7 November 2023

HPC for Elementary Particle Physics

Antonio Rago

QUANTUM THEORY CENTER

SDU

Particle Physics …

QUANTUM THEORY CENTER
SDU

- The question: What are the fundamental constituents and fabric of the universe and how do they interact?

Many open questions

What are the laws of physics operating in the **early Universe**?

What are the **fundamental particles** and fields?

What is the **nature of space-time**?

What is the nature of **dark matter** and **dark energy**?

Why is there **more matter than antimatter**?

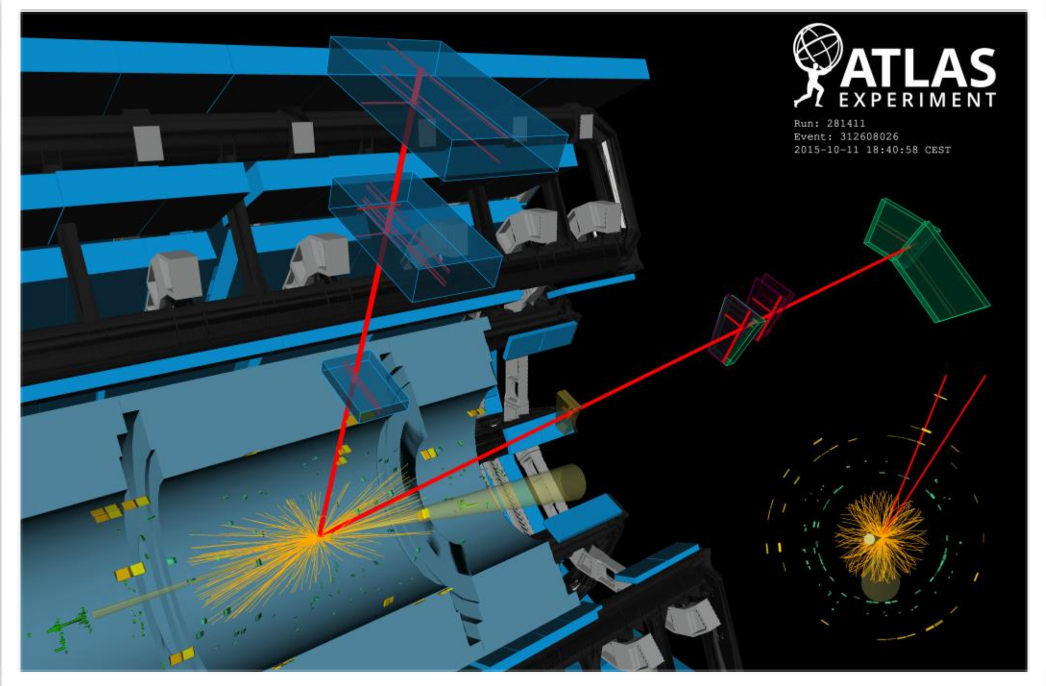

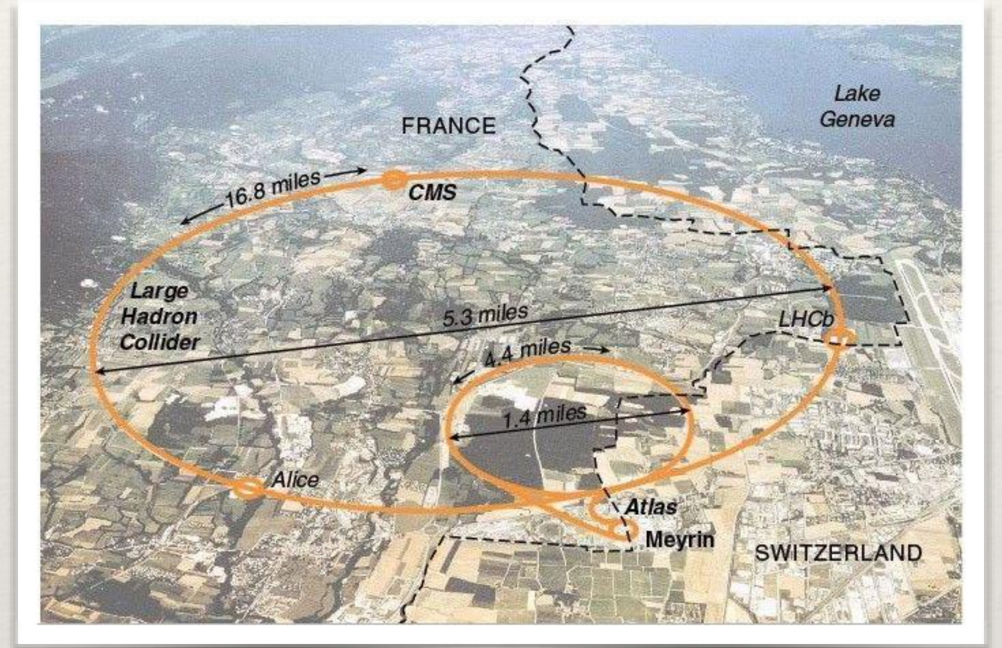

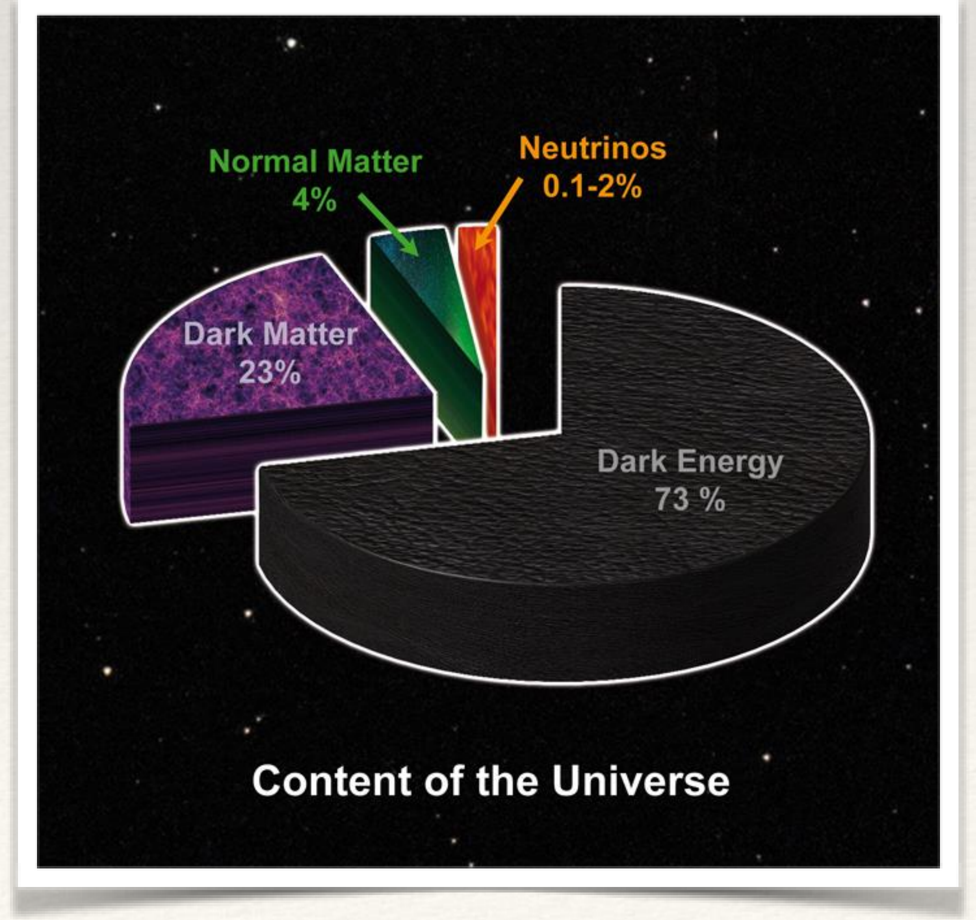

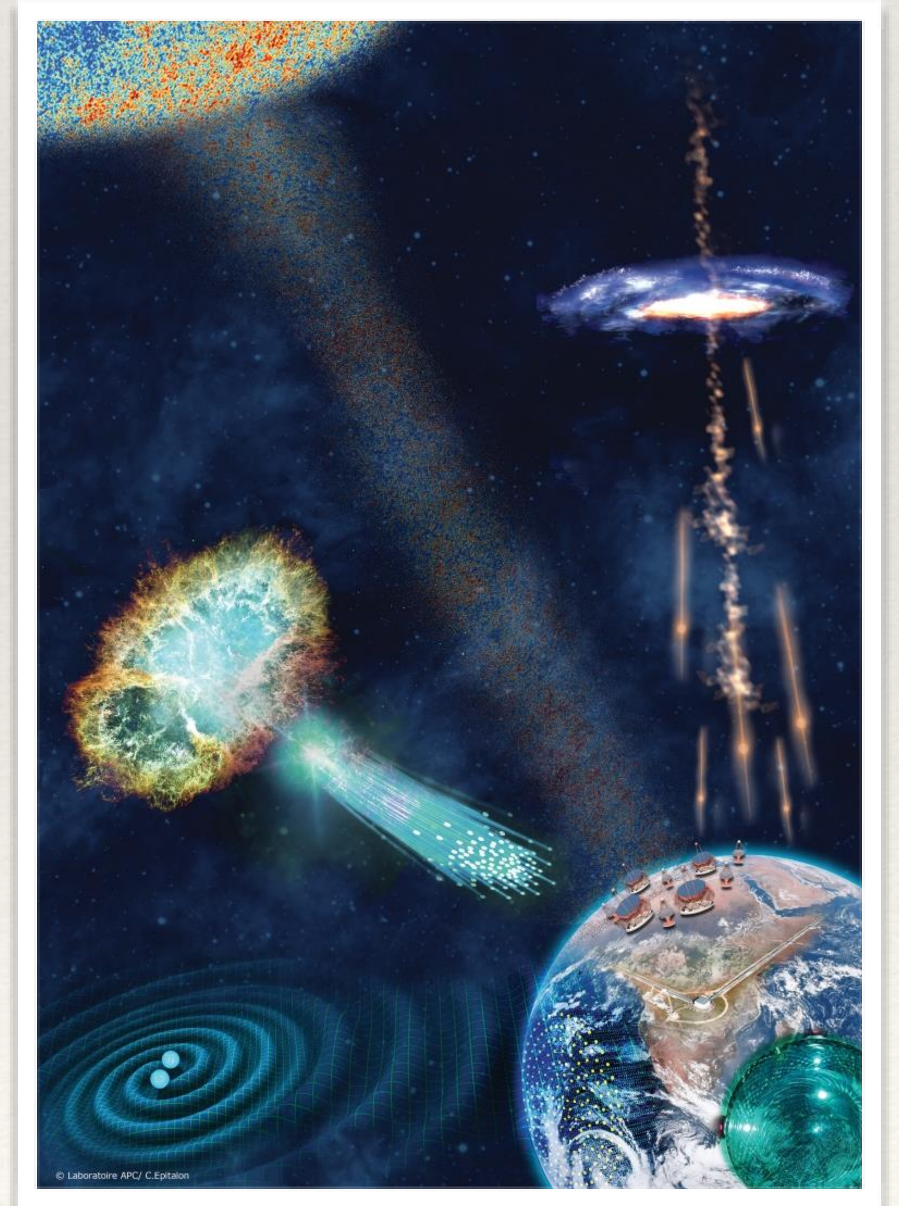

Einstein's dream: unify all forces

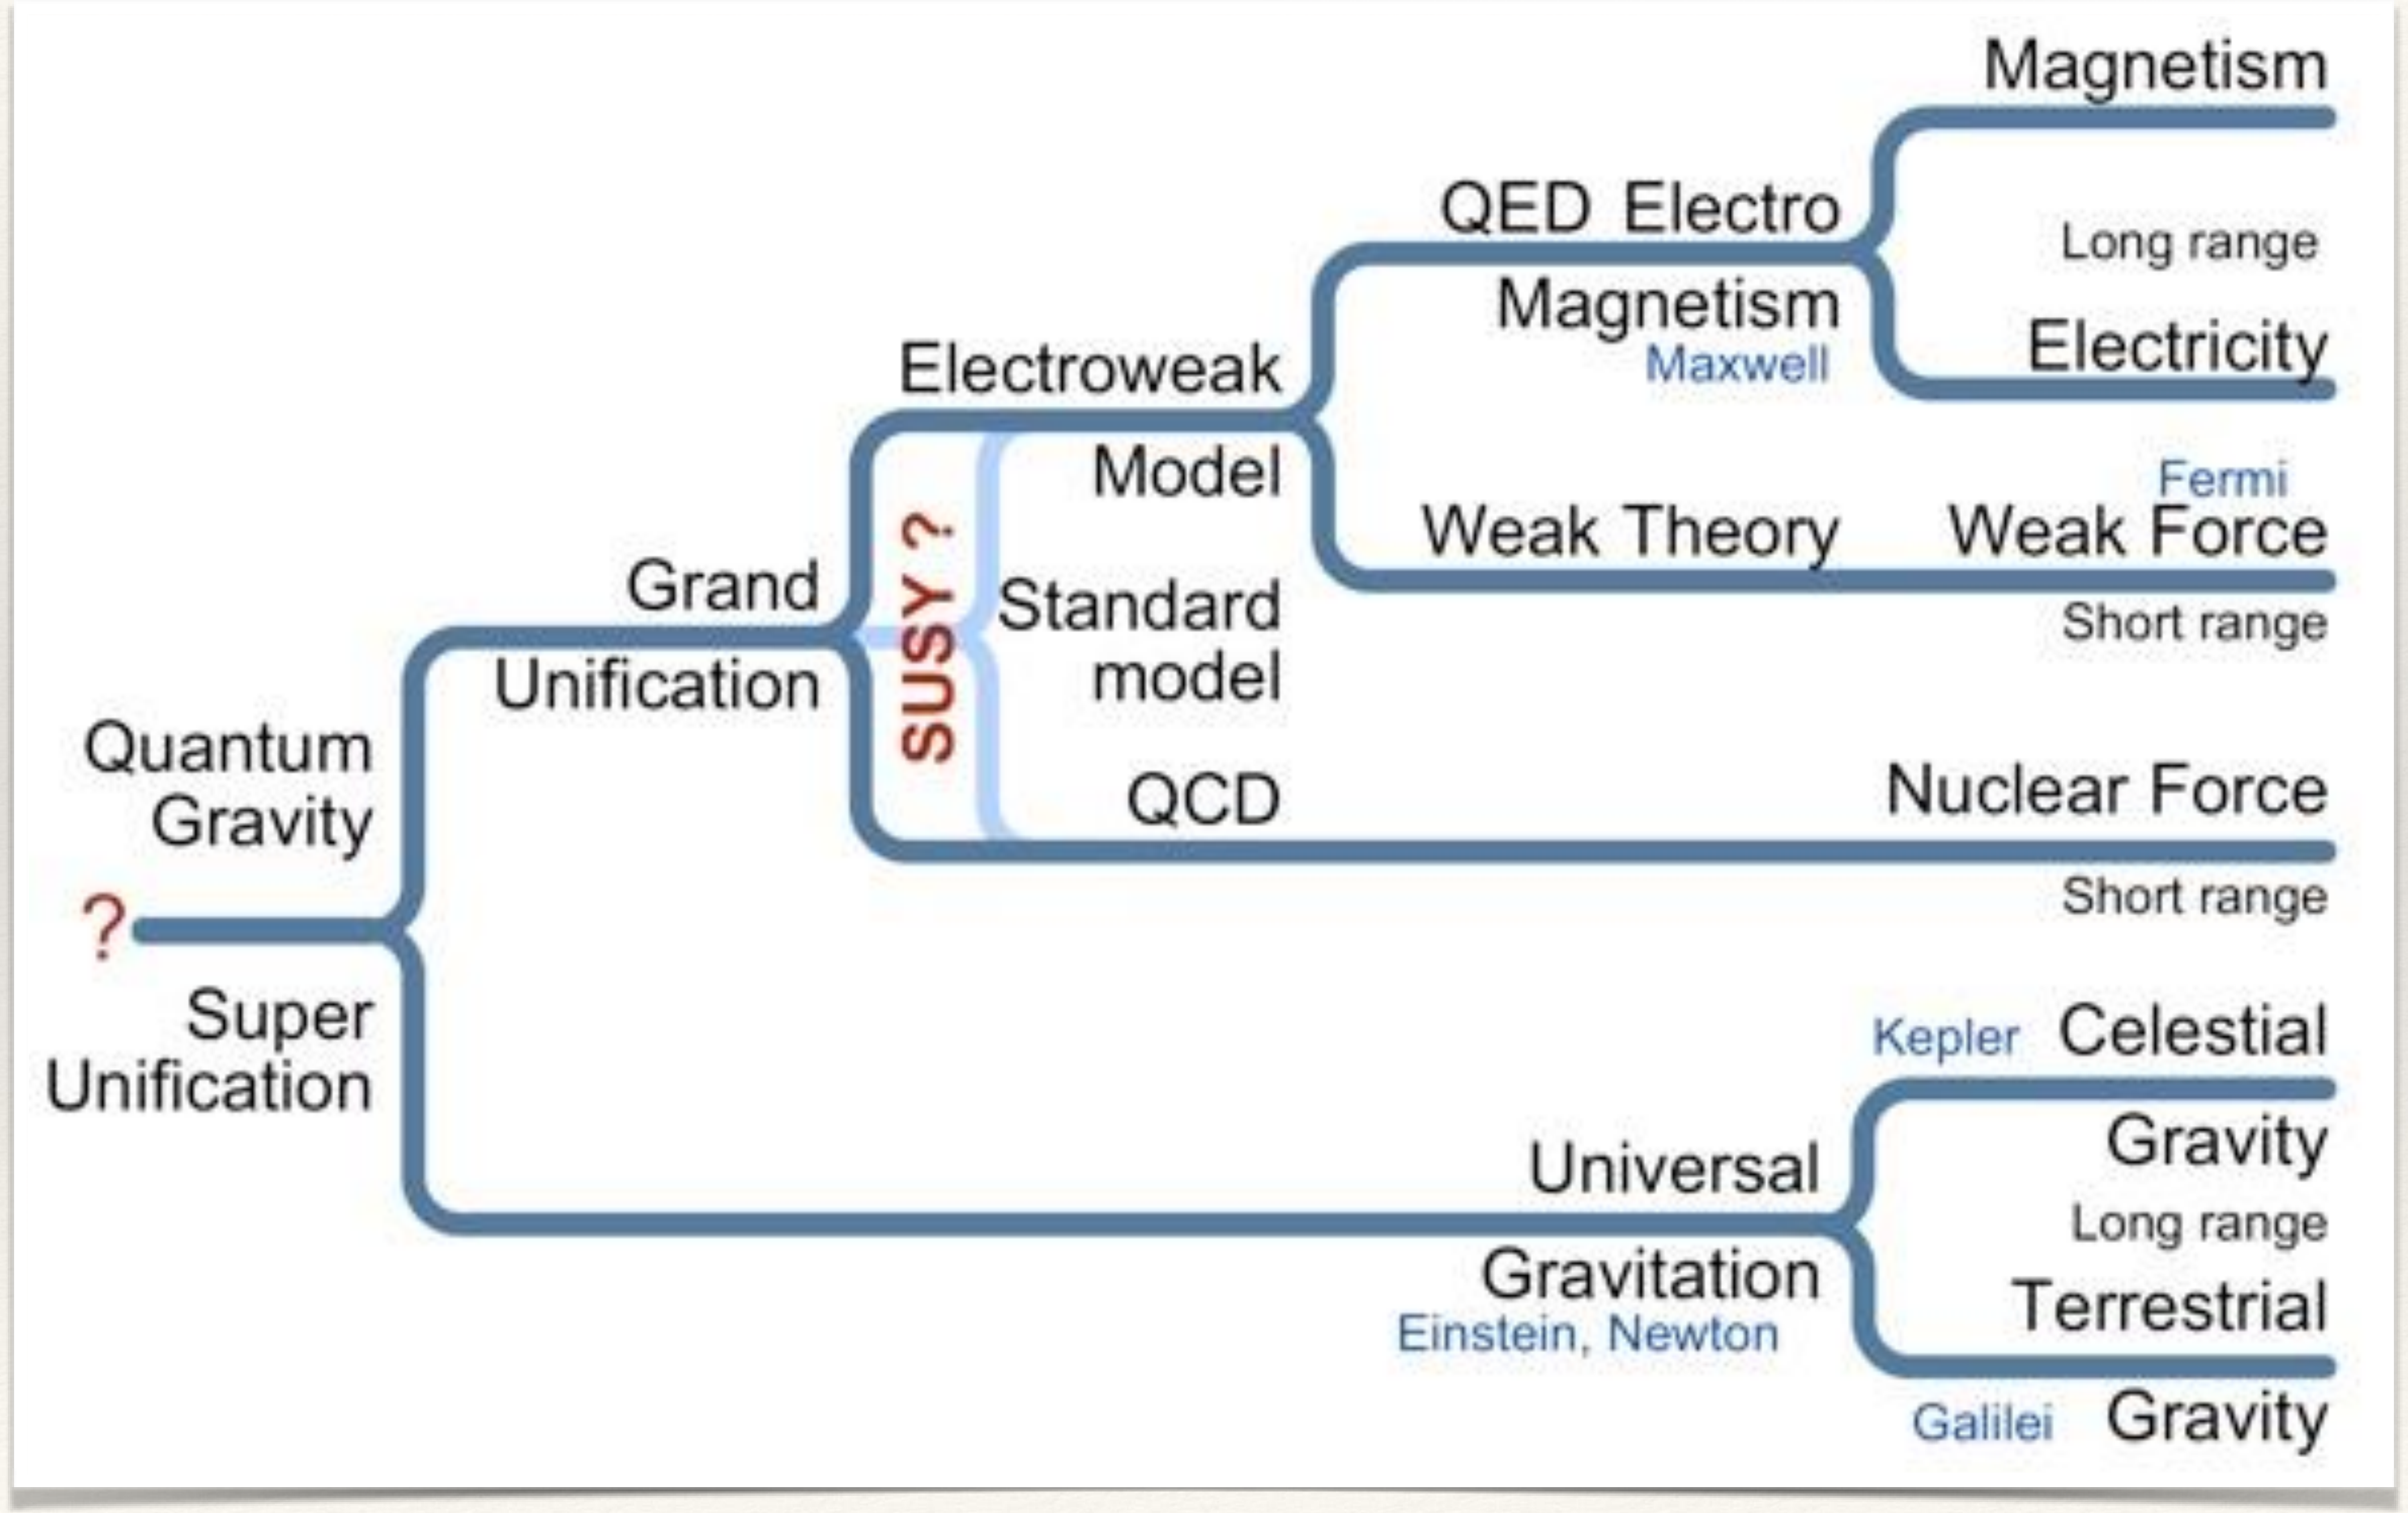

Parker 2013

Elementary particles … in theory

QUANTUM THEORY CENTER
SDU

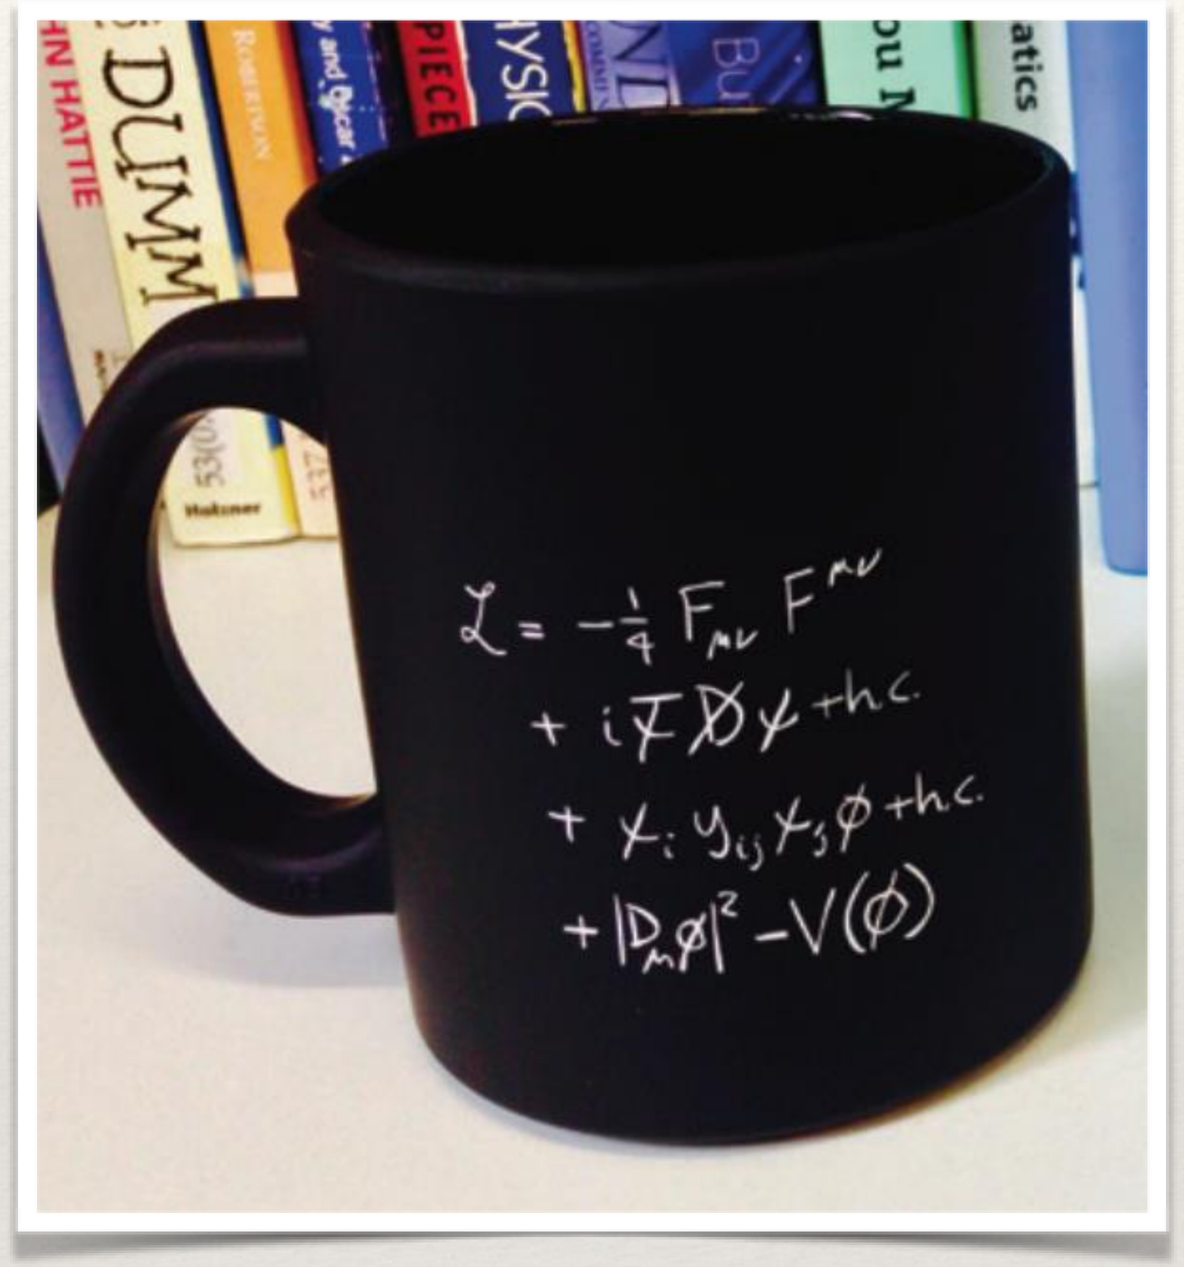

A simple and elegant theory:

the **Standard model of Elementary Particles**

19+7 free parameters (masses and couplings)

