




— IN2P3 —



Computing challenges

École de GIF 2025
Strasbourg

Sabine Crépe-Renaudin




A data-driven science

The tools


3 main ingrédients



Accelerators




Detectors




Computing

Standard model





→ the ability to fully simulate experiments








$$\begin{aligned} \mathcal{L}_{SM} = & -\frac{1}{2}\partial_\nu g_\mu^a \partial_\nu g_\mu^a - g_s f^{abc} \partial_\mu g_\nu^a g_\mu^b g_\nu^c - \frac{1}{4}g_s^2 f^{abc} f^{ade} g_\mu^a g_\nu^b g_\mu^c g_\nu^d - \partial_\nu W_\mu^+ \partial_\nu W_\mu^- - M^2 W_\mu^+ W_\mu^- - \\ & \frac{1}{2}\partial_\nu Z_\mu^0 \partial_\nu Z_\mu^0 - \frac{1}{2}M^2 Z_\mu^0 Z_\mu^0 - \frac{1}{2}\partial_\mu A_\nu \partial_\nu A_\mu - ig_{cw}(\partial_\nu Z_\mu^0(W_\mu^+ W_\nu^- - W_\mu^- W_\nu^+) - Z_\mu^0(W_\mu^+ \partial_\nu W_\nu^- - \\ & W_\mu^- \partial_\nu W_\nu^+) + Z_\mu^0(W_\mu^+ \partial_\nu W_\nu^- - W_\mu^- \partial_\nu W_\nu^+)) - ig_{sw}(\partial_\nu A_\mu(W_\mu^+ W_\nu^- - W_\mu^- W_\nu^+) - A_\nu(W_\mu^+ \partial_\nu W_\nu^- - \\ & W_\mu^- \partial_\nu W_\nu^+) + A_\nu(W_\mu^+ \partial_\nu W_\nu^- - W_\mu^- \partial_\nu W_\nu^+)) - \frac{1}{2}g^2 W_\mu^+ W_\mu^- W_\mu^+ W_\mu^- + \frac{1}{2}g^2 W_\mu^+ W_\mu^- W_\mu^+ W_\mu^- + \\ & g^2 c_w^2 (Z_\mu^0 W_\mu^+ Z_\nu^0 W_\nu^- - Z_\mu^0 Z_\nu^0 W_\mu^+ W_\nu^-) + g^2 s_w^2 (A_\mu W_\mu^+ A_\nu W_\nu^- - A_\mu A_\nu W_\mu^+ W_\nu^-) + g^2 s_w c_w (A_\mu Z_\nu^0 (W_\mu^+ W_\nu^- - \\ & W_\mu^- W_\nu^+) - 2A_\mu Z_\nu^0 W_\mu^+ W_\nu^-) - \frac{1}{2}\partial_\mu H \partial_\mu H - 2M^2 \alpha_h H^2 - \partial_\mu \phi^+ \partial_\mu \phi^- - \frac{1}{2}\partial_\mu \phi^0 \partial_\mu \phi^0 - \\ & \beta_h \left(\frac{2M^2}{\alpha_h^2} + \frac{2M}{\alpha_h} H + \frac{1}{2}(H^2 + \phi^0 \phi^0 + 2\phi^+ \phi^-) \right) + \frac{2M^4}{\alpha_h} - \alpha_h - g\alpha_h M (H^3 + H\phi^0 \phi^0 + 2H\phi^+ \phi^-) - \\ & \frac{1}{8}g^2 \alpha_h (H^4 + (\phi^0)^4 + 4(\phi^+ \phi^-)^2 + 4(\phi^0)^2 \phi^+ \phi^- + 4H^2 \phi^+ \phi^- + 2(\phi^0)^2 H^2) - gM W_\mu^+ W_\mu^- H - \\ & \frac{1}{2}g \frac{M}{c_w} Z_\mu^0 Z_\mu^0 H - \frac{1}{2}ig (W_\mu^+ (\phi^0 \partial_\mu \phi^- - \phi^- \partial_\mu \phi^0) - W_\mu^- (\phi^0 \partial_\mu \phi^+ - \phi^+ \partial_\mu \phi^0)) + \\ & \frac{1}{2}g (W_\mu^+ (H \partial_\mu \phi^- - \phi^- \partial_\mu H) + W_\mu^- (H \partial_\mu \phi^+ - \phi^+ \partial_\mu H)) + \frac{1}{2}g \frac{1}{c_w} (Z_\mu^0 (H \partial_\mu \phi^0 - \phi^0 \partial_\mu H) + \\ & M (\frac{1}{c_w} Z_\mu^0 \partial_\mu \phi^0 + W_\mu^+ \partial_\mu \phi^- + W_\mu^- \partial_\mu \phi^+) - ig \frac{1}{c_w} M Z_\mu^0 (W_\mu^+ \phi^- - W_\mu^- \phi^+) + ig s_w M A_\mu (W_\mu^+ \phi^- - W_\mu^- \phi^+) - \\ & ig \frac{1-2c_w^2}{2c_w} Z_\mu^0 (\phi^+ \partial_\mu \phi^- - \phi^- \partial_\mu \phi^+) + ig s_w A_\mu (\phi^+ \partial_\mu \phi^- - \phi^- \partial_\mu \phi^+) - \frac{1}{2}g^2 W_\mu^+ W_\mu^- (H^2 + (\phi^0)^2 + 2\phi^+ \phi^-) - \\ & \frac{1}{8}g^2 \frac{1}{c_w^2} Z_\mu^0 Z_\mu^0 (H^2 + (\phi^0)^2 + 2(2s_w^2 - 1)^2 \phi^+ \phi^-) - \frac{1}{2}g^2 \frac{s_w^2}{c_w} Z_\mu^0 \phi^0 (W_\mu^+ \phi^- + W_\mu^- \phi^+) - \\ & \frac{1}{2}ig^2 \frac{s_w^2}{c_w} Z_\mu^0 H (W_\mu^+ \phi^- - W_\mu^- \phi^+) + \frac{1}{2}g^2 s_w A_\mu \phi^0 (W_\mu^+ \phi^- + W_\mu^- \phi^+) + \frac{1}{2}ig^2 s_w A_\mu H (W_\mu^+ \phi^- - W_\mu^- \phi^+) - \\ & g^2 \frac{s_w^2}{c_w^2} (2c_w^2 - 1) Z_\mu^0 A_\mu \phi^0 \phi^- - g^2 s_w^2 A_\mu A_\mu \phi^+ \phi^- + \frac{1}{2}ig s_w \lambda_\nu^a (\bar{q}_\nu^a \gamma^\mu q_\nu^a) g_\mu^a - e^\lambda (\gamma^\mu + m_\nu^\lambda) e^\lambda - \bar{\nu}^\lambda (\gamma^\mu + \\ & m_\nu^\lambda) \nu^\lambda - \bar{u}_j^\lambda (\gamma^\mu + m_u^\lambda) u_j^\lambda - \bar{d}_j^\lambda (\gamma^\mu + m_d^\lambda) d_j^\lambda + ig s_w A_\mu (-\bar{e}^\lambda \gamma^\mu e^\lambda) + \frac{2}{3}(\bar{u}_j^\lambda \gamma^\mu u_j^\lambda) - \frac{1}{3}(\bar{d}_j^\lambda \gamma^\mu d_j^\lambda) + \\ & \frac{ig}{4c_w} Z_\mu^0 \{ (\bar{\nu}^\lambda \gamma^\mu (1 + \gamma^5) \nu^\lambda) + (\bar{e}^\lambda \gamma^\mu (4s_w^2 - 1 - \gamma^5) e^\lambda) + (\bar{d}_j^\lambda \gamma^\mu (\frac{2}{3}s_w^2 - 1 - \gamma^5) d_j^\lambda) + (\bar{u}_j^\lambda \gamma^\mu (1 - \frac{2}{3}s_w^2 + \\ & \gamma^5) u_j^\lambda) \} + \frac{ig}{2\sqrt{2}} W_\mu^+ ((\bar{\nu}^\lambda \gamma^\mu (1 + \gamma^5) U^{lep}_{\lambda\kappa} e^\kappa) + (\bar{u}_j^\lambda \gamma^\mu (1 + \gamma^5) C_{\lambda\kappa} d_j^\kappa)) + \\ & \frac{ig}{2\sqrt{2}} W_\mu^- ((\bar{e}^\kappa U^{lep\lambda}_{\kappa\lambda} \gamma^\mu (1 + \gamma^5) \nu^\lambda) + (\bar{d}_j^\kappa C_{\lambda\lambda}^\lambda \gamma^\mu (1 + \gamma^5) u_j^\lambda)) + \\ & \frac{ig}{2M\sqrt{2}} \phi^+ (-m_\nu^\lambda (\bar{\nu}^\lambda U^{lep}_{\lambda\kappa} (1 + \gamma^5) e^\kappa) + m_\nu^\lambda (\bar{\nu}^\lambda U^{lep}_{\lambda\kappa} (1 - \gamma^5) e^\kappa) + m_\nu^\lambda (\bar{u}_j^\lambda U^{lep}_{\lambda\kappa} (1 + \gamma^5) e^\kappa) + \\ & \frac{ig}{2M\sqrt{2}} \phi^- (m_\nu^\lambda (\bar{e}^\lambda U^{lep}_{\lambda\kappa} (1 + \gamma^5) \nu^\kappa) - m_\nu^\lambda (\bar{e}^\lambda U^{lep}_{\lambda\kappa} (1 - \gamma^5) \nu^\kappa) - \frac{g}{2} \frac{m_\nu^\lambda}{M} H (\bar{\nu}^\lambda \nu^\lambda) - \frac{g}{2} \frac{m_\nu^\lambda}{M} H (\bar{e}^\lambda e^\lambda) + \\ & \frac{ig}{2} \frac{m_\nu^\lambda}{M} \phi^0 (\bar{\nu}^\lambda \gamma^5 \nu^\lambda) - \frac{ig}{2} \frac{m_\nu^\lambda}{M} \phi^0 (\bar{e}^\lambda \gamma^5 e^\lambda) - \frac{1}{4} \bar{\nu}_\kappa M_{\lambda\kappa}^R (1 - \gamma_5) \bar{\nu}_\kappa - \frac{1}{4} \bar{\nu}_\kappa M_{\lambda\kappa}^L (1 - \gamma_5) \bar{\nu}_\kappa + \\ & \frac{ig}{2M\sqrt{2}} \phi^+ (-m_d^\lambda (\bar{u}_j^\lambda C_{\lambda\kappa} (1 - \gamma^5) d_j^\kappa) + m_\nu^\lambda (\bar{u}_j^\lambda C_{\lambda\kappa} (1 + \gamma^5) d_j^\kappa) + \\ & \frac{ig}{2M\sqrt{2}} \phi^- (m_d^\lambda (\bar{d}_j^\lambda C_{\lambda\kappa}^\dagger (1 + \gamma^5) u_j^\kappa) - m_\nu^\lambda (\bar{d}_j^\lambda C_{\lambda\kappa}^\dagger (1 - \gamma^5) u_j^\kappa) - \frac{g}{2} \frac{m_\nu^\lambda}{M} H (\bar{u}_j^\lambda u_j^\lambda) - \frac{g}{2} \frac{m_\nu^\lambda}{M} H (\bar{d}_j^\lambda d_j^\lambda) + \\ & \frac{ig}{2} \frac{m_\nu^\lambda}{M} \phi^0 (\bar{u}_j^\lambda \gamma^5 u_j^\lambda) - \frac{ig}{2} \frac{m_\nu^\lambda}{M} \phi^0 (\bar{d}_j^\lambda \gamma^5 d_j^\lambda) \end{aligned}$$

