Lattice Quantum Chromodynamics Project

SciDAC-4

Robert Edwards

Jefferson Lab

LQCD ASCR/NP SciDAC-4

Computing the Properties of Matter with Leadership Computing Resources

PI: Robert Edwards (JLab)

Co-PIs: Will Detmold (MIT), Balint Joo (JLab), Swagato Mukherjee (BNL)

Senior Investigators:

Andrei Alexandru (GWU)
Saman Amarasinghe (MIT)
Alexei Bazavov (MSU)
Kate Clark (NVIDIA)
Rob Fowler (UNC)
Dhiraj Kalamkar (Intel)
Xu Liu (W&M Computer Sci)
Kostas Orginos (W&M Phys)
Sergey Panitkin (BNL)
Andrew Pochinsky (MIT)
Kenneth Roche (PNNL)
Martin Savage (UW)
Frank Winter (JLab)
Boram Yoon (LANL)

Team members

Long term collaborations with ASCR supported community and Industry

- Gauge generation (co-PI: Balint Joo)
  - Balint Joo - Sparse linear system solvers on emerging LCF-s (Cori, Summit, NERSC-9)
    - Siva Rajamanickam - KokkosKernels FASTMATH
  - Boram Yoon - Sympletic PDE Integrators
  - Kate Clark (NVIDIA) - comms reduced solvers, Integrators
  - Kostas Orginos - 1 flavor methods
  - Xu Liu - Memory optimizations (QPerf and HPC Toolkit)
  - Robert Fowler - QUARC/DSL interface to Clang/LLVM - automatic code generation

- Correlation functions/contractions (co-PI: Will Detmold)
  - Saman Amarasinghe - (TACO) code generation, auto-tuning for contractions (& gauge gen)
  - Andrew Pochinsky - Halide for QCD
  - Kenneth Roche - workflow, data reductions/sparsification/SVD approximations for contractions
  - Andrei Alexandrou - overlap analysis campaigns for KNL-s

- Thermodynamics and Workflow (co-PI: Swagato Mukherjee)
  - Sergey Panitkin: PanDA/ATLAS workflow for LCF systems, multi-site campaigns, scheduling, file transfers & data integrity
  - Alexei Bazavov - transport coefficients

LQCD/NP Science & connection to Expt.

QCD stands for *Quantum Chromo-Dynamics*
the theory that describes the interactions of quarks and gluons that constitute the matter of the visible universe

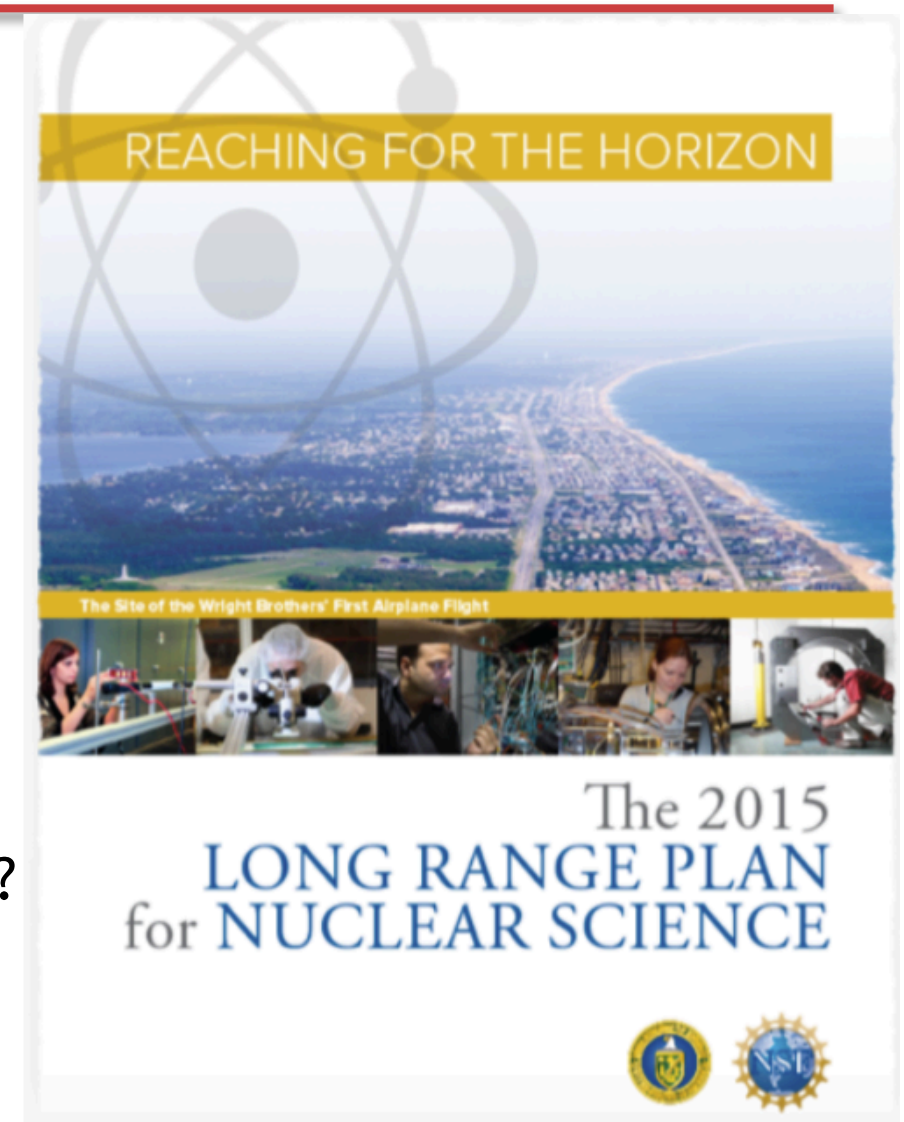

- What observable states does QCD allow?
  - What is the role of the gluons? Is there exotic matter?
  - Focus of GlueX@JLab experiment
- How do nucleons (protons & neutrons) arise?
  - How are quarks & gluons distributed in a proton or neutron?
  - Focus of 12 GeV@JLab, RHIC-spin@BNL and future EIC
- QCD must predict properties of light nuclei
  - Nuclear reaction properties. Are there new fundamental symmetries?
  - FRIB@Mich. State. will investigate nuclear structure and interactions
- How does QCD behave under extreme temperatures & pressures such as in supernovae or shortly after Big-Bang?
  - Studied in RHIC@BNL & ALICE@LHC

http://science.energy.gov/~/media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf

