



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking




ÉIRE2026.EU
IRELAND2026.EU

Lattice QCD: An HPC Community Case Study

Antonio Rago - Chair EuroLFT

Computational Physics: From Particles to Cosmos

Major achievements of lattice field theory



QCD and the strong interaction

Hadrons, quark-gluon structure and fundamental parameters




Figure 12.1: Sketch of the unitarity triangle.


Weak and electroweak physics

Flavour, decays and precision tests of the Standard Model



Finite-temperature and finite-density QCD

Phase diagram, early Universe and neutron-star matter, gravitational waves.




Beyond the Standard Model

New strongly coupled gauge theories and mechanisms



Algorithms and computational advances

Low modes, scalable solvers and exascale systems





Broader impact beyond QCD

Quantum materials, nuclear theory and quantum algorithms

From measurements to precision tests

Experiments measure observables. Lattice QCD calculates the non-perturbative inputs needed for precision tests.

Experimental measurements



Experiments measure how often each final state appears at each collision energy.

Precision tests of the Standard Model

Quantity	Measured (GeV)	SM prediction (GeV)
Mass of W boson	80.387 ± 0.019	80.390 ± 0.018
Mass of Z boson	91.1876 ± 0.0021	91.1874 ± 0.0021

Illustrative values: A. Rago, DEIC Conference (2023), slides 6-7

... a race for **precision!**

Lattice QCD turns measured events into quantitative tests by calculating the strong-interaction contribution from first principles.

Exact on the given lattice

Controlled extrapolations make lattice QCD a precision method

“Numerical simulations have the reputation of being an approximate method that mainly serves to obtain qualitative information on the behaviour of complex systems.

This is, however, not so in Lattice QCD, where the simulations produce results that are exact (on the given lattice) up to statistical errors. The systematic uncertainties related to a non-zero lattice spacing and finite volume, still need to be investigated, but these effects are theoretically well understood, and can usually be brought under control.”

Martin Lüscher, 2002

Fixed lattice parameters

Exact result with quantified statistical uncertainty

Finite spacing and volume


Systematic effects are understood and controlled

Physical prediction

Continuum and infinite-volume extrapolations


Lattice QCD: a controlled computation

Discretize space-time, sample quantum fields, and remove lattice artifacts




Discretize QCD

Replace continuous space-time with a finite lattice that a computer can sample.



Generate ensembles

Measure observables across many statistically representative field configurations.



Recover continuum physics from observables

Take the lattice spacing towards zero to control discretization effects.





Massively parallel HPC

Thousands of nodes make controlled field-theory simulations feasible.


Validation by design

Independent formulations must agree in the continuum limit

The LHC analogy



FLAG2024


 $f_+(0)$ 

The strong interaction becomes quantitative

Lattice QCD determines the running strong coupling across energy scales

Lattice formulation

Discrete space-
time makes QCD
computable



Precision result


One running
coupling across
many scales

A typical lattice-QCD workflow

A representative workflow combines ensemble generation, GPU-accelerated measurements, and statistical postprocessing

1

Simulation



Markov-chain Monte Carlo generates ensembles.
1-4 long CPU/GPU chains, $O(10^5)$ cores/
 $O(10^3)$ GPUs.
10-100 M core-hours. Data: $O(10)$ TB.

2

Analysis



Sparse solves use simulation output.
Many parallel GPU jobs, $O(10^3)$ GPUs.
1-10 M GPU-hours. Data: $O(100)$ TB.

3

Postprocessing




Statistical fits, errors, and visualization.
 $O(100)$ CPU cores or one local GPU.
Curated results: $O(10)$ GB.

The workflow shifts workload, hardware, and data volume at every stage.

The sparse Dirac operator

A local operator that must be solved repeatedly at extreme scale



Representative sparsity pattern

Locality makes the matrix sparse

Each lattice site couples only to its nearest neighbours. The resulting matrix has a narrow structure even for an enormous global system.


Realistic calculations reach $O(10^{10})$ degrees of freedom.

Repeated sparse solves

Generation and measurement both return to the same sparse linear system.

Spatial decomposition and halo exchange

Each node owns a local sub-lattice and shares boundary data with its neighbours



Neighbouring nodes exchange only the boundary layer

The local stencil crosses boundaries

Each node computes on its own block. Sites near a neighbour need a copy of that neighbour's boundary values.



Halo exchange supplies those values

The exchange transfers thin boundary layers between adjacent subdomains before the next computation.

Smaller local volumes increase communication relative to computation.

Dirac-operator scaling and the multi-scale problem

The local operator scales well. The full calculation needs better algorithms.



Machine — LUMI-G — JUWELS

HiRep v2, Fig. 31: strong and weak scaling to 1,024 GPUs

Why the local operator scales well

Local computation $\propto L^4$

Halo exchange $\propto L^3$

For a local four-dimensional lattice of extent L , work grows with the volume and boundary communication with the surface.

The full calculation has many scales


Light quarks, large volumes and long-distance modes increase solver iterations. More hardware alone does not remove that cost.

Algorithms must improve problem scaling alongside machine scaling.