The LHC




- 40 millions collisions per second
- hundreds of particules from collisions

Detectors




~150 millions channels




ATLAS
EXPERIMENT
Candidate Event:
 $pp \rightarrow H(\rightarrow b\bar{b}) + W(\rightarrow \mu\nu)$
Run: 338712 Event: 335908183
2017-10-19 23:31:18 CEST

→ 3 Mo/collision

A flood of data to process



40 millions de
collisions par seconde




150 millions de canaux,
3Mo/collision




40 MHz
~100 To/s

A flood of data to process



40 millions de collisions par seconde



150 millions de canaux,
3Mo/collision




40 MHz
~100 To/s

Temps de fonctionnement :
7 millions de secondes/an




$40 \cdot 10^6 \times 7 \cdot 10^6 \times 3 \cdot 10^6$
 $= 600 \cdot 10^{18}$ octets
 $= 600$ Exa-octets / an /exp

A flood of data to process



40 millions de collisions par seconde



150 millions de canaux,
3Mo/collision



40 MHz
~100 To/s



Niveau 1, hardware
~ 2.5 μ s

Stockage
Reconstruction
Analyse


Système de
déclenchement

Niveau 2, software (60k coeurs CPU)
~ 500ms

3 kHz
6 Go/s




100 kHz
300 Go/s




Raw data


Current data taking period (run 3)



Simulate the data




40 millions de collisions par seconde




150 millions de canaux,
3Mo/collision



40 MHz
~100 To/s



génération




simulation de l'interaction
avec le détecteur



digitisation

Data processing at LHC





Typical for LHC run 1-2

Data processing at LHC

Reconstruction



- from signal in the detectors to particles and their characteristics



Data processing at LHC

Reconstruction

- from signal in the detectors to particles and their characteristics




Data processing at LHC

Analysis

- select collisions, particles of interest, often with a high level of background noise
 - ML used for decades
- measure, compare with theory (simulation)

Software that matches the complexity of the detectors

- Developed by around 100 people
 - 4 million lines of C++ code
 - 1 million lines of Python code
- Constantly evolving: multiprocess => multithreaded




Constantly evolving analysis techniques

Use of AI for decades

- Signal/background noise separation, particle identification: BDT widely used since the 1990s

Use of AI at all levels




- Development of AI to accelerate simulation (digital twins, surrogate models), for selection, reconstruction, accelerator control, triggering and online analysis
- Use of various AI techniques: (BDT), conventional NN, variational autoencoder, Graph NN, DNN, generative adversarial NN
- developments to adapt them to our specific requirements: large data sets, highly accurate simulations
 - collaborations with AI researchers
- implementation on CPU, GPU or FPGA


Important gain thanks to AI

Examples :

- Flavour tagging
 - GN2 state-of-the-art jet flavour tagging algorithm,
 - Transformer-based (<https://arxiv.org/abs/1706.03762>)



<https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2023-01/>



<https://arxiv.org/pdf/2404.01071>

Important gain thanks to AI

Examples :

- Top tagging





Figure 5: ROC curves for all algorithms evaluated on the same test sample, shown as the AUC ensemble median of multiple trainings. More precise numbers as well as uncertainty bands given by the ensemble analysis are given in Tab. 1.