Lattice QCD & the excited hadron spectrum

LQCD SciDAC project has pioneered novel techniques in lattice QCD

Light quark meson + "exotics" & "hybrids" spectrum

Exotics

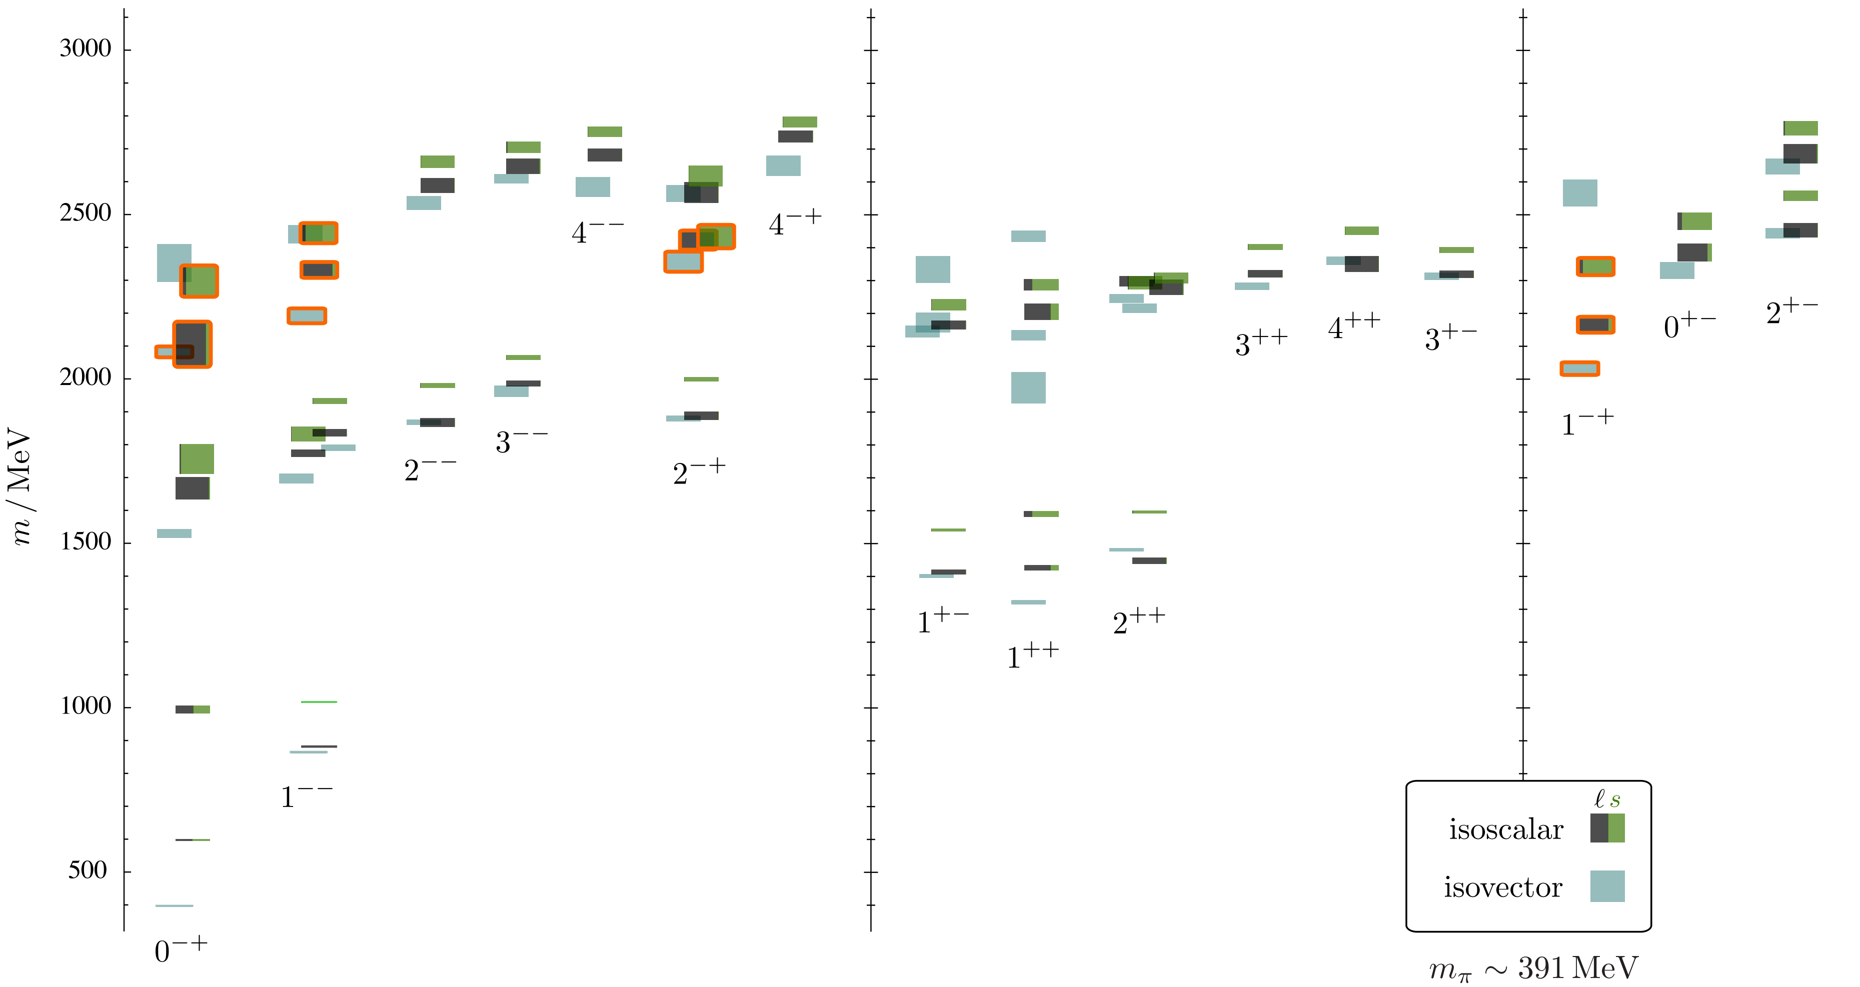

'Toward the excited isoscalar meson spectrum from lattice QCD'
PRD 88 094505 (2013)

clear indication of a spectrum of **hybrid mesons**

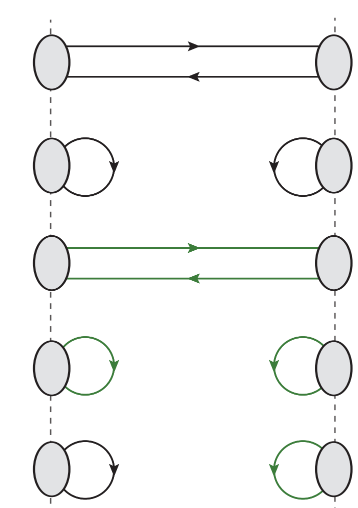

high precision calculation of **disconnected diagrams**

Jefferson Lab

Hadron Spectroscopy - role of the glue

Focus of GlueX & CLAS12 @ JLab & COMPASS, LHCb @ CERN, BES @ Beijing

Light quark meson + "exotics" & "hybrids" spectrum

Exotics

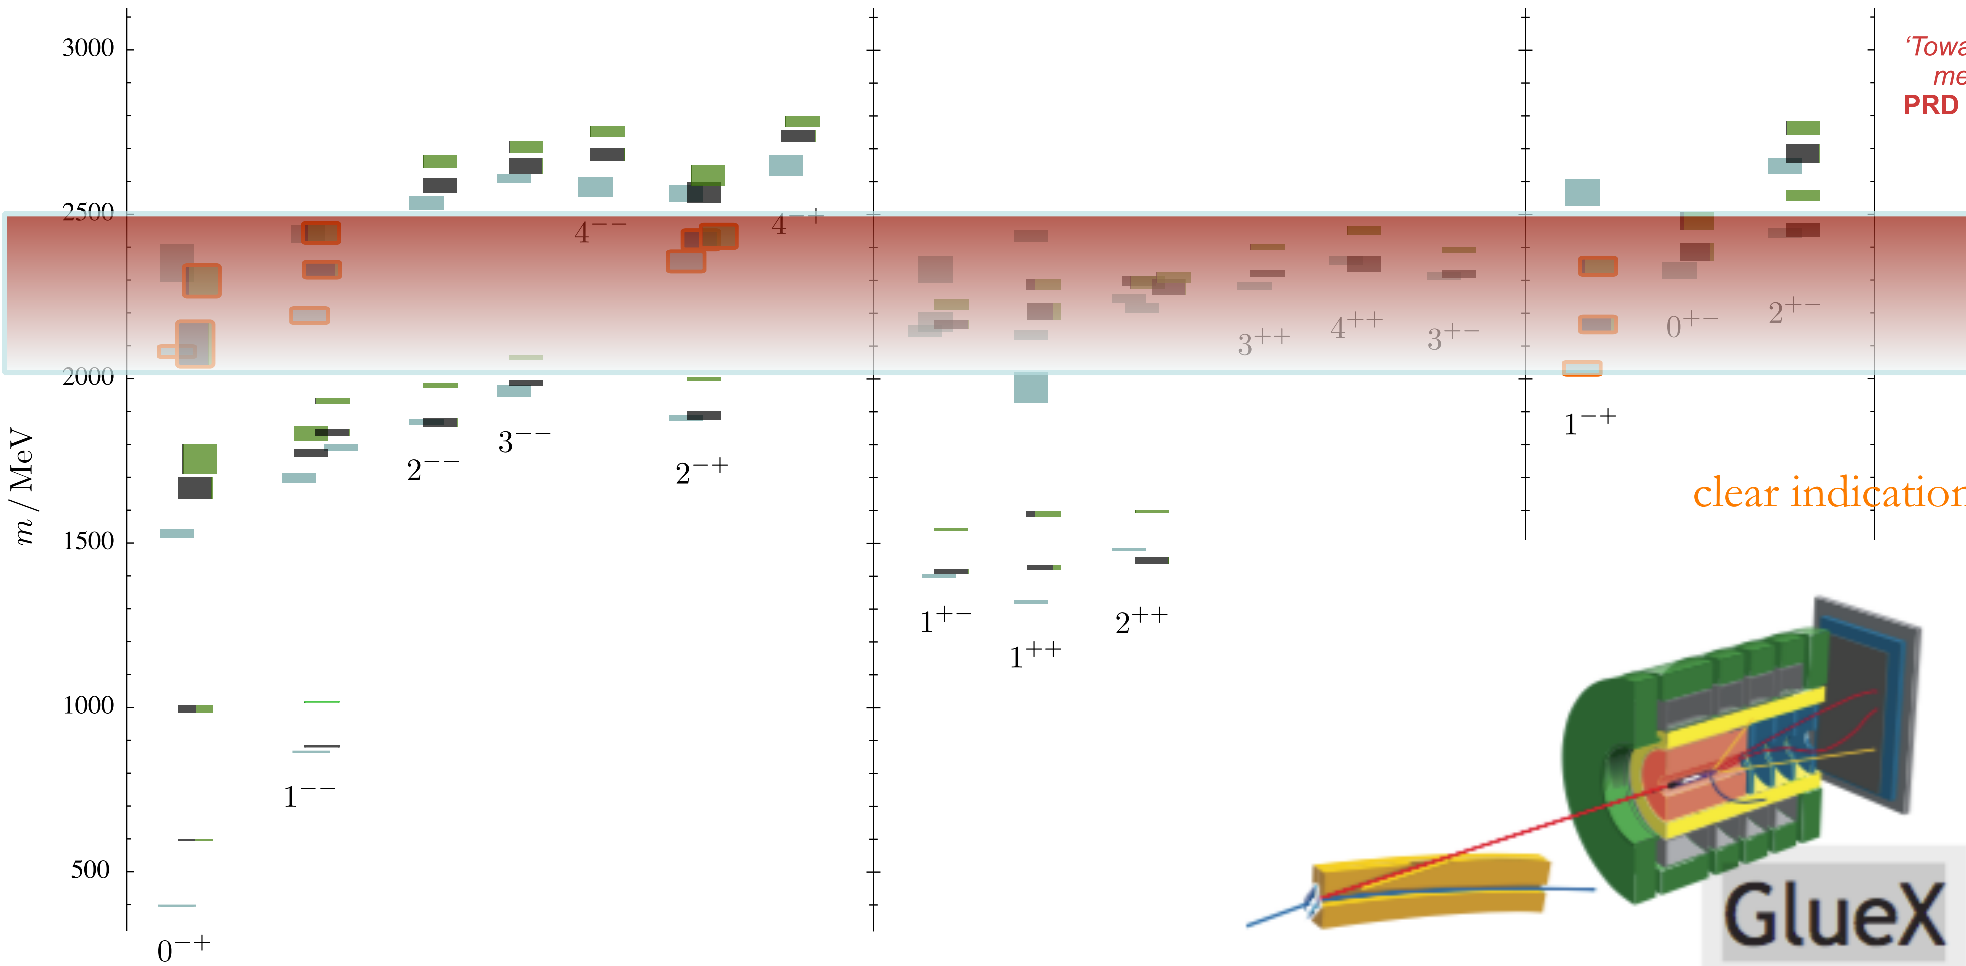

'Toward the excited isoscalar meson spectrum from lattice QCD'
PRD 88 094505 (2013)

clear indication of a spectrum of **hybrid mesons**

Pattern of states suggest gluonic excitations

➡ Need to know decay modes and rates to compare to expt.

