Silvano Simula

INFN - Roma Tre

MHPC Workshop on
High Performance Computing
SISSA, Trieste
February 24th -26th, 2016

Studies of Flavor Physics using Lattice QCD simulations with modern HPC hardware

* brief summary of the theoretical INFN activities requiring HPC

* the case of Flavor Physics and Lattice QCD

* computing machines presently used by INFN and modern HPC with accelerators (GPUs and MICs)

Scientific Structure of INFN: National Committees

I: Particle Physics

II: Astroparticle Physics

III: Nuclear Physics

IV: **Theoretical Physics**

V: Technological Research

research line 1: Field and String Theories **(~ 170 FTE)**

research line 2: Phenomenology of Elementary Particles **(~ 100 FTE)**

research line 3: Hadronic and Nuclear Physics **(~ 75 FTE)**

research line 4: Mathematical Methods **(~ 75 FTE)**

research line 5: Astroparticle Physics **(~ 105 FTE)**

research line 6: Statistical Physics and Applied Field Theory **(~ 75 FTE)**

* the research lines are organized into **specific projects** with a national PI (~ 40 projects for a total of ~ 600 FTE)

* each project is evaluated by external referees (also international ones) each 3 years

*** ~ 11 specific projects (~ 80 researchers) need HPC**

LATTICE QCD AND FLAVOR PHYSICS

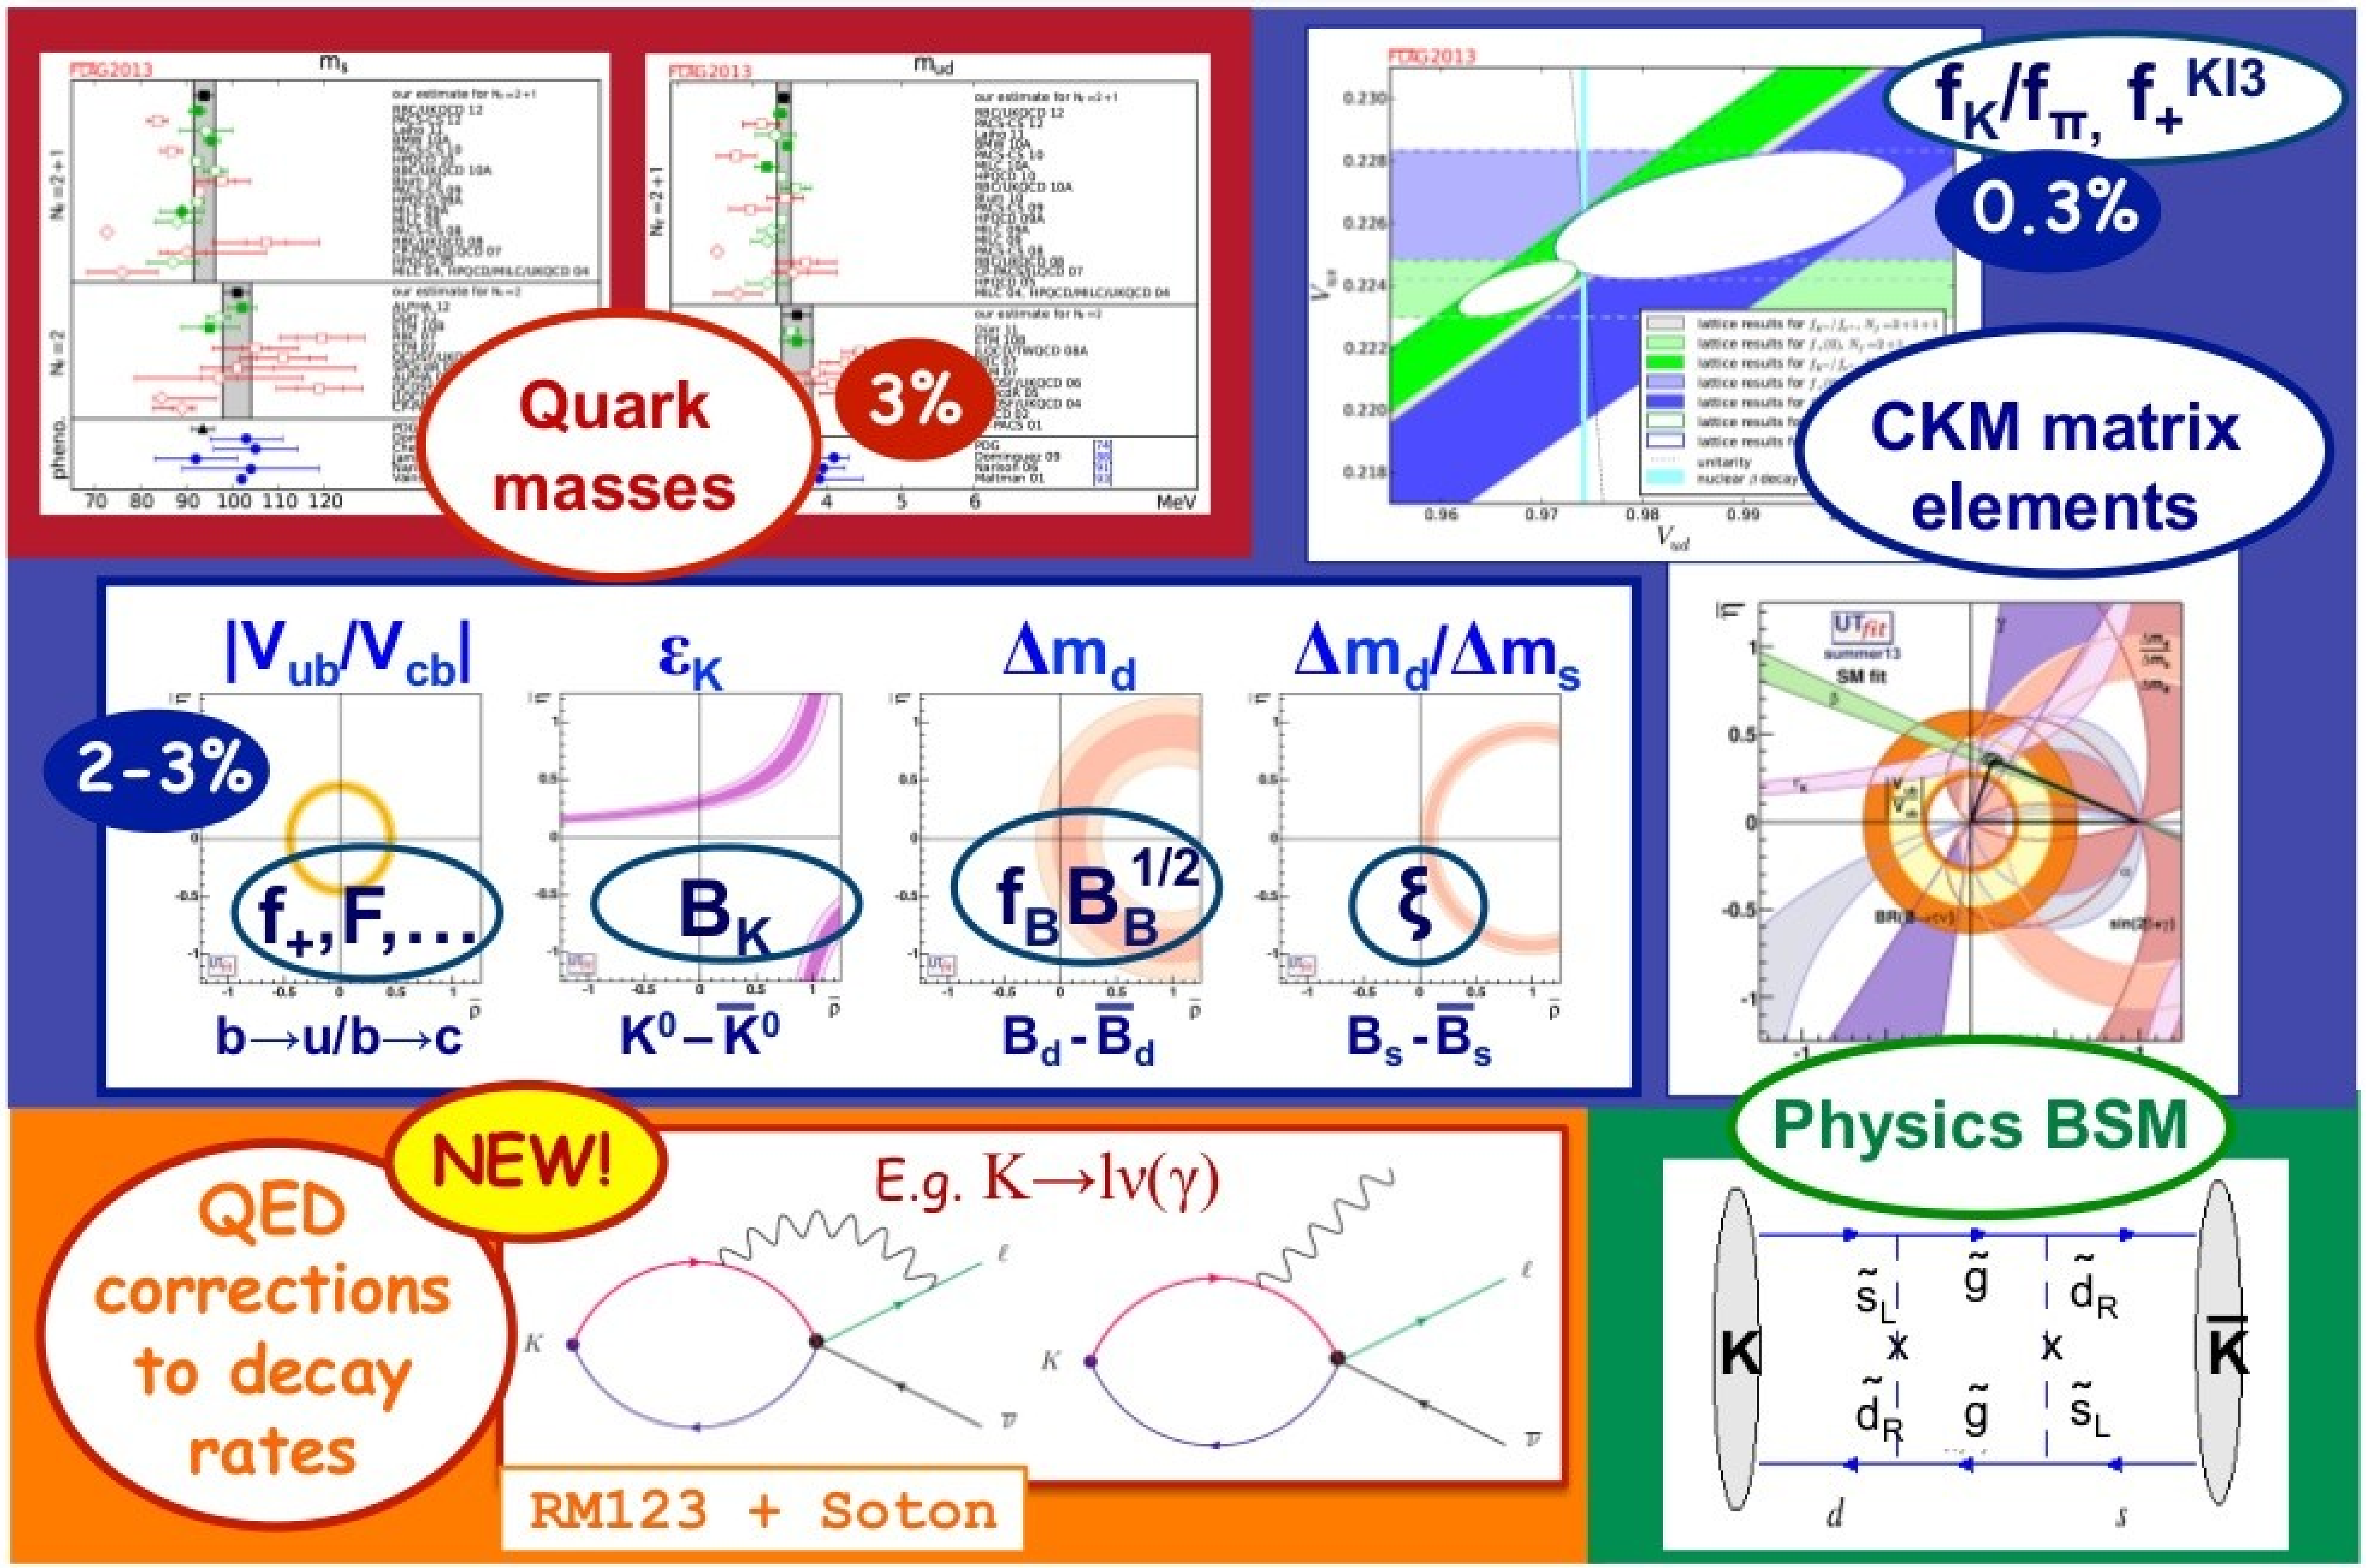

more info at https://web.infn.it/CSN4/IS/Linea2/LQCD/index.html

courtesy of V. Lubicz

Quark Gluon Plasma and Strong Interactions

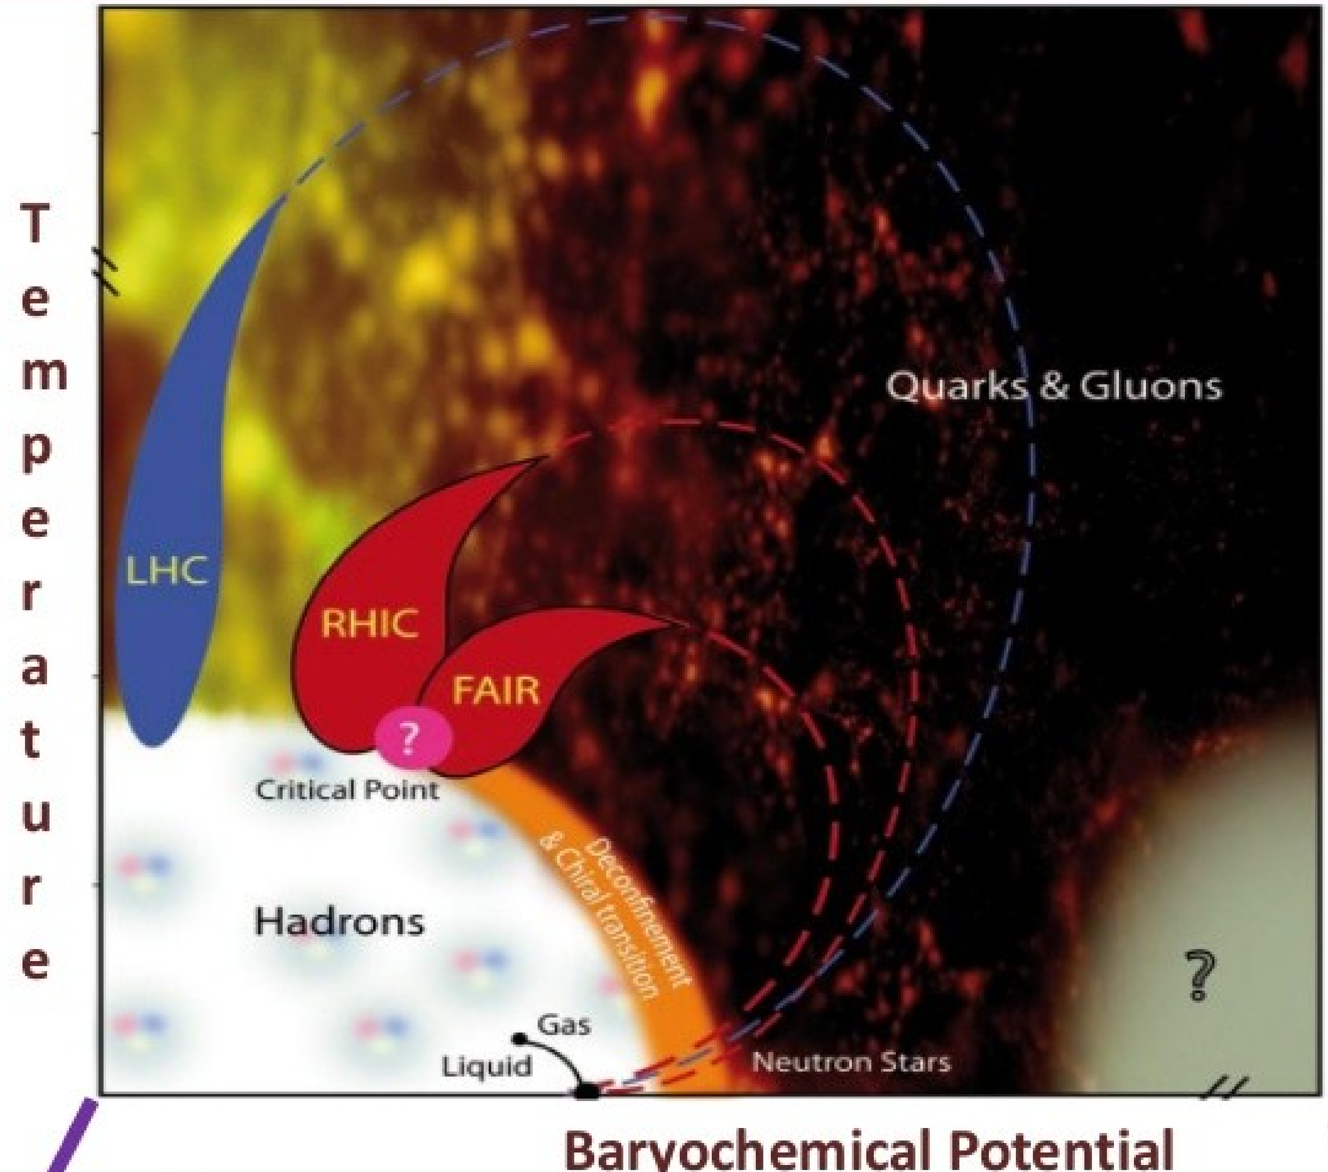

- Pseudocritical line
- Equation of state
- Bottomonium in the plasma
- Magnetic and Chromomagnetic Fields
- New Methods for EoS
- New methods for finite density
- Confinement and Topology
- AdS/CFT approach to QGP

Your favorite parameter here

Effective String models , 3 and 4d
Yang Mills Theory
Chiral transition with two flavors
Chiral Transition with many flavors

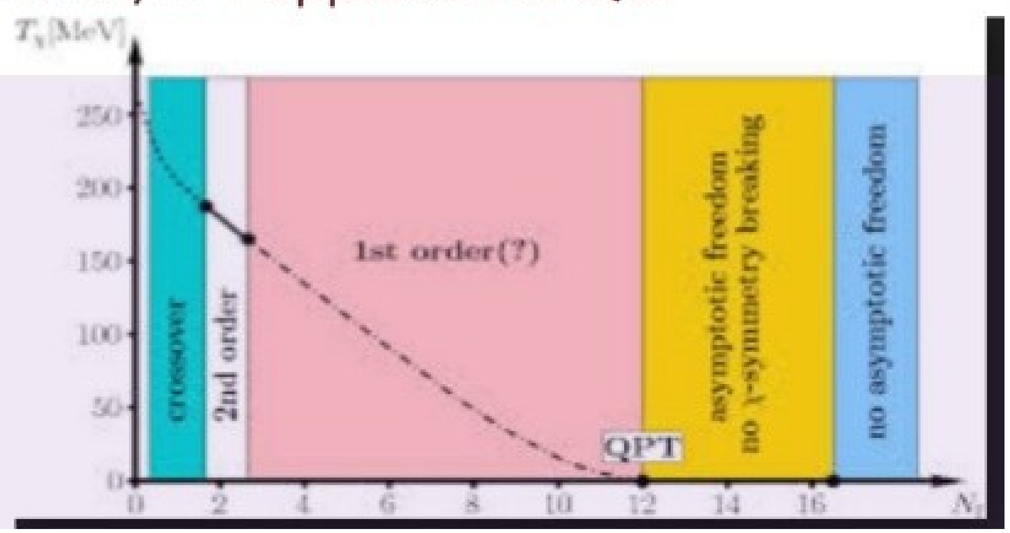

more info at https://web.infn.it/CSN4/IS/Linea1/FT_CP/index.html

Nuclear Physics

•Light Nuclei: test of the nuclear interaction, study of reactions of astrophysical interest, energy production, Efimov physics & universal properties, radioactive beams, hypernuclei,...

Methods: direct solution of the Schroedinger equation, Green Function Monte Carlo techniques

•Medium & Heavy Nuclei: dark-matter & antinucleon cross section, double beta decay matrix elements, hypernuclei, study of nuclei with an excess of protons or neutrons,...

Methods: shell model & coupled cluster methods, Auxiliary Field Diffusion Monte Carlo techniques

•Infinite nuclear matter: equation of state, compact stars, strange stars,...

Methods: Configuration Interaction Monte Carlo techniques, cluster expansion

Computation requirements very different:

•Shell model & other techniques: diagonalization of big matrices (very efficient infrastructure of communication and memory management)

•Monte Carlo techniques: very high CPU consumption

more info at https://web2.infn.it/CSN4/IS/Linea3/FBS/index.html

General Relativity

Studies of compact objects such as Black Holes and Neutron Stars, focussing on their role as possible sources of Gravitational Waves of interest for interferometric detectors (Virgo and LIGO)

Selected for a Viewpoint in *Physics*

PRL **116,** 061102 (2016) PHYSICAL REVIEW LETTERS week ending 12 FEBRUARY 2016

Observation of Gravitational Waves from a Binary Black Hole Merger

B. P. Abbott *et al.**

(LIGO Scientific Collaboration and Virgo Collaboration)

(Received 21 January 2016; published 11 February 2016)

On September 14, 2015 at 09:50:45 UTC the two detectors of the Laser Interferometer Gravitational-Wave Observatory simultaneously observed a transient gravitational-wave signal. The signal sweeps upwards in frequency from 35 to 250 Hz with a peak gravitational-wave strain of 1.0×10^{-21}. It matches the waveform predicted by general relativity for the inspiral and merger of a pair of black holes and the ringdown of the resulting single black hole. The signal was observed with a matched-filter signal-to-noise ratio of 24 and a false alarm rate estimated to be less than 1 event per 203 000 years, equivalent to a significance greater than 5.1σ. The source lies at a luminosity distance of 410^{+160}_{-180} Mpc corresponding to a redshift $z = 0.09^{+0.03}_{-0.04}$. In the source frame, the initial black hole masses are $36^{+5}_{-4} M_\odot$ and $29^{+4}_{-4} M_\odot$, and the final black hole mass is $62^{+4}_{-4} M_\odot$, with $3.0^{+0.5}_{-0.5} M_\odot c^2$ radiated in gravitational waves. All uncertainties define 90% credible intervals. These observations demonstrate the existence of binary stellar-mass black hole systems. This is the first direct detection of gravitational waves and the first observation of a binary black hole merger.

DOI: 10.1103/PhysRevLett.116.061102

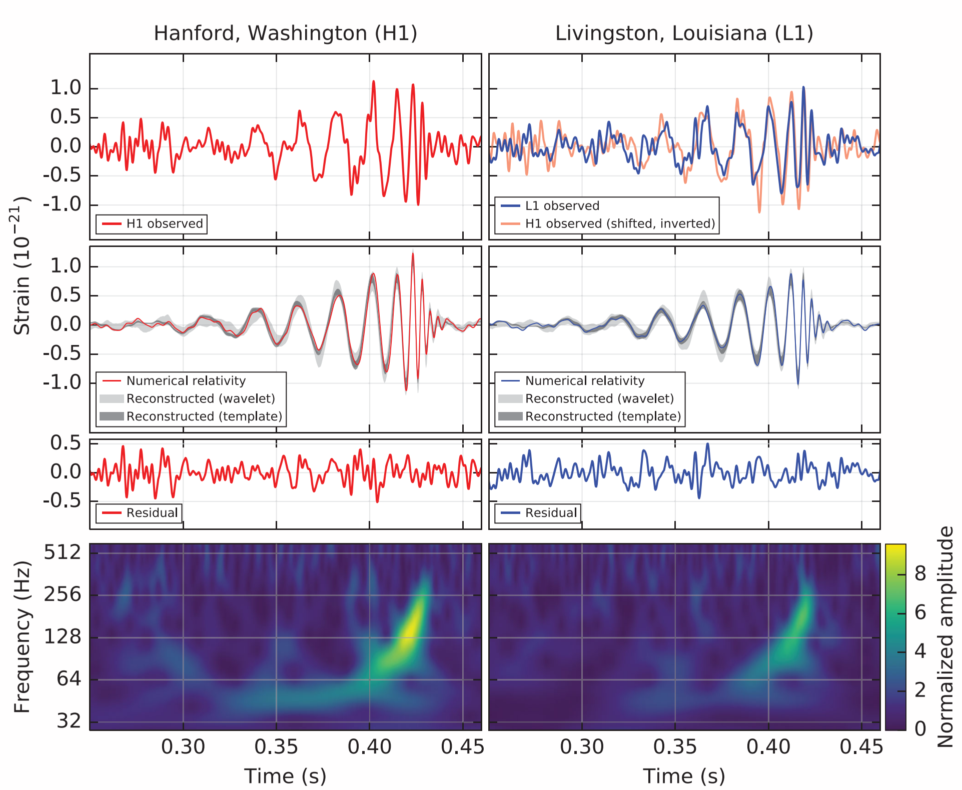

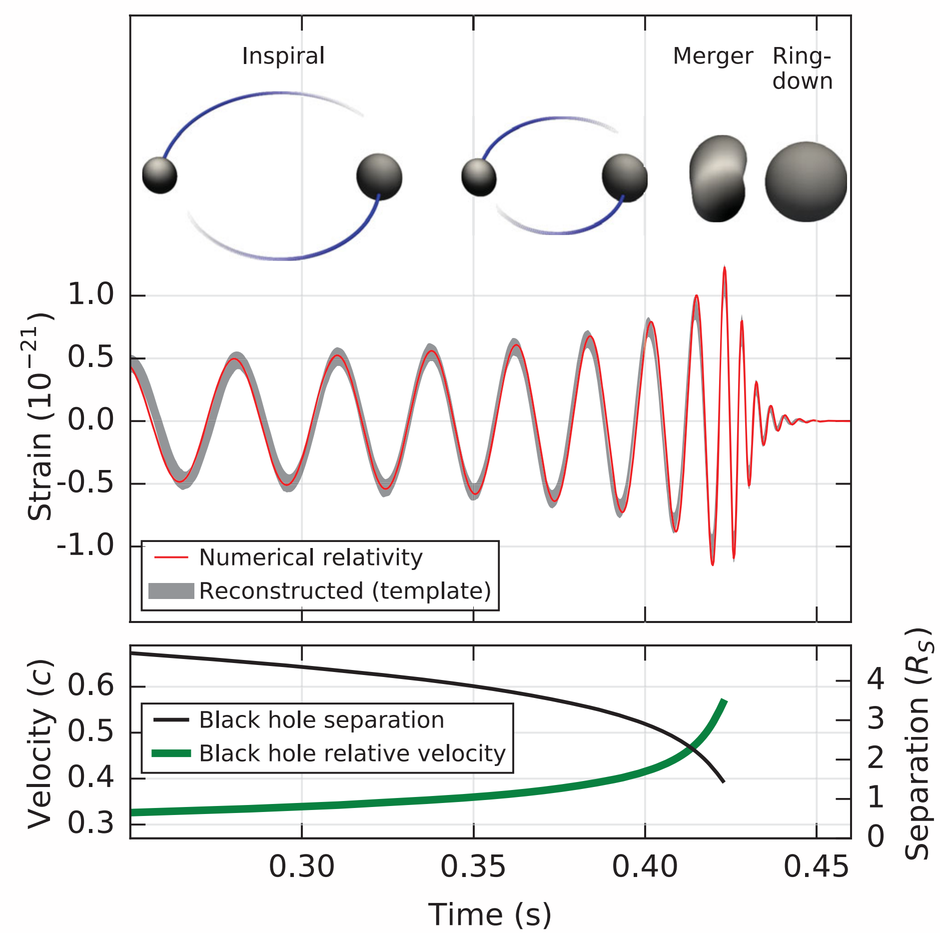

FIG. 2. *Top:* Estimated gravitational-wave strain amplitude from GW150914 projected onto H1. This shows the full bandwidth of the waveforms, without the filtering used for Fig. 1. The inset images show numerical relativity models of the black hole horizons as the black holes coalesce. *Bottom:* The Keplerian effective black hole separation in units of Schwarzschild radii ($R_S = 2GM/c^2$) and the effective relative velocity given by the post-Newtonian parameter $v/c = (GM\pi f/c^3)^{1/3}$, where f is the gravitational-wave frequency calculated with numerical relativity and M is the total mass (value from Table I).

