QCD and the Testing of Fundamental Symmetries in Nuclei

André Walker-Loud for CalLat

- Science Motivation
- Quarks to Nuclei: tensor contractions and threads
- Parallel I/O: HDF5 for lattice QCD

CalLat Team

LLNL

Bronis R. de Supinski	Performance	SUPER	
Robert Falgout	Solvers	FASTMath	
Tom Luu	LQCD, effective theory	Co-Dir., Science	(LLNL ➡ Bonn/Jülich)
Ron Soltz	LQCD		
Pavlos Vranas	LQCD		

LBNL/UC Berkeley

Wick Haxton	Effective theory	PI	
Esmond Ng	Linear algebra	Co-Dir., Computing; FASTMath	
Andre Walker-Loud	LQCD		(LBNL ➡ William and Mary/Jefferson Lab)
Sam Williams	Performance	SUPER	
Chao Yang	Linear algebra	FASTMath	

Nvidia

Michael Clark	LQCD/GPUs

LBNL/UC Merced

Juan Meza	Performance

Young researchers: Abhinav Sarje (*LBNL*), Sergey Syritsyn (*LBNL* ➡ *BNL*)
Michael Buchoff (LLNL ➡ Seattle),
Chris Schroeder, Joe Wasem *LLNL*
Thorsten Kurth (Wupertal ➡ LBNL)

Science Motivation

Our goal is to solve the theory of QCD

QuantumChromoDynamics (QCD) is the fundamental theory of (nuclear) strong interactions

- fusion (nuclear reactions in stars)
- fission (radioactive decay)

Understanding QCD is necessary for searching for physics beyond the Standard Model - low-energy precision tests of fundamental symmetries (eg. Parity)

Our goal is to solve the theory of QCD

QCD describes the interactions between quarks and gluons in 3 space and 1 time dimension

- set of strongly coupled, non-linear differential equations

We need BIG computers

Our goal is to solve the theory of QCD

Fun facts about QCD:

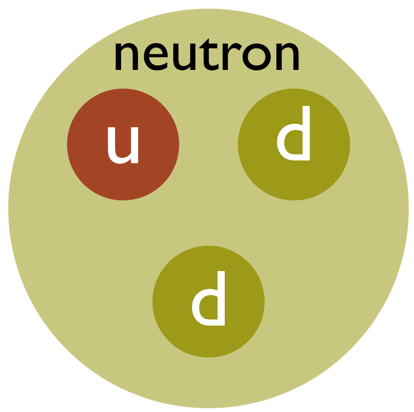

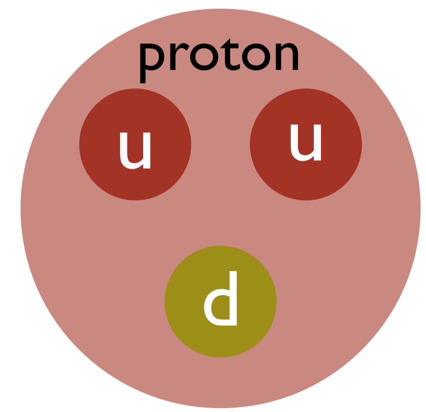

$$m_d \simeq 4.7 \text{ MeV} \qquad m_u \simeq 2.2 \text{ MeV}$$

$$M_n = 939.6 \text{ MeV} \qquad M_p = 938.3 \text{ MeV}$$

$$\lim_{m_{u,d} \to 0} M_{n,p} \simeq 900 \text{ MeV}$$

~95% of known mass in universe is from QCD

Our goal is to solve the theory of QCD

Fun facts about QCD:

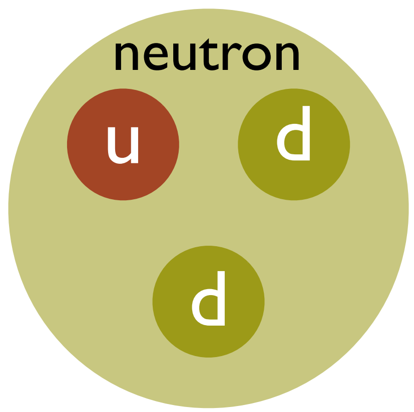

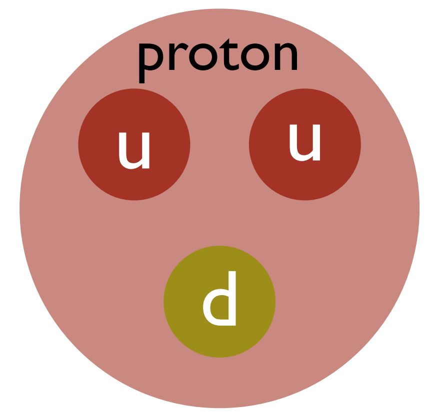

$$m_d \simeq 4.7 \text{ MeV} \qquad m_u \simeq 2.2 \text{ MeV}$$

$$M_n = 939.6 \text{ MeV} \qquad M_p = 938.3 \text{ MeV}$$

$$\lim_{m_{u,d} \to 0} M_{n,p} \simeq 900 \text{ MeV}$$

~95% of known mass in universe is from QCD

NOTE $M_n > M_p$

very important consequences for our universe

Our goal is to solve the theory of QCD

Fun facts about QCD:

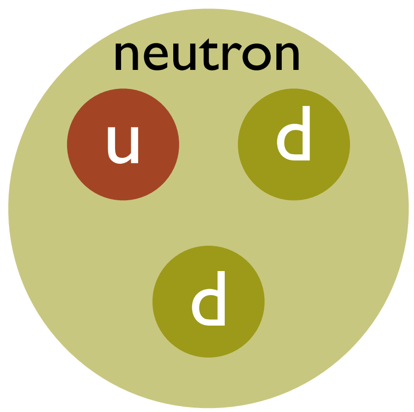

$$M_n > M_p$$

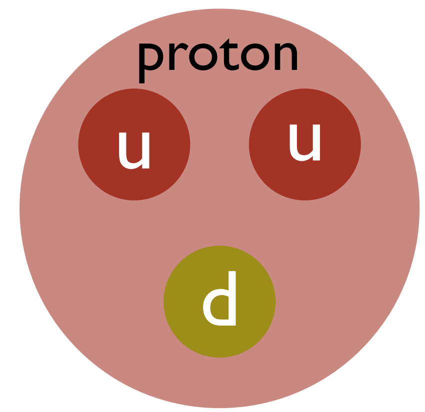

free neutrons

$$\mathrm{n} \longrightarrow \mathrm{p}$$

$$\tau_n \sim 15 \text{ min}$$

Our goal is to solve the theory of QCD

Fun facts about QCD:

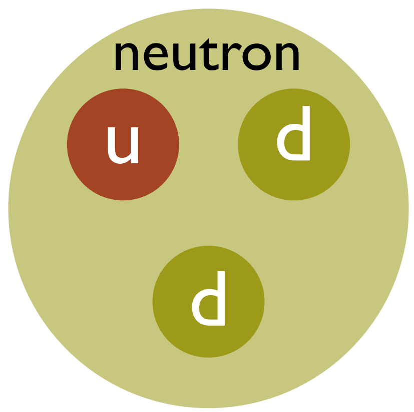

$$M_n > M_p$$

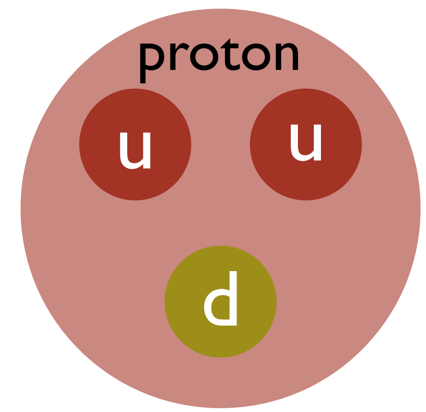

n $\xrightarrow[\tau_n \sim 15 \text{ min}]{\text{free neutrons}}$ p

Why where there any neutrons left in the primordial universe?

Formation of nuclei

Our goal is to solve the theory of QCD

Fun facts about QCD:

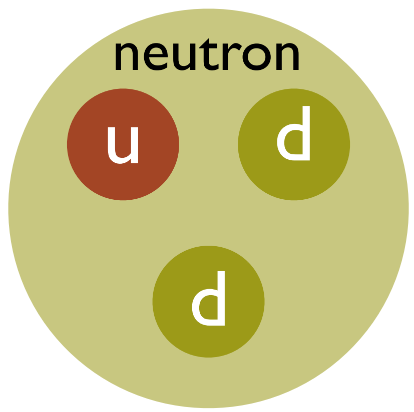

$$M_n > M_p$$

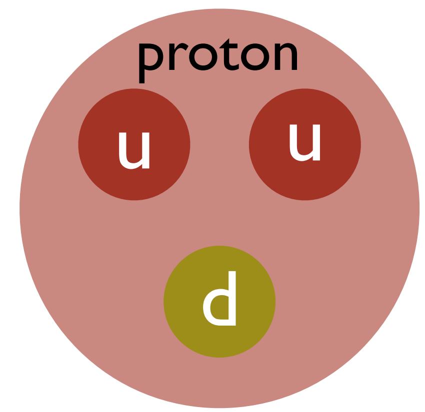

$$\mathrm{n} \xrightarrow[\tau_n \sim 15\ \mathrm{min}]{\text{free neutrons}} \mathrm{p}$$

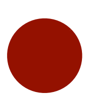

Hydrogen

$$M_p = 938.3\ \mathrm{MeV}$$

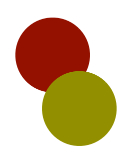

Deuterium

$$B_d \simeq 2.2\ \mathrm{MeV}$$

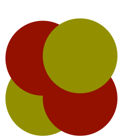

Helium

$$B_{^4He} \simeq 28.3\ \mathrm{MeV}$$

$$M_n - M_p = 1.29333217(42)\ \mathrm{MeV}$$

Our goal is to solve the theory of QCD

Fun facts about QCD:

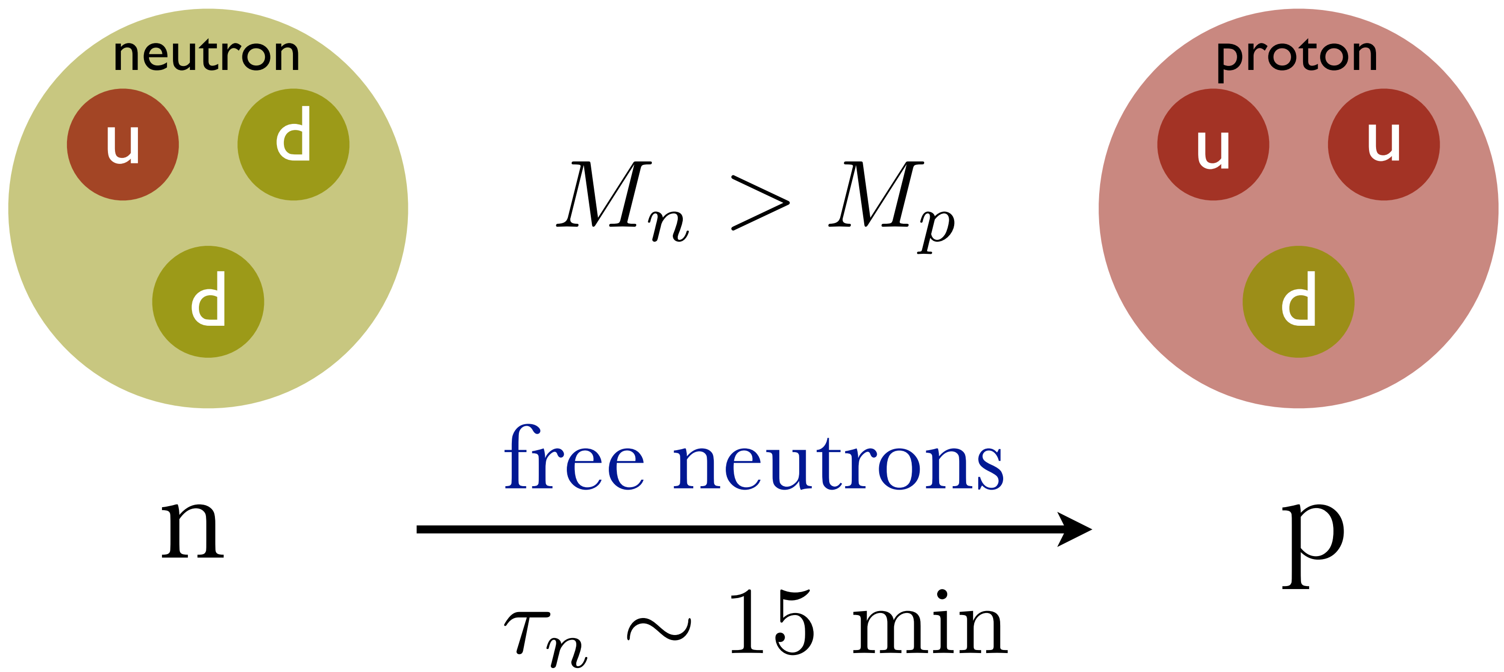

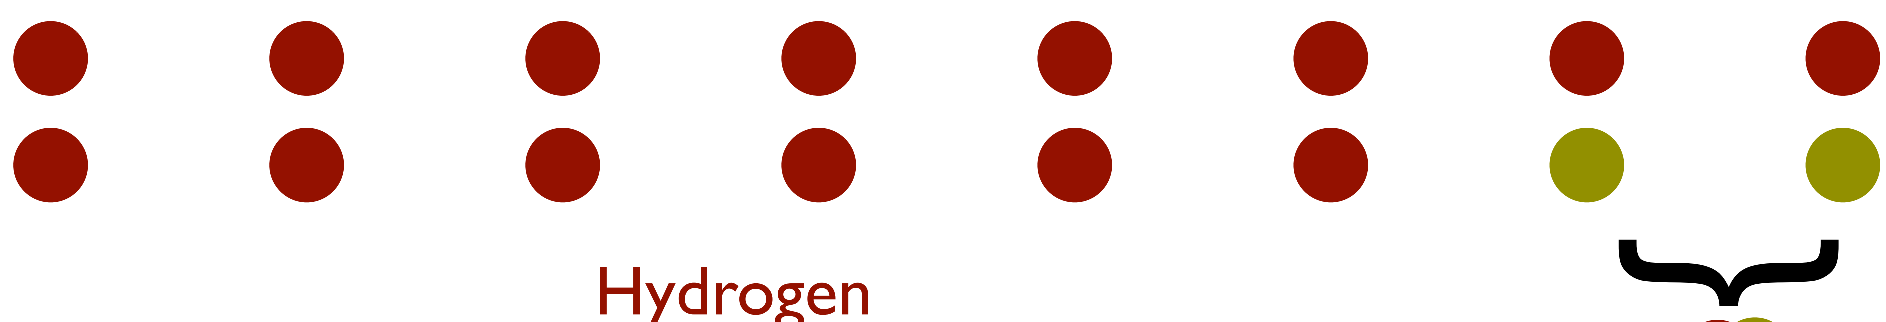

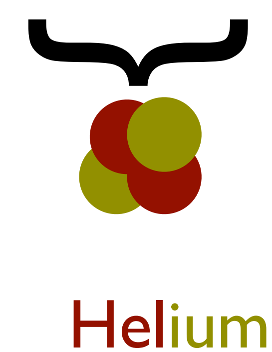

The early universe contains 75% H and 25% ^{4}He by mass fraction ("all" deuterium converted to ^{4}He)

