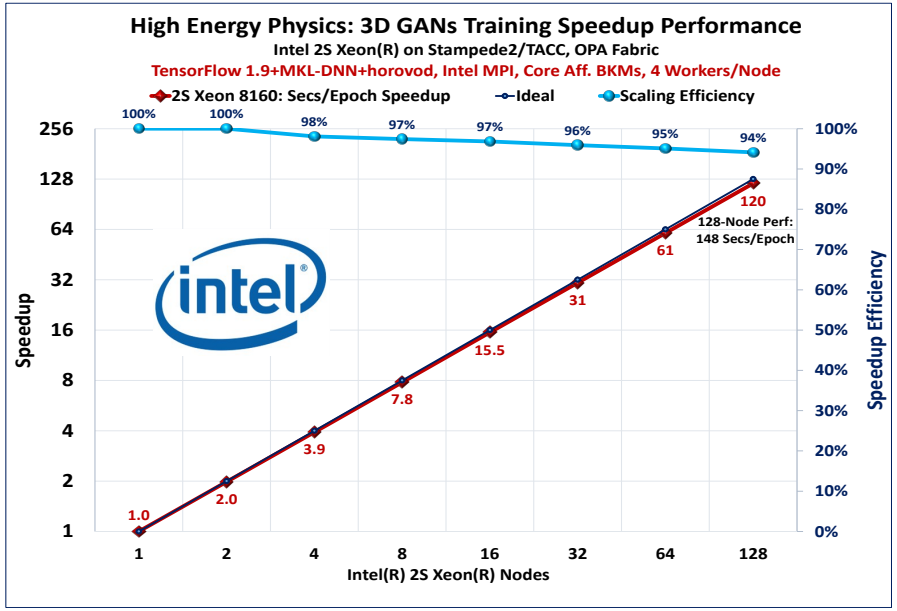

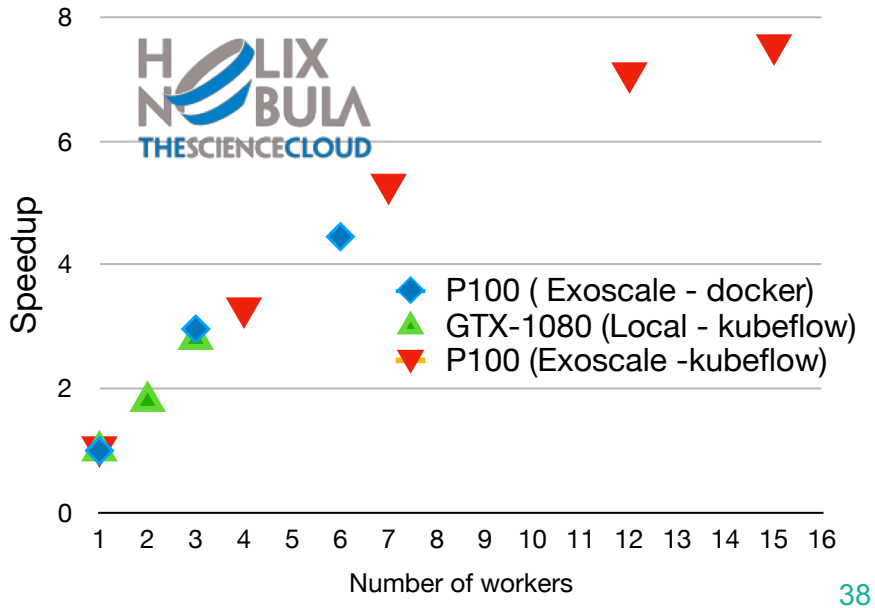

Sofia.vallecorsa@cern.ch

PID at the LHCb RICH detector

- Tracks extrapolated to the RICH surface
- Hit pattern in ring around the extrapolation point is fed to a CNN
- Classification problem
- Classical solution computationally expensive