more info at https://web2.infn.it/CSN4/IS/Linea5/TEONGRAV/index.html

Fluid Dynamics: the physics of turbolence

Quantitative Biology

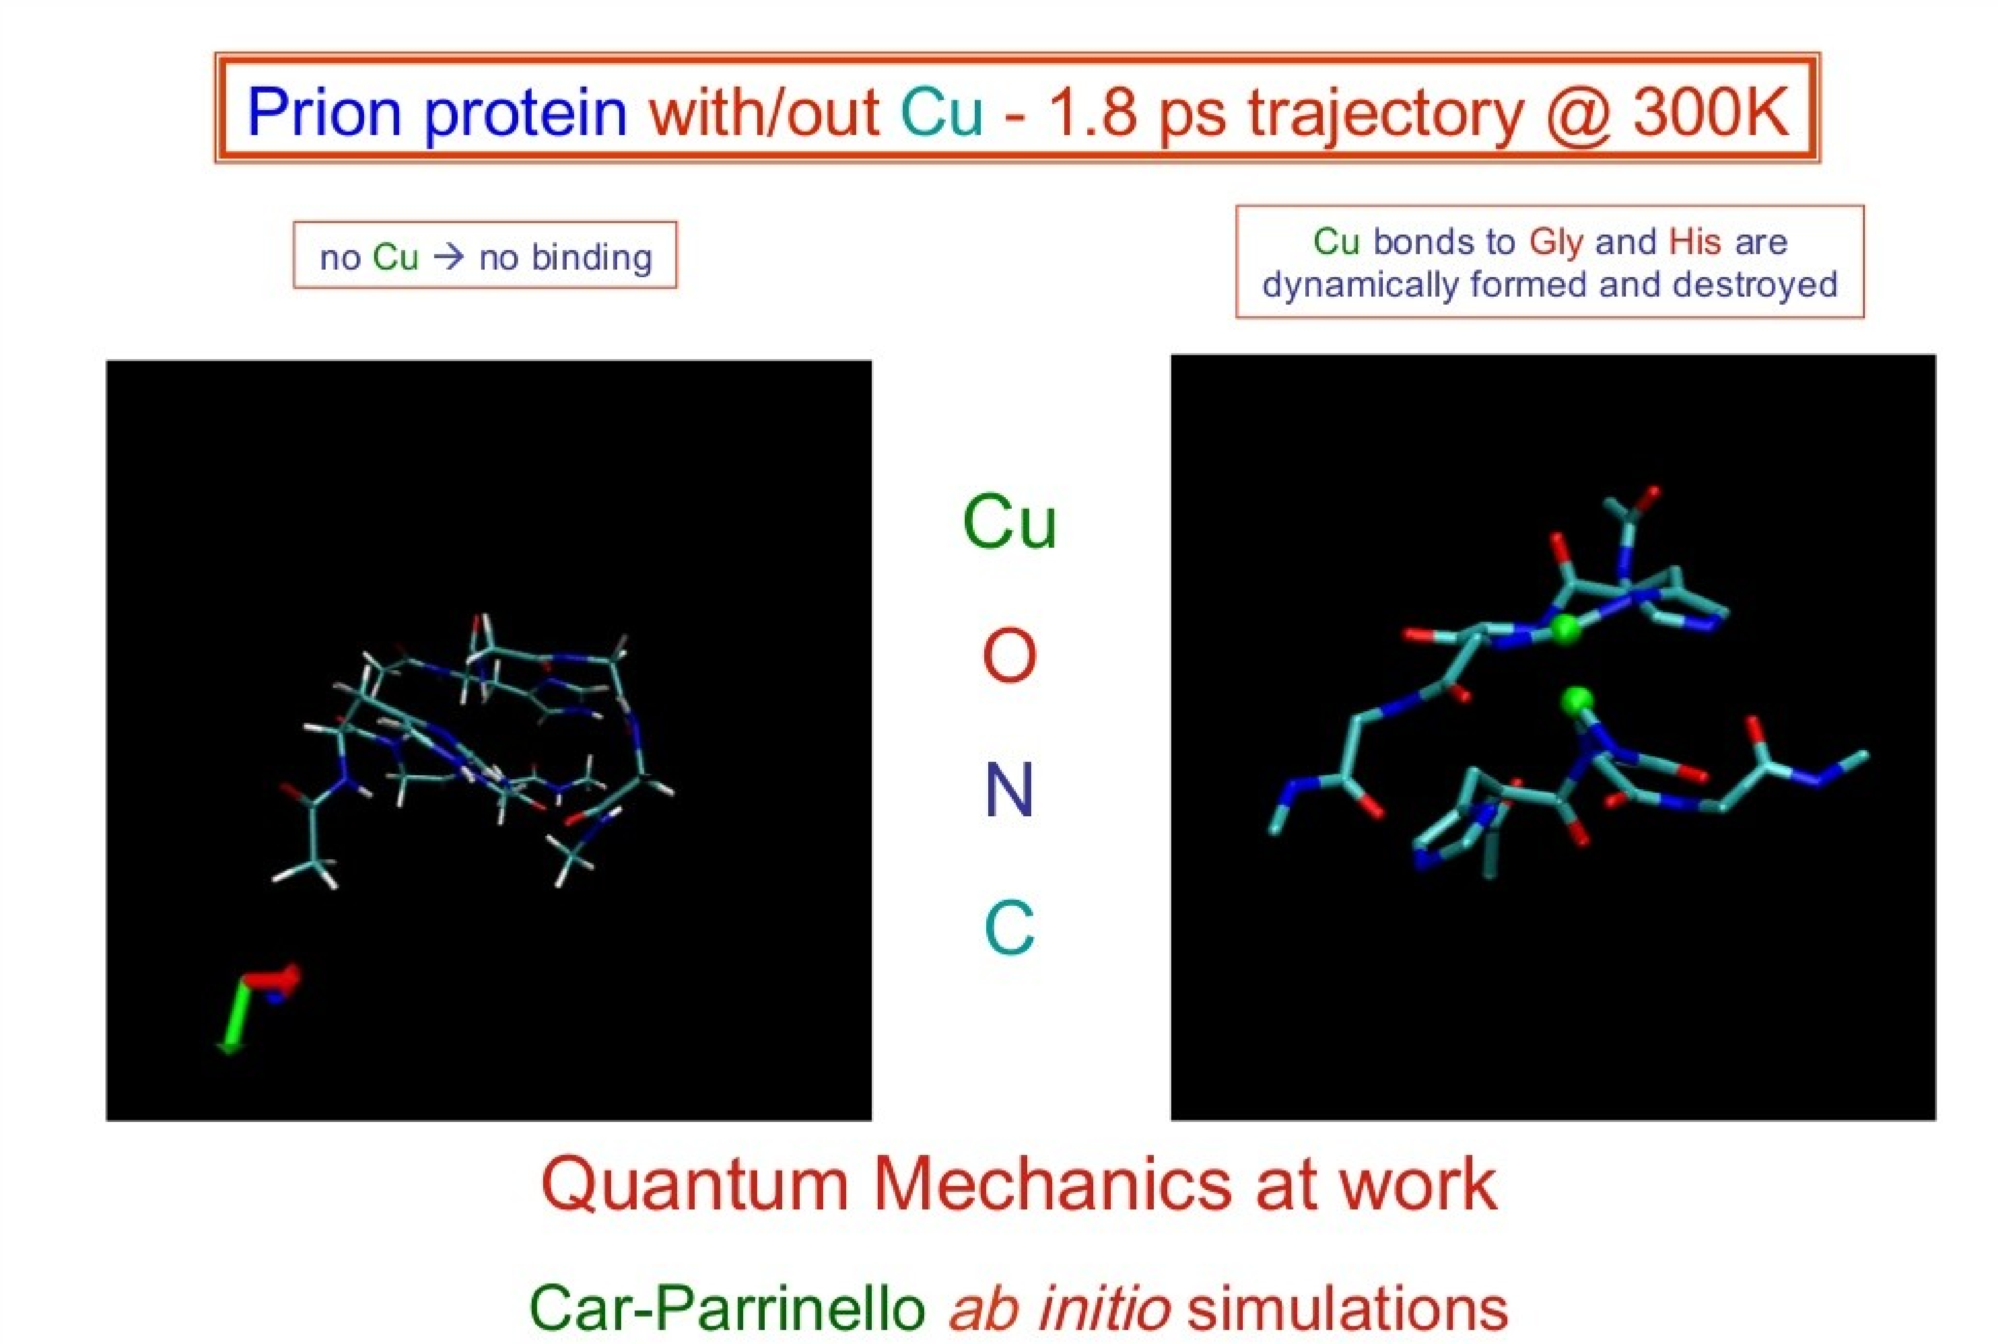

more info at

https://web.infn.it/CSN4/IS/Linea6/FIELDTURB/index.html and

https://web.infn.it/CSN4/IS/Linea6/BIOPHYS/index.html

The case of Flavor Physics

New Physics energy scale Λ_{NP} still unknown

$\Lambda_{NP} \sim 1$ TeV

direct search of NP at LHC possible
indirect search of NP complementary

$\Lambda_{NP} >> 1$ TeV

only indirect search of NP
(through virtual loop effects)

precision studies of New Physics are an important part of existing and planned experiments

LHCb at CERN, Belle2 at KEK, NA62 at CERN, KOTO at J-PARC, MEG at PSI, ...

the goal is to achieve a theoretical precision comparable to the experimental one

* the interpretation of experimental data typically requires the precise knowledge of some hadronic parameters governed by the strong interactions

* **lattice QCD** is a tool encoding all the non-perturbative effects of the strong interactions starting from first principles only

QCD simulations on the lattice

Lattice is a way to regularized a field theory in Euclidean space

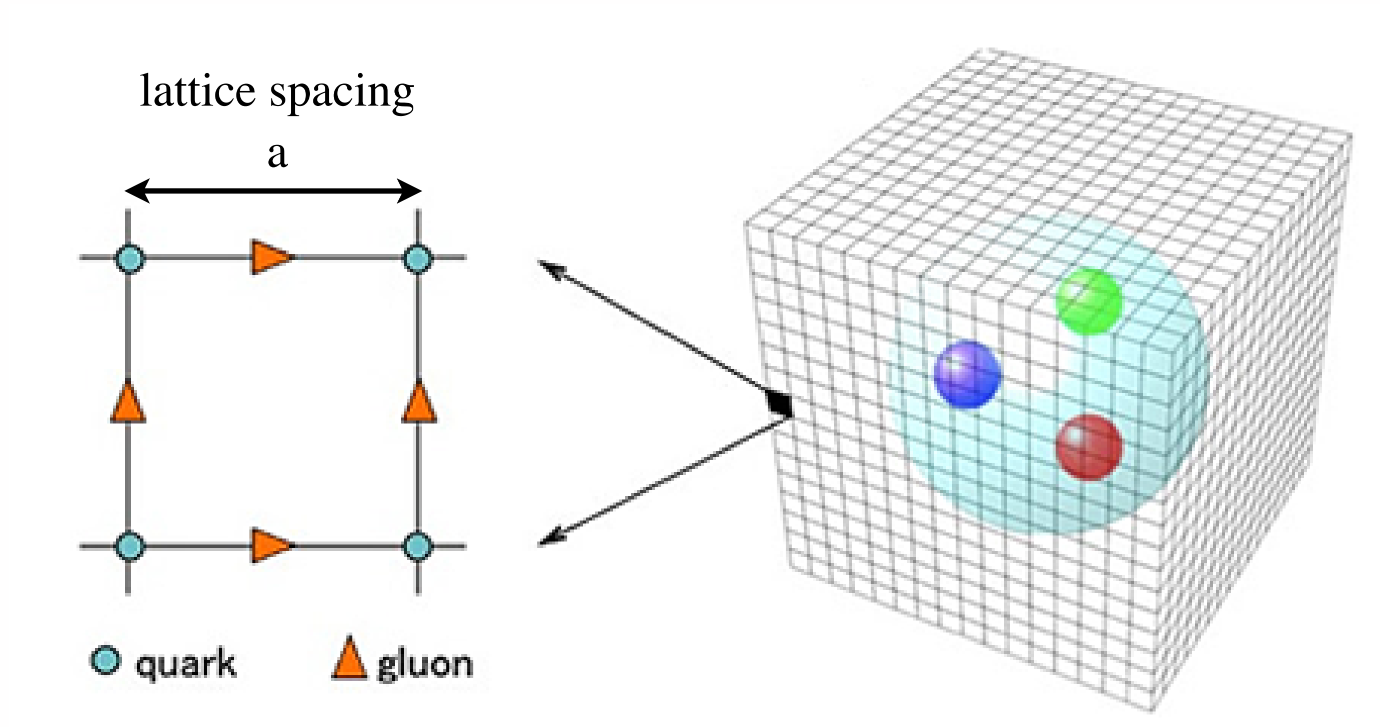

action: $S_{lattice} = S_{QCD} + O(a)$

observable P: $\langle P \rangle_{lattice} = \langle P \rangle_{continuum} + O(a^2)$

Monte Carlo numerical simulations

$$\langle P \rangle_{lattice} = \frac{\sum P\, e^{-S_{lattice}}}{\sum e^{-S_{lattice}}} \xrightarrow{\text{importance sampling}} \sum_{\{U_i\}} P[U_i]$$

* statistical errors → larger samples
* discretization errors → several lattice spacings → continuum limit
*finite size effects → several lattice volumes → infinite volume limit
* renormalization → non-perturbative schemes (RI-MOM, ...)