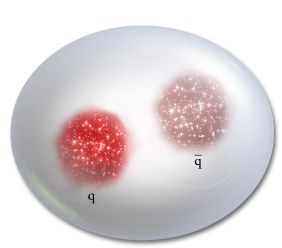

Conventional Meson

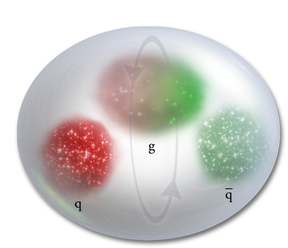

Hybrid Meson

Jefferson Lab

First calculations of decay modes

Light quark meson + “exotics” & “hybrids” spectrum

Exoti

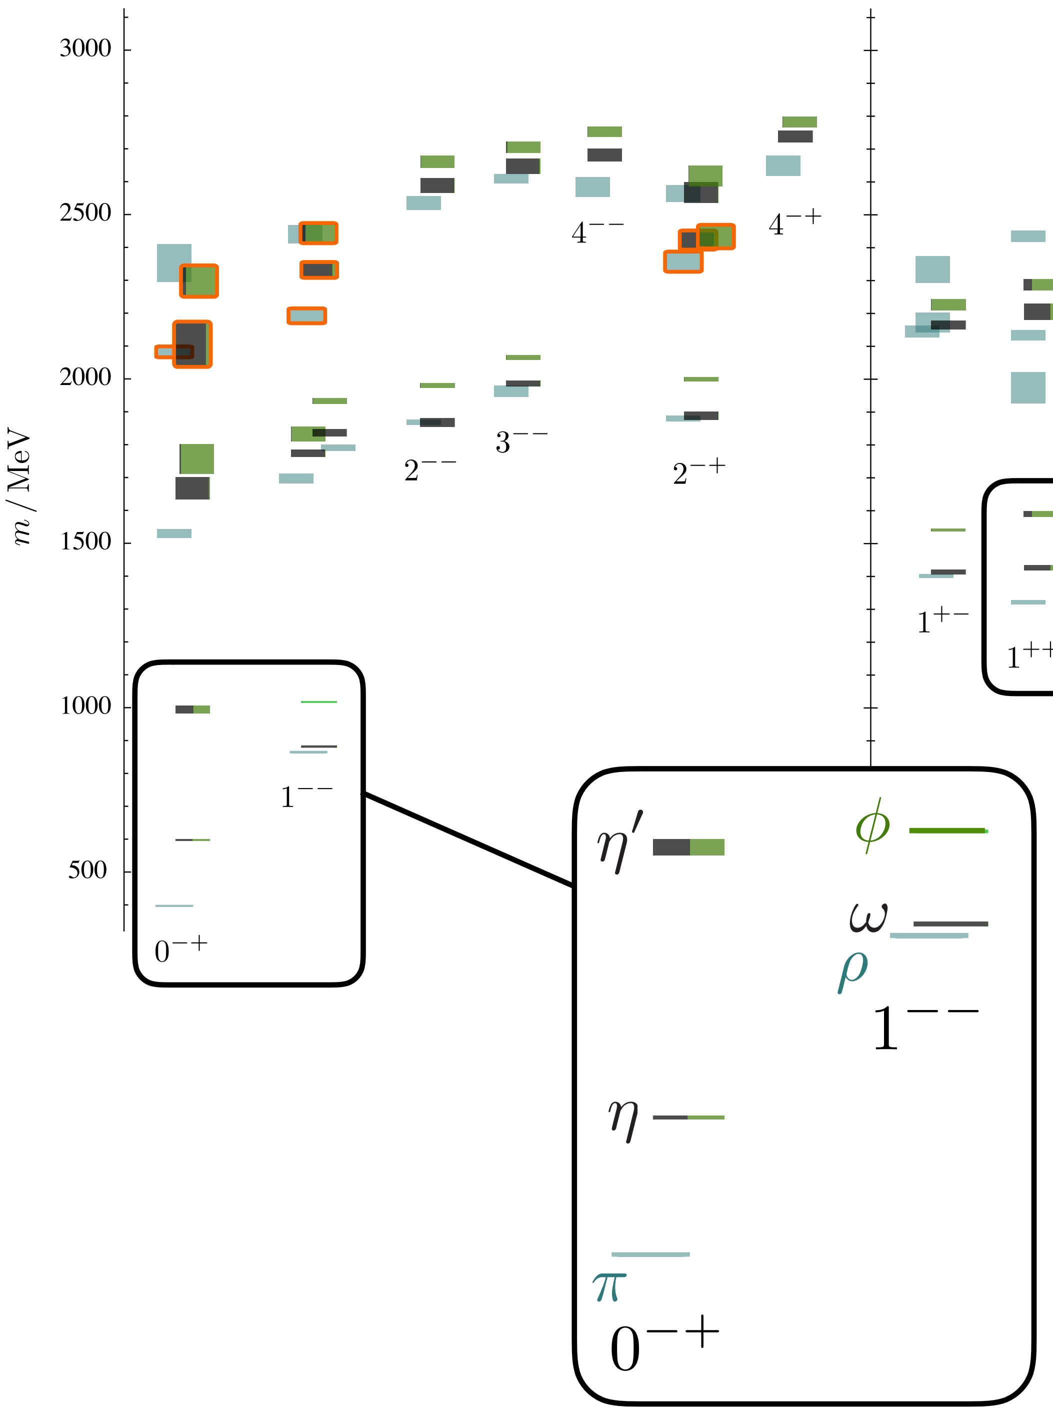
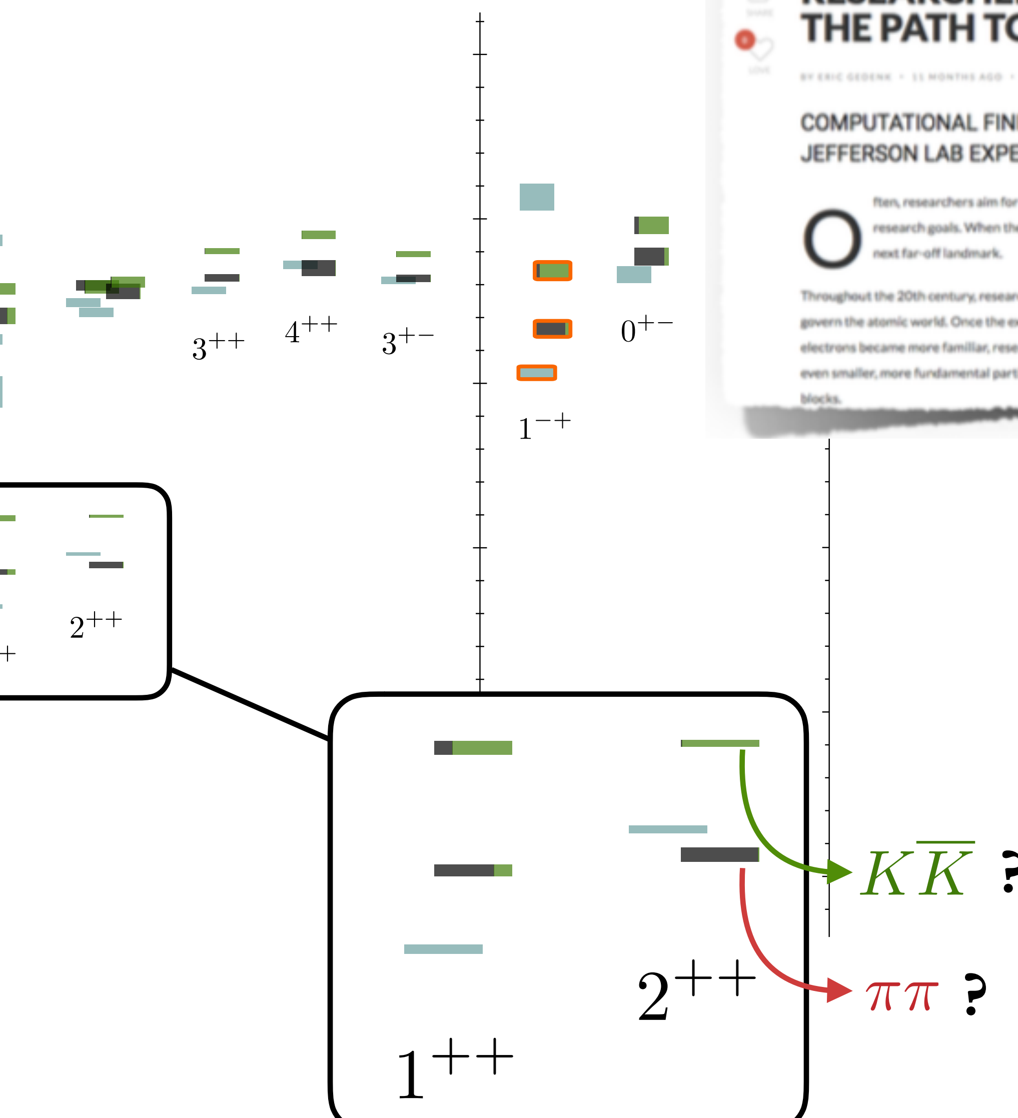

flavor tagging by decays

PDG

$f_2(1525)$ ↪ $K\bar{K}$ 89%

$f_2(1270)$ ↪ $\pi\pi$ 84%

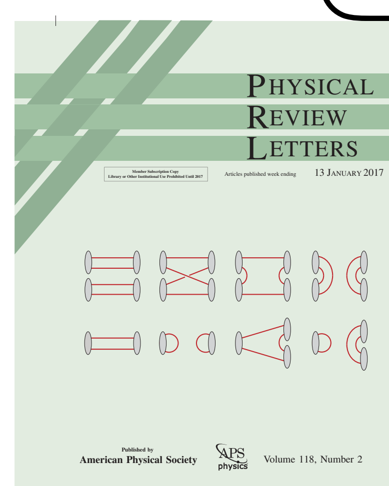

Used as motivation in
GlueX detector upgrade proposal

HadSpec: Phys.Rev.D97 (2018);
Phys.Rev.Lett.118 (2017)

Proton structure

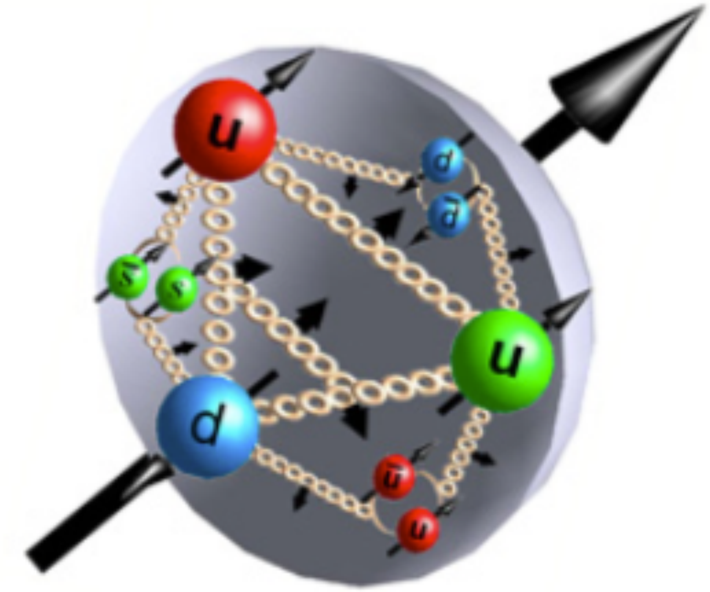

- Proton has intrinsic spin J = 1/2
- How is this distributed amongst the constituents
- "Proton spin crisis"(1988): spin of quarks accounts ~ 1/8?