Our goal is to solve the theory of QCD

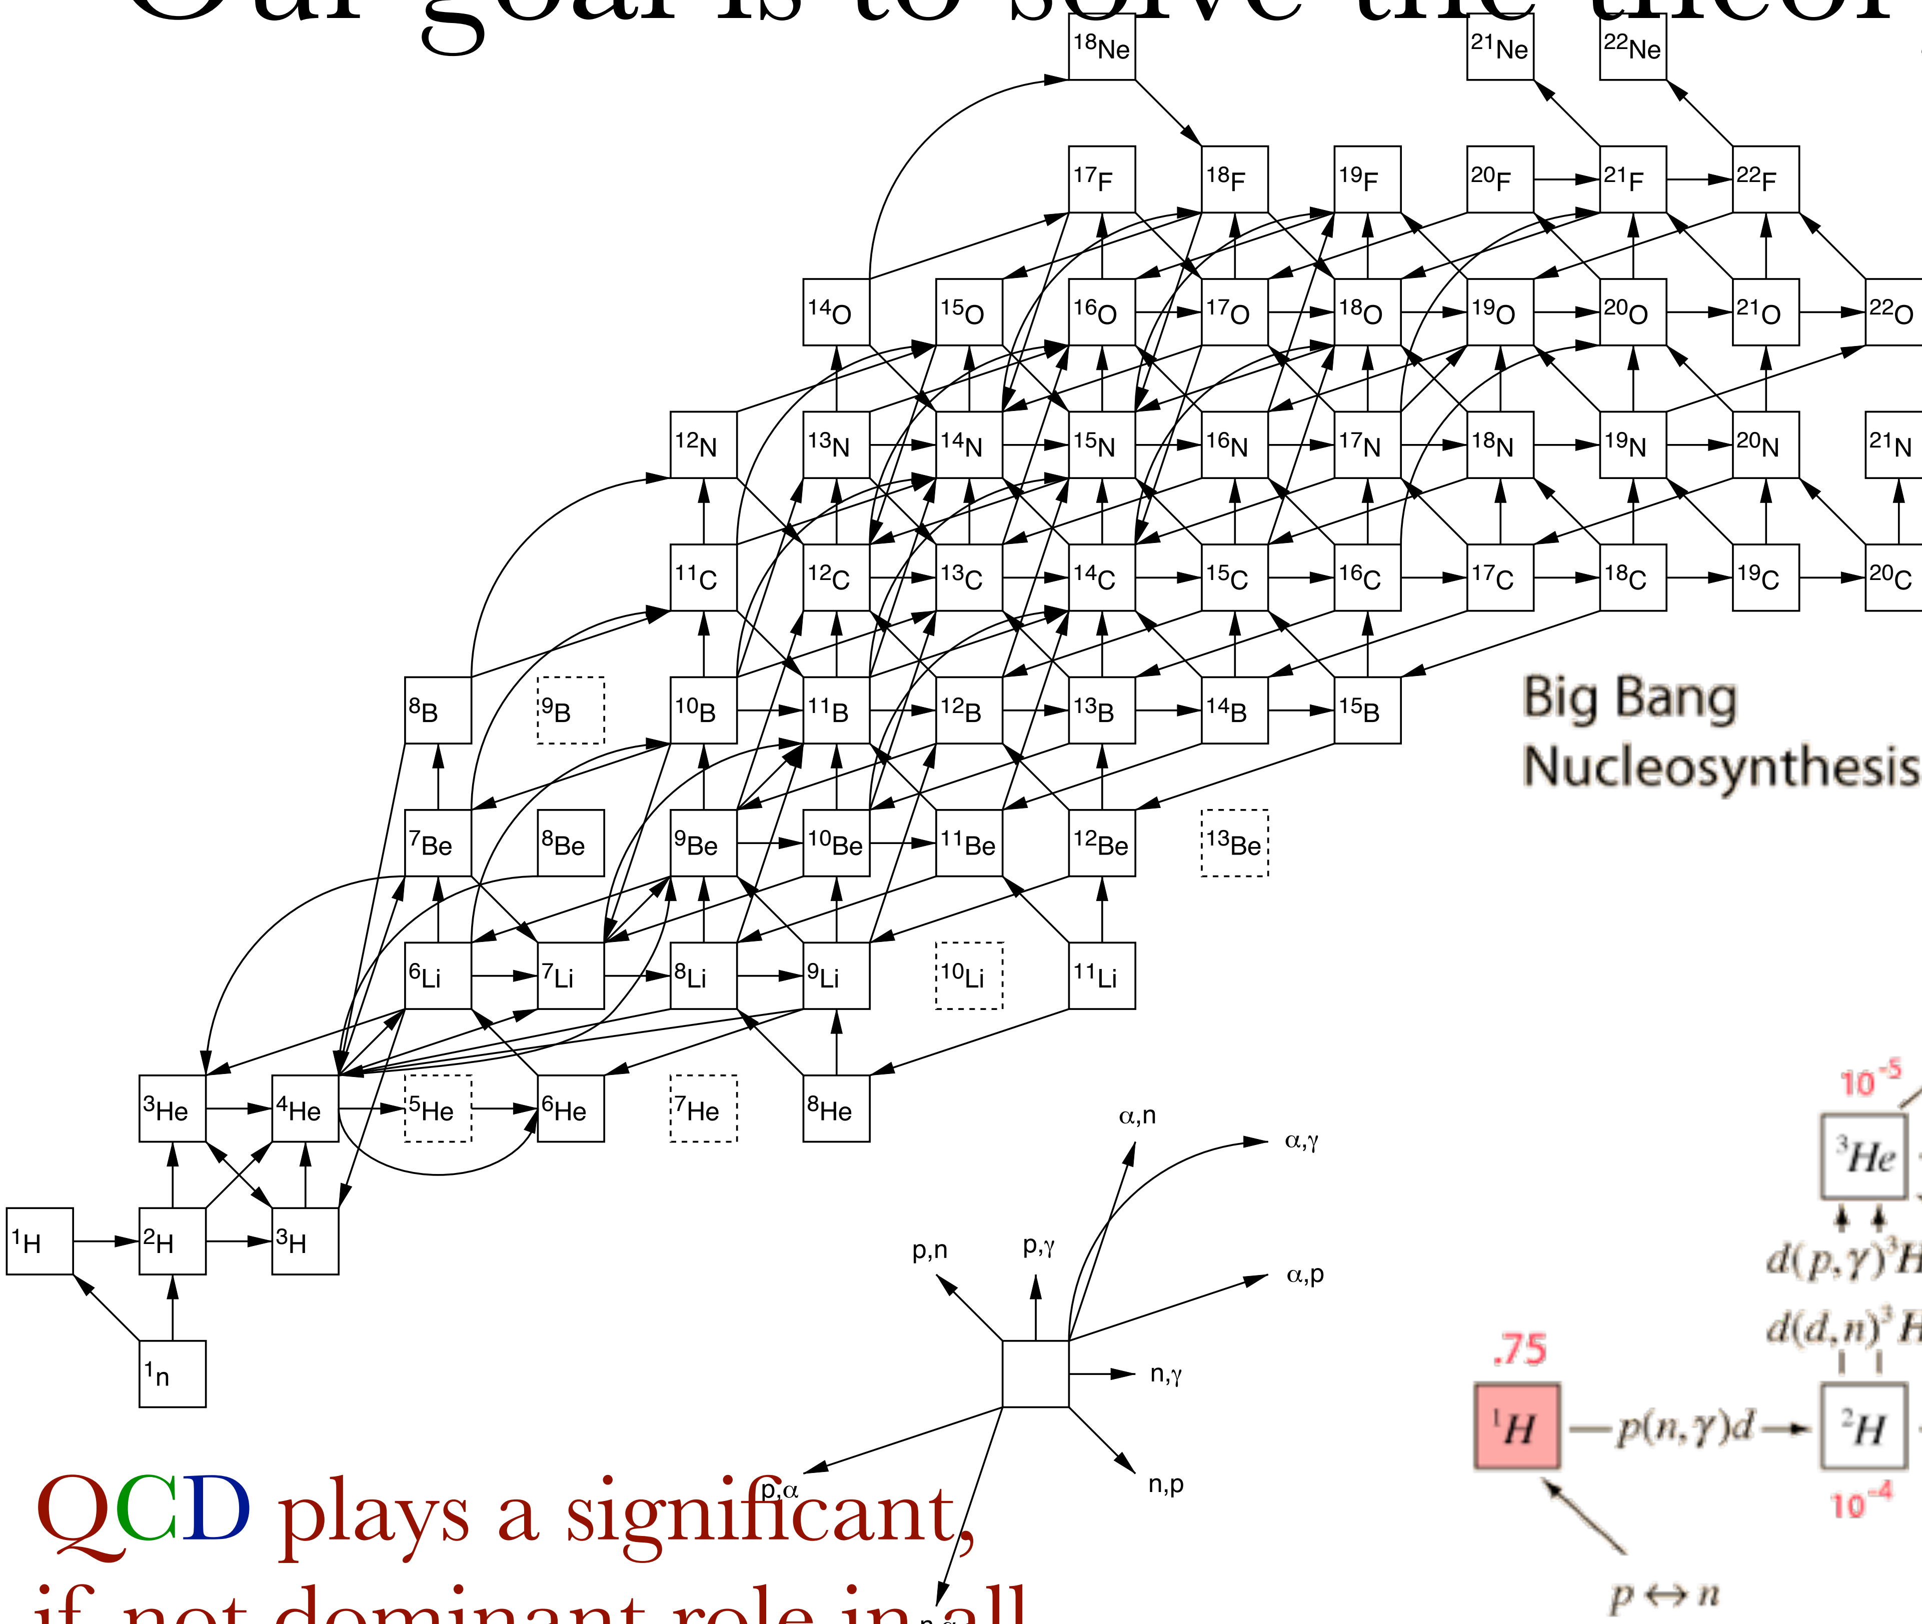

Big Bang Nucleosynthesis

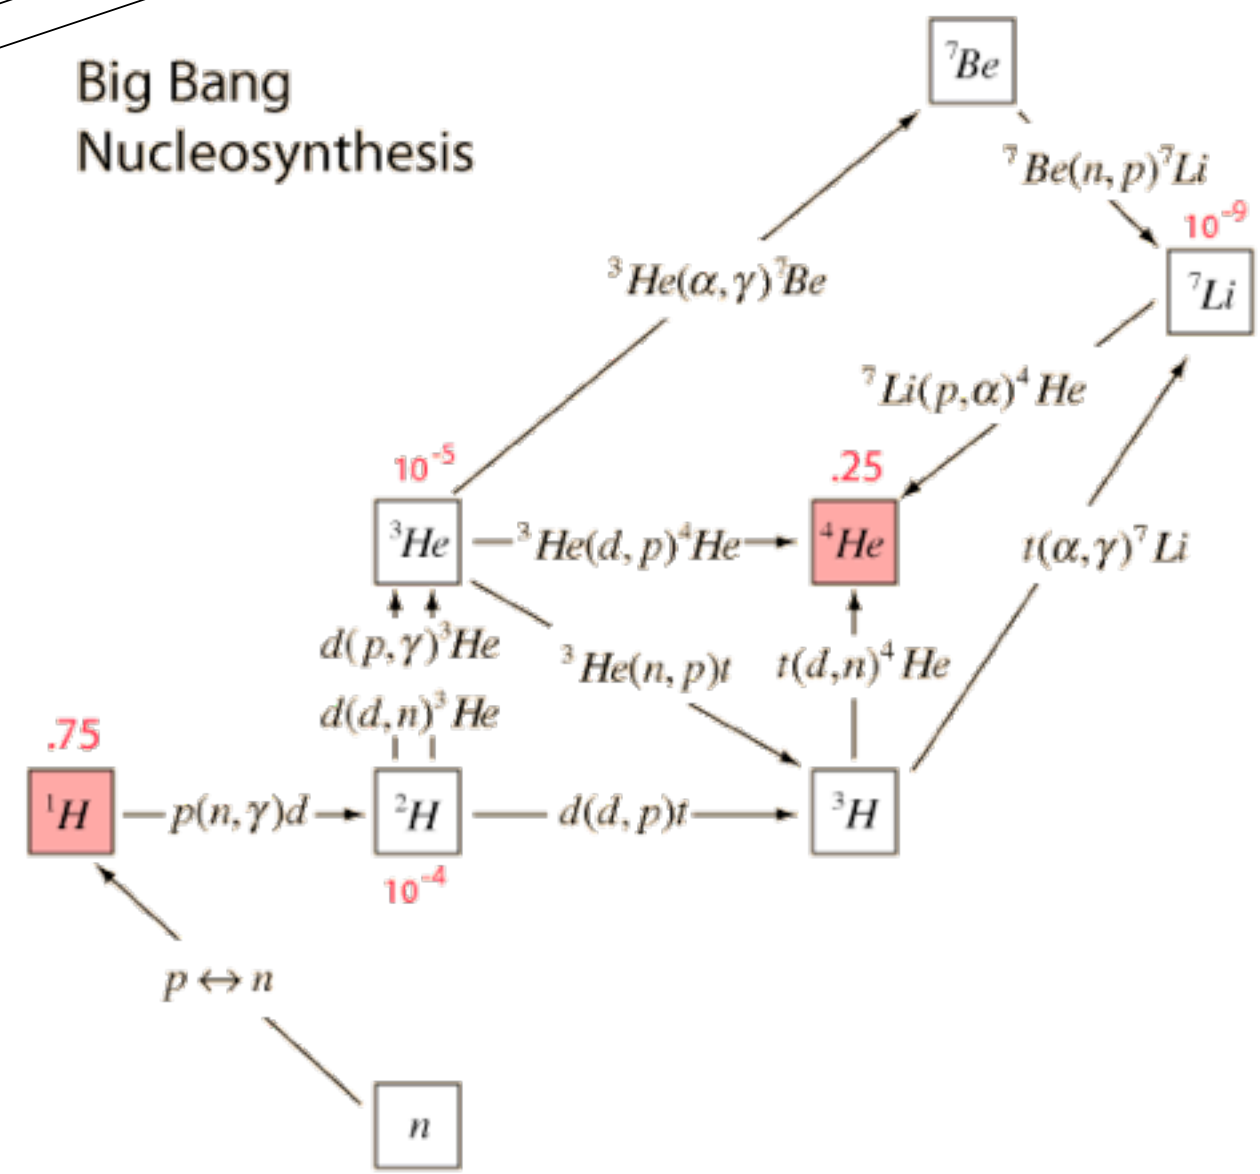

QCD plays a significant, if not dominant role in all these reactions

Our goal is to solve the theory of QCD

- M_n - M_p plays an extremely significant role in the evolution of the universe as we know it

Initial conditions for Big Bang Nucleosynthesis (BBN)

$$\frac{X_n}{X_p} = e^{-\frac{M_n - M_p}{T}}$$

- The neutron lifetime is highly sensitive to the value of this mass splitting

$$\frac{1}{\tau_n} = \frac{(G_F \cos\theta_C)^2}{2\pi^3} m_e^5 (1 + 3g_A^2)\, f\left(\frac{M_n - M_p}{m_e}\right)$$

Point Nucleons $f(a) \simeq \frac{1}{15}\left(2a^4 - 9a^2 - 8\right)\sqrt{a^2-1} + a\ln\left(a + \sqrt{a^2-1}\right)$

Griffiths "Introduction to Elementary Particles"

10% change in $M_n - M_p$ corresponds to ~100% change neutron lifetime

Our goal is to solve the theory of QCD

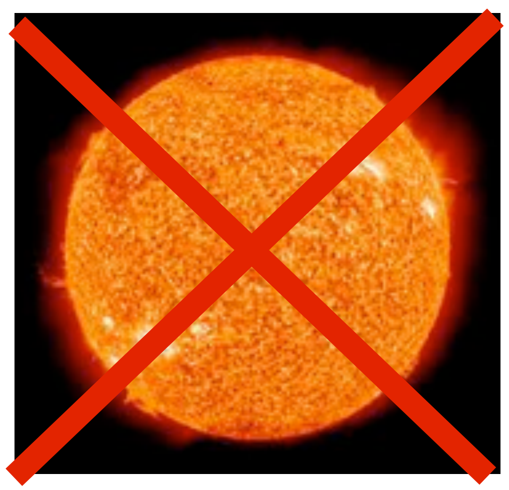

No Sun!

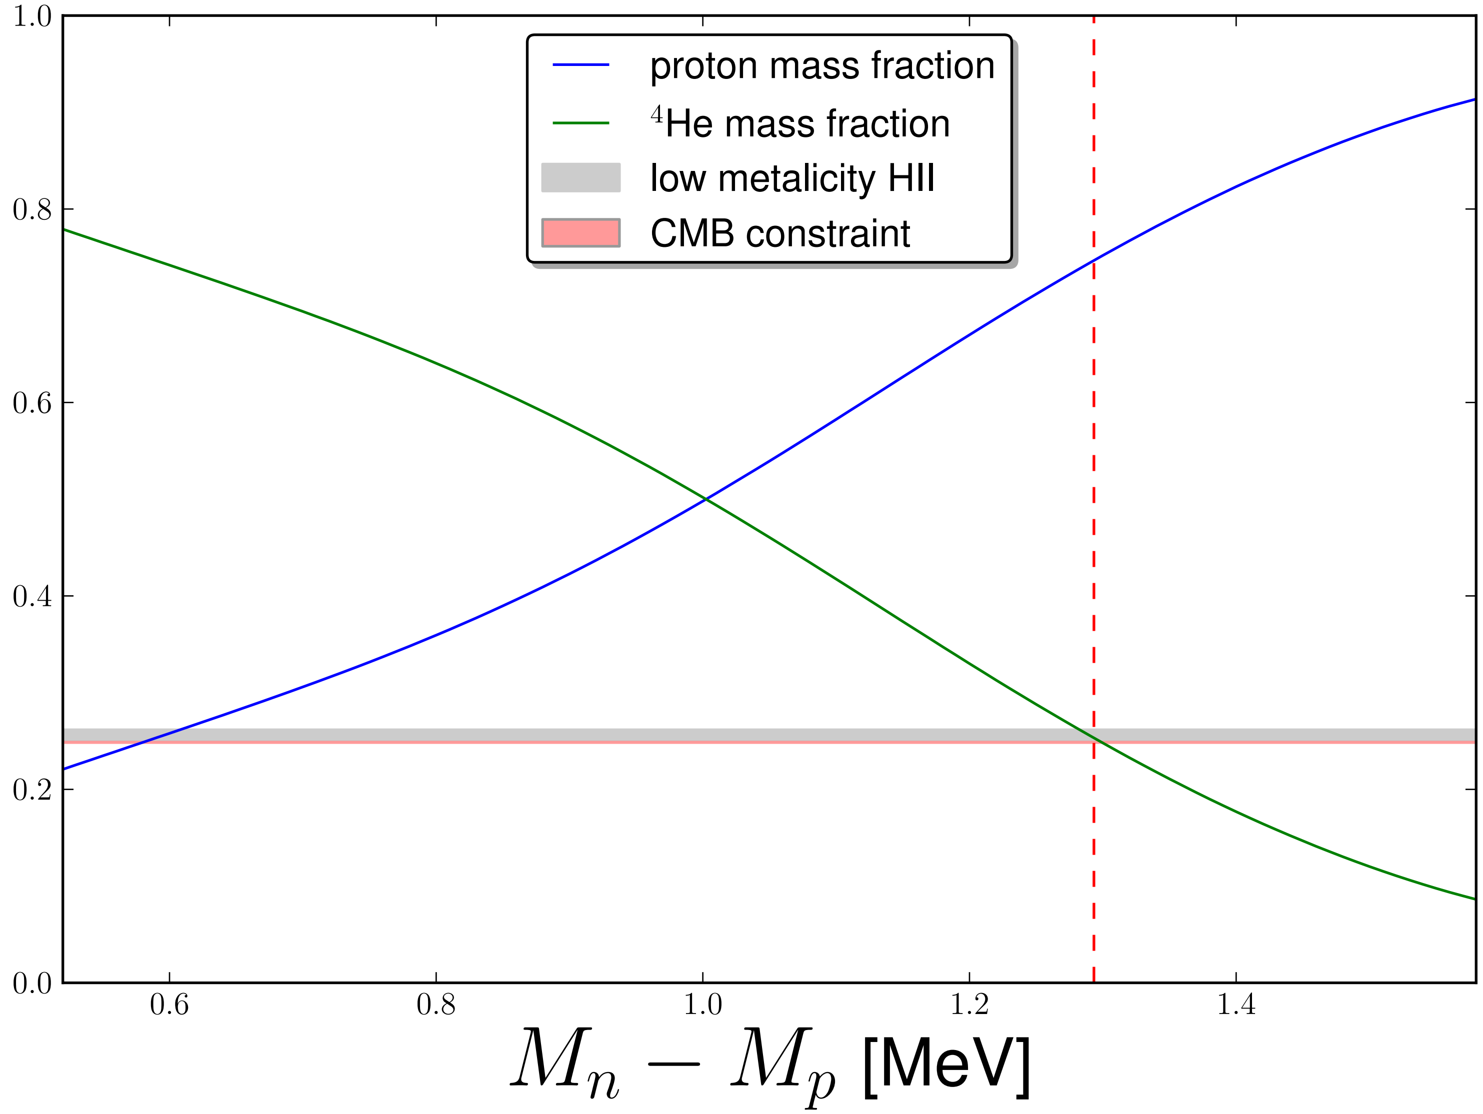

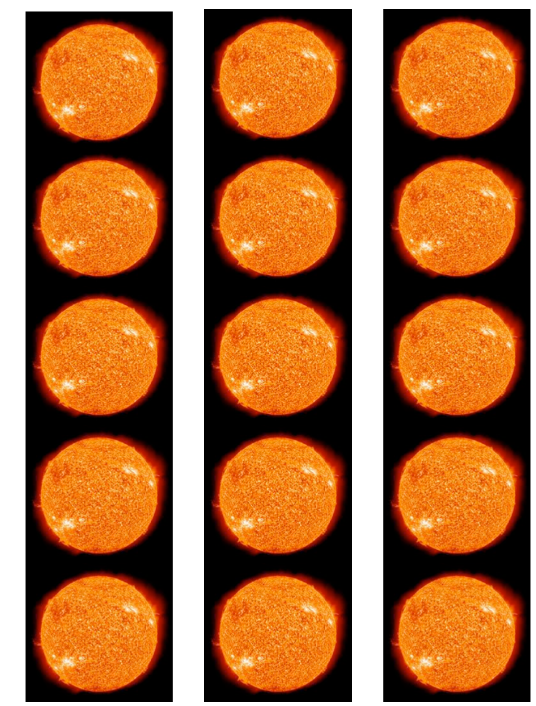

Too many suns?

What is QCD (in slightly more detail)?

What is QCD (in slightly more detail)?