* chiral extrapolations (simulations at massive pions) → gauge ensembles produced at the physical pion point
* dynamical sea quarks → $N_f = 2 + 1 + 1$: two mass-degenerate u/d quarks, strange and charm at their physical masses

Flavor Lattice Averaging Group

Eur. Phys. J. C (2014) 74:2890
DOI 10.1140/epjc/s10052-014-2890-7

THE EUROPEAN PHYSICAL JOURNAL C

Review

Review of lattice results concerning low-energy particle physics

FLAG Working Group

S. Aoki[1], Y. Aoki[2,3], C. Bernard[4], T. Blum[3,5], G. Colangelo[6,a], M. Della Morte[7,8], S. Dürr[9,10], A. X. El-Khadra[11], H. Fukaya[12], R. Horsley[13], A. Jüttner[14], T. Kaneko[15], J. Laiho[16,28], L. Lellouch[17,18], H. Leutwyler[6], V. Lubicz[19,20], E. Lunghi[21], S. Necco[6], T. Onogi[12], C. Pena[22], C. T. Sachrajda[14], S. R. Sharpe[23], S. Simula[20], R. Sommer[24], R. S. Van de Water[25], A. Vladikas[26], U. Wenger[6], H. Wittig[27]

The aim of FLAG is to answer the question:

" What is currently the best lattice value for a particular quantity ? "

FLAG-1 review in 2011
FLAG-2 review in 2014

FLAG-3 review in progress

quantity	FLAG-2 average	FLAG-2 error (%)	relevance	"expected" error (%)
$\alpha_{\overline{MS}}^{(5)}(M_Z)$	0.1184 (12)	1.0	QCD parameter	~ 0.5
m_{ud} (MeV)	3.42 (9)	2.6	QCD parameter	OK (=> ~ 1)
m_s (MeV)	93.8 (2.4)	2.6	QCD parameter	OK (=> ~ 1)
f_{K^+}/f_{π^+}	1.195 (5)	0.4	V_{us} from $K_{\ell 2}$	~ 0.2
$f_+^{K\pi}(0)$	0.9661 (32)	0.3	V_{us} from $K_{\ell 3}$	~ 0.2
$\hat{B}_K$	0.766 (10)	1.3	$K-\overline{K}$ oscillations	~ 0.1
f_{D_s} (MeV)	248.6 (2.7)	1.1	V_{cs} (V_{cd})	~ 0.2
f_{B_s} (MeV)	224 (5)	2.2	$B_s \to \mu^+\mu^-$	~ 0.5
$f_{B_s}\sqrt{\hat{B}_{B_s}}$ (MeV)	266 (18)	6.8	$B-\overline{B}$ oscillations	~ 1
ξ	1.268 (63)	5.0	$B-\overline{B}$ oscillations	~ 1

extract Cabibbo's angle from $K_{\ell 3}$ decays

$K_{\ell 3}$ decays: $\mathrm{K} \to \pi \ell \nu_\ell$ mediated by the weak vector current $\mathrm{V}_\mu = \bar{s}\gamma_\mu u$

matrix elements: $\langle \pi(p_\pi) | V_\mu | K(p_K) \rangle = f_+(q^2)(p_K + p_\pi)_\mu + f_-(q^2)(p_K - p_\pi)_\mu$ $\qquad q^2 = (p_K - p_\pi)^2$

from experiments: $\Gamma(K \to \pi \ell \nu_\ell) \propto |V_{us}| f_+(0) = 0.2165(4) \quad [\sim 0.2\%]$ [FlaviaNet '14]

CKM matrix element $\mathrm{V}_{us} = \sin(\theta_C)$

calculations of 3-point correlation functions

$$C_\mu^{K\pi}(t,t';p_K,p_\pi) \equiv \sum_{x,y,z} \left\langle O_\pi(\vec{y},t_y)\hat{V}_\mu(x,t_x)O_K^+(\vec{z},t_z)\right\rangle \times \delta_{t,t_x-t_z}\delta_{t',t_y-t_z} \times e^{-ip_K(x-z)}e^{ip_\pi(x-y)}$$

$$\xrightarrow[t'-t\to\infty]{t\to\infty} \frac{\sqrt{Z_K Z_\pi}}{4E_K(p_K)E_\pi(p_\pi)} \boxed{\langle \pi(p_\pi)|\hat{V}_\mu|K(p_K)\rangle} e^{-E_K(p_K)t}e^{-E_\pi(p_\pi)(t'-t)}$$

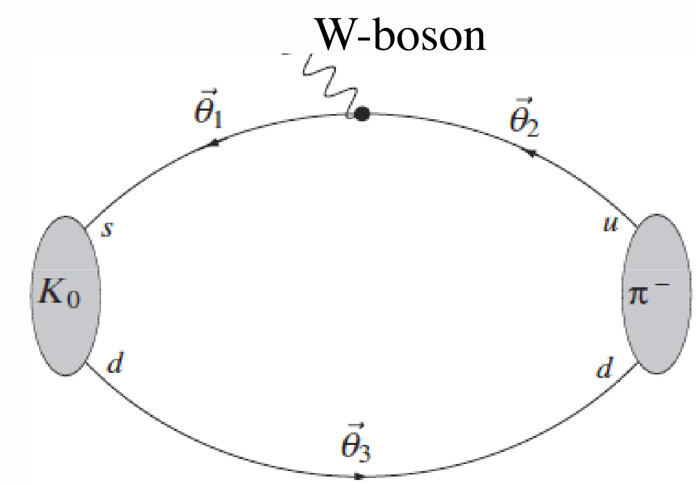

t = time distance between the W-boson insertion and the source (K_0)
t' = time distance between the source (K_0) and the sink (π^-)
O_π and O_K: interpolating fields

$$\vec{p}_K = \frac{2\pi}{L}(\vec{\theta}_3 - \vec{\theta}_1) \quad \vec{p}_\pi = \frac{2\pi}{L}(\vec{\theta}_3 - \vec{\theta}_2)$$

$$\vec{q} = \vec{p}_K - \vec{p}_\pi = \frac{2\pi}{L}(\vec{\theta}_2 - \vec{\theta}_1)$$

* **main ingredient**: the propagator for a moving quark, $S_\theta(x, y)$, which requires the inversion of a sparse, large matrix

$$\sum_z D(x,z)S(z,y) = \delta(x,y)$$

a lattice volume of $48^3 * 96$ corresponds to ~ 10 million lattice sites

first neighbors

modern hardware can accelerate the inversions !

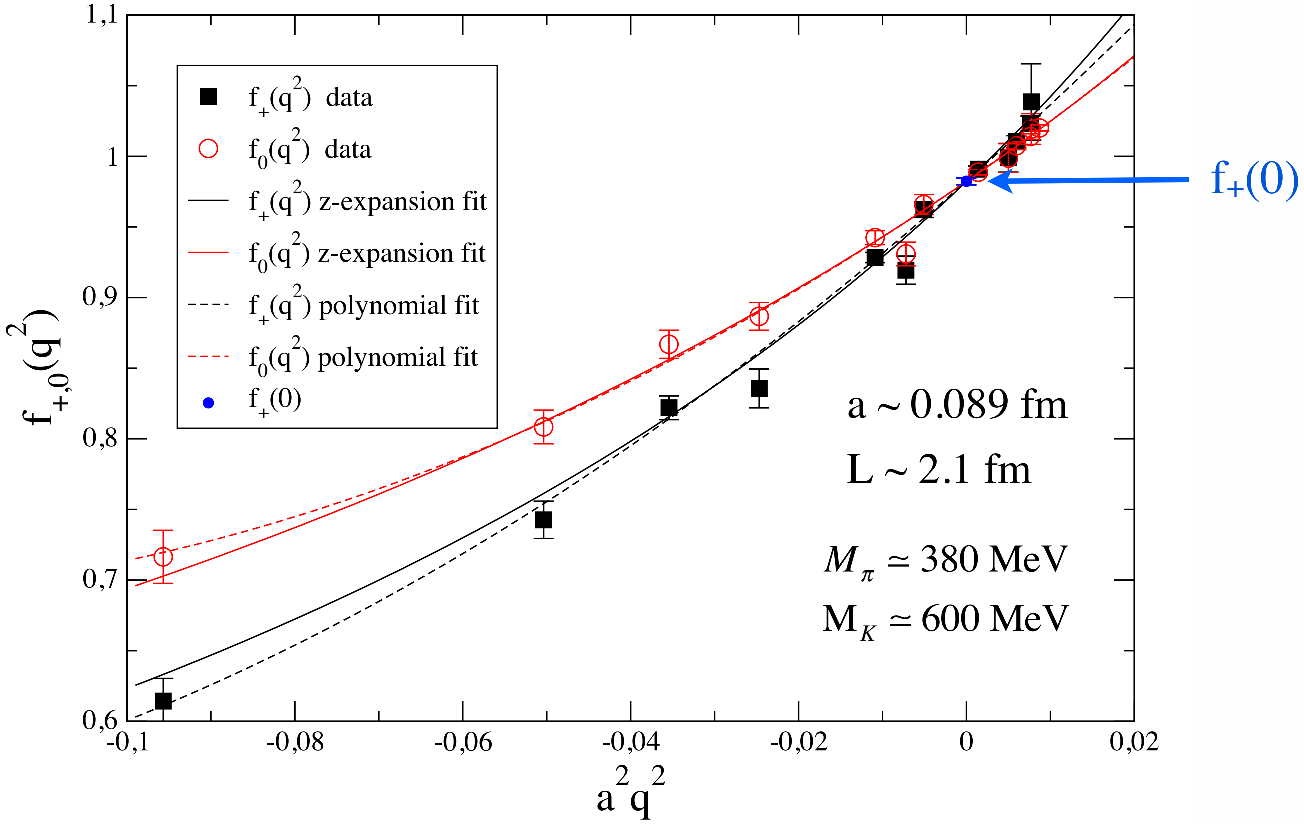

Roma Tre: arXiv1602.04113

15 gauge ensembles with $200 < M_\pi < 450$ MeV
3 lattice spacings $a \sim (0.09, 0.08, 0.06)$ fm
various lattice sizes, L, between 2 and 3 fm

$M_\pi \to 140$ MeV
$a \to 0$
$L \to \infty$

$$f_+(0) = 0.9709\,(46) \qquad [4.7\%]$$

$$|V_{us}|\, f_+(0) = 0.2165\,(4)$$

$$|V_{us}| = 0.2230\,(11) \qquad [4.9\%]$$

$|V_{ud}| = 0.97417\,(21)$ [nuclear β decay]
$|V_{ub}| = 0.00413\,(49)$ [B decays]

$$|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 0.99875\,(64)$$

test of CKM unitarity (~ 2σ tension)

need to improve the theoretical precision by a factor > 2

more computing power and better algorithms

comparison with experiments

dispersive parameterization:

$$f_+(q^2) = f_+(0)\, e^{\frac{q^2}{M_\pi^2}\left[\Lambda_+ + H(q^2)\right]}$$

$$f_0(q^2) = f_+(0)\, e^{\frac{q^2}{M_K^2 - M_\pi^2}\left[\log(C) - G(q^2)\right]}$$

calculable from Kπ phase shifts

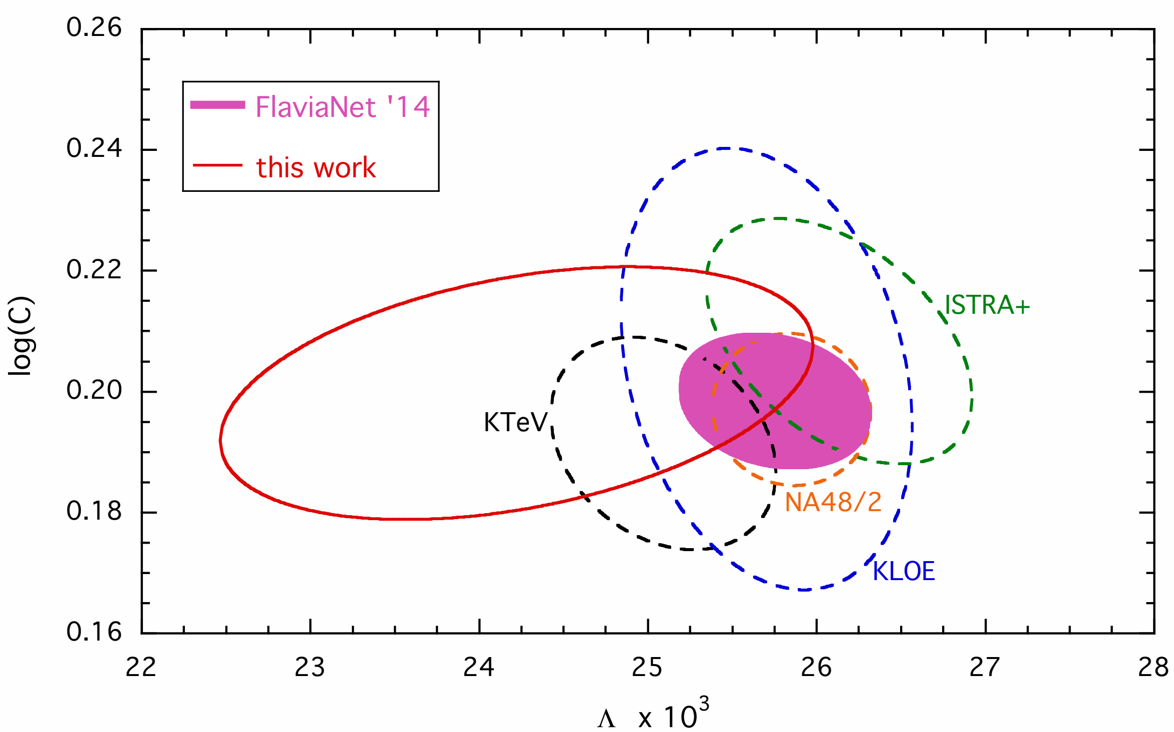

HPC machines @ Cineca used by INFN

- **FERMI** IBM-BlueGene /Q

Model: IBM-BlueGene /Q
Architecture: 10 BGQ Frame with 2 MidPlanes each
Front-end Nodes OS: Red-Hat EL 6.2
Compute Node Kernel: lightweight Linux-like kernel
Processor Type: IBM PowerA2, 1.6 GHz
Computing Nodes: 10.240 with 16 cores each
Computing Cores: 163.840
RAM: 16GB / node; 1GB/core
Internal Network: Network interface with 11 links ->5D Torus
Disk Space: more than 2PB of scratch space
Peak Performance: 2.1 PFlop/s