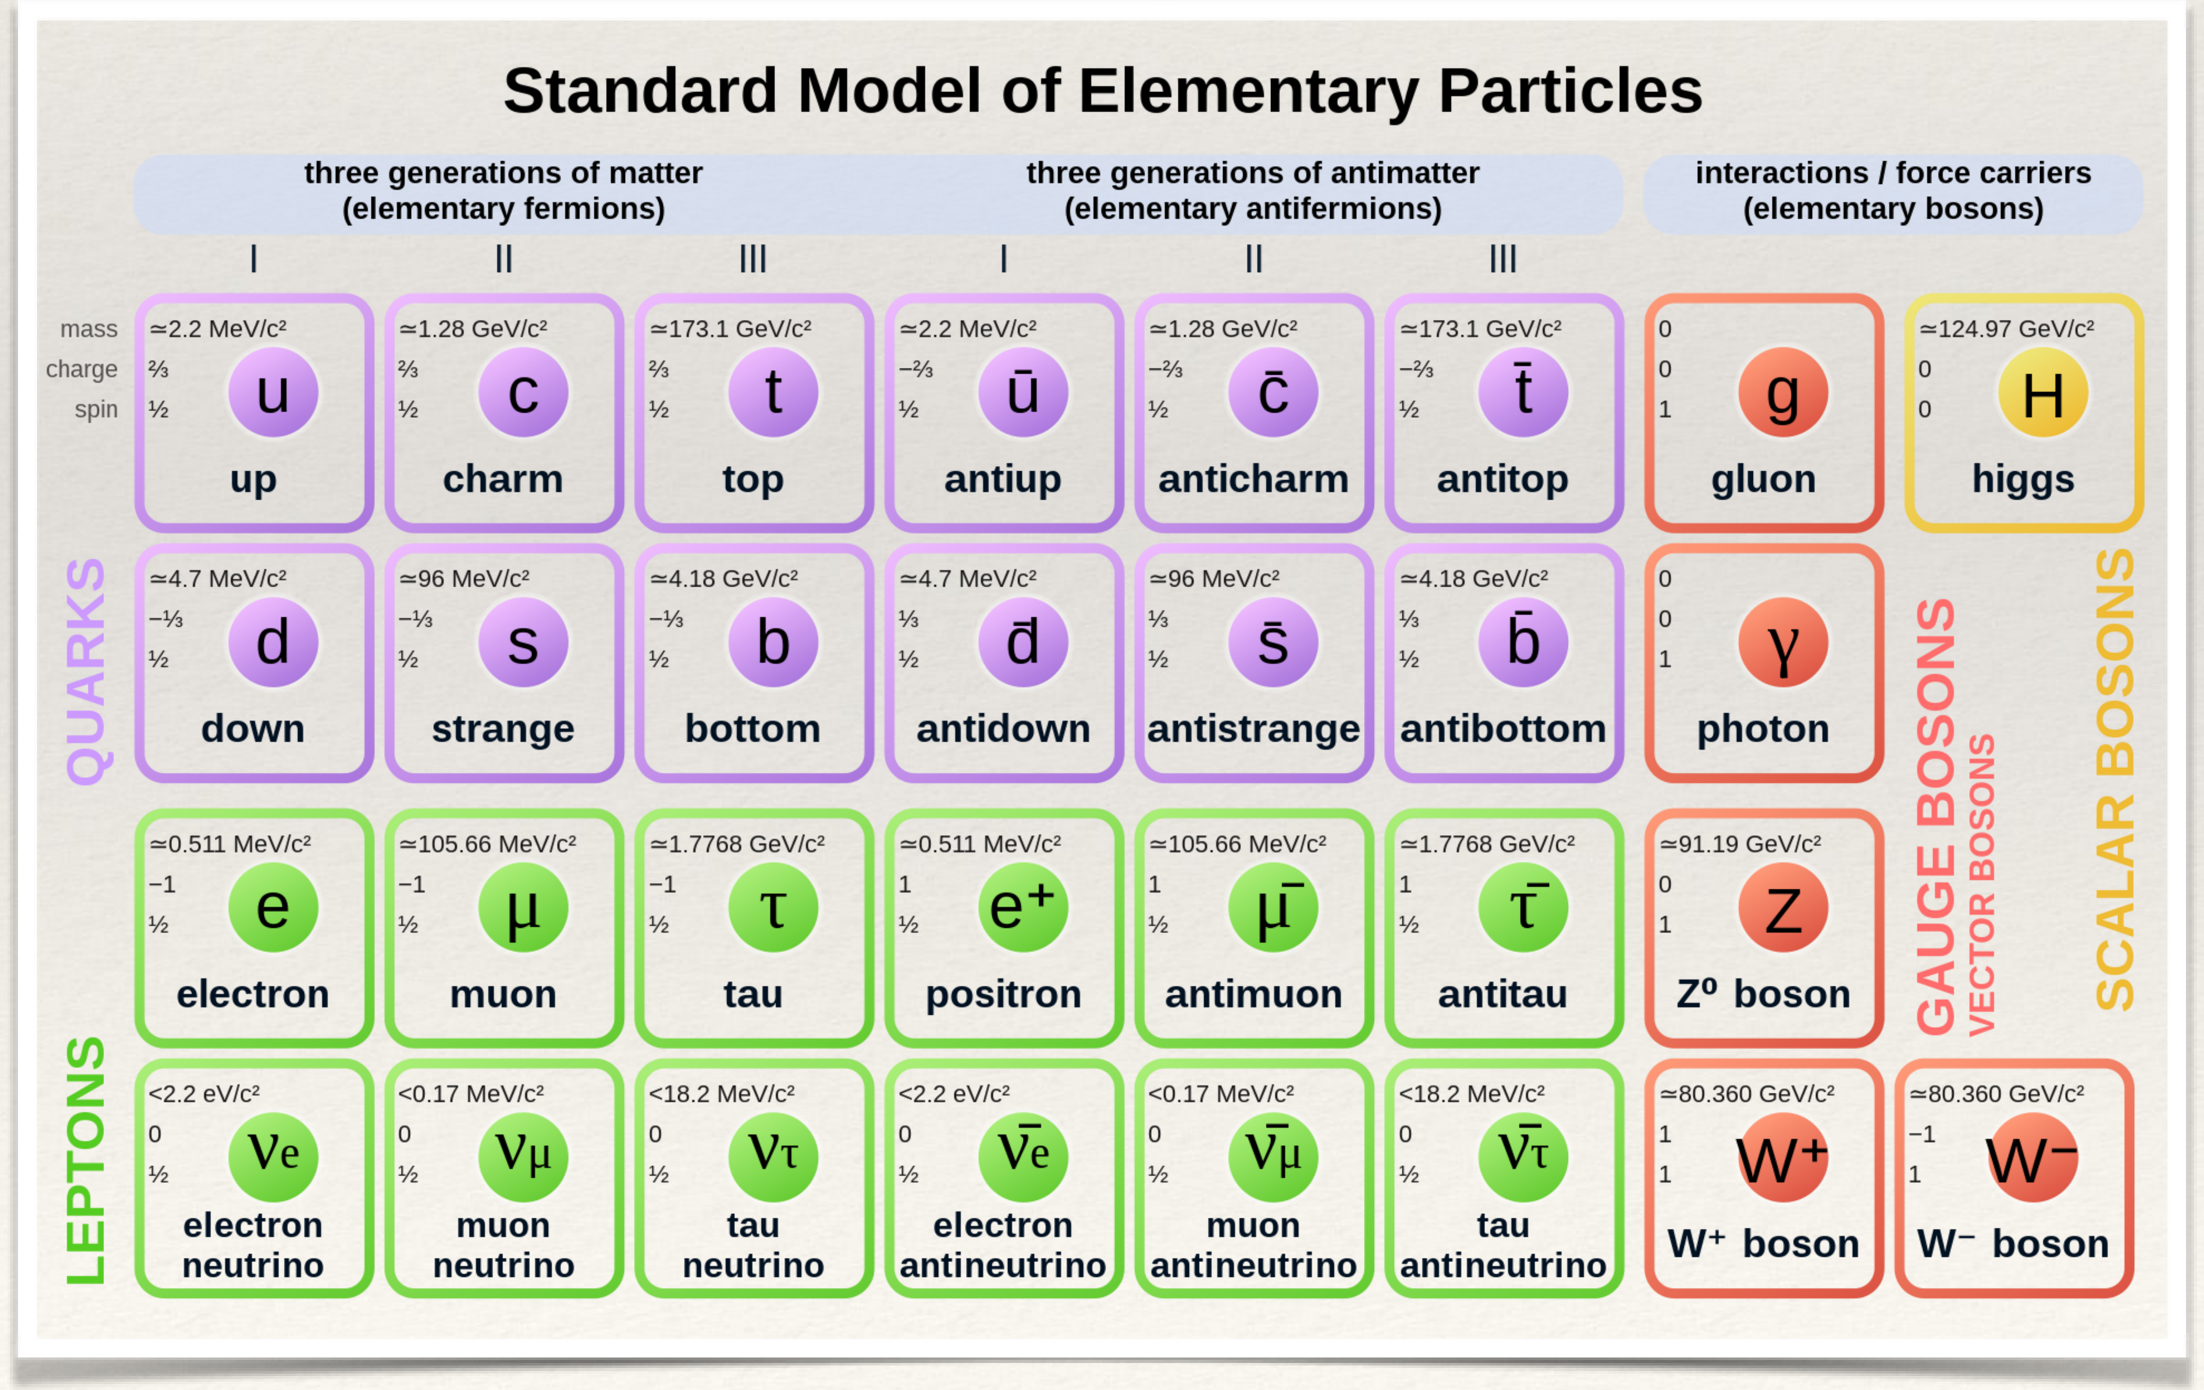

Standard Model of Elementary Particles

	three generations of matter (elementary fermions) I	II	III	three generations of antimatter (elementary antifermions) I	II	III	interactions / force carriers (elementary bosons)	
QUARKS (mass / charge / spin)	≈2.2 MeV/c², ⅔, ½ — u up	≈1.28 GeV/c², ⅔, ½ — c charm	≈173.1 GeV/c², ⅔, ½ — t top	≈2.2 MeV/c², −⅔, ½ — ū antiup	≈1.28 GeV/c², −⅔, ½ — c̄ anticharm	≈173.1 GeV/c², −⅔, ½ — t̄ antitop	0, 0, 1 — g gluon	≈124.97 GeV/c², 0, 0 — H higgs
QUARKS	≈4.7 MeV/c², −⅓, ½ — d down	≈96 MeV/c², −⅓, ½ — s strange	≈4.18 GeV/c², −⅓, ½ — b bottom	≈4.7 MeV/c², ⅓, ½ — d̄ antidown	≈96 MeV/c², ⅓, ½ — s̄ antistrange	≈4.18 GeV/c², ⅓, ½ — b̄ antibottom	0, 0, 1 — γ photon	SCALAR BOSONS
LEPTONS	≈0.511 MeV/c², −1, ½ — e electron	≈105.66 MeV/c², −1, ½ — μ muon	≈1.7768 GeV/c², −1, ½ — τ tau	≈0.511 MeV/c², 1, ½ — e⁺ positron	≈105.66 MeV/c², 1, ½ — μ̄ antimuon	≈1.7768 GeV/c², 1, ½ — τ̄ antitau	≈91.19 GeV/c², 0, 1 — Z Z⁰ boson	GAUGE BOSONS / VECTOR BOSONS
LEPTONS	<2.2 eV/c², 0, ½ — νe electron neutrino	<0.17 MeV/c², 0, ½ — νμ muon neutrino	<18.2 MeV/c², 0, ½ — ντ tau neutrino	<2.2 eV/c², 0, ½ — ν̄e electron antineutrino	<0.17 MeV/c², 0, ½ — ν̄μ muon antineutrino	<18.2 MeV/c², 0, ½ — ν̄τ tau antineutrino	≈80.360 GeV/c², 1, 1 — W⁺ W⁺ boson	≈80.360 GeV/c², −1, 1 — W⁻ W⁻ boson