Algorithms reduce cost across the workflow


Algorithms now reduce the cost of solves, ensembles, and measurements

Early solver progress



2001-04: the practical light-quark cost barrier moved

1987–2010: reducing the cost of a solve



Fewer iterations and a smaller critical-slowing-down penalty

2010s–today: algorithmic gains span the full workflow

Configuration generation

Integration, preconditioning, and sampling reduce the cost of independent ensembles

Observable measurement

Low-mode and stochastic methods, AMA, and distillation improve precision per configuration

AI expands the workflow

A growing sector, subject to validation, alongside HPC

Active directions

Generative sampling
Variance reduction
Surrogate models
Parameter tuning
Analysis

Scientific guardrails

Exactness and bias control
Validated observables
Error budgets
Provenance

Expanding sector

Rapid progress in
sampling, measurement,
and analysis—paired
with HPC and validation.

Four transferable lessons

What computational physics can take from lattice QCD

Algorithms can rival hardware generations.

Methods shift the feasible science frontier.

Data outlive machines.

Metadata and access multiply the return on simulation.

Communities are infrastructure.

Shared software and validation create durable capability.

Extreme workloads are co-design partners.

They reveal what future systems must do well.

Physics, algorithms, software, data, and supercomputers progress together.



EuroHPC
User Days

Dublin Ireland
22→24
September 2026




ML-Based Event Reconstruction for Experiments at the Future Circular Collider

Lena Herrmann, CERN




Preparation for the future of EU high energy physics



[fcc.web.cern.ch]

Proposed Future Circular Collider


- 91 km ring, built at CERN
 - First stage: FCC-ee
- **Precision measurements:**
- Higgs Couplings
 - Flavour Observables
 - Electroweak Parameters



[ATL-PHYS-PUB-2025-018]




Preparation for the future of EU high energy physics



[fcc.web.cern.ch]

Proposed Future Circular Collider

- 91 km ring, built at CERN
 - First stage: FCC-ee
- **Precision measurements:**
- Higgs Couplings
 - Flavour Observables
 - Electroweak Parameters




[ATL-PHYS-PUB-2025-018]




Precision Measurements central to FCC Physics Program

under development:

High Performance Detectors



High Quality Reconstruction



Interplay

Needs flexible algorithms with fast turn-around

Precision Measurements central to FCC Physics Program


under development:

High Performance Detectors

High Quality Reconstruction


Interplay

Needs flexible algorithms with fast turn-around



The Subdetectors ...

... onion-like **barrel structure** instrumented with
... **optimized detection layers** that
... yield **complementary particle properties**




- **Tracker:** trajectories and momenta of charged particles
- **ECAL:** electromagnetic showers, energies of electrons and photons
- **HCAL:** hadronic showers, energies of hadrons

Precision Measurements central to FCC Physics Program


under development:

High Performance Detectors



Interplay

High Quality Reconstruction




Needs flexible algorithms with fast turn-around

Particle Flow (PF)

→ reconstruct stable particles

Categories	Particles
Charged Hadrons (CH)	$h^\pm$: $\pi^\pm$, $K^\pm$, p
Neutral Hadrons (NH)	h^0 : n, K^0
Photons	γ
Electrons	$e^\pm$
Muons	$\mu^\pm$

- Characteristic signature per particle category and subdetector
- Combined subdetector information crucial for particle identification




Particle Flow (PF)

→ reconstruct stable particles

Categories	Particles
Charged Hadrons (CH)	$h^\pm$: $\pi^\pm$, $K^\pm$, p
Neutral Hadrons (NH)	h^0 : n, K^0
Photons	γ
Electrons	$e^\pm$
Muons	$\mu^\pm$

- Characteristic signature per particle category and subdetector
- Combined subdetector information crucial for particle identification




Particle Flow (PF)

→ reconstruct stable particles

Categories	Particles
Charged Hadrons (CH)	$h^\pm$: $\pi^\pm$, $K^\pm$, p
Neutral Hadrons (NH)	h^0 : n, K^0
Photons	γ
Electrons	$e^\pm$
Muons	$\mu^\pm$

- Characteristic signature per particle category and subdetector
- **Combined subdetector information crucial for particle identification**




Particle Flow (PF)

→ reconstruct stable particles

Categories	Particles
Charged Hadrons (CH)	$h^\pm$: $\pi^\pm$, $K^\pm$, p
Neutral Hadrons (NH)	h^0 : n, K^0
Photons	γ
Electrons	$e^\pm$
Muons	$\mu^\pm$

- Characteristic signature per particle category and subdetector
- **Combined subdetector information crucial for particle identification**




Particle Flow (PF)

→ reconstruct stable particles

Categories	Particles
Charged Hadrons (CH)	$h^\pm$: $\pi^\pm$, $K^\pm$, p
Neutral Hadrons (NH)	h^0 : n, K^0
Photons	γ
Electrons	$e^\pm$
Muons	$\mu^\pm$

- Characteristic signature per particle category and subdetector
- **Combined subdetector information crucial for particle identification**



Particle Flow (PF)

Complex Signatures!