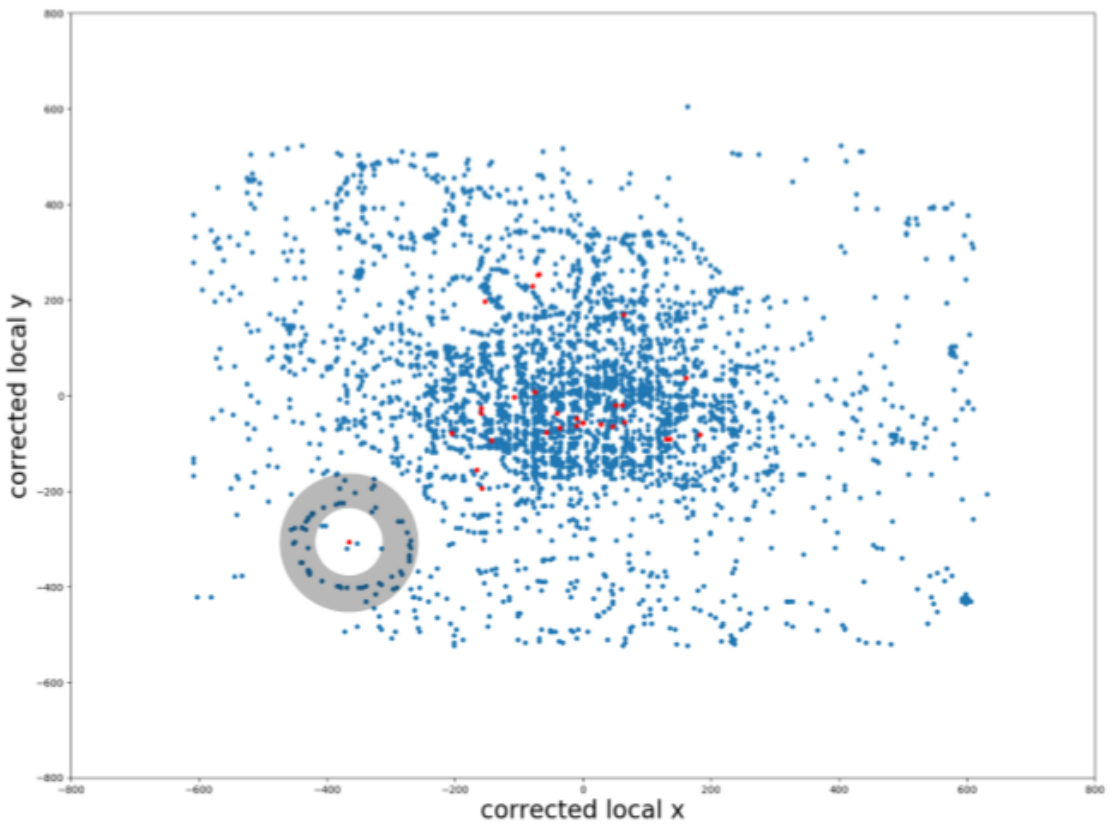

More info: D. Campora Perez

Reconstruction in CMS HGCal

Starting from Run4, High Granularity Calorimeter, a 5-D detector

- usage of unsupervised clustering algorithms
  - optimized to arrange the recorded particle hits into candidate particle showers
- usage of supervised algorithms to accomplish tasks such as particle identification and energy regression on the clustered showers
- Development of de-noising techniques
- Integration of fast inference in CMSSW