~ 1400 Mcorehours / anno

- **FERMI** → *nuovo Tier-0 "MARCONI"* (~ 5 x FERMI ~ seconda parte del 2016)
accelerated with MICs

- **GALILEO** (CINECA + INFN + Univ. Milano Bicocca)

Model: IBM NeXtScale
Architecture: Linux Infiniband Cluster
Nodes: 516
Processors: 2 8-cores Intel Haswell 2.40 GHz per node
Cores: 16 cores/node, 8256 cores in total
Accelerators: 2 Intel Phi 7120p per node on 384 nodes (768 in total);
2 NVIDIA K80 per node on 40 nodes (80 in total, 20 available for scientific research)
RAM: 128 GB/node, 8 GB/core
Internal Network: Infiniband with 4x QDR switches
Disk Space: 2.000 TB of local scratch
Peak Performance: 1.000 TFlop/s (to be defined)
Operating system: CentOS 7.0

~ 70 Mcorehours / anno

our jobs at the largest volume require ~1024 nodes of Fermi (1/10) or ~128 nodes of Galileo (1/5)

~ 25% of Fermi and ~ 45% of Galileo used by projects in particle physics

courtesy of L. Cosmai

LQCD on GPUs

LQCD GPU libs

QUDA

- "QCD on CUDA" `http://lattice.github.com/quda`
- Effort started at Boston University in 2008, now in wide use as the GPU solver backend for Chroma, MILC, and various other codes.
- Various solvers for several discretizations, including multi-GPU support and domain-decomposed (Schwarz) preconditioners.

cuLGT

- "CUDA Lattice Gauge Theory" `http://www.cuLGT.com`
- Evolved since 2010, developed in Graz and Tübingen.
- Main focus lies on lattice gauge fixing (Coulomb, Landau and maximally Abelian gauge) but a very general, object oriented infrastructure for lattice QCD calculations on GPUs is offered.

Performance

on Eurora @ Cineca: NVIDIA Tesla K20

- lattice size 32^4
- Quda: twisted mass conjugent gradient inverter (nondegenerate doublet of quark flavours)
- cuLGT: Landau gauge fixing with the overrelaxation algorithm

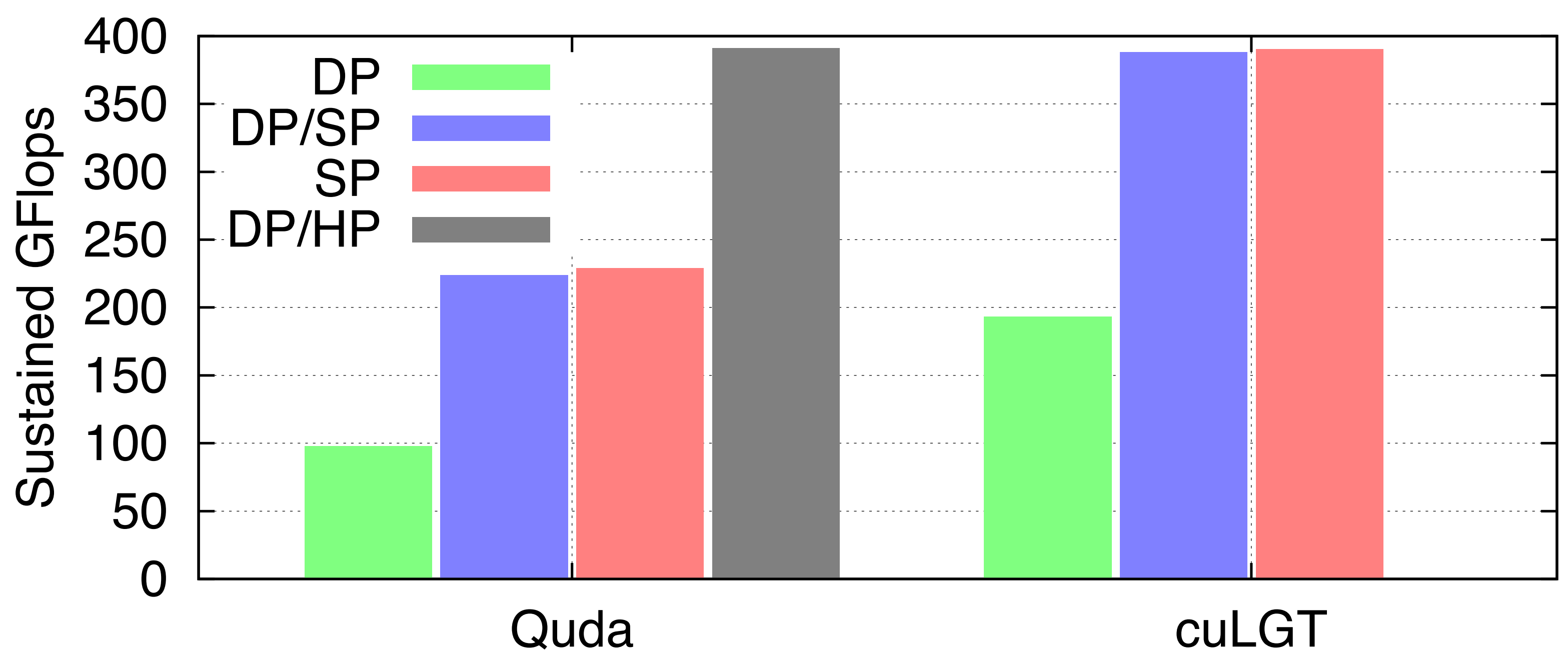

Figure 1: The different colors correspond to double precision (DP), mixed double/single precision (DP/SP), single precision (SP) and mixed double/half precision (DP/HP) (only Quda).

Comparison to previous setups

- Fermi @ Cineca: BlueGene/Q (10.240 nodes with 16 cores each)
- Processor Type: IBM PowerA2, 1.6 GHz
- setup which is in production at Roma Tre:
  - $32^3 \times 64$ lattice: twisted mass inverter with mixed double/single precision and SSE vector instructions
  - **128 Fermi nodes (2048 CPU cores)**
- compare to:
  - same lattice, equivalent inverter (Quda) with mixed double/half precision
  - **one Eurora node (two GPUs)**
- ratio of average time per inverter iteration:

$$\frac{t_{\text{Fermi}}}{t_{\text{Eurora}}} = \frac{0.009898s}{0.017016s} = 0.58$$

Multi-GPU performance

on Eurora @ Cineca: NVIDIA Tesla K20

- Quda twisted mass inverter with one flavour of quarks
- double/half mixed precision and 12 parameter reconstruction

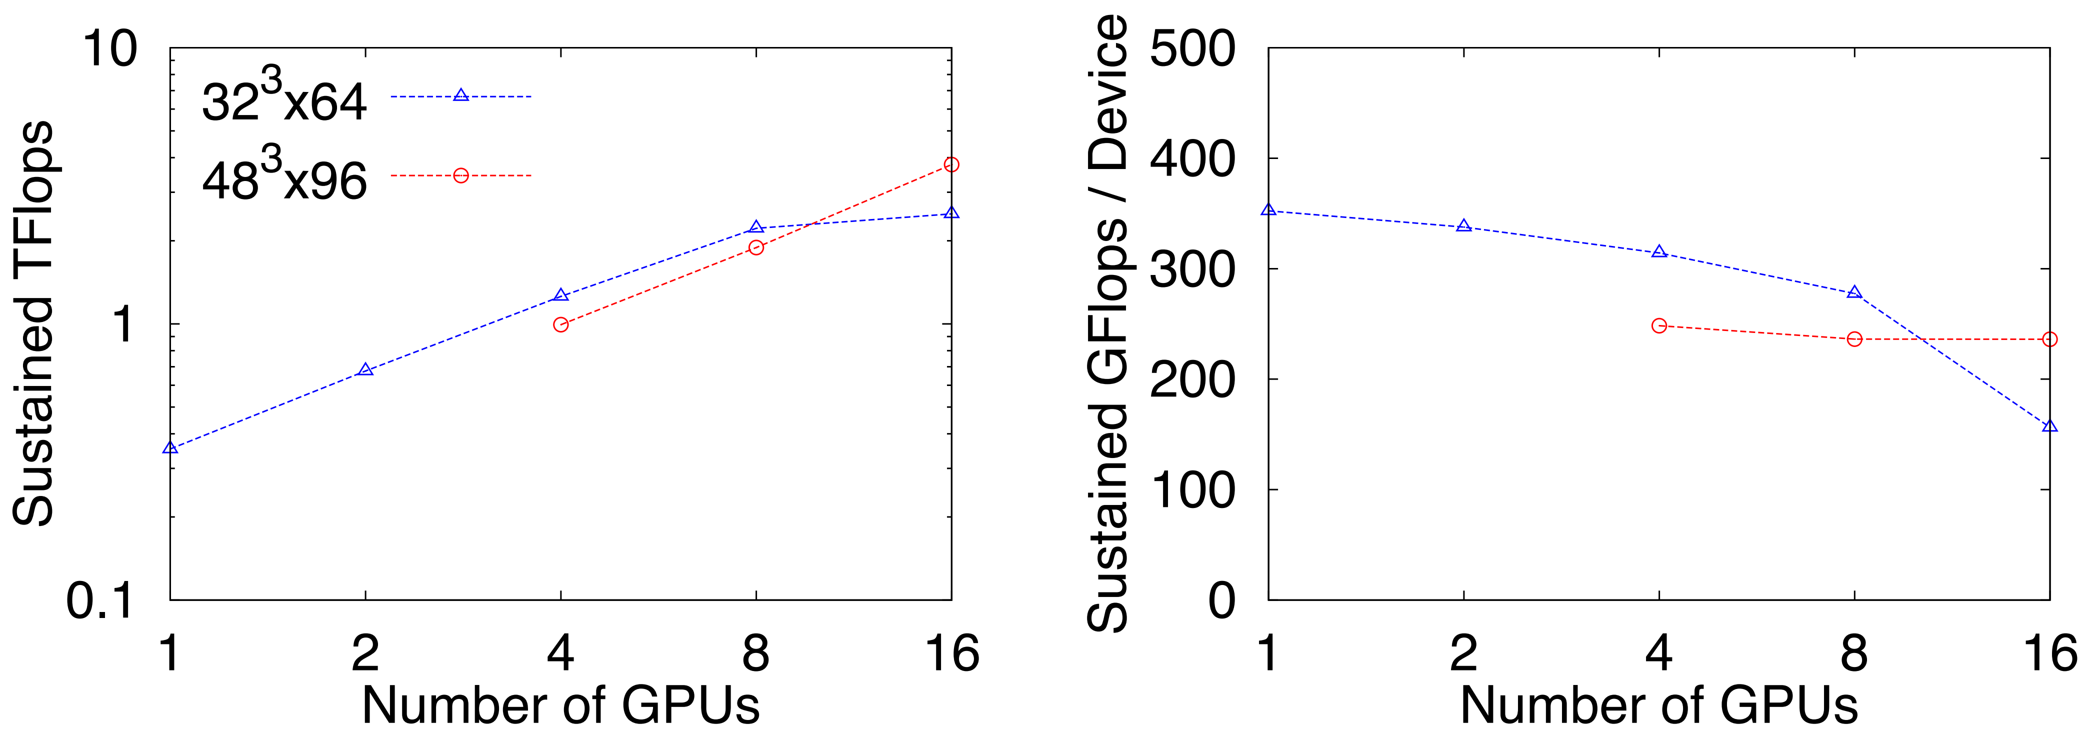

Figure 2: Strong scaling for two different lattice sizes. On the left the total performance and on the right the performance per GPU.

better performance on large local volumes

(similar results on Galileo @ Cineca: 2 NVIDIA K80 per node)

LQCD on Intel MICs

Galileo @ Cineca

516 Compute nodes:

- 2 octa-core Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz per Compute node
- 128 GB RAM per Compute node
- 2 x 16GB Intel Xeon-Phi 7120P Accelerator per Compute node (on 384 nodes)

