SM&FT 2022

THE XIX WORKSHOP ON STATISTICAL MECHANICS AND NON PERTURBATIVE FIELD THEORY

Frontiers in Computational Physics

Computational Challenges in Lattice QCD

Francesco Sanfilippo, INFN RM3

Istituto Nazionale di Fisica Nucleare

20 December 2022

Outline

1 - Lattice QCD, why is it so though?

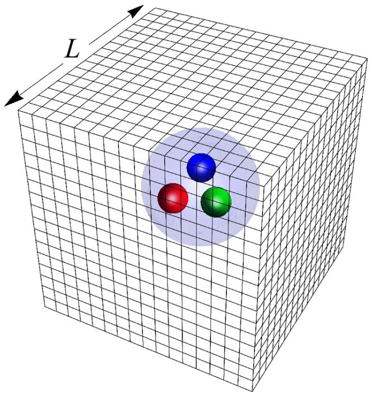

2 - Recent achievements of LQCD

3 - <u>Open challenges</u>

LATTICE QCD SIMULATIONS

First principle simulation of strong interactions

Quantum Chromodynamics on a Lattice

4D (spacetime) with $O(10^{10})$ degrees of freedom

Hybrid Monte Carlo + Molecular Dynamics simulations

Numerical solution of the discrete Dirac Equation
(partial derivative equation → large sparse matrix)

A long list of scientific achievements: reconstruction of the hadron spectrum, thermodynamics of strong interactions, calculation of hadronic vacuum polarization...

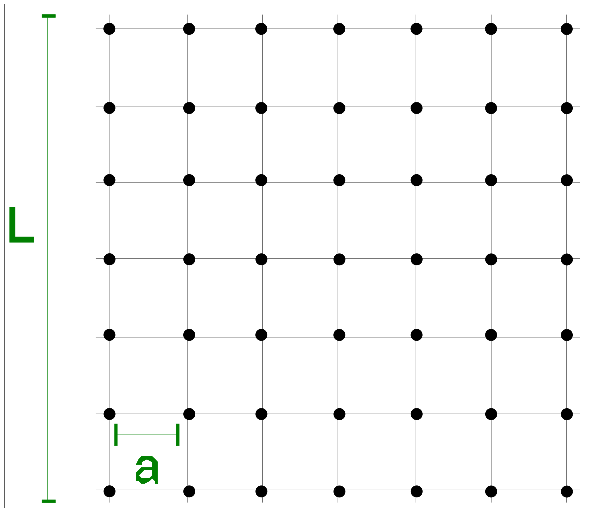

Typical Lattice QCD Simulation/Measurement Scheme

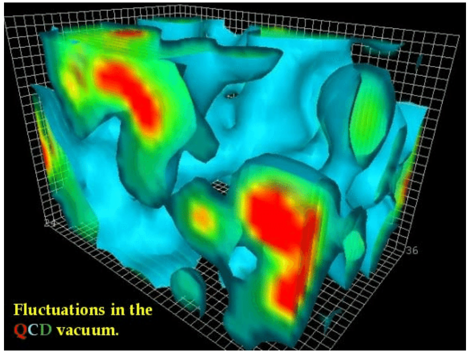

Producing O(100-1000) “**configurations**” of gluonic fields.

TOOL: Molecular Dynamics + Monte Carlo to evolve configurations of **gluonic fields,** the background in which quark particles “move”.

RESOURCES: 1 configuration ~ O(1-50 GB data) ~ 1 day of simulation on **O(5000) cores**.
Hundreds of MCorehours gained through national, European & worldwide supercomputing calls.
Similar in spirit to the production of collisions at particle accelerators (tens of PB of data).

A **few large collaborations** (“big experiments”) with important difference on the discretization.
Multi-year “Runs”, with statistics & systematics improving in time.

Typical Lattice QCD Simulation/Measurement Scheme

Propagating O(100) quark on the gluon field backgrounds, take some algebraic combination:

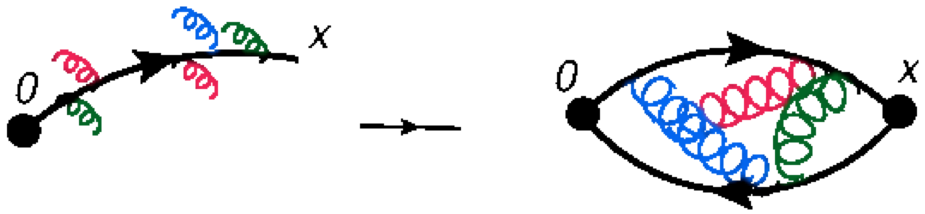

TOOL: Numerical solution of Dirac Equation, tensor algebra to manipulate many spin and color degrees of freedom.

RESOURCES: 100 propagator ~ 1 hour of simulation on O(5000) cores/few GPUS.
Similar in spirit to data analysis of collision events. “Smaller” national, European calls.

PORTING: Several efficient numerical solvers for CPU & GPU, tensor algebra more tricky.

More collaborations of smaller scale with more specific problems & more code platforms.

GOOD: the critical task is the same for everybody, solved thanks to efficient libraries.
BAD: the remaining part of the code can still have a significant cost and is not homogeneous.