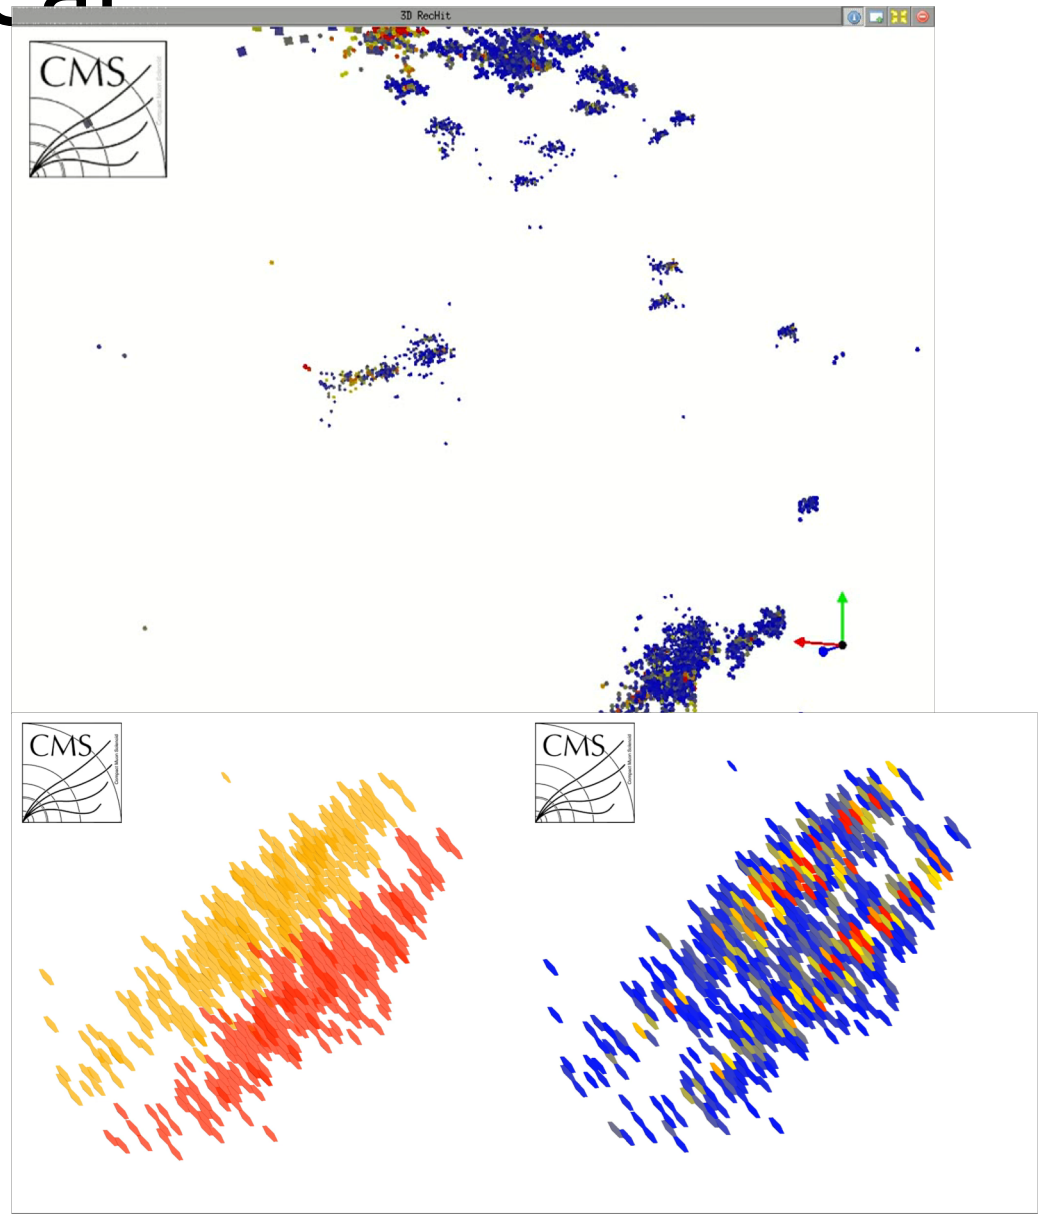

More info: [M. Rovere, V. Innocente, J. Kieseler]

Generative models for HGCal

- Stick to classic calorimeter generation problems
- Use HGCAL as benchmark
- Explore alternative/complementary directions to computing vision
  - Graph Networks for GANs
  - VAEs (with Graph or recurrent models)
- Understand feasibility
- Optimise the model
- Customize solutions