~ 40 particles per event (Z(jj))


→ ~ 4000 hits (energy depositions)

→ ~ 20 tracks

Geometric overlapping signatures


Large Datasets:

~ 6×10^{12} events ($\sqrt{s} = 91$ GeV)



End-to-end HitPF approach


Inputs:



calorimeter/muon hits
+ recon. tracks

- Energy
- Coordinates (x,y,z)
- Type (ECAL, HCAL, muon, track)


Machine Learning Model



Embedded as

- scalar
- points
- binary one-hot scalar features

DPC Density Peak Clustering




PF Object

Output

- Per hit: latent coordinates
- Condensation scores


Problem: Instance segmentation

- identify particles from hits
- assign unique identifiers and properties
- Target: Simulated particles entering the calorimeter (Unambiguous training label, unlike stable generator particles)



End-to-end HitPF approach


Inputs:



calorimeter/muon hits
+ recon. tracks

- Energy
- Coordinates (x,y,z)
- Type (ECAL, HCAL, muon, track)

Machine Learning Model



Embedded as


- scalar
- points
- binary one-hot scalar features

Output

- Per hit: latent coordinates
- Condensation scores

DPC


Density Peak Clustering



PF Object


Problem: Instance segmentation

- identify particles from hits
- assign unique identifiers and properties
- Target: Simulated particles entering the calorimeter (Unambiguous training label, unlike stable generator particles)



End-to-end HitPF approach


Inputs:



calorimeter/muon hits
+ recon. tracks

- Energy
- Coordinates (x,y,z)
- Type (ECAL, HCAL, muon, track)

Machine Learning Model



Embedded as


- scalar
- points
- binary one-hot scalar features

Output

- Per hit: latent coordinates
- Condensation scores


DPC

Density Peak Clustering



PF Object


Properties



The Dataset

Physics Dataset

- back-to-back jets
- High particle multiplicity




Specifications:

- $e^+e^- \rightarrow Z(q\bar{q}), q = u, d, s$ at $\sqrt{s} = 91\text{GeV}$
- 1.1M events
- Full simulation CLD (CLIC-like detector)
- ~25 particles per event
- Broad spectrum of charged and neutral hadrons
- Broad energy range: $0 < E < \sim 50\text{ GeV}$


Application:

- “clustering model”



Gun Dataset

- isolated particles and prompt γ




Specifications:


- Isolated particles
- 200k events
- Full simulation CLD (CLIC-like detector)
- 10 particles per event
- Balanced in CH, NH, e, γ, μ
- $0.5 < E < 50\text{ GeV}$

Application:

- “property models”



Particle Identification: Efficiency and Fake Rate




MC

High Efficiency:

$$\epsilon_i = \frac{\#_{\text{rec},i}}{\#_{\text{true},i}}$$

Fraction of correctly reconstructed particles relative to the true particles

Comparison Benchmark:
classical **Pandora** approach



RECO


Low fake rate:

$$\text{fake rate}_i = \frac{\#_{\text{fake},i}}{\#_{\text{rec},i}}$$

Fraction of reconstructed particles without **matching true particle** ('fake'), relative to all reconstructed particles


Truth Matching
via Intersection Over Union
cost and Hungarian matching

Event Level Observable: Invariant Mass



Rough Estimate for ZH run:

**averaged numbers, simplified physics assumptions*



- Important observable for physics analyses
- Integrates resolution/ reconstruction efficiencies/ fake rate

Generic particle flow algorithm with short turn around

- better performance than the classical approach Pandora on benchmark dataset
- Gains are strongest in complex, overlapping topologies, especially for neutral hadrons

HitPF development ongoing:

- other detector concepts
- different sub-detector composition
- CH property regression
- Considering timing information
- Other model architectures (MaskFormer)
- Other energies/ more physics processes



Publication: [\[arxiv:2603.04084\]](https://arxiv.org/abs/2603.04084)

- Under review at Nature Comm. Phys.


Other ML Reconstruction Projects

ML based tracking finding

- Integrate in HitPF
- Quantify beam induced background effects




Downstream **jet quark tagger**
based on HitPF inputs



ML solutions for specialized tasks

- Looper identification
- Photon-Pion separation



... and more!

Access Calls

AI Factories Access Call


for Science and Collaborative EU Projects

- EHPC-AIF-2025SC02-04:
Event Building at the Future Circular Collider ML-Based Track Reconstruction and Particle Flow
- October 2025 - June 2026
- 20 700 node hours
- MN5 ACC (BSC)


Regular Access Call

- EHPC-REG-2026R01-012:
Event Building at the Future Circular Collider: Machine Learning based high level reconstruction
- October 2026 – October 2027
- 20 750 node hours
- MN5 ACC (BSC)

BSC Performance of HitPF



- Wall clock time to complete 10 000 training steps
- Global batch size increased linearly
- Variations below 1% relative to baseline



- Wall clock time to complete 10 000 training steps
- Total workload fixed (per device batch size: 24,12,6,3)
- Parallel efficiency of 99.3%, 95%, 85% (2,4,8GPUs)

4 GPUs chosen for balance between speedup and efficient use of resources

Experience

- Rapid access to resources following grant approval
- Responsive technical support from BSC team
- Occasional long queueing times
e.g. 7h queueing for 5h on one node
- Allocation extension granted
 - simplified paper submission to journal
 - no resource access during review stage
- Lack of WAN on MN5 complicates
 - maintenance of software environment
 - use of logging tools i.e. *Weights and Biases*
- EuroHPC federation platform: passport required for registration
 - complication for students from abroad

Overall positive!