- Disentangling this had take many years

Proton structure

- Proton has intrinsic spin J = 1/2
- How is this distributed amongst the constituents
- "Proton spin crisis"(1988): spin of quarks accounts ~ 1/8?

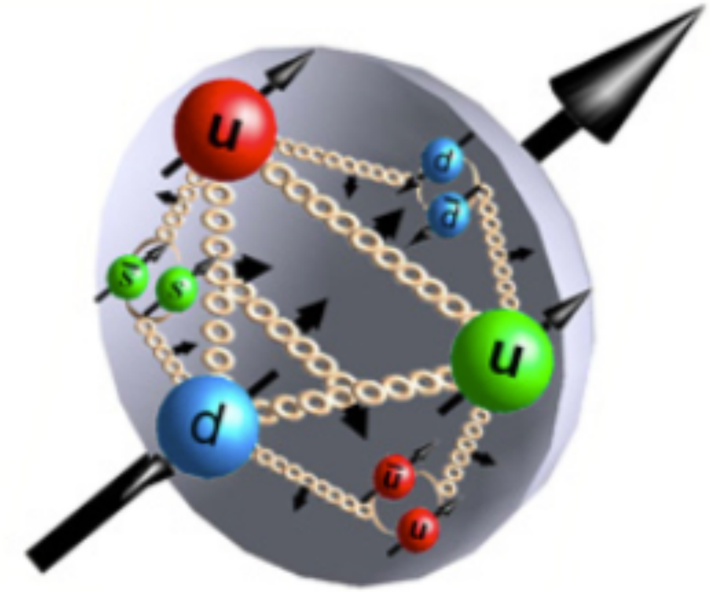

Spin is 1/2 !

First direct calculations
[caveats - only 2 dynamical flavors]

Improving under SciDAC-4

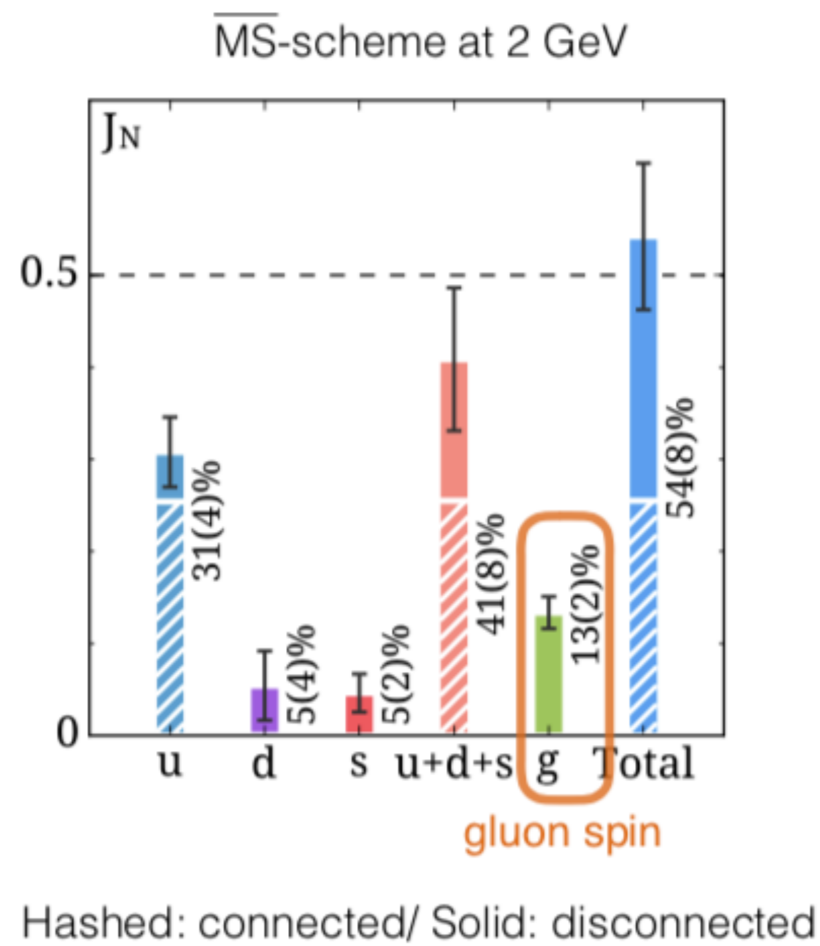

ETMC: Phys.Rev.Lett.119 (2018)

Nuclei

- NPLQCD: interactions of nuclei
- np ➔ dγ neutron capture [PRL 2016]
- pp ➔ deν fusion process [PRL 2017]
  - starts the solar fusion cycle
  - difficult to access in expt.
  - first LQCD calc. - competitive to extractions
- Coupling to scalar currents [dark matter] [PRL 2018]
  - significant nuclear effects

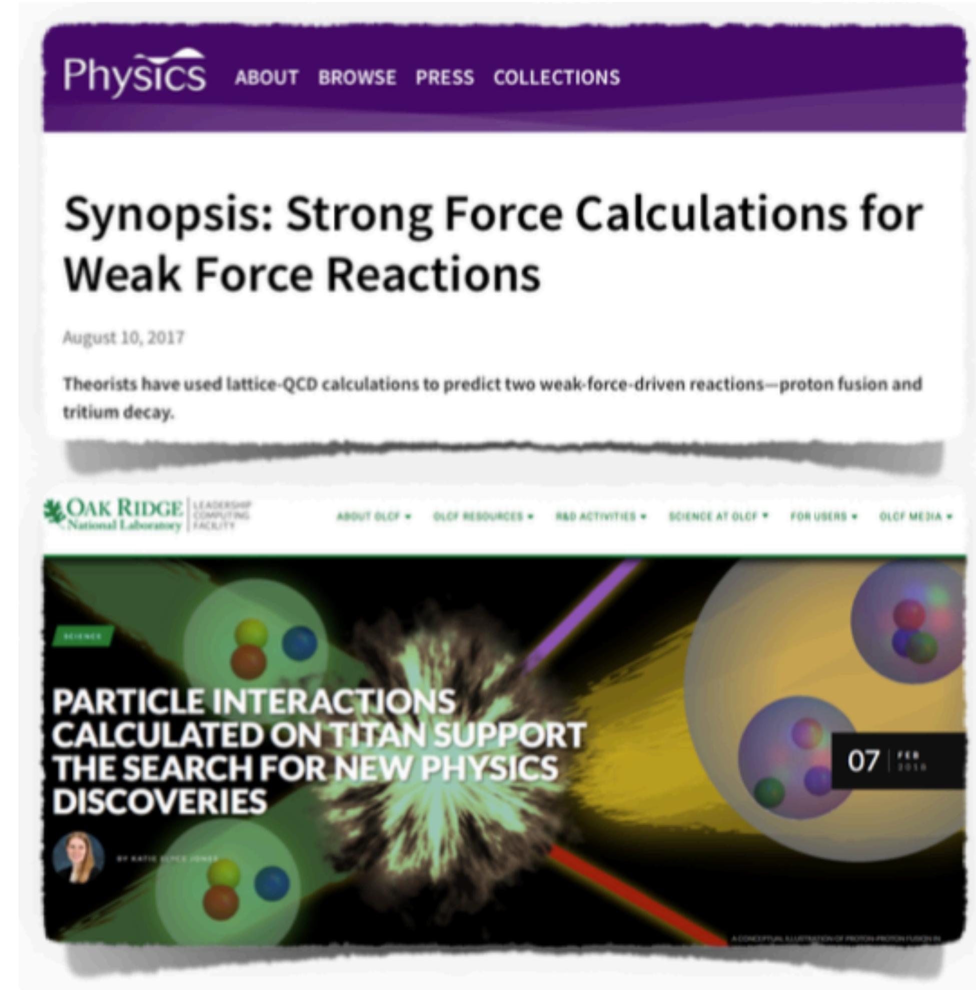

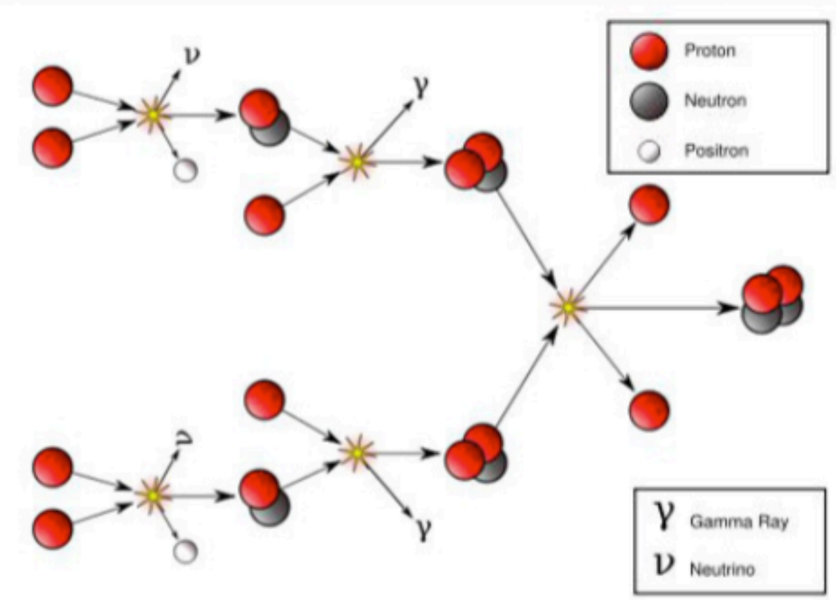

Hot-Dense matter

properties of quark gluon plasma (QGP)