More info: M. Pierini, F. Pantaleo

Distributed Training and Optimization

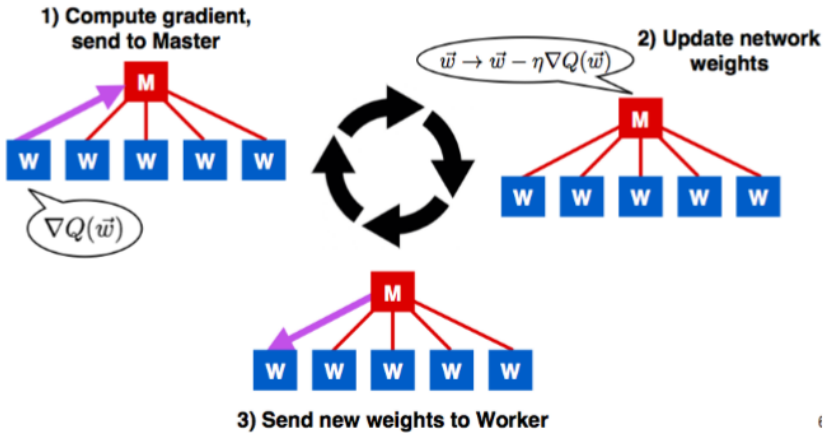

Full parameter scan is resource/time consuming.

Hence looking for a way to reach the optimum hyper-parameter set for a provided figure of merit (e.g. loss)

Utilize all directions of parallelism in model training and model optimization

Enable exploitation of large HPC facilities as multi-node training machines

More info: J-R. Vlimant, M. Pierini

Quantum Computing

Research paths in Quantum Computing

- Tools and methodology development on emulators and simulators
  - Proof-of-concept algorithms for HEP workloads
  - Compare results on real devices
- Understand the role that CERN can play as part of broader QC development initiatives
  - APIs and user interfaces to access QC systems
  - Engineering aspects of QC installation (cryogenics, material science, ..

Examples

1. Quantum SVM for Higgs boson searches (University of Wisconsin)
   Higgs coupling to top quark
   Higgs decay to muons
   Double Higgs production
   Dark Matter
2. Quantum SVM for boosted Higgs (University of Washington)
   Identify jets originating from Higgs boson decay in order to study Higgs coupling to Standard Model particles
3. Classical optimisations
   Particle trajectory reconstruction (Intel)
   Grid workload optimisation (Google)

Quantum Support Vector Machine

Quantum SVM for ttH ($H \rightarrow \gamma\gamma$) classification

More in W. Guan's talk

QSVM is simulated on IBM Qiskit simulator

- different numbers of qubits and events

Entanglement is used to encode relationships between features

Apply PCA to input data features

- Reduced from 45 to 8,10 or 20 (limited by number of qubits)

Running full training with quantum simulators requires large computing resources

- Memory increases with qubit, training events and complexity

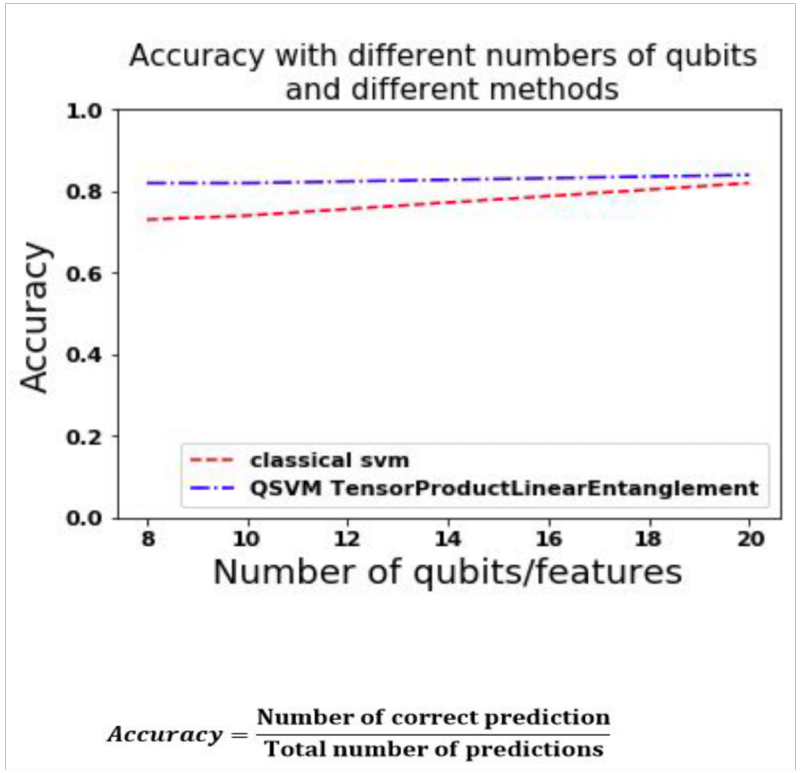

$$Accuracy = \frac{\text{Number of correct prediction}}{\text{Total number of predictions}}$$