Team of > 15 including
students and growing!

Bottleneck: *available in-house GPUs*

SOLVED by the support of
EuroHPC and
MN5 APP

Thank you!




Appendix

Reconstruction Pipeline

CLUSTER BASED

Sequential, fine-tuned, **intermediate clustering**, detector specific, complicates detector development




HIT BASED

End-to-end approach, flexible, **low-level inputs**, detector agnostic, facilitates detector development

Truth Matching


- Find the best *reconstructed-to-truth* assignment via Hungarian matching
- Intersection over Union (IoU) cost matrix: match reconstructed particle i to truth particle j



	i_1	i_2	i_3
j_1	0.1	0.4	0.8
j_2	0.3	0.5	0.2
j_3	0.6	0.15	0.7


Find assignment a that
minimizes total cost

$$\text{IoU} = \frac{\text{intersection}}{\text{union}}$$



$$C = \underset{i}{\operatorname{argmin}} \sum C_{i,a(i)}$$

PID: Efficiency and Fake Rate



CH:

- Up to 20% higher eff.
- Fake rate $G(2)$ smaller
- Eff. limited by track reco.

Photons:

- ~5% higher eff. at high E
- Improvement by up to factor 2 for low E
- Fake rate up to ~ $G(1)$ red.

NH:

- Up to 20% higher eff. with fake rate reduced by ~ factor 5 in this region
- Better clustering (no fake clusters from splitting)

PID: Confusion Matrix

CH	77
	66

Top row: HitPF
Bottom row: Pandora

HitPF achieves ...
higher efficiency and lower
fake rate than Pandora
across all energies

1 GeV < E < 10 GeV

True		1 GeV < E < 10 GeV					
		CH	γ	NH	e	μ	missed
CH	CH	77	1	15	1	2	4
	γ	66	4	21	3	0	6
γ	CH	0	95	1	0	0	3
	γ	0	93	2	0	0	5
NH	CH	1	4	63	0	0	32
	γ	2	6	49	0	0	43
e	CH	1	50	2	42	0	4
	γ	2	62	16	15	0	5
μ	CH	4	0	36	0	58	1
	γ	37	0	36	0	23	4
fake	CH	0	0	7	0	0	0
	γ	1	2	10	1	2	0

Predicted

10 GeV < E < 100 GeV

True		10 GeV < E < 100 GeV					
		CH	γ	NH	e	μ	missed
CH	CH	84	0	13	1	0	1
	γ	75	2	19	1	0	3
γ	CH	0	98	1	0	0	1
	γ	0	96	2	0	0	2
NH	CH	3	3	90	0	0	3
	γ	4	8	78	0	0	9
e	CH	0	51	2	45	0	2
	γ	1	77	2	18	0	3
μ	CH	0	0	10	0	86	3
	γ	2	1	20	0	71	6
fake	CH	0	0	0	0	0	0
	γ	0	0	1	1	1	0

Predicted

Higher PID purity:

→ Stronger population of
diagonal for HitPF

CH:

- Often misidentified as NH
- Tracking inefficiencies

Electrons:



- Often misidentified as γ
- Driven by secondary e with
low tracking efficiency

Energy and Angular Resolution



Energy Resolution: half-width of 16th-84th percentile normalized to median

Angular resolution: 68th percentile of the opening angle between rec and true



Charged Hadron



Photon



Neutral Hadron



CH:

Comparable by construction

→ solely tracker
information used

NH:

Largest gain in energy/
angular resolution

→ Less shower merging,
better hit capture

Cost & CO_2 Saving for ZH run – Slide 19

- **Runtime:** 3 years $\rightarrow \Delta t = 0.22$ F36 months $\approx$ 8 months [\[arXiv:2505.00274\]](https://arxiv.org/abs/2505.00274)
- Total stated **running cost:** 400MCHF/a (electricity + personnel) [\[arXiv:2505.00274\]](https://arxiv.org/abs/2505.00274)
Extrapolated ZH cost: 400MCH F3a = 1200 MCHF
Saving: 0.22 F1200MCHF $\approx$ 250 MCHF

- **CO_2 savings:**

Annual power Z : 1070 GWh/a, 222 MW colliding mode power

tt : 1770 GWh/a, 357 MW colliding mode power

$$\text{Slope: } \frac{1770-1070 \text{ GWh}}{357-222 \text{ MW}} \approx 5.2 \frac{\text{GWh}}{\text{MW}}$$

ZH colliding mode power: 294 MW $\rightarrow \sim 1444$ GWh/a $\rightarrow$ 3a: 4.33 TWh