Massive parallelization scheme: O(100) NODES with O(50) CORES with O(16) AVX REG

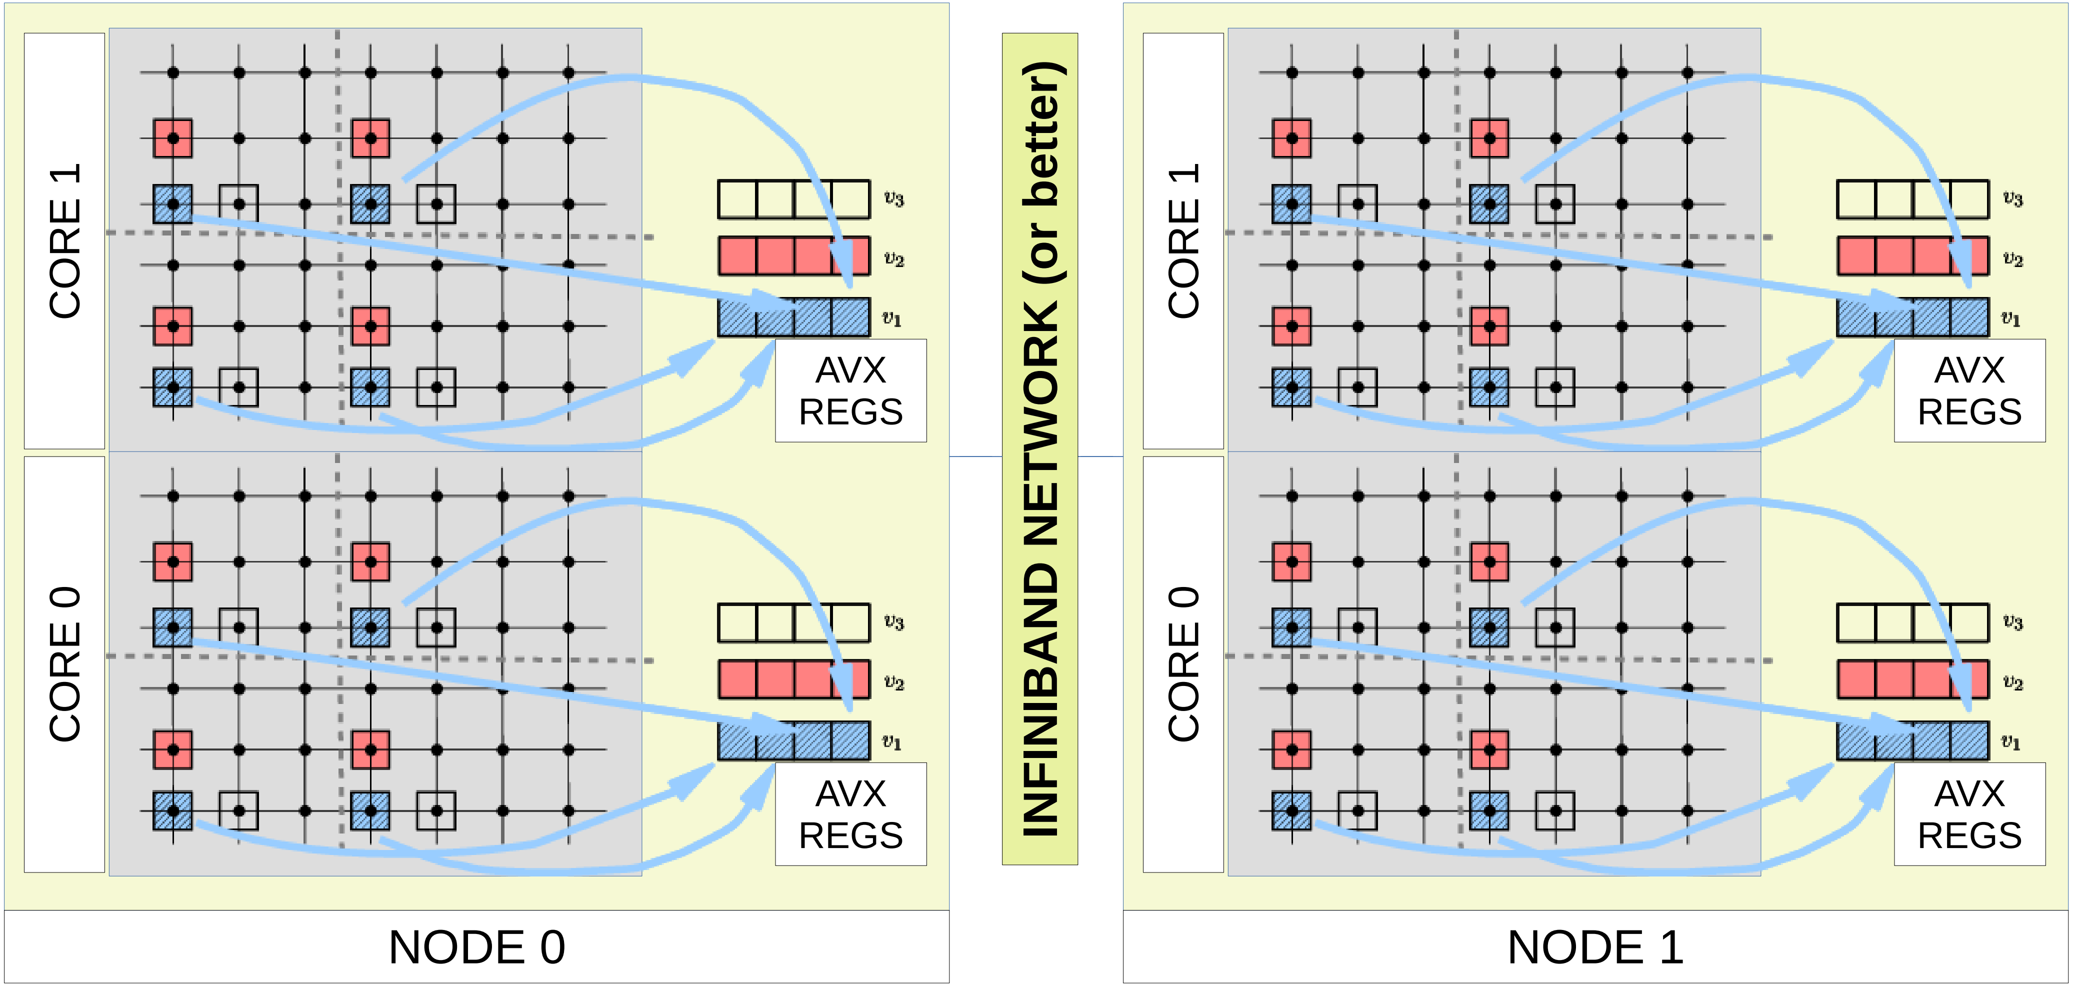

In some dedicated architectures (APE, BG/Q) even the network adapters form a 4D network!!!

Why is Lattice QCD so computationally demanding?

#1 Issue: Quark masses dependency

- Simulation cost: rapidly grows as quark masses are lowered
- Early solution: quenching = drop virtual pair contributions from partition function

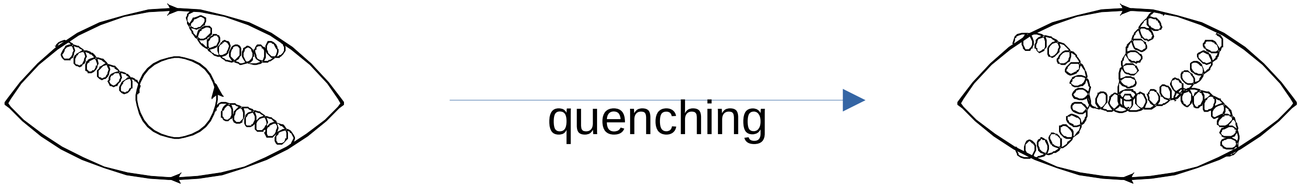

- Intermediate solution: consider unphysical light quarks $M_\pi \sim 300 \div 500 \text{ MeV}$
- Nowadays: many collaborations (CP-PACS, FERMILAB/MILC, BMW, RBC/UKQCD, TMLQCD...) use pions of physical mass

Why is Lattice QCD so computationally demanding?

<u>#2 Issue</u>: Lattice size dependence

Small UV cut-off to resolve heavy hadrons

$$a \ll 1/M_H$$

Large IR cut-off to accommodate pions

$$L \gg 1/M_\pi$$

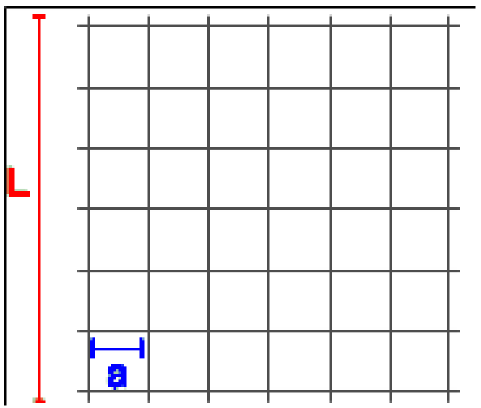

Therefore one needs to take $L/a \gg M_H/M_\pi \sim 20$

$$\#points = (L/a)^3 \times T/a = 64^3 \times 128 \div 128^3 \times 256$$

$$\#internal\ d.o.f \sim 100$$

Total number of degrees of freedom:
$10^8 \div 10^{10}!$

State of the art

Kenneth G.Wilson prophecy (father of Lattice QCD in 1974)

Thirty years will be necessary for computational resources and algorithms to reach proper maturity [**Lattice conference 1989**]

Nowadays (~thirty years later)

- Physical light quarks and large volumes $\gtrsim (6\,\mathrm{fm})^3$
- Simulations performed at several lattice spacings
- Isospin & Electromagnetic corrections accounted

PRECISION ERA!!!

What helped these improvements?

Increase in computing power

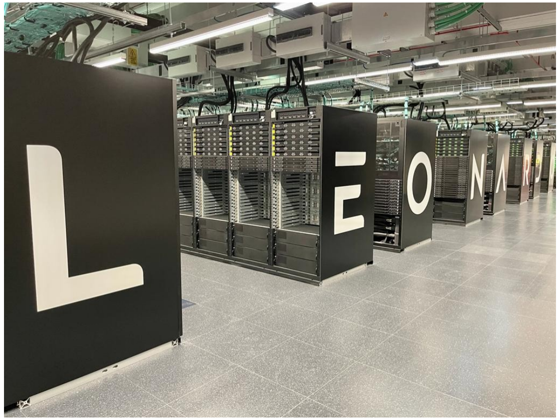

Conceptual developments

- Improved regularizations of LQCD (Stout smearing, Dynamic Clover, Twisted Mass…)
- Better understanding of behavior of Monte Carlo simulations

Algorithm breakthroughs

- Multiple timescale Molecular Dynamic integrators
- Deflation, Multigrid, Domain Decomposition solvers, etc.

Multigrid solvers

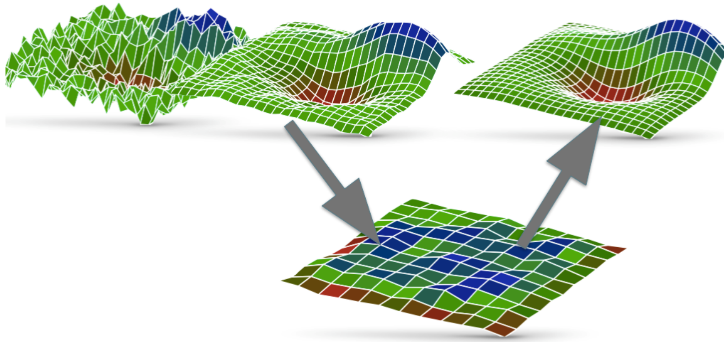

<u>An old idea</u>: treat separately the coarse scale and fine scales of the lattice...