Neutrons and Protons are composed of confined quarks and gluons

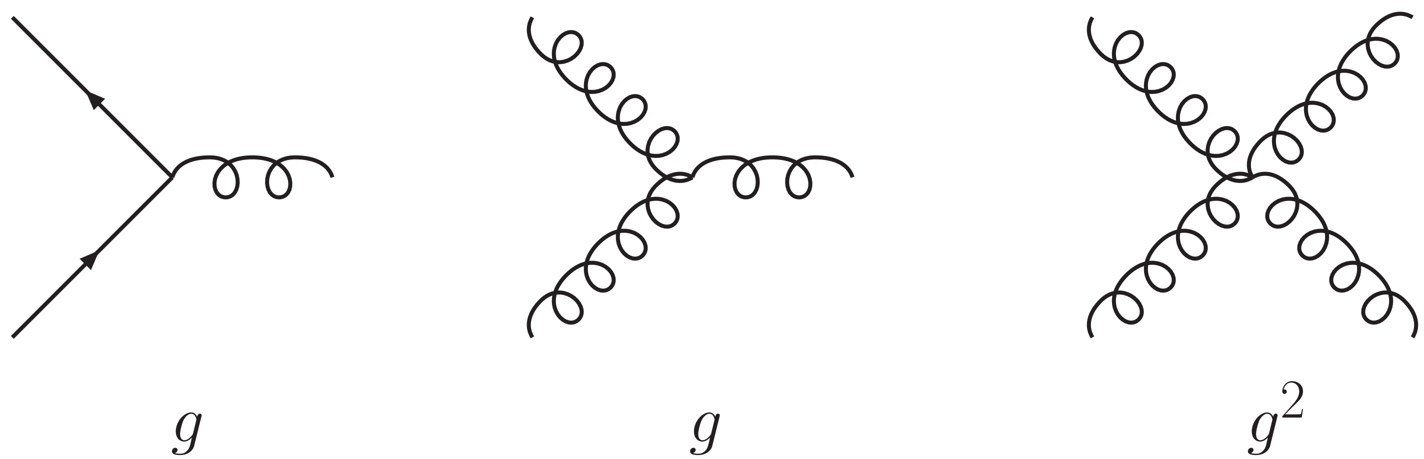

quark-gluon interactions

3- and 4-gluon interactions

like electromagnetism

responsible for non-perturbative nature

What is QCD (in slightly more detail)?

low energy → high energy

$g \sim \mathcal{O}(1)$

Strong Coupling

$g \to 0$

Asymptotic Freedom

$$\alpha_s(Q) = \frac{g^2(Q)}{4\pi}$$

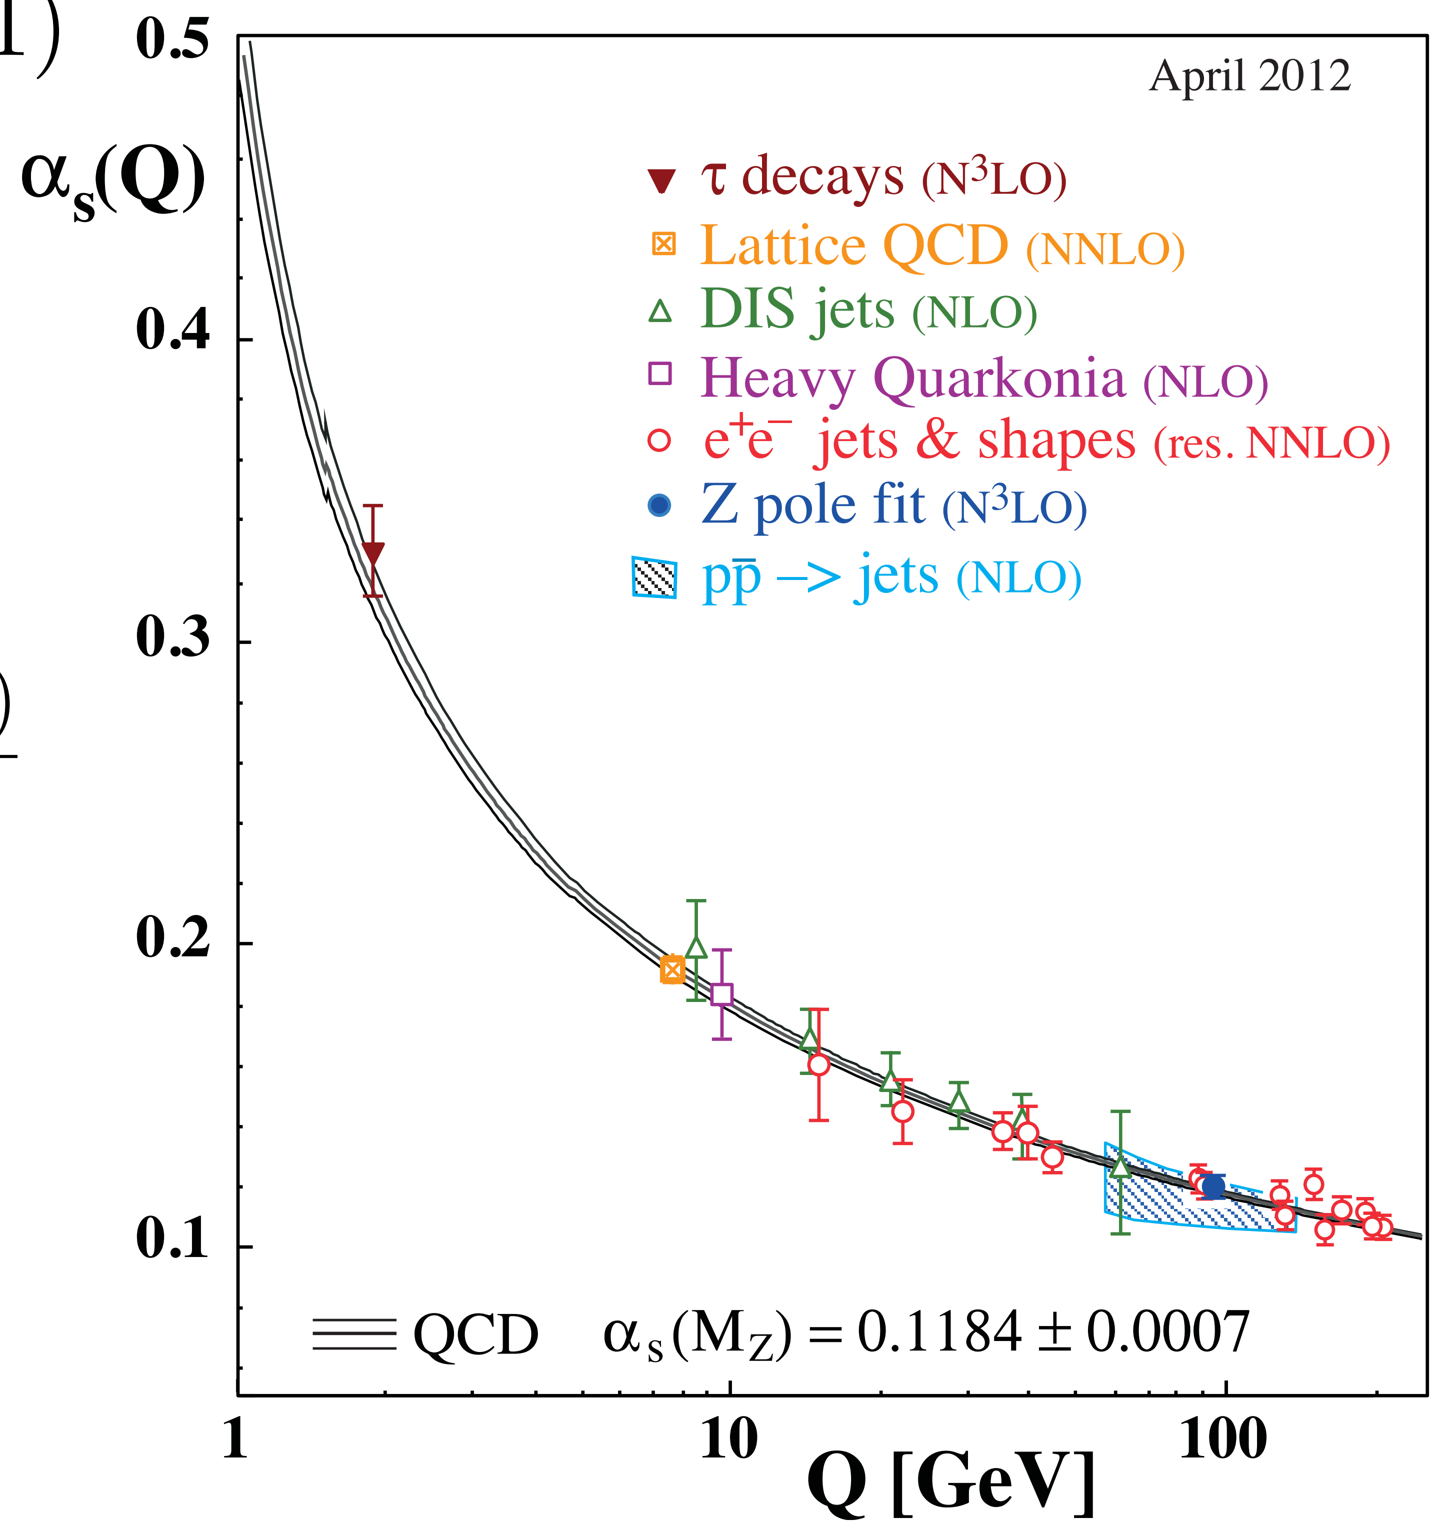

What is QCD (in slightly more detail)?

$g \sim \mathcal{O}(1)$

Strong Coupling

$g \to 0$

Asymptotic Freedom

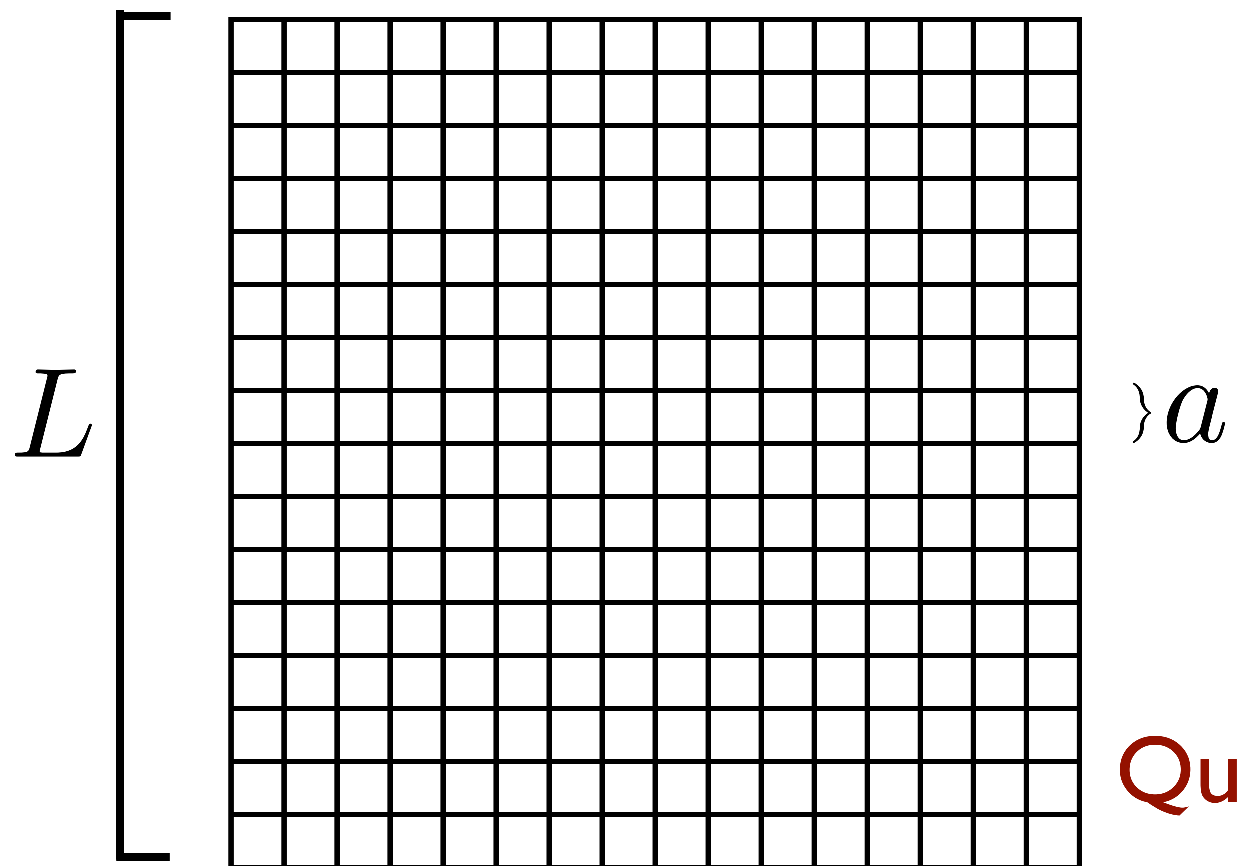

universe

$\sim L^4$

$$\Delta x \Delta p \geq \frac{\hbar}{2}$$

Quantum Mechanics:
uncertainty principle
small distance
=
large energy

take discretization scale to zero $a \to 0$
we know the *exact* theory to put
on the computer

What is QCD (in slightly more detail)?

$g \sim \mathcal{O}(1)$

Strong Coupling

$g \to 0$

Asymptotic Freedom

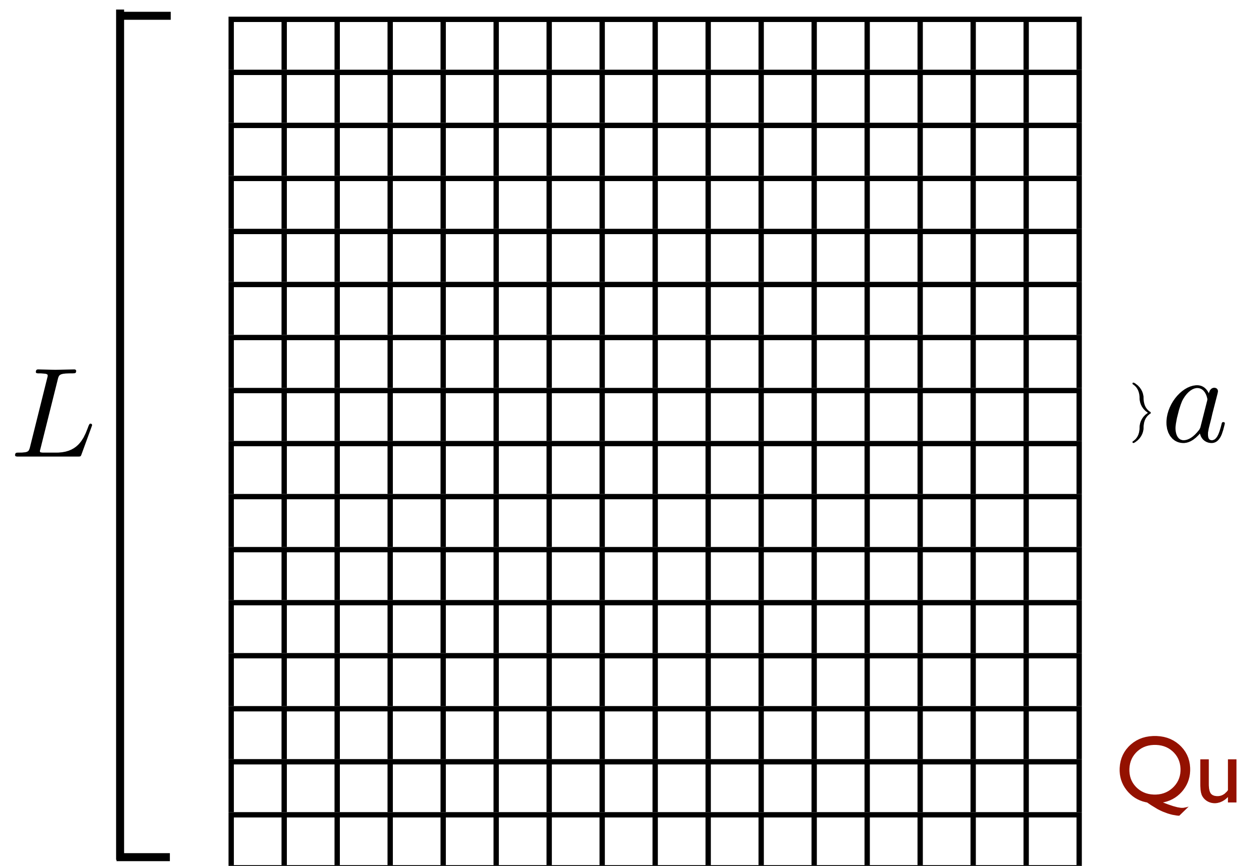

universe

$\sim L^4$

$$\Delta x \Delta p \geq \frac{\hbar}{2}$$

Quantum Mechanics:

uncertainty principle

small distance

=

large energy

take discretization scale to zero $a \to 0$

we know the *exact* theory to put

on the computer

but $t_{comp} \sim \frac{1}{a^6}$

What is QCD (in slightly more detail)?

$g \sim \mathcal{O}(1)$

Strong Coupling

$g \to 0$

Asymptotic Freedom

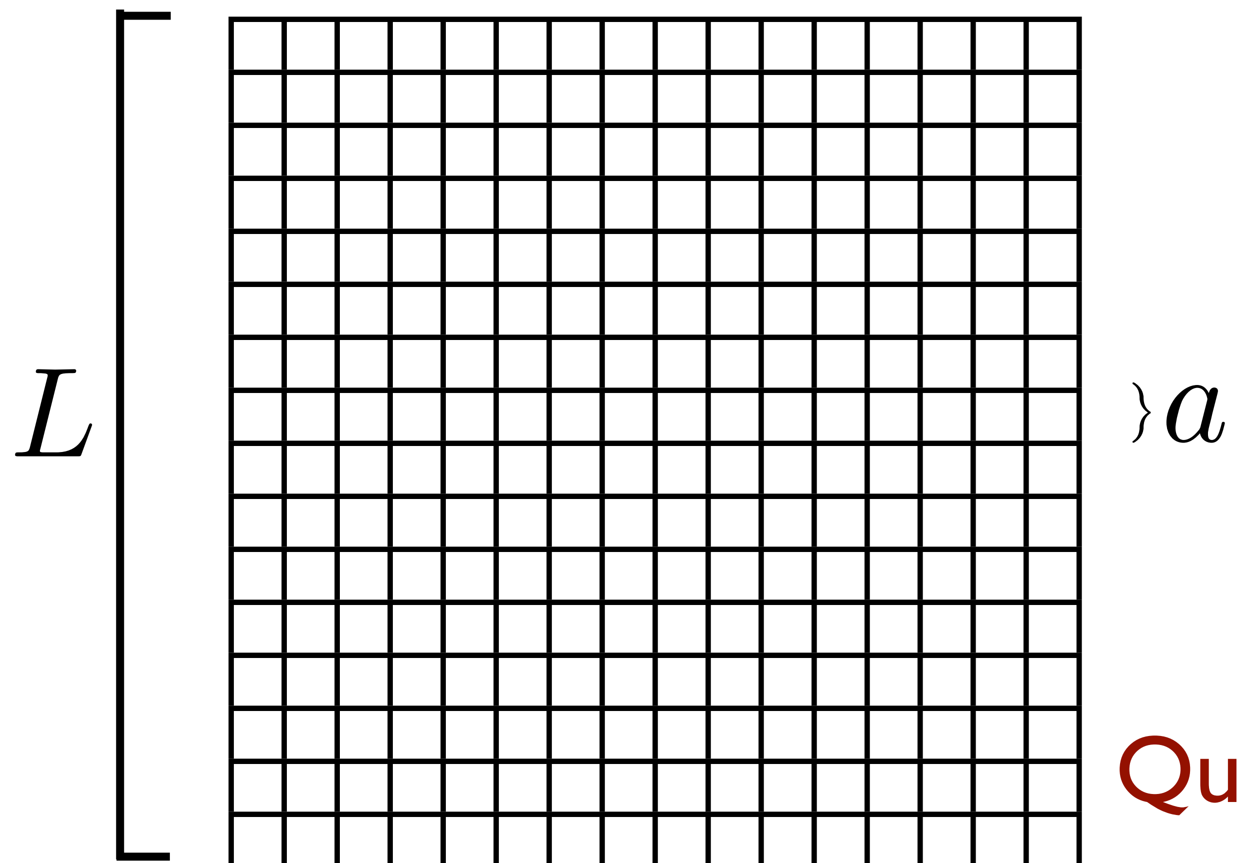

universe

$\sim L^4$

$$\Delta x \Delta p \geq \frac{\hbar}{2}$$

Quantum Mechanics:

uncertainty principle

small distance

=

large energy

take discretization scale to zero $a \to 0$

we know the *exact* theory to put

on the computer

We need BIG computers

What is QCD (in slightly more detail)?