<https://arxiv.org/abs/1902.09914>

The *frontier* of classification

Top tagger comparison: R_{30} = BG rejection for 30% efficiency vs. #parameters



[G. Kasieczka, EuCAIF 2024]

Transformers rule the world

Inductive bias (Lorentz invariance, symmetries,...):

- More parameter efficient
- **BUT** less performant



The bitter lesson vs. heroic domain-specific modeling efforts

Exact symmetries in latent space – hard to learn
Only approximate symmetries in data space

SOS 2024 Tobias Golling



Today usage: data

Total: ~ 3 Exaoctets data stored : disk + tapes



Data challenges pour préparer l'avenir

Transferts





Today usage: computing

Evolution of computing requirements

- HS23 unit = HEPScore23
 - 1 recent core ~10HS23 => 1.5 million cores
- New benchmark deployed to better account for heterogeneous resources

Resources used

- WLCG T0+T1+T2 and 3 pledged or opportunistic (Cloud/HPC/HLT/Volunteer Computing not taken into account unless part of a WLCG site)




Our infrastructure

A distributed infrastructure

WLCG computing grid



- A network of computing and storage centres providing transparent and secure access for thousands of physicists around the world who analyse this data
- ~ 150 computing and storage sites in 42 countries across 4 tiers
 - Tier 0 CERN: initial data backup, initial calibration and processing
 - Tier 1 duplication of raw data (cold storage on tape), reprocessing and all types of processing
 - Tier 2 all types of processing
 - Tier 3 analysis only
- much less hierarchical than at the beginning
- services for managing data, processing, referencing and monitoring



A distributed infrastructure

WLCG computing grid

- A network of computing and storage centres providing transparent and secure access for thousands of physicists around the world who analyse this data
- ~ 150 computing and storage sites in 42 countries across 4 tiers
 - Tier 0 CERN: initial data backup, initial calibration and processing
 - Tier 1 duplication of raw data (cold storage on tape), reprocessing and all types of processing
 - Tier 2 all types of processing
 - Tier 3 analysis only
- much less hierarchical than at the beginning
- services for managing data, processing, referencing and monitoring







Running jobs: 365644
Active CPU cores: 807139
Transfer rate: 21.54 GiB/sec

An essential ingredient: the network

2 private networks: LHCOPN for T1s and LHCONE for most of the grid

- from 10 to 400Gb/s
- LHCONE now open to other experiments: Belle2, Auger, Juno, Xenon, etc.



Middleware, databases and management tools

Des logiciels et bases de données pour orchestrer l'ensemble

- Software and databases to orchestrate the whole system
- Distribute data dynamically: FTS, xrootd, webdav, DDM, [Rucio](#), [DIRAC](#)
- Distribute computing tasks: [Panda](#), [DIRAC](#), ...
- Distribute software: [CVMFS](#) = CernVM File System provides a scalable, reliable software distribution service.
- Databases: data acquisition conditions, detectors, software, datasets, sites, etc.
- Monitoring: sites, computing tasks, storage, transfers, network, etc.
- Ticket systems: GGUS, JIRA



French participation

LCG-France project

- Tier 1: CC-IN2P3
- Tier 2s: (8 IN2P3/CNRS, 1 CEA)
- network : Renater et CC-IN2P3
- France provides 10% to 13% of computing and storage ressources

SITES LCG-FRANCE



Le CC-IN2P3
et quelques
Tier 2



IN2P3 contributes also to

- experiment software
- data reconstruction analysis
- data and job and tasks management
- R&D activities
- ...

CC-IN2P3


Description and missions

- National research infrastructure for IT resources dedicated to our research fields
 - main experiments LHC/HL-LHC T1, LSST, Belle II, CTA, KM3NET, DUNE,...)
- Provides storage (disk+tape) and computing resources with the more appropriate architecture
 - Mainly HTC but increasing part of GPU and small HPC resources included
- Provides IT related services
 - IN2P3 sites connexion in relation with Renater
 - Tools for software developments, set of collaborative tools




CC-IN2P3 Datacenter

- 2 computer rooms: 1700 m² up to 390 racks
- > 1 MHS23 ~ 1400 cores
~ 62 000 tasks
- >100 PB disks + > 200 PB tapes
- 80 staffs
- Users :
 - CPU 104 teams/930 users
 - STO 221 teams/4596 users



CC-IN2P3 access


<https://portail.cc.in2p3.fr/>




<https://doc.cc.in2p3.fr/>



<https://support.cc.in2p3.fr/>



<https://id.cc.in2p3.fr/>



Overall

A notable success

- Needs of the LHC experiments successfully met by the construction of a distributed Exascale federation with resources, services, software and dedicated teams worldwide

Within a flat budget

Collaborations strong enough to go through major crises:

- Covid pandemic: WLCG ran smoothly during the pandemic period
- Ukraine war: Russian resources mostly compensated, Ukrainian sites back online
- and consequences:
 - cope with delay in delivery => Ex: 1 year delay to get network components, other hardware ~x2 wrt to pre-COVID times
 - energy cost increase




Future Challenges

Future challenges: the HL-LHC

You are here !

HL-LHC: a new phase of data

- accelerators and detectors upgrades
- a tenfold increase in the number of collisions
- greater proton concentration => much larger collision pile-ups
- more and more complex data




HL-LHC computing challenge

Data throughput at HL-LHC

- from detectors
 - LHC: 100 Go/s => HL-LHC 1-15 To/s
- Running time : 5Ms/year
 - Data volume: 5-50 EB/year
 - Drastic reduction required at trigger level
 - Expected throughput from CERN to Tiers1 = 4.8Tb/s !

Also an analysis challenge


	LHC run1-2	HL-LHC run
analysis dataset	10 TB	1000 TB
used ressources	laptop	analysis facility



HL-LHC computing needs


Needs multiplied by 10

- Experiments will collect exabytes of data each year
- Significant investment in R&D needed to meet the challenges while staying within budget and energy constraints
- Actually, R&D continuously progressing: the graphs below have already benefited from R&D efforts in previous years



The result of ongoing and continuous efforts

HL-LHC computing resource needs evolution



2015 projections: resource needs = 10x more than budget allows

2022 projections: resource needs compatible with budget (optimistic scenario)

Simone Campana CHEP 2023

FCC estimation

→ see [ESPPU G Ganis talk](#)

Hypothesis


- Sub-detector design : far from being frozen
 - use reference detector for scaling: CLD, FCC version of the CLIC detector
 - Event sizes (leptonic-hadronic): parametrized 1.2-18 kB, full sim 0.16-2.3 MB
- Algorithm design
- Workflows
- Computing technology
 - Processing time (leptonic-hadronic): parametrized 0.5-30 ms, full sim 1.6-23 s
- Statistics needs
 - Baseline assumption: MC $\approx$ expected data sample

FCC Z run estimation

Illustrative Storage Projection for Z Run




4 experiments, 4 equal years at Z



Evolution of WLCG disk resources
+10%/y, +15%/y, +20%/y
(starting point: 500 PB in 2020
≈ ATLAS+CMS + 10%)

LHC **tape** at the end of HL (≈ 5 EB)

LHC **disk** at the end of HL (≈ 5 EB)



[ESPPU G Ganis talk](#)