<u>Break-through</u>: Data-Driven (a.k.a. Adaptive) Algebraic Multigrid Methods [A.Frommer et al, 2011]

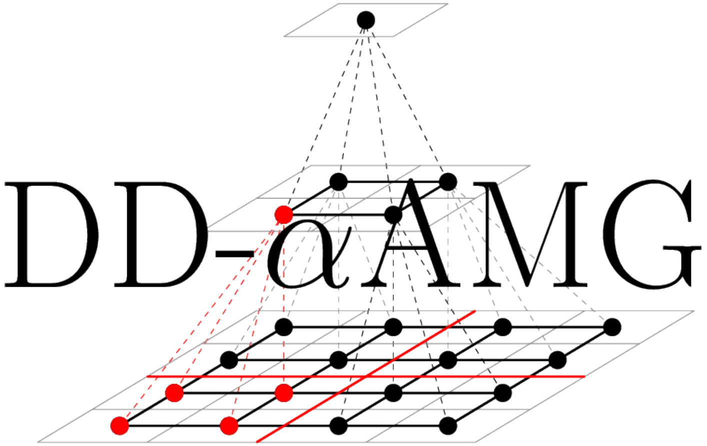

NVIDIA CUDA

...nicely implemented for GPU architecture [M.A.Clark et al. 2018]

RECENT ACHIEVEMENTS

of the lattice QCD community

Hadron spectrum including QED

Neutron-Proton mass difference

[BMW coll. Science (347) 2014]

$$M_n - M_p = (1.293\,\mathrm{MeV})^{exp}$$

$$M_n - M_p = (1.51(16)(23)\,\mathrm{MeV})^{lat}$$

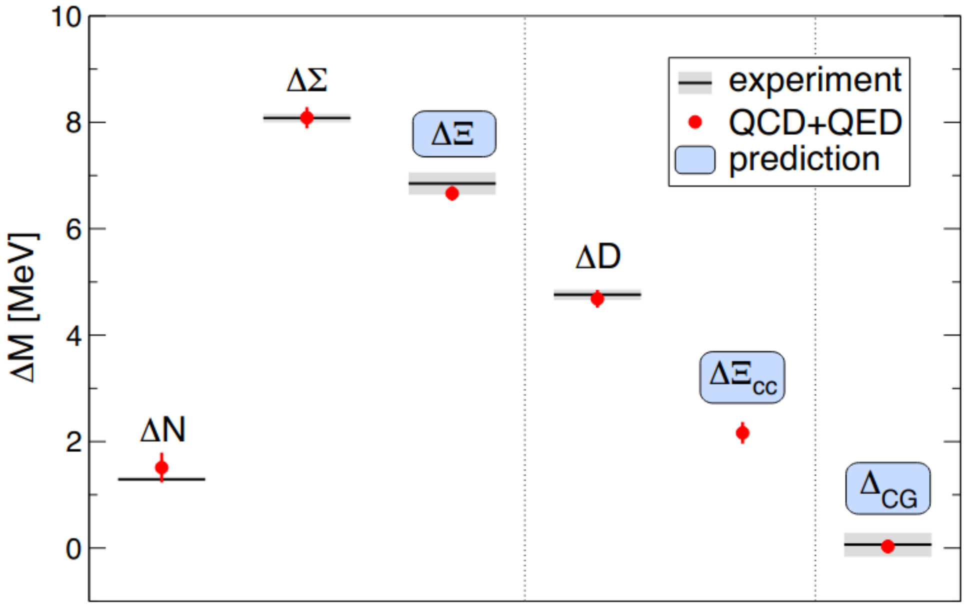

Sub % accuracy in the reconstruction of the fine-grained structures of hadron spectrum

EoS and magnetic properties of QCD

Equation of State, also at finite density
[S.Borsanyi et al, PRL (2021) 126]

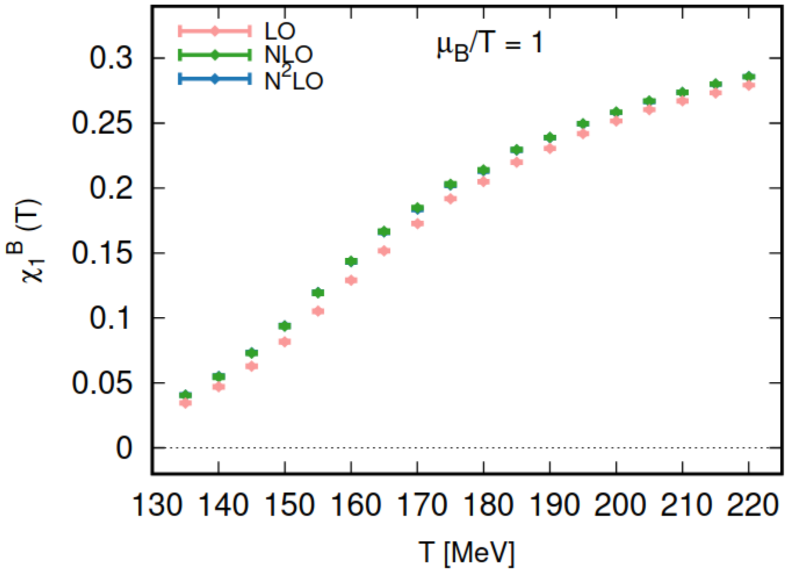

Magnetic Susceptibility of QCD matter
[C.Bonati et al, PRL (2013) 111]