the big questions …

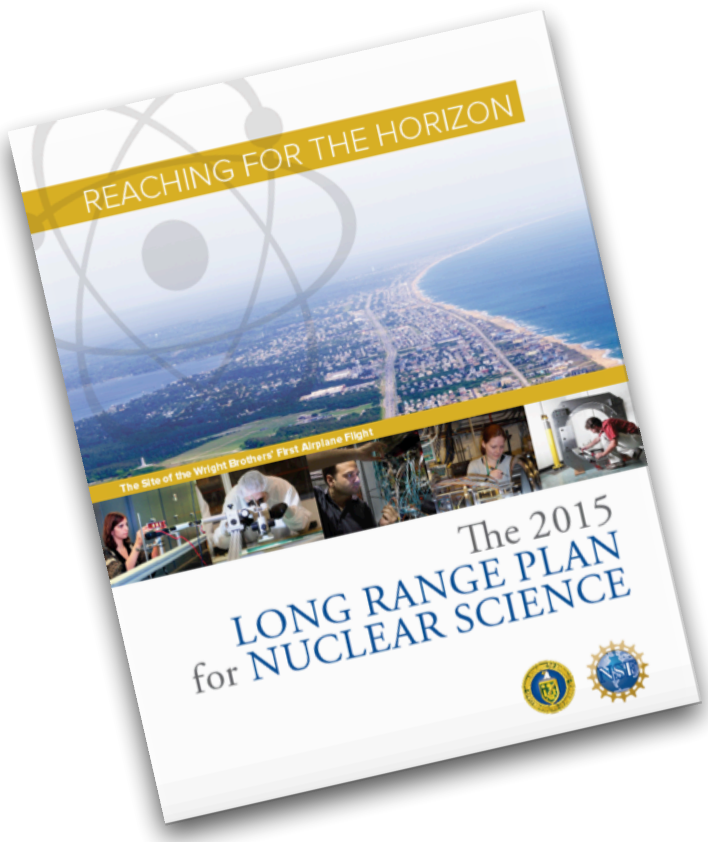

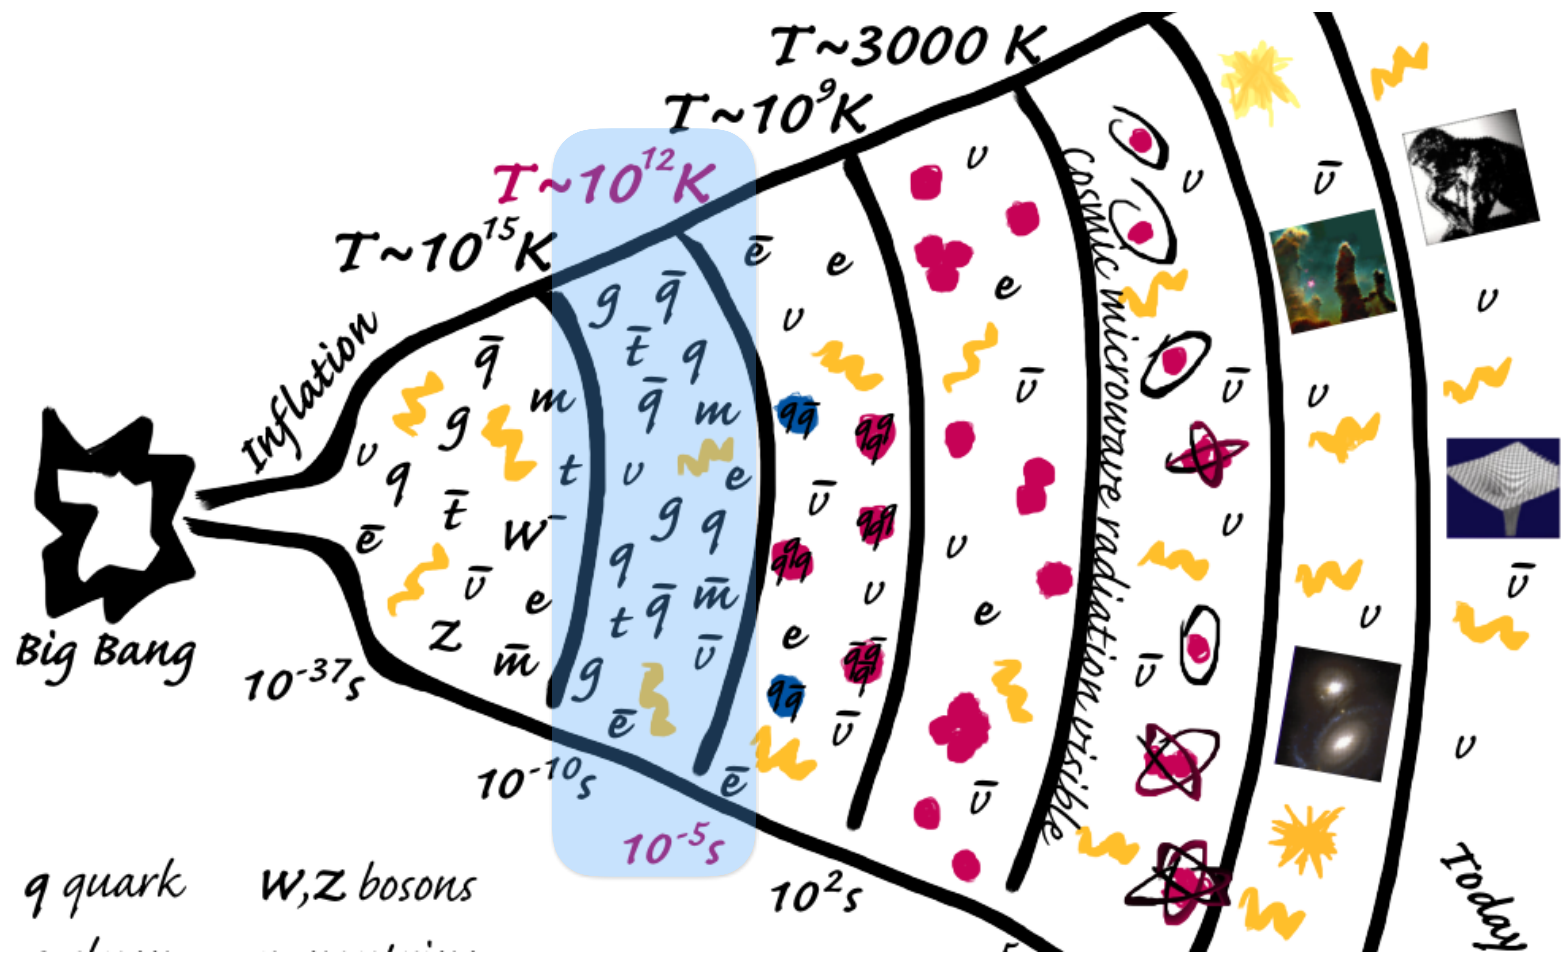

characterize the matter that existed
~ micro-seconds after the Big Bang

Hot-Dense matter

properties of strongly interacting matter

the big questions …

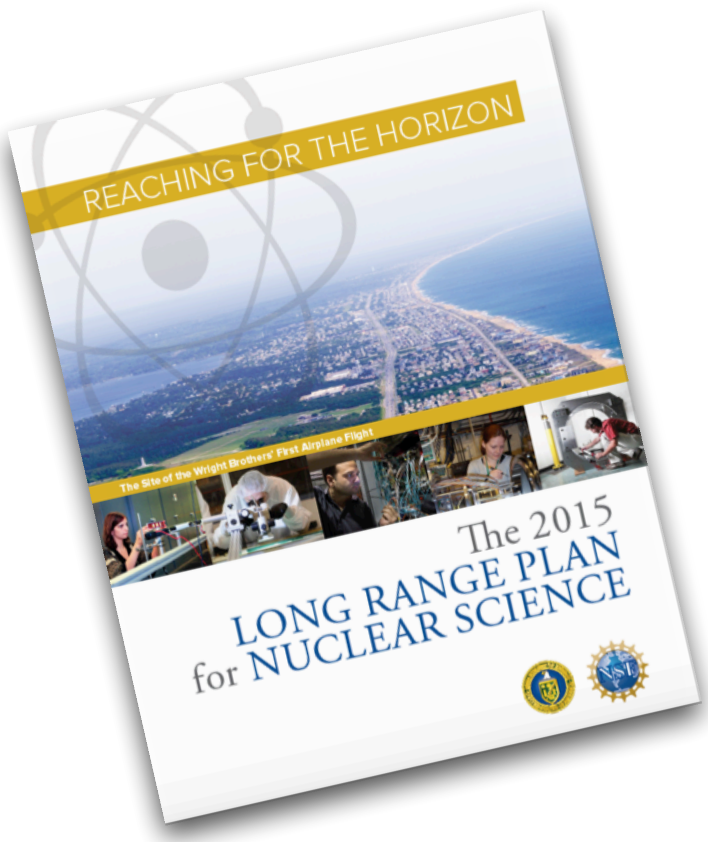

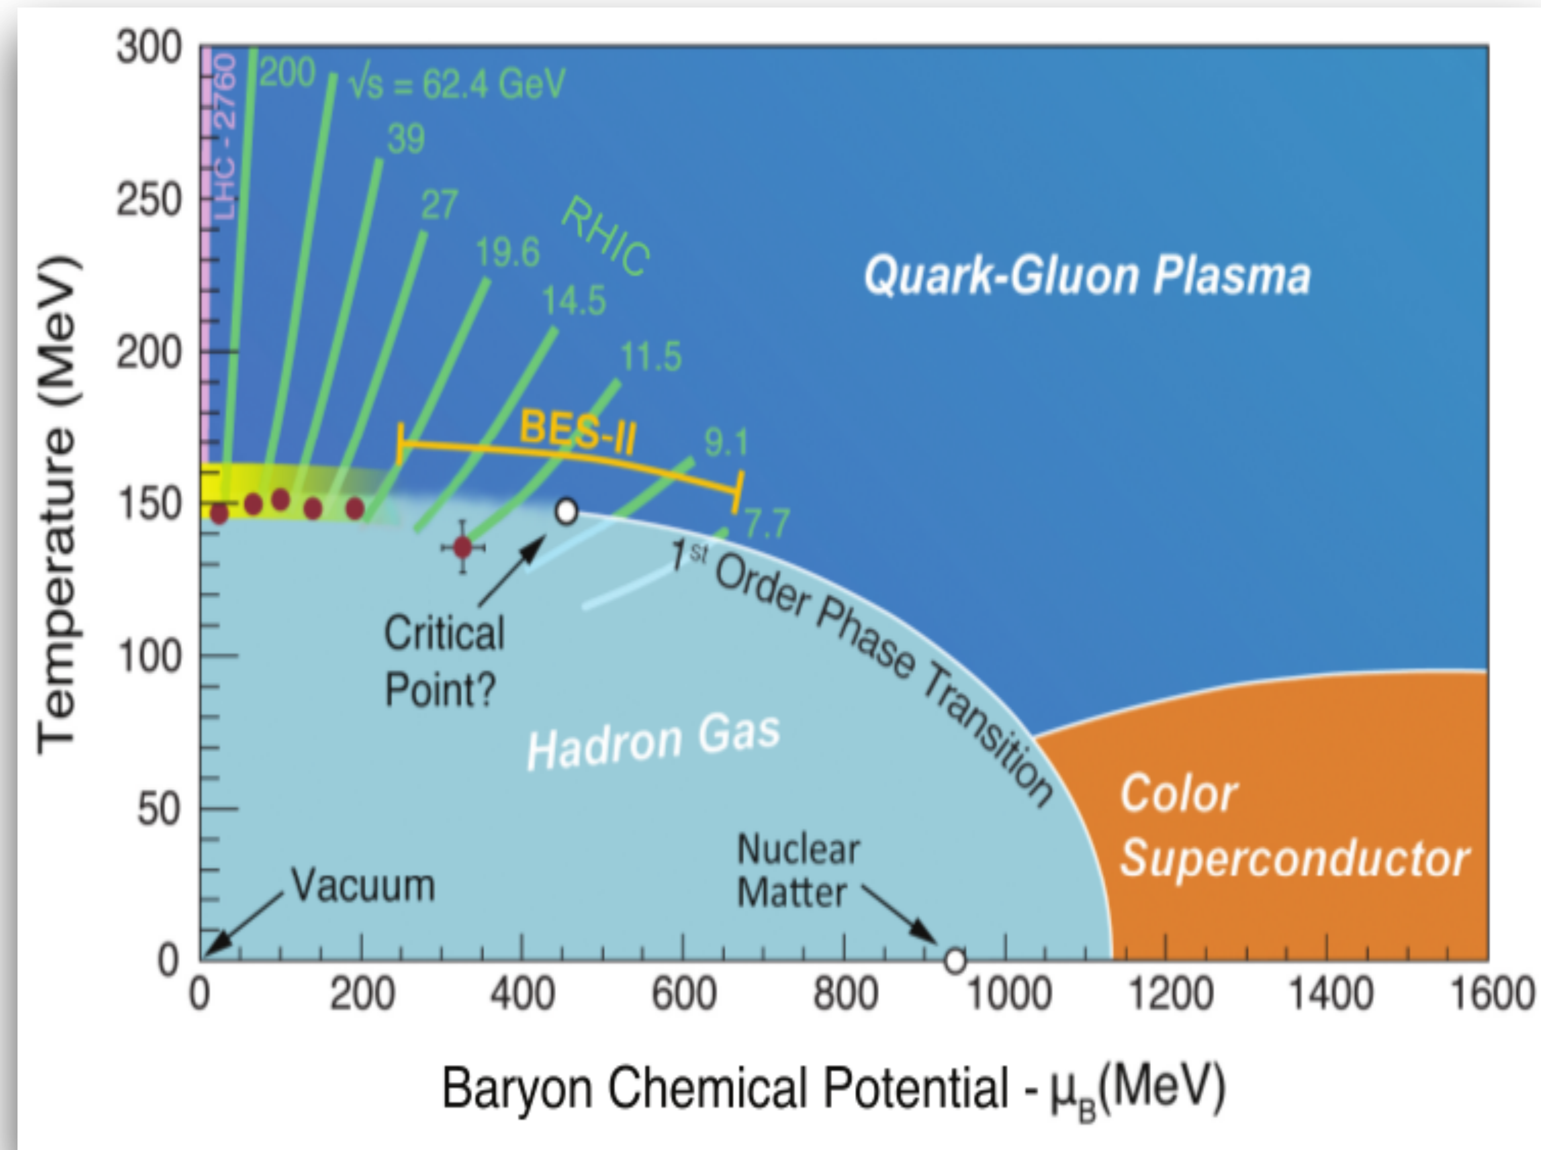

- varying temperatures and densities

Hot-Dense matter - next gen. expt. at BNL

Beam Energy Scan-II @ RHIC

2019-2021

Phase diagram? Critical point?

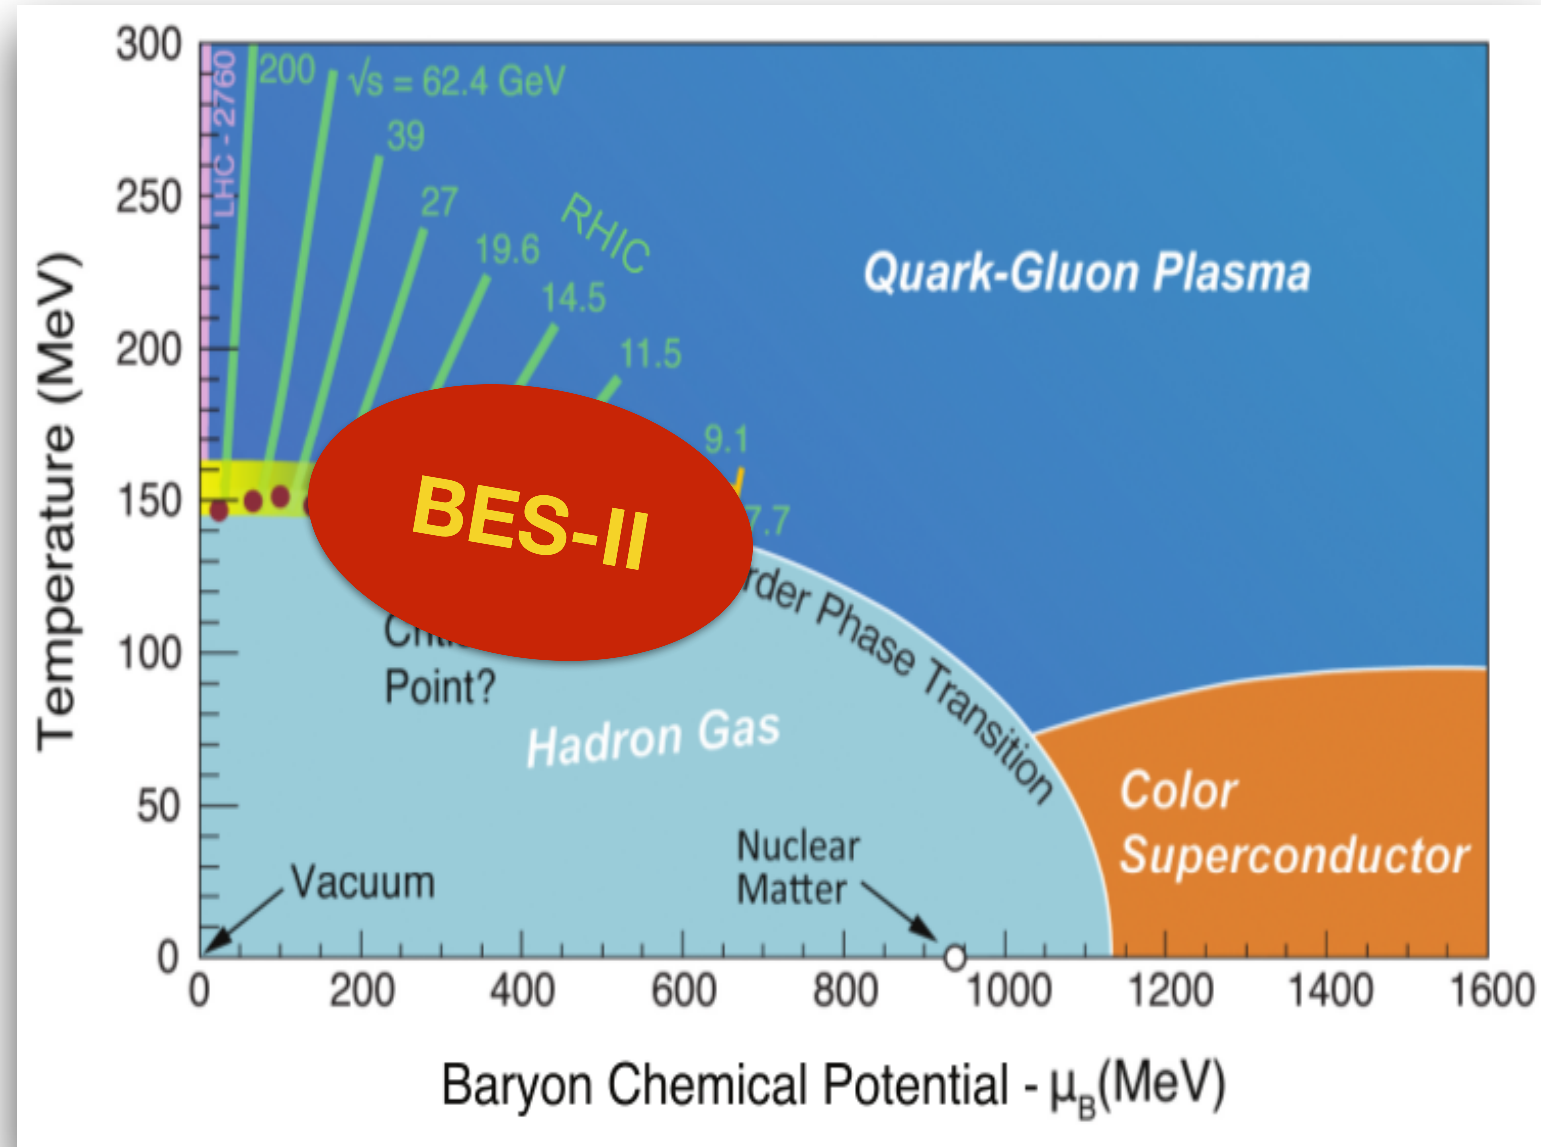

Computational Requirements

Gauge generation
- Leadership level
- Strong scaling

Strong scaling limited
- Comms/compute less balanced in recent machines
- Latency in comms important
- Require comms reduced linear solvers
- 1-flavor solvers

Sympletic PDE integrators

Data-parallel code gen. with comms

Analysis
- Throughput level
- O(1M) RHS-s/cfg

Throughput challenge
- Solvers still important
- Large problems / node -> minimize memory traffic
- Job coordination -> Grid based

Contractions

Correlators
- O(1M)/cfg

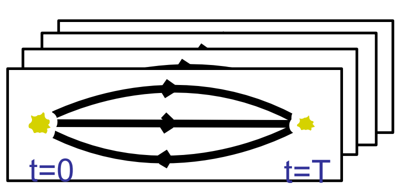

"Workflow" challenge
- Combinatorics -> improved graph theoretic methods
- Code generation for (sparse?) tensor contractions
- Job coordination -> Grid based

How to produce gauge fields?