$\Delta E = 0.22$ F4.33 = 0.95 TWh

Grid Carbon Intensity: 29.0 g CO_2 eq/kWh [\[assets.rte-france.2026-04\]](https://assets.rte-france.2026-04)

ΔCO_2 : 0.95 TWh F29.0 g CO_2 eq/kWh ~ 27.5 10kt CO_2 eq

Significance

$$Z = \frac{s}{\sqrt{B}} \propto \sqrt{\frac{\mathcal{L}}{\sigma_m}}$$

$$Z_1 \stackrel{!}{=} Z_2$$

$\rightarrow \mathcal{L}$ scales
linearly



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



An tAire Riaracháin agus an tAire Dáilí,
Dáil Éireann, Na hOileáirí agus Éireann,
Dáil Éireann agus Éireann,
Research, Innovation and Science

ICHEC
INSTITUTE OF COMPUTATIONAL HIGH ENERGY PHYSICS



ÉIRE2026 EU
IRELAND2026.EU

Determination of the strong interaction scale

Probing the Standard Model of particle physics at new levels of precision

Zoltan Fodor

Budapest–Marseille–Wuppertal collaboration (BMW)

BMW/DMZ, Nature 653 (2026) 373–377, arXiv:2407.10913v2

Muon g-2

A precision magnetic fingerprint


$$a_\mu = (g_\mu - 2)/2$$

$$a_\mu^{\text{SM}} =$$

$$11659205.2(3.6) \times 10^{-10}$$


total uncertainty: 0.31 ppm – the precision scale
for new-physics searches

- ▶ Quantum fields shift the muon magnetic response away from exactly 2.
- ▶ Experiment measures spin precession; theory predicts the same number from known physics.
- ▶ A persistent mismatch, after all uncertainties are controlled, could signal physics beyond the Standard Model.



Prediction and measurement now agree within 0.5 σ , tightly constraining any new-physics contribution to a_μ .

The precision limit



Nearly a decade of sustained effort: uncertainty reduction from the 2017 calculation to the latest HVP result.

The strong-force cloud sets today's theory limit

- ▶ Experiment has reached 0.127 ppm:
 $\sigma_{\text{exp}} \simeq 1.5 \times 10^{-10}$.
- ▶ The current HVP uncertainty is 3.4×10^{-10} , **about 2.3× the experimental error**.
- ▶ In the dominant light-quark HVP window, the lattice-ruler w_0 term is 1.72×10^{-10} —already larger than the full experimental error.

In that window, w_0 accounts for 39% of the total variance.

Reducing it is a necessary step toward experimental precision.

Lattice QCD

The equation is compact but not perturbative


$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G_{\mu\nu}^a G^{a\mu\nu} + \sum_q \bar{q} (i\gamma^\mu D_\mu - m_q) q,$$

$$D_\mu = \partial_\mu - i g_s T^a A_\mu^a$$

- ▶ The first term is gluon fields; the second is quarks moving through them.
- ▶ The free inputs, shown in red, are the strong coupling and quark masses; measured hadrons fix them.
- ▶ At low energy perturbation theory fails, so supercomputers must sample the fields numerically.

Space-time becomes a finite four-dimensional grid

quark fields $\psi(x)$ on sites gluon gauge links $U_\mu(x)$



Repeat at several a and L , then extrapolate to $a \rightarrow 0$ and $L \rightarrow \infty$.

Bare β and scale

The bare gauge coupling controls the grid spacing; a physical observable calibrates it.

SIMULATION INPUT

$$\beta = \frac{6}{g_0^2} = 4.3032$$

- ▶ $\beta = 6/g_0^2$ is the dimensionless bare gauge coupling fixed alongside the bare quark masses.
- ▶ For a fixed action, larger β gives a finer—but still uncalibrated—grid.

NONPERTURBATIVE CALIBRATION

$$\underbrace{(w_0/a)_{\text{lat}}}_{\text{Wilson flow}} \quad \text{and} \quad \underbrace{w_0^{\text{phys}}}_{\text{physical anchor}}$$

$$a = \frac{w_0^{\text{phys}}}{(w_0/a)_{\text{lat}}}$$

- ▶ Wilson flow supplies w_0/a , an inexpensive and precise ruler in lattice units.
- ▶ After the quark masses are tuned, a physical anchor fixes w_0^{phys} and therefore a .

$\beta = 4.3032$ targets $a \simeq 0.036$ fm: **25% finer** than the current 0.048 fm endpoint.

Why these anchors?

Both are experimentally precise; their lattice systematics are complementary.

BARYONIC ANCHOR

Ω^- mass

$$M_{\Omega}^{\text{exp}} = 1672.45 \text{ MeV}$$

- ▶ Precisely measured and stable under the strong interaction.
- ▶ Its sss content suppresses leading strong-isospin sensitivity; the QED component in the Ω -based w_0 decomposition is about 0.06% and controlled.
- ▶ An independent baryonic anchor, but the three-quark signal is costly and needs excited-state control.