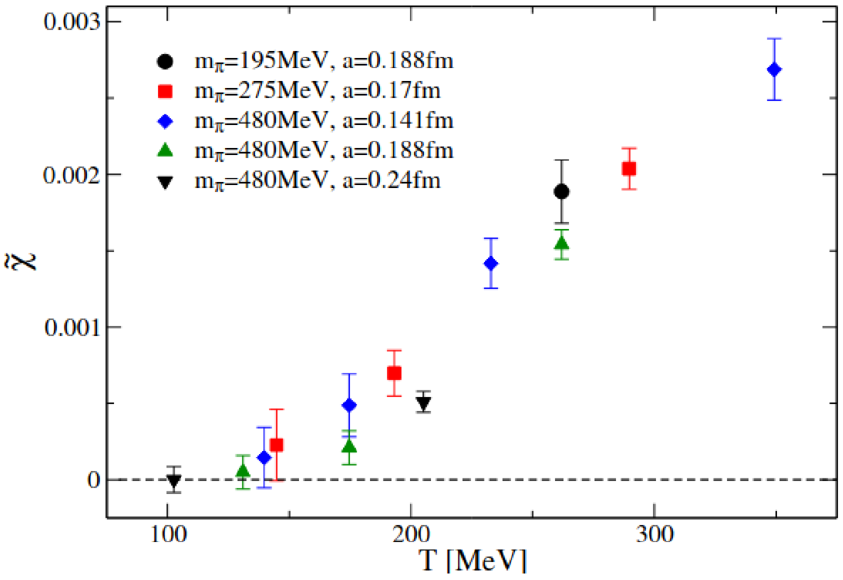

See talk by **J.Guenther @ 12:00 today**

Window contribution to the Hadronic vacuum polarization of muon

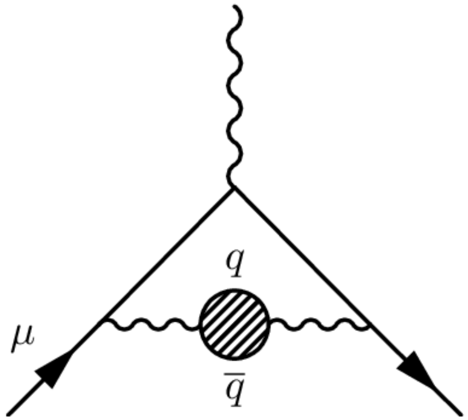

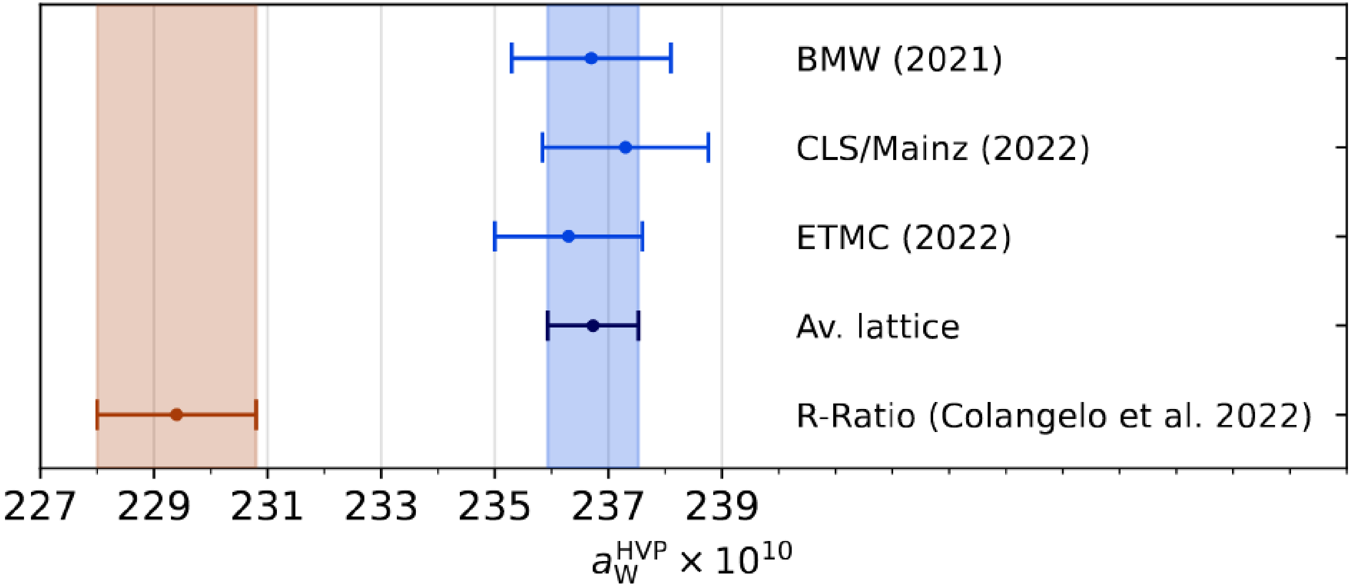

Several σ discrepancy: “**The new g_μ-2 puzzle**” [cfr L.Di Luzio et al., Phys.Lett.B 2022]

OPEN CHALLENGES

1) Topological freezing
2) Continuation from Euclidean to Minkowsky
3) Signal to noise ratio deterioration

BONUS: technological issues

CHALL
1/3

Topological freezing

See also MP Lombardo talk @12:30

As the continuum limit is approached, simulations don’t tunnel properly topological sectors

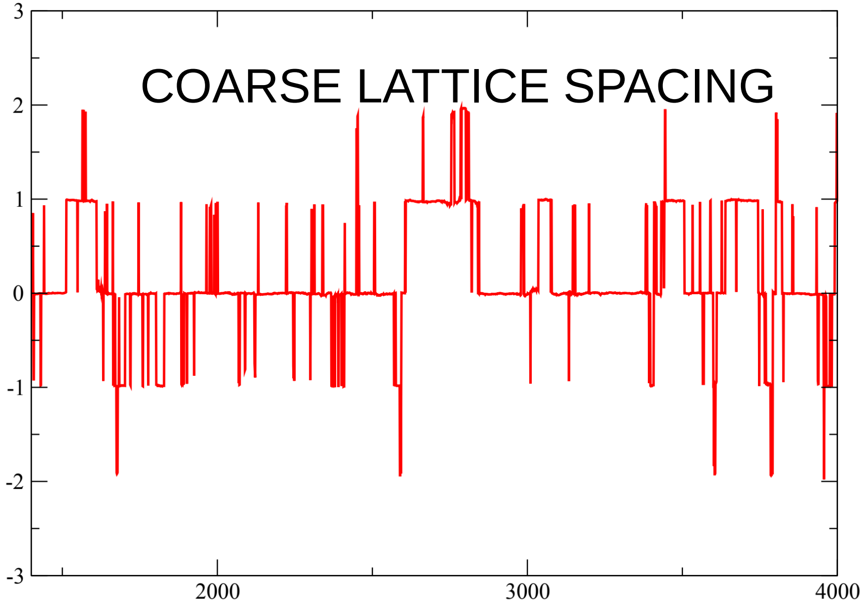

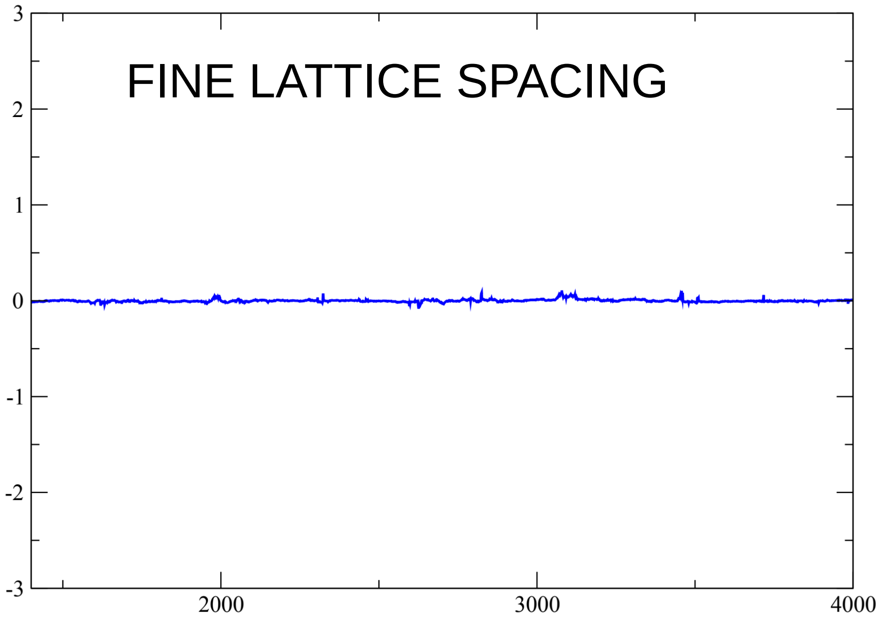

- First observed 20 years ago [L. Del Debbio, H. Panagopoulos and E. Vicari, JHEP 2002].
- Well studied since more than 10 years [M. Luscher, PoS LATTICE 2010 (2010)]

Emergence of topological barriers

<u>Topological sectors</u>

get more and more separated as one proceeds towards the continuum limit

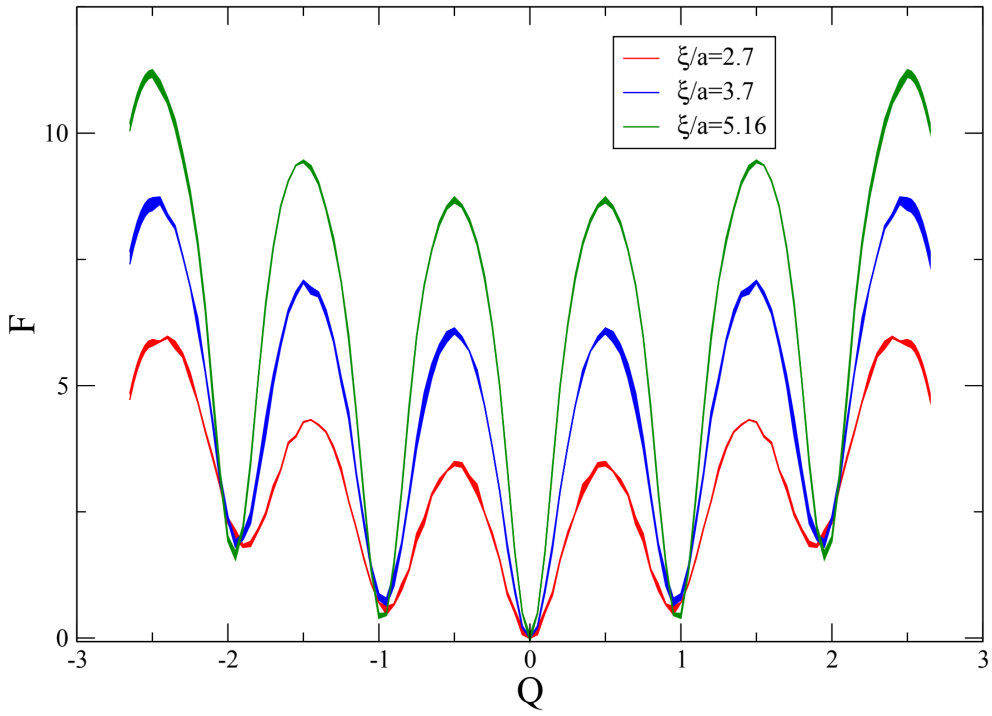

A twofold problem

- Phenomenological issues:
  - Thermodynamics of the early universe,
  - Cold and Hot Axion phenomenology,
  - Singlet particle properties.
- Simulation issues: how to simulate all topological sectors with the proper weight?