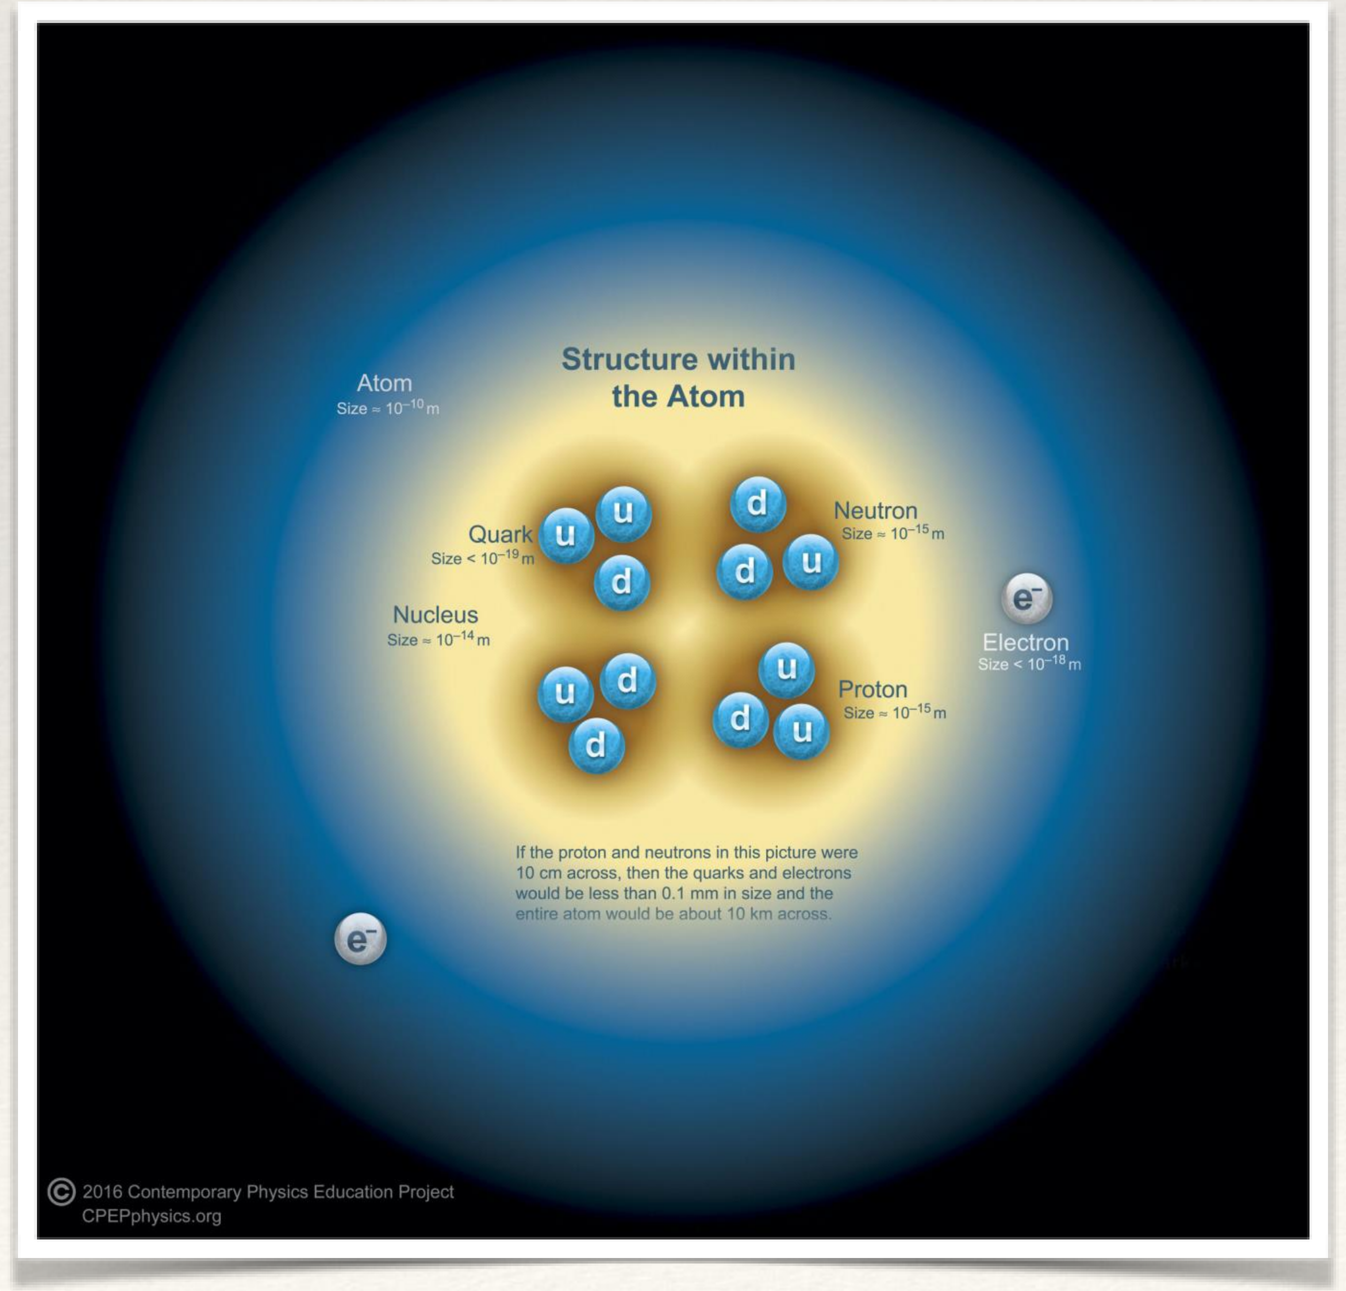

Compare with experiments in a nutshell

QUANTUM THEORY CENTER
SDU

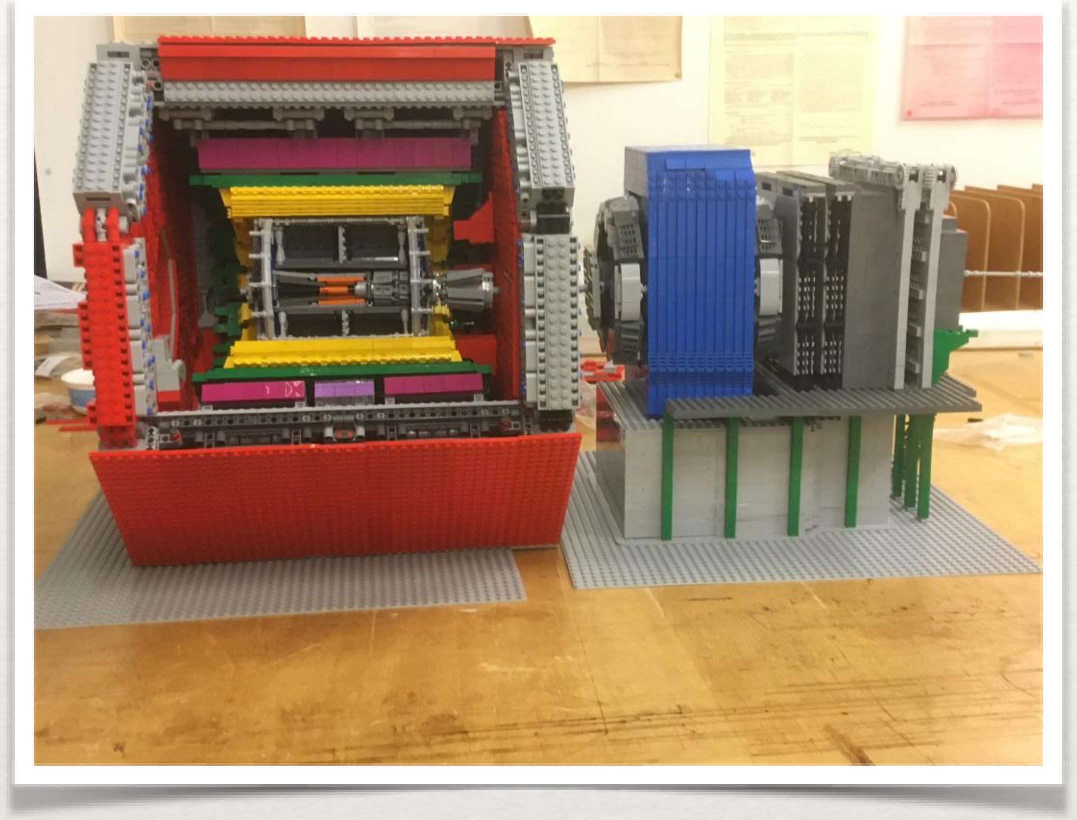

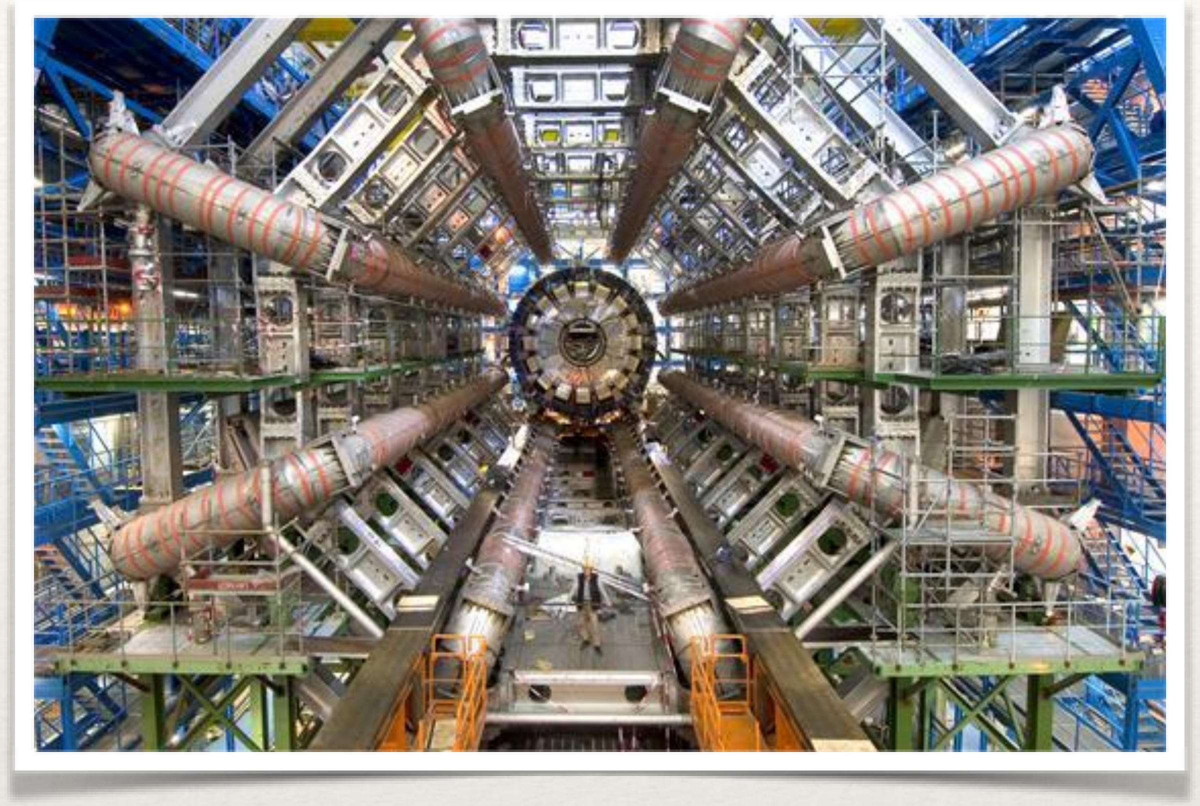

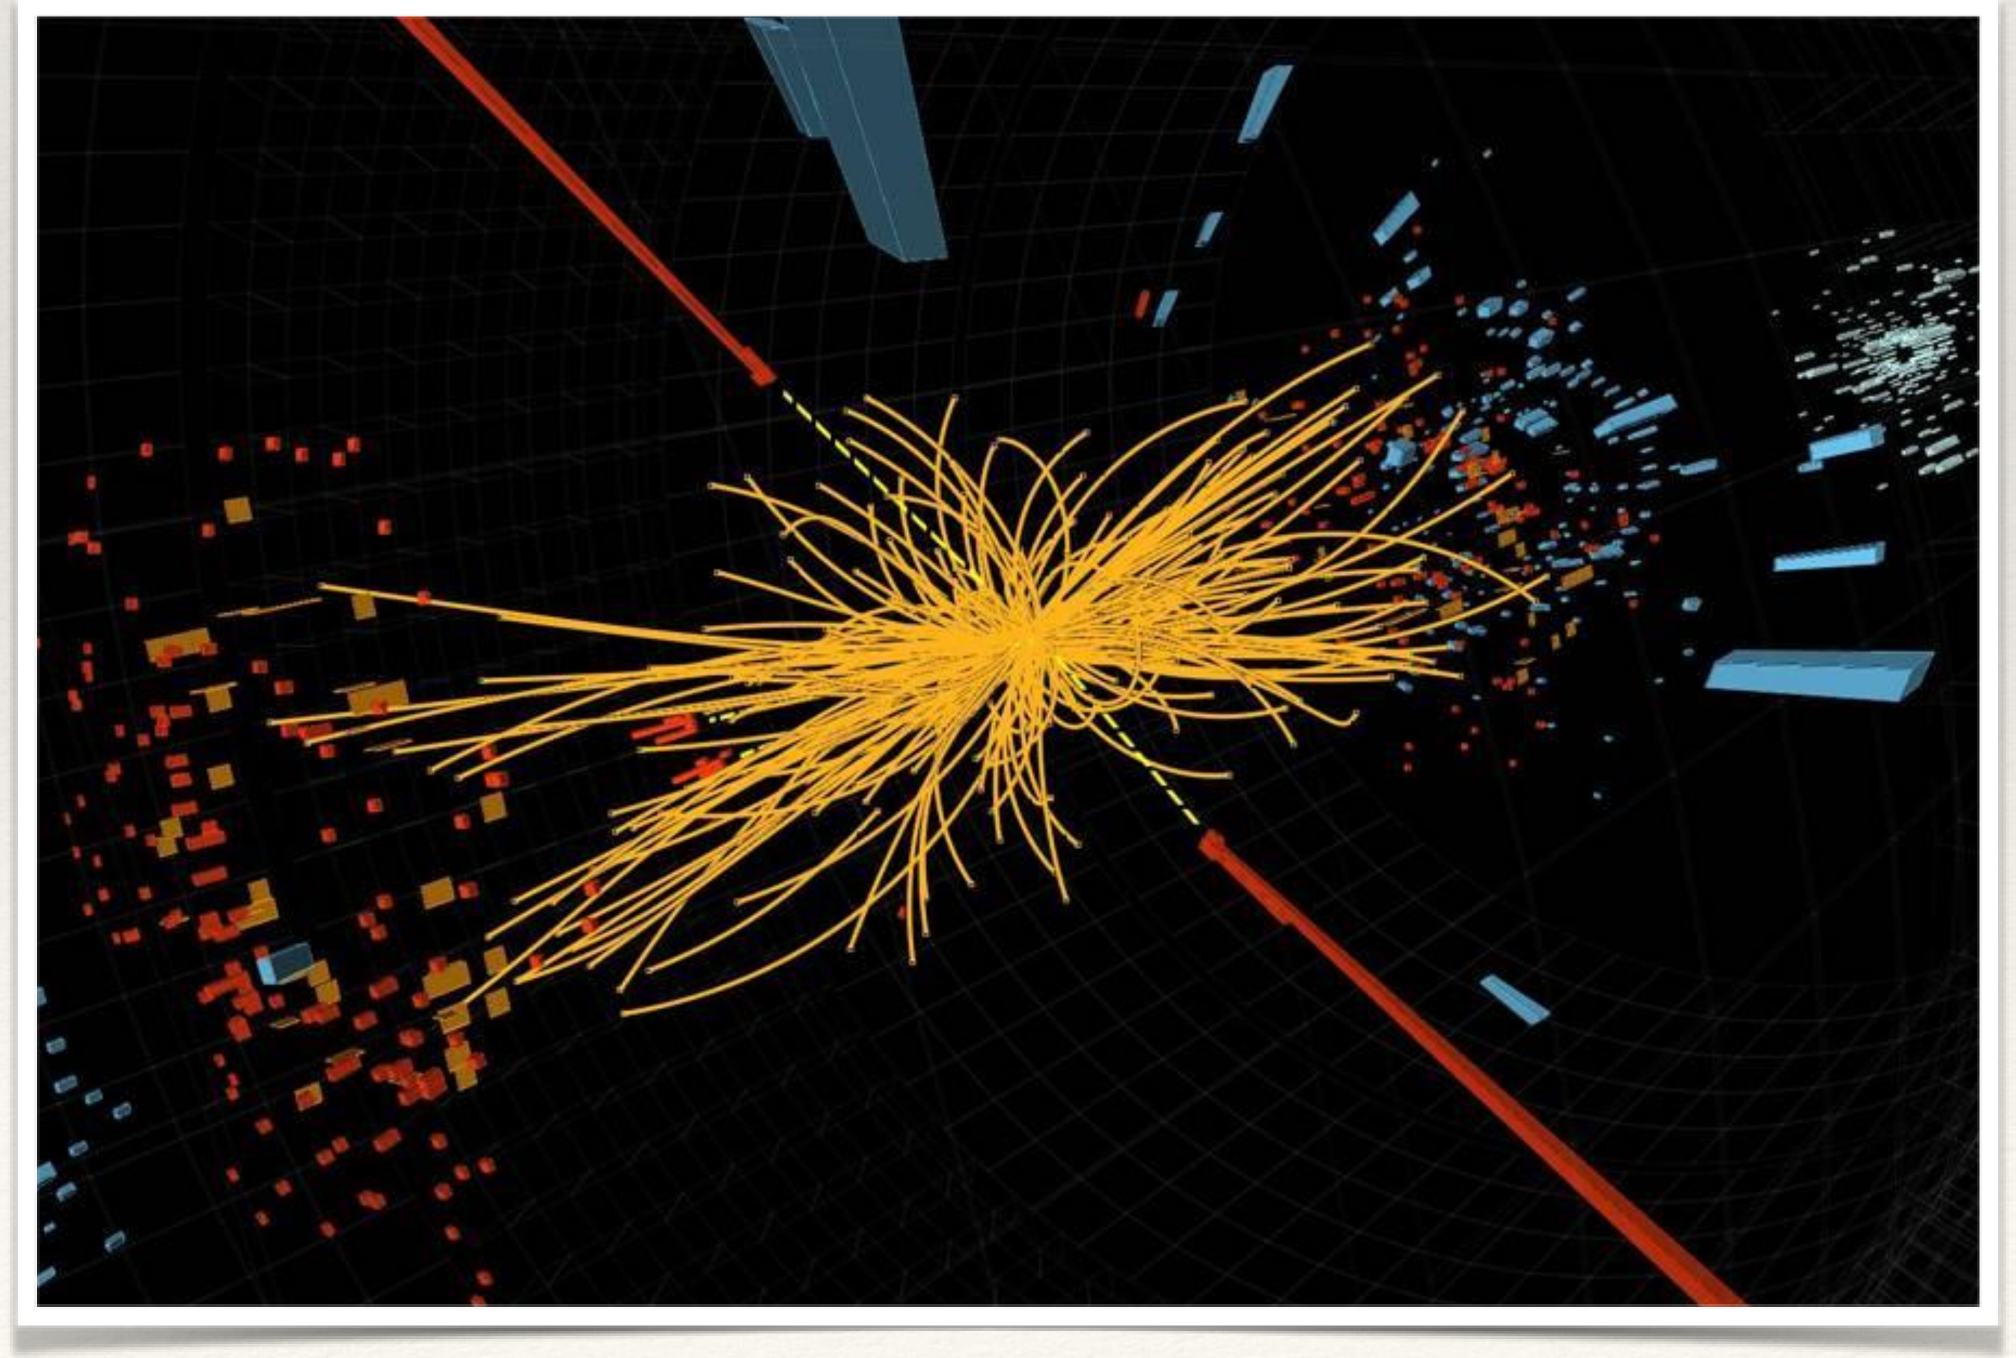

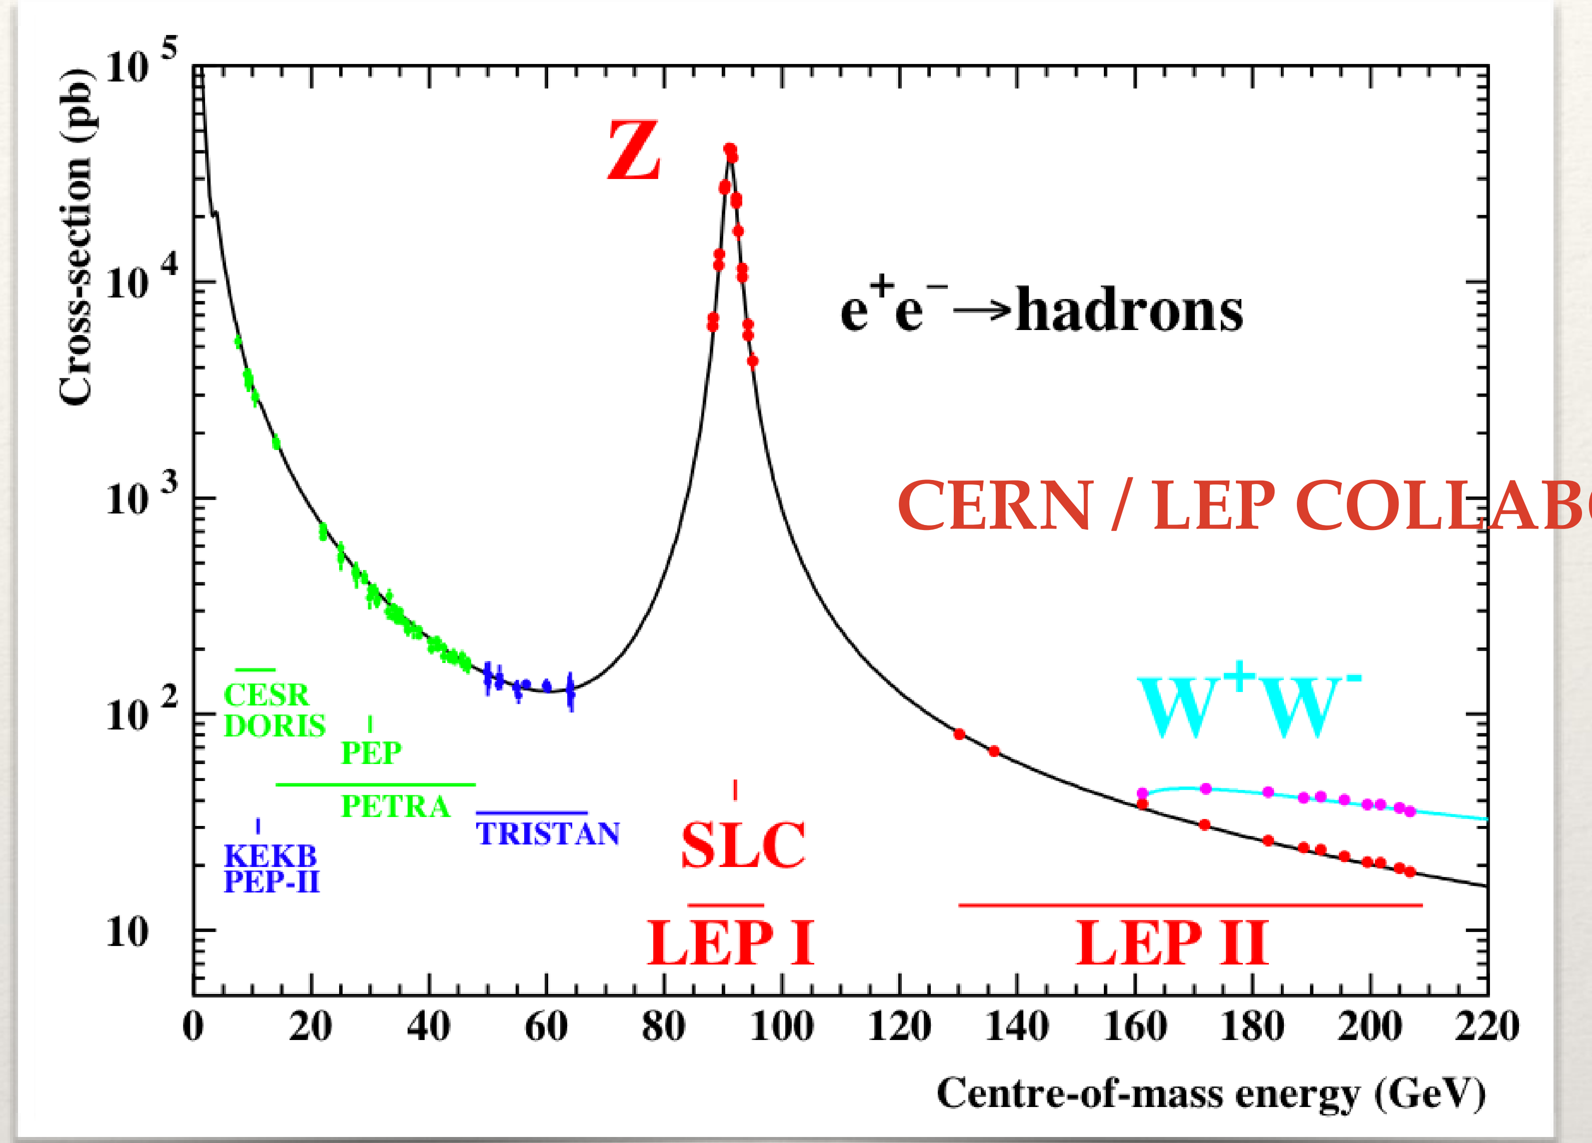

Switch on your experiments and start to count.

For each possible set of decay products, plot the fraction of collisions in which those decay products are produced against the total energy of the particles coming into the collision.

… and compare with your predictions

QUANTUM THEORY CENTER
SDU

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

… a race for **precision**!

… and compare with your predictions

QUANTUM THEORY CENTER SDU

Quantity	Measured (GeV)
Mass of W boson	80.387 ± 0.019
Mass of Z boson	91.1876 ± 0.0021

… a race for **precision**!

Only a fine-grained control of all the contributions included in the Standard Model can show verification or tension in our predictions!

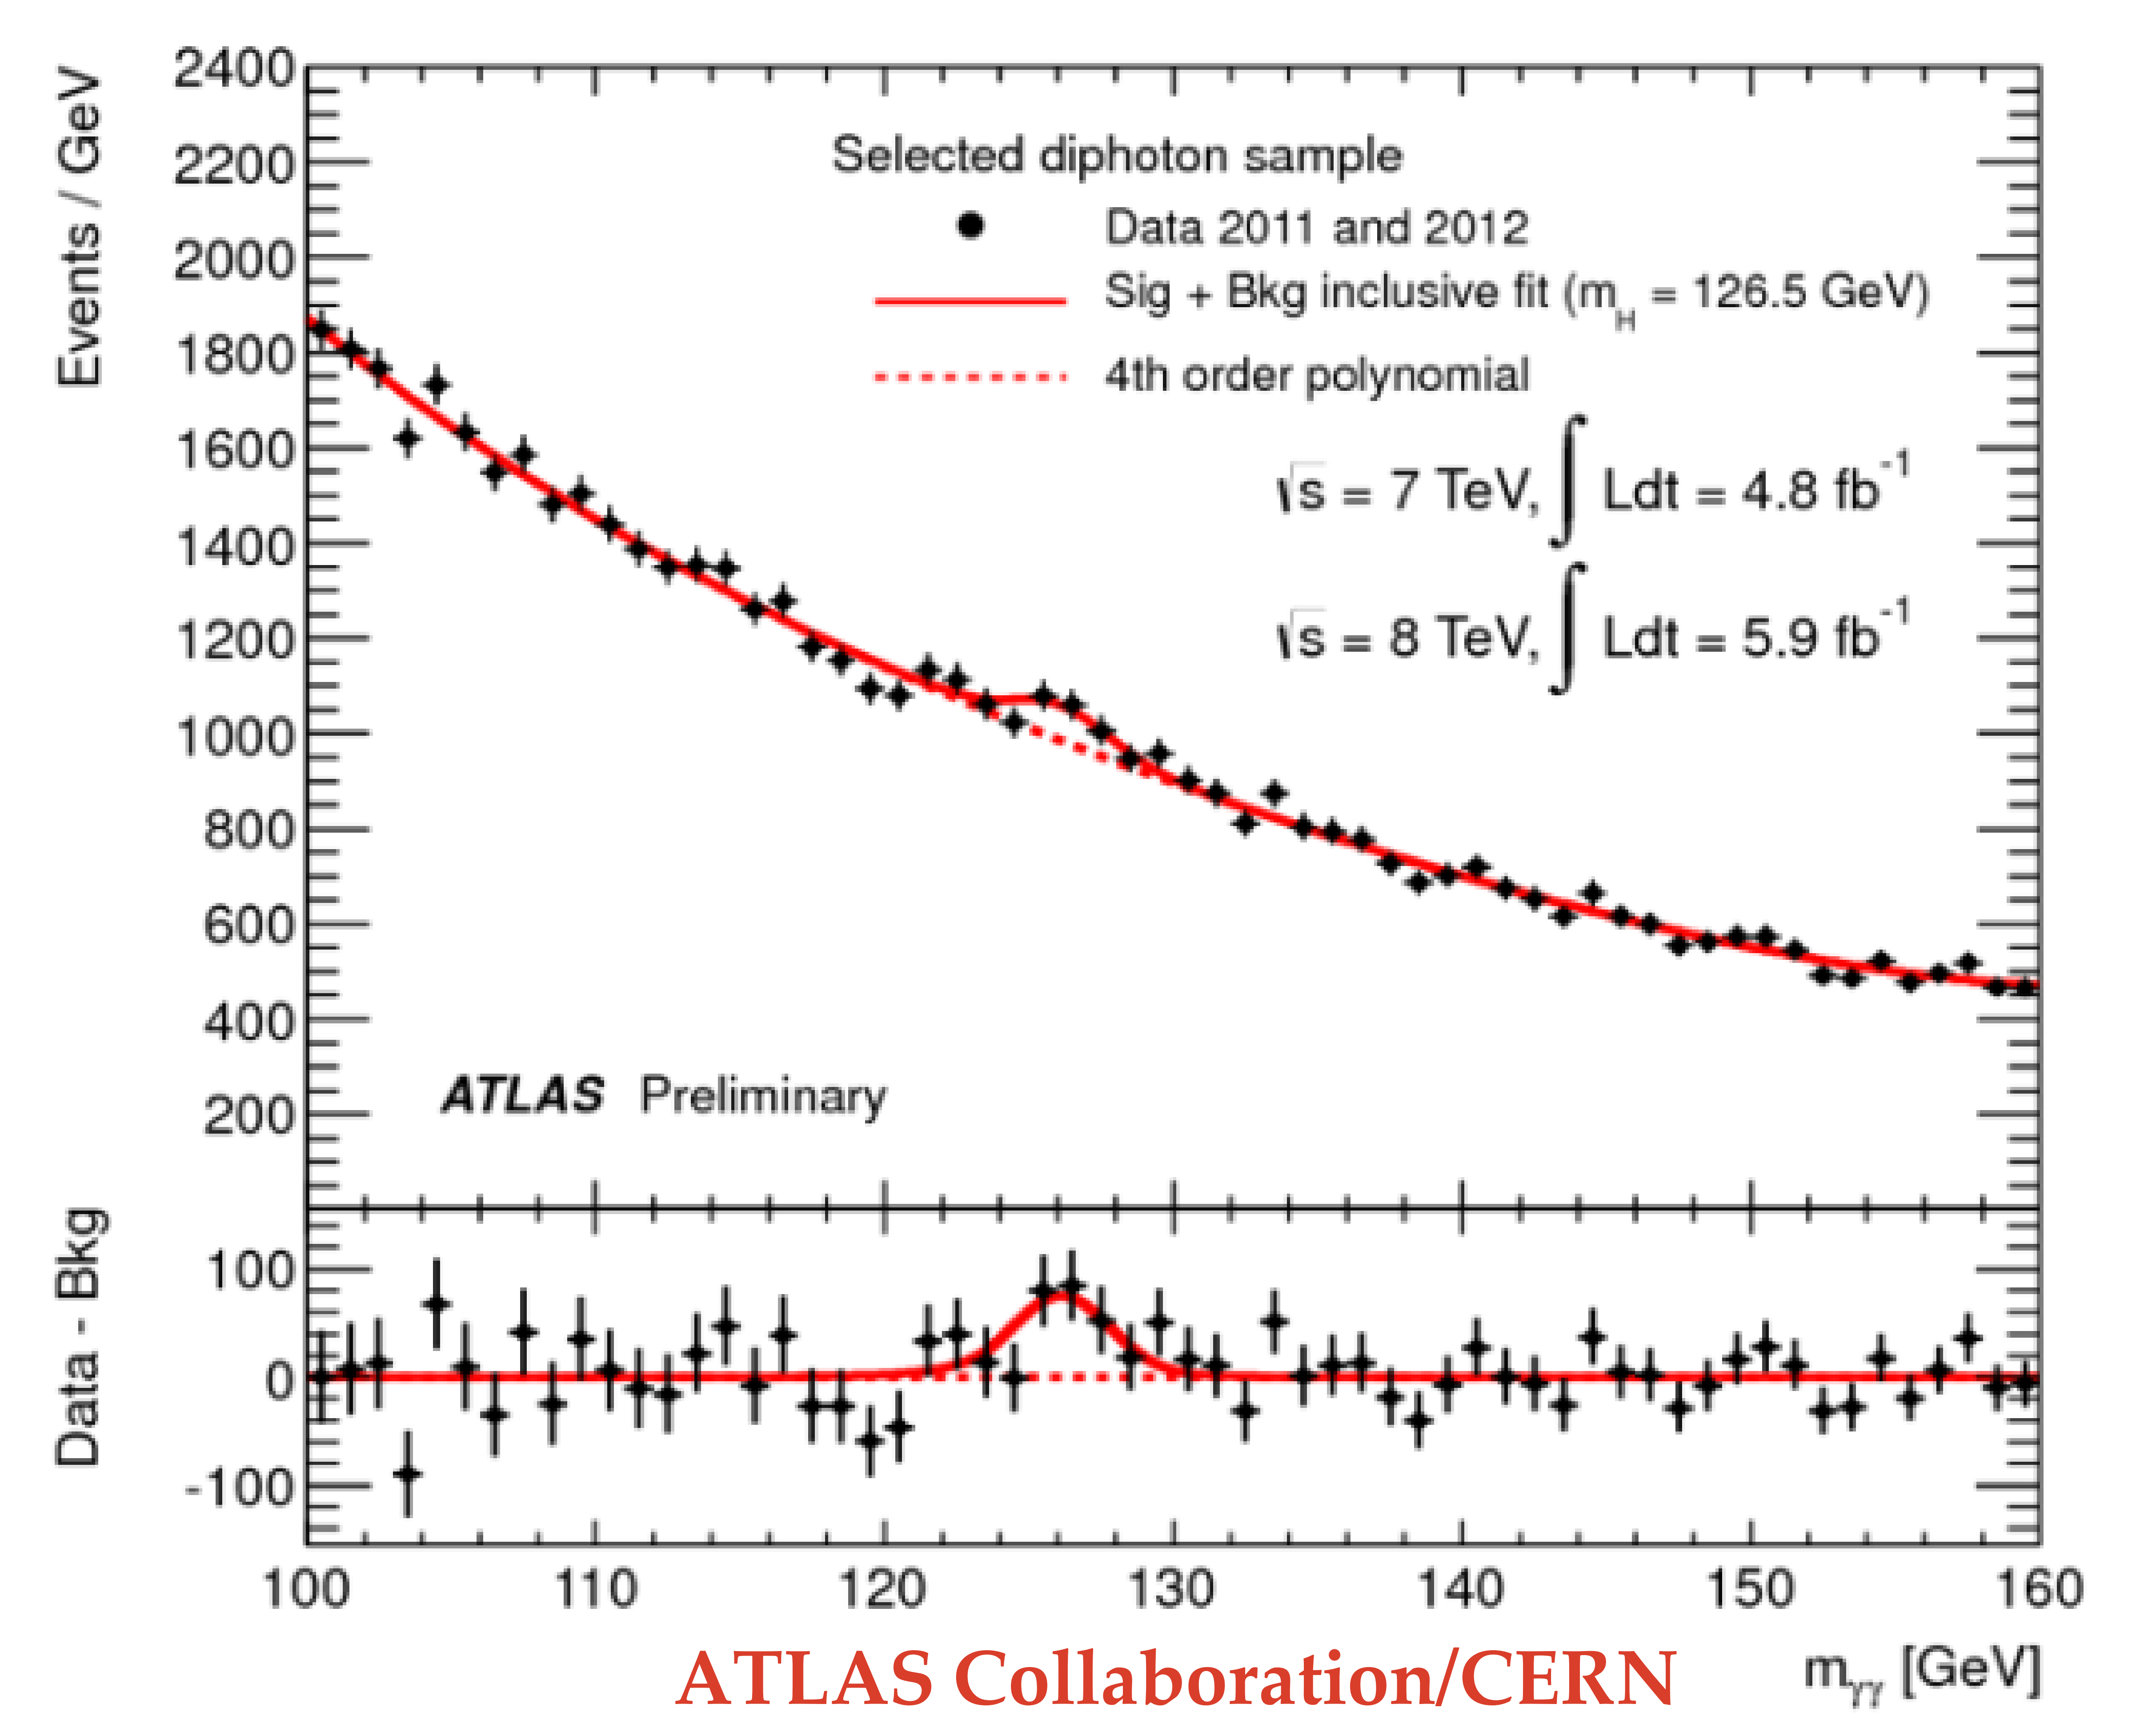

ATLAS Collaboration/CERN

different frameworks

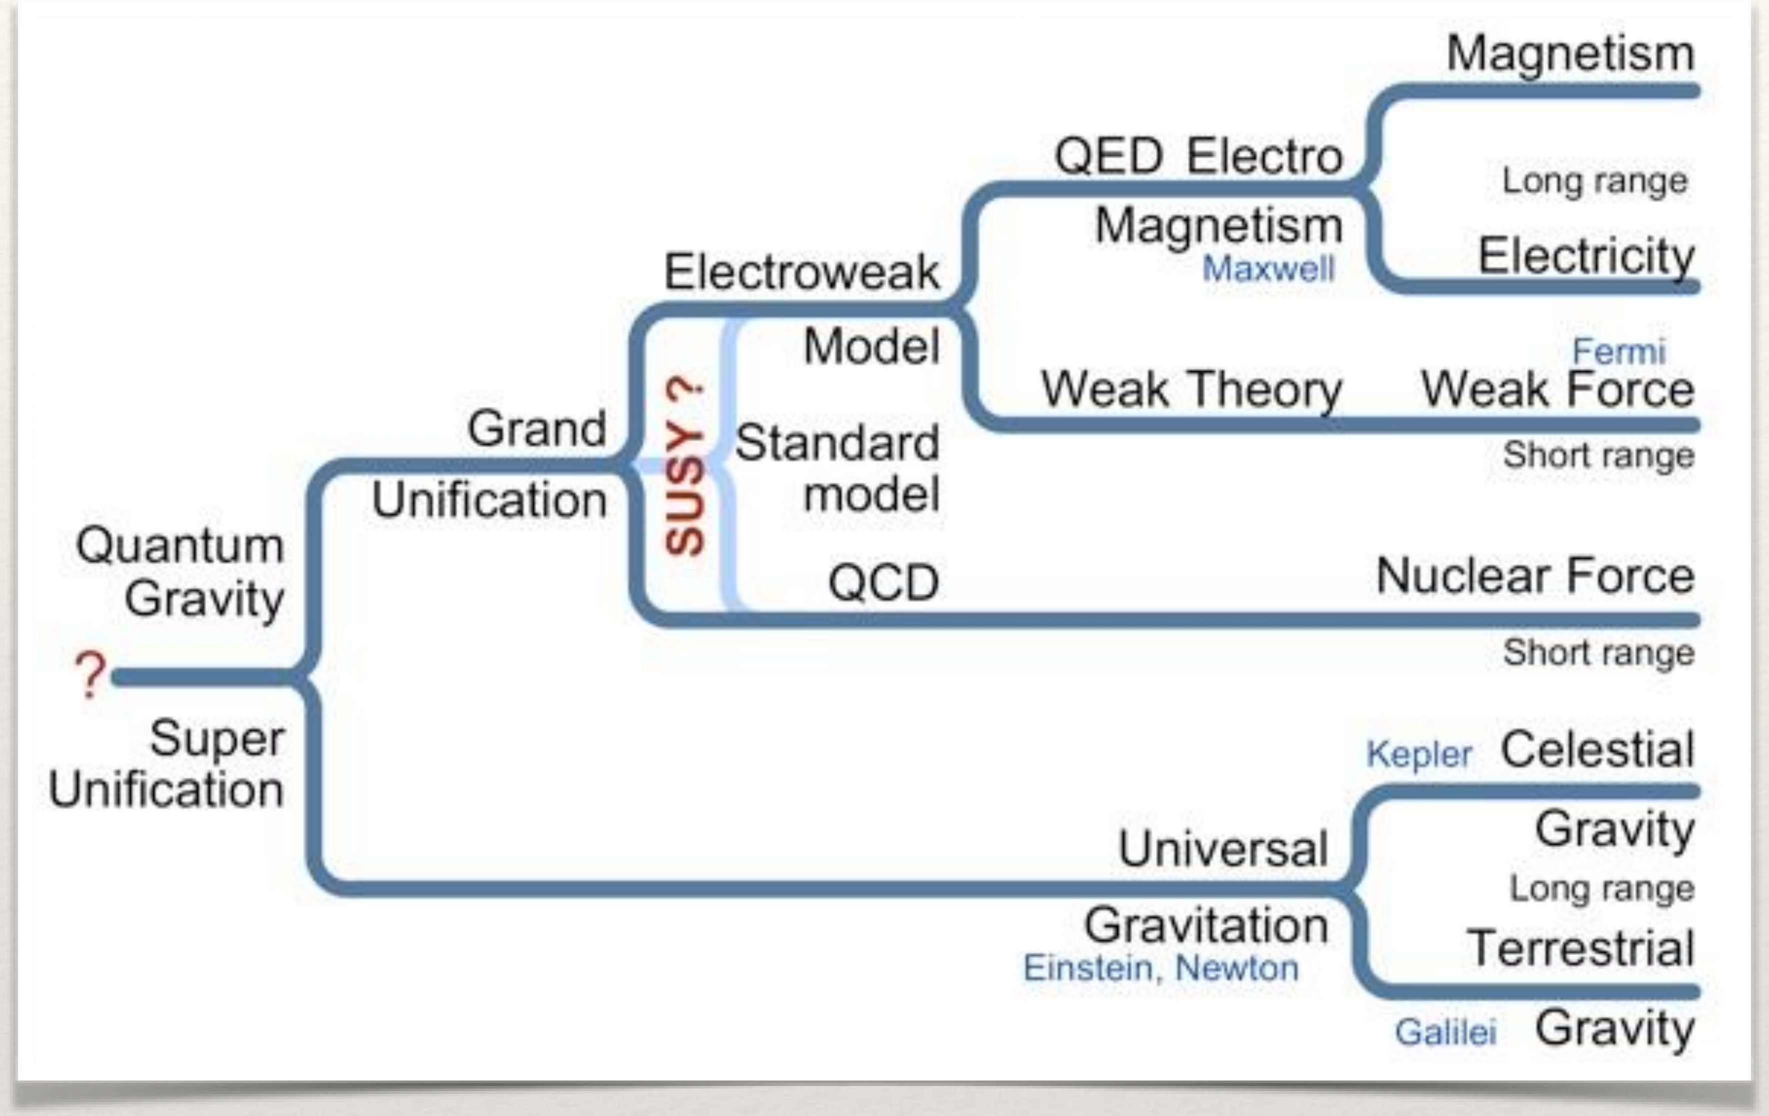

Can be studied in **perturbative**

Quantum Field Theory

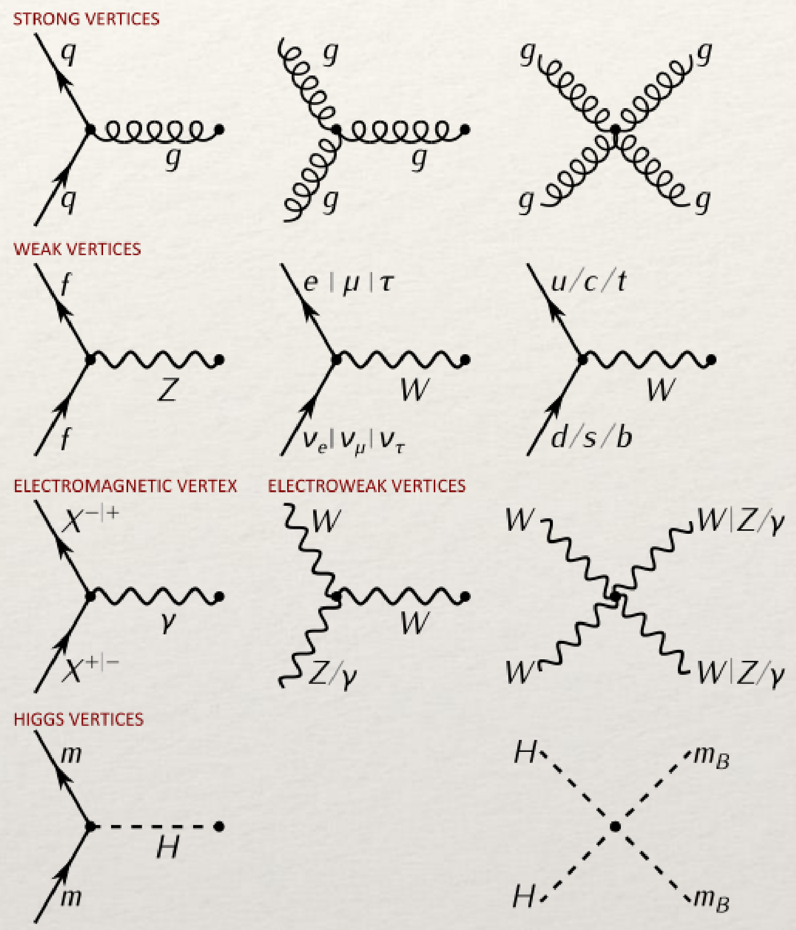

The quantum observables are formal power series in the coupling constant which measures the strength of the interaction.

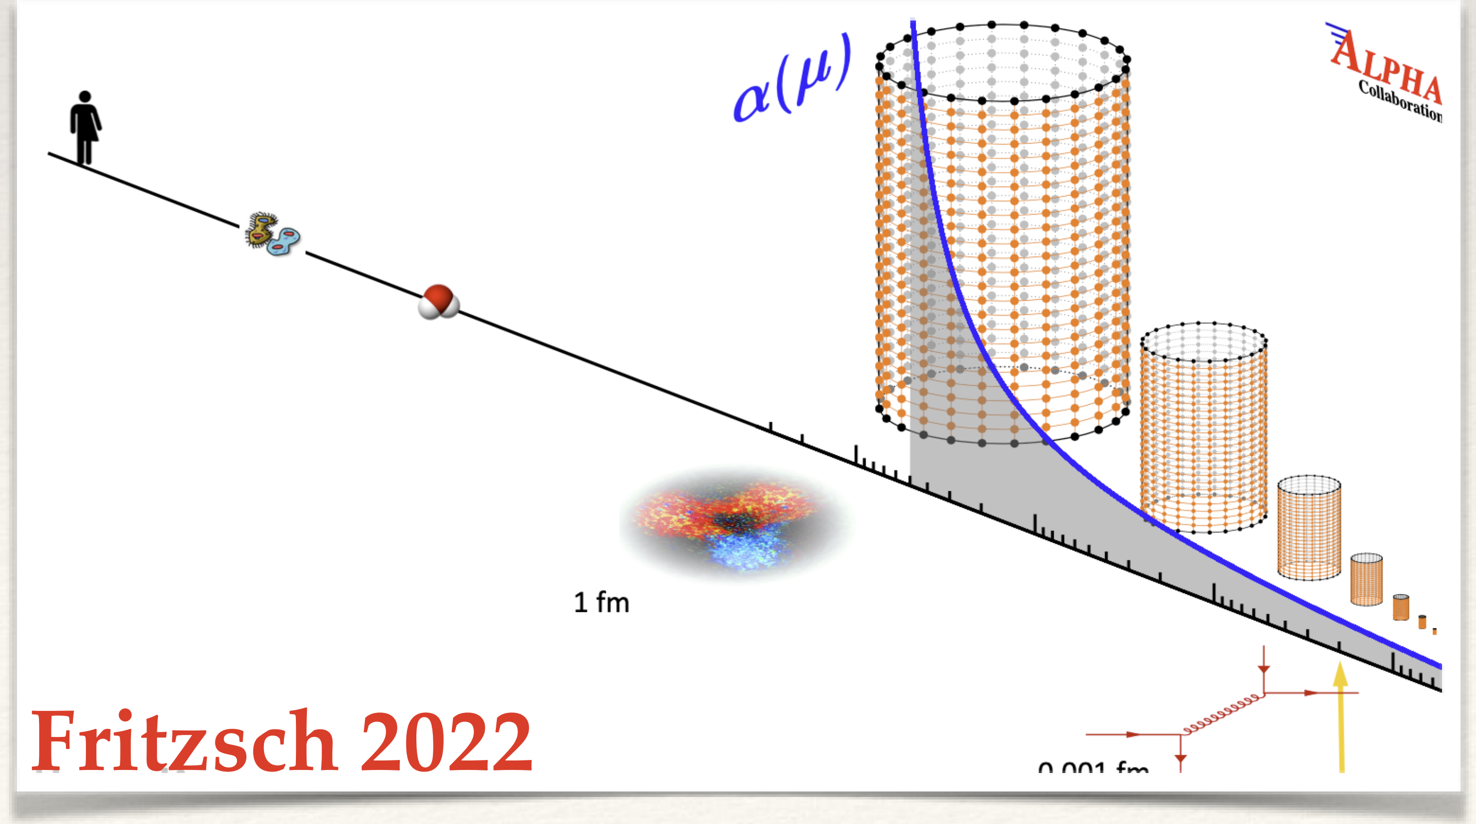

Fritzsch 2022

There are however known non-perturbative effects which are not captured in perturbation theory.

