SPECTRUM

Data Processing Needs and Trends in High Energy Physics

CERN-related HEP

David Southwick, Maria Girone, CERN

Funded by the European Union

SPECTRUM is funded by the European Union
- Grant Agreement Number 101131550

HEP at CERN

Overview

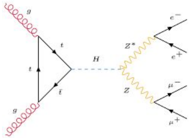

THEORY

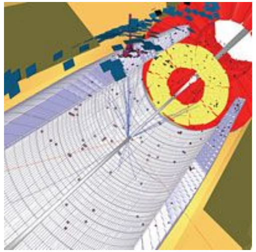

\+

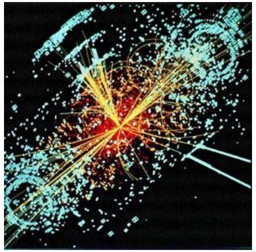

Simulation

Data

RESULTS

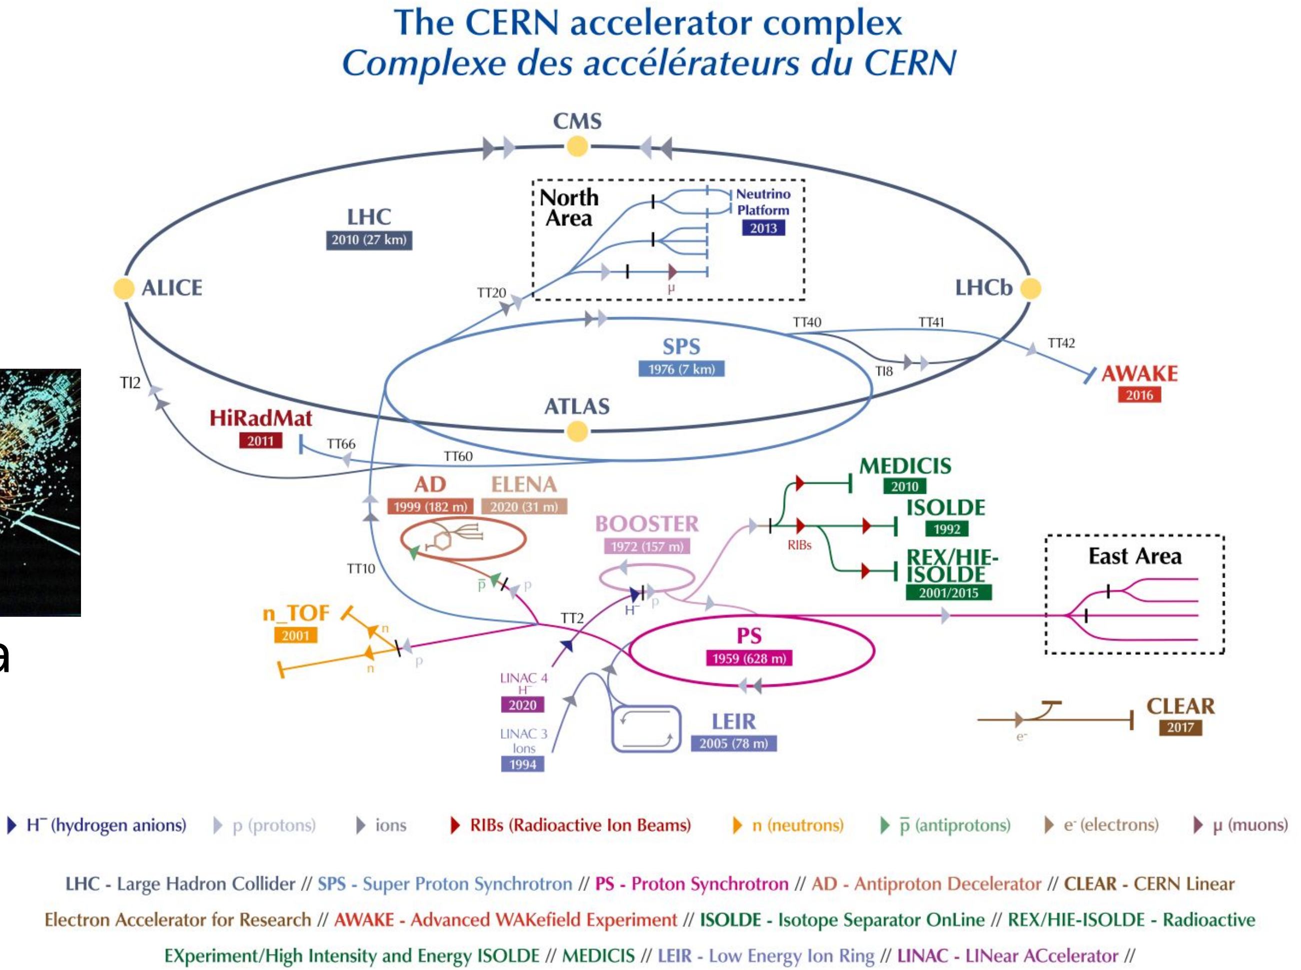

CERN Computing

Historical Computing Trends for LHC Experiments

The Worldwide LHC Computing Grid (WLCG) is the distributed computing grid that provides ~12,000 physicists with ~local access to LHC data:

- Around 1.5 Million CPU cores running 24/7
- 1 Exabyte disk, 1.4 Exabyte tape
- CERN provides ~20% of WLCG resources

WLCG sites provide a common environment:

- Authorization/Accounting