FCC Storage projection



4 equally demanding experiments

More realistic sequence (15 y operations)

1. 1-2 m 0.25 * nom_lumi run at Z
2. 1-2 m 0.25 * nom_lumi run at WW
3. 3.5 y HZ run
4. 4 y Z run (peak and off-peak)
5. 2 y WW run
6. 5 y ttbar run

1 day Z calibration runs every month
during WW, HZ, ttbar runs

Based on [P Janot et al](#) (and P Janot private communication)





→ Ultimate FCCee computational needs are o(HL-LHC)

[ESPPU G Ganis talk](#)

A changing environment: challenges and opportunities

Hardware costs





Increase of WLCG ressources within flat budget

- WLCG model relies on decrease of CPU, disk and tape cost => need ~15% increase of ressource every year for same budget to fulfil experiments needs
- Works well since beginning of LHC (one exception in 2017 with the outstanding performance of LHC)
- Observed average ressource increase in last 5 years (several countries report the evaluation)
 - CPU: +14%/year, disk: +15%/year, tape media: +20%/year

Start to follow up GPU

- Trend in GigaFLOPS/USD is favourable for video games but flattened for HPC cards price
- Very volatile markets, long procurement times (52+ weeks of delivery time). High demand worldwide (AI/ML/ChatGPT)





Hardware cost

Concerns for the next years

- Much more fluctuations of the market => less predictable prices
 - Increase of demand and crises tend to increase prices (in 2024 CPU price increased wrt 2023) and hardware time delivery (x2 these last years)
- choose to reduce from 15% ressources growth/year to 5% for CPU and 10% for disk in WLCG prevision model

Ex : Disk



Hardware evolution

Computing architecture evolution

- Moore's Law: transistor density still doubles every two years
- Clock speed stalls since 2000 too much power used
- Dennard's scaling: power used by silicon device is independent on the number of transistor but proportional to the transistor area

New processors


- GPU: multi-core servers with co-processors and complex memory configuration
- in 2023 70% of Top500 machine power is from accelerators
- power consumption controlled
- multiple competing infrastructure and different programming langage
- non-x86 CPU architectures share increasing (AMD, ARM)
- more energy efficient than x-86 CPU
- HPC FUGAKU is based on ARM

Consequences

- evolution towards more parallelism
- evolution towards heterogeneous system

Next revolution: Quantum Computing


50 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

<https://github.com/karlrupp/microprocessor-trend-data?tab=readme-ov-file>

Accelerator performance share in Top 500



CPU share



<https://www.digitimes.com/news/a20211007GS400.html&chid=2>

<https://www.top500.org/statistics/overtime/>

Les défis de l'informatique

Sabine Crépe-Renaudin

20/11/2025

43

Heterogenous resources

HPC

- Huge investments of countries in HPC machines, entering exascale area
- Challenges
 - very heterogeneous in hardware and policies
 - mostly GPU now
 - not generally suited for data-intensive processing
 - also a network issue !
 - security policies

Commercial clouds

- Clouds flexible, large ressources available
- Challenges:
 - interfaces
 - networking
 - procurement, economic model and vendor locking
- Cost effectiveness ? potentially interesting for special tasks or peak needs

Different hardware

- CPU, GPU, FPGA, ASIC
 - Vendors: Intel, AMD, ARM, Power, NVIDIA
 - and different programming libraries !
- Need portable code
- Portability libraries with abstraction layer to hide the backend implementation and use their parallelism efficiently




Growing computing needs

In our research fields

- Several experiments have or will have non negligible computing needs in the near future
- Also in astro-particle and cosmology: Vera Rubin/LSST, CTA, KM3NET, Euclid, ET and SKA and multi-messenger approach
- Nuclear physics: dynamic structuration ongoing, EuroLabs, FAIR, EIC, Lattice QCD...

In other research fields and in society


- Importance of data growing everywhere: improved technologies easier to use and computing resources more easily available, developments of AI etc



<https://www.itu.int/itu-d/reports/statistics/2023/10/10/ff23-internet-traffic/>

Taking into account Open Science movement


- Open source software and FAIR (findable, accessible, interoperable, reproducible) data
- Important investments for instance in Europe with the construction of the EOSC (European Open Science Cloud)



Slide credit: Richard Dubois

O'Mullane et al. 2021 (RTR-003.000.00)

Slide credit: M. Graham



Slide credit: M. Bouwhuis HEPiX 2024

26-4-2024

Computing for the KM3NeT research infrastructure - M. Bouwhuis

12

Computing from Nuclear structure, reactions & astrophysics

HPC at the Exascale

Predictive capabilities:

- Structure & reactions of light nuclei
- Nuclear matrix elements
- Neutrino & electron inc.
- Nuclear matter; fission
- Nucleosynthesis, explosions

AI/ML


- UQ, extrapol. in structure, reactions, astrophysics
- Fast, accurate emulators
- Experiment design
- Detect. design, particle id, beam tuning, upgrades
- ML for 'multi-sensu' data
- Training of large models

QC

- Ideal for many-body syst.
- Solve NP problems with known 'classical' issues
- Algorithms & demos: state prep; time evolv.; special density
- Guide QC hardware design

Needs

- Support for refactoring/optimizing codes
- Increased access: Capacity, HPC, many-core (CPUs, tensor proc, untp) computer resources
- Interdisciplinary Collabs: Math, CS, AI/ML, QIS



Taylor, R. et al. Big Data Research Infrastructure Collaboration Toward the SKA (BRICSKA). doi:10.1590/0001-3765202120201027. PMID: 34076206

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT




Figure 6.4: Estimated size of various disk samples in PB for DUNE and CMS at the HL-LHC for comparison. This estimate includes retention policies and multiple copies. The points show actual use in 2021 which was lower than planned due to delays in distributing second copies of samples to remote sites.

Artificial Intelligence


Use of AI since decades in HEP

- signal/background separation, particle identification, multivariate analysis... ex BDT widely used since 90'

Modern AI and computing ressources allow new usage

- generative models, unsupervised classification, low latency inference, transformers
- CPU, GPU, FPGA implementation
- AI assisted code generation
 - AI usage and developments at all stages of computing in our fields


ML Publications in Science



Ben Blaiszik, "2021 AI/ML Publication Statistics and Charts", Zenodo, Sep. 07, 2022, doi: 10.5281/zenodo.7057437.

[CHEP 2023](#)


3



© David Rousseau


Sustainability

Evolution of Greenhouse Gas Emissions



IPCC 2023 Synthesis Report, [Ref. \[7\]](#)

The energy gap Evolution primary energy consumption



<https://ourworldindata.org/grapher/electricity-prod-source-stacked>

The contribution of data-centers to greenhouse-gas emission is sizeable and growing, no exception in our fields of research

A lot of other impacts (rare earth elements...

Paris agreement

- Zero emission around 2100 => -50% by 2030 !
- To get the 50% in 2030 you have equivalently
 - to expand CO₂ free energies by a factor 12
 - to increase energy efficiency by a factor 2
 - to save energy by a factor 2
- will have a mix of these

→ What translation in our computing models and infrastructure ?