- Hamilton's equations - 1st order coupled differential eqns.
- Each integration step: sparse matrix solution
- Strong scaling challenge

- Bummer!
  - Must be "reversible"
  - No adaptive time steps

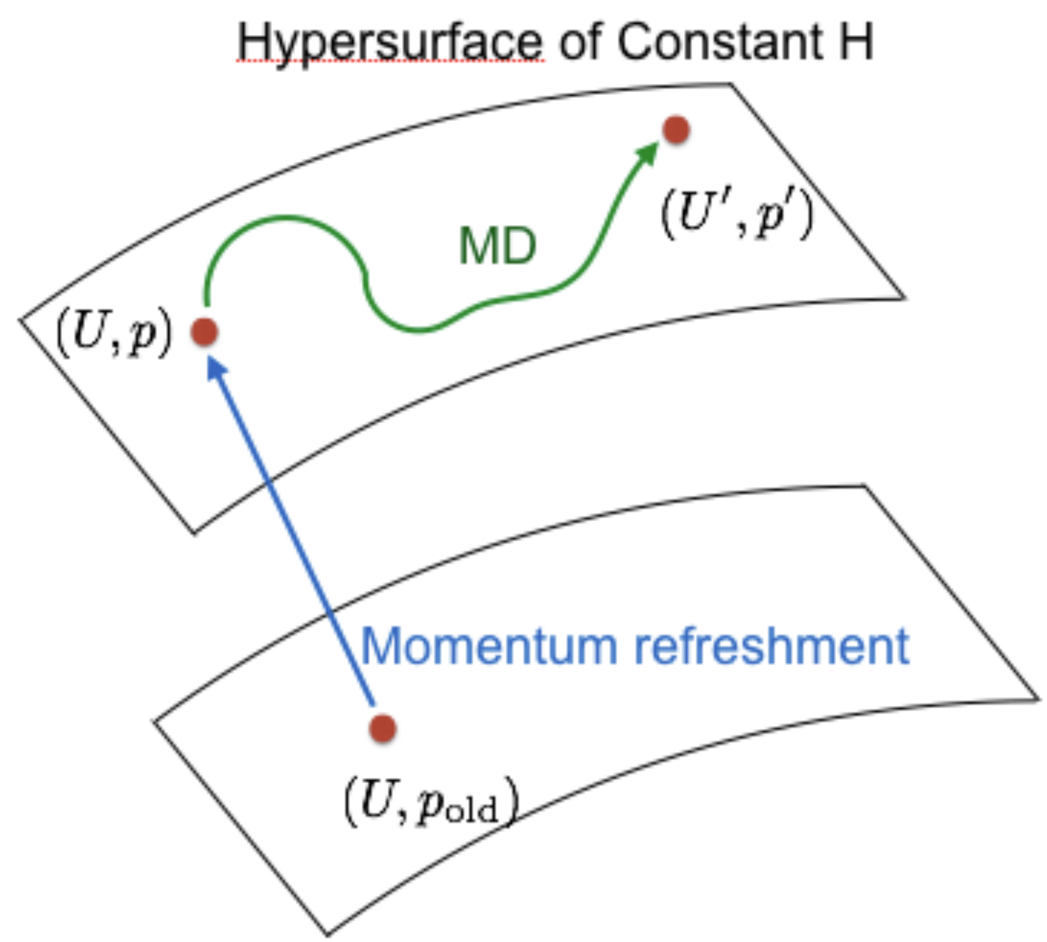

Accelerating QCD gauge generation on GPUs

Collaboration involving ASCR support and Industry partners

[See poster by Balint Joo]

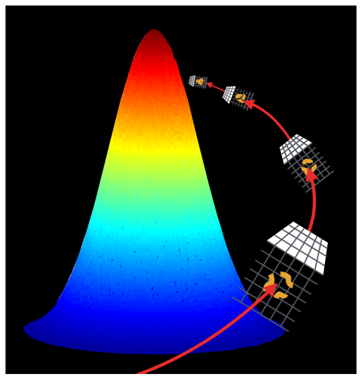

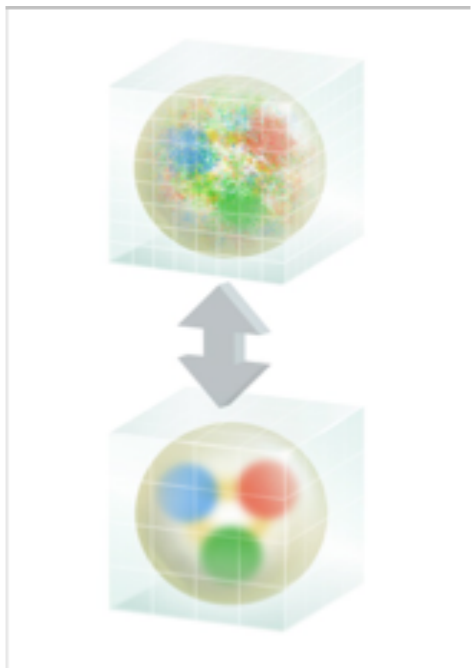

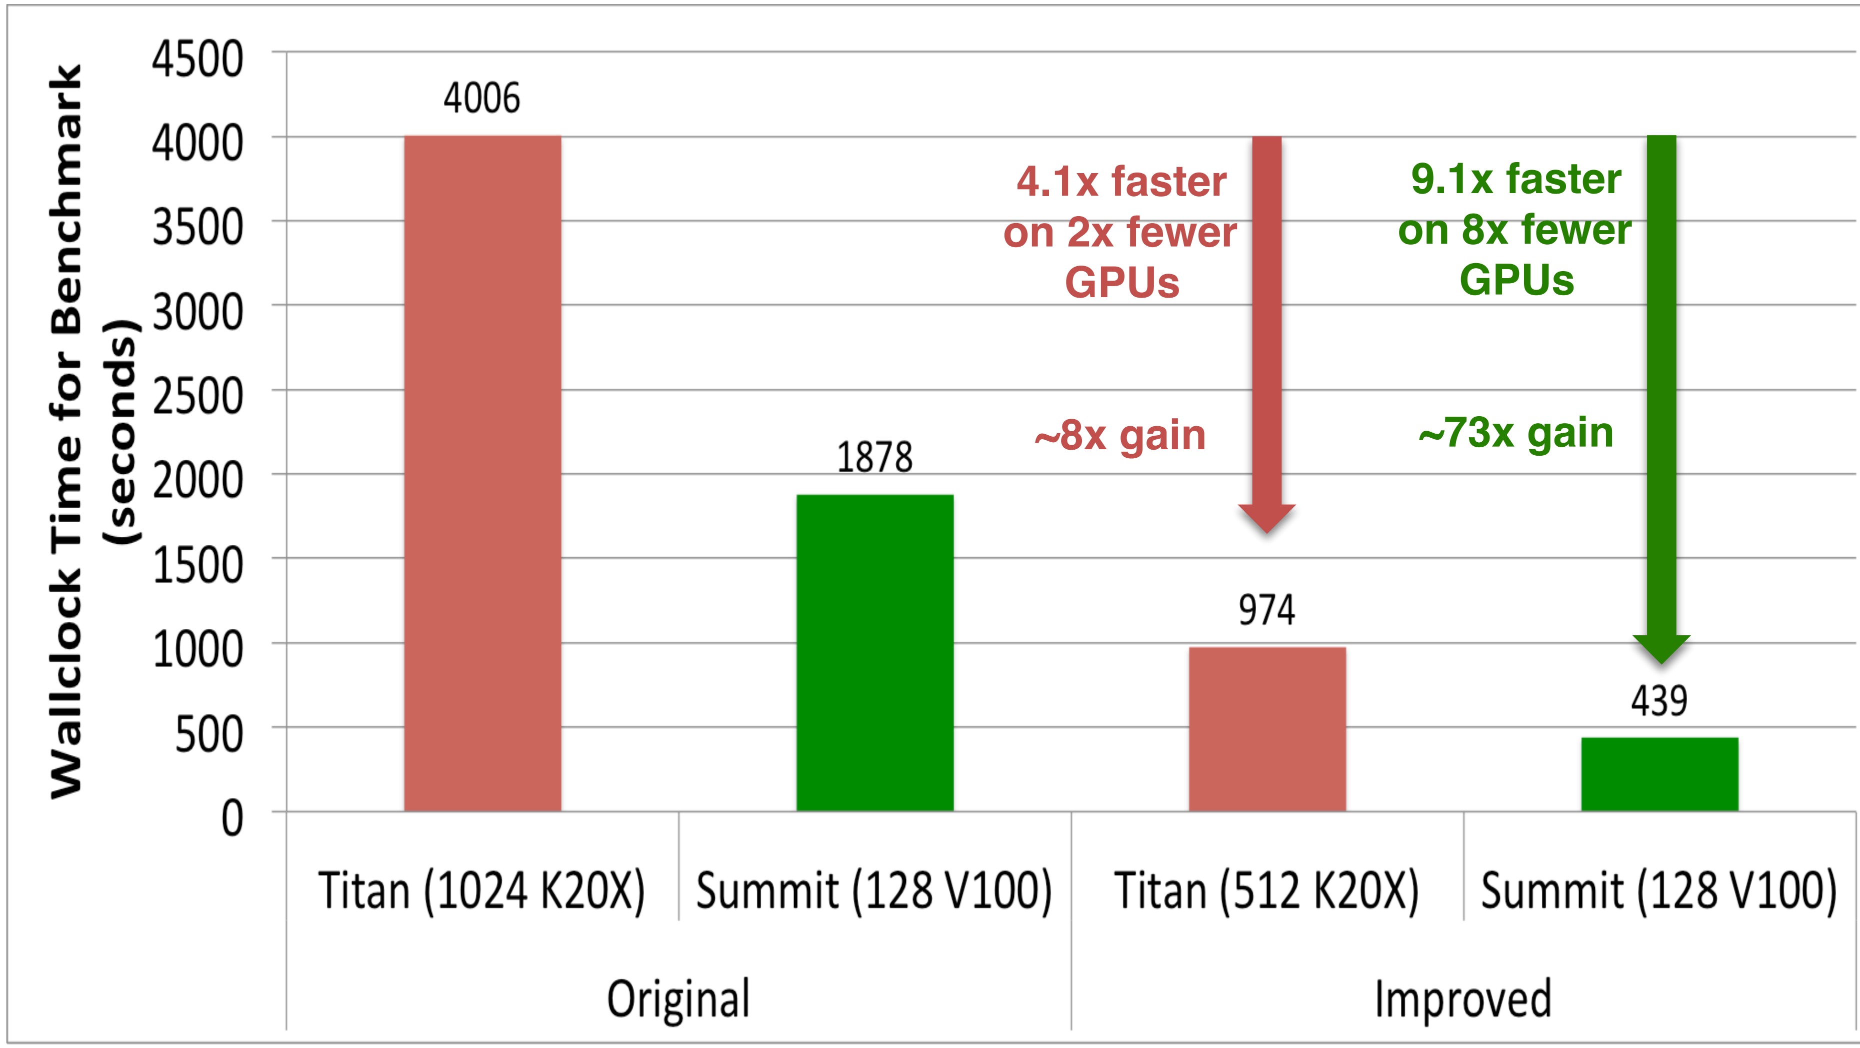

- ~9x wallclock speed-up on Summit using 8x fewer GPUs than Titan: **~73x improvement in computational efficiency**
- **Allows previously unaffordable calculations**