$g \sim \mathcal{O}(1)$

Strong Coupling

$g \to 0$

Asymptotic Freedom

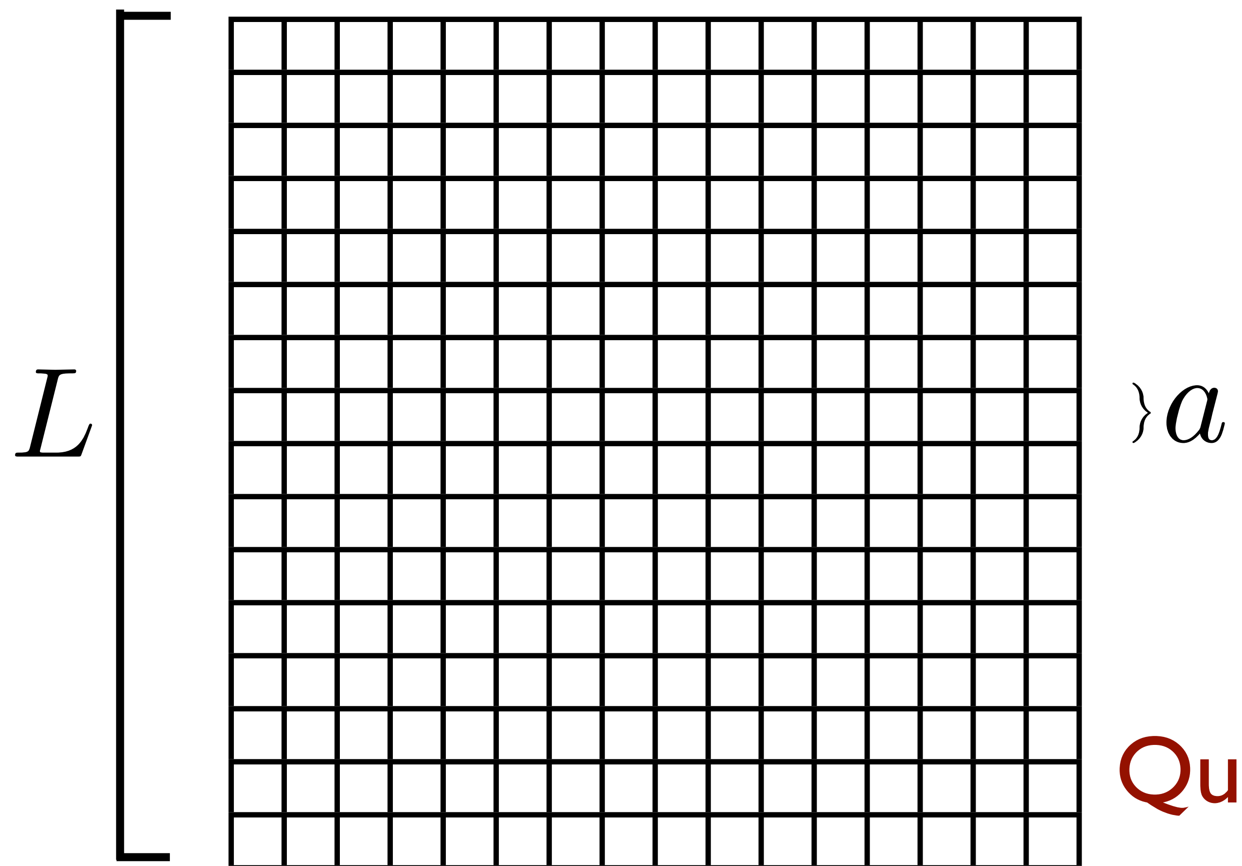

universe

$\sim L^4$

$$\Delta x \Delta p \geq \frac{\hbar}{2}$$

Quantum Mechanics:
uncertainty principle
small distance
=
large energy

state of the art today: $L = 64 - 128$

lattice QCD calculations will
really flourish in the exa-scale era

Generating Lattice Gauge Configurations on Sequoia and Vulcan (BG/Q) with further calculations on Edge (GPU)

CalLat Collaboration

Abstract: The BG/Q machines, Sequoia and its successor Vulcan, are producing the next generation of cold QCD, isotropic, clover-improved Wilson lattice gauge configurations at roughly 400 MeV pion-mass, with dimensions of 48^3 spatial x 96 temporal lattice points and 64^3 spatial x 96 temporal lattice points, corresponding to physical spatial size of 6.7 and 9.0 fm (10^{-15} m), respectively. Additionally, the GPU enabled machines are being heavily utilized for matrix-inversions necessary for calculations of physical correlation functions. In this poster we show the current status of these calculations, and describe efforts to improve certain aspects of them.

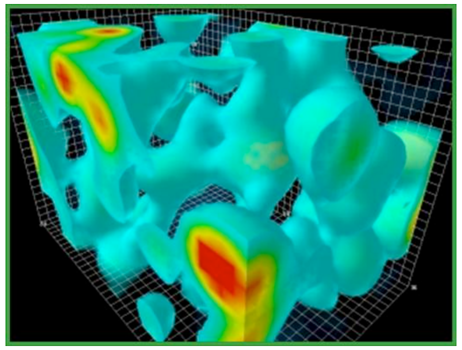

BG/Q

Acceptance Rate

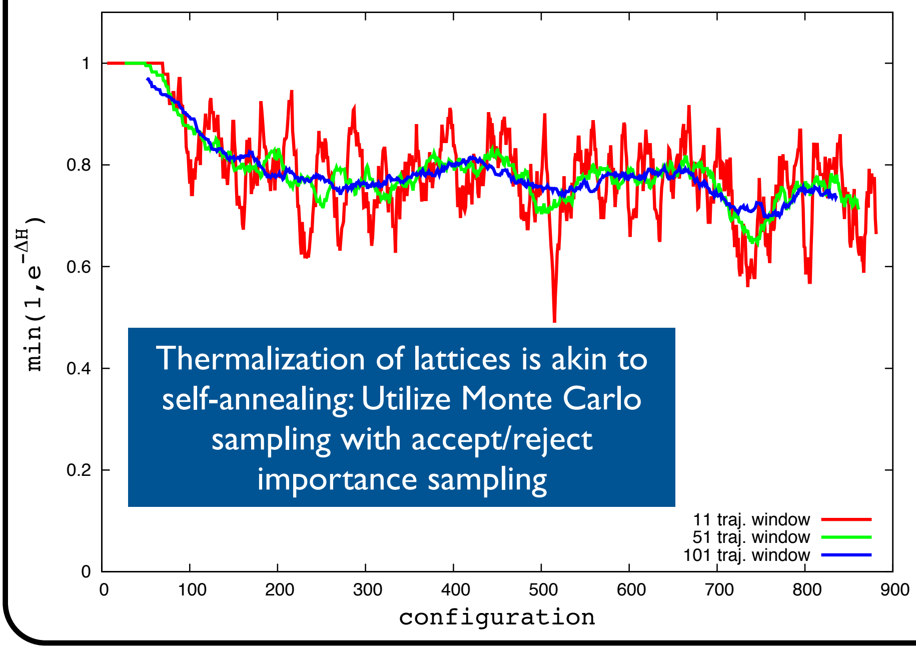

Weak Scaling on BG/Q

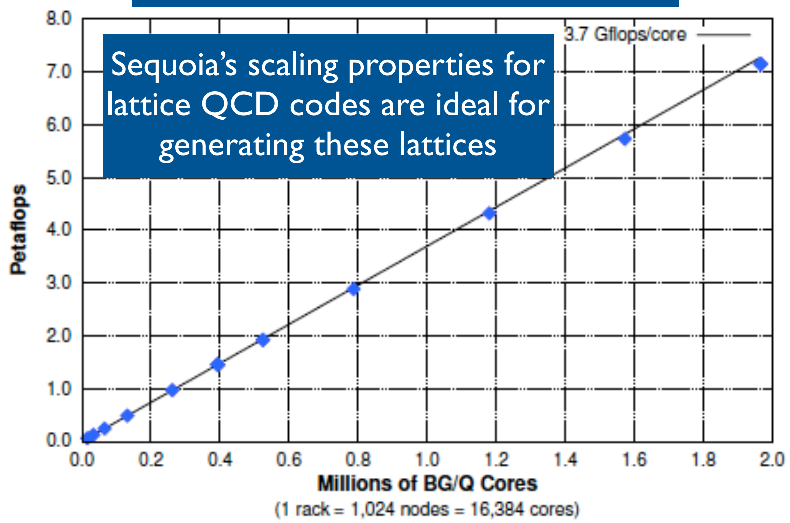

Plaquette Thermalization

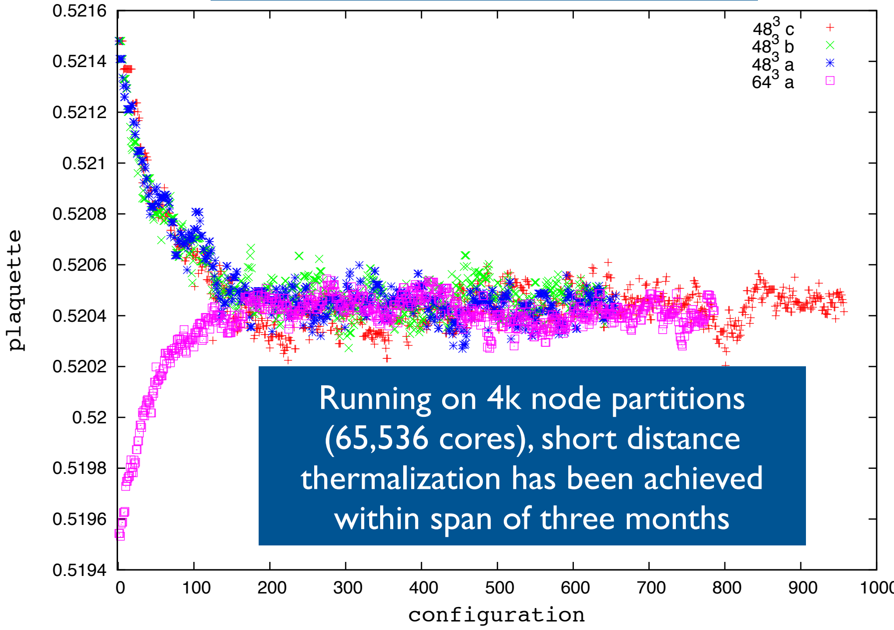

GPU

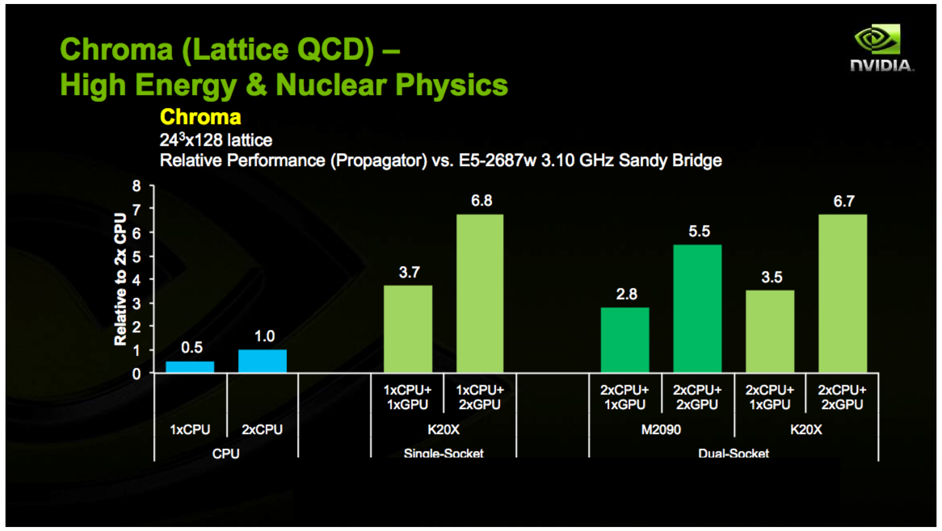

Kepler Wilson-Dslash Performance

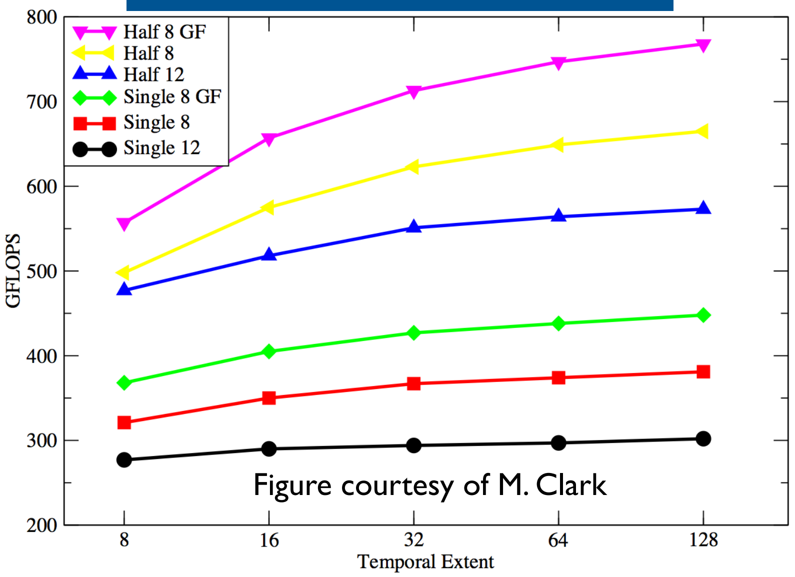

Next steps:
Calculation of the physical correlation functions require complex tensor contractions, which naively scale factorially with the number of "quark" lines. We are investigating whether the multi-threaded environments on both gpu and BG/Q architectures can be utilized to speed up these contractions.

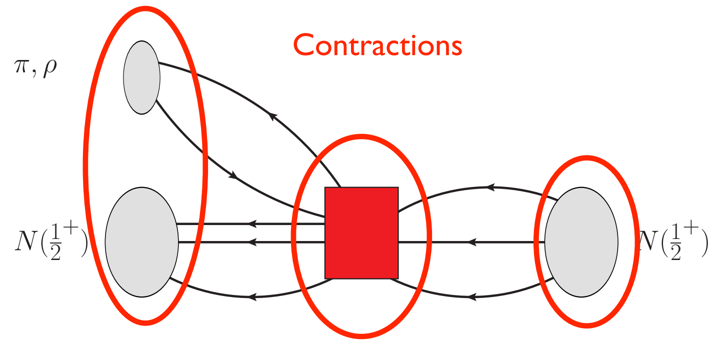

Our goal is to solve the theory of QCD

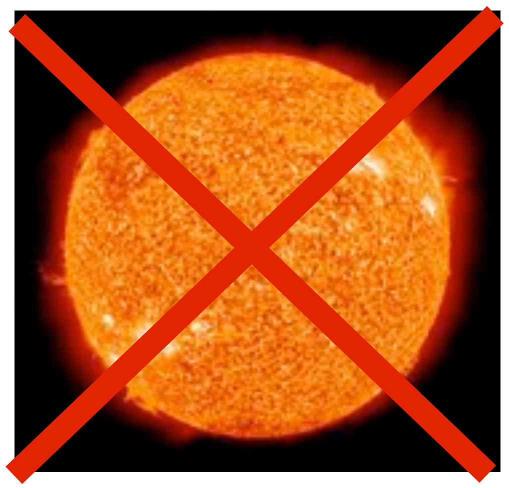

No Sun!

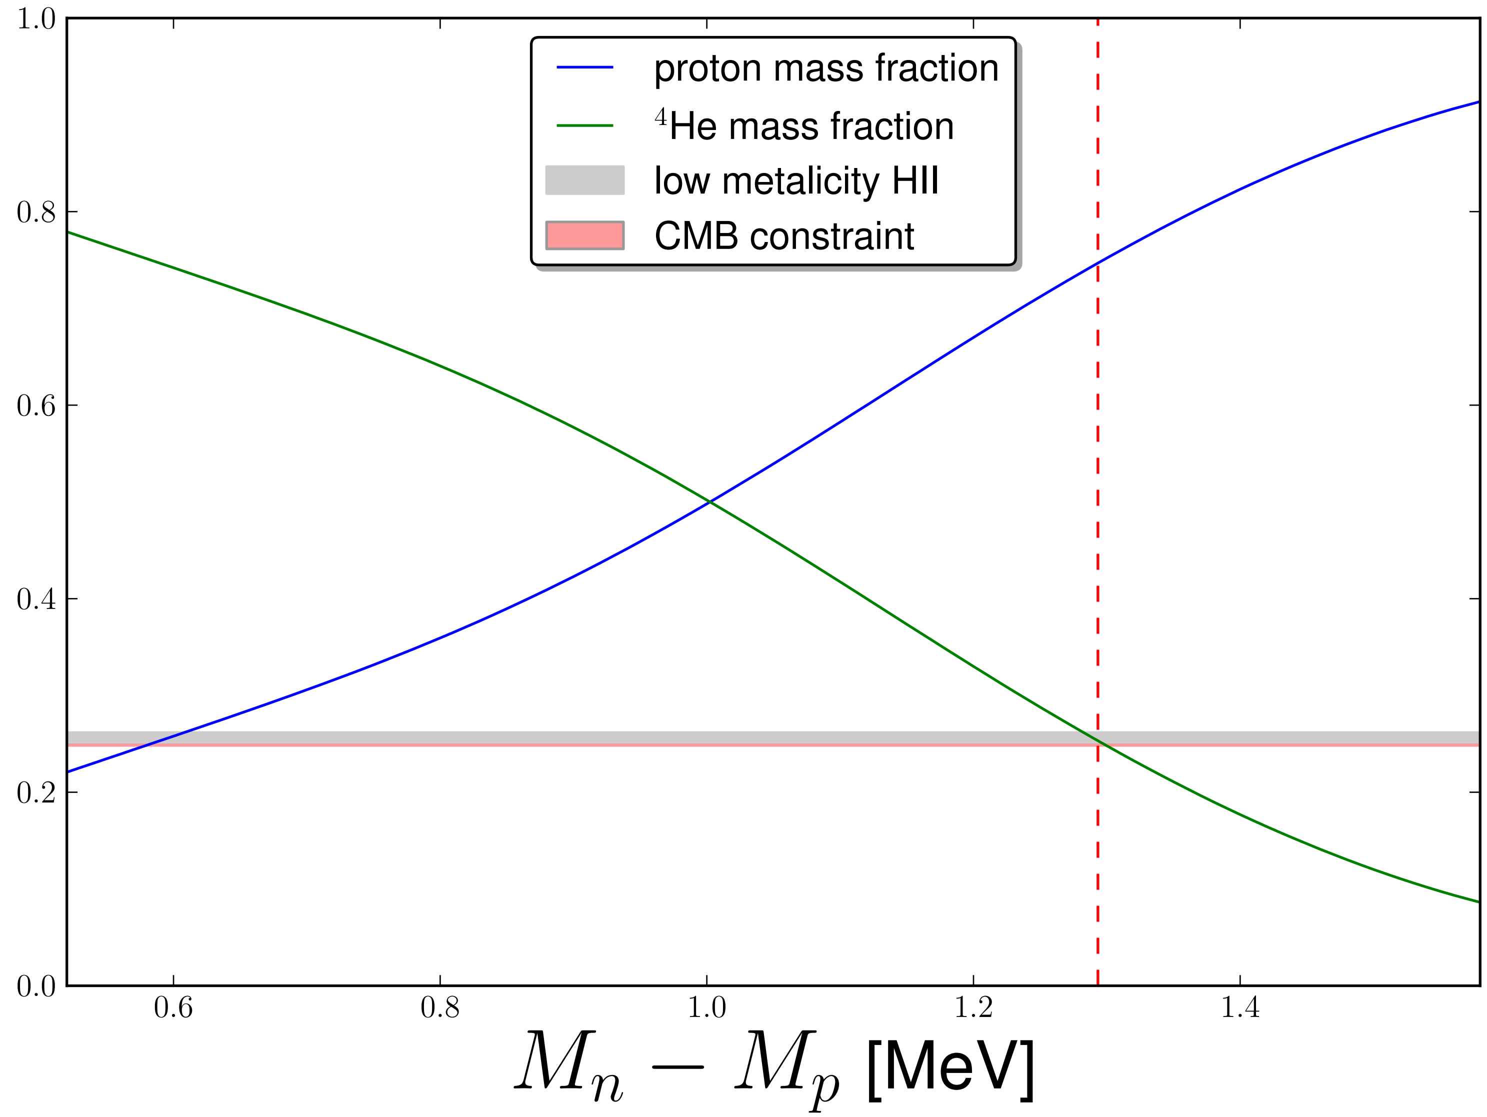

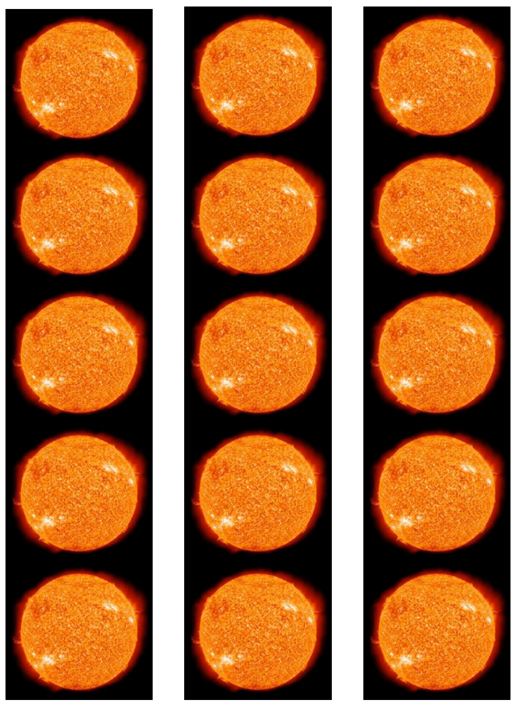

Too many suns?