- ~Homogenous Hardware / disk space
- Edge service (CVMFS, etc)
- Network and disk speed policies

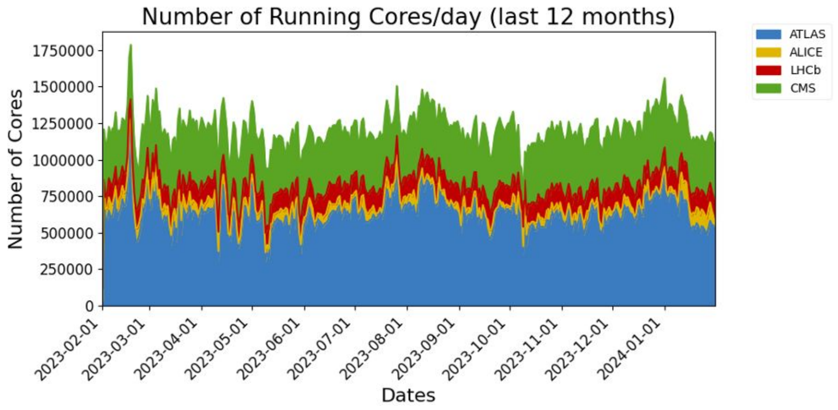

https://wlcg.web.cern.ch/using-wlcg/monitoring-visualisation/monthly-stats

HEP Motivation

LHC Experiments at CERN

LHC expects more than exabyte of new data for each year of HL-LHC era from ~2029-2040.

This data must be exported in ~real time from CERN to compute sites.

CERN is not alone: SKAO expects similar requirements during similar period; other big-data sciences to follow

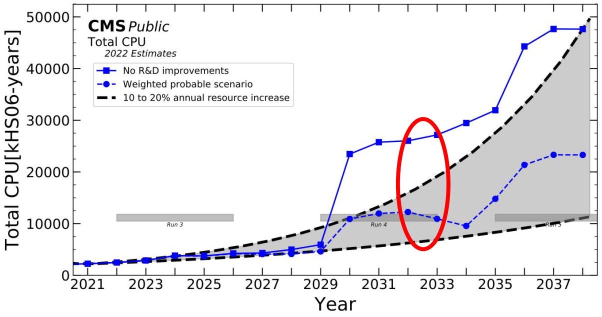

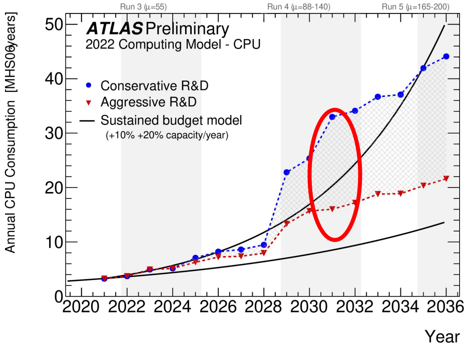

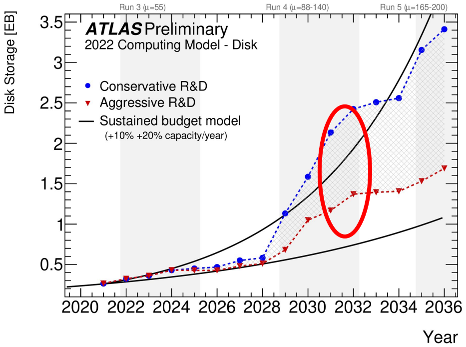

CMS and ATLAS Compute Resource projections

CERN Computing

Computing outside LHC Experiments

Computing physics theories exactly from first principles, without any* assumptions, to the desired precision: Lattice QCD