What about (super)computers?

QUANTUM THEORY CENTER
SDU

… a long lasting relation

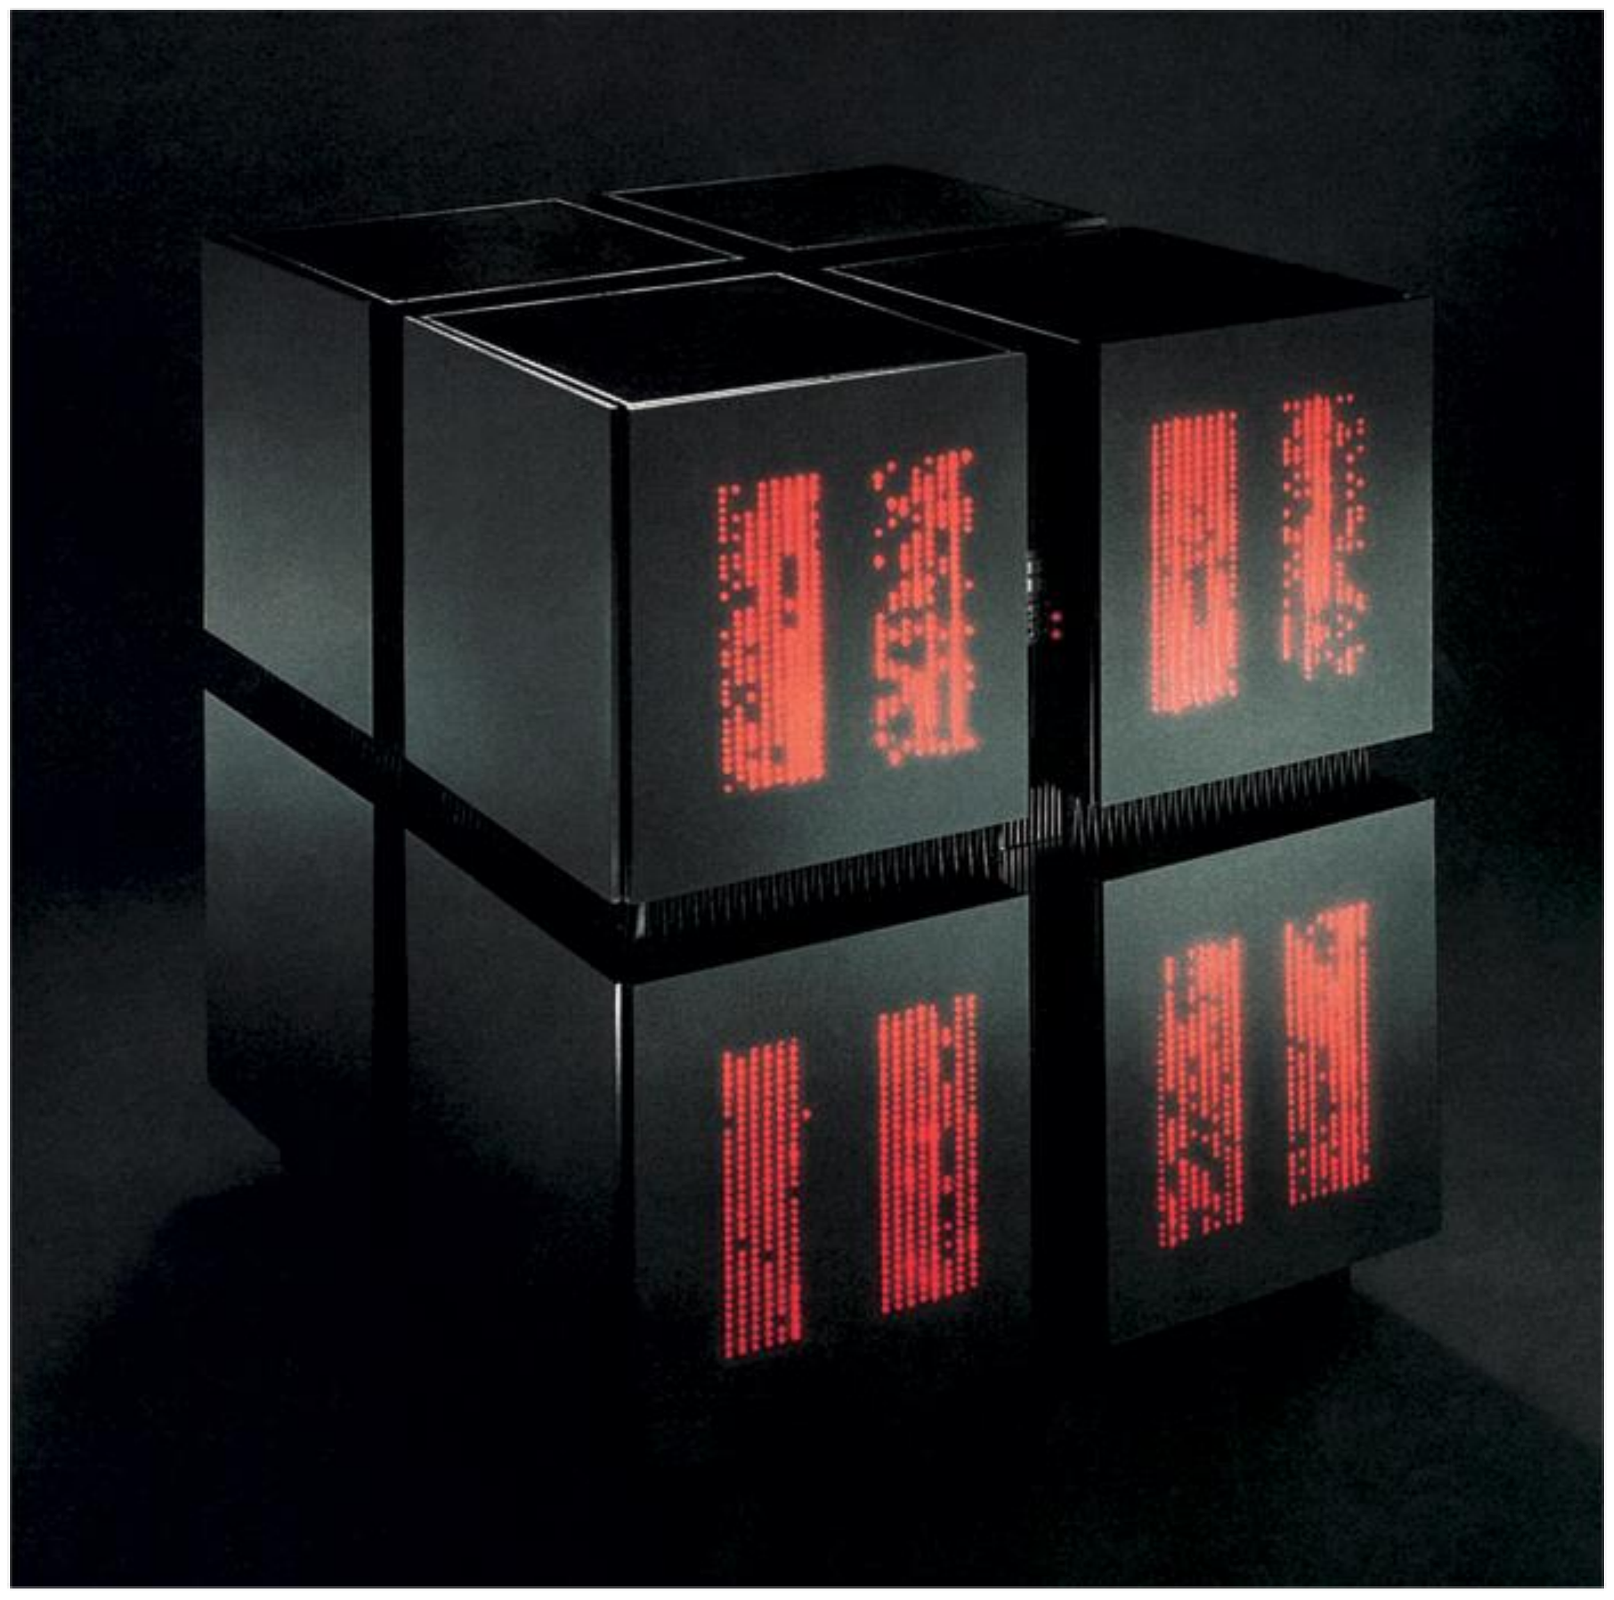

"1986: First commercial AI supercomputer modelled after the human brain

"The Connection Machine CM-1 was the first commercial supercomputer designed expressly for problems of Artificial Intelligence (AI). A massively parallel supercomputer with 65,536 processors, it was the brainchild of Danny Hillis, conceived in the early 1980s while he was a doctoral student with Marvin Minsky at the MIT Artificial Intelligence Lab, and built at his start-up Thinking Machines Corporation. Departing from conventional computer architecture of the time, the CM-1 was modeled on the structure of a human brain: rather than relying on a single powerful processor to perform calculations one after another, the data was distributed over the tens of thousands of simple 1-bit processors, all of which could perform calculations simultaneously, an architecture known as Single Instruction Multiple Data (SIMD).

"What enabled the processors to communicate faster than previous SIMD designs was the internal network, a 12-dimensional boolean n-cube structure suggested by Nobel Prize physicist Richard Feynman, who spent his summers working with us. Within this hardwired physical structure, the software data structures for communication and transfer of data between processors could change as needed depending on the nature of the problem. The connections between processors were more important than the processors themselves, hence the name "Connection Machine."

fromTamiko Thiel's website: http://tamikothiel.com/cm/

What about (super)computers?

QUANTUM THEORY CENTER
SDU

The equations of SM can be put on a computer. The theory must be described in a discrete space-time, a lattice is introduced.

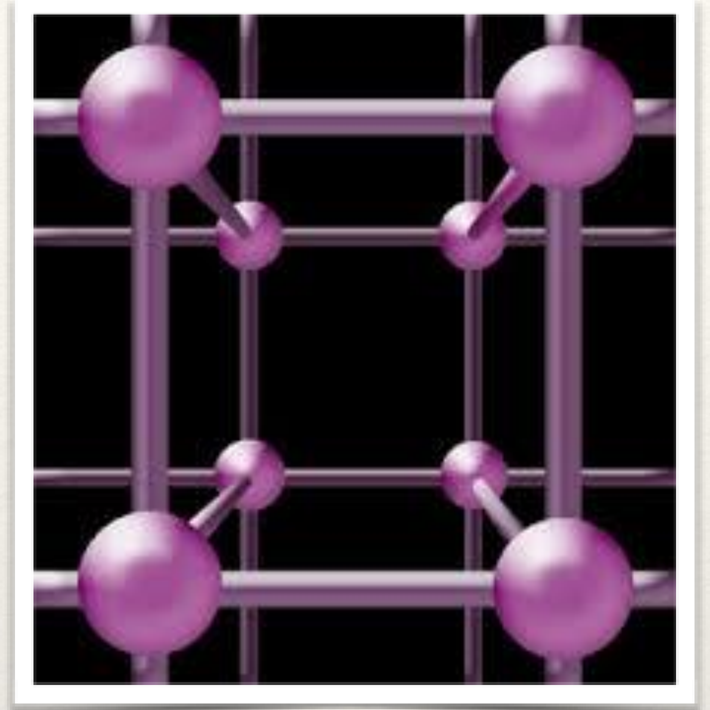

The original theory can be recovered in the limit for the lattice spacing going to zero

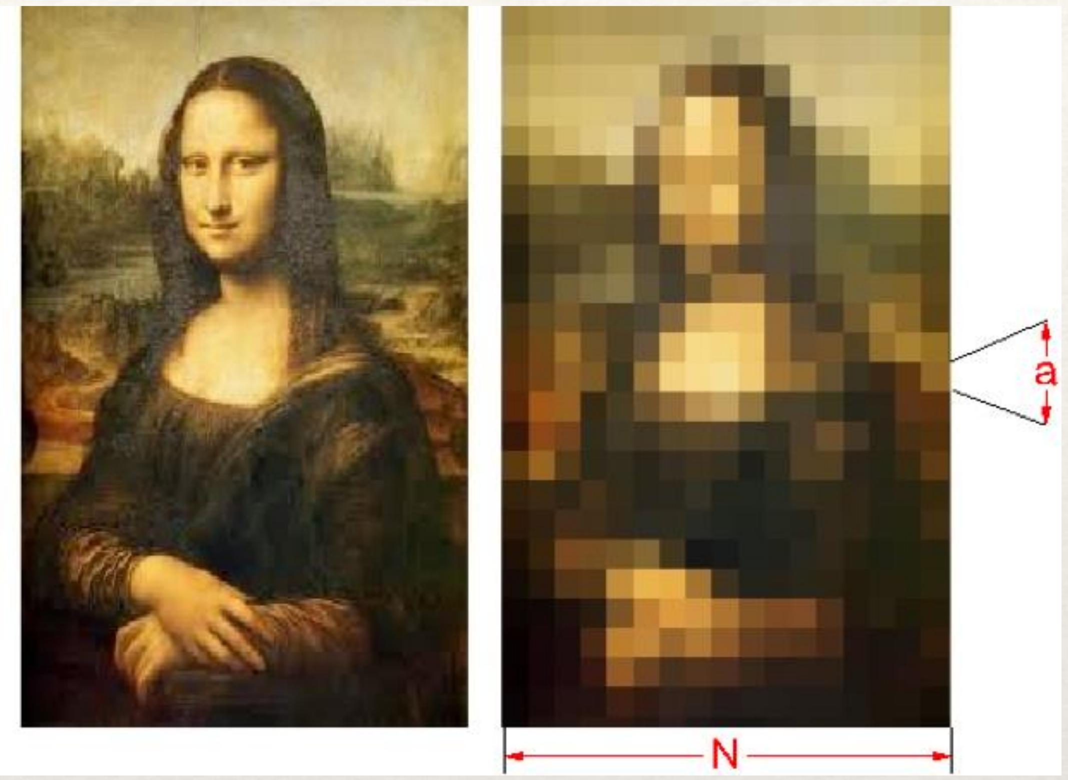

daVinci 1503

Measurements of phenomenological quantities are conducted by evaluating observables across extensive sets of statistically relevant snapshots.

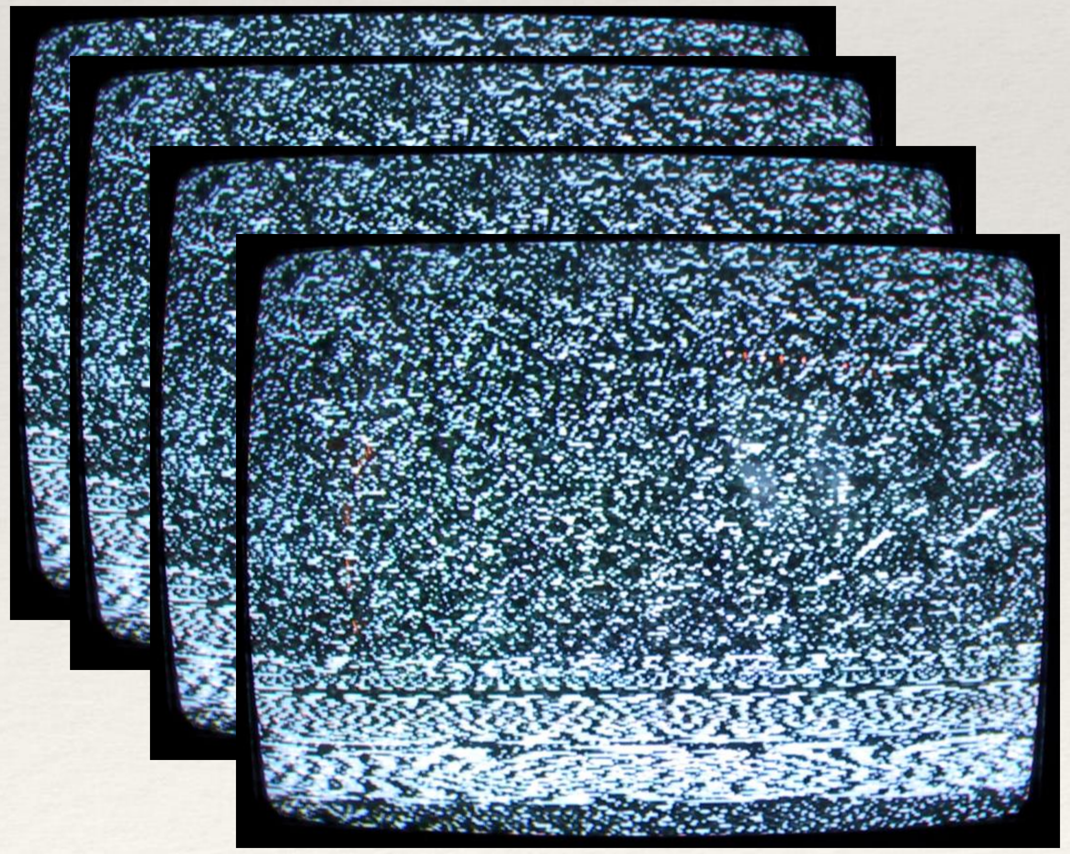

Typically the equations are solved using massively parallel supercomputers with many thousands of nodes to solve the equations.

bubbles in the vacuum

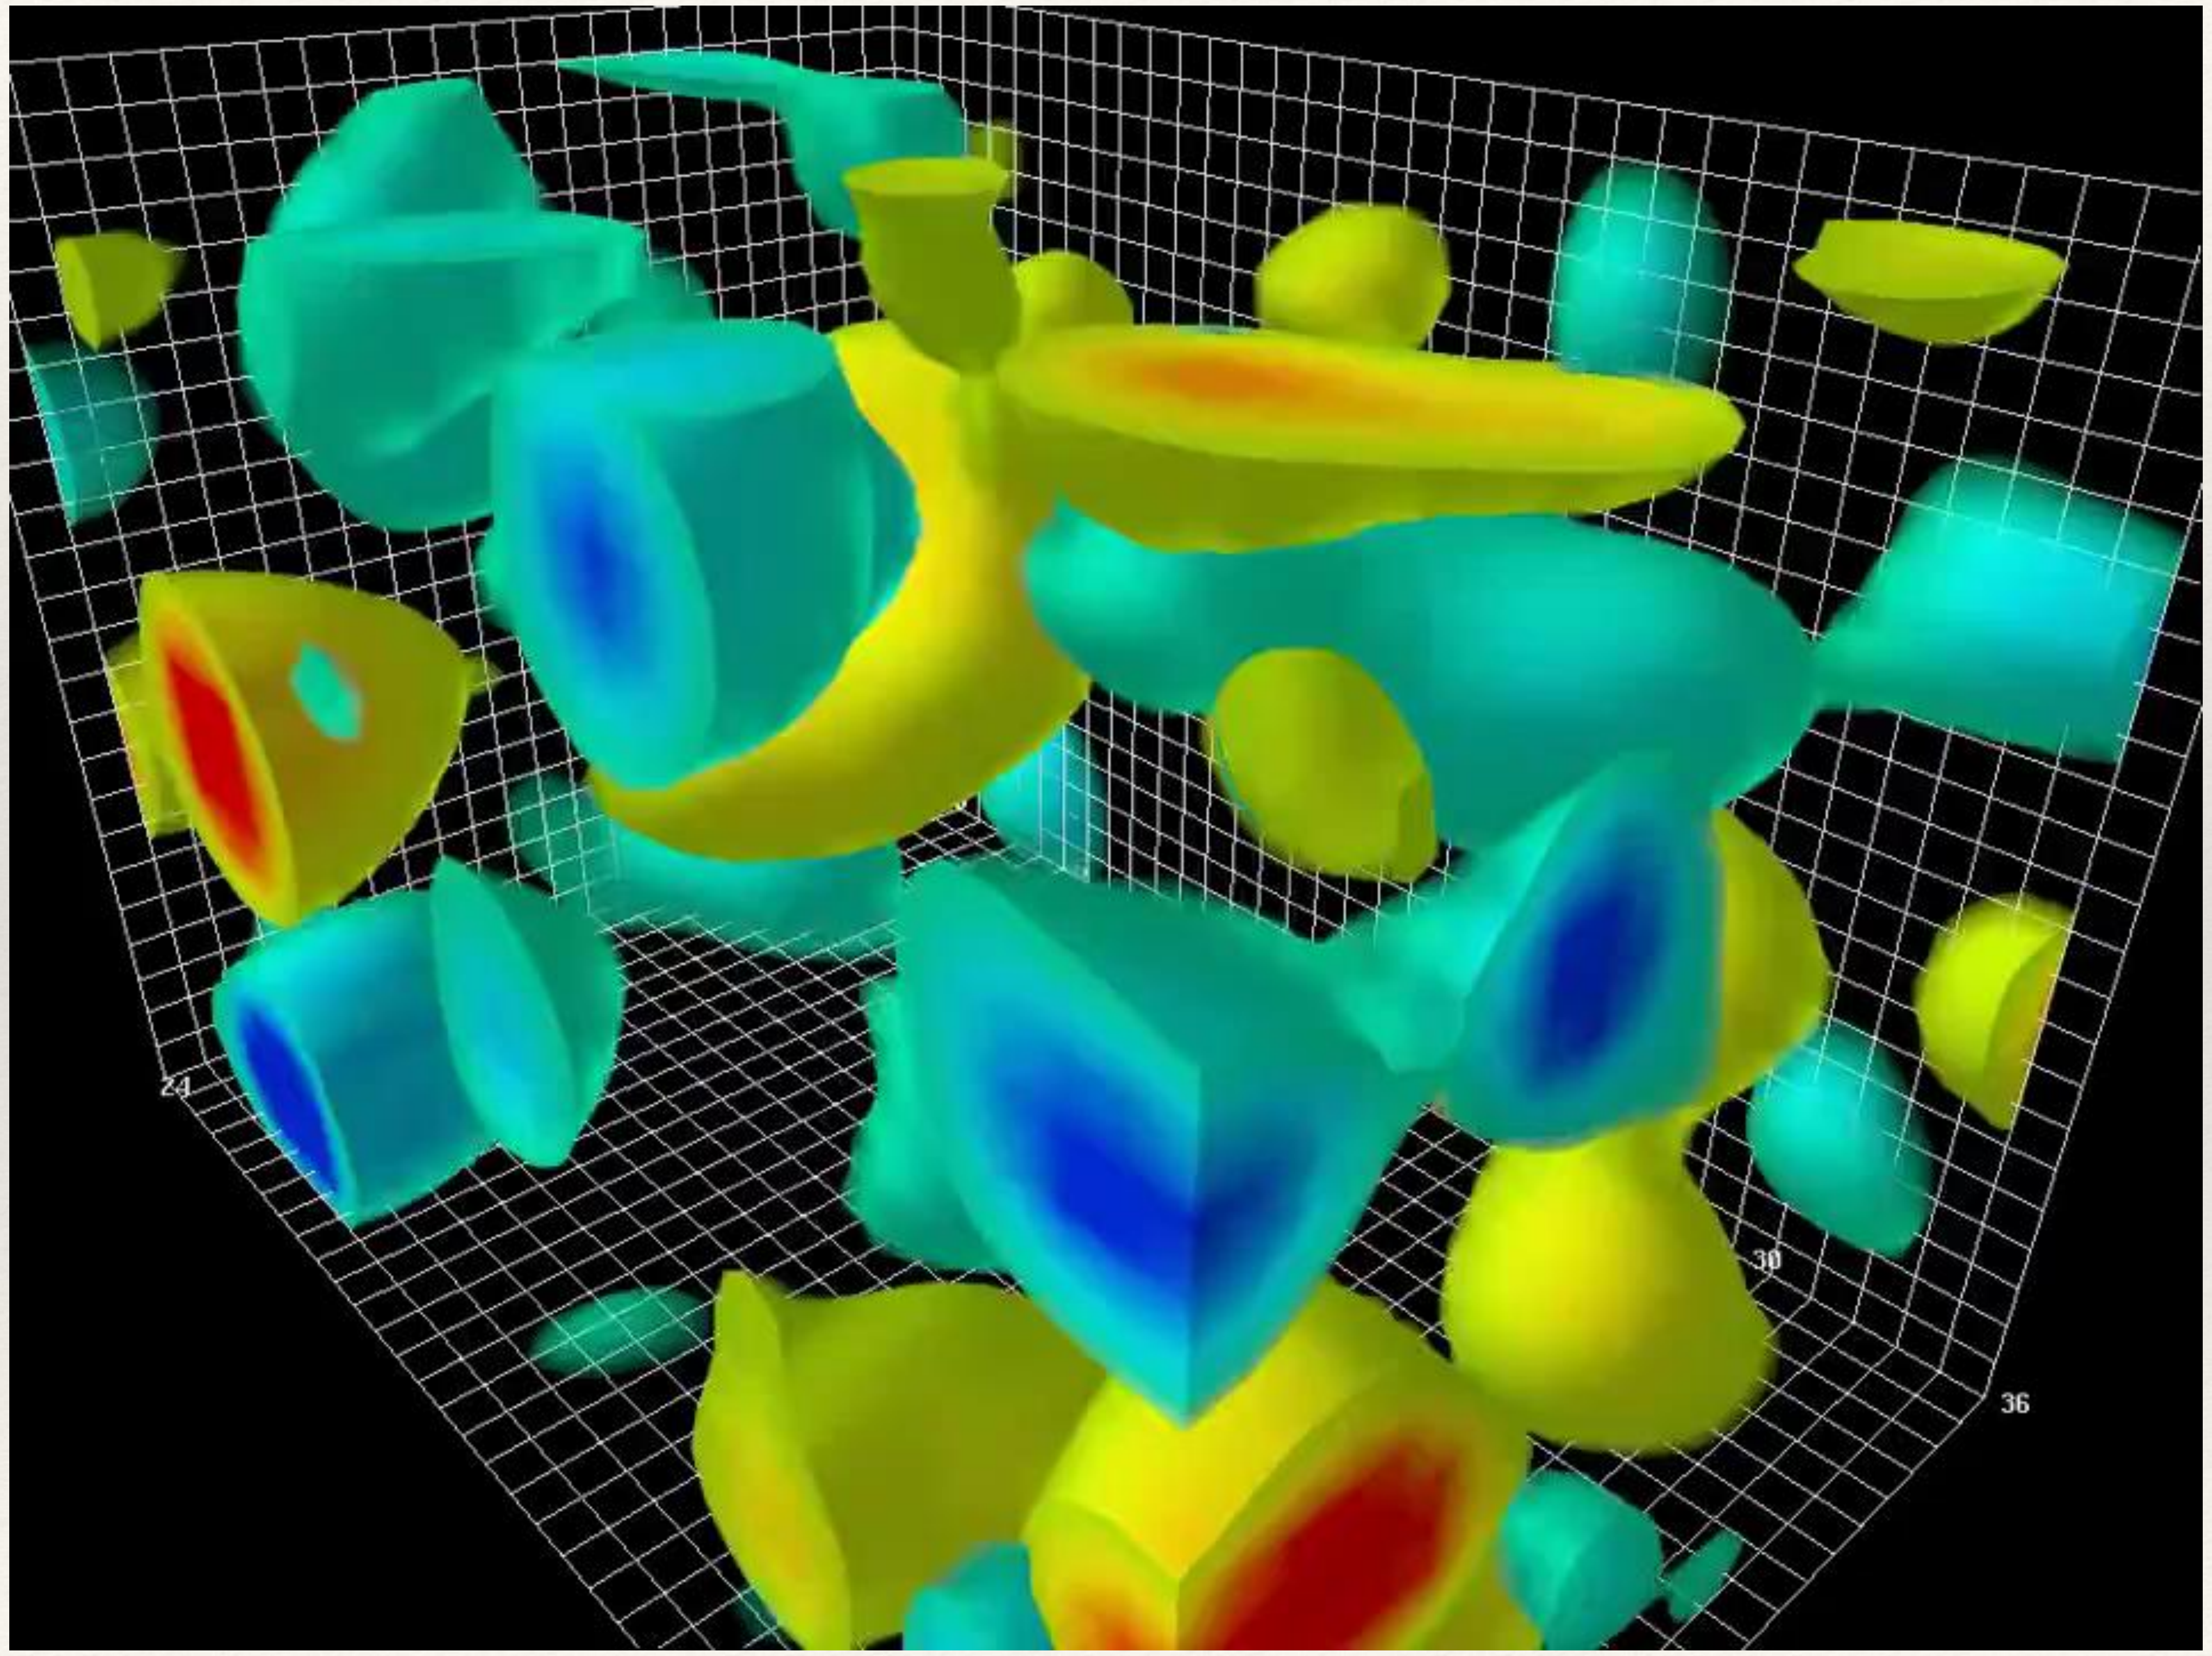

http://www.physics.adelaide.edu.au/cssm/lattice/

A multiscale problem

QUANTUM THEORY CENTER
SDU

Table of some baryons

Particle	Symbol	Quark Content	Mass MeV/c^2	Mean lifetime (s)	Decays to
Proton	p	uud	938.3	Stable	Unobserved
Neutron	n	ddu	939.6	885.7±0.8	$p + e^- + \bar{\nu}_e$
Delta	Δ^{++}	uuu	1232	6×10^{-24}	$\pi^+ + p$
Delta	Δ^+	uud	1232	6×10^{-24}	$\pi^+ + n$ or $\pi^0 + p$
Delta	Δ^0	udd	1232	6×10^{-24}	$\pi^0 + n$ or $\pi^- + p$
Delta	Δ^-	ddd	1232	6×10^{-24}	$\pi^- + n$
Lambda	Λ^0	uds	1115.7	2.60×10^{-10}	$\pi^- + p$ or $\pi^0 + n$
Sigma	Σ^+	uus	1189.4	0.8×10^{-10}	$\pi^0 + p$ or $\pi^+ + n$
Sigma	Σ^0	uds	1192.5	6×10^{-20}	$\Lambda^0 + \gamma$
Sigma	Σ^-	dds	1197.4	1.5×10^{-10}	$\pi^- + n$