Our goal is to solve the theory of QCD

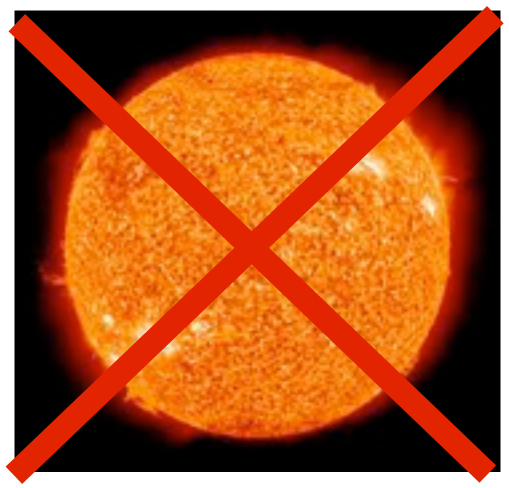

No Sun!

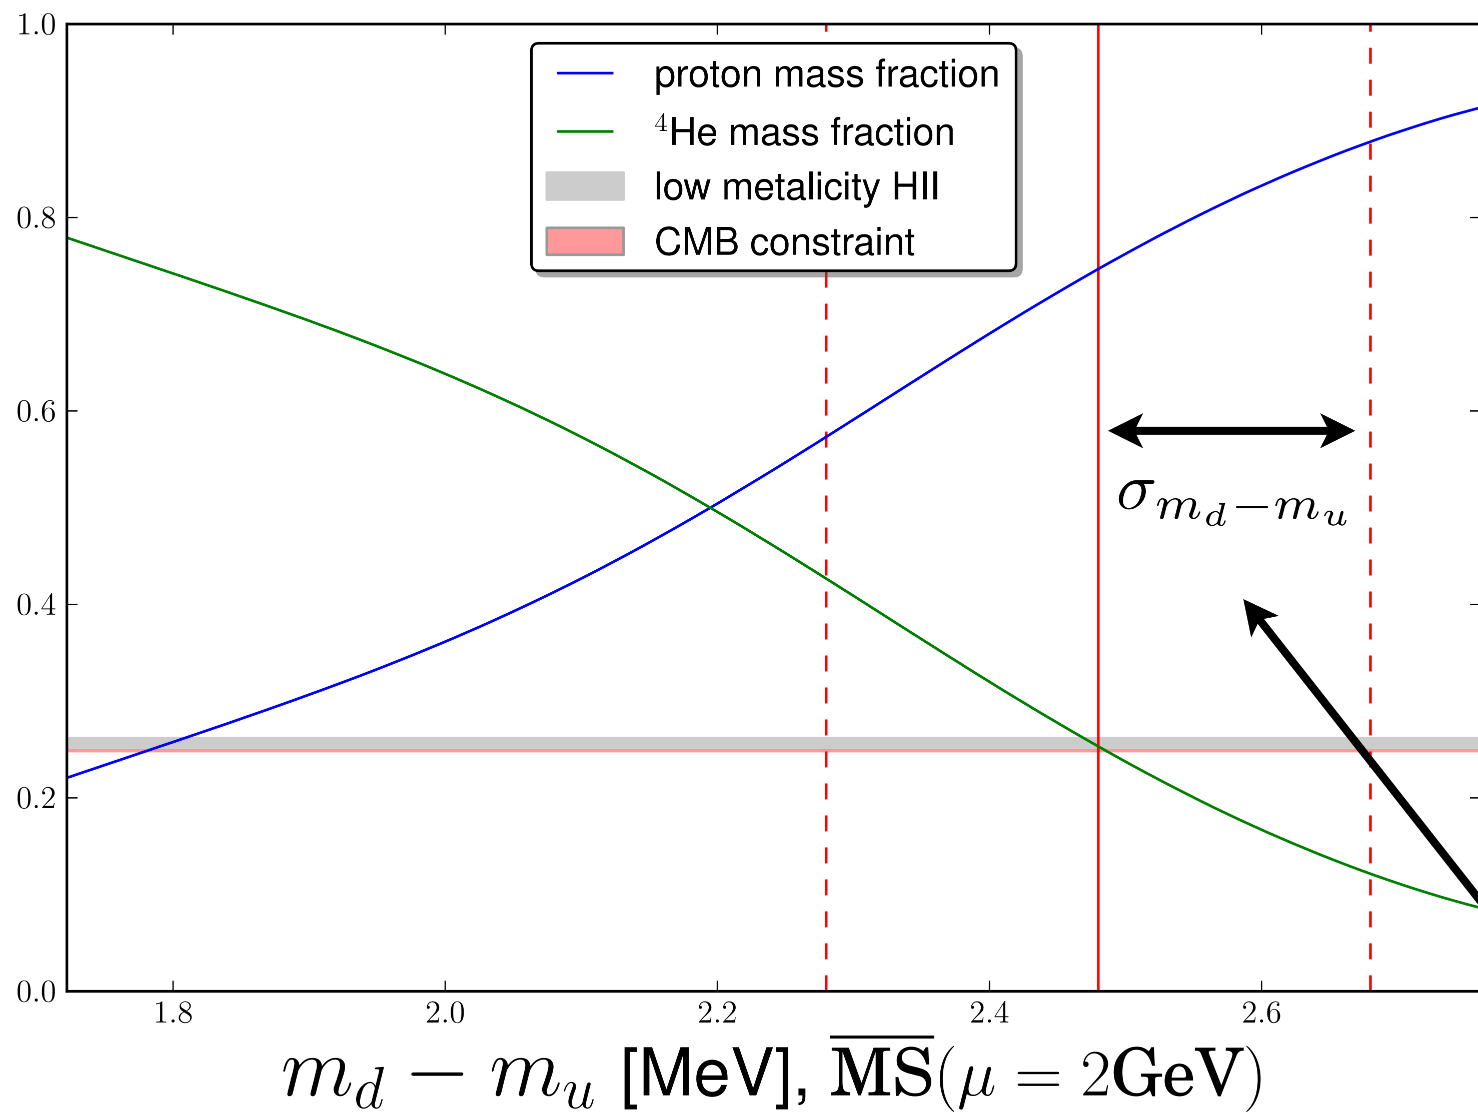

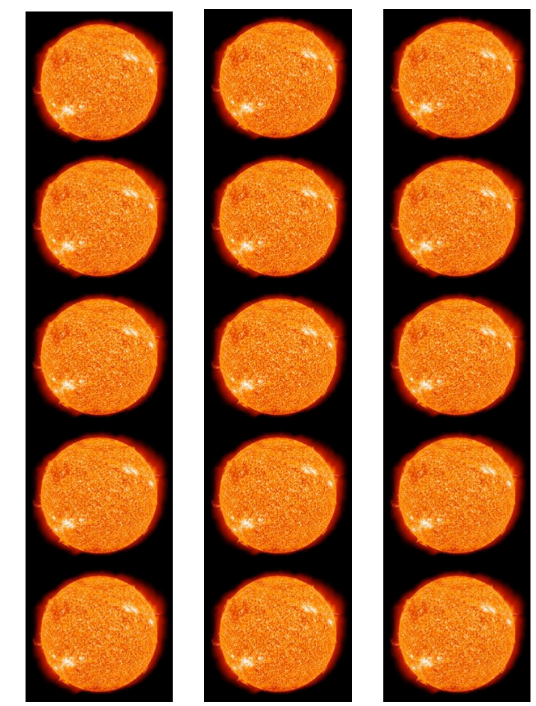

Too many suns?

Lattice QCD

Our goal is to solve the theory of QCD

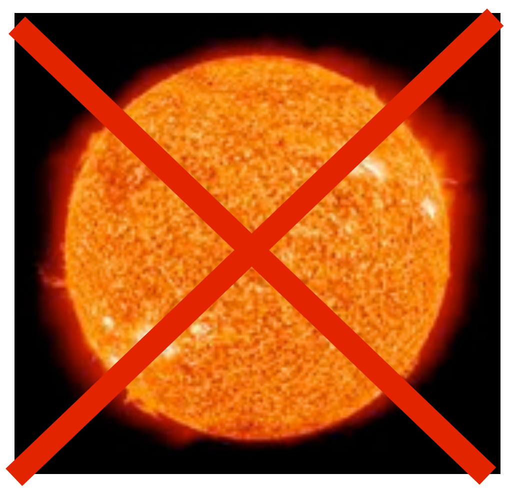

No Sun!

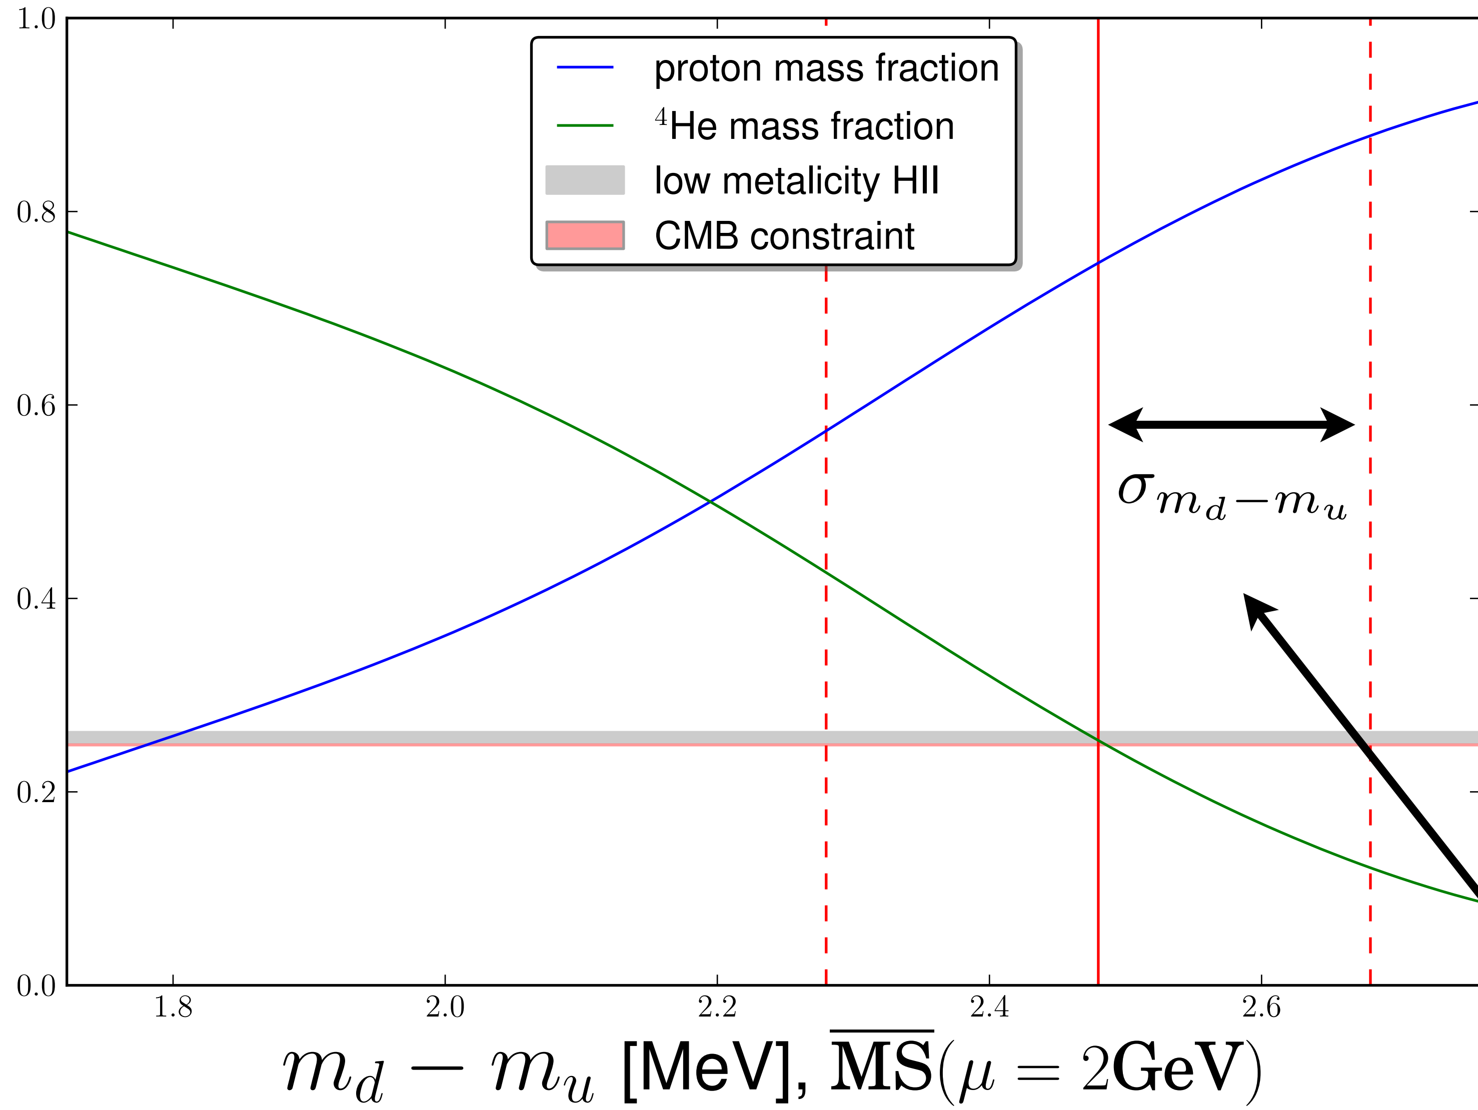

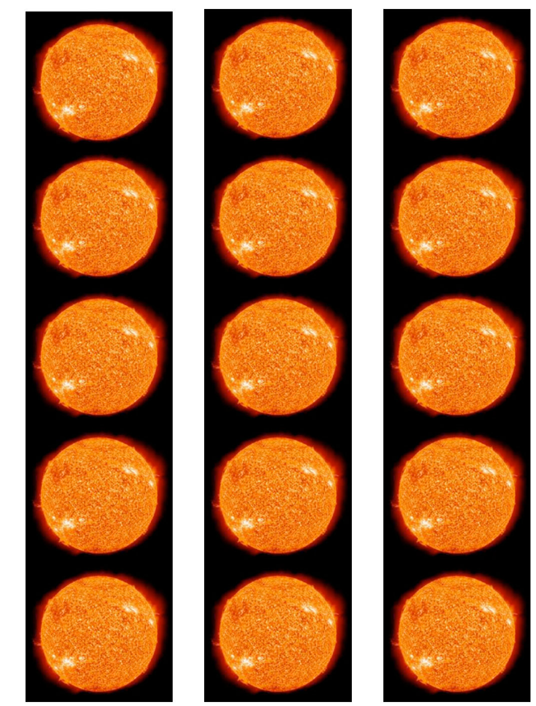

Too many suns?

Lattice QCD

We are now able to quantitatively *connect the quarks with the cosmos*

Hadronic Parity Violation and Lattice QCD

CalLat Collaboration

Abstract: The strangeness-conserving neutral weak current has not been isolated experimentally, but is the subject of a major experiment now underway on the cold neutron beamline of the SNS. This interaction dominates the parity non-conserving (PNC) long-range weak interaction between nucleons. A CalLat goal is to evaluate the isovector πNN coupling corresponding to this interaction as well as the isotensor ρNN coupling: without the latter result, the most significant experiment in the field, the asymmetry in p+p scattering, loses impact. We describe our global analysis motivating this work, a recent Lattice QCD calculation of the PNC πNN coupling, and plans to improve this result and undertake similar LQCD calculations for the isotensor ρNN coupling.

Hadronic PNC

The elementary couplings of weak currents to quarks are determined by the Standard Model (SM). The strangeness-conserving hadronic weak interaction is only accessible through nucleon-nucleon and nuclear experiments, with PNC the filter that allows one to identify the weak interaction despite the dominance of the strong interaction (SI). However the SI acting in the nucleon will alter couplings. A key goal of the field has been to isolate the neutral current contribution and to evaluate associated SI renormalization effects.

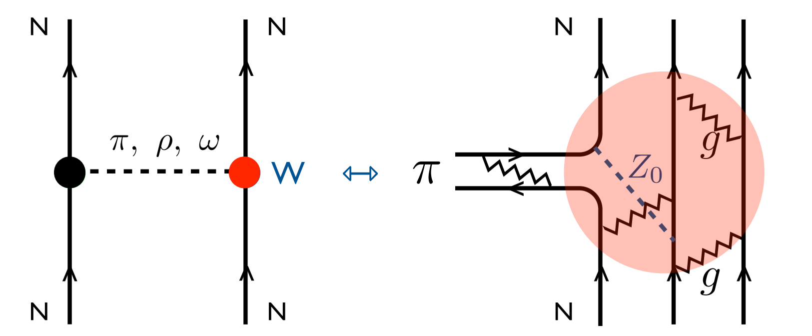

The low-energy PNC NN interaction can be viewed as arising from meson exchange, with one strong and one weak vertex, with the latter modified by SI effects. A controlled calculation of the PNC meson-nucleon coupling in lattice QCD would provide the basic information needed to calculate PNC observables for more complex nuclear systems in perturbation theory.

This program is timely because of the ongoing experimental effort at the cold neutron beamline at the SNS to measure the PNC asymmetry in n+p → D+γ, which is expected to be at the level of one part in 10^8. This observable isolates the isovector pion-exchange contribution to NN PNC, a channel that should be dominated by the so-far unobserved hadronic neutral current. There are strong indications from PNC measurements in ^{18}F that the PNC neutral current effects are much smaller than expected. Initial lattice QCD results confirm this result. Thus the SNS experiment with need to reach its design sensitivity to see a signal.

Global PNC Analysis

There are four sets of interpretable measurements constraining PNC: the longitudinal asymmetries for p+p and p+α, the photon circular polarization in ^{18}F, and the photon asymmetry in ^{19}F. For some time it has been believed that these measurements lead to an inconsistent description of PNC. The situation greatly improved in a recent global analysis in which identified and corrected inconsistencies in previous analyses of the p+p measurements.