- has been a major **driver** and **consumer** of global HPC resources for the past 30 years

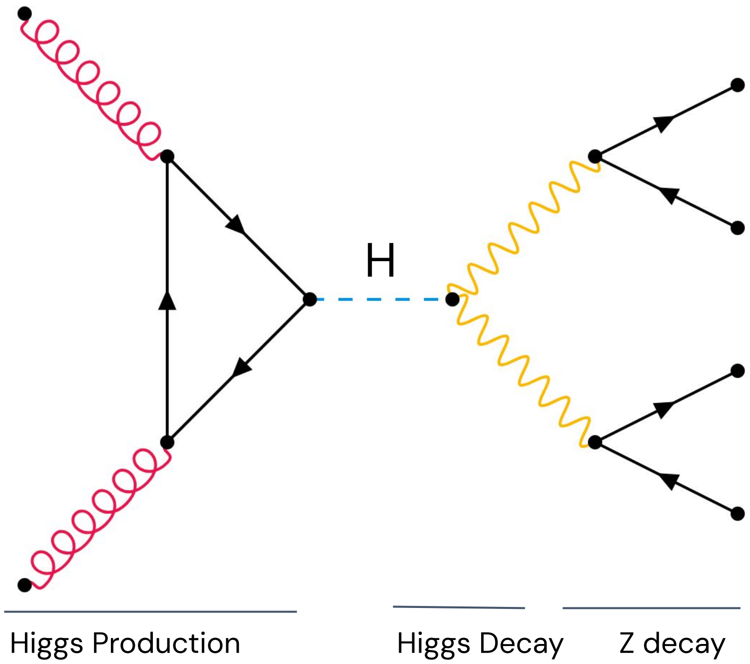

Recent Trends

HEP at CERN

Evolutions in

- Hardware
- Software
- Wetware

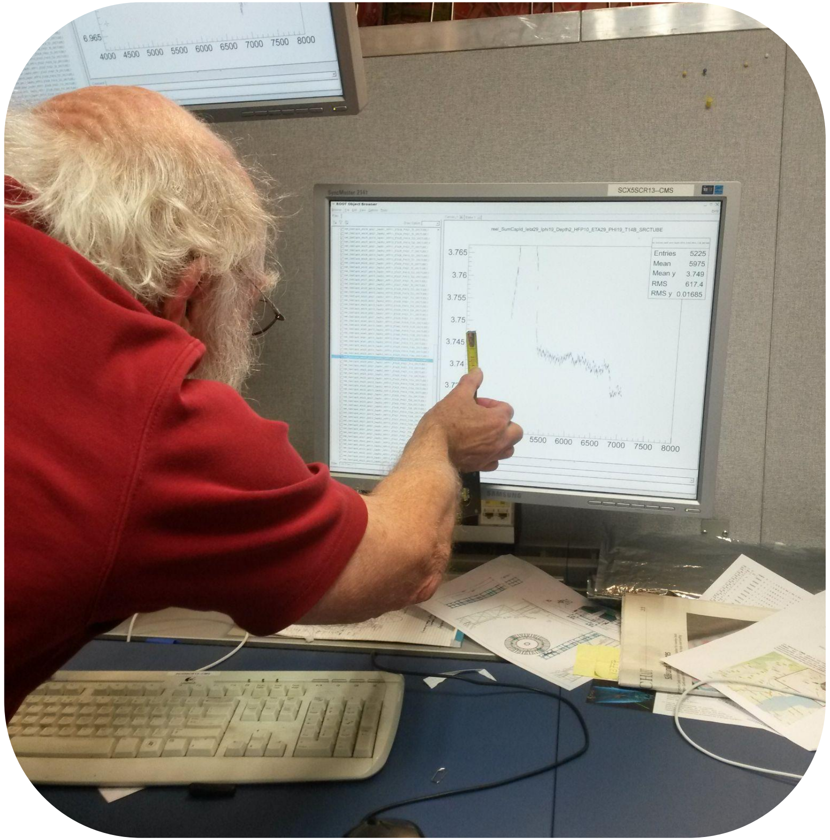

Hardware

Homogeneous to Heterogeneous

Moore’s law continues its iterations of death and rebirth:

- Specialized computing centers
    - Specialized hardware
        - Specialized hardware structures (vectors)
            - *Specialized languages*

Computing today driven by demand for AI hardware

HEP is closely following this trend

Adoption of Machine Learning

General Trends

Today, nearly all experiments at CERN are developing AI/ML to increase efficiency

Industry drove the convergence of AI and HPC with large model development and the need for faster insights to data

HEP (and big-data sciences) have been investing in ML/AI development in diverse areas

Common theme: Need for resources!

Status: AI is here to stay

ATLAS:

- Most simulation is still classical (but **Fast ML based on GAN is in production**)
- **Tagging is fully ML**, tracking classical, trigger mostly classical.
- Analysis is mostly classical or simple ML models
- **Expect 50% of ATLAS algorithms accelerated by GPU-based ML by 2030s**

ALICE:

- Multiple ML workloads with different data, training, deployment patterns
- **So far, smaller scale and simpler models than in ATLAS and CMS**

CMS:

- Multiple ML-based reconstruction **already in production**
- **Advanced use cases, highly customized**
- **Moving toward larger models (transformer - based)**
- **Extensive work** at the level of ML optimisation, frameworks (ML fully integrated in CMSSW),
- **At least 30% of CMS algorithms are ML-based today**

LHCb:

- Main use cases for online operations and trigger
- **Requirements at the analysis level are lower, given the data is simpler and luminosity lower than at ATLAS or CMS**

ATS:

- **Automation of the accelerators infrastructure** is the main scope for ML research
- In addition: **accelerator design** and **AI assistants (LLMs)**

2nd IT Machine Learning Workshop

HPC Opportunities and Challenges

HPC Adoption

Enormous computing resources that are far more heterogeneous than typical Grid sites

- Early adopters of technology, including accelerators
- Advanced low-latency networking
- Green computing becoming important

Complex to migrate from homogenous grid computing:

- Software and architecture adoption (workloads, schedulers, benchmarking, data handling infrastructures…)
- Authorization, Authentication, Accounting
- Networking
- Provisioning (opportunistic vs Pledged resources)

First outlined for HEP in 2020:

Common challenges for HPC integration, M.Girone

Storage

Storing Science

HPC storage is typically built from a common set of commercial building blocks.

Although standard, they are uniquely implemented at each site:

- Variable composition of replications, metadata nodes, interconnect capabilities
- Little to no visibility into capabilities, usage, accounting, etc.

Lots of moving parts! Break down HPC storage in three areas:

1. Data ingress/egress from HPC centre
2. Efficient usage of storage systems on site
3. Dynamic scaling interaction between (1) and (2)

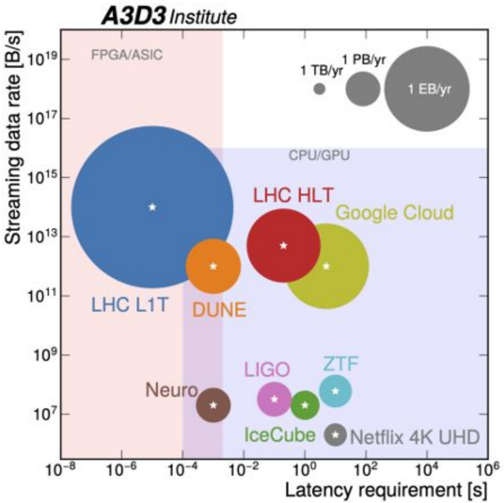

Job Scheduling

Scheduling Science

SLURM scheduler used by HPC sites not immediately compatible with HEP job scheduler (HTcondor)

SLURM – push only, BATCH pull (pilot jobs)

Two ongoing efforts to extend batch schedulers to HPC:

- Extending HTCondor service (tested on connectivity-restricted sites)
- Dask + slurm plugin for submission/translation

Connectivity

Moving Science

CERN current target ~**5Tbps** connectivity by time of HL-LHC from CERN Tier0 to compute sites.