Table of some mesons

Particle	Symbol	Anti-particle	Quark Content	Mass MeV/c^2	Mean lifetime (s)	Principal decays
Charged Pion	π^+	π^-	$u\bar{d}$	139.6	2.60×10^{-8}	$\mu^+ + \nu_\mu$
Neutral Pion	π^0	Self	$u\bar{u} - d\bar{d}$	135.0	0.84×10^{-16}	2γ
Charged Kaon	K^+	K^-	$u\bar{s}$	493.7	1.24×10^{-8}	$\mu^+ + \nu_\mu$ or $\pi^+ + \pi_0$
Neutral Kaon	K^0	$\bar{K}^0$	$d\bar{s}$	497.7		

Table of Quarks

Name	Symbol	Charge (e)	Spin	Mass MeV/c^2	Strangeness	Baryon number	Lepton number
up	u	+2/3	1/2	1.7-3.3	0	1/3	0
down	d	-1/3	1/2	4.1-5.8	0	1/3	0
strange	s	-1/3	1/2	101	-1	1/3	0
charm	c	+2/3	1/2	1270	0	1/3	0
bottom	b	-1/3	1/2	4190-4670	0	1/3	0
top	t	+2/3	1/2	172000	0	1/3	0

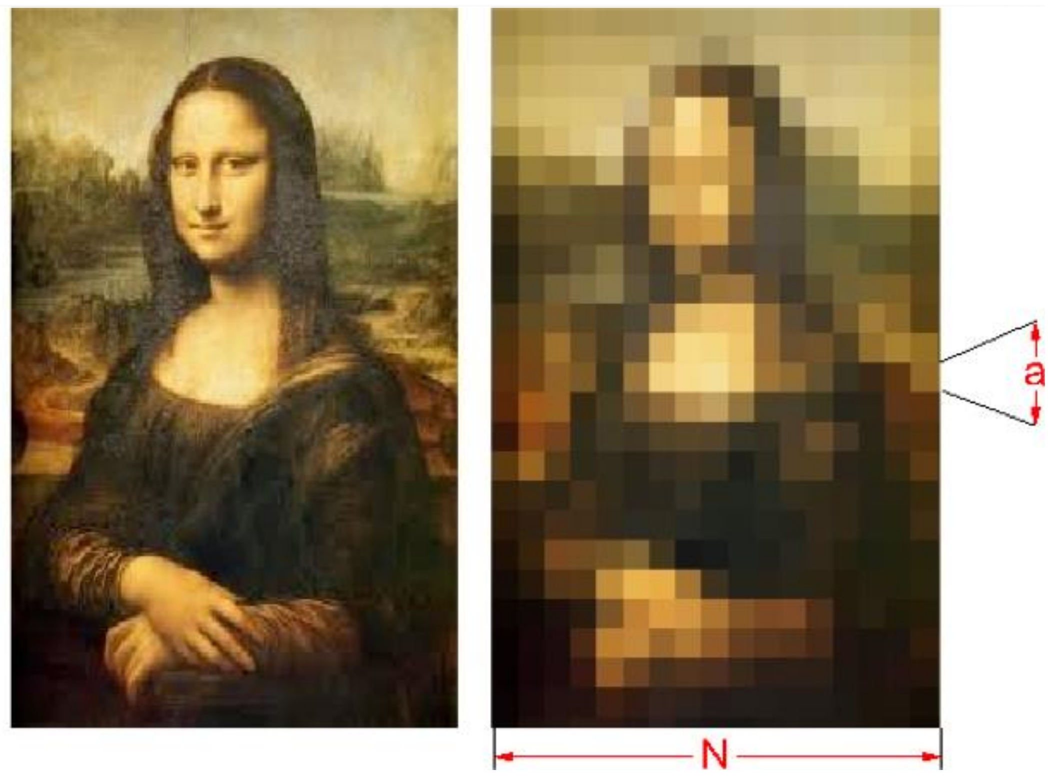

Where N and a are problem dependent

the cost in one task: invert the Dirac operator

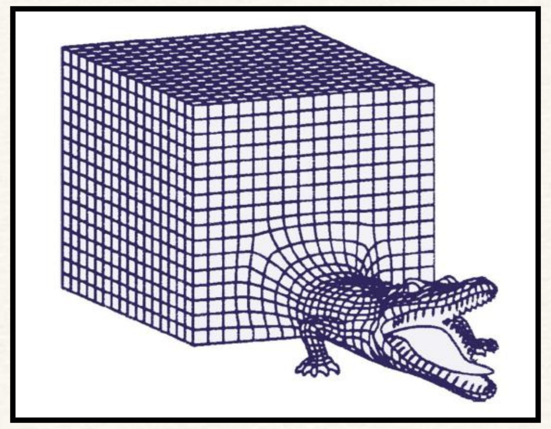

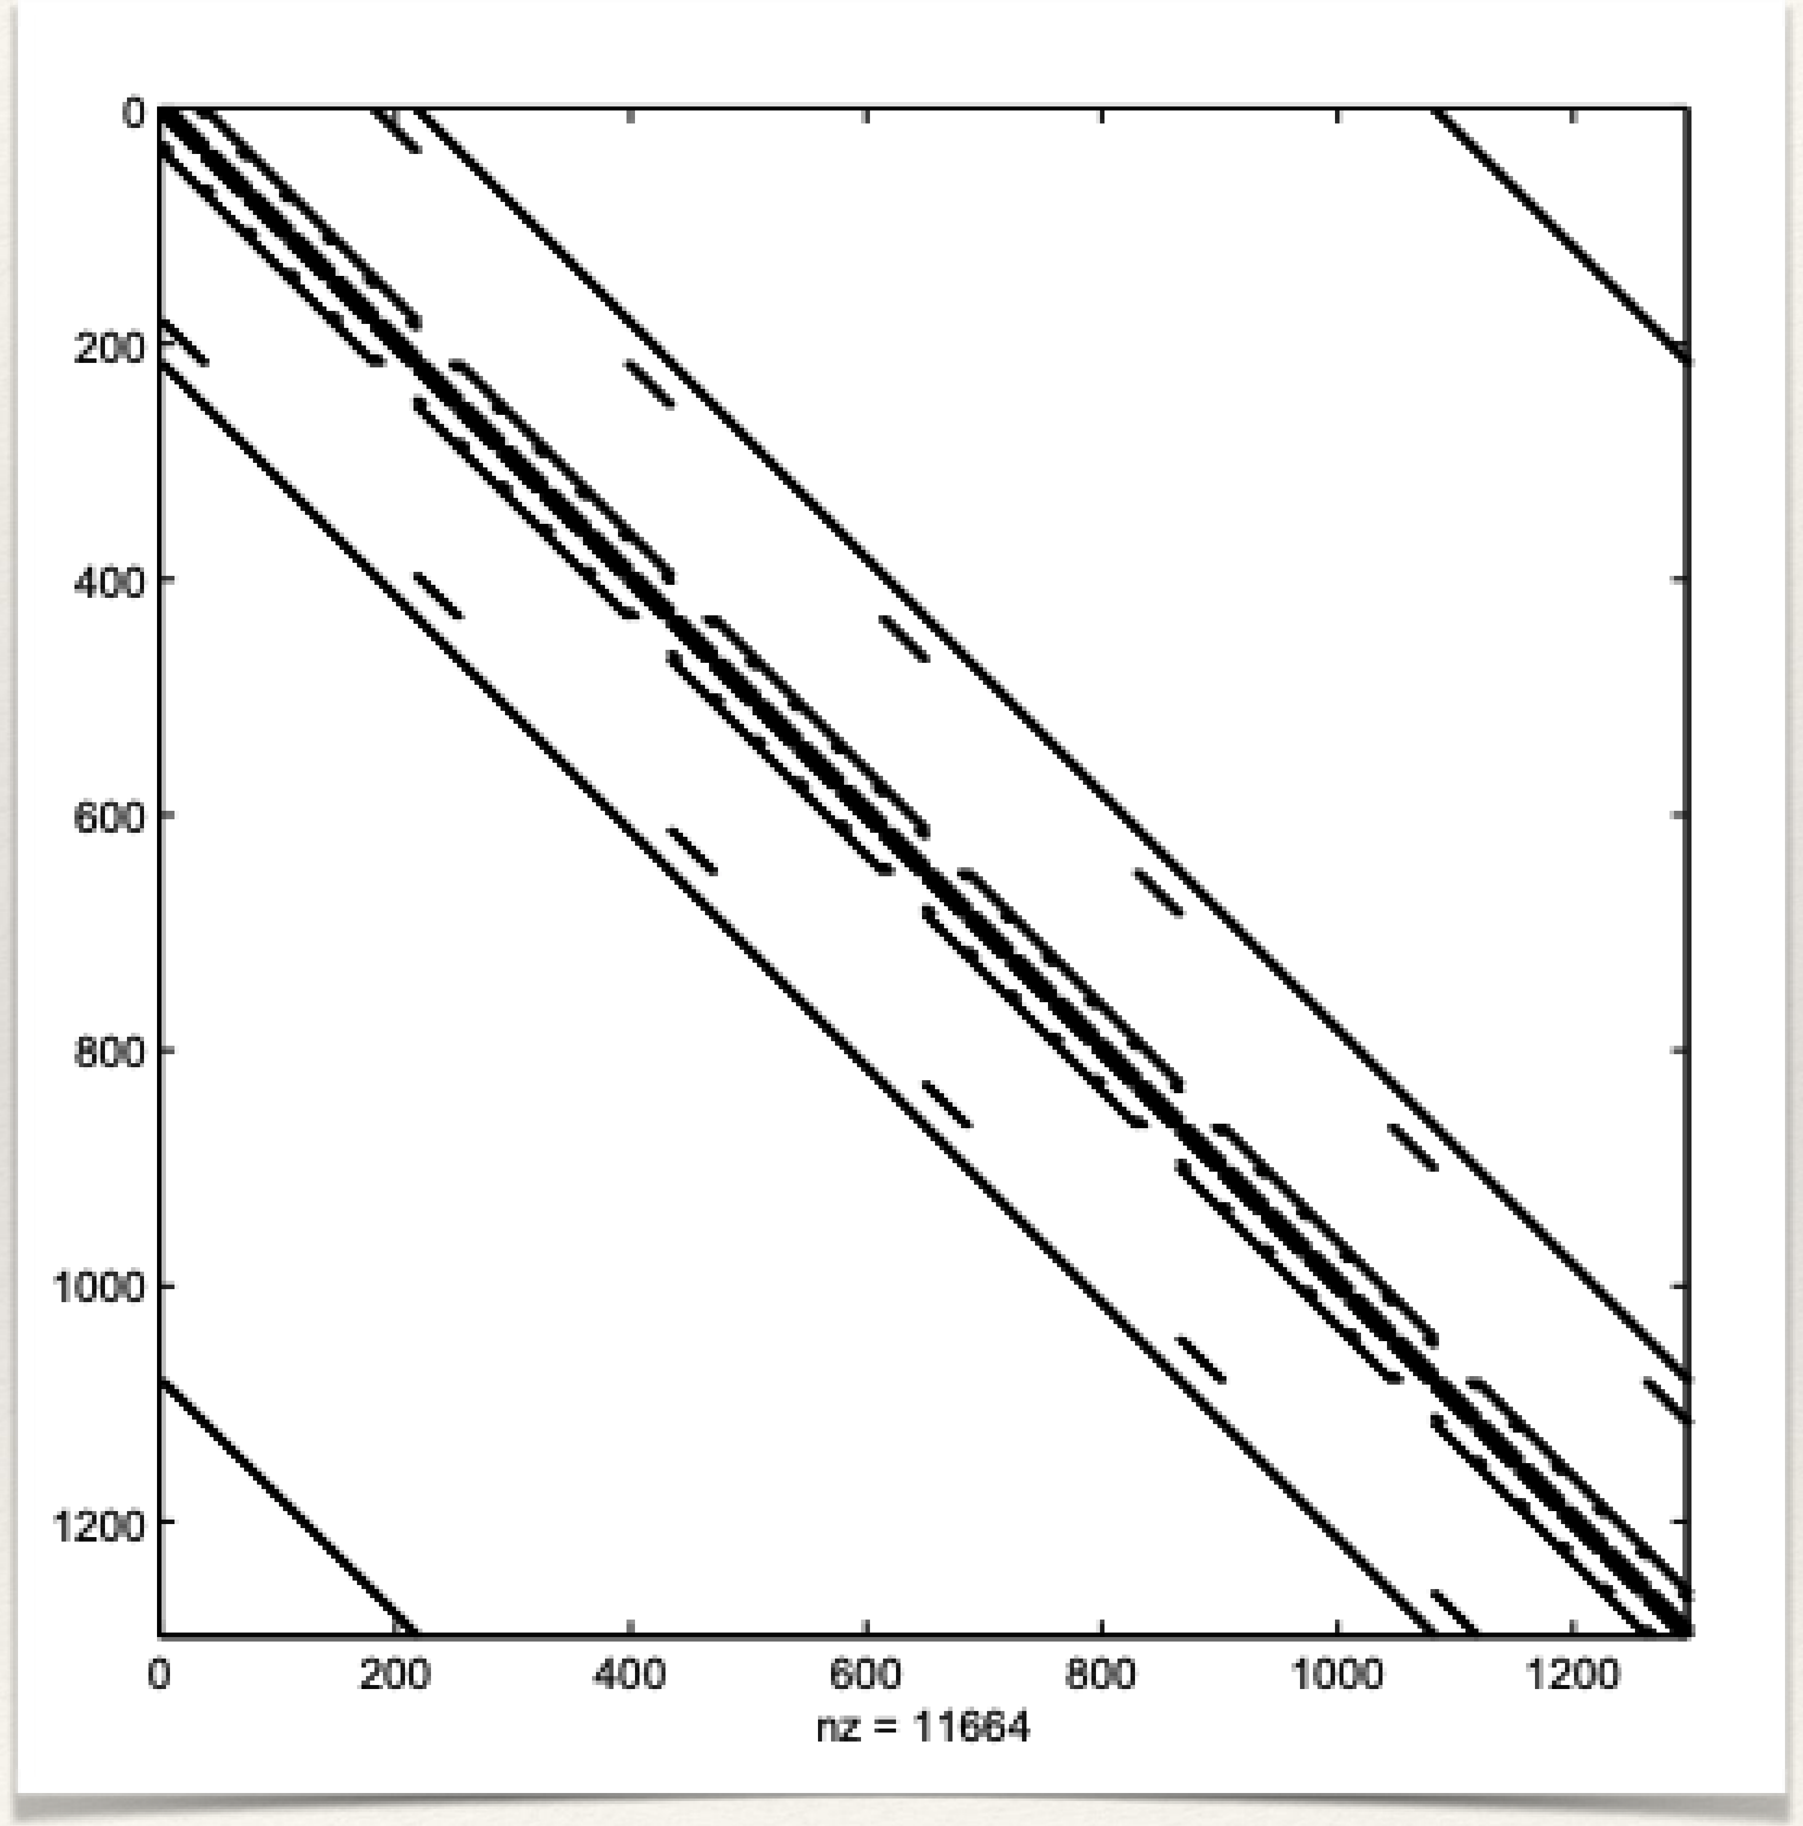

In mathematics and quantum mechanics, a **Dirac operator** is a differential operator that is a formal square root, or half-iterate, of a second-order operator such as a Laplacian.

It describes all spin-1/2 massive particles, called "Dirac particles", such as electrons and quarks

$N \sim 10^{10}$ elements

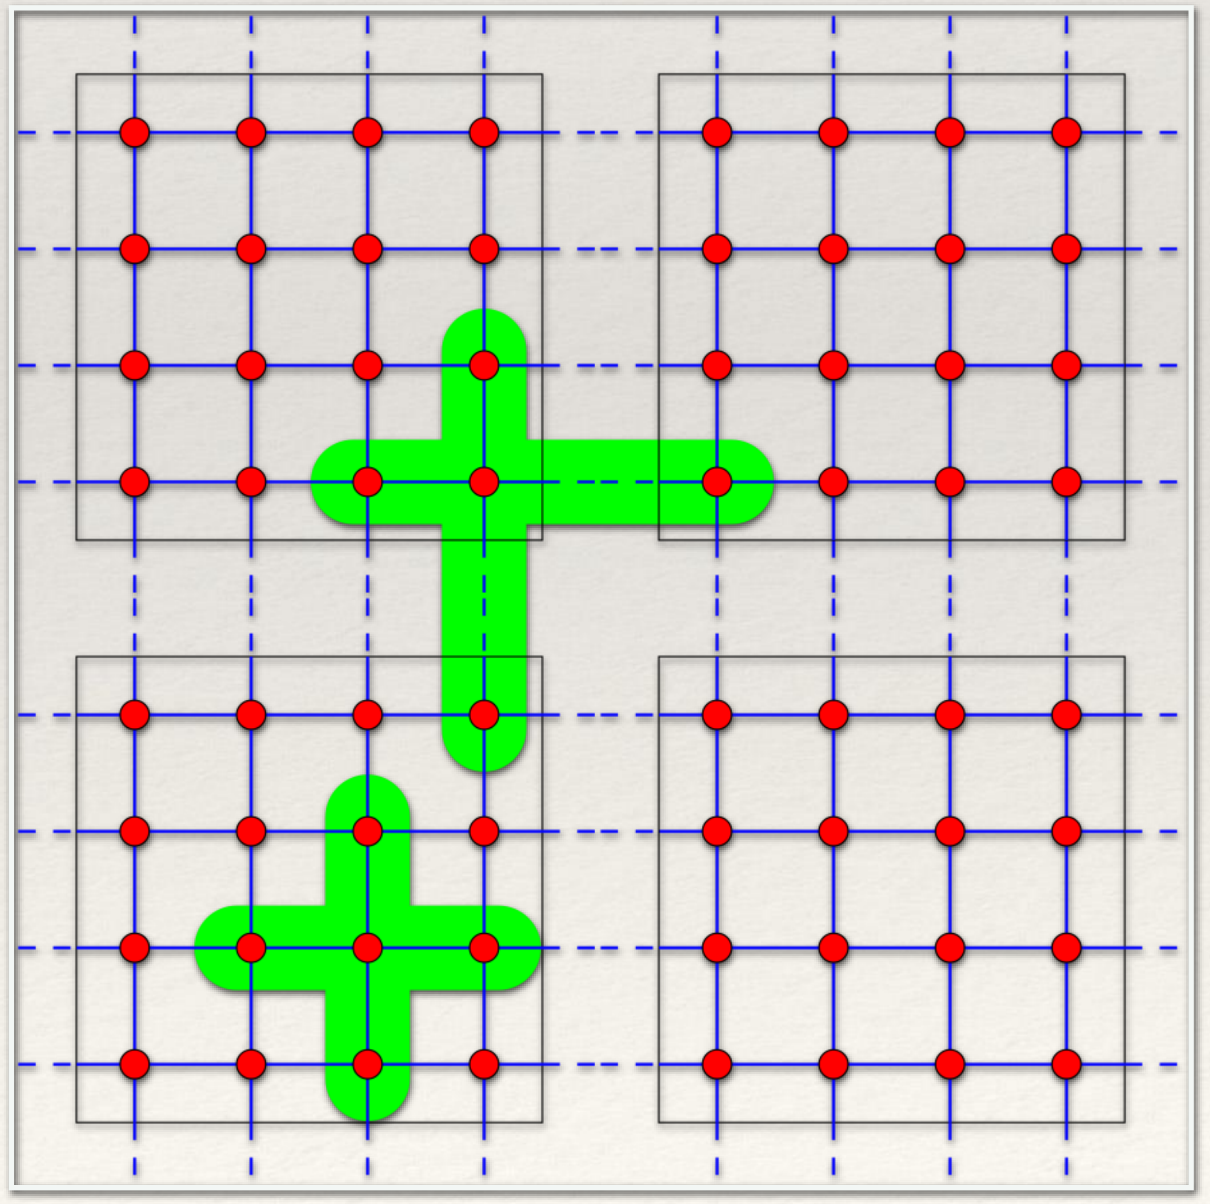

"Numerical simulations have the reputation of being an approximate method that mainly serves to obtain qualitative information on the behaviour of complex systems.

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,

"Lattice QCD -- From Quark Confinement to Asymptotic freedom"
Plenary presentation at International Conference on Theoretical Physics , Paris UNESCO, July 2002
http://luscher.web.cern.ch/luscher/lectures/Paris02.pdf

Lattice QCD today

- Competitive Lattice simulations take place in the framework of collaborations
  - Between 3 and 10 research groups
  - USA, Europe, Japan
  - Among the most successful HPC consumers all over the world
- Coordinated to have access to substantial amounts of CPU time and human resources to cope with:
  - Code development
  - Implementation of simulations
  - Data Analysis
- Enormous ingenuity is needed to progress on the algorithms encompassing the physics of the problem
- Remains a task for individuals and small teams

- USQCD (USA) : https://www.usqcd.org/
- RBC-UKQCD (USA, UK, Japan): http://rbc.phys.columbia.edu/
- MILC (USA + Europe): http://www.physics.utah.edu/~detar/milc/
- HPQCD (USA + Europe): http://www.physics.gla.ac.uk/HPQCD/
- ALPHA (EU): https://wwwzeuthen.desy.de/alpha/
- BMW (EU): http://www.bmw.uniwuppertal.de/
- CP-PACS, JLQCD (Japan - Tsukuba)
- openQCD (EU+USA+Taiwan): https://openlat1.gitlab.io/