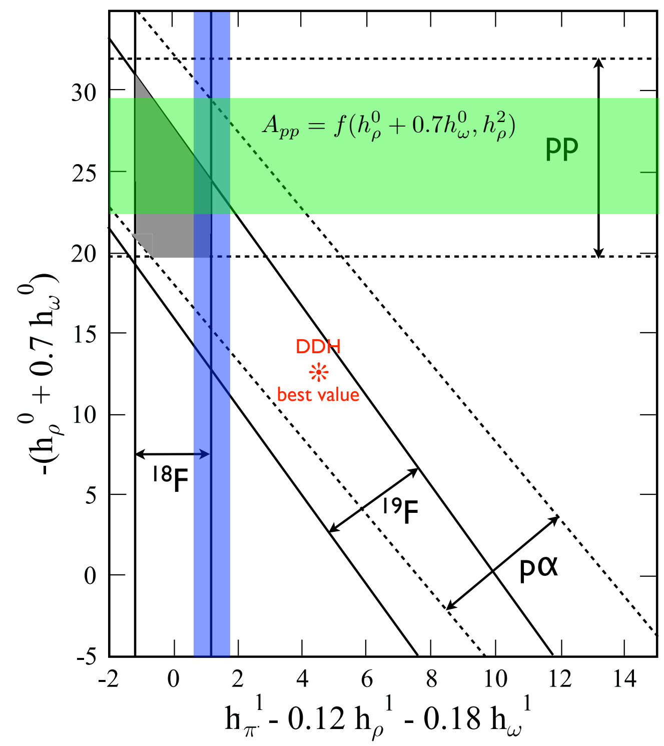

The analysis shows that the isovector pion coupling h_π is suppressed by at least a factor of three and the isoscalar ρ/ω coupling is enhanced by a factor of two relative to SM estimates (DDH best values). The factor-of-six difference in the ratio of isospin I=0 to I=1 couplings is superficially reminiscent of the ΔI=1/2 rule describing the analogous I=1/2 to I=3/2 ratio in strangeness-changing weak hadronic decays.

First LQCD Efforts on PNC

The blue band on the global analysis graph shows the first LQCD calculation of h_π^1. The contractions performed by Wasem required approximately 6 months of running on LLNL's Edge GPU cluster. This initial calculation was performed on a single ensemble of n_f=2+1 dynamical anisotropic clover gauge configurations for a lattice of dimension 2.5 fm, spacing 0.123 fm, and pion mass 389 MeV, and includes only connected diagrams. The error band is an estimate of the systematic uncertainties associated with these lattice parameters. The result is consistent with the ^{18}F bound.

PNC Lattice QCD: Future Work

The improvement of the h_π^1 calculation and the extension of the work to the isotensor PNC ρNN coupling h_ρ^2 is one of the initial goals of the CalLat Collaboration. An LQCD calculation of h_π^1 at nearly physical pion masses with the inclusion of quark-loop contributions appears within reach, given new machines such as LLNL's 5-pflop Vulcan, Titan and Oak Ridge and Edison at NERSC. h_ρ^2 contributes to the precisely measured p+p asymmetry but is otherwise unconstrained by experiment. Consequently the p+p band between the dashed horizontal lines in the figure has been enlarged due to "marginalizing" against h_ρ^2. A lattice calculation of h_ρ^2, which requires no costly quark-loop contributions, could narrow the band to the area indicated in green.

Improving the Lattice Calculation

In order to improve the calculation of h_π^1 and perform the calculation of h_ρ^2, new techniques must be implemented. The interpolating fields can be projected onto definite quantum numbers, such as total angular momentum, J, and parity, P, but they couple to all the eigenstates of QCD with those definite quantum numbers.

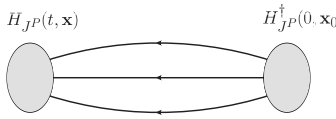

$$C_{J^P}(t) = \sum_{\mathbf{x}} \langle 0|H_{J^P}(t,\mathbf{x})H^\dagger_{J^P}(0,\mathbf{x}_0)|0\rangle$$
$$= \sum_{\mathbf{x}} \langle 0|H_{J^P}(t,\mathbf{x}) \sum_n |n\rangle\langle n|H^\dagger_{J^P}(0,\mathbf{x}_0)|0\rangle$$
$$= \sum_n A_{n,J^P}\, e^{-E_n^{J,P} t}$$

The first calculation (Wasem) used a simple interpolating field for the Nπ final state:

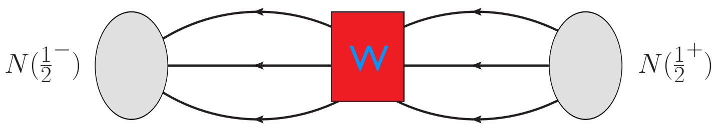

In order to improve the coupling to the matrix element of interest, we need to construct interpolating fields which more closely resemble the state of interest. A large basis of such operators will allow us to "filter" the excited states out of the calculation. To accomplish this goal, we will implement several new features into the calculation:

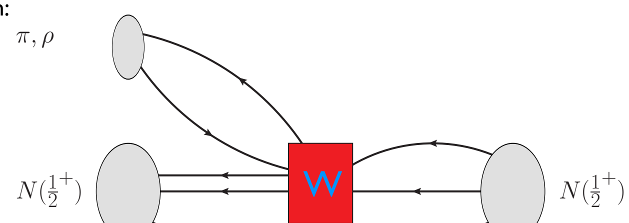

1) a large basis of interpolating fields ("wave-functions"), used to perform a variational analysis
2) disconnected quark-diagrams connected to the source or sink operators
3) fast calculation of all the necessary quark level Wick-contractions (tensor-contractions)

Conclusions: Implementation of these new methods will be more computationally demanding, but will allow for significantly improved calculations of Parity-Non-Conserving processes in low-energy nucleon interactions. The first step will allow us to compute the Isospin=2 coupling, which we plan to accomplish within a year. Steps 2 and 3 are necessary for the Isospin=1 coupling and more complex matrix elements.

Two projects in the beginning stages

Quarks to Nuclei: tensor contractions

Typical LQCD Workflow

(David Richards)

Generate the configurations
- Leadership level

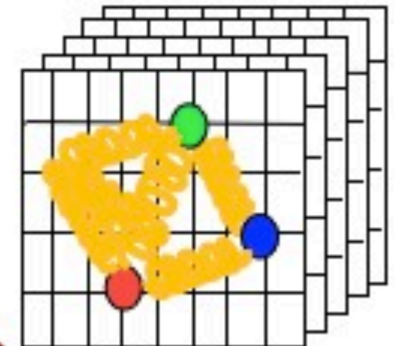

Few big jobs
Few big files

Analyze
- Typically mid-range level

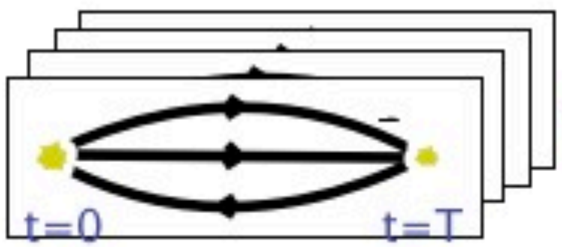

IMPORTANT FOR NP

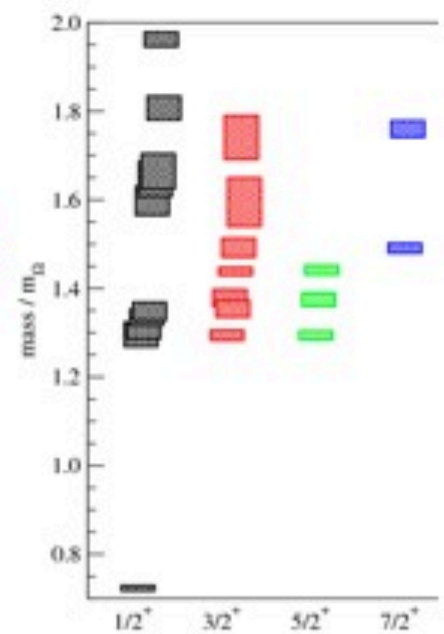

Extract
- Extract information from measured observables

Many small jobs Many big files

I/O movement

Generate the Gauge Fields

$$\mathcal{Z}_{QCD} = \int DU_\mu D\bar{\psi} D\psi \ e^{-\int d^4x \bar{\psi} D_W \psi} e^{-S[U_\mu]}$$

$$= \int DU_\mu \mathrm{Det}(D_W) e^{-S[U_\mu]}$$

$$= \int DU_\mu D\phi^* D\phi \ e^{-\int d^4x \phi^* D_W^{-1} \phi} e^{-S[U_\mu]}$$

This Matrix Inversion is one of the most challenging (costly) aspects of lattice calculations

$\mathrm{Det}(D_W) e^{-S[U_\mu]}$ probability weight used for Monte Carlo estimation of integral

⇒ $\{U_i \equiv U_{\mu,i}; i \in [0, N)\}$

$$\langle O \rangle = \lim_{N \to \infty} \frac{1}{N} \sum_{i=0}^{N-1} O[U_i, D_W^{-1}[U_i]]$$