Optimisation of classical problems

Track reconstruction in dense environment

Track candidates are identified via combinatorial search, then run through Kalmnan filters

Use QC to speed up initial combinatorial searches

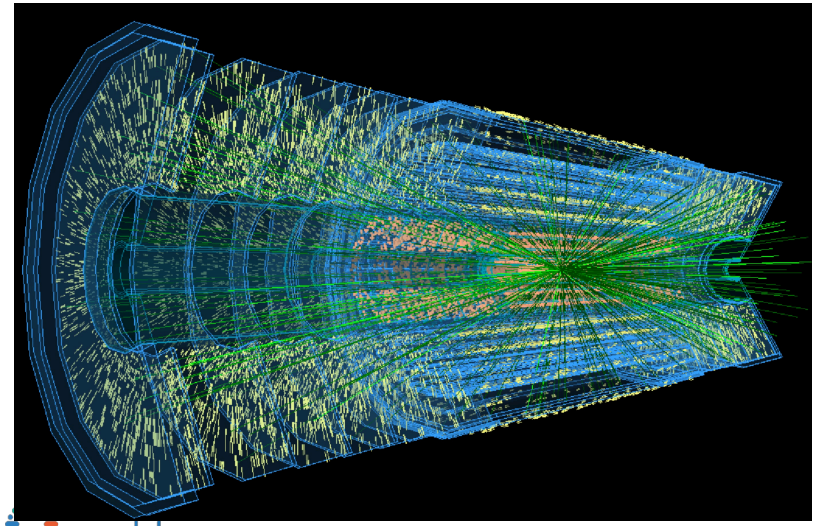

Grid storage optimisation

ALICE Grid runs ~140.000 jobs 24 x 7 x 365

- 70 computing centres in 40 countries
- 150,000 CPU cores and 120 PB of storage

Use QC to optimize storage location in a dynamic environment

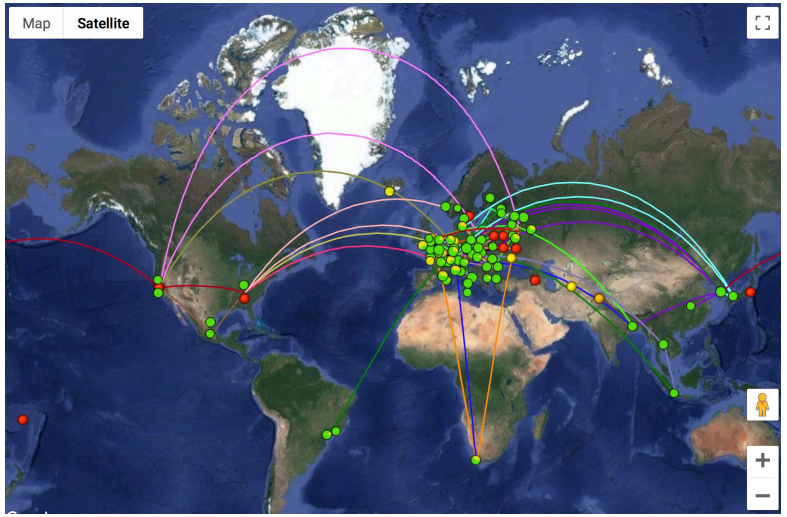

Request

- Taking into account written inputs submitted to the European Strategy Update as well as your own experience, we would welcome your thoughts on what are

- the general trends and opportunities in the computing industry, which areas requires R&D efforts in order to address future computing needs,
- how the required R&D activities can be coordinated;
- what possible synergies/collaborations in computing R&D can exist between HEP collider, accelerator and non-accelerator programs, and also with non-HEP programs and industry