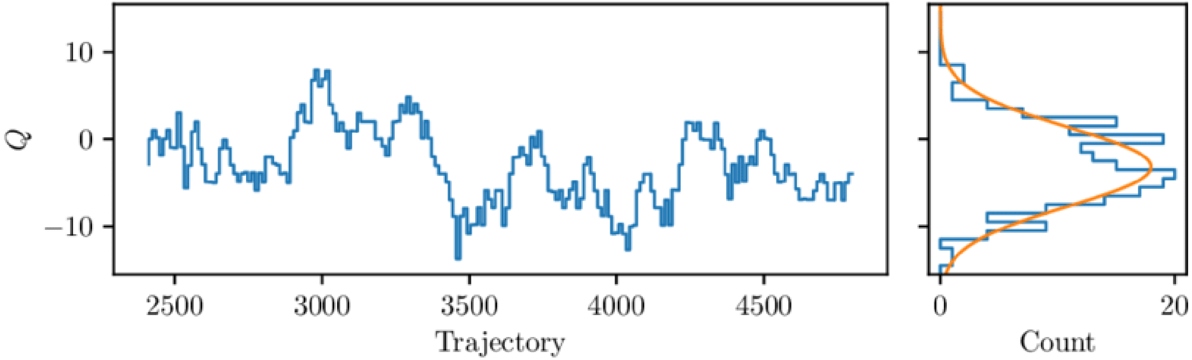

Various solutions proposed...

Open boundary conditions [M. Luscher and S. Schaefer, JHEP 1107 (2011)]

bypass the quantization of the topology, avoiding to close one of the boundaries

✓ Topological objects free to flow in and out from the lattice ✗ Boundary effects?

Master field simulations [P.Fritzsch et al., PoS Lattice 2021]

replace classical (Markov chain) ensemble with a single master-field

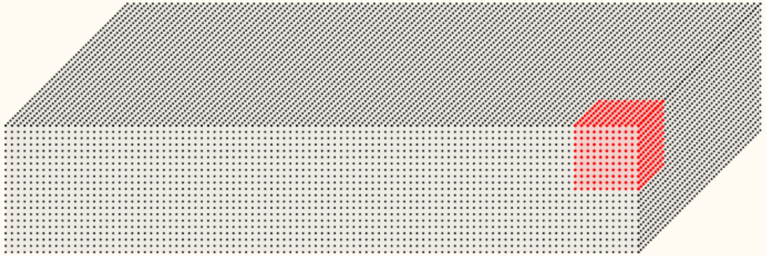

$$\frac{V_4^{\mathrm{mf}}}{V_4} = \prod_{i=0}^{3} N_i \simeq 100 - 1000$$

✓ Self-averaging ? Thermalization? Hergodicity? Under investigation...

Parallel tempering

Simulate several temperature/boundary conditions simultaneously, swapping 1 physical and N “eased”

✓ Works in pure gauge

? Fermionic determinant?

? Fine tuning of the tower of simulations

? change of simulation paradigm

[C.Bonnanno, C.Bonati, M.D’Elia, JHEP 2021]

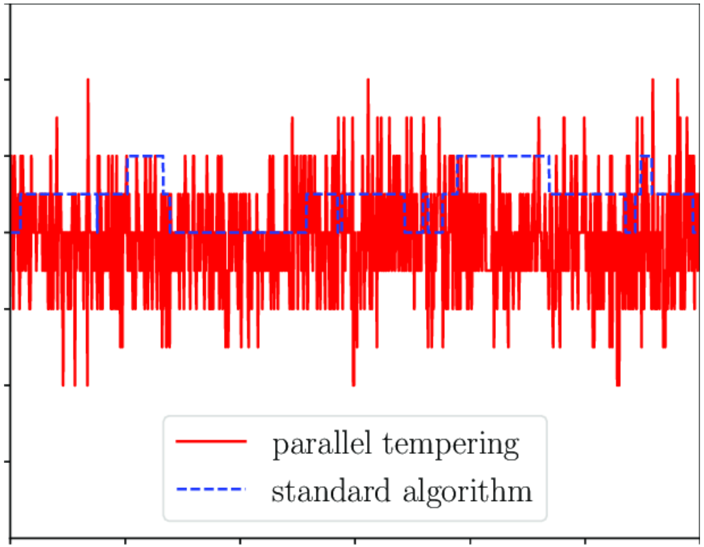

See also presentation by **C.Bonanno @17:35 Monday**

Metadynamics

Self-constructed bias potential to contrast the development of the topological barriers

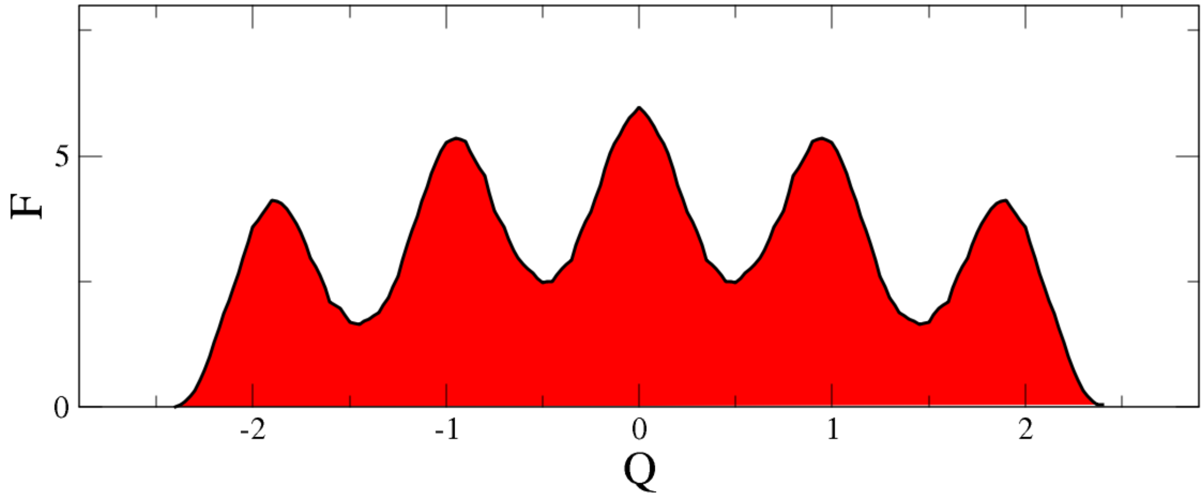

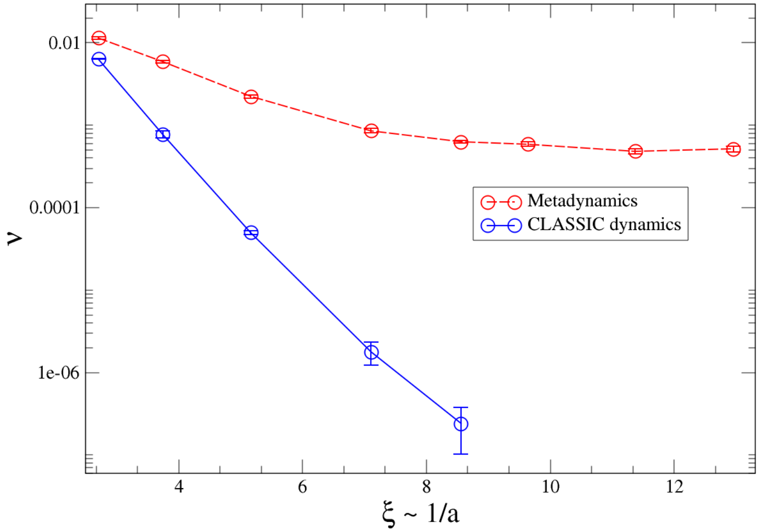

First studied for CP(N) [A.Laio, G.Martinelli, FS, JHEP 2016]

Metadynamics

- Currently explored in pure gauge simulations

[T.Eichhorn, C.Hoelbling, P.Rouenhoff, L.Varnhorst, PoS Lattice 2022]

V Self-adaptative, effective
? Reweighting, overhead

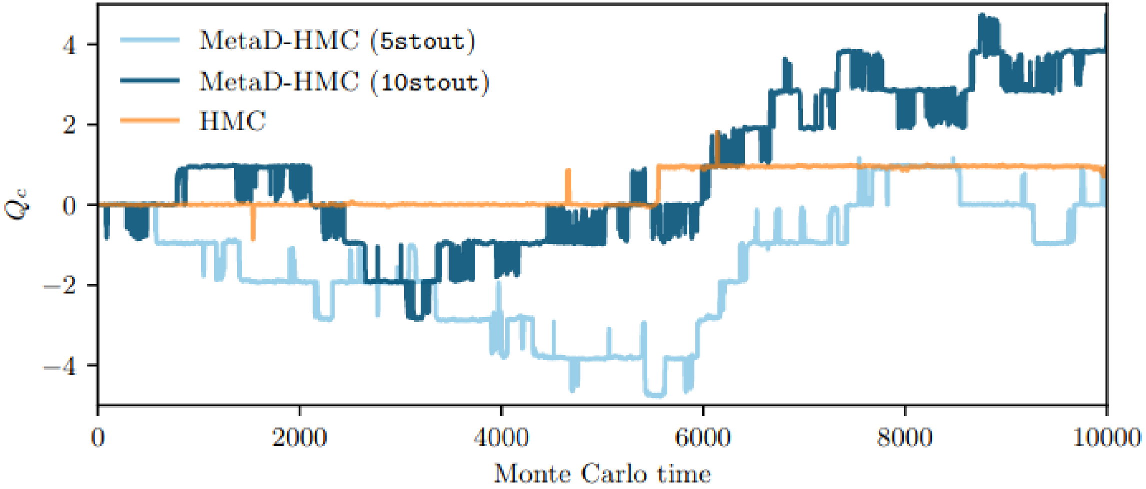

- “Static” potential studied in full QCD simulations

[C.Bonati et al., JHEP (2018) 170]