Given a set of gauge fields - one performs *measurements*

consider a proton

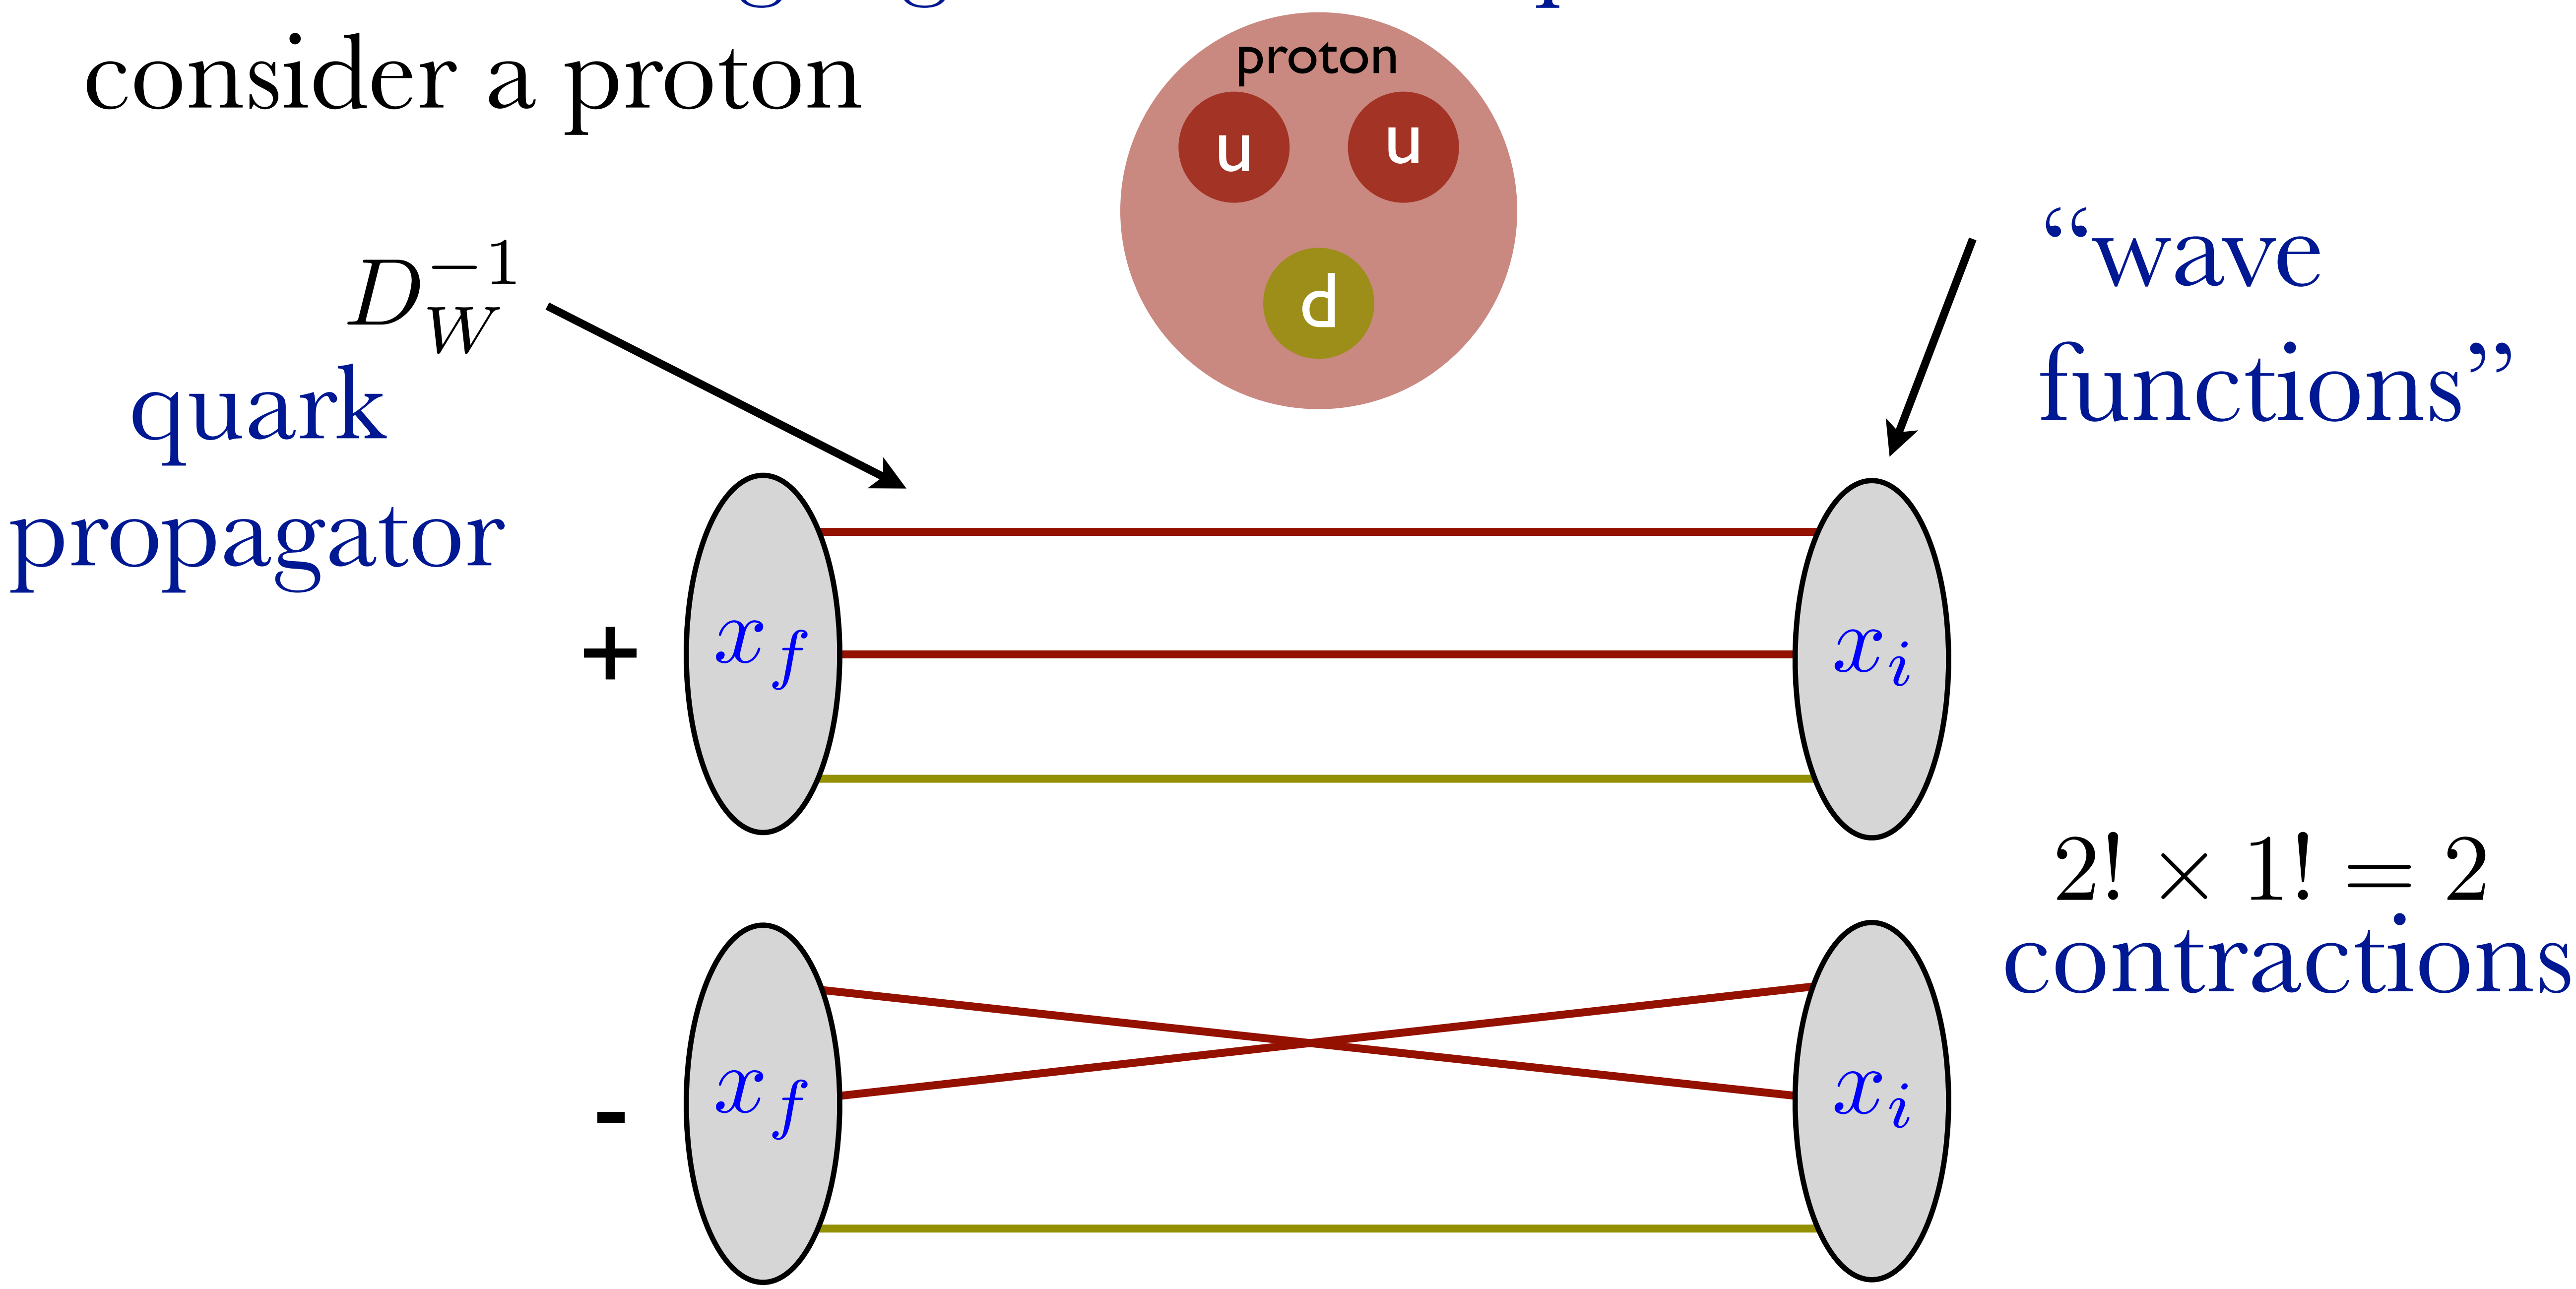

quark-exchange diagrams are source of fermion sign problem

Given a set of gauge fields - one performs *measurements*

consider a deuteron

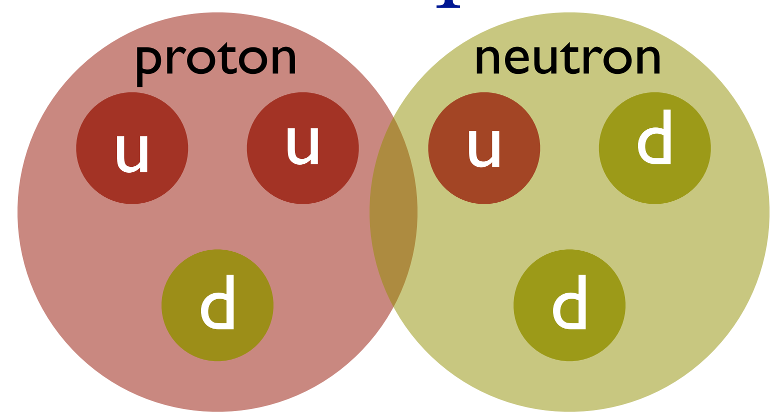

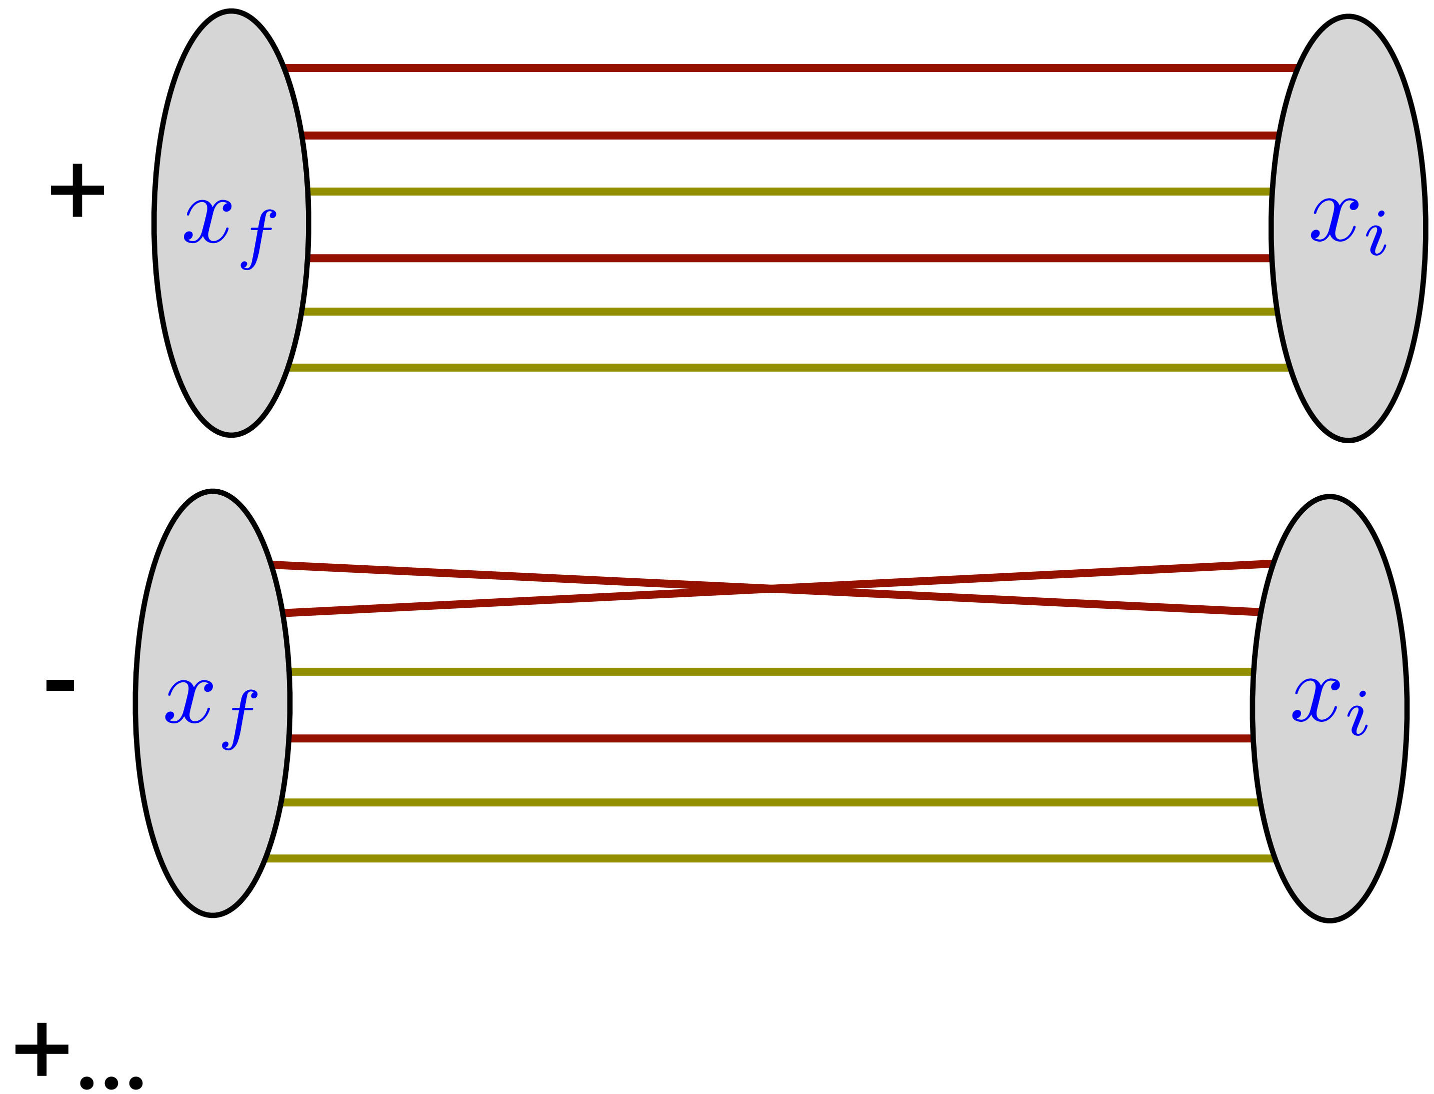

$3! \times 3! = 36$

contractions

Given a set of gauge fields - one performs *measurements*

consider a ^{4}He

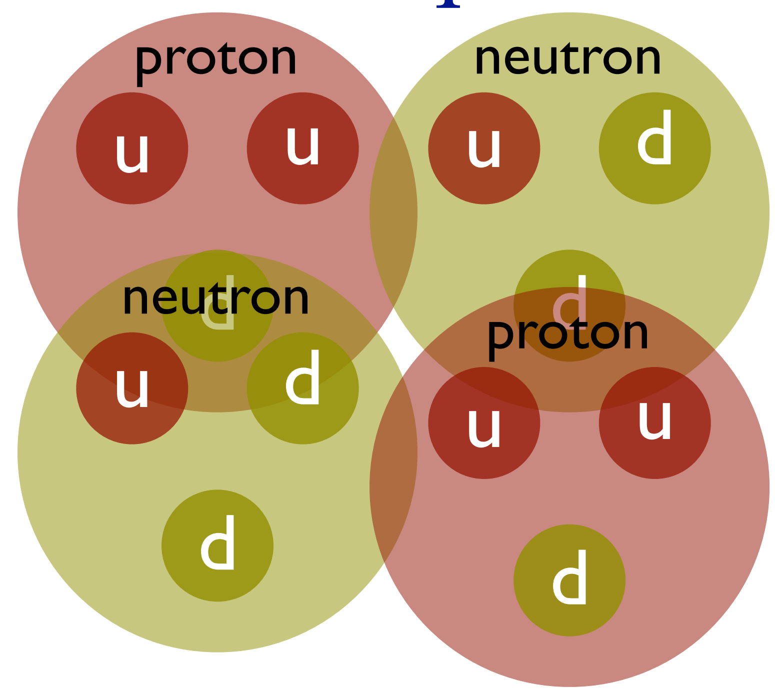

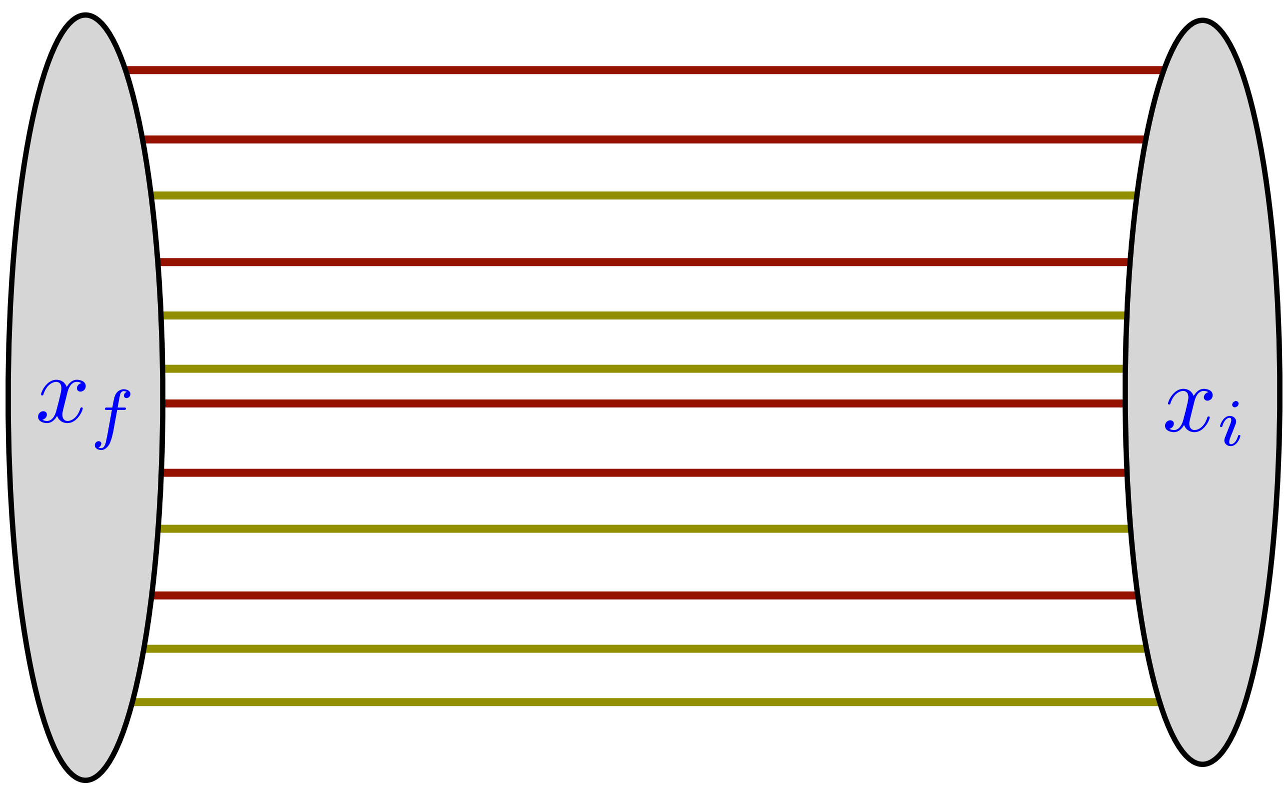

+...

$6! \times 6! = 518400$

contractions

Not entirely fair - lots of symmetry to reduce the number of contractions

the point is these contraction costs are quite significant

Given a set of gauge fields - one performs *measurements*

$$S = D_W^{-1}$$

$$S(x_f, \alpha', a' : x_i, \alpha, a)$$

"spin": 0,1,2,3 "color": 0,1,2

New methods use basis (eg. 3-d laplacian) to approximate the low (important) eigen-modes of this matrix

$$(x_f : x_i) \rightarrow \mathcal{O}(100) \quad \text{eigen-vectors}$$

Constructing "measurements" amounts to performing complex tensor contractions of the quark-propagator indices

Calculation of Matrix Elements relevant to precision tests in nuclei are more complex - an additional set of contractions must be performed

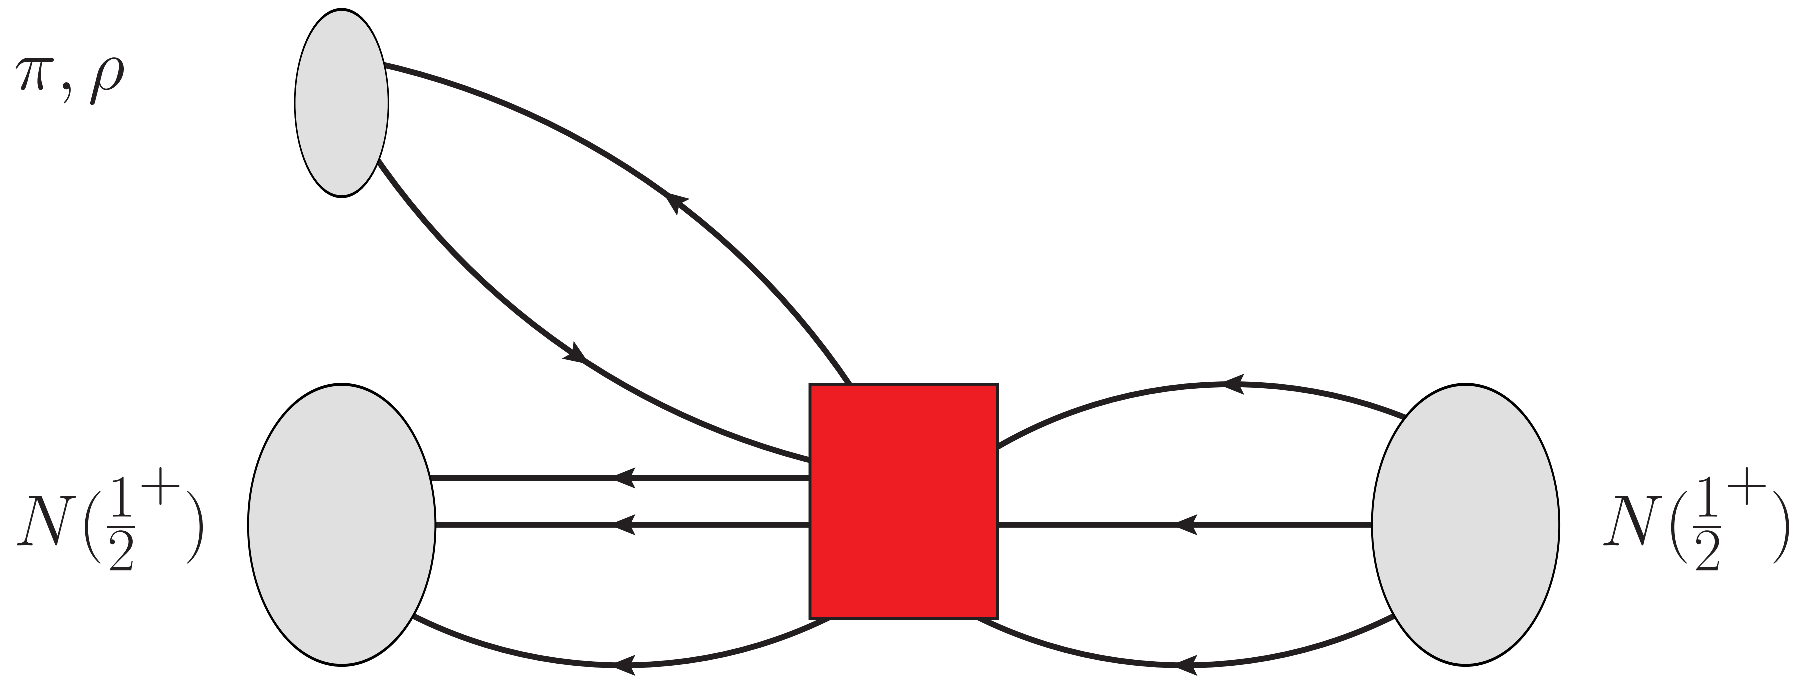

Can we take advantage of the multi-threaded cores on new architectures (GPUs, BG/Q, ...) to improve the performance of these tensor contractions?

Sam Williams and Abhinav Sarje
LBNL

Can we take advantage of the multi-threaded cores on new architectures (GPUs, BG/Q, ...) to improve the performance of these tensor contractions?

- For a single, complicated contraction, can multiple threads improve the performance?
  - needs large loops, not large number of small loops
- For a large set of different contractions, can we develop an automated contraction code which distributes the next contraction to idle thread?
  - we are optimistic this should work

Sam Williams and Abhinav Sarje
LBNL

Parallel I/O: HDF5 for Lattice QCD

Parallel I/O: HDF5 for Lattice QCD

Over the past 1-2 years, I have come to love the HDF5 library. My experience with HDF5 is almost completely through Python (pytables). This part of my talk, is largely free advertising for HDF5, and also an attempt to convince my lattice colleagues we should adopt it, as well as expose any of you who are unfamiliar.