MESONIC ANCHOR

Pion muonic decay



$$F_{ud}^{\text{QCD+QED}} = 131.711(45) \text{ MeV}$$

from $\Gamma(\pi^+ \rightarrow \mu^+ \nu_{\mu} [\gamma])$ and PDG V_{ud}

- ▶ The phenomenological input is known to 0.034% relative precision.
- ▶ On the lattice, f_{π} comes from a clean, inexpensive pseudoscalar two-point function.
- ▶ Leading QED and isospin-breaking effects are included, with precise continuum control.

Pion decay sets the production scale; Ω^- provides an independent baryonic cross-check.

Why 0.036 fm matters

25% finer0.048 $\rightarrow$ 0.036 fm**44% smaller**nominal a^2 factor **$\sim 2\times$ tighter**0.28% $\rightarrow$ 0.15% projected scale error

Existing fits spread
near the continuum
extrapolation.

The target fine point
directly constrains the
continuum trend
where they diverge
most.

Continuum fits for $w_0 M_\Omega$ (left) and blinded f_K / f_{π^+} (right); the red arrow marks the target $a \simeq 0.036$ fm point.

The target ensemble

$$\beta = 4.3032$$

bare gauge coupling targeting $a \simeq 0.036$ fm

$$176^3 \times 264$$

1.44 billion sites on a four-dimensional space-time grid


$$\simeq 6.3 \times 9.5 \text{ fm}$$

large physical volume retained while the grid becomes finer

2,000 target snapshots


from 20,000 trajectories

A new endpoint extends an ensemble program built over a decade



Existing $\beta \leq 4.1479$; target $\beta = 4.3032$ extends the sequence.

44-node production



QUADA CG scaling for $176^3 \times 264$: solid LUMI-G and dashed Adastra curves agree; the red marker is 44 nodes.

The target measurement fits,
scales and stays within the job
limit

44 nodes / 352 GCDs

smallest partition for Ω ; the 18.0 TB working set
uses about 85% of GPU memory

10.6 h

Adastra measurement time per configuration;
CG comparison projects LUMI-G below 24 h

Planned EuroHPC node-hours

73% Ω 23% mesons 4% flow

Projected gains

The 0.036 fm program targets one shared systematic across observables.

Quantity	Starting point	Program target	Scientific gain
Lattice ruler w_0	Published: 0.28%	0.15%	Reusable calibration across observables
LO-HVP in $g - 2$	Published: 0.48%	about 0.36%	Sharper SM/new-physics test
f_{K^+}/f_{π^+}	2024 baseline: 0.18%	below 0.10%	Sharper weak-decay and CKM-unitarity test

The first two starting values are final-v2; f_{K^+}/f_{π^+} is a 2024 baseline. Program targets are 2024 planning projections, not completed results.

EuroHPC resources deliver one ultrafine spacing—and reusable precision across observables.



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



An tAire Forasairdeachais agus Ardteicneolaíocht,
Dáilghle, Na hEilinn agus an Fódairíocht,
Is gairneoir do Riaráil agus Léimneachas,
Eolas, Innealtas agus Saol



ÉIRE2026.EU
IRELAND2026.EU

Thank You



EuroHPC
User Days

Dublin Ireland
22→24
September 2026




Simulating the Intergalactic Medium

Laura Keating, University of Edinburgh

THE
ROYAL
SOCIETY



THE UNIVERSITY *of* EDINBURGH
School of Physics
& Astronomy





100 Mpc



Intergalactic
gas in a
cosmological
simulation

60
kpc


The EAGLE Simulation
Schaye+14








1D Power Spectrum

*Small scale information:*

Temperature of the intergalactic gas
Probe of the matter power spectrum
Nature of dark matter

3D Power Spectrum

*Large scale information:*

Matter clustering
Neutrino masses
Baryon Acoustic Oscillations


Simulating the Intergalactic Medium: Main Ingredients

Gravitational interaction of Dark Matter, Gas and Stars:
Tree-Particle Mesh N-body solver


Hydrodynamics of Gas:
Smooth Particle Hydrodynamics
(Lagrangian treatment)

Cosmology:
Initial Conditions
Expansion of Universe

Other Physics:
Star formation, Radiative
processes (heating/cooling)




Bolton+16



The Sherwood Simulation Suite

35 Million CPU hours in two
previous PRACE allocations to
perform high-resolution
intergalactic medium simulations





New spectroscopic surveys like
DESI and the future Extremely Large
Telescope demand new simulations
in large volumes and at high
resolution

Simulating cosmic structure formation with the GADGET-4 code

Volker Springel^{1★}, Rüdiger Pakmor¹, Oliver Zier¹, and Martin Reinecke¹

¹*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85740 Garching bei München, Germany*

Plus some customisation:

Star formation routine optimised
for intergalactic medium
simulations (factor 2 speed-up)

On-the-fly calculation of quasar
absorption spectra

Non-equilibrium chemistry

Updated model for the
metagalactic ionizing background

New EuroHPC allocation: Probing early galaxy formation and cosmology with large-volume simulations of the intergalactic medium