LQCD MIC libs

- QPhiX Library implements [Balint Joo (JLAB) & Intel]
  - Wilson & Clover Dirac Operators
    - D, $A^{-1}D$ — "Dslash"
    - A χ - b D ψ — "aChiMBDPsi" -variant
    - Use these to assemble: (1 - DD) or ($A - DA^{-1}D$) Even-Odd Schur prec operators
  - Basic Solvers
    - CG, BiCGStab, Iterative Refinement (for mixed precision)
- Written in C++
  - Currently mostly geared to interfacing with QDP++/Chroma

- support for IMCI, AVX, AVX2, AVX12, SSE and QPX instructions

- threads via openMP and multi-processing via MPI

tests of computing time needed for one "Dslash" inversion

single node

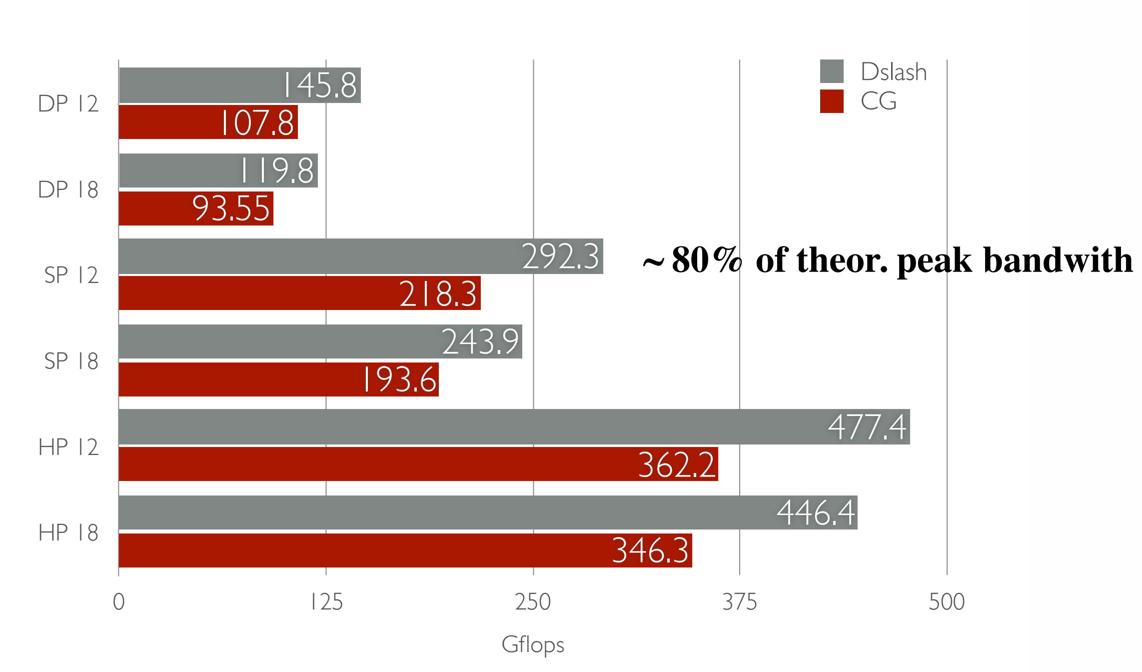

Figure 1: Twisted Mass $32^3 \times 64$ Xeon-Phi 7120P.

weak scaling is OK thanks to Intel's MPI proxy

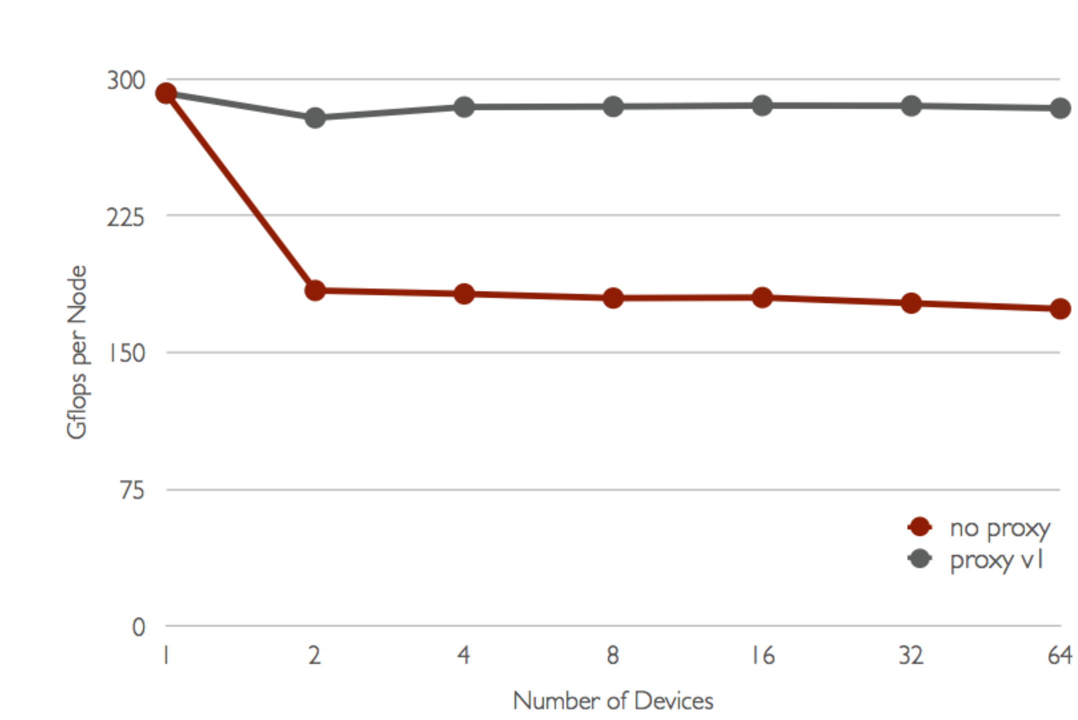

Figure 3: Weak Scaling Twisted Mass Dslash SP $48^3 \times 96$ per device Xeon-Phi 7120P.

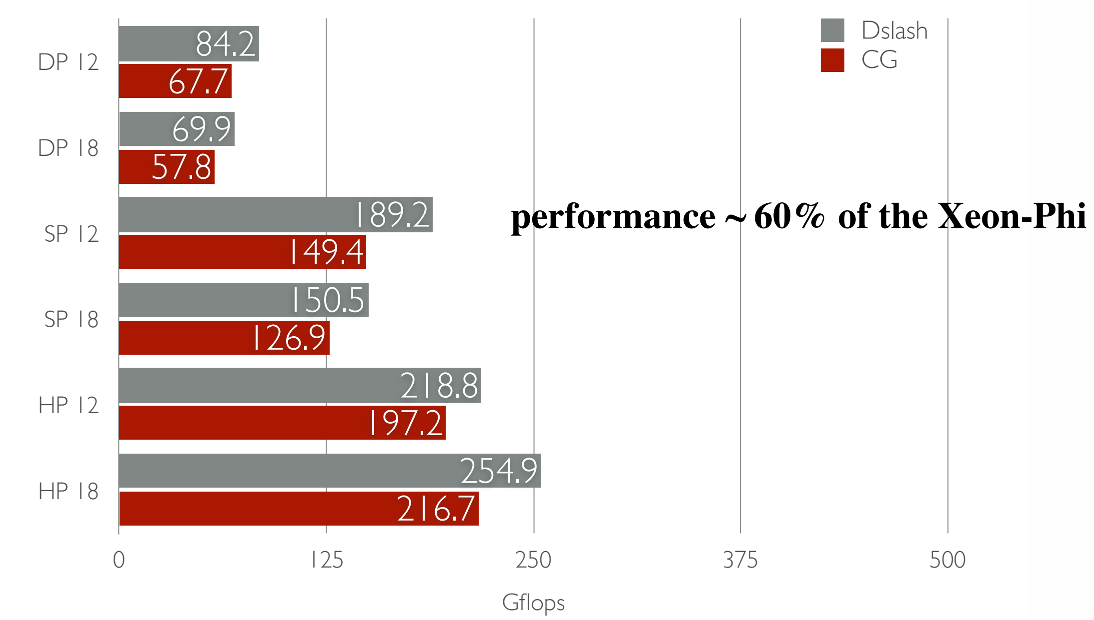

Figure 2: Twisted Mass $32^3 \times 64$ Dual Socket Xeon Haswell E5-2630 2.4GHz AVX2.

strong scaling is OK for large local volumes

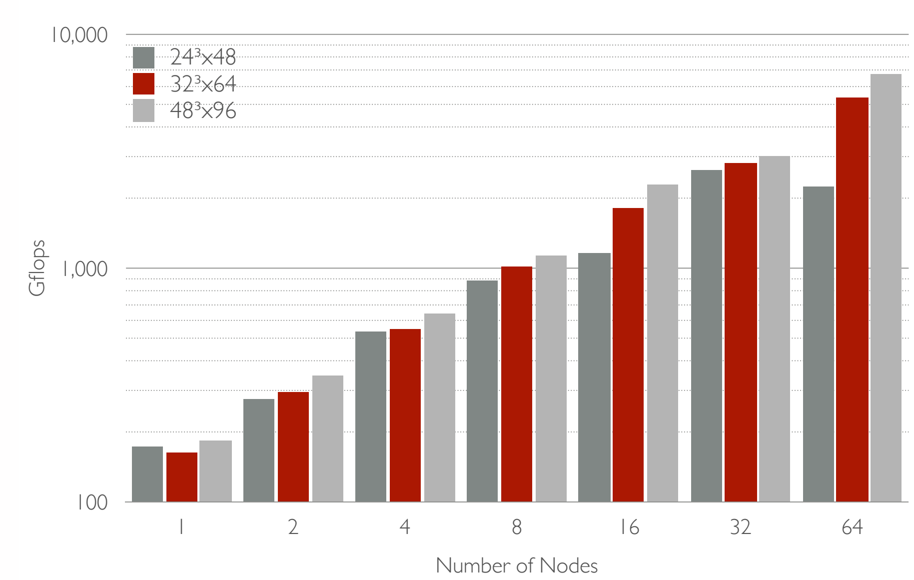

Figure 5: Strong Scaling Twisted Mass Dslash SP Compression12 Dual Socket Xeon Haswell E5-2630 2.4GHz AVX2.

Roma Tre & FNAL: arXiv 1510.08879

The SUperMAssive (SUMA) INFN project for HPC

"progetto premiale" 2013-2015: ~ 1.5 ML euro

website: https://web2.infn.it/SUMA/

An effort going into two different directions:

1) Agreements with large computing centres (CINECA)

Starting early 2010, access to 100 Mcore-hours / year on the CINECA Tier-0 machine (BG/Q) - (nicely matches ~100 Mcore-hours / year from PRACE projects with INFN PIs)

SUMA has cofunded the CINECA Tier-1 Cluster (traditional multi-core CPUs + GPUs + Intel/MIC, 1 Pflops peak)

Reserved use of 20% of the machine, i.e. ...
(additional ~30 Mcore-hours BG/Q-equivalent)

Several post-docs hired for (on average) 2 years

M. Brambilla	*(Parma, LGT)*
E. Calore	*(Ferrara, LBE)*
L. Scorzato	*(Trento, LGT)*
A. Feo	*(Parma, Grav.)*
G. Engel	*(Milano, LGT)*
M. Schrock	*(Roma 3, LGT)*
F. Negro	*(Pisa, LGT)*
G. Caruso	*(Pisa, Programming)*
F. Stellato	*(Roma 2, Q-Bio)*
P. Vilaseca Mainar	*(Roma, LGT)*
L. Riggio	*(Roma3, LGT)*

OUTLOOK

******* new INFN project, involving experiments and HPC, recently approved by CIPE *******

possible scenario

*** 2-year project (2017-2018)**

*** ~ 3.5 - 4 ML euro for HPC machines**

*** ~ 0.9 ML euro for grants devoted to porting and optimization of parallelized codes**

backup slides

Where we and the world have been, are and will be ...