→ See for instance HECAP+ initiative

Initiatives and ongoing work

Main R&D subjects

Distributed computing models

- Dynamic bandwidth usage, third-party copying, data lakes, monitoring and network usage priorities
- Integration of cloud or HPC infrastructures
- New analysis platforms, virtual research environments

Artificial intelligence

- To accelerate simulation and reconstruction, for the analysis of increasingly numerous and complex data

Heterogeneous and specialised architectures and software

- x86, GPUs, FPGAs, APUs, XPU, Superchips, etc.
- Online data processing (real-time analysis)
- Porting software to different architectures
- Investigating library portability and unified programming models
- Software efficiency and energy efficiency
- December workshop: <https://indico.cern.ch/event/1526077/overview>

Reducing environmental impact

- New dedicated working group + workshop <https://indico.cern.ch/event/1450885/timetable/>
- Measuring overall impact and defining a 'physical per unit of energy consumption' metric
- See work on software, new technologies
- Flows based on the type of energy used, etc.

New disruptive technologies

- Quantum computing: QML, reconstruction

A collaborative effort beyond HL-LHC

Enlarge collaborations


- Common developments allow to share expertise, to share effort and find common solution when possible
 - ease deployments and allow shared resources and infrastructure
- Eased by international collaborations and programs
 - HEP Software Foundation (HSF), IRIS-HEP for software developments
 - WLCG/DOMA for WLCG infrastructure expands beyond LHC with DUNE, Belle-2, JUNO and VIRGO as WLCG observers
 - European programs for the development of the European Open Science Cloud (EOSC) => ESCAPE project for HEP, astronomy and nuclear physics, EVERSE project for software, Spectrum for HPC
 - [JENA computing workshop](#) to discuss synergies across the 3 communities



DOMA



EVERSE





SOFTWARE

Event Generation

→ will represent 10-20% of HL-LHC Computing


R&D in HSF Physics Event Generator WG

- Main Objectives [Ref]
 - Accounting and profiling to understand CPU costs
 - Move to GPUs and vectorised code
 - Optimise phase space sampling and integration algorithms, including AI use
 - see [MadJax](#) (differentiable per-event MEs)
 - Research on reducing the cost associated with negative weight events
 - see New Sherpa configuration by ATLAS that achieves 50% reduction of negative weights [Ref]
 - Promote collaboration, training, funding and career opportunities


Example: MadGraph 5

- Matrix Element calculation in event generator can be efficiently parallelised using vector CPU or GPUs
 - Vector CPU: Speed ups of 6x for AVX512
 - GPU: 20x to 80x depending on the process using Tesla A100
- reengineering of MG5aMC is a functional alpha release for LO QCD / EM processes, but weak interactions need more work

ATLAS Preliminary
2022 Computing Model - CPU: 2031, Conservative R&D
Tot: 33.8 MHS06*y




ATLAS Preliminary
2022 Computing Model - CPU: 2031, Aggressive R&D
Tot: 16.6 MHS06*y



Improving the Madevent Side

- Improved handling of MLM
- GPU-assisted unweighting
 - Use GPU to for parallel weight computation
 - Helps FORTRAN unweighting routine to discard events faster
- More investigation on madevent side possible



[MadGraph CHEP 2023](#)

Simulation

Simulation impact

- Simulations are dominating Run 3 CPU usage: ALICE 50%, LHCb 90%, ATLAS 5
- Lot of R&D towards new technique and use of accelerators
- Need
 - to improve accurate full simulation for detector performance and training of last simulati
 - to make more use of fast simulation

Fast simulation


- originally based on parameterisations of detector response, (generative) ML techniques now widely used
 - ATLAS AtlFast3, CMS FastSim and FlashSim
 - Porting fast simulations on heterogeneous resources, exploiting portability models Cf FastCaloSim

Full simulation




- Continuous progress to speed up Full Simulation:
 - same accuracy in the physics description, but faster
 - recent 20% improvement for ATLAS simulation (Woodcock Tracking of gamma – avoid to stop at volume boundaries)
- R&D with GPU prototypes: EM shower simulation on GPU
 - [AdePT](#) (Accelerated demonstrator of electromagnetic Particle Transport), [Celeritas](#)

Full vs fast Simulation




FlashSim

FlashSim means skipping all intermediate steps




Adept




Reconstruction

R&D on track reconstruction with high pile-up


- non linear increase of needed computing ressource with event complexity for classical reconstruction algorithms




CKF chain




GNN chain




6



2010 $\langle \mu \rangle \sim 5$



2018 $\langle \mu \rangle \sim 40$




2029 $\langle \mu \rangle \sim 200$

ACTS project

- ACTS open source project: toolbox for charged particle reconstruction for HEP/NP
- R&D on method and algorithm, 2 main chains developed: Kalman Filter or Graph Neural Net (GNN)
- R&D on parallelisation and performance portability
- On track to be used by ATLAS for track reconstruction by the LHC's Run-4

Event throughput for ttbar



GPUs become competitive at high pile-up. Highest performance observed on **NVIDIA[®] CHEP 2023** workstation GPUs so far.

Analysis


Analysis challenge

- Typical dataset size 10 TB (run 1-2) => 1000 TB (run 4)
- Current workflows do not scale:
 - disk/tape are limited/expensive and analysis data are projected to occupy ~30% (ATLAS model)
 - will need to go from Laptop to Analysis Facility


Analysis R&D

- Analysis workflow evolution: Columnar Analysis going from event loop to array programming and more use of Python ecosystem and industrial standards
 - Coffea: python ecosystem -> [ScikitHEP](#), [Awkward array](#), [DASK](#)
 - array programming solution for quick insights delivery
 - Developments of more compact formats and reducing data copies and intermediate formats
 - NanoAOD and ATLAS [PHYSLITE](#)
 - target ~10kB/event for fast and lightweight analysis
 - ROOT's Tree -> RNTuple, RDataFrame high level API for data analysis
 - encodings and I/O upgrade of the event data file format and access
 - Access to novel and future storage technologies: native support for HPC and cloud, object stores, direct disk-to-GPU data transfer
 - Collaboration with experiment for Event Data Model integration and optimisation
- 10-20% smaller files, x3 to x5 better single core performance
- enables fast adaptation to modern technologies, like object stores
- 2024 goal: release RNTuple 1.0 supporting the EDM of all experiments


Columnar Analysis



Python ecosystem




File size



[CHEP 2023](#)

RDataFrame



[WLCG-HSF2024](#)

Analysis facility

→ Several initiatives

IRIS-HEP Analysis Grand Challenge [Link](#)

- binned analysis, reinterpretation and end-to-end optimisation of physics analysis use cases
- includes the development of the required cyber infrastructure

Analysis Facility prototype at CERN [Link](#)


- both batch and interactive
- DASK, HT-Condor integration
- open for user tests (documentation)

Virtual Research Environment [Link](#)


- development in EU-funded project in the line of EOSC building: ESCAPE
- open source platform with access to all the digital content needed to develop, share and reproduce an end-to-end scientific result in compliance with FAIR (findable, accessible, interoperable, reproducible) principles
- All the building blocks from storage to computing with AAI and notebook to analysis preservation
- Reana: reproducible research analysis platform [Link](#)
 - developed at CERN
 - allow to preserve, distribute and reproduce
 - flexible scalable and reusable => container

And more like INFN Napoli initiative [Link](#)


IRIS-HEP Analysis Grand Challenge




CERN facility software component



ESCAPE VRE Analysis workflow



HEP software : language



Adapted from J. Pivarski

			
Metric	C++	Python	Julia
Performance	✓		✓
Ease-of-use		✓	✓
Learning Curve		✓	✓
Safety (memory)		✓	✓

© E Rodrigues ESSPU 2025


Evolution of usage

- Past essentially fortran, now C++ & Python (Python increasingly used in user analysis)
- Julia sneaking into the arena, especially in smaller experiments (high performance + productivity)
- Present & future: increasing mix of languages

Key4HEP

What HEP software for the future ?