LUMI-C: 120k Node Hours in
Regular Access Mode

*Chosen for system size and
memory:*

Largest simulations run on 512
nodes


Full node memory available
through /dev/shm
(Necessary for shared memory
parallelisation)




Weak scaling of GADGET-4
on LUMI-C:

Starts to deviate from perfect
scaling above 8 nodes, but
our proposed simulations
were memory-limited


Strong scaling of GADGET-4 on LUMI-C: Excellent performance up to 512 nodes



One additional modification we needed as a new implementation of the domain decomposition, driven by the relatively low memory per node on LUMI-C



One additional modification we needed as a new implementation of the domain decomposition, driven by the relatively low memory per node on LUMI-C



Outcome of LUMI-C allocation:

Two new flagship intergalactic medium simulations

320 cMpc/h and 640 cMpc/h with 4096^3 gas particles reaching a maximum of 3 - 6 ckpc/h spatial resolution

64 times larger particle number than our previous simulations

Analysis in progress: results coming soon!

Conclusions:

- The intergalactic medium is a powerful tool for studying the composition of our Universe
- Simulations are essential for interpreting this dataset, and as the observations improve we need the models to keep up
- Facilities like EuroHPC machines are allowing us to reach this goal

Numerically Simulating Plasmas in Extreme Environments


PARES - Pair Acceleration and Radiation Emitting Sources

Bernardo Barbosa

GoLP / Instituto de Plasmas e Fusão Nuclear, Universidade de Lisboa, Lisbon, Portugal

bernardo.barbosa@tecnico.ulisboa.pt

golp.tecnico.ulisboa.pt



Acknowledgments & collaborators



Marija Vranić, Luís Silva, Jorge Vieira, Thales Silva, Thomas Grismayer, Lucas Gamiz, Chiara Badiali, Róbert Babjak, Óscar Amaro, Sofia Pustova, Rafael Fleming, André Beverido



EuroHPC
Joint Undertaking



DEUCALION

VŠB TECHNICKÁ
UNIVERZITA
OSTRAVA

IT4INNOVATIONS
NÁRODNÍ SUPERPOČÍTAČOVÉ
CENTRUM



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Activities supported by FCT – Fundação para a Ciência e Tecnologia, I.P. through projects UID/50010/2025, PTDC/FIS-PLA/3800/2021, and PLASMA 2023.16184.ICDT

fct
Fundação
para a Ciência
e a Tecnologia

Plasmas permeate the Universe

Most of the observable Universe is in the **plasma** state of matter

- Multiple orders of magnitude of density and temperature mean wildly different physics






Plasmas permeate the Universe

Most of the observable Universe is in the **plasma** state of matter

- Multiple orders of magnitude of density and temperature mean wildly different physics


Can house extreme fields, beyond the Schwinger limit ($\sim 10^9$ T) $\rightarrow$ **Pair creation**

- Exhibit behavior which remains to be understood such as γ -Ray and X-ray bursts




What is a plasma?

A plasma is a state of matter composed of charged particles




What is a plasma?

A plasma is a state of matter composed of charged particles




Each individual particle generates an electric and magnetic field




What is a plasma?

A plasma is a state of matter composed of charged particles




Each individual particle generates an electric and magnetic field





What is a plasma?

A plasma is a state of matter composed of charged particles




Each individual particle generates an electric and magnetic field





What is a plasma?

A plasma is a state of matter composed of charged particles




Each individual particle generates an electric and magnetic field




Realistic first principle simulations are impossible

E and B field are long range forces — **every particle affects every other particle**



Realistic first principle simulations are impossible

E and B field are long range forces — **every particle affects every other particle**




Realistic first principle simulations are impossible

E and B field are long range forces — **every particle affects every other particle**

$$\vec{F}_i = \sum_{j \neq i}^N \vec{F}_{j \rightarrow i}^{EM} \longrightarrow \mathcal{O}(N^2)$$

Even if N were small, this would quickly become impossible to manage




Realistic first principle simulations are impossible

E and B field are long range forces — **every particle affects every other particle**


$$\vec{F}_i = \sum_{j \neq i}^N \vec{F}_{j \rightarrow i}^{EM} \longrightarrow \mathcal{O}(N^2)$$

Even if N were small, this would quickly become impossible to manage




A way of simplifying the problem while maintaining kinematic effects is the **PIC algorithm**:

- Particles are now extended to the size of a cell, instead of point-like
- Particles represent more than one real particle (macroparticles)




A way of simplifying the problem while maintaining kinematic effects is the **PIC algorithm**:

- Particles are now extended to the size of a cell, instead of point-like
- Particles represent more than one real particle (macroparticles)
- EM forces are mediated via the EM fields at the grid points
- Particles deposit current only in the grid points adjacent




A way of simplifying the problem while maintaining kinematic effects is the **PIC algorithm**:

- Particles are now extended to the size of a cell, instead of point-like
 - Particles represent more than one real particle (macroparticles)
 - EM forces are mediated via the EM fields at the grid points
 - Particles deposit current only in the grid points adjacent
- $\mathcal{O}(N_p)$
 $\mathcal{O}(N_c \log(N_c))$