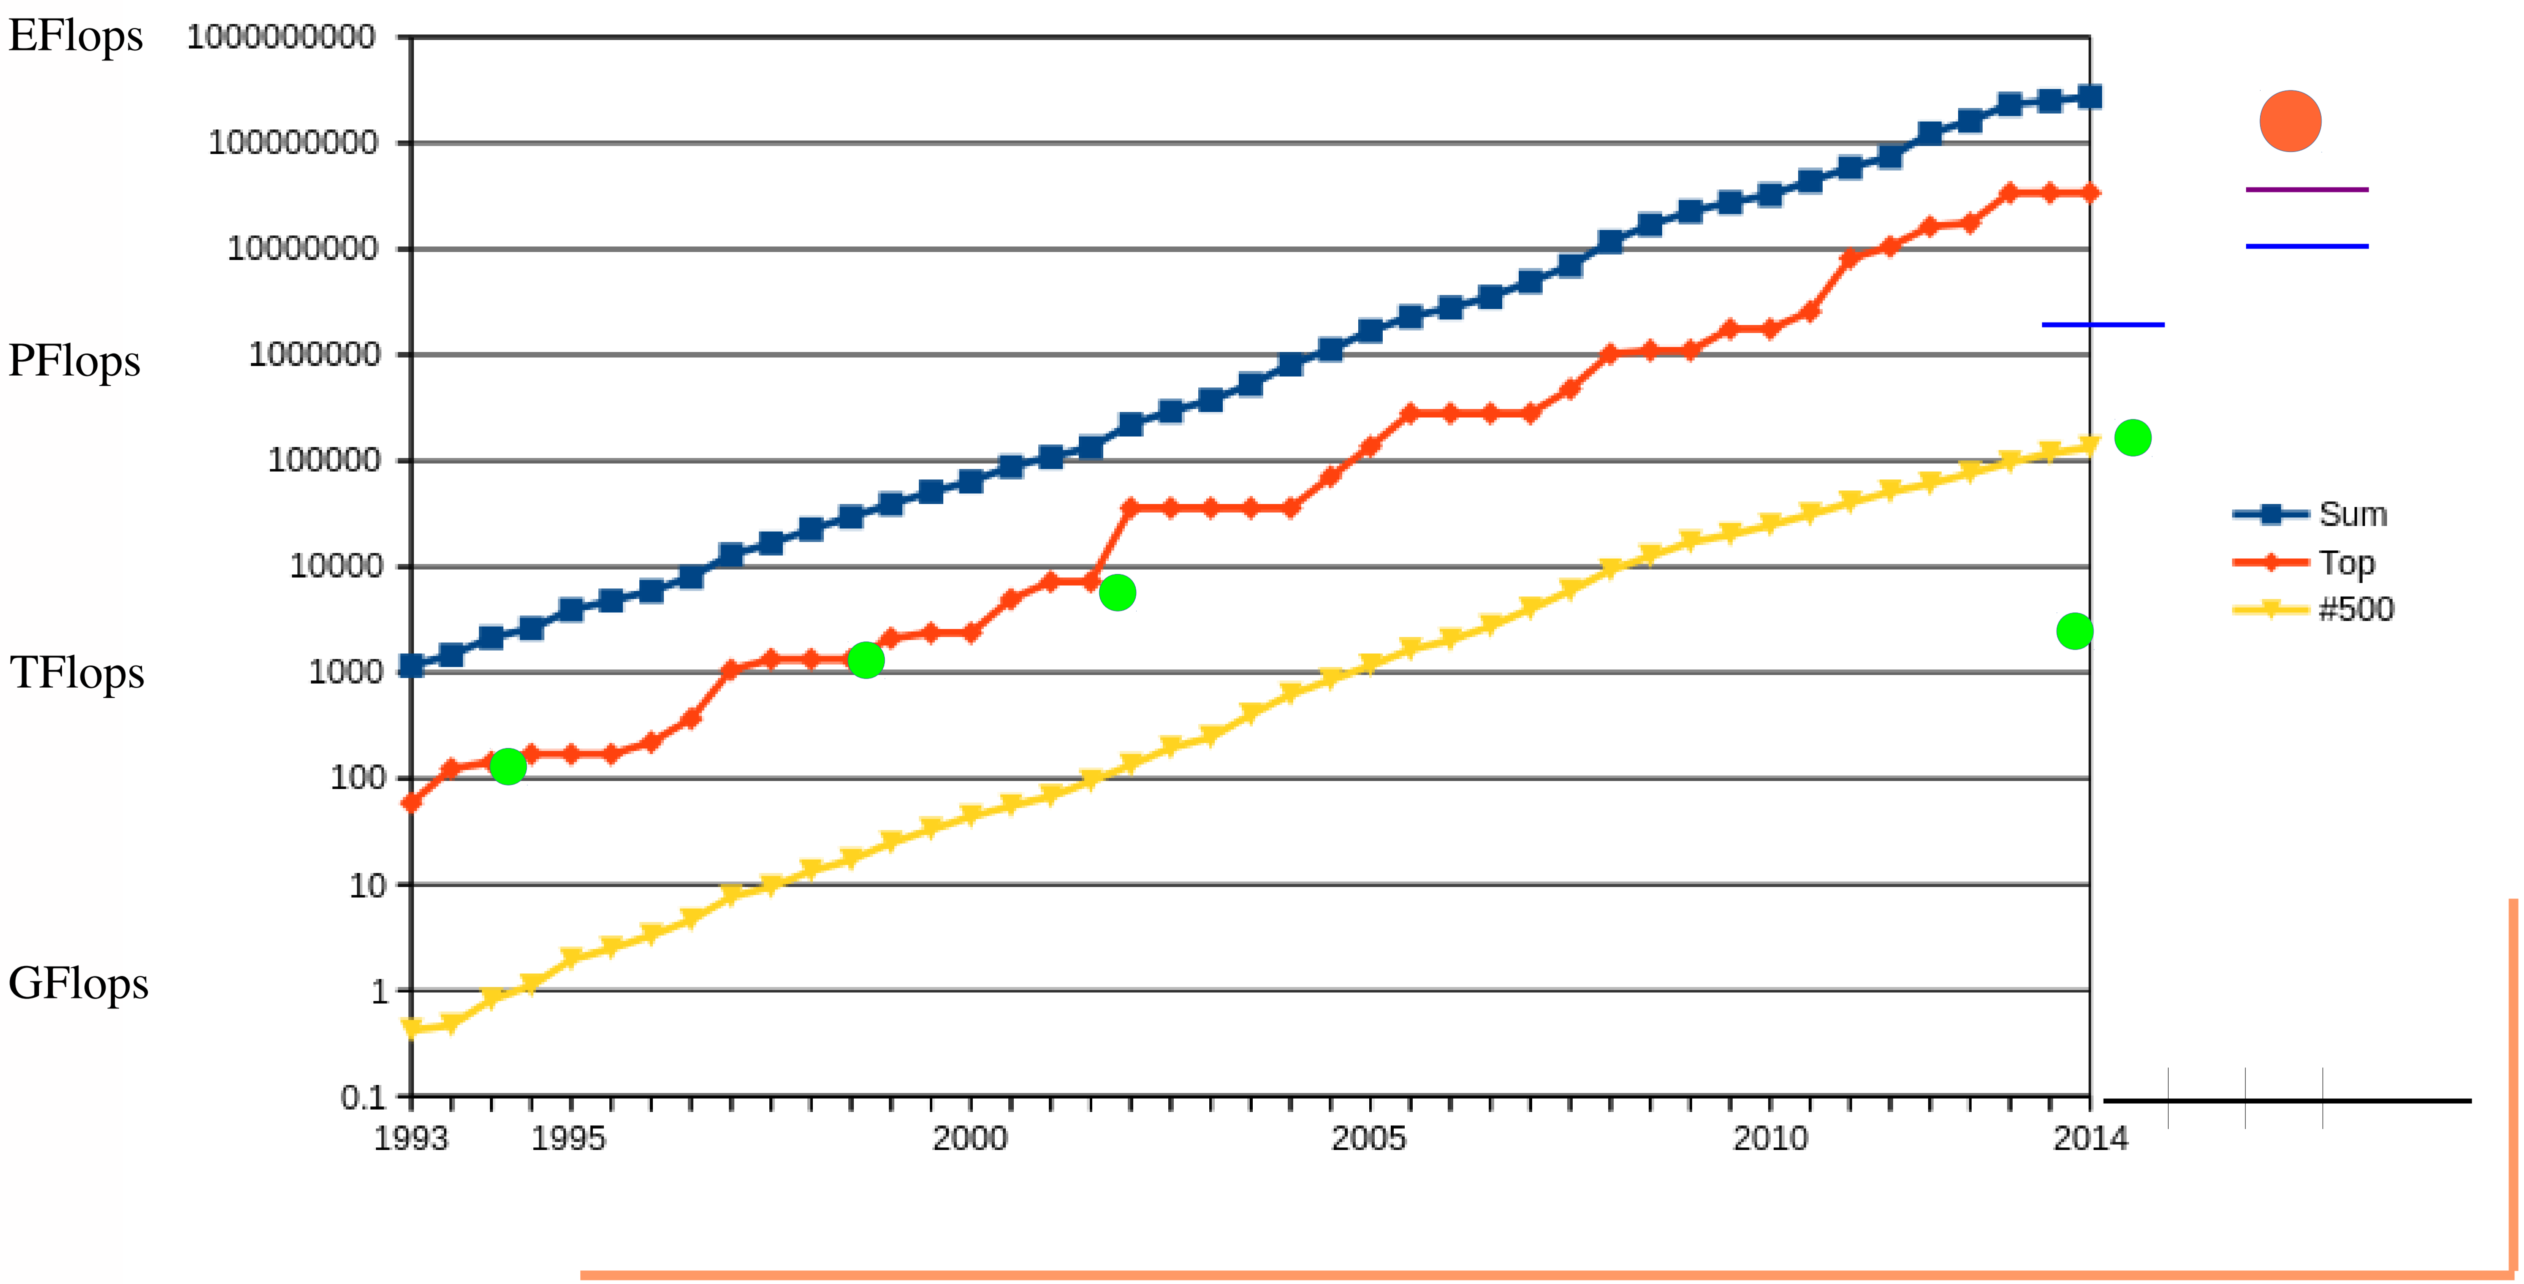

courtesy of R. Tripiccione

Al CINECA, upgrade da FERMI → MARCONI (from China with love)

Giugno 2016: *2 Pflops (Cluster Intel Broadwell)*

Settembre 2016: *2 Pflops (BW) + 11 Pflops (MIC[new gen])*

Giugno 2017: *7 Pflops (Cluster) + 11 Pflops (MIC)*

Warning: 4 Pflops "reserved" for Eurofusion

Money: 18 Pflops @ 26 M€ → ~ 1.5 M€ / Pflops
[+300 K€ /(Pflops year) electricity bill]