http://www.hdfgroup.org/HDF5/

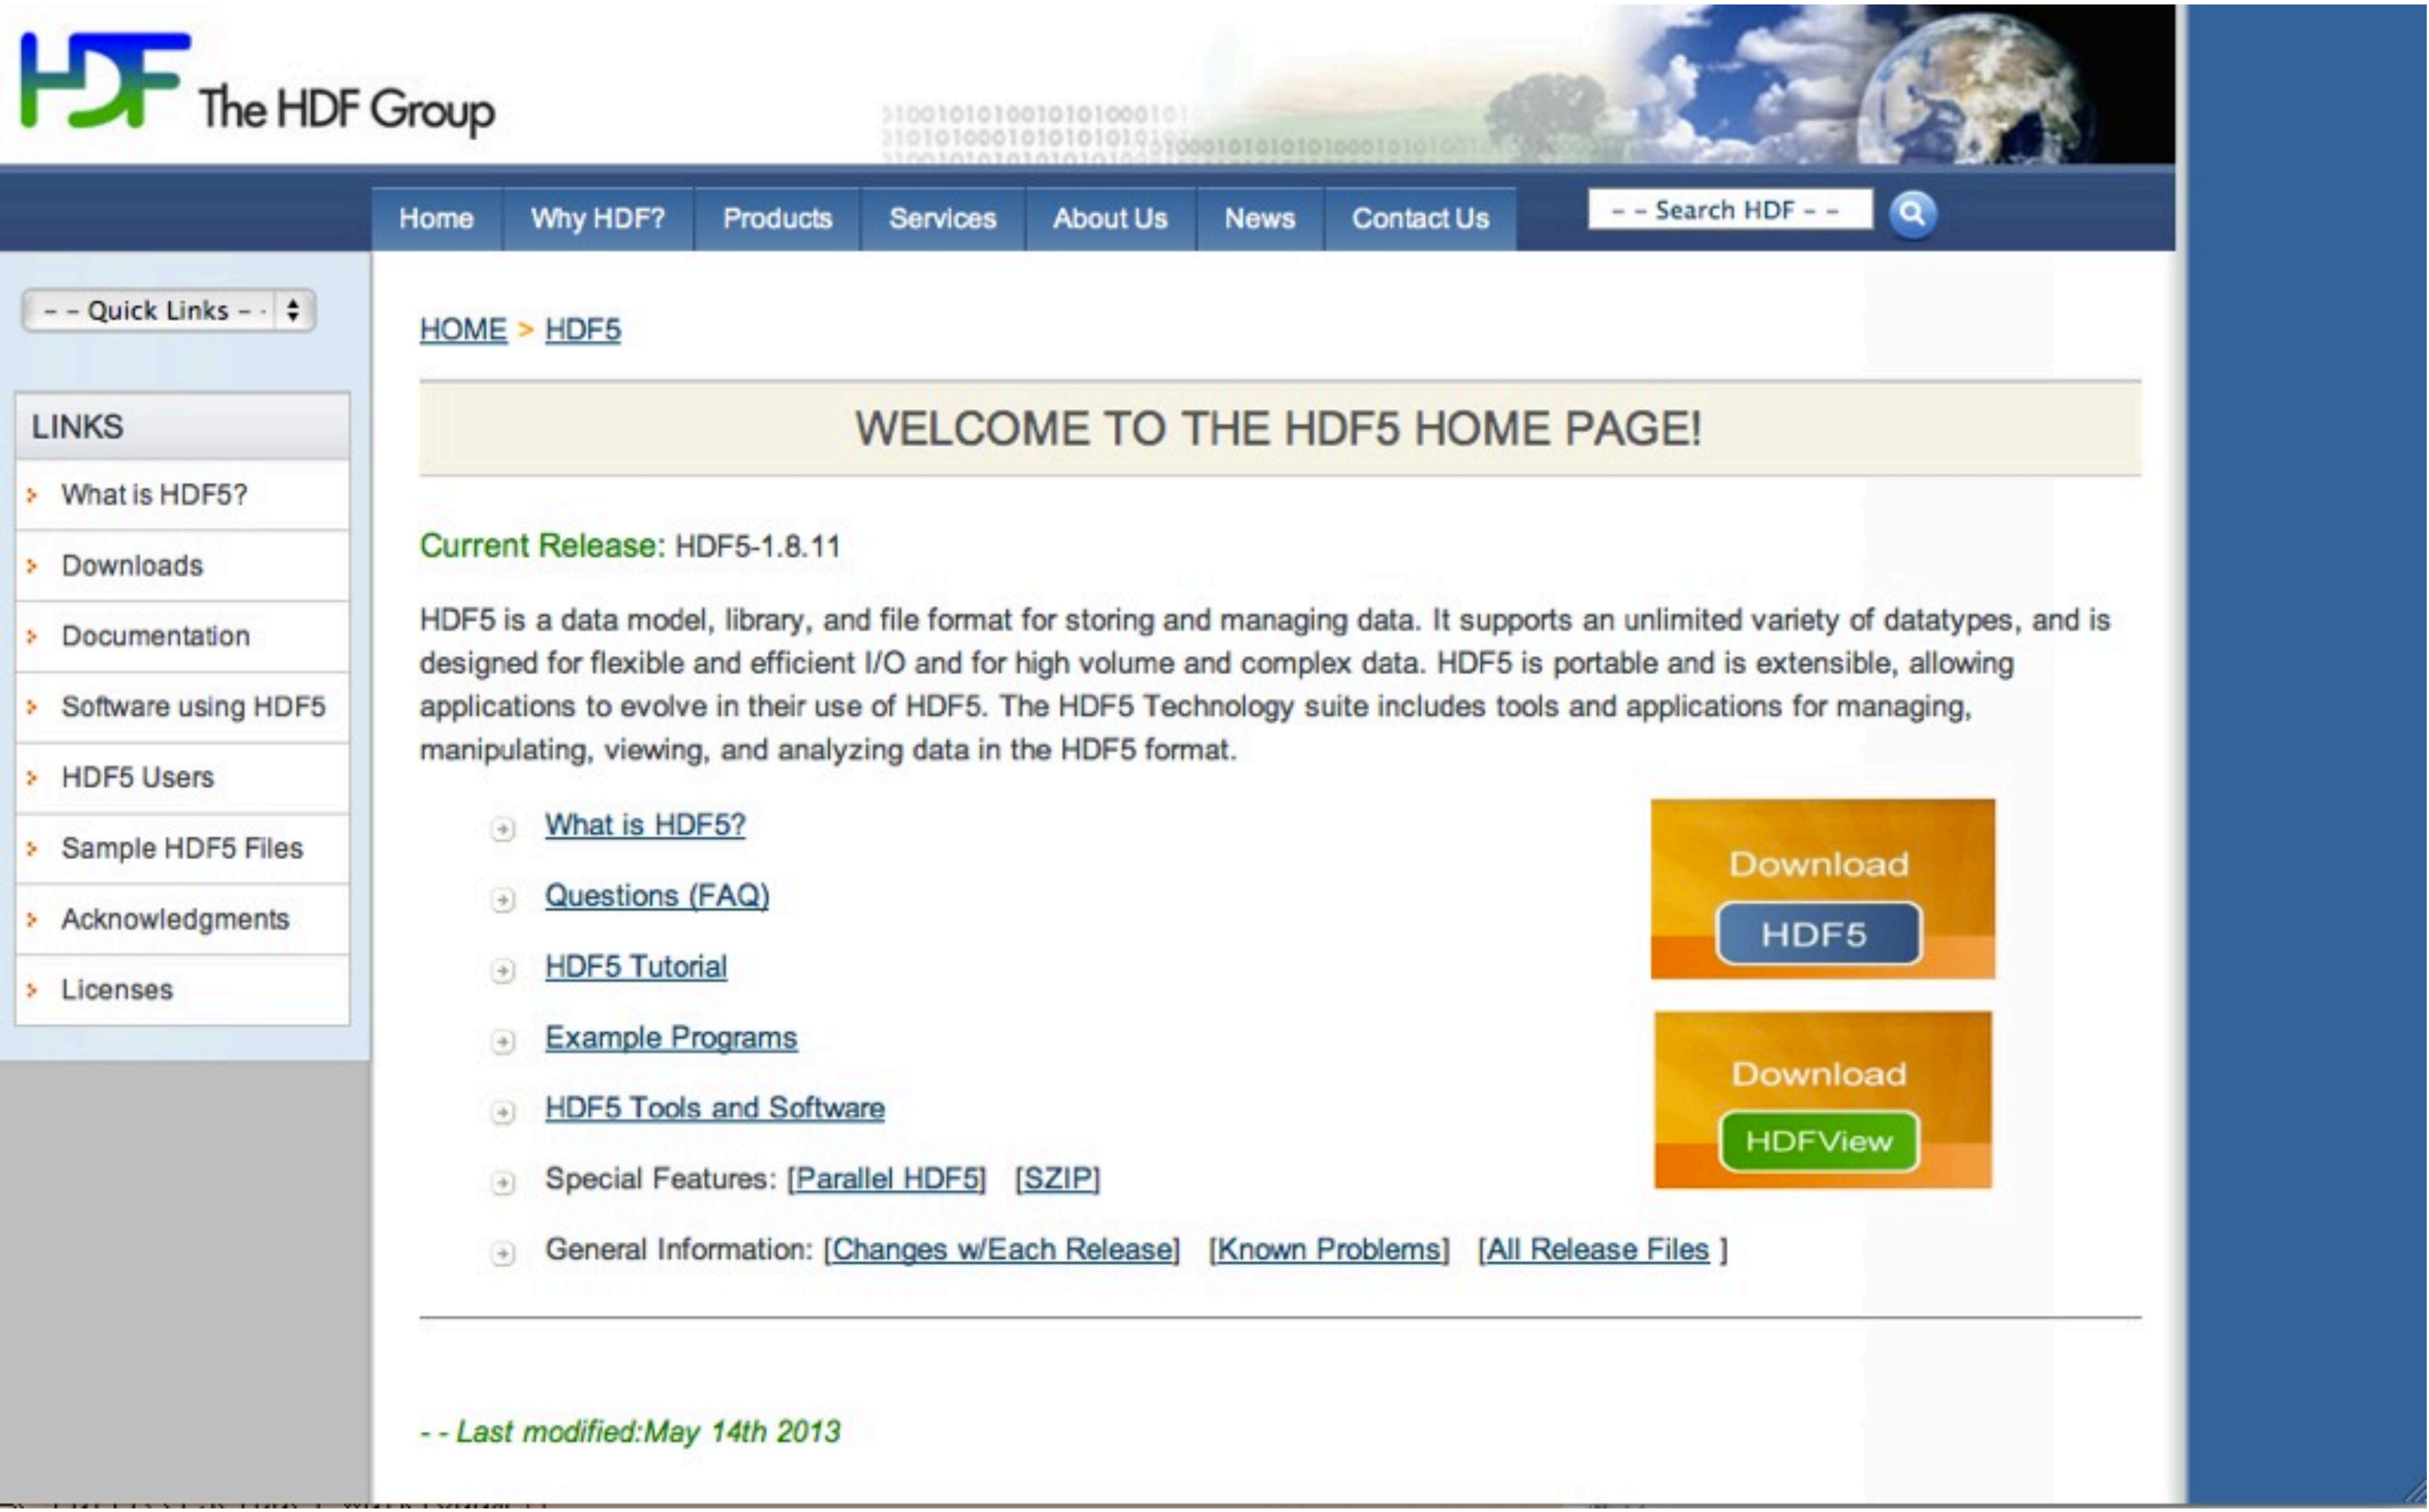

What is HDF5 (in my words)?

Hierarchical Data Format version 5

HDF5 is a "smart" meta-data tree on top of binary (compressible) data files, with fast access to the data. Only the tree needs to be loaded in memory.

HDF5 is basically open-source. It is not gnu public licensed, but anyone can download/install the source code. Also, the HDF5 group manages/controls code updates (as opposed to truly open source code), which ensures a professional reliability of the code.

Why HDF5?

HDF5 is a maintained by a private, non-profit company, *The HDF5 Group* (started at NCSA 1988, spun off in 2005).

- professionally maintained software, stable, freely available
- hooks into HDF5 from basically all common software languages, C, C++, Fortran, Python, Matlab, Java, Mathematica (sort of)

HDF5 is available on all the major US supercomputers

HDF5 supports parallel I/O via MPI-IO

Why HDF5?

HDF5 is used to stress test new large computing/file systems:

Trillion Particles, 120,000 cores and 350 TBs: Lessons Learned from a Hero I/O run on Hopper

Byna, Uselton, Prabhat, Knaak and He

- among other achievements, they were able to write HDF5 files at a sustained rate of 27 GB/s, with I/O dumps of 30-42 TB a time
- collective writes to single file can work as well as file-per-process-writes

The HDF5 Group is involved in the Exascale Fastforward Storage Project

HDF5 is adding a new Virtual Object Layer which adds even more flexibility to the data structures, including being able to store data in other formats (netCDF, HDF4)

Why HDF5 for Lattice QCD?

The HDF5 file structure is organized very similar to a linux/unix file system, where one can have groups/nodes which are like directories, that ultimately store binary data

```
# pytables example
import tables
f = tables.openFile('my_data_file.h5')
f.getNode('/proton')
    children := ['spin_up' (Group), 'spin_down' (Group)]
f.getNode('/proton/spin_up')
    children := ['psq_0' (Array), 'psq_1' (Array)]
```

The data files are naturally stored as multi-dimensional arrays

```
f.getNode('/proton/spin_up/psq_0')
    Array(291, 256, 2, 1)
    atom := Float64Atom(shape=(), dflt=0.0)
    flavor := 'numpy'
    byteorder := 'little'
    chunkshape := None
```

HDF5 supports (arbitrarily) complicated meta-data descriptors - no need for multiple files to describe binary data

Why HDF5 for Lattice QCD?

Adopting HDF5 will further standardize our data files and alleviate the need for us to worry about I/O issues (at least at the level we do now)

```
#include “hdf5.h”
```

With tools to access HDF5 files in Matlab, Python, etc. we will be able to significantly improve our data visualization (which is currently nearly non-existent)

You have seen this figure at this workshop (as well as other conferences) - it was made in 2003 (too be fair, it is challenging to make visualizations of imaginary-time Quantum Mechanical Concepts)

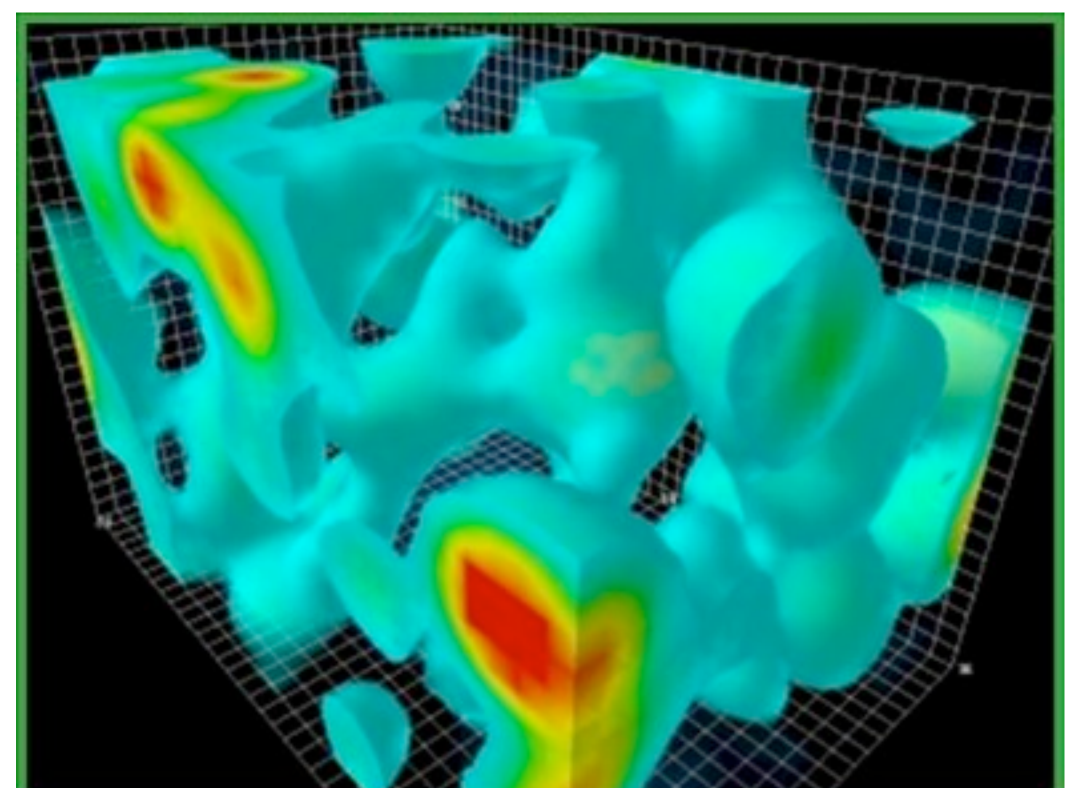

Adding Optional Support for HDF5 in USQCD Software working with Abhinav Sarje (LBNL) (coordinated with Bálint Joó)

We are currently adding optional compile of QDP++ with HDF5

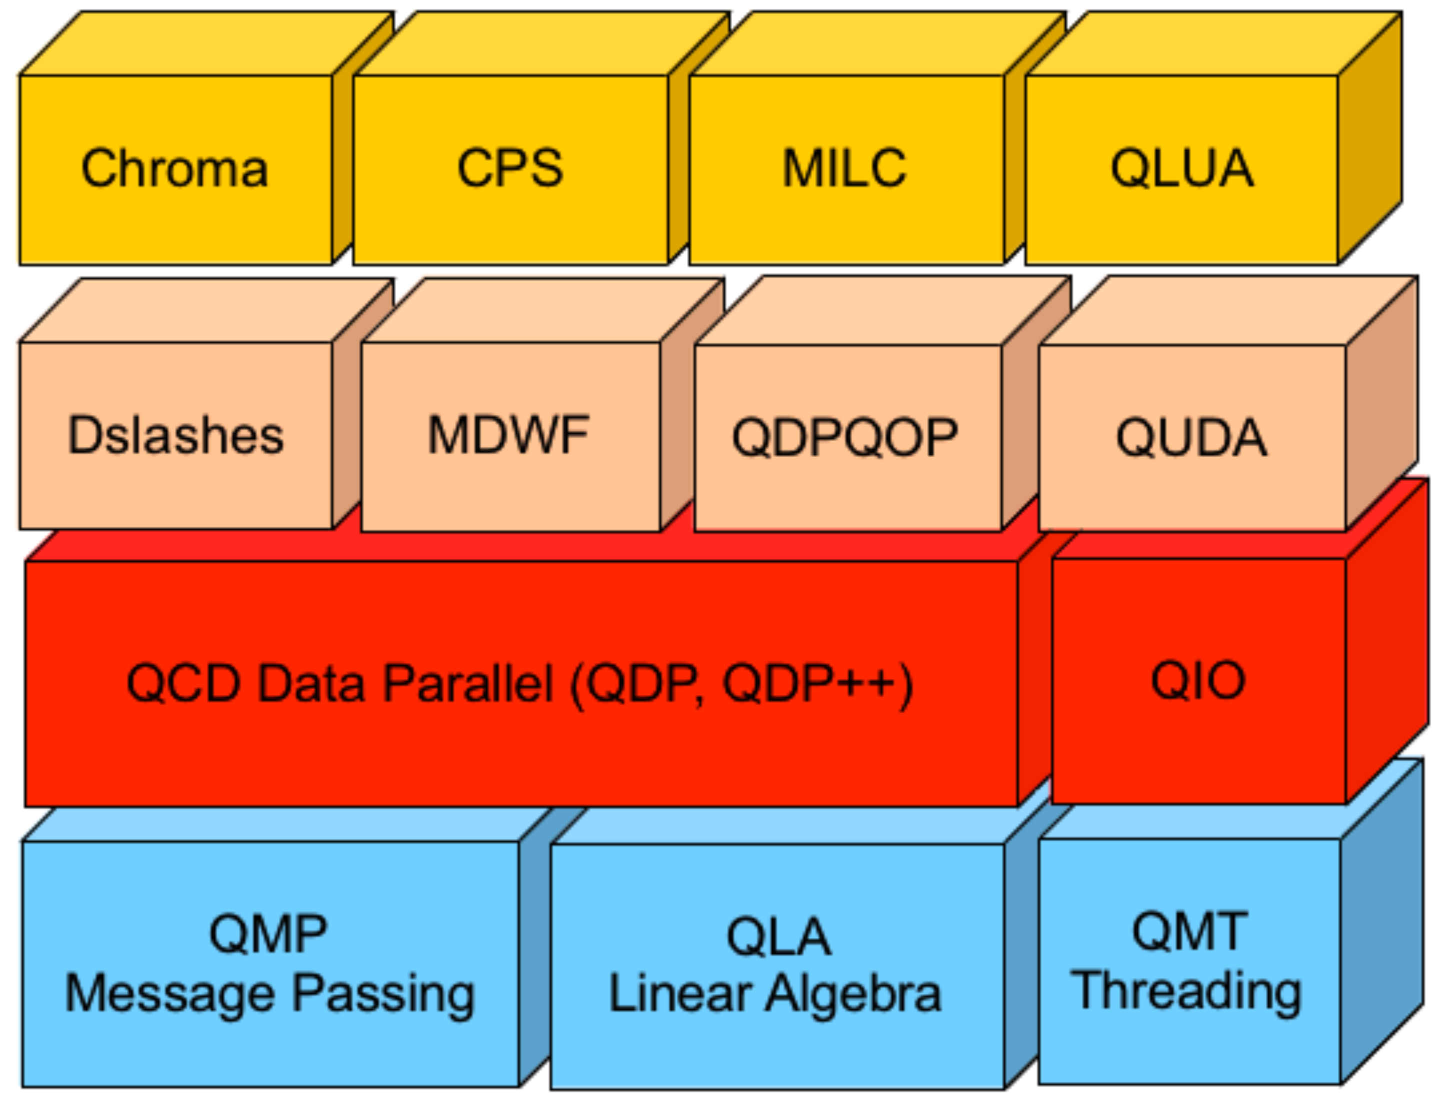

This will include I/O for all the basic file structures we currently write

configs, propagators,
correlation functions

Andrew Pochinsky has already implemented HDF5 into QLUA

Adding Optional Support for HDF5 in USQCD Software
working with Abhinav Sarje (LBNL)
(coordinated with Bálint Joó)

At a personal level, this is quite a fun collaboration.
My coding experience/exposure so far has mostly been
by the seat of my pants
Now I get to work with people who have years of proper coding experience, and so my own abilities are improving (hopefully significantly)
Hopefully, those I am working with are getting a better understanding of the science we are doing also.

Thank You