- we don't know
- but we already need software to build the next machine
 - HEP community has agreed on one common software stack : ESPPU contribution #240
 - "members of the ILC, CLIC, FCC, CEPC and the Muon Collider communities ... decided to pool their resources to develop a common and shared software stack ... Key4hep project which has become the de facto standard for studying detector concepts in any currently foreseen future collider
 - <https://key4hep.github.io/key4hep-doc/main/index.html>
 - [CHEP2020](#)




Heterogenous computing

GPU usage already in run 3 (I)

ALICE continuous readout

- Heavy ion data taking rate increased from 10 to 50 kHz at run 3
- With the high TPC latency, decision to operate in continuous readout mode **without trigger**.
 - Time-frame (TF) of 2.5-20ms instead of event acquisition




O2 facility

- First Level Processing (FLP): Readout + FPGA corrections to compress data and build the time-frames
- Event Processing Node (EPN): 8 GPUs AMD MI50 – 2 CPUs 32 core AMD Rome 7452
- online synchronous reconstruction and data reduction when beam on
- asynchronous reconstruction when beam off
 - Offloading 60% of the processing to the GPU induces a speedup ~2.5x, expect 5x with 80%
 - 2023 first use of EPN farm: 1/3 of data processed on EPN (CPU+GPU) and 2/3 on GRID (CPU)

Performance

- 55 CPU 3.3GHz Cores equivalent to a single AMD-Mi50 GPU

ALICE in RUN 3: THE O² PROJECT





[WLCG-HSF 2023](#)

[CHEP 2023](#)

GPU usage already in run 3 (II)

Full Software trigger at LHCb

- For Run 3, LHCb is getting $5 \times$ more collisions per second => trigger becomes the bottleneck
- Goal evolved from rejecting (trivial) background to categorise different 'signals'
- => Drop the L0 trigger, reconstruct 30 MHz of events before making trigger decision




RTA: Real-Time Analysis

- FPGAs-based clustering for Silicon Pixel detector
- HLT1: GPU based reconstruction
 - 163 RTX A5000s (one per Event Builder node)
 - software HLT1 sequence has been implemented in CUDA
- HLT2: CPU-based full reconstruction

Successfully commissioned

- In 2022: 40 Tbits/s network throughput achieved
 - ➔ equivalent to 4% of the internet
- Achieving 30 MHz with less than 200 GPUs !





GPU usage already in run 3 (III)

CMS heterogenous HLT farm

- HL-LHC will require ~one order of magnitude faster trigger
 - L1T 100kHz -> 750kHz (7.5x), pile-up 60 -> 200 (~3x) + more complex detectors
- => include GPUs at the HLT already for run 3
 - 200 nodes 2CPU+2GPU = ~26k CPU cores AMD Milan 7633 + 400 NVIDIA T4 GPUs
- HCAL, ECAL, pixel local reconstruction and pixel tracking on GPU => a lot of work needed to optimise software on these heterogenous architectures
- Allow scouting strategy = increasing event rate (lower thresholds) and decreasing event size to perform analysis directly with objects reconstructed at the HLT (no RAW data is stored)

R&D in Performance Portability

- Porting Heterogenous Code to Alpaka (performance portability library) – to reduce dependency on hardware; kokkos and sycl also studied
- Aiming to offload 10% of (Run-3 and Phase-2) offline reconstruction by end of 2023



- Event throughput +80%
- Average time per event -40%
- Power consumption: -30%
- Experience gained for offline computing

[WLCG-HSF 2023](#)

[CMS CHEP 2023](#)

Data Challenges


Objective

- progressively test the infrastructure by increasing data rate on top of production
- allows to test concurrently in real condition on a large set of sites end-to-end transfer capacities:
 - network, storage, middleware
- allows to test new technologies


Campaigns

- DC21: target 10% of HL-LHC needs
 - 1Tb/s reached ! and lessons learned
 - [2021 Data Challenges Wrap Up](#)
- DC24: target 25% of HL-LHC needs
 - 1,2 Tbps **minimal** scenario, 2,4 Tbps **flexible** scenario
 - including the 4 LHC experiments + Belle2 + DUNE
 - New monitoring following the DC21 post mortem, tokens usage tests
 - Token AAI tested
- DC27 50% of HL-LHC


DC21 WLCG data transfers (Gbps) – 15 days



DC24 WLCG data transfers (Gbps) – 15 days




Site network monitoring



Shawn McKee / U Michigan, Marian Babik / CERN
 Data Utilization / U Michigan, Ilija Vukotic / U Chicago
 Site Network Utilisation - computed from aggregated utilisation (SNMP counters) provided by sites via simple API. Screenshot shows network utilisation during DC24 as seen by the sites.
 2024
 Spring 2024 HEPiX, Paris, France
<https://indico.cern.ch/event/1377701/timetable/#20240417>

Token utilisation in transfers



[Introduction to token](#)

[IAM datachallenge WLCG-HSF 2024](#)



Storage and data management

Difficult challenge

- No opportunistic storage and storage hard to operate => storage needs are the hardest resource to fulfil
- Constant work to optimise data distribution [ATLAS CHEP 2023](#)
 - lifetime model, regular/automatic cleaning of unused or superseded data
- Early R&D and test through DOMA collaboration and ESCAPE EU project and experiment

Better usage of tape: data carousel

- tape driven workflow to allow job to get data directly from tape
=> disk space saving
- close collaboration between sites, experiments workflow and data management system
- in production in ATLAS with peak rate 20-25 GB/s
- develop now derived reconstruction data on demand
- Tape smart writing: find intelligent algorithms for file placement on tape



[CHEP 2023](#)

Storage and data management


Data Lake

- Put data in a distributed federated storage that minimise replication, assure availability, with different storage quality level depending on the needs
- Serve data to remote (or local) compute ressources => grid, cloud, HPC, ..
- Simple caching is all that is needed at compute site (or none, if fast network)


ESCAPE Data Lake

- Based on CERN File Transfer Service (FTS), Rucio data management and orchestration open-source tools
- Security assured through token based AAI
- HEP development benefited to astronomy and nuclear physics collaborations, now collaborative developments within ESCAPE new collaboration (CTA, SKA, LSST, KM3NET, LOFAR, FAIR, HL-LHC etc)
- Part of EOSC-Future Virtual Research Environment, successfully exploited with ESCAPE Dark Matter Science Project CHEP 2023 CHEP 2023

→ See also Data lake implementation in the Nordic countries CHEP 2023



Rucio instance



CHEP 2023


HPC usage

Increased use of HPC ressources

- Mostly for MC simulation, some MC reconstruction
- Exploiting HPCs could be
 - relatively easy for those which have policies similar to HEP sites ones
 - very complicated In other cases: workload management systems had to be creatively adapted in order to use these resources


Challenges

- low network connectivity on the worker nodes
- security rules prevent access to standard software repositories (CVMFS)
- Ingress/egress requirements
- ressources booking



CMS Public

Number of Running CPU Cores on HPCs - Monthly Average



Cloud ressources

→ Clouds in WLCG sites have been used for a long time


Commercial Cloud usage

- Clouds have a lot of different ressources and allows for flexibility, on demand ressources
- Some challenges: interface, network, provider specificities, vendor locking
- Also procurement and economic model
- Cost effectiveness to be demonstrated