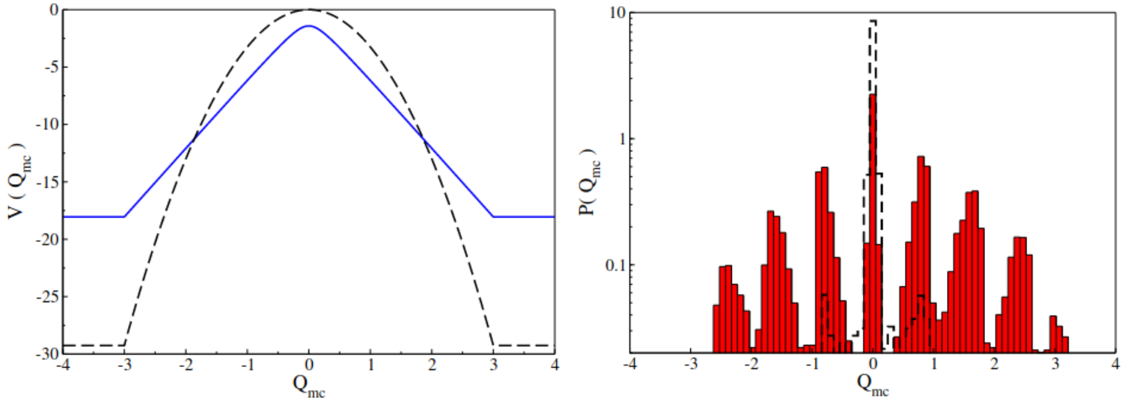

Analytic continuation from Euclidan to Minkowsky

Matrix elements are related to correlators by inverse Fourier transform

$$C(t) = \int dE \exp[-iEt] \langle P|O|P' \rangle$$

in Minkowsky time (real time), but Lattice calculations are carried out in Euclidean time, $t \to -i\tau$ which means solving inverse Laplace transform

$$C(\tau) = \int dE \exp[-E\tau] \langle P|O|P' \rangle$$

in presence of finite sample (few tens of lattice sites)
and noise (statistical fluctuation due to finite sample size)

A number of limitations to lattice

- Decay of hadron particle: the lowest lying state in each channel dominates the correlator (Maiani-Testa no go theorem)

$$C(\tau) = \int dE \exp[-E\tau]\langle P|O|P'\rangle \xrightarrow{\tau\to\infty} \exp[-E_0\tau]\langle P_0|O|P_0\rangle$$

which means typically one can study only single particle states!

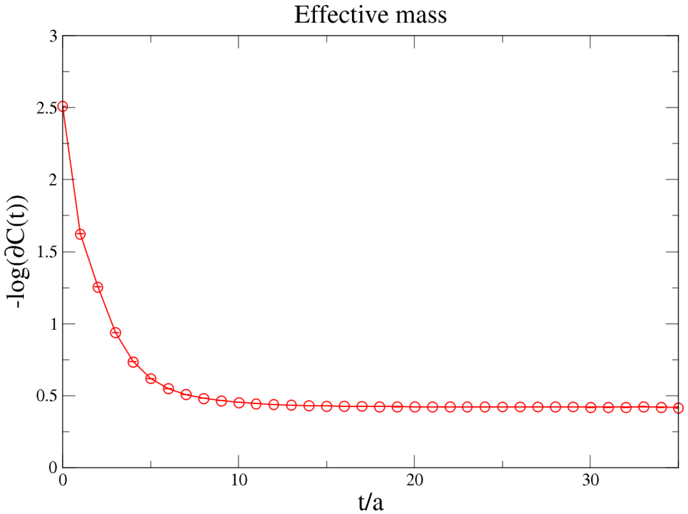

- Real time dynamics cannot be studied either (e.g. conductivity)

Lellousch-Luscher formalism

Reconstruct matrix elements from energy shift in a finite box

$$E(P_1, P_2) = E(P_1) + E(P_2) + \Delta E(L)$$

particles interact due to finite box L.

Quantization condition relates
<u>energy, box size and scattering lengths</u>
(assuming partial wave expansion etc).

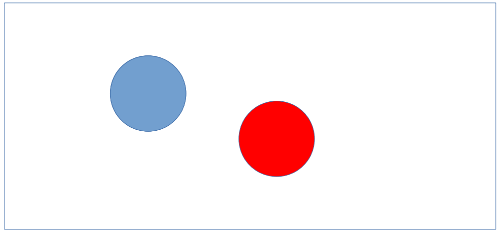

- Needs to know the quantization condition for multiple particle in a box
- Well studied for 2-body decays

  e.g: $K \to \pi\pi$

- Beyond 2-particles is
  <u>much more involved!</u>

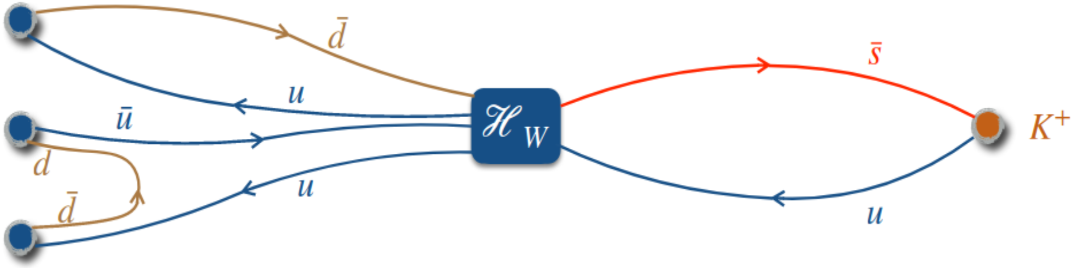

Smooth the problem [since ~2020]

$$C(\tau) = \int dE \exp[-E\tau]\rho(E)$$

Ease the inverse Laplace problem smoothing the corners:

$$C_\sigma(\tau) = \int dE \exp[-E\tau] \underbrace{\int dE_0 \sigma(E, E_0)\rho(E_0)}$$

Solve for the convolution of the original solution: $\rho_{\sigma(E)}$

- Old methods [G.Backus, F.Gilbert, Geophys. J. Int. 1968] to new grounds
- The smoothing might be extrapolated away, or kept and incorporated with the experimental comparison.

A number of recent applications

"Variations on the Maiani-Testa approach and the inverse problem"

[M. Bruno and M.T.Hansen, JHEP 2021]

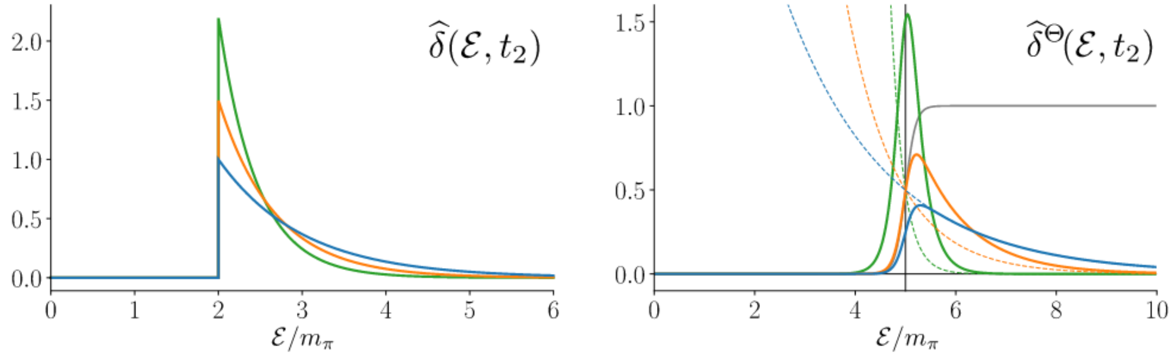

"Lattice study of EM conductivity of quark-gluon plasma in magnetic field"