- LHC Network currently being updated to meet this target
- WAN from HPC sites may be limiting factor for resource allocation without pre-placed data

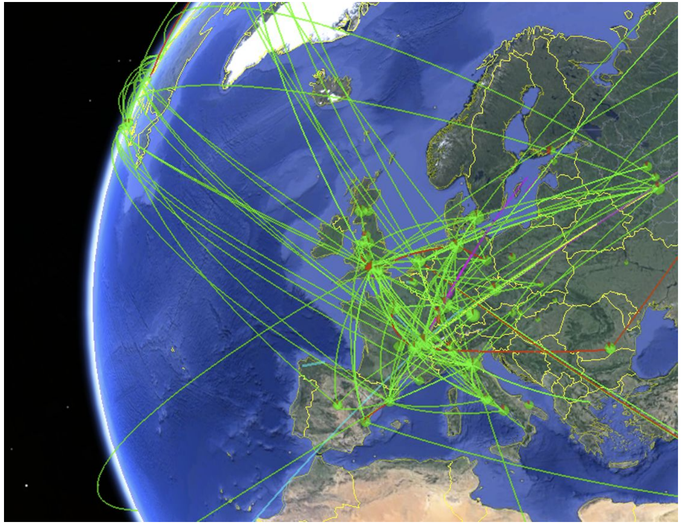

AAI Transformation

Accounting Science

WLCG transition from certificate-based authorization to token-based carries through into HPC

- Among several components of the ESCAPE project, AAI aims to bridge CERN AAI to HPC
- OIDC-token Authentication migration from X.509 Certificate – faster, easier for institutional trust
- Federated login AuthN/AuthZ for HPC via EduGAIN federation/Puhuri
- ESCAPE IAM has been integrated into the EOSC AAI federation in collaboration with GÉANT

Portability Frameworks

D.R.Y. for Heterogeneous Hardware

Advances in wetworks – changing how we approach new hardware

Portable languages for modern heterogeneous hardware

Portable deployment

(via OCI containers)

	CUDA	Kokkos	SYCL	HIP	OpenMP	alpaka	std::par
NVIDIA GPU			*intel/llvm compute-cpp*	*hipcc*	*nvc++ LLVM, Cray GCC, XL*		*nvc++*
AMD GPU			*openSYCL intel/llvm*	*hipcc*	*AOMP LLVM Cray*		
Intel GPU			*oneAPI intel/llvm*	CHIP-SPV: *early prototype*	*Intel OneAPI compiler*	*prototype*	*oneapi::dpl*
x86 CPU			*oneAPI intel/llvm computecpp*	*via HIP-CPU Runtime*	*nvc++ LLVM, CCE, GCC, XL*		
FPGA				*via Xilinx Runtime*	*prototype compilers (OpenArc, Intel, etc.)*	*protytype via SYCL*	

CHEP 2023 https://indico.jlab.org/event/459/contributions/11807

Ramping Up

A complex problem with many moving parts – All feasible methods to close the future computing gap are being pursued

- Including HPC!

Substantial technical investment, both for production(CPU) and development in past years

HEP and Big-Data sciences can leverage potentially large benefits by exploiting HPCs