ATLAS Google project

- Full integration of Google into ADC activities and infrastructure: important work to adapt data and job management tools (Panda, Rucio)
 - PanDA & Rucio developments are cloud-independent
 - Available for user analysis with Dask & Jupyter
- Work to understand cost of commercial cloud with respect to owning our ressources
- Look for interesting use cases
 - Cloud bursting: Dynamic/on-demand slot allocation
 - Network offloading: Use Google network for transfers
 - GPUs: On-demand GPU hardware
 - Special analysis workflows: ML, fitting, special MCs

Running various w/flows on 5k job slots at the Google site



Bursting up to 100k slots for simulation and reconstruction jobs



[ATLAS CHEP 2023](#)

Sustainability


Remember: Paris agreement goal

- Zero emission around 2100 => -50% by 2030 !
- To get the 50% in 2030 you have equivalently
 - to expand CO2 free energies by a factor 12
 - to increase energy efficiency by a factor 2
 - to save energy by a factor 2

→ A lot of initiatives to decrease our computing carbon footprint

- Understand our current and future footprint
- Reduce everywhere as soon as we know how to do it
 - reduce use of older nodes
 - reduce CPU clock speed when bad electricity footprint
- R&D and studies

• And we are big consumers!



- The CERN energy consumption is ~1.25TWh/year during LHC runs
- The CERN IT infrastructure contributes <5% in average
- CERN provides ~20% of the WLCG resources


[CHEP 2023](#)

Heterogenous computing and Carbon footprint: ARM


ARM architecture

- Promising architecture with low power consumption
- some sites planning to provide resources
- => most of the LHC workloads ported to ARM.
ATLAS software successfully validated and running. Ongoing validation for the 3 other LHC experiments
- In depth studies at Gridpp Glasgow site with very interesting results


ARM and AMD comparison




HEPScore/Watt workflow comparison



HEP-Score/Power vs frequency



Energy/thread vs frequency




Heterogenous computing and Carbon footprint: ARM

ARM architecture



- Promising architecture with low power consumption
 - some sites planning to provide resources
 - => most of the LHC workloads ported to ARM.
ATLAS software successfully validated and running. Ongoing validation for the 3 other experiments
- ➔ In depth studies at Gridpp interesting results

Other processor brands are joining this level of performance.


ARM and AMD comparison




HEPScore/Watt workflow comparison



HEP-Score/Power vs frequency




Energy/thread vs frequency



Estimating the WLCG footprint for HL-LHC

Simone Campana CHEP 2023


Distribution of power consumption



+


Watt/HS06

Watts per kHS06




+

Power Usage Effectiveness (PUE) of site




→ Effect of new computing center @Preveessin

CPU needs




→ Optimistic/pessimistic R&D scenarii

CPU energy needs per year



GWh and GWh/fb-1 estimation for WLCG




→ Energy needs in Run-4 and Run-5:

- +100% compared to Run-2 in the pessimistic scenario
- +10% in the optimistic scenario

Skilled and motivated people

Recognition


- Computing and software activities are fundamental in our research activities
 - they have huge impact on our physics results
 - their cost is similar than the one of the experiment
 - they are part of our experiments and are no more a « service »
- We need skilled and motivated people
 - Their work should be shared and publicised
 - Not only in dedicated meeting and conference !
 - Their expertise should be recognise also when physicists



M. Ballroom CHEP 2023

Training

- Students often lack of software and computing skills
- Computing and software are evolving quickly => continuous learning is needed
 - ➔ tutorials, training are important and there are lot's of opportunities:
 - tutorials, school and training locally, by their institution, in collaborations, organised by the HEP Software Foundation (HSF)...
 - HSF as a forum to build the community and share knowledge
 - [Software Training in HEP](#)



Open science


Context

International context

- **UNESCO**: recommendations for open Science
- **OCDE** : recommendations on research data
- **Europe**
 - The European Union requires open access to publications and data from research it funds;
 - since 2021, it has defined open science as a criterion for scientific excellence.
 - Promotion of EOSC – European Open Science Cloud
 - COARA: research evaluation
 - International working groups: RDA, GO FAIR, etc

Initiatives nationales

- Loi pour une République numérique (2016)
 - « *données aussi ouvertes que possible et aussi fermées que nécessaire* »
- [Plan national pour la science ouverte](#) (2018)
- [deuxième Plan national pour la science ouverte](#) (2021)
 - publications en accès ouverts pour les appels à projet
 - plateforme de données ouvertes : [data-gouv.fr](https://data.gouv.fr)
- [Feuille de route 2021-2024](#) (2021)
- **CNRS**
 - Feuille de route de la science ouverte (2019)
 - Plan des données de la recherche (2020)




Open science

Objectives

- Public research => knowledge has to be shared back
- Data and analysis preservation
- Quality and reproducibility
- Allow reuse or new usage => new science
- valorisation

Does not come for free


- need thinking, work (the earlier the better and the more efficient) and skills
- need infrastructures
- => need HR



Open

- publications
- data
- software
- hardware (CERN specificity)


Open Science in HEP: a long story



S. Dallmeier-Tiessen | CERN OS Policy

CERN tools

Open Science at CERN



Publications



SCOAP³ agreement

- <https://scoap3.org/>
- Partnership between 3,000 libraries, research organisations in 44 countries and 3 intergovernmental organisations (CERN, IAEA, JINR)
- Agreement with leading particle physics publishers for open publication at no cost to authors
- 7 editors, 11 journals
- Copyright belongs to authors and CC-BY licence
- Covers approximately 87% of scientific output

→ 6,500 articles/year

→ New => books: <https://scoap3.org/scoap3-books/>

INSPIRE-HEP

- INSPIRE (<http://inspirehep.net/>) : trusted community hub that helps researchers to share and find accurate scholarly information in high energy physics.
- Information platform for HEP community, comprising 8 interlinked databases on literature, conferences, institutions, journals, researchers, experiments, jobs and data.
- CERN, DESY, Fermilab, IHEP, IN2P3 collaboration (SLAC was one of the pioneer but quit),
- Serves the scientific community for almost 50 years
 - previously known as SPIRES
 - first website outside Europe and the first database on the web.
- Close interaction with the user community and with arXiv, ADS, HEPData, ORCID, PDG and publishers is the backbone of INSPIRE's evolution.

Publications à l'IN2P3

Partenariat INSPIRE

- [INSPIRE](#) : bibliothèque numérique ouverte pour la physique des hautes énergies
- **l'IN2P3** participe au pilotage d'Inspire avec CERN, Fermilab, DESY et IHEP et développe des outils et de services pour les besoins spécifiques de l'IN2P3
- 2 ETP (1 expert dédié + support INSPIRE et 4 curateurs)

Traitement des publications

- Moisson des (pré-)publications depuis arXiv et les éditeurs
- Exportation automatique INSPIRE → HAL
 - pré-publications, articles, actes de conférences publiés, chapitres de livre...
 - toutes les publications avec affiliation française prise en compte même sans affiliation IN2P3
- Vérification, enrichissement des métadonnées et validation par un catalogueur de l'IN2P3

- 3600 publications* traitées par an dans INSPIRE dont 2000 pour l'IN2P3
- portail HAL-IN2P3 : 64 000 notices, 2 500 notices par an (tous types de documents confondus)
- ~90% en accès ouvert

*en physique avec au moins un auteur français



FAIR Data: preamble

Open and FAIR data

- FAIR = findable, accessible, interoperable, reusable
- does not just mean opening up data to the outside world
- but also good practices for data management throughout the entire cycle:
 - collecting high-quality data → describing it with rich metadata
 - identifying it (not just with a DOI) and referencing it → storing it in reliable storage facilities
 - process and analyse it → delete, clean or archive it → open it... or not
- software is part of the process and is also a type of data that follows the same rules: it must be referenced, versioned, stored and opened (if the conditions are met) + licence
 - the data and the corresponding software must be associated

Consequences

- Well-referenced data and software with well-defined metadata facilitate high-quality processing and analysis and enable checks, reproducibility, sharing and reuse to avoid duplication.