[N.Astrakhantsev et al., PRD 2021]

see M.Naviglio @19:00 Monday

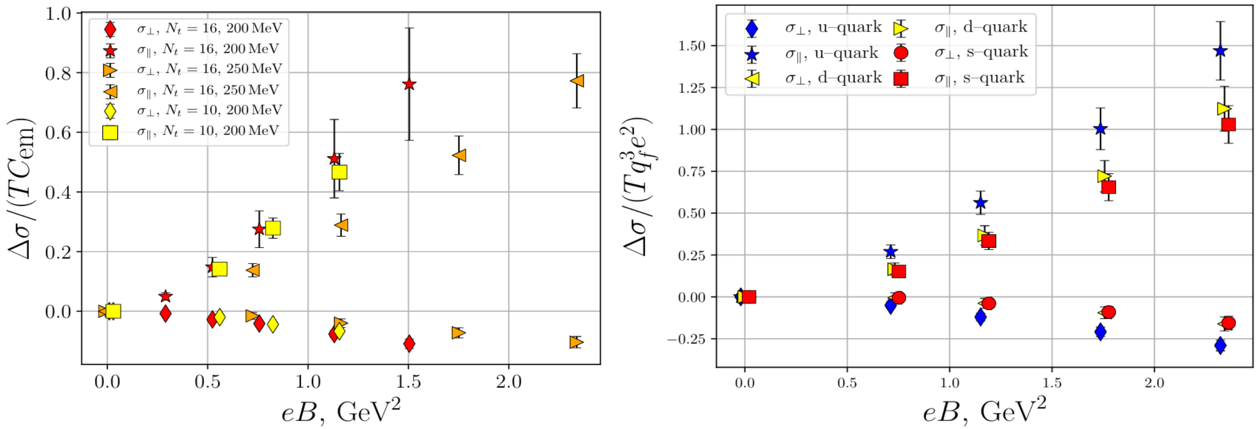

Inclusive Semileptonic decays of heavy mesons

[P.Gambino et al, JHEP 07 (2022)]

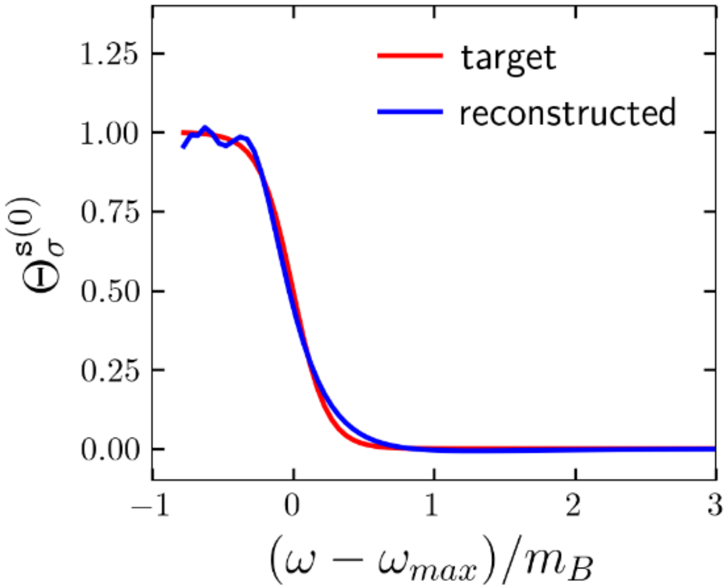

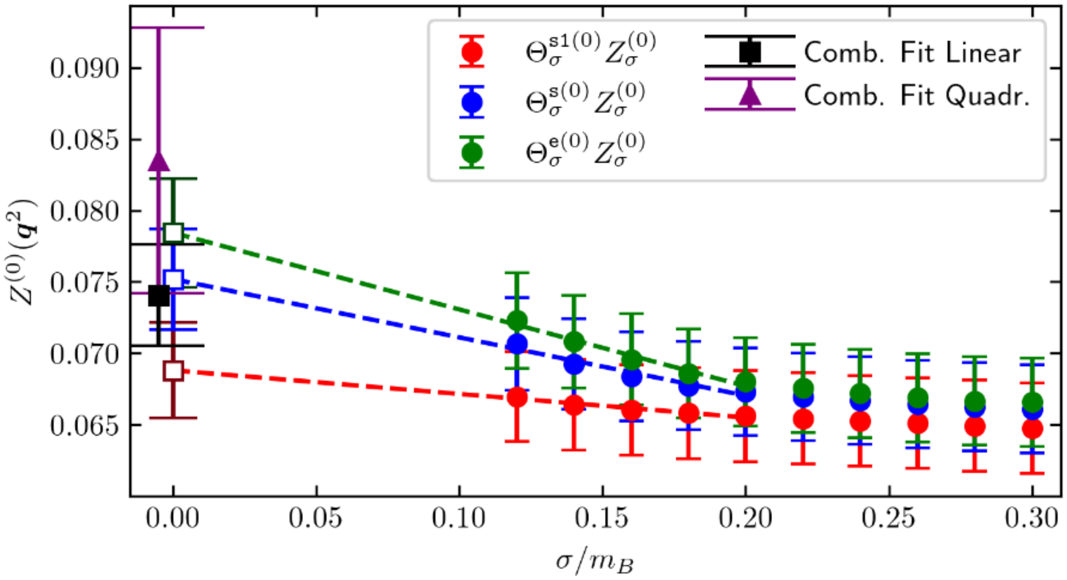

[See also contribution of **A.Smecca, @18:35 today**]

R-Ratio of $e^+e^- \to hadr$ scattering

[ETM collaboration, arXiv:2212.08467]

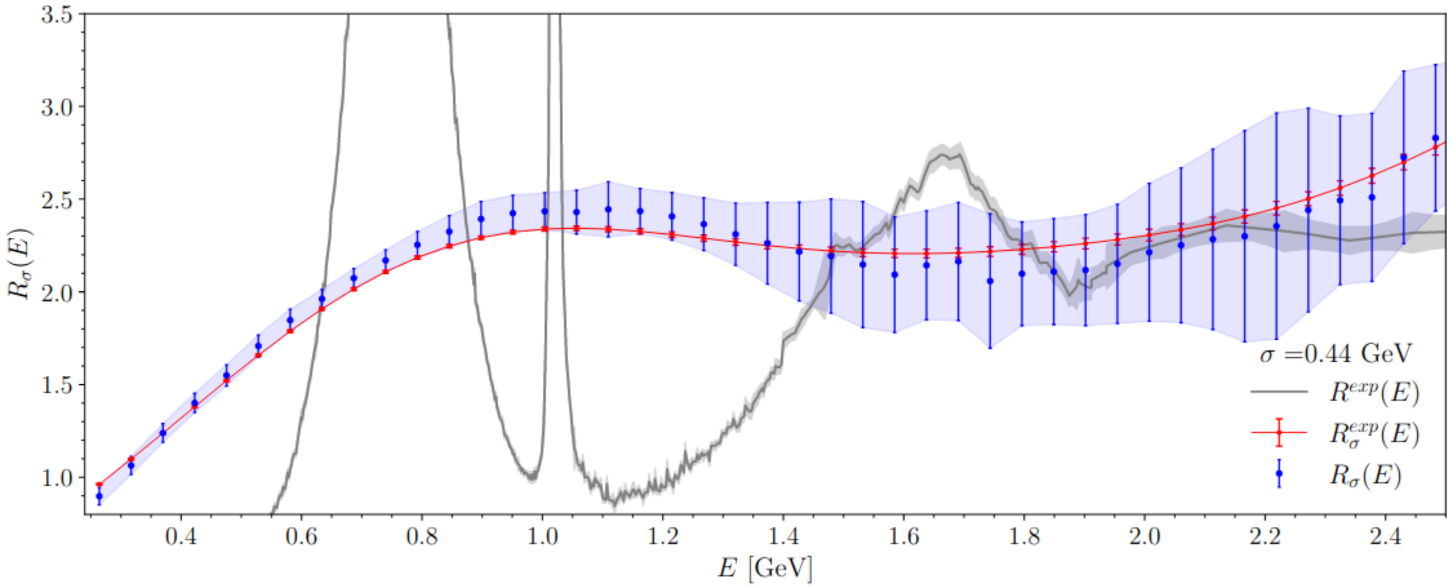

With more statistics, longer euclidean time

→ <u>finer resolution</u>, more interesting phenomenology

...AND MORE!!!