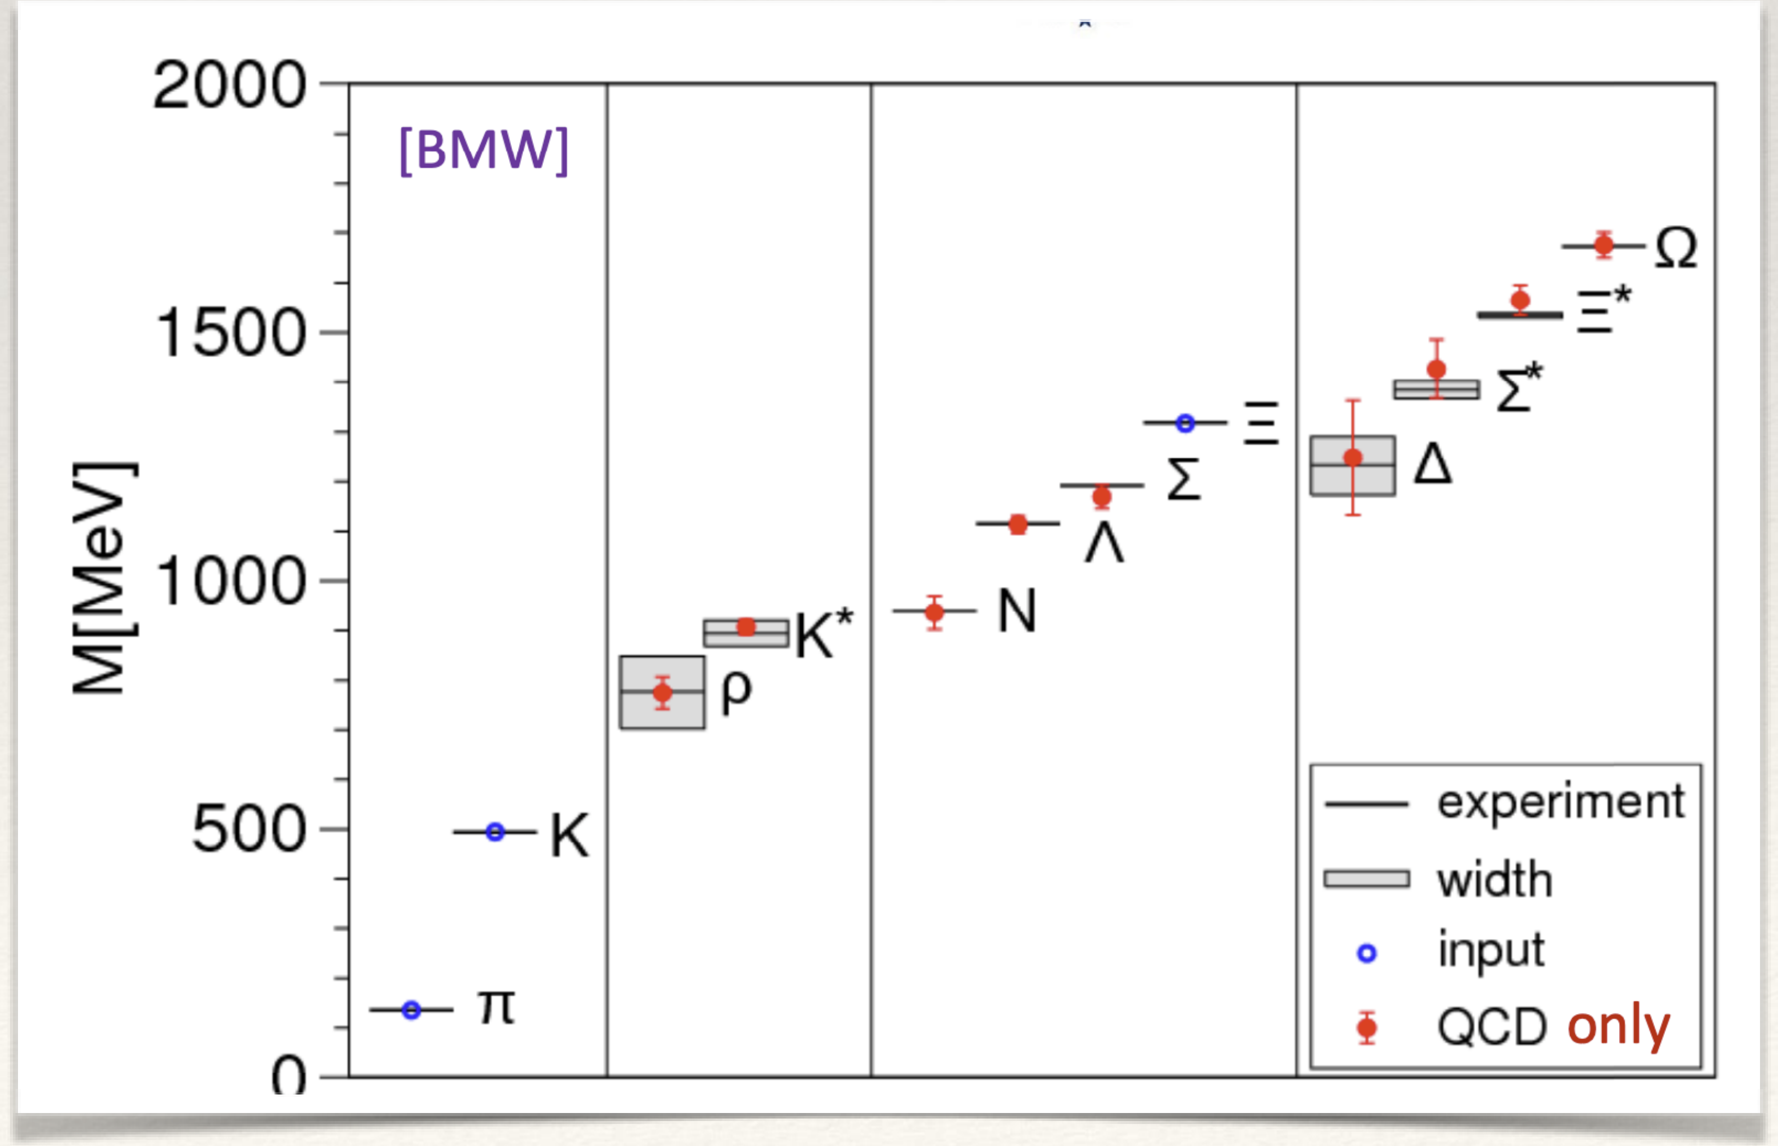

efficiency has always been the key

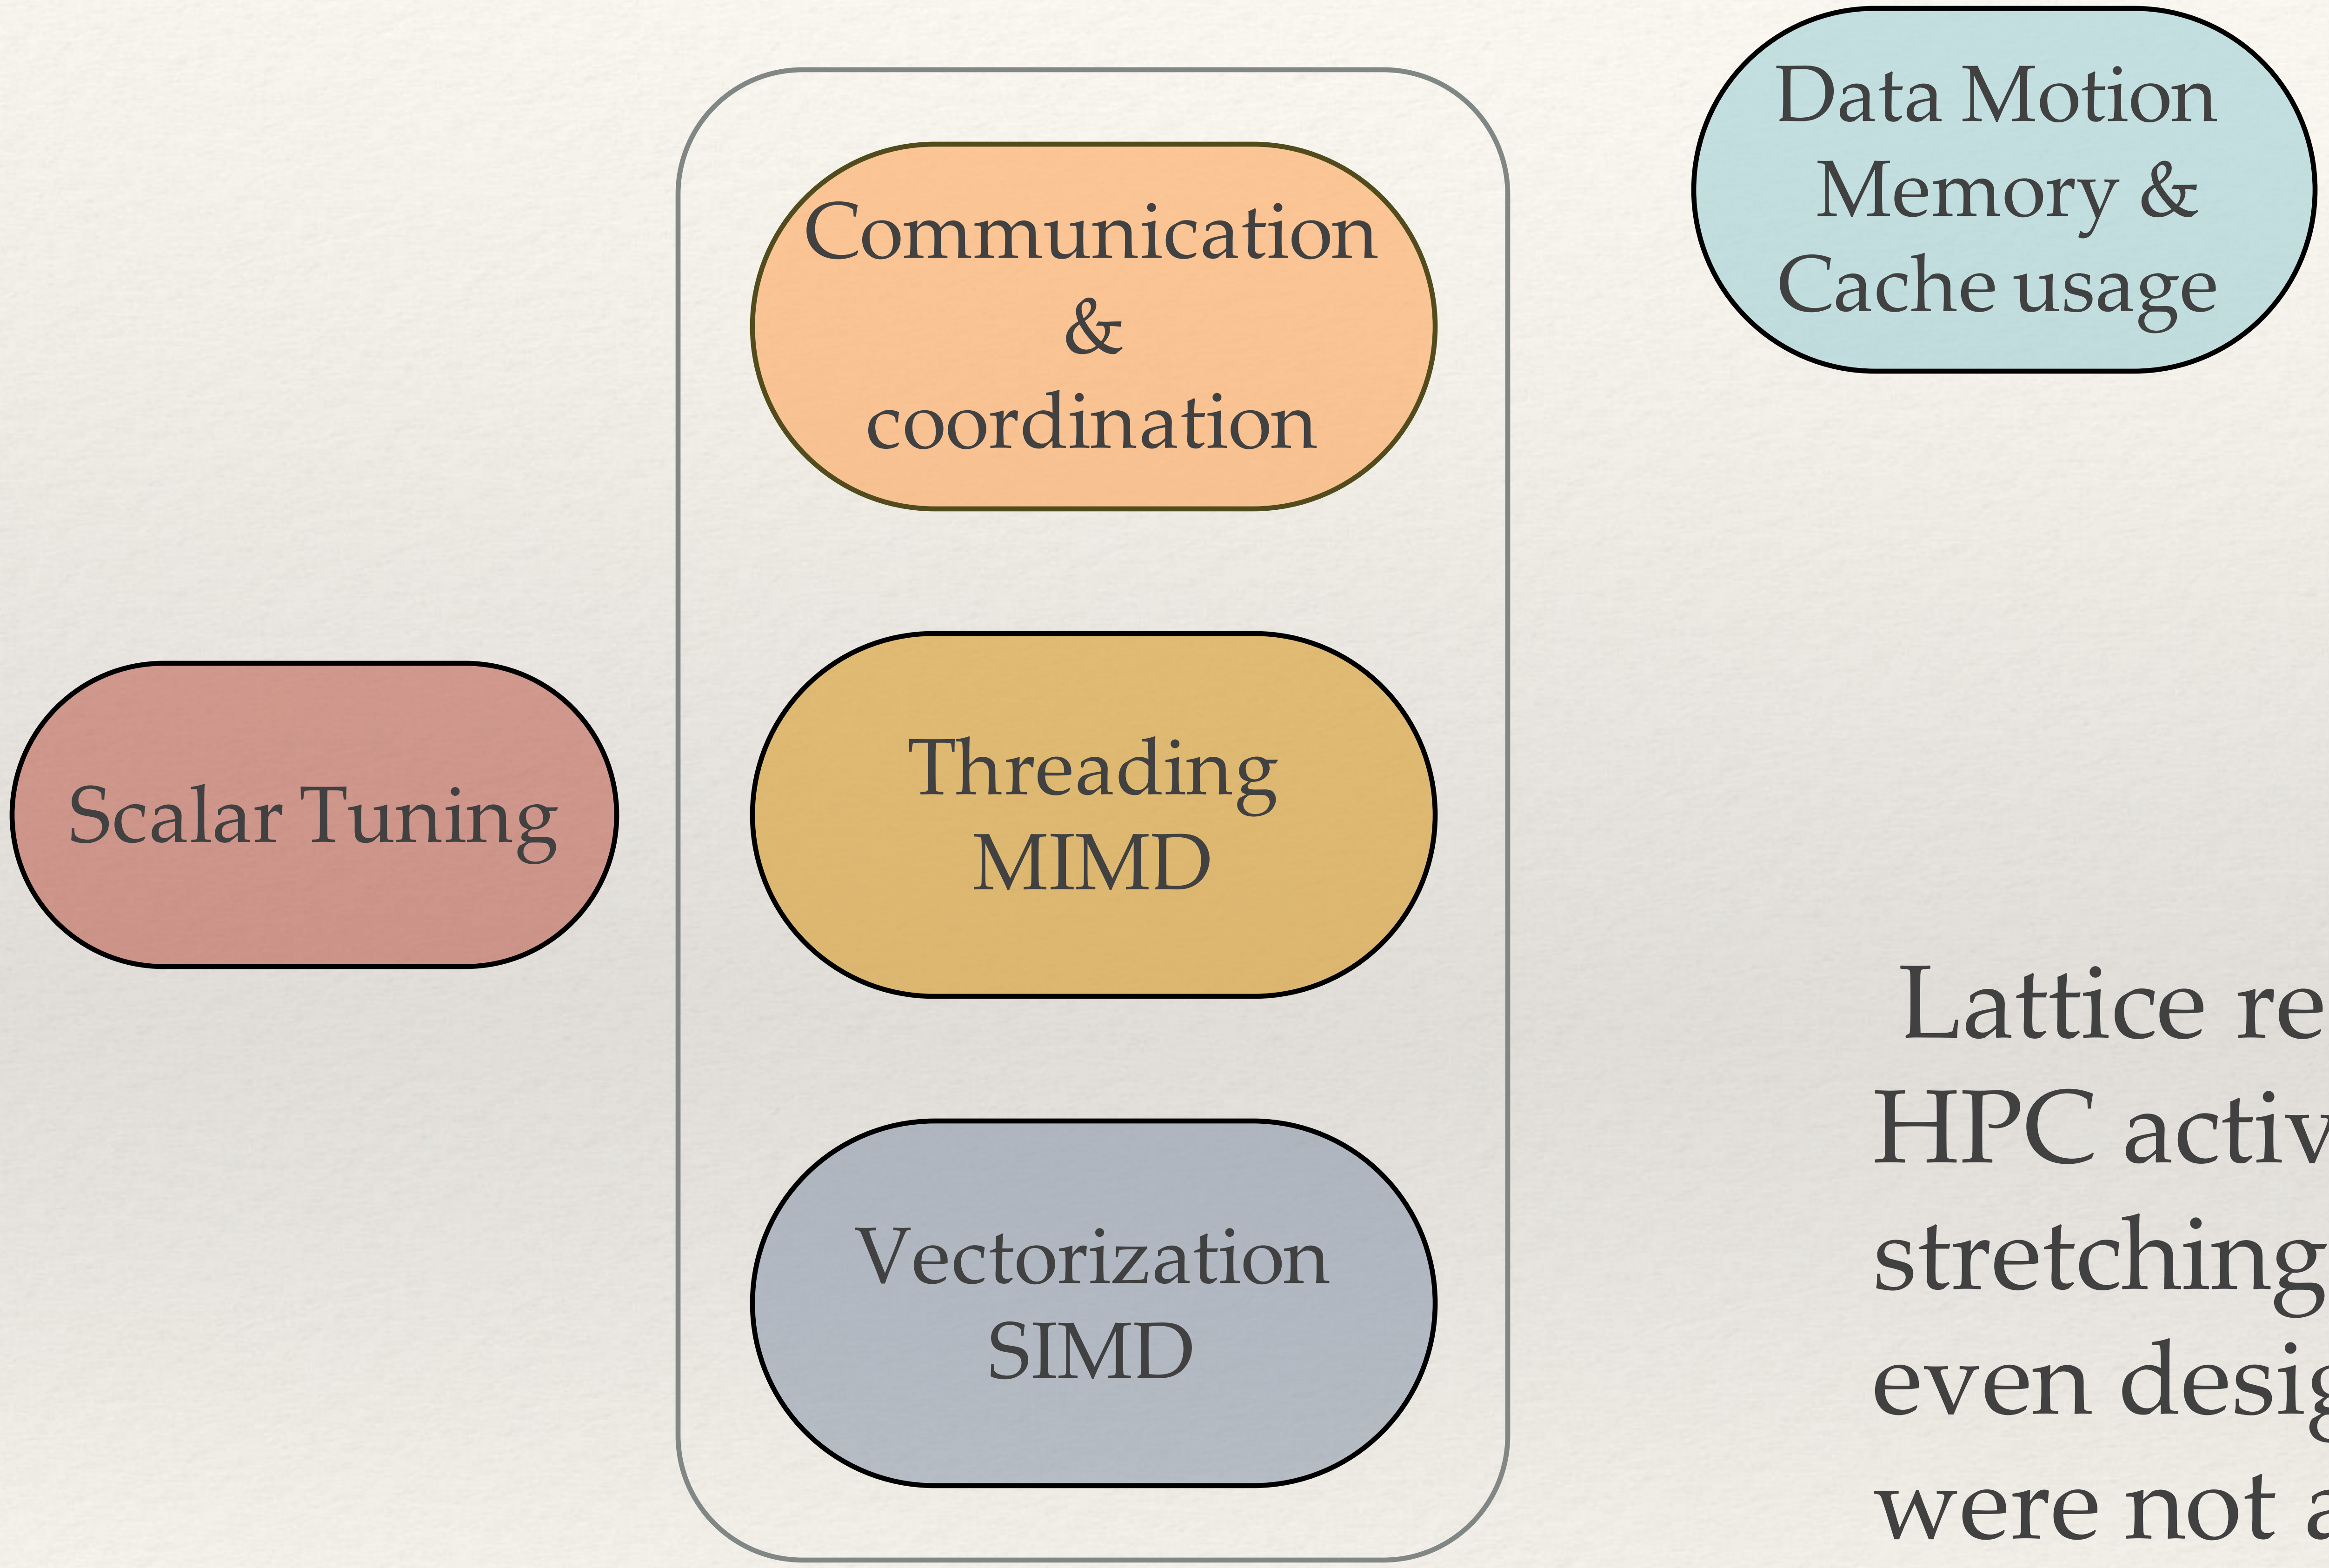

The good programmer will saturate one of the bounds (Cache Memory or Network) … and then be unable to do any better

Lattice researchers have always been at the forefront of HPC activities designing innovative algorithms, stretching hardware and numerical capabilities and even designing new hardware when other solutions were not attainable.

Performances of various codes

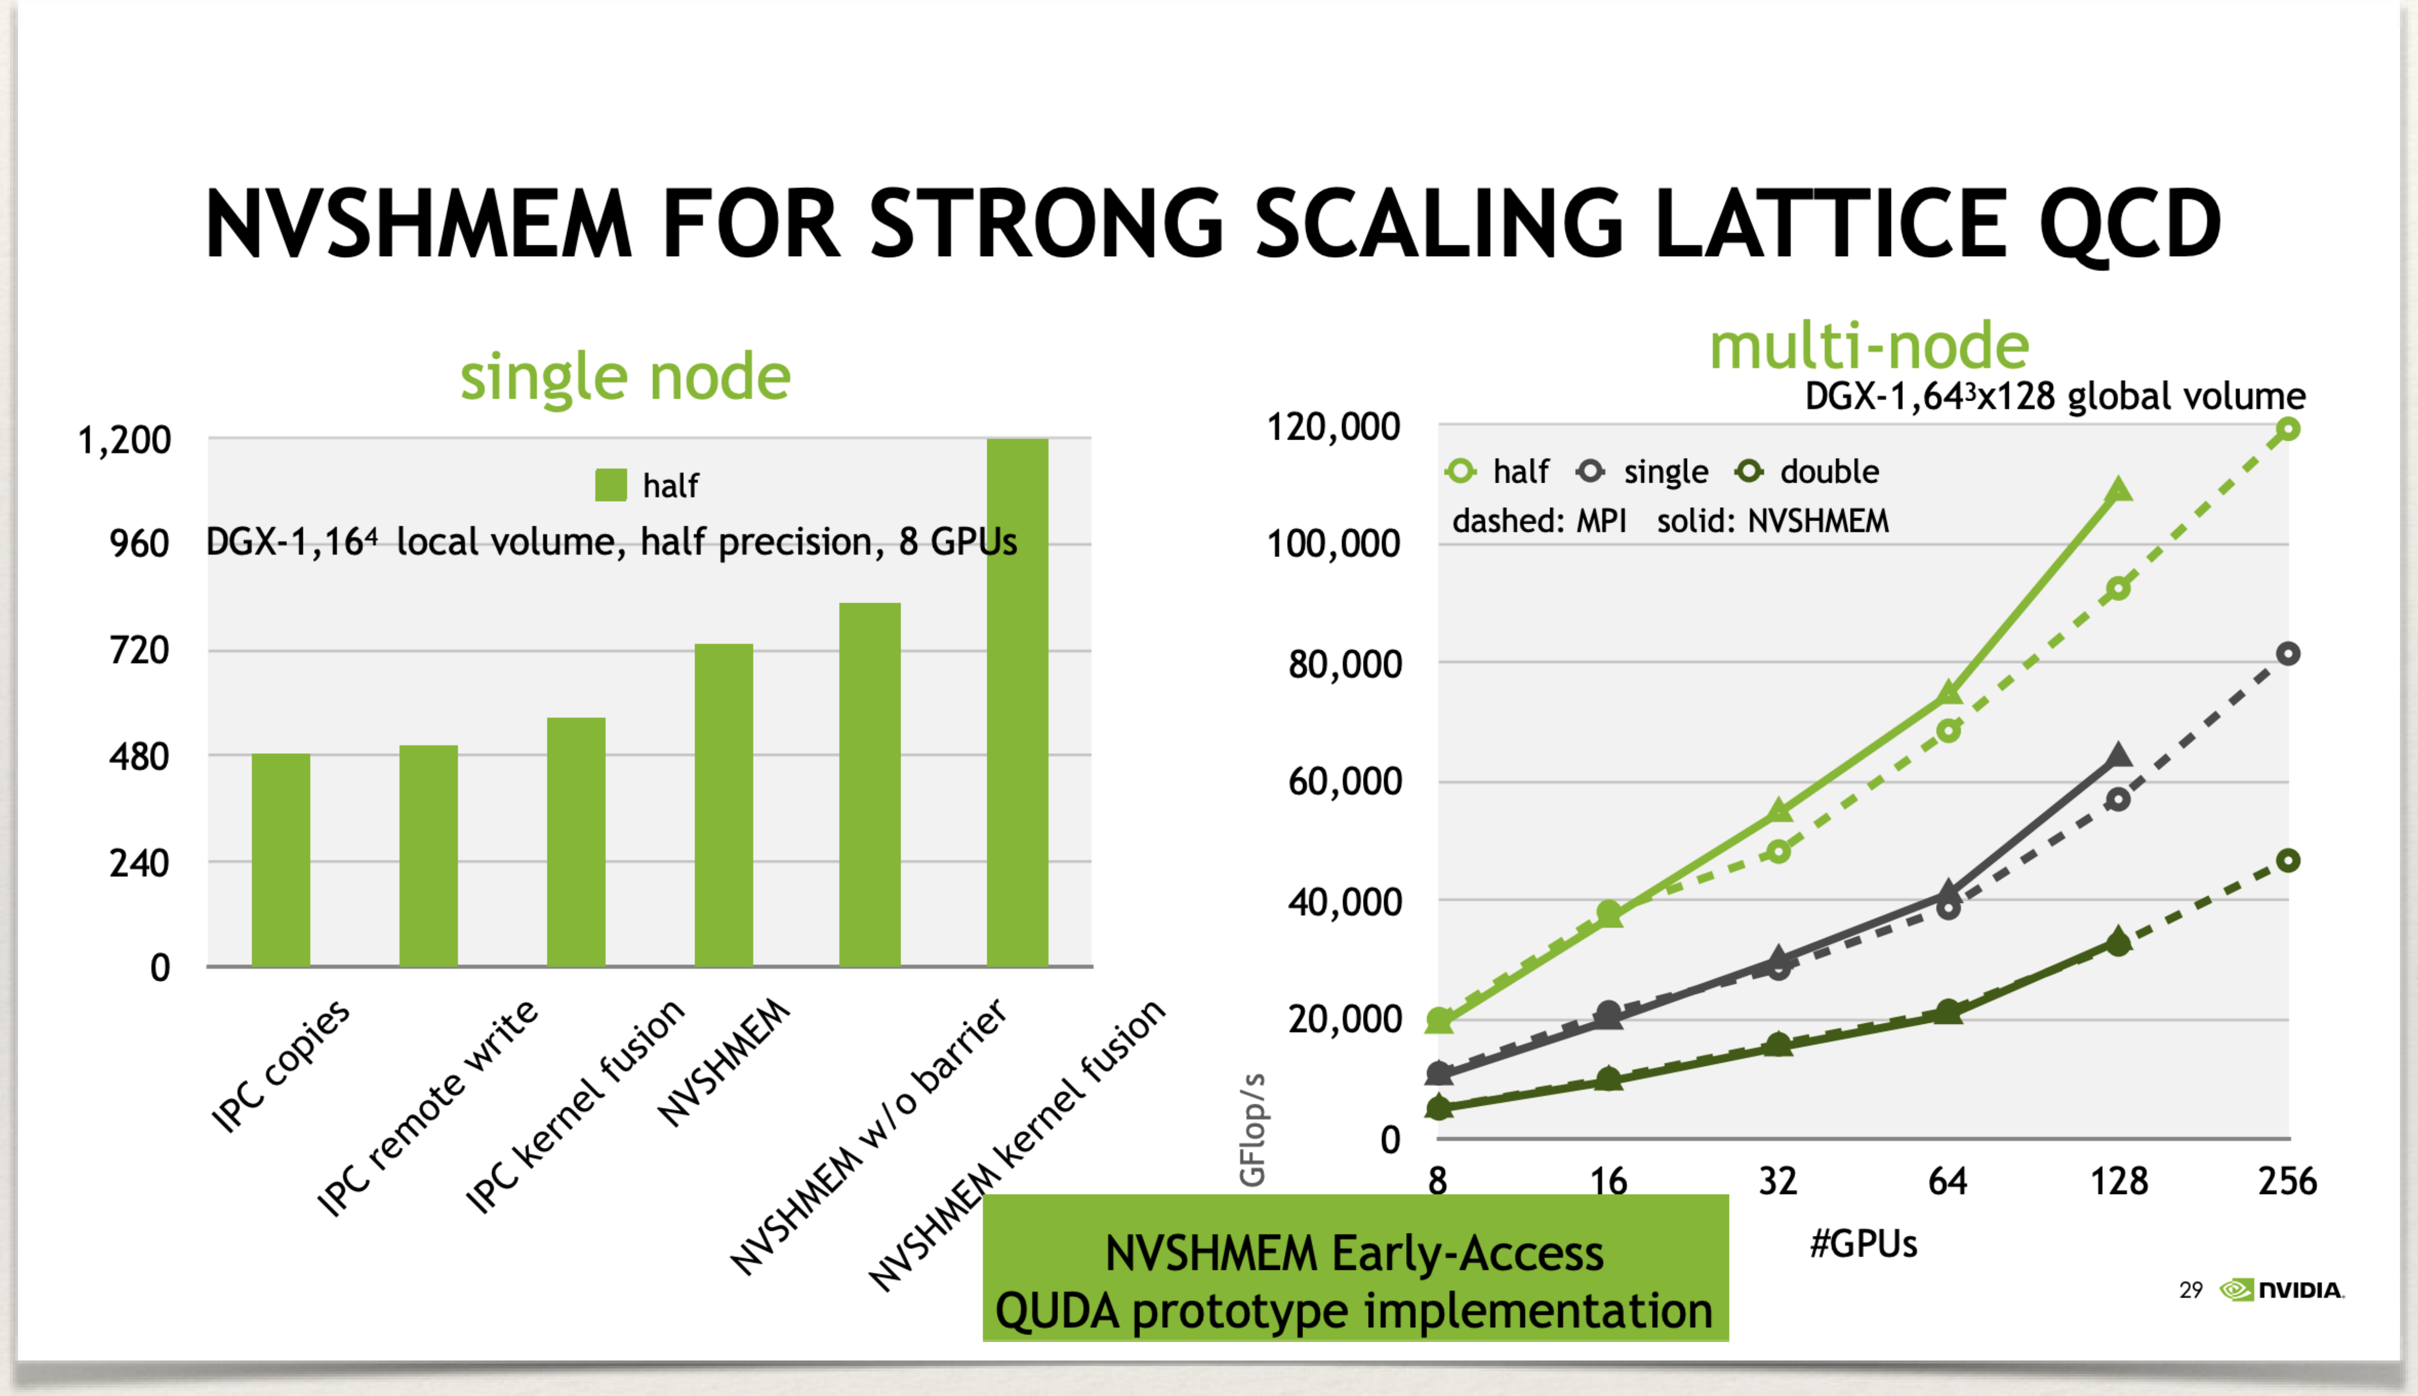

QUDA

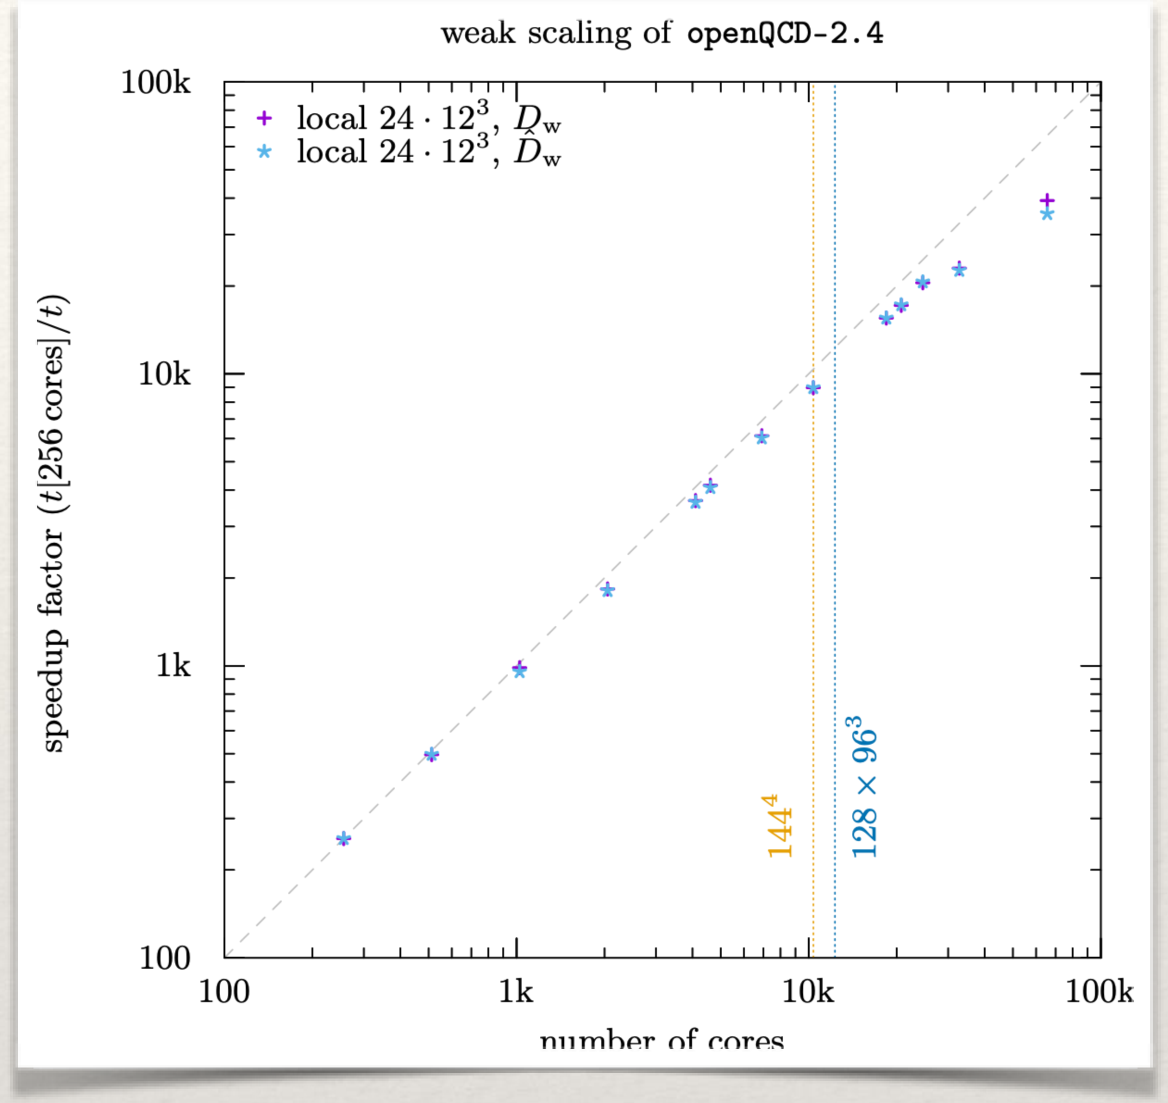

openQCD

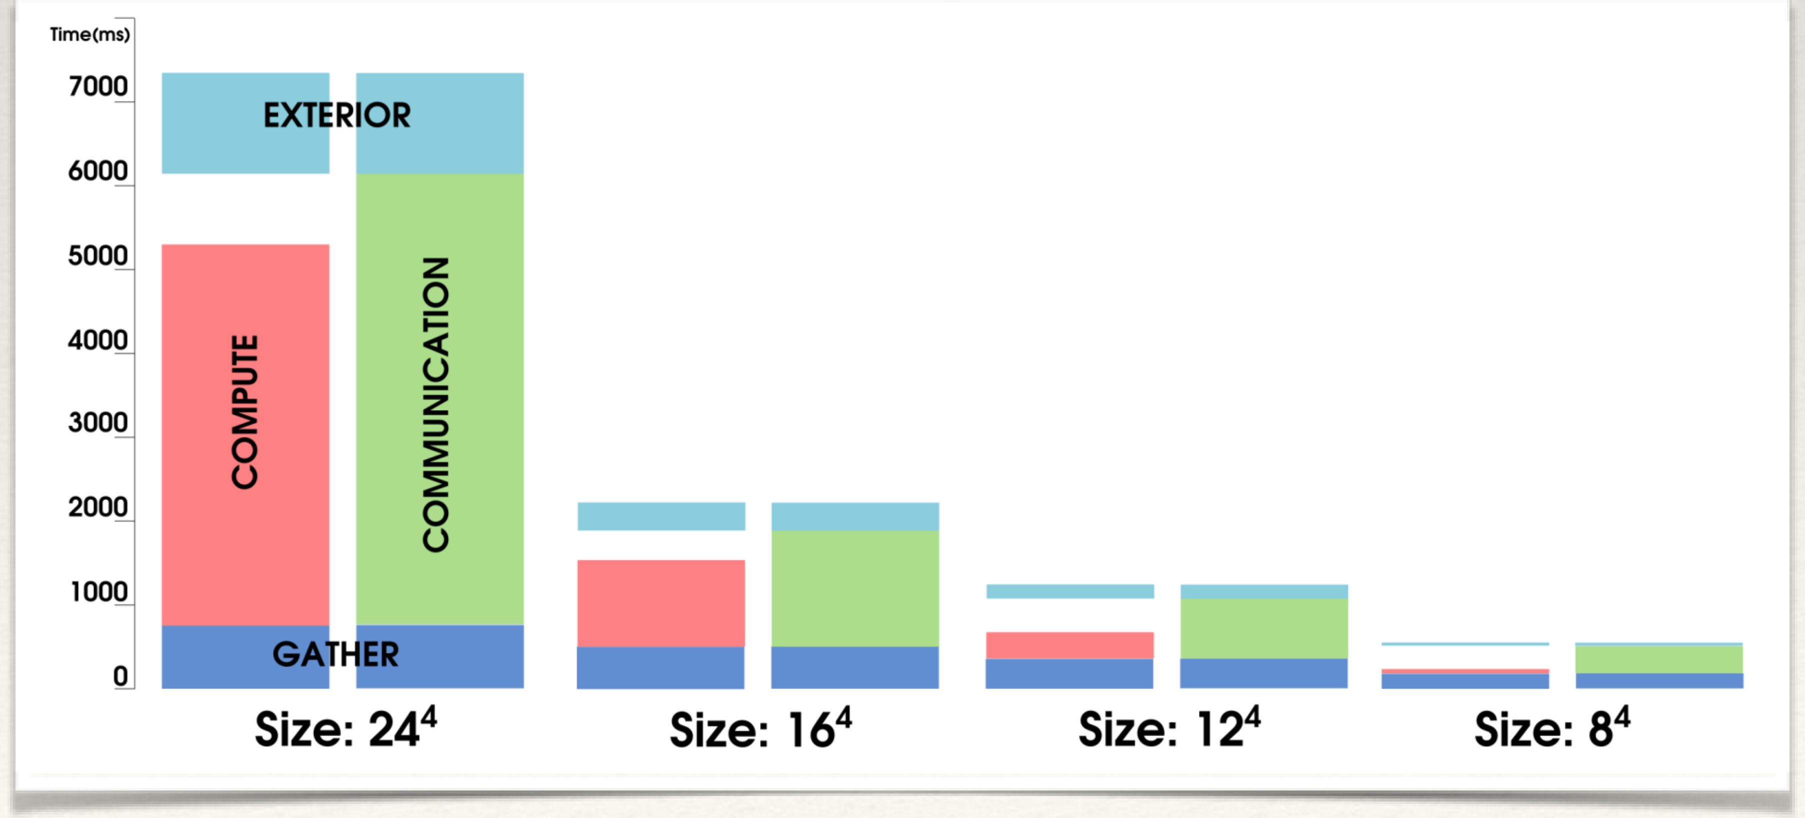

GRID

hardware for QCD

QUANTUM THEORY CENTER
SDU

CERN-TH.4283/85
ROM2F/85/28

THE APE PROJECT: A GIGAFLOP PARALLEL PROCESSOR FOR LATTICE CALCULATIONS(*)

P.Bacilieri
INFN-CNAF, Bologna, Italy

S.Cabasino, F.Marzano, P.Paolucci, S.Petrarca, G.Salina
Dipartimento di Fisica, Università di Roma "La Sapienza"
INFN - Sezione di Roma, Roma, Italy

N.Cabibbo ... F.Marinari, G.Parisi
Dipartimento di
INFN

F.Costantini, G.Fio
Diparti
INFN

Dipartim
INFN

Dipartim
INFN

Rockf

Status of the *QCDSP* project

Dong Chen,[a *†] Ping Chen,[a *] Norman H. Christ,[a *] Robert G. Edwards,[b *] George R. Alan Gara,[c ‡] Sten Hansen,[d] Chulwoo Jung,[a *] Adrian L. Kaehler,[a *] Anthony D. Kenn Gregory W. Kilcup,[e *] Yubing Luo,[a *] Catalin I. Malureanu,[a *] Robert D. Mawhinney,[a] John Parsons,[c ‡] ChengZhong Sui,[a *] Pavlos M. Vranas[a *§] and Yuri Zhestkov[a *]

[a]Columbia University, New York, NY

[b]SCRI, Florida State University, Tallahassee, FL

[c]Nevis Laboratories, Columbia University, Irvington, NY

[d]Fermi National Accelerator Laboratory, Batavia, IL

[e]Ohio State University, Columbus, OH

We describe the completed 8,192-node, 0.4Tflops machine at Columbia as well as the 12,288-node, 0.6Tflops machine assembled at the RIKEN Brookhaven Research Center. Present performance as well as our experience in commissioning these large machines is presented. We outline our on-going physics program and explain how the configuration of the machine is varied to support a wide range of lattice QCD problems, requiring a variety of machine sizes. Finally a brief discussion is given of future prospects for large-scale lattice QCD machines.