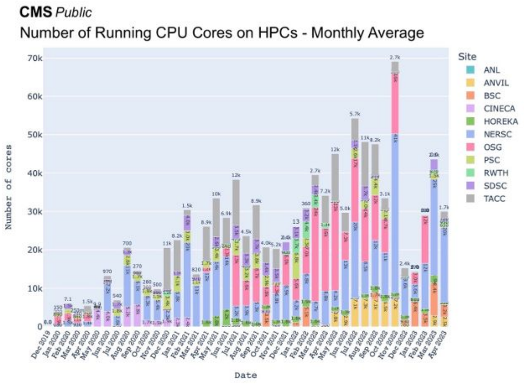

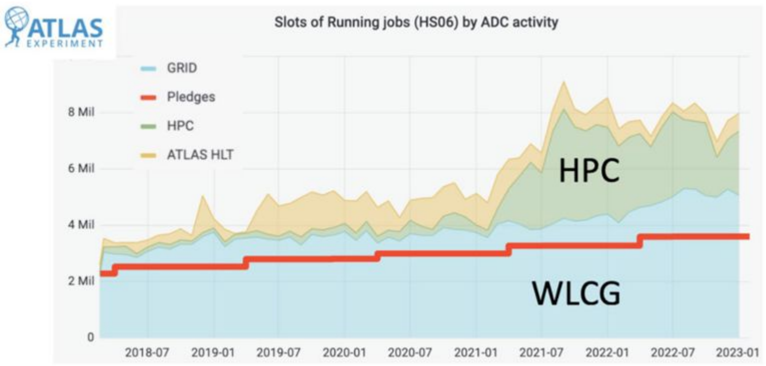

Challenges

Outlook

Much effort has been invested into bridging the future computing gap in the past years, but challenges remain:

- Integrating independent HPC machines as single entities (time/effort intensive)
- No common framework for Access/Usage policies, services, machine-lifetime
  (SPECTRUM will help!)
- Software deployment, edge services for data and workflow management
- Workflow/job orchestration – integration with data locality tracking, HTcondor, etc

  e.g. "opportunistic" Data ingress/egress based on locality, compute resource & time constraints

SPECTRUM

Thank you!
Questions?

spectrumproject.eu

linkedin.com/company/spectrum-project-eu

Funded by
the European Union

SPECTRUM is funded by the European Union
- Grant Agreement Number 101131550

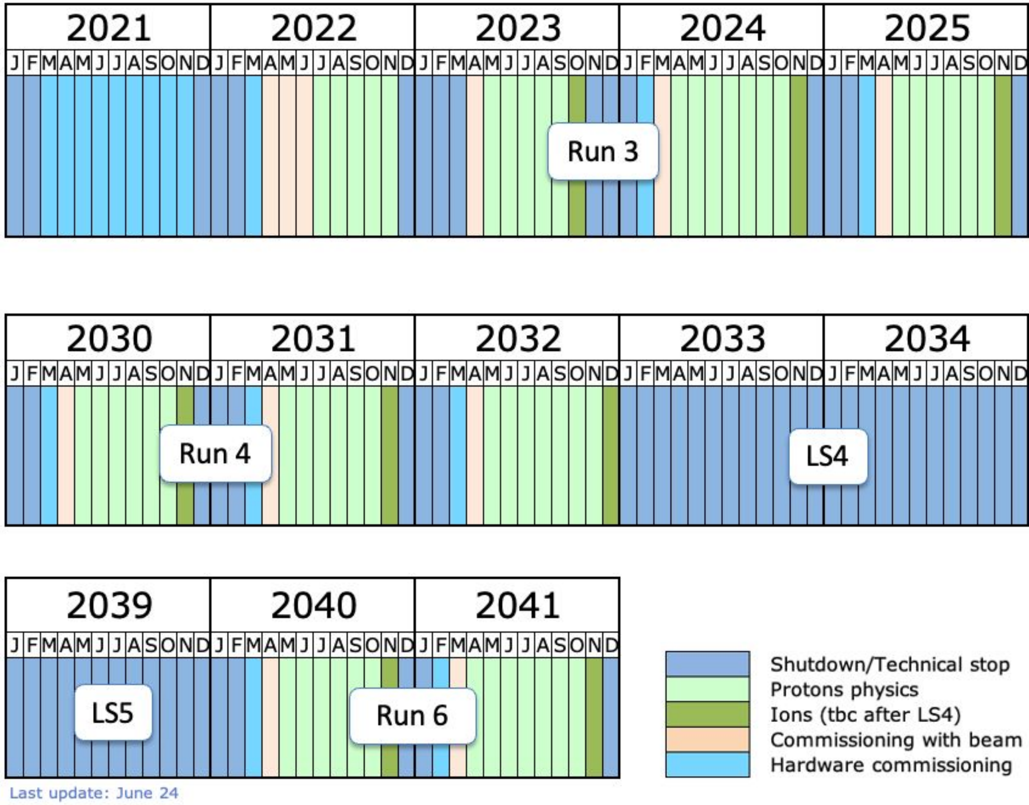

https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm

Quantum Computing in HEP

European Quantum via HPC

HPC essential for quantum computing, massive computing needs for simulation & analysis research

2 quantum simulator sites (100+qubits each) at GENCI(FR), JSC(DE)

6 sites selected to host first European quantum computers

CERN | QUANTUM TECHNOLOGY INITIATIVE