Software

Software needs also good practice

- quality
- documentation
- deployments
- forge
- licence
- archiving => Software heritage



Open Source/free software

Structuration for help

- Open Source Program Offices (OSPOs)
 - CERN <https://opensource.web.cern.ch/>
 - GRICAD Grenoble <https://ospo-uga.sciencesconf.org/>

Recent recommendations

Recommendations for Best Practices for Data Preservation and Open Science in HEP

- <https://arxiv.org/abs/2508.18892v1>

1 Executive Summaries

1.1	Executive Summary for Host laboratory	...
1.2	Executive Summary for Experiment management	...
1.3	Executive Summary for Home institute	...
1.4	Executive Summary for WG leaders	...
1.5	Executive Summary for Funding agency	...
1.6	Executive Summary for Tool developers	...
1.7	Executive Summary for Analysts	...
1.8	Executive Summary for Data management	...
1.9	Executive Summary for Open data group	...

arXiv:2508.18892v1 [hep-ex] 26 Aug 2025

Recommendations for Best Practices for Data Preservation and Open Science in HEP

Simone Campana¹, Irakli Chakaberia², Gang Chen³
Cristinel Diaconu^{4,5}, Caterina Doglioni⁶, Dillon S. Fitzgerald⁷
Vincent Garonne⁸, Anne Gentil-Beccot¹, Fleur Heiniger¹
Michael D. Hildreth⁹, Julie M. Hogan¹⁰, Hao Hu³
Eric Lancon⁸, Clemens Lange¹¹, Kati Lassila-Perini^{12*}
Olivia Mandica-Hart¹, Zach Marshall², Thomas McCauley⁹
Harvey Newman¹³, Mihoko Nojiri¹⁴, Ianna Osborne¹⁵
Fazhi Qi³, Salomé Rohr¹, Stefan Roiser¹
Thomas Schörner¹⁶, Ulrich Schwickerath¹, Elizabeth Sexton-Kennedy¹⁷
Seema Sharma¹⁸, Tibor Šimko¹, Michael Sparks⁹
Graeme Andrew Stewart¹⁶, Nicola Tarocco¹, Giacomo Tenaglia¹
Gustavo Valdivieso¹⁹, Antonia Winkler^{20,1}, Christoph Wissing¹⁶

¹ CERN, Switzerland

² Lawrence Berkeley National Laboratory, USA

³ Institute of High Energy Physics, China

⁴ Centre de Physique des Particules de Marseille CPPM, France

⁵ CNRS/IN2P3 and Aix-Marseille Université, France

⁶ University of Manchester, UK

⁷ University of Michigan, USA

⁸ Brookhaven National Laboratory, USA

⁹ University of Notre Dame, USA

¹⁰ Bethel University, USA

¹¹ Paul Scherrer Institute, Switzerland

¹² Helsinki Institute of Physics, Finland

¹³ California Institute of Technology, USA

¹⁴ KEK, Japan

¹⁵ Princeton University, USA

¹⁶ Deutsches Elektronen-Synchrotron DESY, Germany

¹⁷ Fermi National Accelerator Laboratory, USA

¹⁸ Indian Institute of Science and Education (Pune), India

¹⁹ Federal University of Alenas, Brazil

²⁰ Humboldt University of Berlin, Germany

August 27, 2025

Version: v1.0-16a371e

These recommendations are best read on the web site:

<https://icfa-data-best-practices.app.cern.ch/>

* Corresponding editor: Kati Lassila-Perini (contact via icfa-data-best-practices-contact@cern.ch)

Conclusion

Conclusion

→ Software and Computing are key elements of our experiments

Starting with a success


- Computing for LHC run1 and 2 with a complete and complex set of software, middleware and WLCG infrastructure at the Exa-scale successfully allows to store, process and analyse LHC data, leading to wonderful scientific results

Take up the HL-LHC software and computing challenge

- Infrastructure and software will have to cope for high luminosity (200 x run 2) + pile-up up to 200 + throughput ~10TB/s
- In a context that is rapidly evolving: hardware, technologies, cost → constraints and opportunities
- Taking into account Open Science and effort to reduce our carbon footprint

Huge amount of work and R&D already done and still a lot ahead

- New developments integrated in production as soon as validated
- Software, hardware and computing models need to be adapted for heterogenous ressources, more parallelisation => flexibility needed to adapt to ressources not build for and by us
- Make best use of new technologies and technics: AI, progress done outside our field
- Be prepared to the next (r)evolution like Quantum Computing
- Training is very important as proper recognition for the work done in the computing and software fields



Infos pratiques

Rappel : canaux d'informations

Développements logiciels

- projet Decalog
- Liste de diffusion : DEV-LOGICIELS-L@IN2P3.FR
- en cours : valorisation logiciel libre Open

Infra et projets européens

- euro-computing-L@IN2P3.FR
- plateformes, LCG, France Grille, EGI, EOSC
- dernière JCAD : <https://jcad2025.sciencesconf.org/>

HPC

- projet FITS, exascale (=> Alice Recoque), JENA HPC WG
- Liste de diffusion : HPC-INFO-L@IN2P3.FR

RI3 <https://ri3.in2p3.fr>

- JI : dernière édition : <https://indico.in2p3.fr/event/31391>
- Webinaires et écoles
- Webcasts : <https://webcast.in2p3.fr/container/ri3>

CC-IN2P3

- <https://portail.cc.in2p3.fr/login>

ML/IA Cf Julien Donini

- Projet ML : <https://machine-learning.pages.in2p3.fr/>
- Liste de diffusion : MACHINE-LEARNING-L@IN2P3.FR
- Workshop ML IN2P3 (+IRFU): Caen 26-28/11/2025
 - <https://indico.in2p3.fr/event/36597/>
- Centre AISSAI CNRS :
 - centre dédié à l'IA : <https://www.cnrs.fr/fr/le-centre-artificial-intelligence-science-science-artificial-intelligence-aissai>
 - semestre thématique pour l'IN2P3, appel pour accompagnement d'ingénieur

Informatique Quantique

- Projet QC2I : <https://qc.pages.in2p3.fr/web/>

Science ouverte

- DOP2I working group

Publication

- Contacts : IST@IN2P3.fr ou Mathieu Grivès

Questions

Cartouche T10000 T2

Marque

- Oracle
- Storagetek T10k T2

Caractéristiques

- cartouche 12,5x11x1,5cm, poids 270g
- longueur de la bande 1147m
- épaisseur 5,2 μm

Capacité

- 8,5 To

Archivage

- 30 ans
- tx d'erreurs non corrigées : 10^{-19}
- T° 15-26C