1. INTRODUCTION

The large computational requirements of lattice QCD coupled with the enormous cost/performance advantages that can be obtained with specially configured computer hardware have encouraged the design and construction of a variety of purpose-built machines over the past nearly 18 years. The *QCDSP* machines now being completed by our collaboration[1] represent a continued development in this direction.

The most recent, 12,288-node machine at the RIKEN Brookhaven Research Center has a construction cost of approximately \$1.8M and a peak speed of 0.6Tflops or a cost per peak perfor- about 25% efficiency on a lattice volume of 4^4 sites per node which corresponds to a dollar per delivered Mflops figure of \$13.6/Mflops[2].

Achieving this level of performance required considerable care when writing the routine which applies the Wilson Dirac operator to a spinor field. (We expect that with more effort, this efficiency may increase somewhat further and that our staggered code will achieve a similar level of performance.) This high performance code is written in assembly language and uses many of the special hardware features provided to boost efficiency. However, the bulk of the conjugate gradient code which applies this Dirac operator, as well as the hybrid Monte Carlo evolu-

N-H
ELSEVIER
Parallel Computing 25 (1999) 1635–1661
PARALLEL COMPUTING
www.elsevier.com/locate/parco

CP-PACS: A massively parallel processor at the University of Tsukuba

Kisaburo Nakazawa[a,*], Hiroshi Nakamura[b], Taisuke Boku[c], Ikuo Nakata[d], Yoshiyuki Yamashita[c]

[a] *Department of Electronics and Computer Science, Meisei University, 2-590 Nagafuchi, Ohme, Tokyo 198-8655, Japan*
[b] *Research Center for Advanced Science and Technology, The University of Tokyo, 4-6-1 Komaba, Meguro-ku, Tokyo 153-8904, Japan*
[c] *Institute of Information Sciences and Electronics, University of Tsukuba, 1-1-1 Tennohdai, Tsukuba, Ibaraki 305-8577, Japan*
[d] *University of Library and Information Science, 1-2 Kasuga, Tsukuba, Ibaraki 305-8550, Japan*

Received 2 October 1998; received in revised form 16 December 1998

Abstract

Computational
parallel processor
the University of T
2048 processing u
PACS is 614.4 Gfl
which at that time
CP-PACS has
each node process

QCDOC: A 10 Teraflops Computer for Tightl Calculations

P.A. Boyle, D. Chen, N.H. Christ, M. Clark, S.D. Cohen, C. Cristian, Z. Dong, A. Gara, B. Joo, C. Jung, C. Kim, L. Levkova, X. Liao, G. Liu, R.D. Mawhinney, S. Ohta, K. Petrov, T. Wettig, A. Yamaguchi

July 26, 2004

Abstract

Numerical simulations of the strong nuclear force, known as quantum chromodynamics or QCD, have proven to be a demanding, forefront problem in high-performance computing. In this report, we describe a new computer, QCDOC (QCD On a Chip), designed for optimal price/performance in the study of QCD. QCDOC uses a six-dimensional, low-latency mesh network to connect processing nodes, each of which includes a single custom ASIC, designed by our collaboration and built by IBM, plus DDR SDRAM. Each node has a peak speed of 1 Gigaflops and two 12,288 node, 10+ Teraflops machines are to be completed in the fall of 2004. Currently, a 512 node machine is running, delivering efficiencies as high as 45% of peak on the conjugate gradient solvers that dominate our calculations and a 4096-node machine with a cost of \$1.6M is under construction. This should give us a price/performance less than \$1 per sustained Megaflops.

Home / Proceedings / SAAHPC / SAAHPC 2011

Application Accelerators in High-Performance Computing, Symposium on

QUonG: A GPU-based HPC System Dedicated to LQCD Computing

Year: 2011, Pages: 113-122
DOI Bookmark: 10.1109/SAAHPC.2011.15

Authors

Roberto Ammendola
Andrea Biagioni
Ottorino Frezza
Francesca Lo Cicero
Alessandro Lonardo
Pier Stanislao Paolucci
Davide Rossetti
Francesco Simula
Laura Tosoratto
Piero Vicini

DOWNLOAD PDF | SHARE ARTICLE | GENERATE CITATION

Abstract

QUonG is an INFN (Istituto Nazionale di Fisica Nucleare) initiative targeted to develop a high performance computing system dedicated to Lattice QCD computations. QUonG is a massively parallel computing platform that lever-ages on commodity multi-core processors coupled with last generation GPUs. Its network mesh exploits the characteristics of LQCD algorithm for the design of a point-to-point, high performance, low latency 3-d torus network to interconnect the computing nodes. The network is built upon the APE net+ project: it consists of an FPGA-based PCI Express board exposing six full bidirectional off-board links running at 34 Gbps each, and implementing RDMA protocol and an