Signal/noise deterioration

Correlation functions decays as: $C(t) = \langle O(t) | O^\dagger(0) \rangle \propto \exp(-Et)$

Noise (variance) decays as: $\sigma^2(t) = \langle O(t) | O^\dagger(0) \rangle^2 \propto \exp(-E't)$

ISSUE: It occurs [Parisi, Lepage, '80] **that** $E' < 2E$

such that:

$$\frac{S}{N} = \frac{C(t)}{\sqrt{\sigma^2(t)}} \xrightarrow{t\to\infty} 0$$

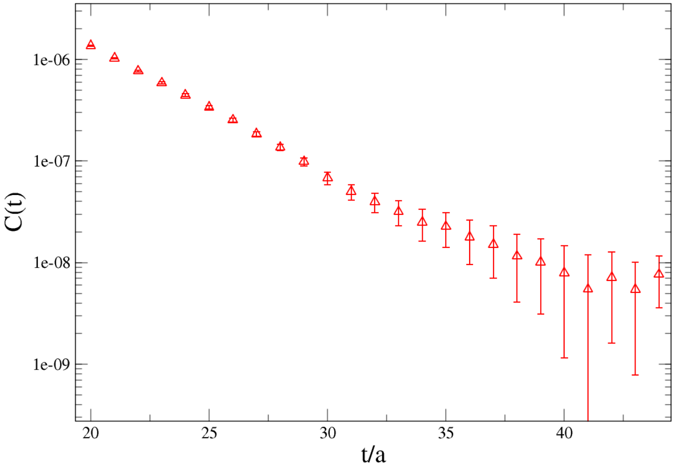

Problems more severe when:

- many quarks are involved
- momenta is transferred
- different flavors enters

Naive solution: brute force

Solutions

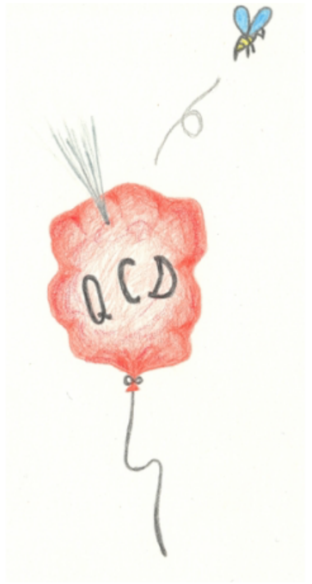

Eigenspace approach: use eigenvectors to compute exactly/approximately part of the solution
→ *Deflation, All-Mode-Averaging, etc.*

Source choice: use stochastic estimators with a reduced overlap with the noise, to reduce the scaling prefactor
→ *Dilution of the source, Hadamard vectors, etc.*

Multilevel integrators: update more frequently long distance factorizing different domains
✓ Fix the poor scaling of signal/noise ratio
? Affordable? Under scrutiny...

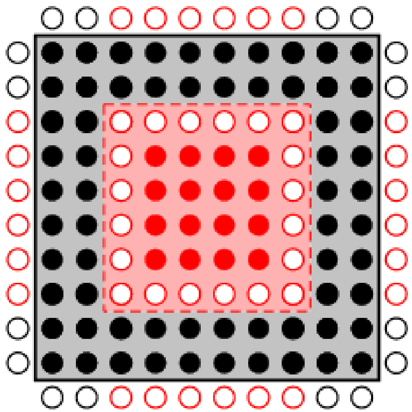

TECHNICAL CHALLENGES

The GPU paradigm

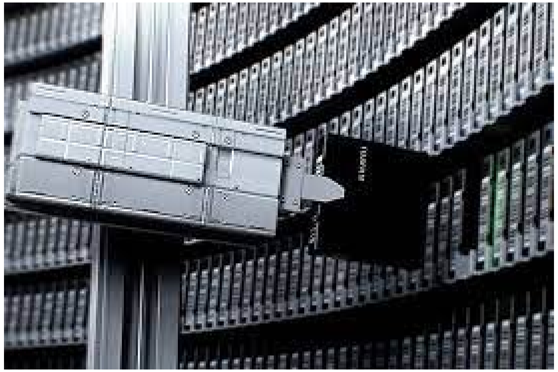

How to store all this data

How to port to GPU (and keep it general)?

… **many approaches** around!!!!

QUDA LIBRARY – M.Clark et al., since 2009

Heterogeneous collection of solvers for the Dirac equation, with a number of modern and adaptative algorithms, supporting various lattice QCD regularizations.

Open Source, actively developed by NVIDIA, through a strong group of former lattice QCD researcher. Makes use of all edge cutting GPU technology available.
Employed by several lattice QCD groups around the world.

PRO
Extremely well performing for the supported tasks

CON
Cannot perform all typical lattice QCD tasks (no full HMC).
Extremely difficult to adapt to different tasks from supported.
No portability.
(+ crazy interface & terrible documentation...)

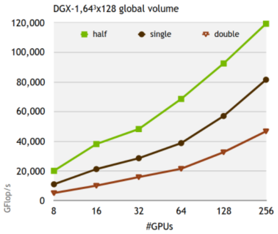

GRID LIBRARY – P.Boyle et al. since 2015

C++ framework for the calculation of correlation functions & full HMC simulations (?)
Targetting a number of Lattice QCD regularization, easy to extend, efficient

Frontend: modern C++ 11 with a bit of metaprogramming + Python interface
Backend: supporting several archtectures: Cuda, HIP, OpenMP, etc (kernel abstraction)

PRO

Intuitive, multiplatform, reasonably efficient on all platforms, relatively lightweight, adopting optimal memory layout transformations to efficiently use the resources.

CON

Reduced community (mostly US/UK oriented), limited expertise available in Italy
Engaging with the developers proved not easy in the past.

MILC software stack from USQCD software stack

Large software stack for HMC simulations & measurements.
Mainly used in the US & UK, a few users in Italy.

PRO

- Large community (in the US),
- Multiplatform.

CON

- Incomplete GPU support (multigpu?)
- Documented? Mhhh...
- Not trivial to setup (quite bloated code),
- Targeting a subset of the lattice interest.

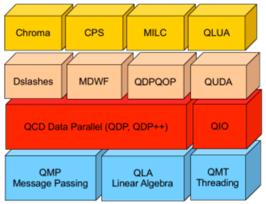

NISSA LIBRARY – F.S. since 2011

In use from two major collaborations (LQCD123, PISA group)
Employed within several PRACE projects (PRA17-4394, PRA20-5171, PRA22-5171…)

Frontend: C++ 11 (envisaged migration to pure abstract C++17 metaprogrammed)
Backend: kernel abstraction, linked to several external libraries (including QUDA)

PRO

- “Large” user platform in Italy.
- Targeting different Lattice QCD regularization, multigpu & multithread.

CON

Missing the memory layout transformations to support more efficiently GPU & vector CPU for non-critical but important tasks

Storing the data

Collab	Public	ILDG	#ens	#cfg	storage (TB)
FASTSUM	1	1	25	22k	40
OpenLat	1	2	8	10k	30
MILC	1	0	>25	75k	1000
JLab/W&M/LANL/MIT/OLCF/Marseille	0	0	13	105k	2000
JLQCD	1	2/3	>230	60k	20
ETMC	1	2/3	21	100k	2500
TWEXT	1	1	60	50k	26
PACS	1	2/3	3	100	60
RBC-UKQCD	1	0	41	20k	500
HotQCD	1	2	58	15M	2250
CLS	1	2	>60	130k	1000
CLQCD $T = 0$	1	1	10	5k	14
CLQCD $T > 0$	1	1	28	150k	120
HAL QCD	1	2	1	1.4k	70
QCDSF-UKQCD-CSSM	1	2/3	60	90k	300

Tens of Petabytes of gluon configurations stored!

ILDG and extensions

- Valuable assets
- Open data? FAIR policy
- Easy of public access
- Backups

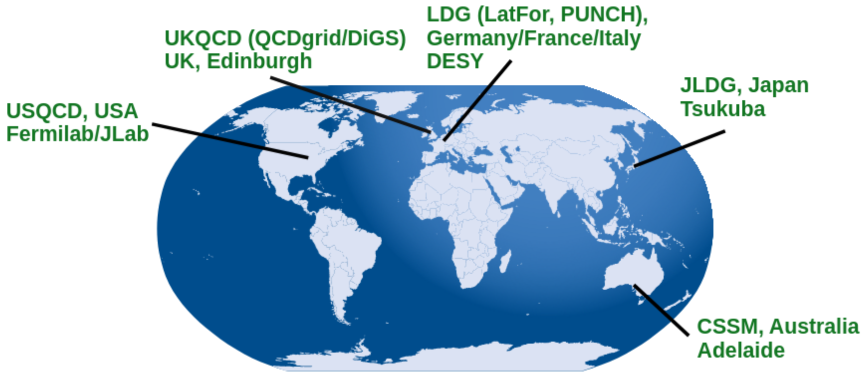

To be addressed by ILDG 2.0:

“The International Lattice Data Grid — towards FAIR Data”

[F.Karsch, H.Simma and T.Yoshie, POS lattice 2022]

Room for help from CNAF & ICSC (data lake)

Conclusions

- After decades of efforts, LQCD has entered **precision era** with supercent accuracy on many quantities.

- Many **new ideas** & algorithms allow to start exploring **new aspects** of strong interactions.

- Multy-year runs, tens of petabyte of storage, ever changing architecture pose still **big challenges!**

…..STAY TUNED!!!!

THANKS!!!