Hardware comparison

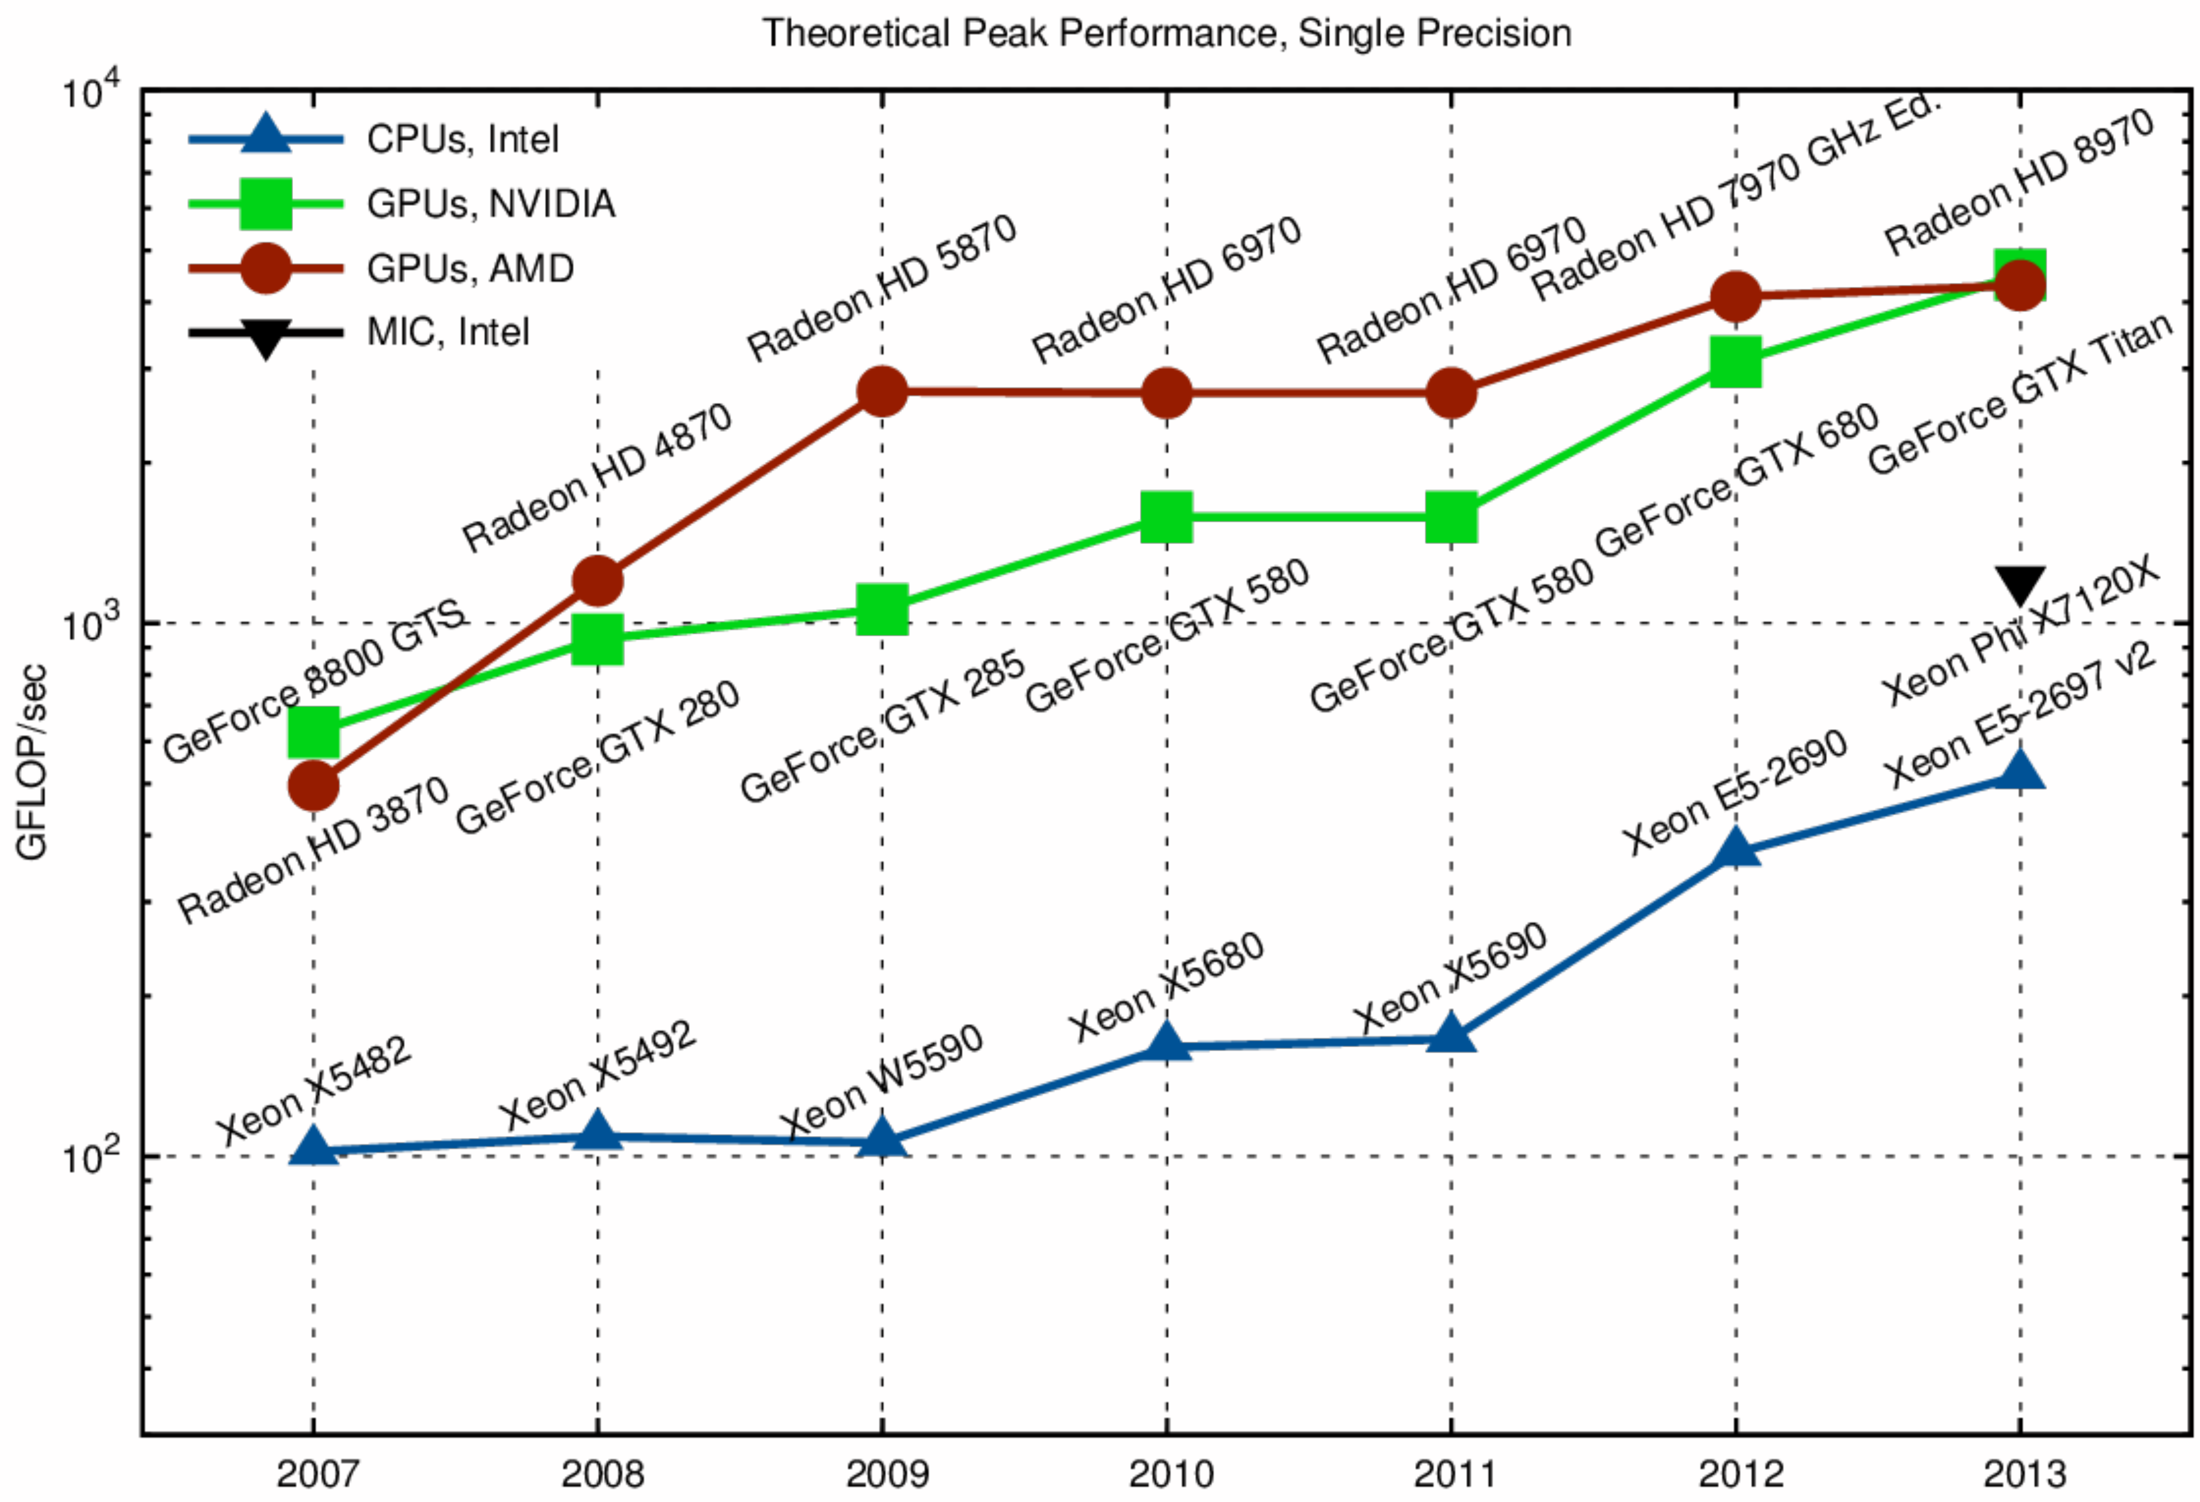

[karlrupp.net]

QPhiX performance

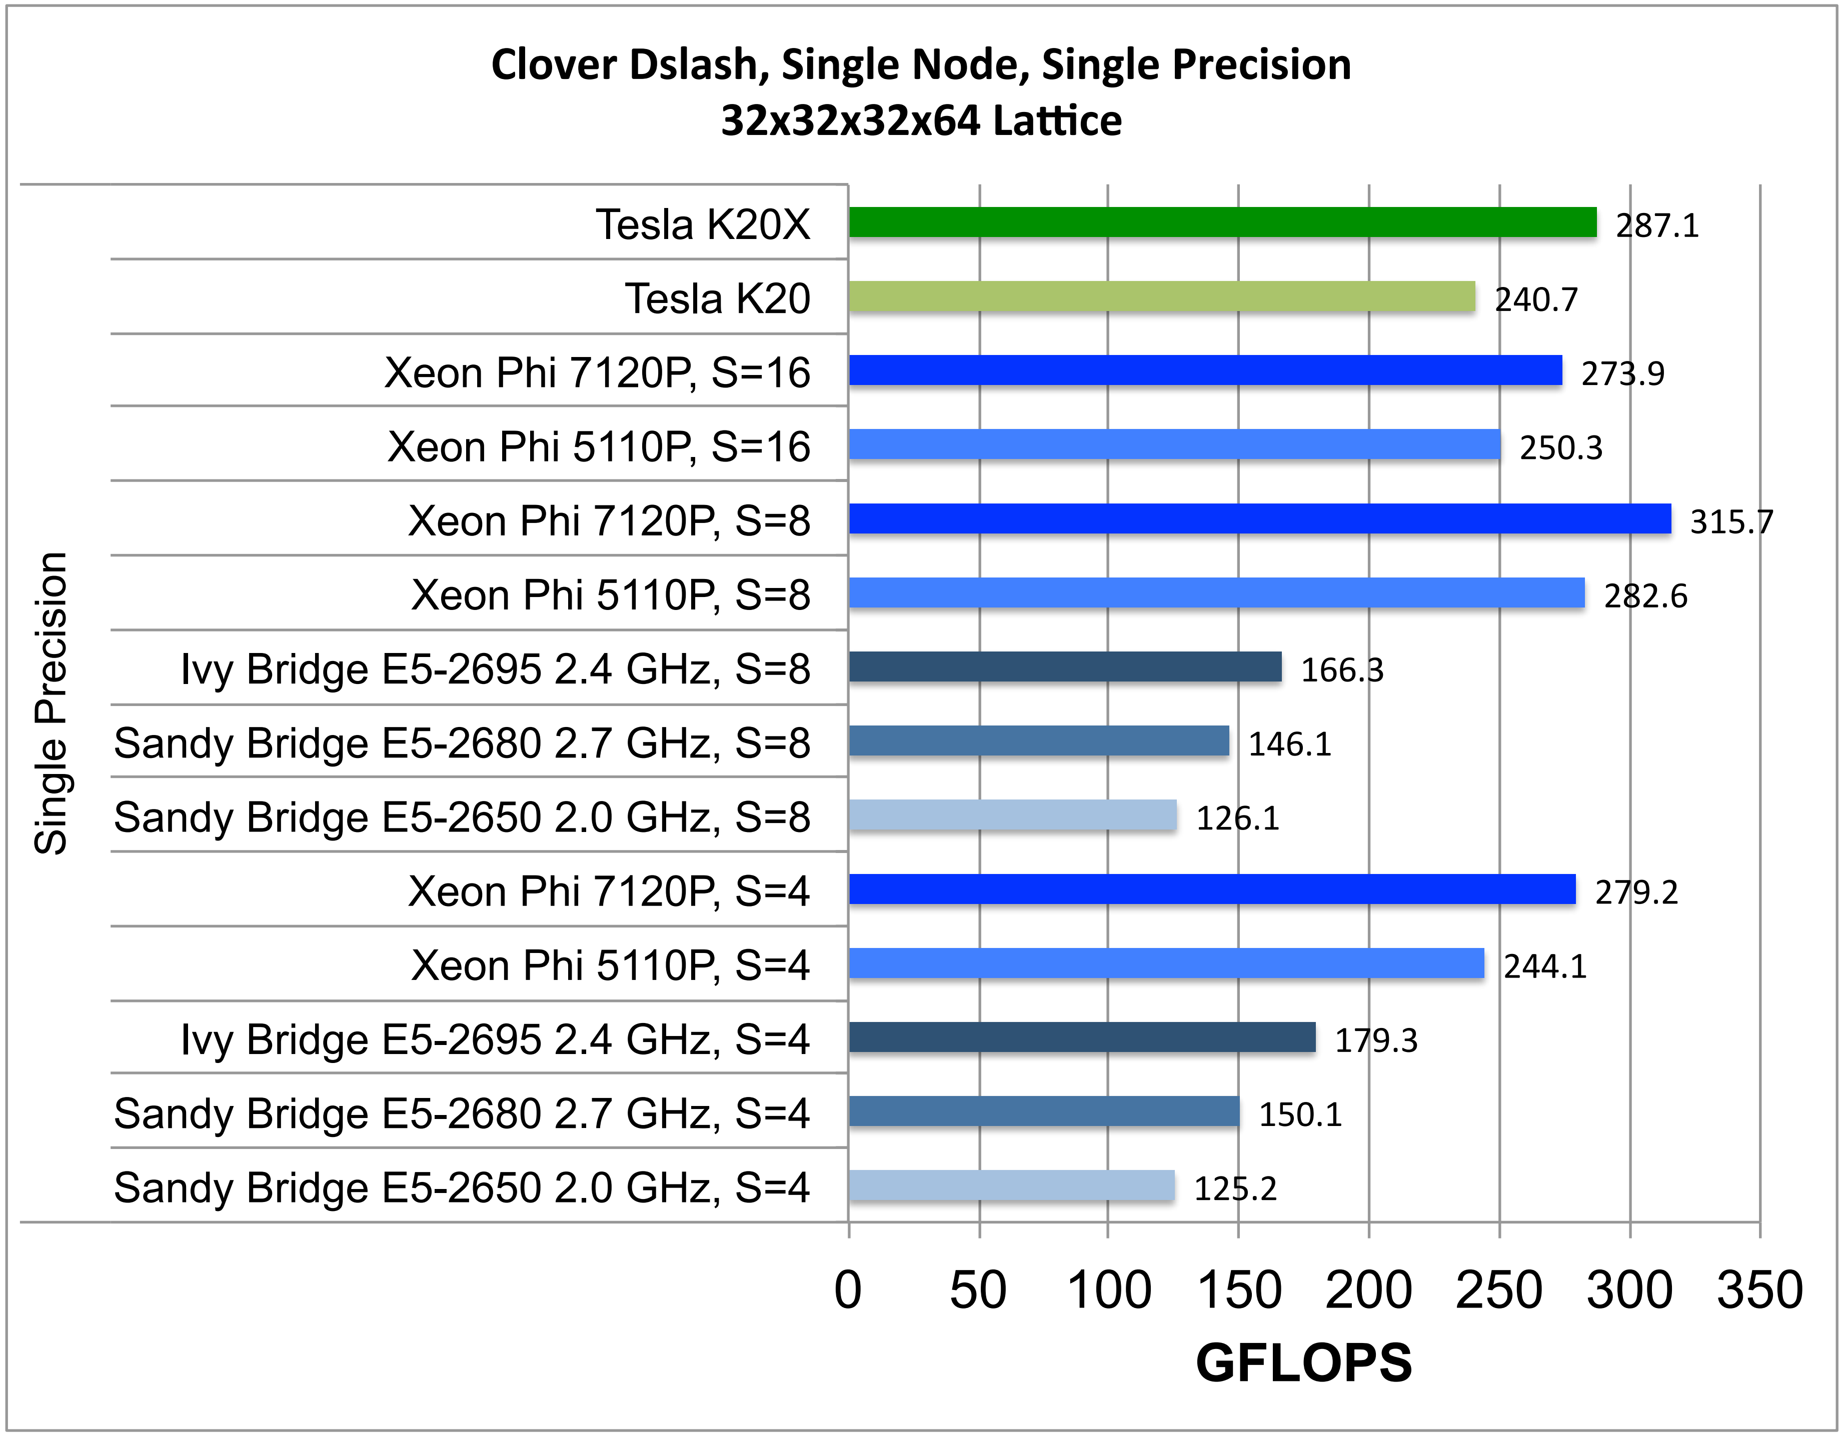

[Balint Joo (JLAB) & Intel]