in the exascale era

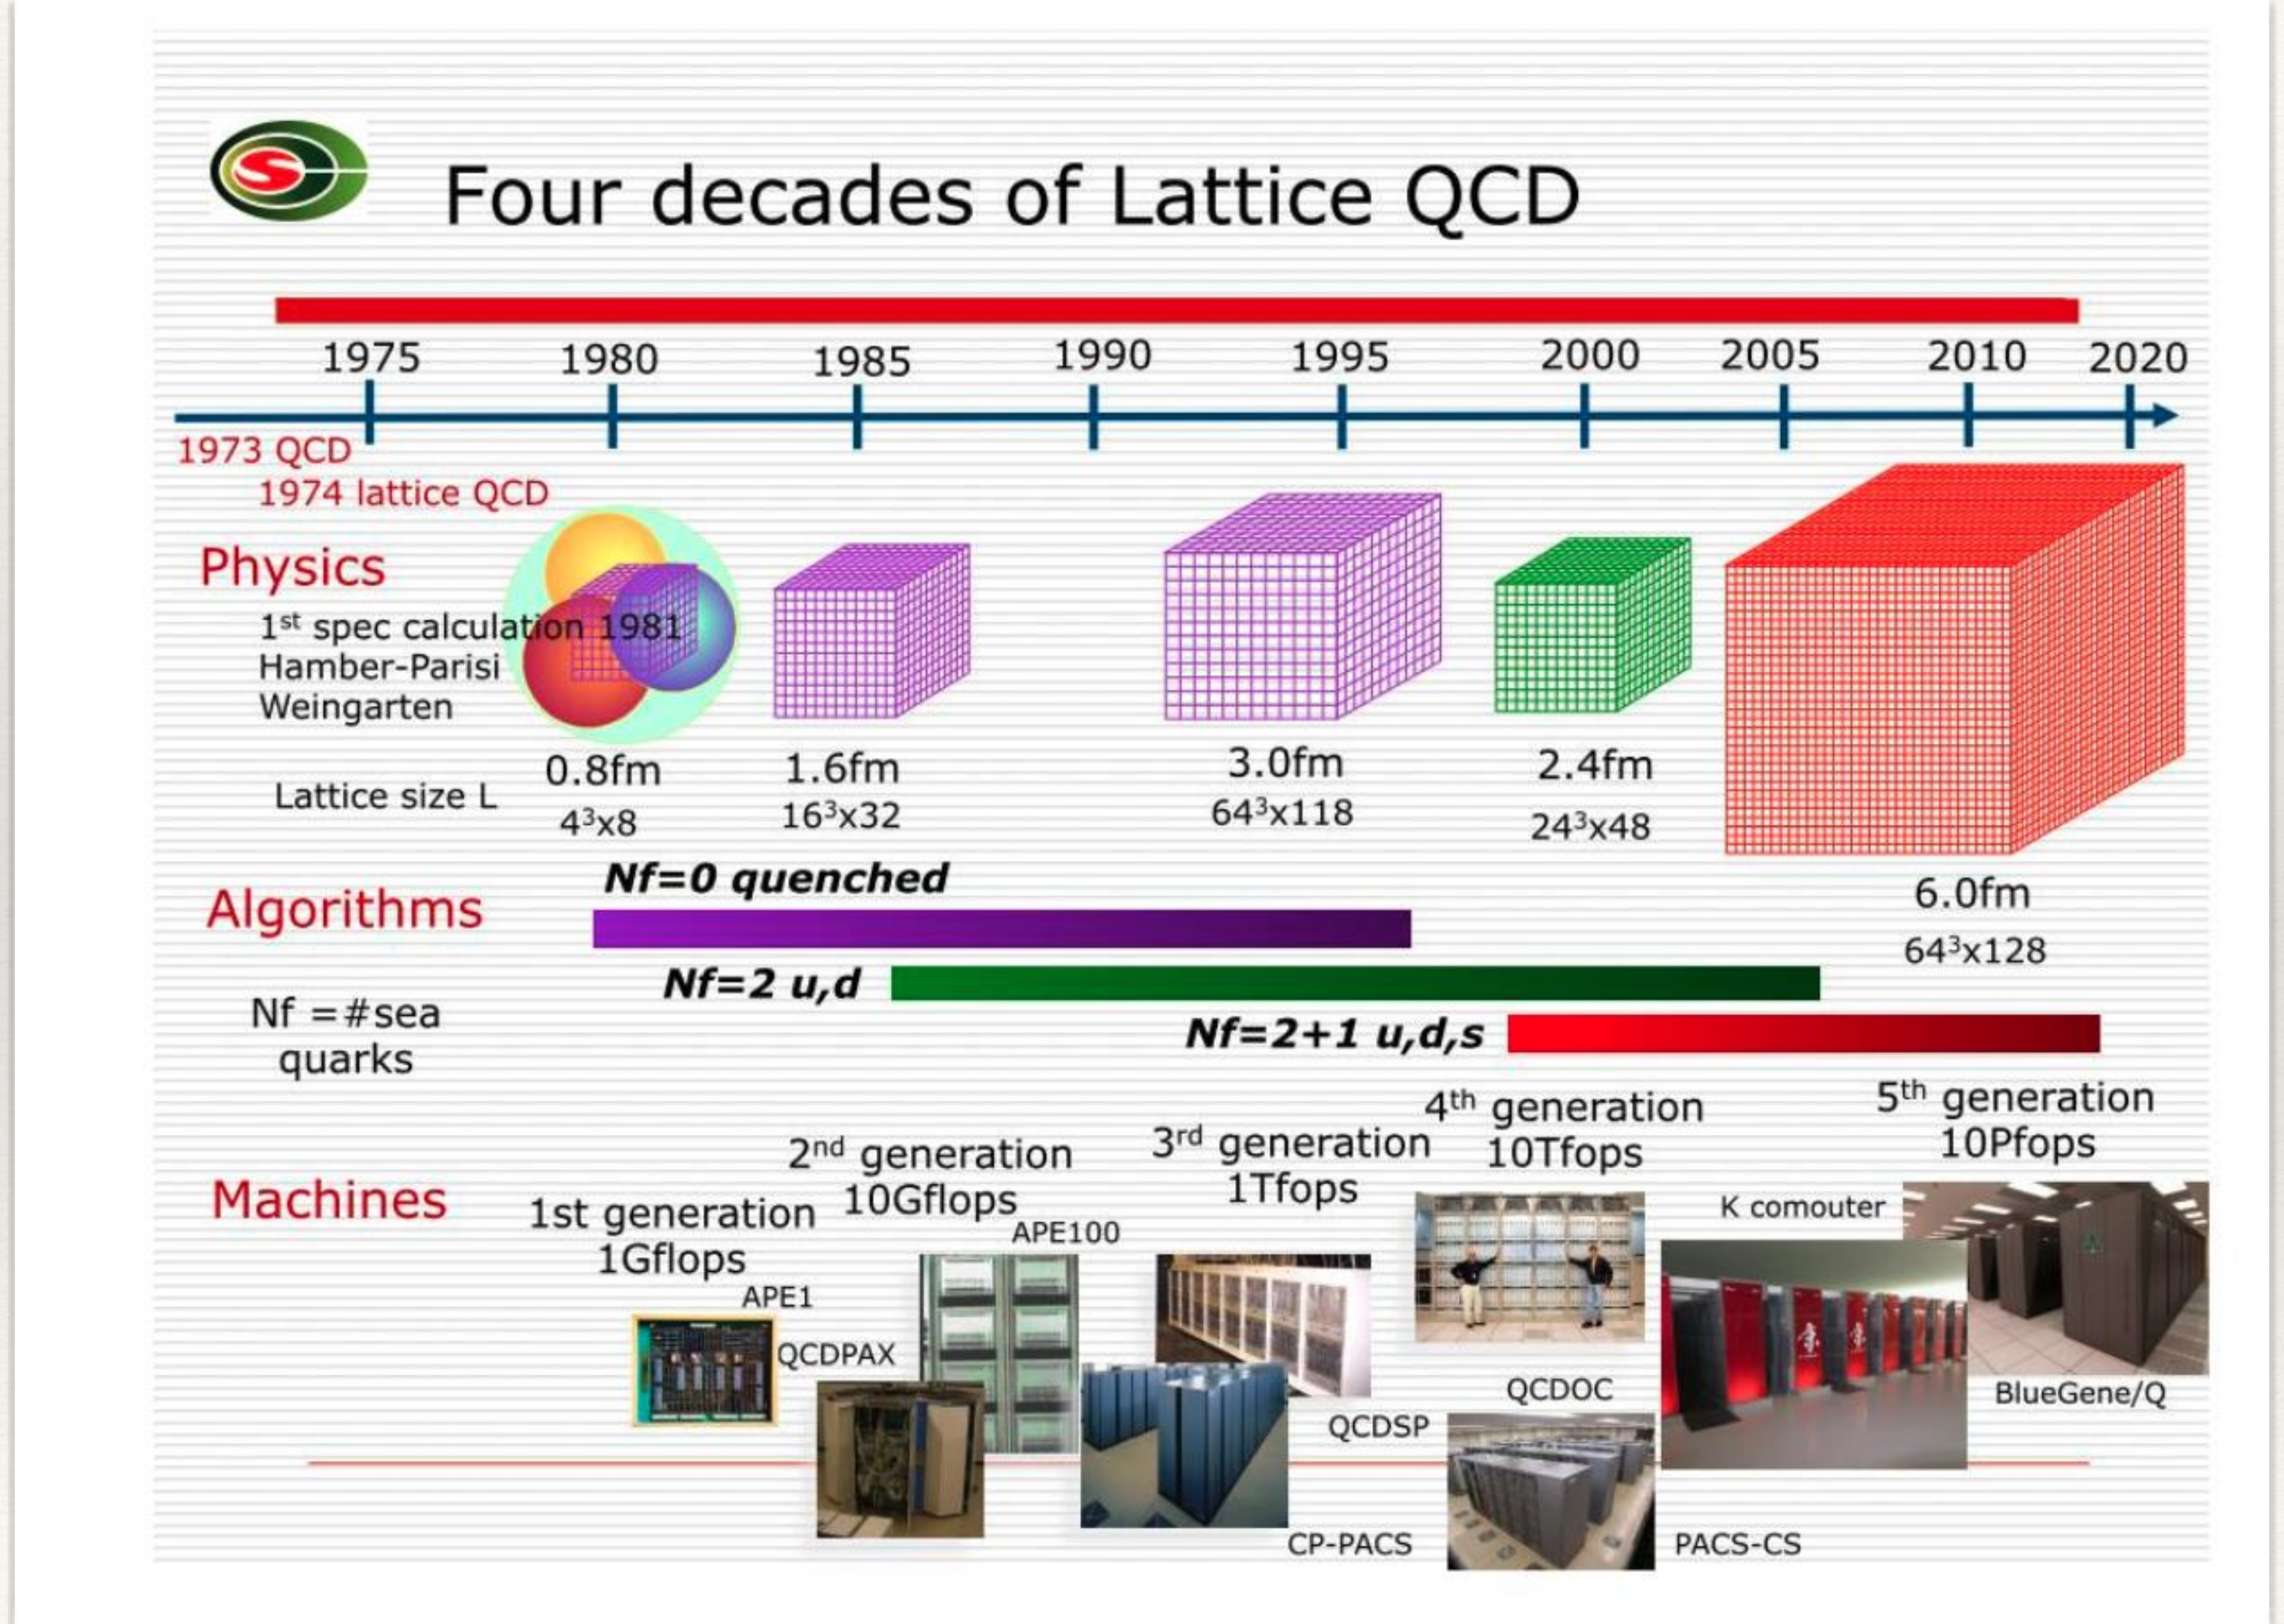

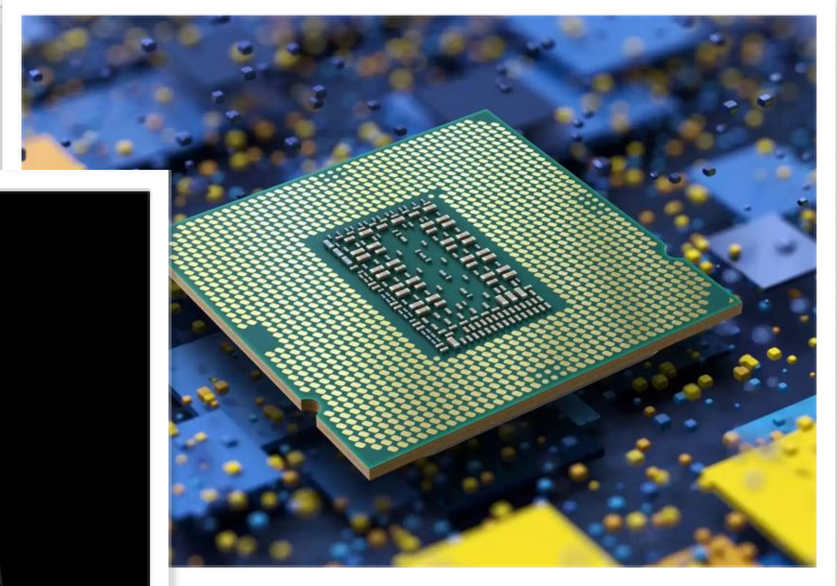

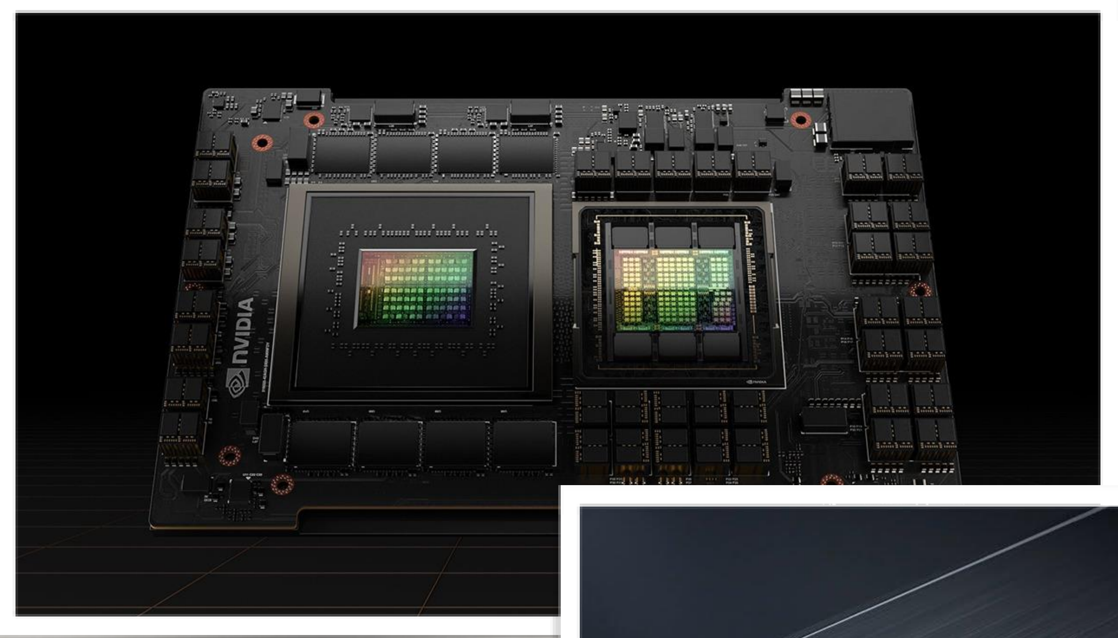

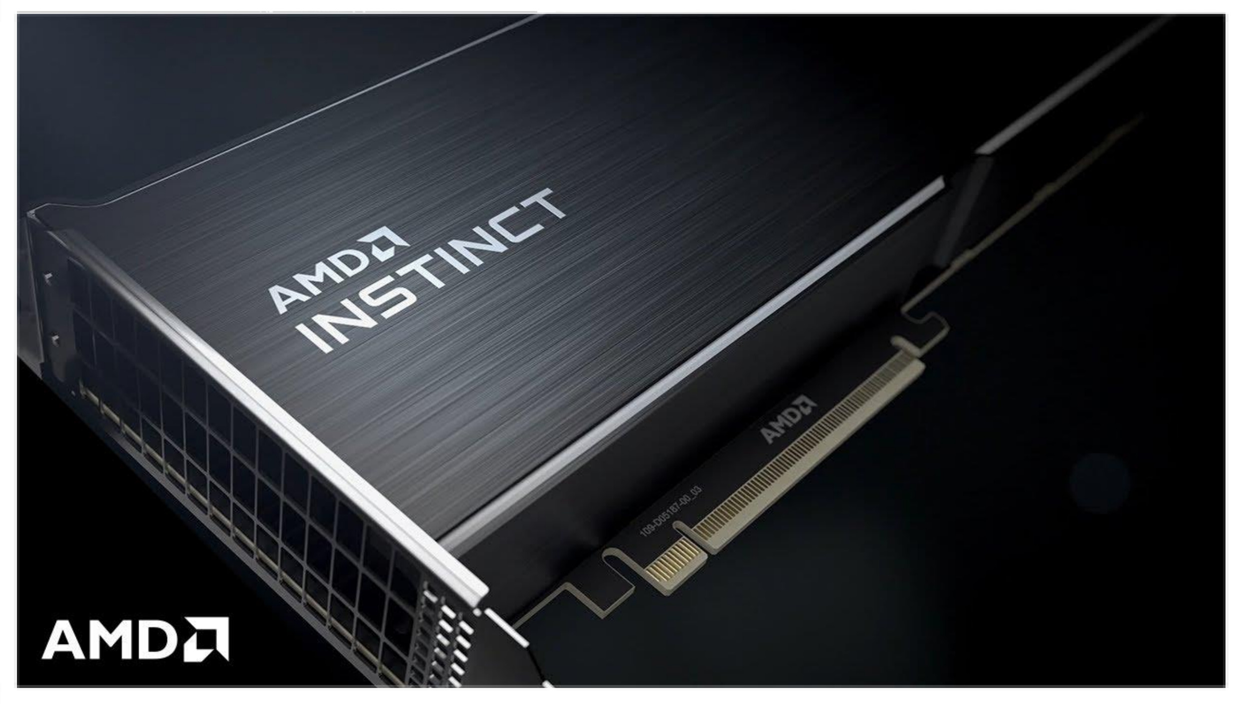

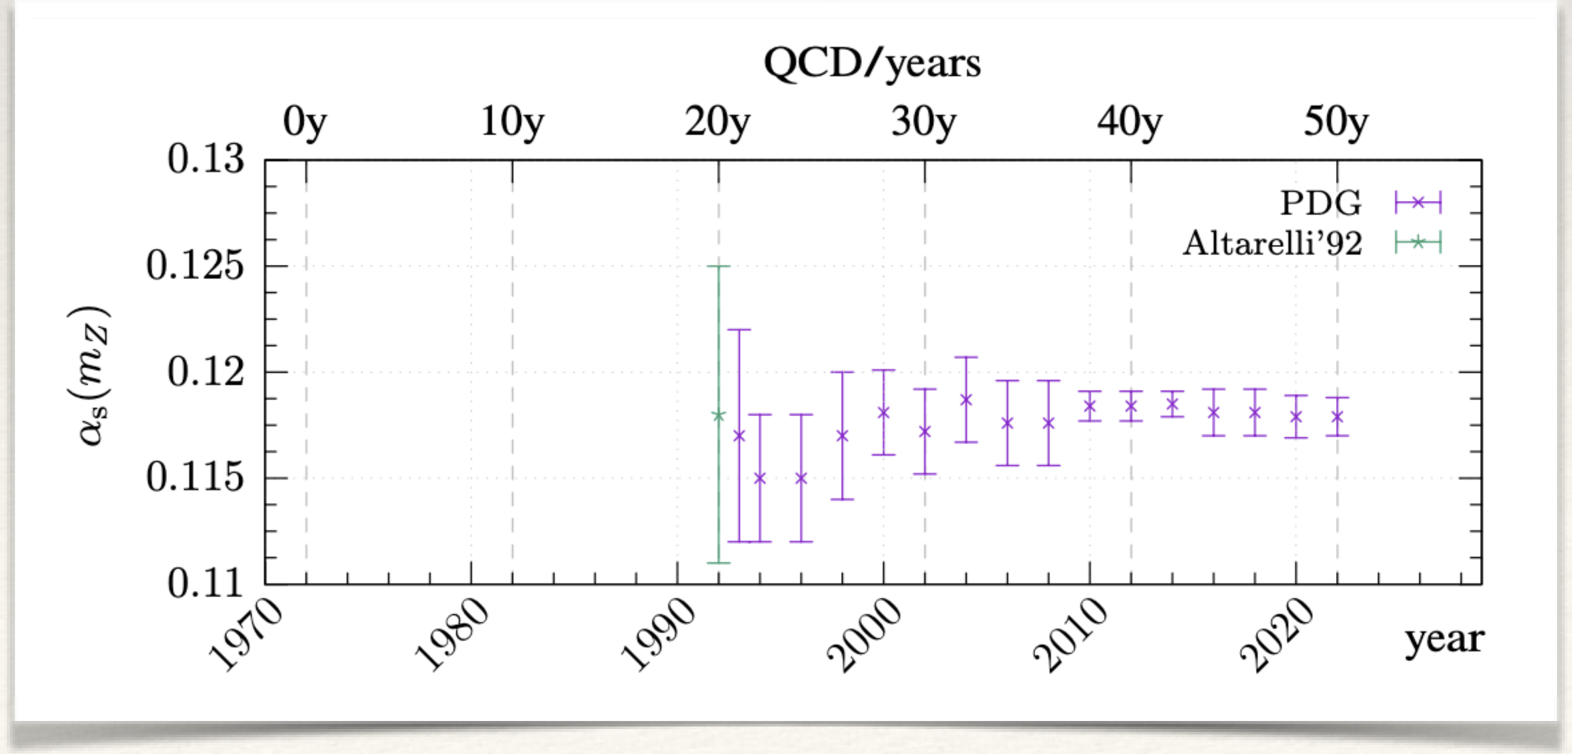

Ukawa Lattice 23

…machines would not suffice

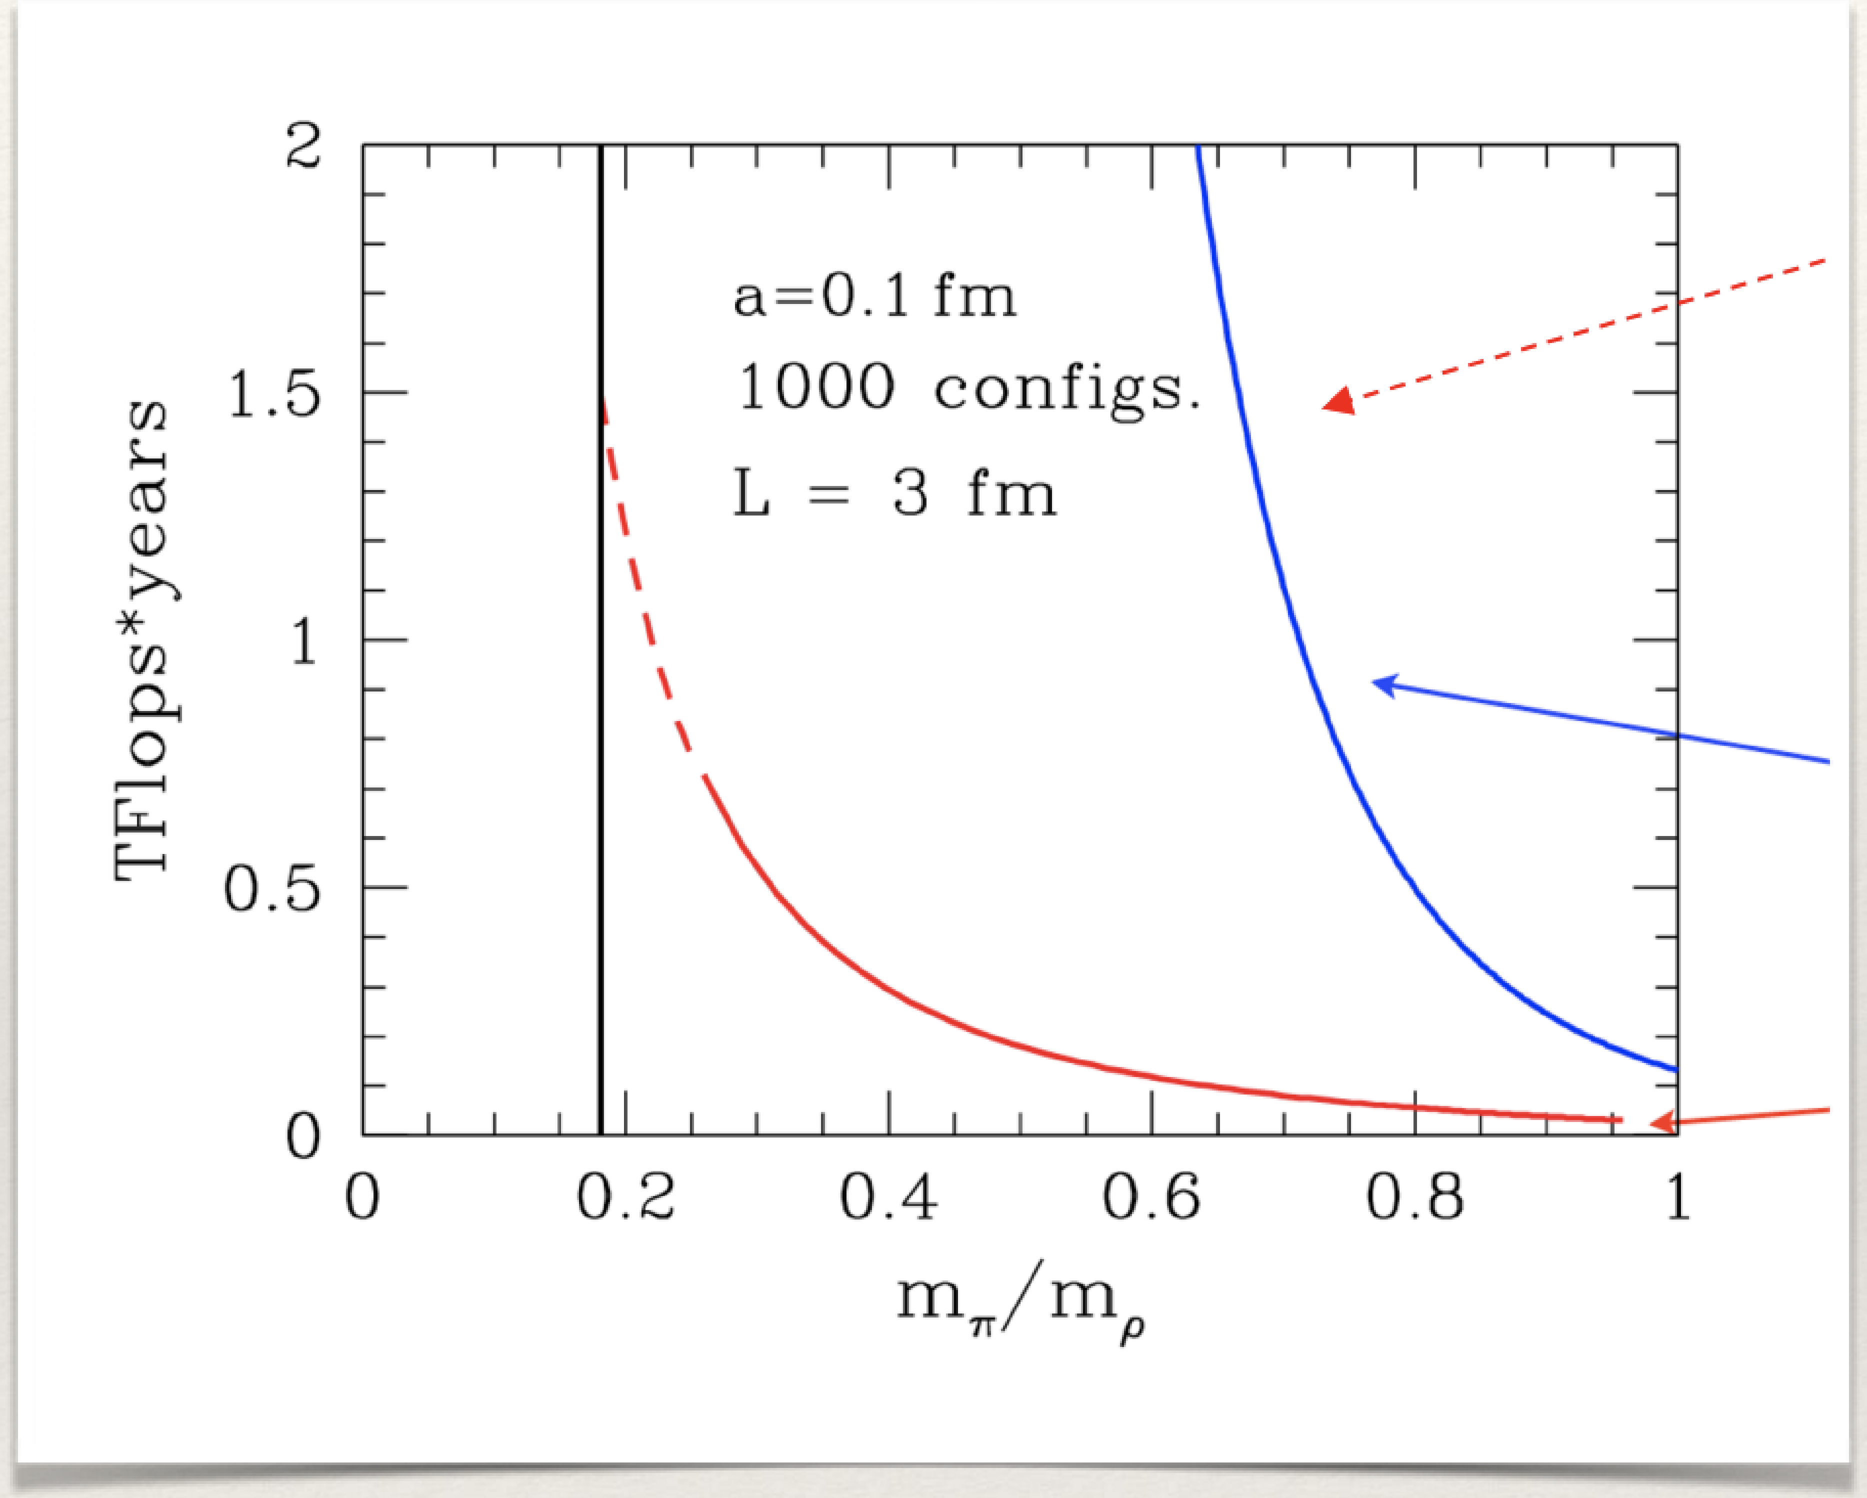

The Berlin Wall (2001)

Ukawa (2001)

Luscher (2004)

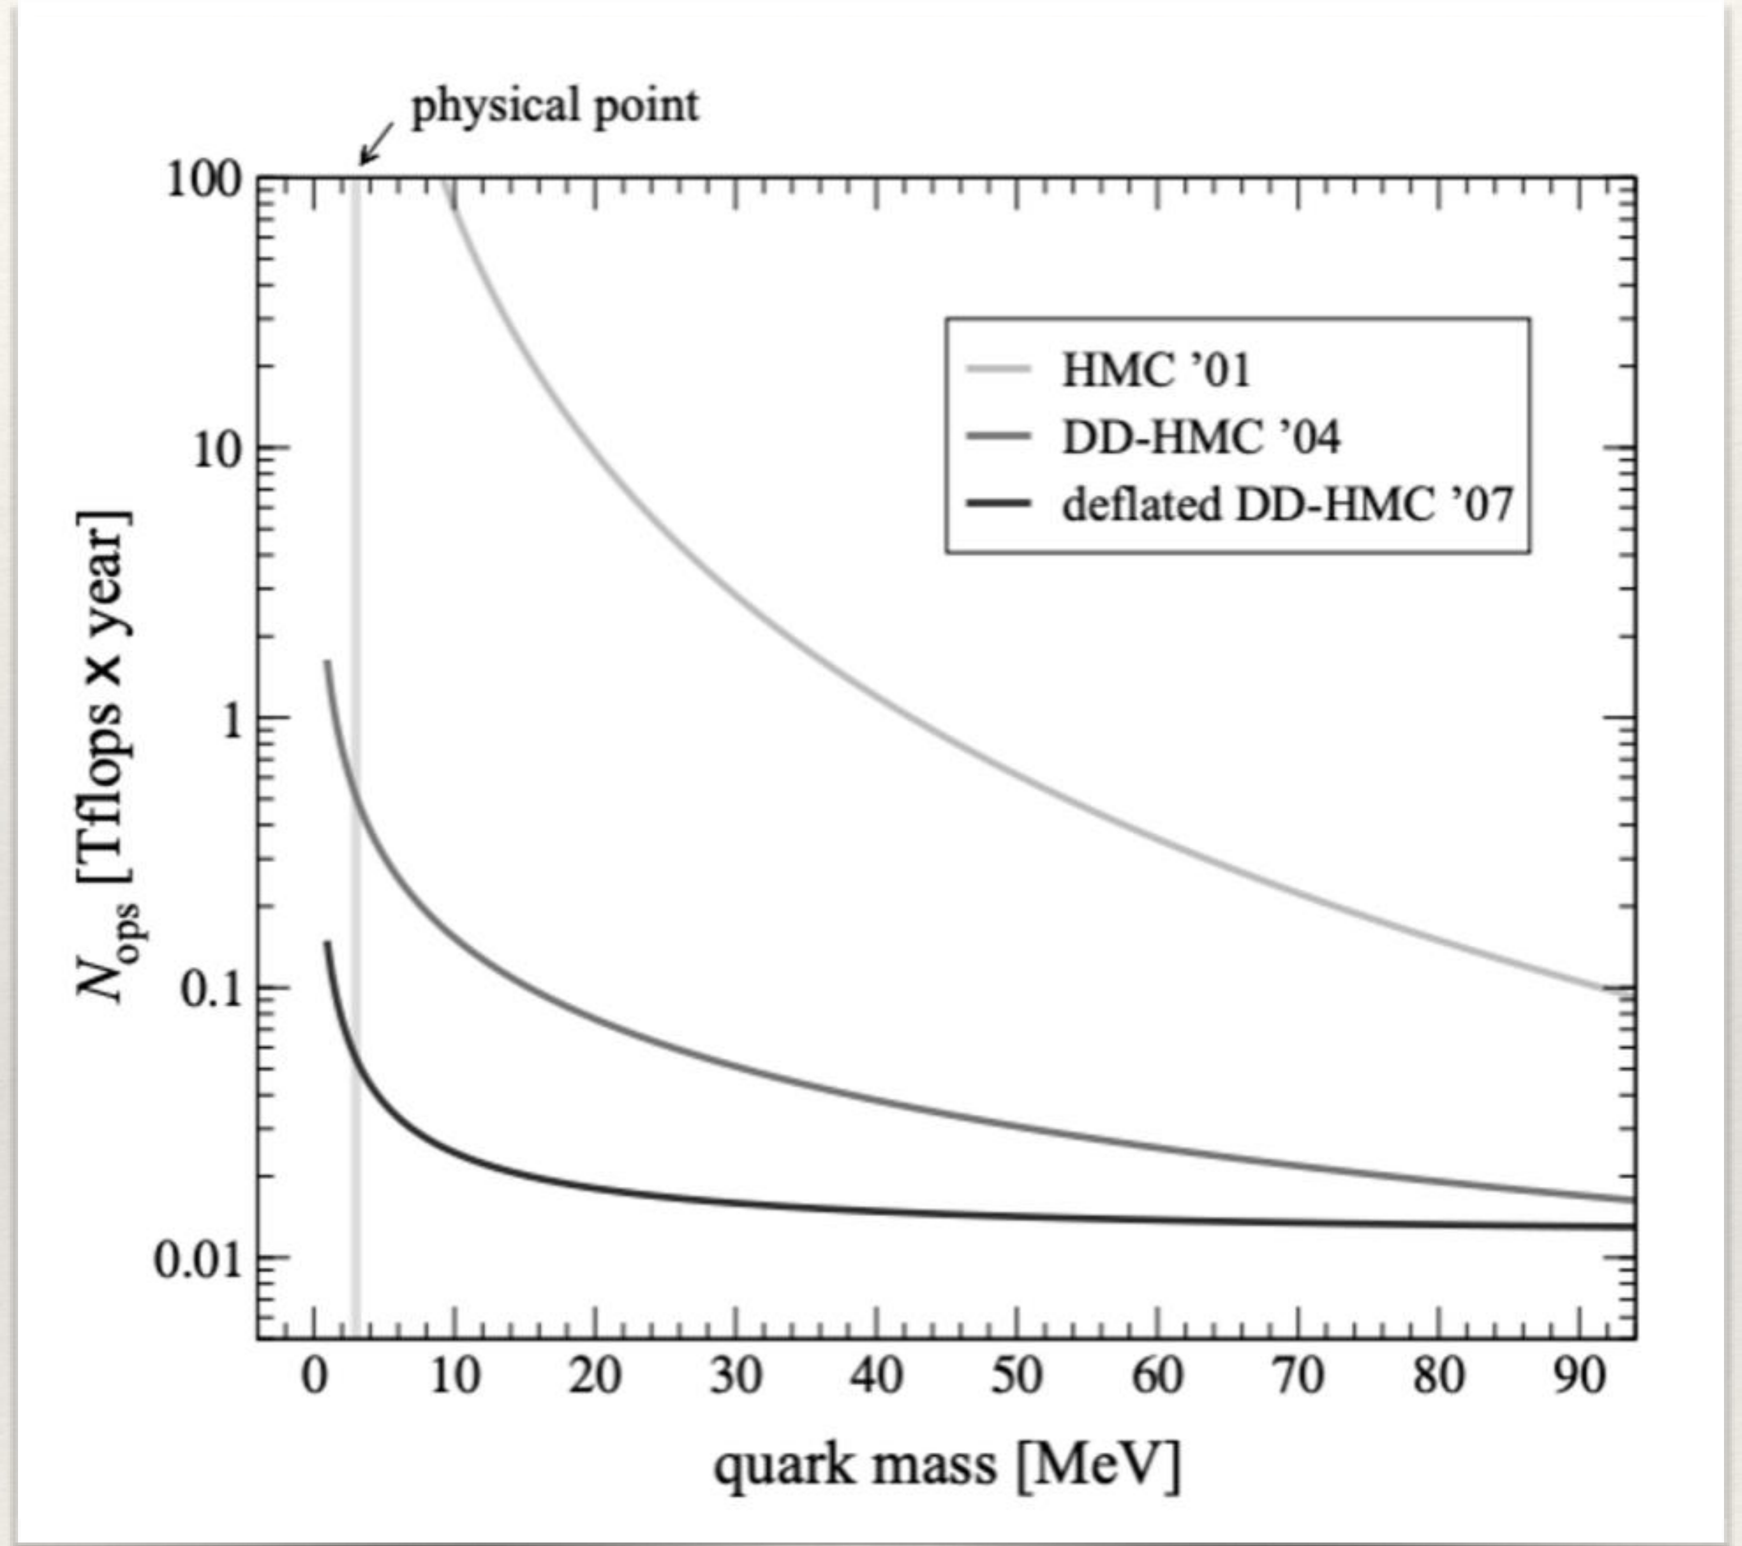

Domain Decomposition in Hybrid MC was implemented in (2004-2006)

Deflated DD implemented on Hybrid MC in 2007

a challenging numerical task

The physical masses of the up and down quarks are much smaller than the typical low-energy hadronic, generic methods will be not efficient.

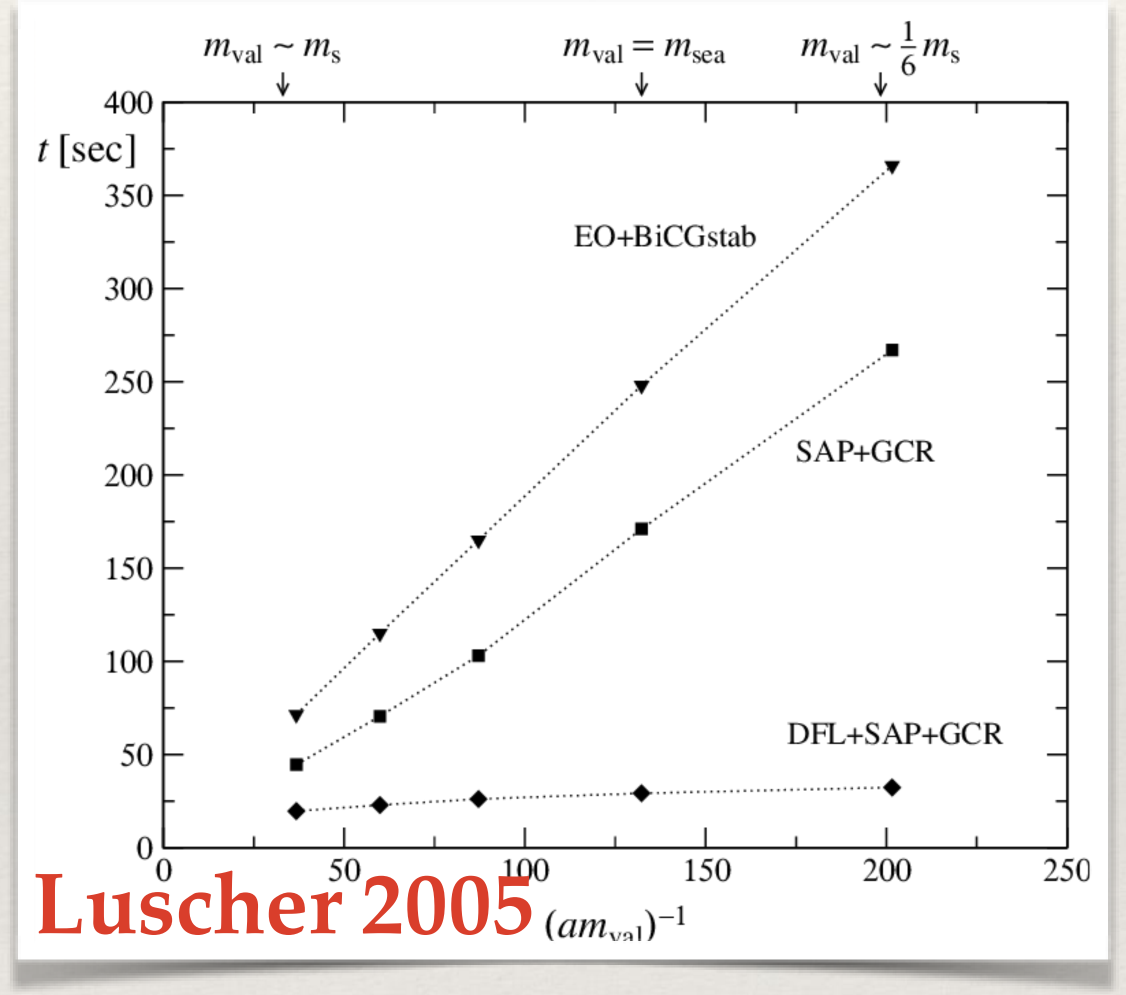

Luscher 2005

A **deflation approach** takes advantage of approximate local coherence of the low modes, that allows highly effective deflation subspaces generation. The numerical effort required for the preparation of the subspace and deflation of the Dirac operator is then only of order of the volume

A **multigrid method** is an algorithm for solving differential equations using a hierarchy of discretizations. It is belongs to a class of techniques called multiresolution methods, efficient in tackling problems exhibiting multiple scales of behavior.

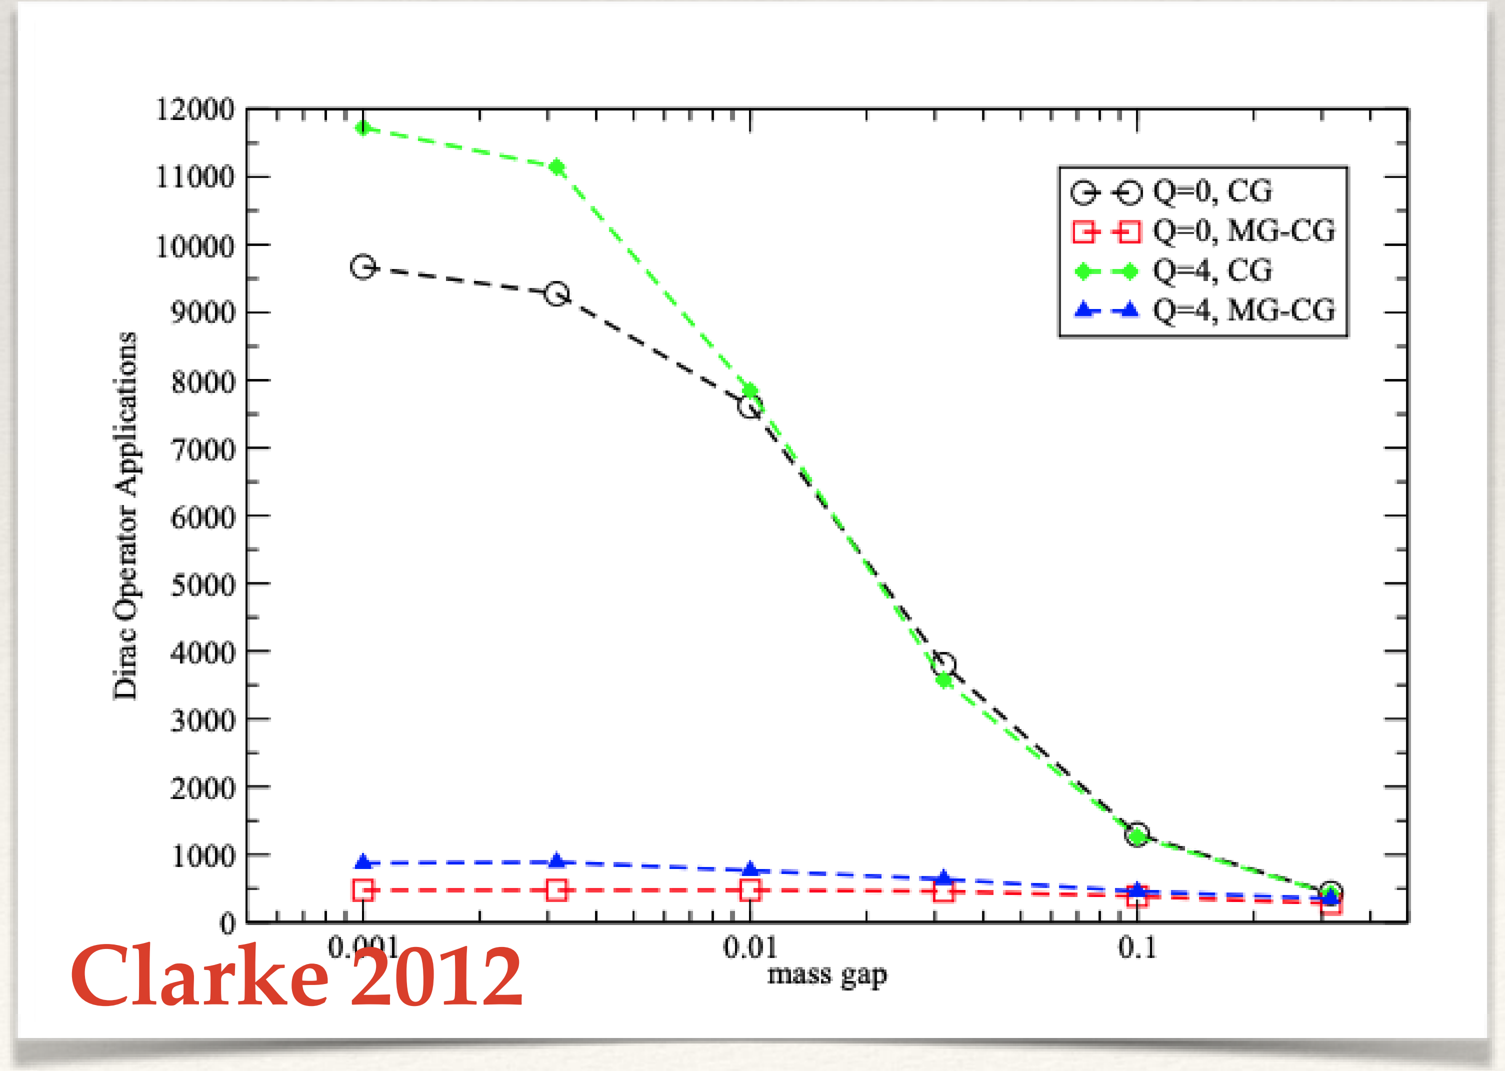

Clarke 2012

FLAG = Flavour Lattice Averaging Group

- The scope of FLAG is to review the current status of lattice results for a variety of physical quantities that are important for flavour physics.
  - Issued every 3 years
  - Composed of experts in Lattice Field Theory, Chiral Perturbation Theory and Standard Model phenomenology.
  - Aim: providing an answer to the frequently posed question “What is currently the best Lattice QCD value for a particular quantity?”
    - In a way that is readily accessible to those who are not expert in lattice methods. Generally not an easy question to answer.
    - Last revision 2021, update few months ago, Feb. 2023 (previous in 2019, 2016, 2013, 2010)
    - **https://arxiv.org/pdf/2111.09849.pdf**

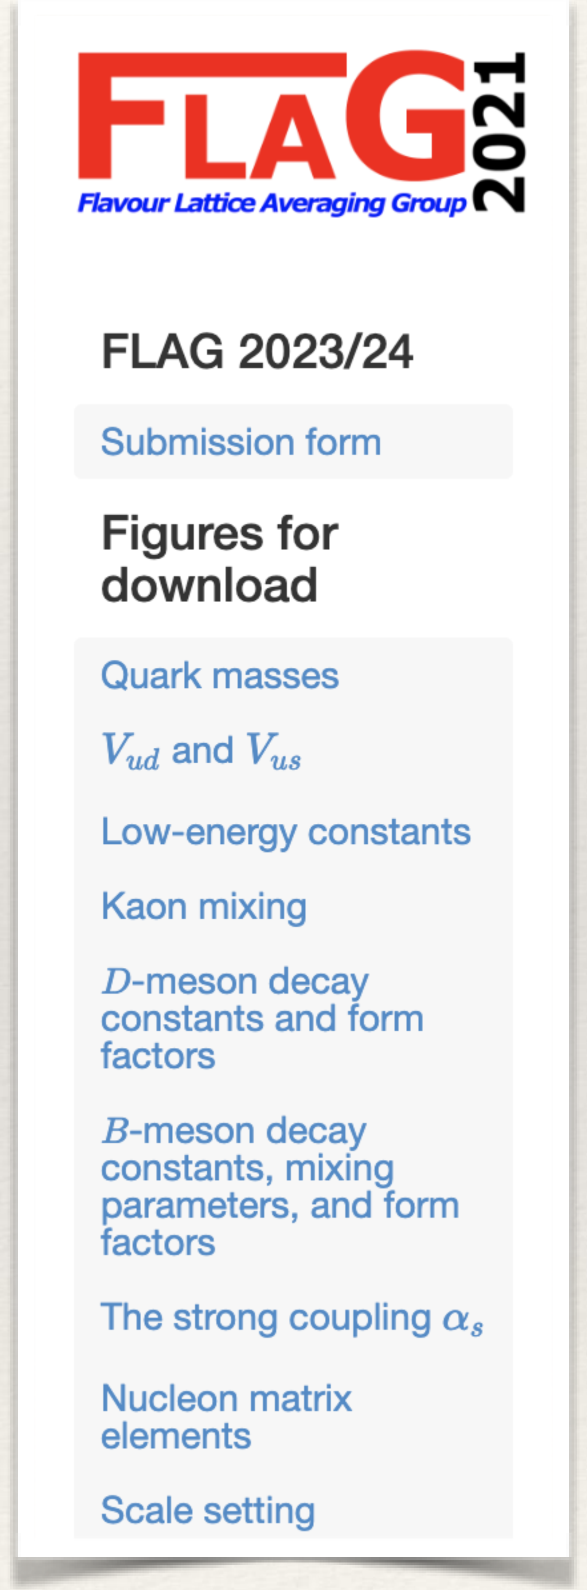

the final message

QUANTUM THEORY CENTER
SDU

- Lattice QCD is a vibrant, modern, and active field of research
- Challenging problems ahead, which might require changing the working assumptions in HEP completely
- Working in Lattice QCD requires a combined understanding of Physics, parallel Supercomputers, modern mathematics,…
- Ideal for education of students