Push particles w/ Lorentz force

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{v} \times \vec{B})$$




Push particles w/ Lorentz force

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{v} \times \vec{B})$$




Particles deposit current in grid




Push particles w/ Lorentz force

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{v} \times \vec{B})$$



Particles deposit current in grid





Push EM fields w/ Maxwell's eqs.

$$\begin{aligned}\frac{\partial \vec{B}}{\partial t} &= -\vec{\nabla} \times \vec{E} \\ \frac{\partial \vec{E}}{\partial t} &= c^2 \vec{\nabla} \times \vec{B} - \vec{J} / \epsilon_0\end{aligned}$$


Push particles w/ Lorentz force

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{v} \times \vec{B})$$

Interpolate EM fields @ particle pos.



Particles deposit current in grid




Push EM fields w/ Maxwell's eqs.

$$\frac{\partial \vec{B}}{\partial t} = -\vec{\nabla} \times \vec{E}$$
$$\frac{\partial \vec{E}}{\partial t} = c^2 \vec{\nabla} \times \vec{B} - \vec{J} / \epsilon_0$$

In order to model quantum effects in strong fields, a Monte Carlo module must run in each push


Nonlinear Compton Scattering (NCS)



Adding strong-field quantum effects via Monte Carlo methods


In order to model quantum effects in strong fields, a Monte Carlo module must run in each push

Nonlinear Compton Scattering (NCS)



$$e^- \longrightarrow e^- + \gamma$$

Nonlinear Breit-Wheeler Pair Creation (NBW)




$$\gamma \longrightarrow e^- + e^+$$

Adding strong-field quantum effects via Monte Carlo methods

In order to model quantum effects in strong fields, a Monte Carlo module must run in each push


Nonlinear Compton Scattering (NCS)



$$e^- \longrightarrow e^- + \gamma$$



Nonlinear Breit-Wheeler Pair Creation (NBW)




$$\gamma \longrightarrow e^- + e^+$$




Adding strong-field quantum effects via Monte Carlo methods

In order to model quantum effects in strong fields, a Monte Carlo module must run in each push


Nonlinear Compton Scattering (NCS)



Nonlinear Breit-Wheeler Pair Creation (NBW)




These processes can be achieved in state-of-the-art **high-power laser** facilities, whose intensities now reach $\sim 10^{23} \text{ W/cm}^2$ (proof of principle done in SLAC [1])



Credits: Photonics Spectra | 1997



Sample quantum rates
for NCS and NBW




Push particles $\frac{d\vec{p}_\gamma}{dt} = 0$

$$\frac{d\vec{p}_q}{dt} = \vec{F}_r + q(\vec{E} + \vec{v} \times \vec{B})$$

Interpolate EM fields @ particle pos.



Particles deposit current in grid




Push EM fields w/ Maxwell's eqs.

$$\frac{\partial \vec{B}}{\partial t} = -\vec{\nabla} \times \vec{E}$$

$$\frac{\partial \vec{E}}{\partial t} = c^2 \vec{\nabla} \times \vec{B} - \vec{J} / \epsilon_0$$

Osiris 4.0



Open-access model

- 40+ research groups worldwide are using OSIRIS
- 300+ publications in leading scientific journals
- Large developer and user community
- Detailed documentation and sample inputs files available

Using OSIRIS 4.0


- The code can be used freely by research institutions after signing an MoU
- Find out more at:
<http://epp.tecnico.ulisboa.pt/osiris>

OSIRIS framework

- Massively Parallel, Fully Relativistic Particle-in-Cell Code
- Parallel scalability to 2 M cores
- Explicit SSE / AVX / QPX / Xeon Phi / CUDA support
- Extended physics/simulation models | **QED rates included via Monte-Carlo module**



Ricardo Fonseca: ricardo.fonseca@tecnico.ulisboa.pt



Laser ionizes the gas

electron-cyclotron maser
instability amplifies EM waves


Radiation recoil is important to
accurately model the phenomena

1 MCPUh




T. Silva *et al.* preprint: arXiv:2412.06783

Particles with higher energy also collapse
to a ring distribution after radiation
Universal attractor for magnetized plasmas



MCPUh


Bichromatic lasers can create and separate pairs



Laser fields can accelerate pairs directly

Laser can push particles to highly relativistic particles over small distances

Direct Laser Acceleration (DLA) uses a plasma to keep particles interacting longer




Muons are essentially heavy electrons

Unstable ($\sim 1 \mu s$)

In relativistic Physics, time dilates

$$t_{lab} \rightarrow \frac{t_{muon}}{\sqrt{1 - \frac{v^2}{c^2}}}$$

Laser Wakefield Acceleration (LWFA) can accelerate the muons quickly enough to become highly relativistic and “stable”



The PIC loop has been a successful tool for unraveling effects of Plasma Physics

Additional Monte Carlo modules can add quantum effects to the description

These simulations pave the way for future experiments to determine the Physics of pair plasmas and astrophysical environments

Laser and plasma based acceleration are an important asset for future particle colliders