To be OLCF highlight

QUARC framework

ASCR support at Univ. North Carolina

- Alternative to "metal" programming

- C++14 and LLVM-based EDSL framework
- High-level array expressions (METL)
- *QOPT*: Domain-specific compiler
- *Automated code generation*
  - Data layout transformations for SIMD vectorization
    - *Vector folding as in QPhiX, GRID*
  - MPI partitioning, communication, and compute overlapping
    - *Using polyhedral analysis*
  - High-level optimizations
    - *Loop fusion, e.g., multi-RHS*

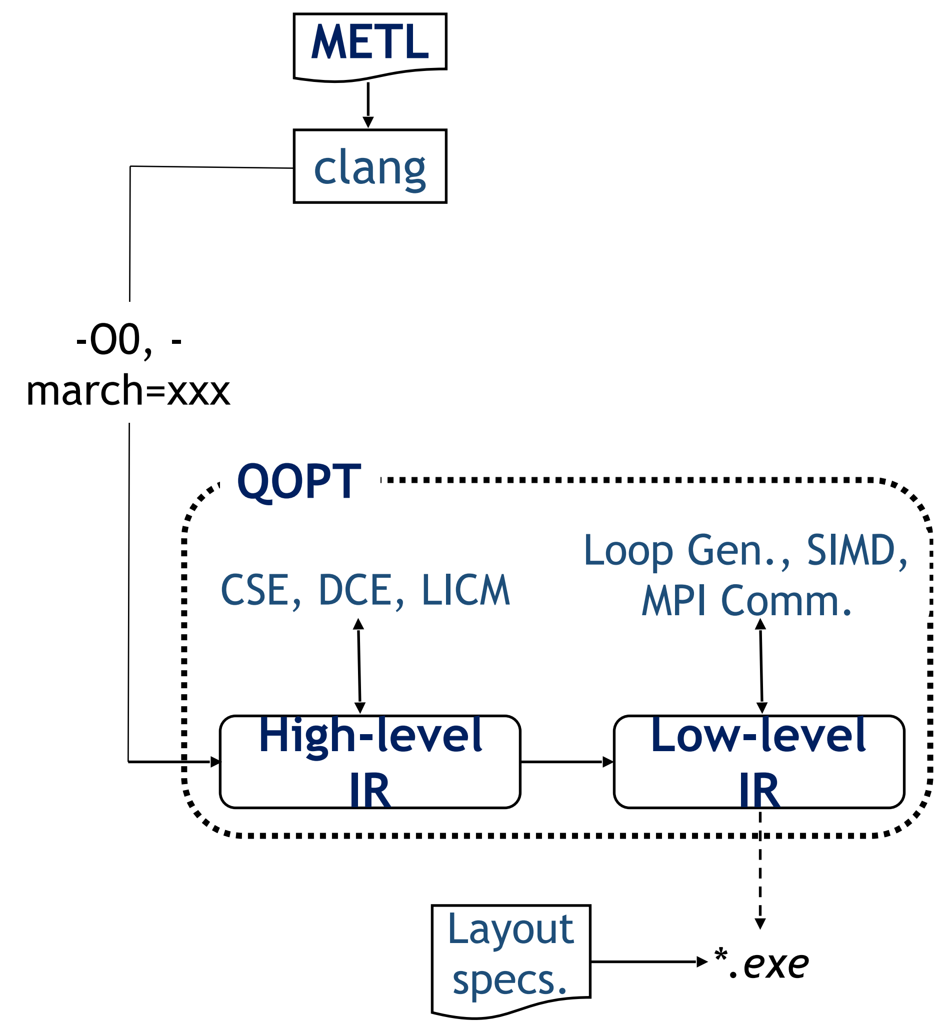

QUARC performance

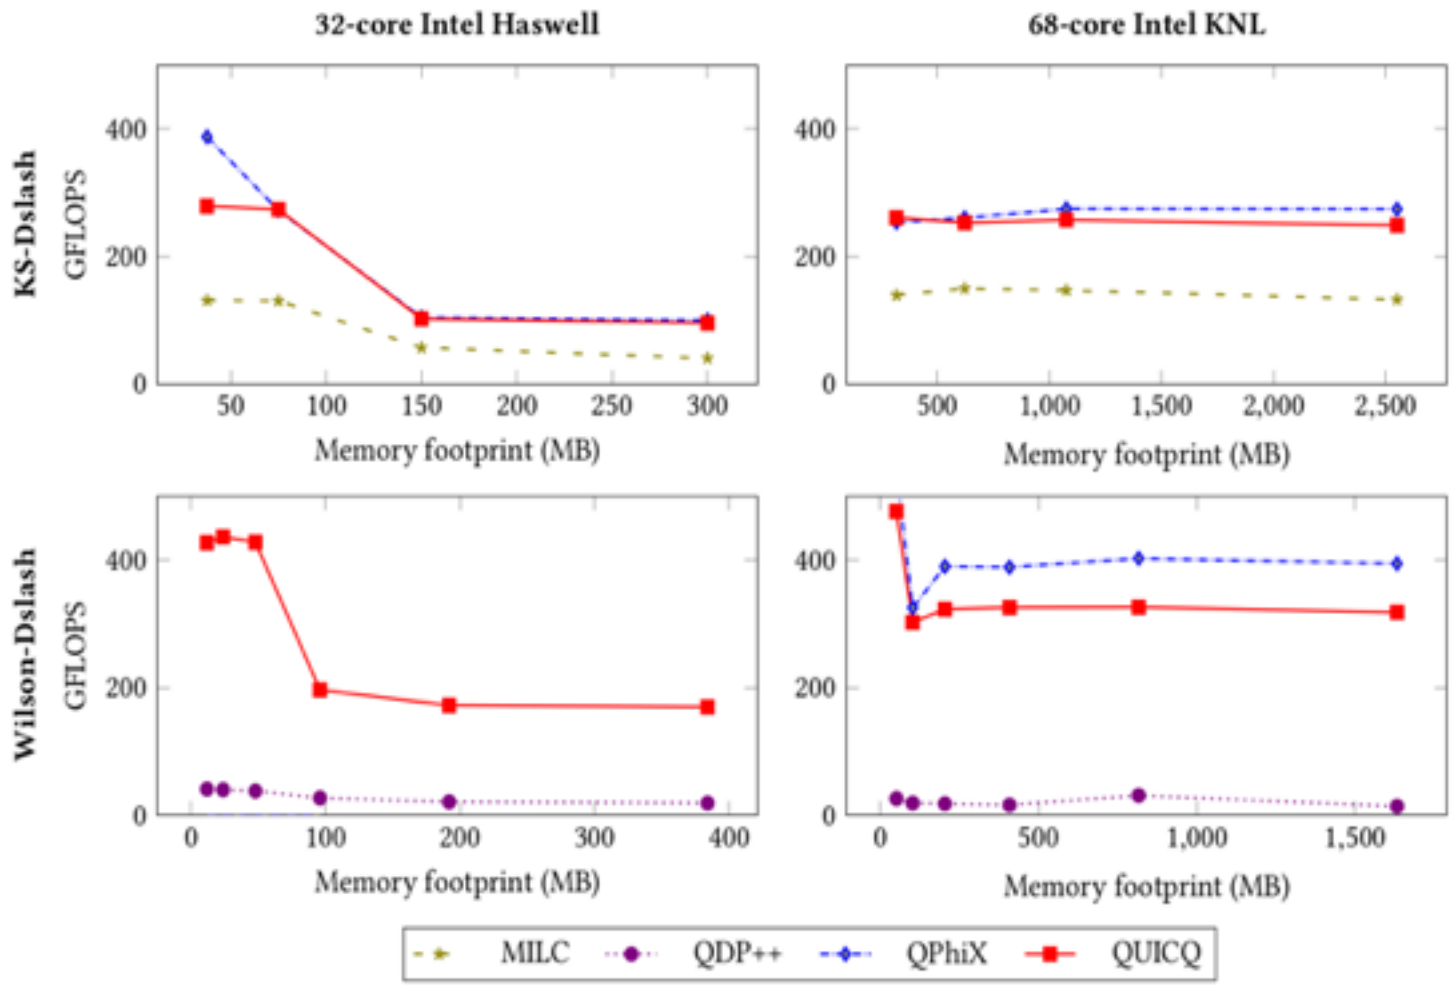

- QUICQ is a prototype EDSL using QUARC
- 10x performance gain over QDP++ (Wilson Dslash)
- 2x performance gain over MILC (KS-Dslash)
- Within 15% of QPhiX on KNL (thread co-location issue)

Next: Improved "workflow"

Gauge generation
- Leadership level
- Strong scaling

Strong scaling limited
- Comms/compute less balanced in recent machines
- Latency in comms important
- Require comms reduced linear solvers
- 1-flavor solvers

Sympletic PDE integrators

Data-parallel code gen. with comms

Analysis
- Throughput level
- O(1M) RHS-s/cfg

Throughput challenge
- Solvers still important
- Large problems / node -> minimize memory traffic
- Job coordination -> Grid based

Contractions

Correlators
- O(1M)/cfg

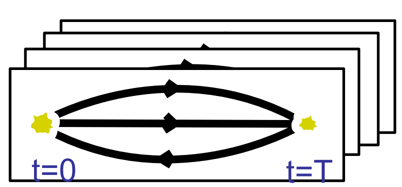

"Workflow" challenge
- Combinatorics -> improved graph theoretic methods
- Code generation for (sparse?) tensor contractions
- Job coordination -> Grid based

PanDA Work Management System

- Developed for ATLAS expt. @ CERN LHC
- Currently processing 1M jobs/day for ATLAS. In use on Titan
- SciDAC-4: "Harvester" - new resource-facing service to Grids, clusters & clouds
- HEP/NP/ASCR partnership [See poster by Sergey Panitkin]

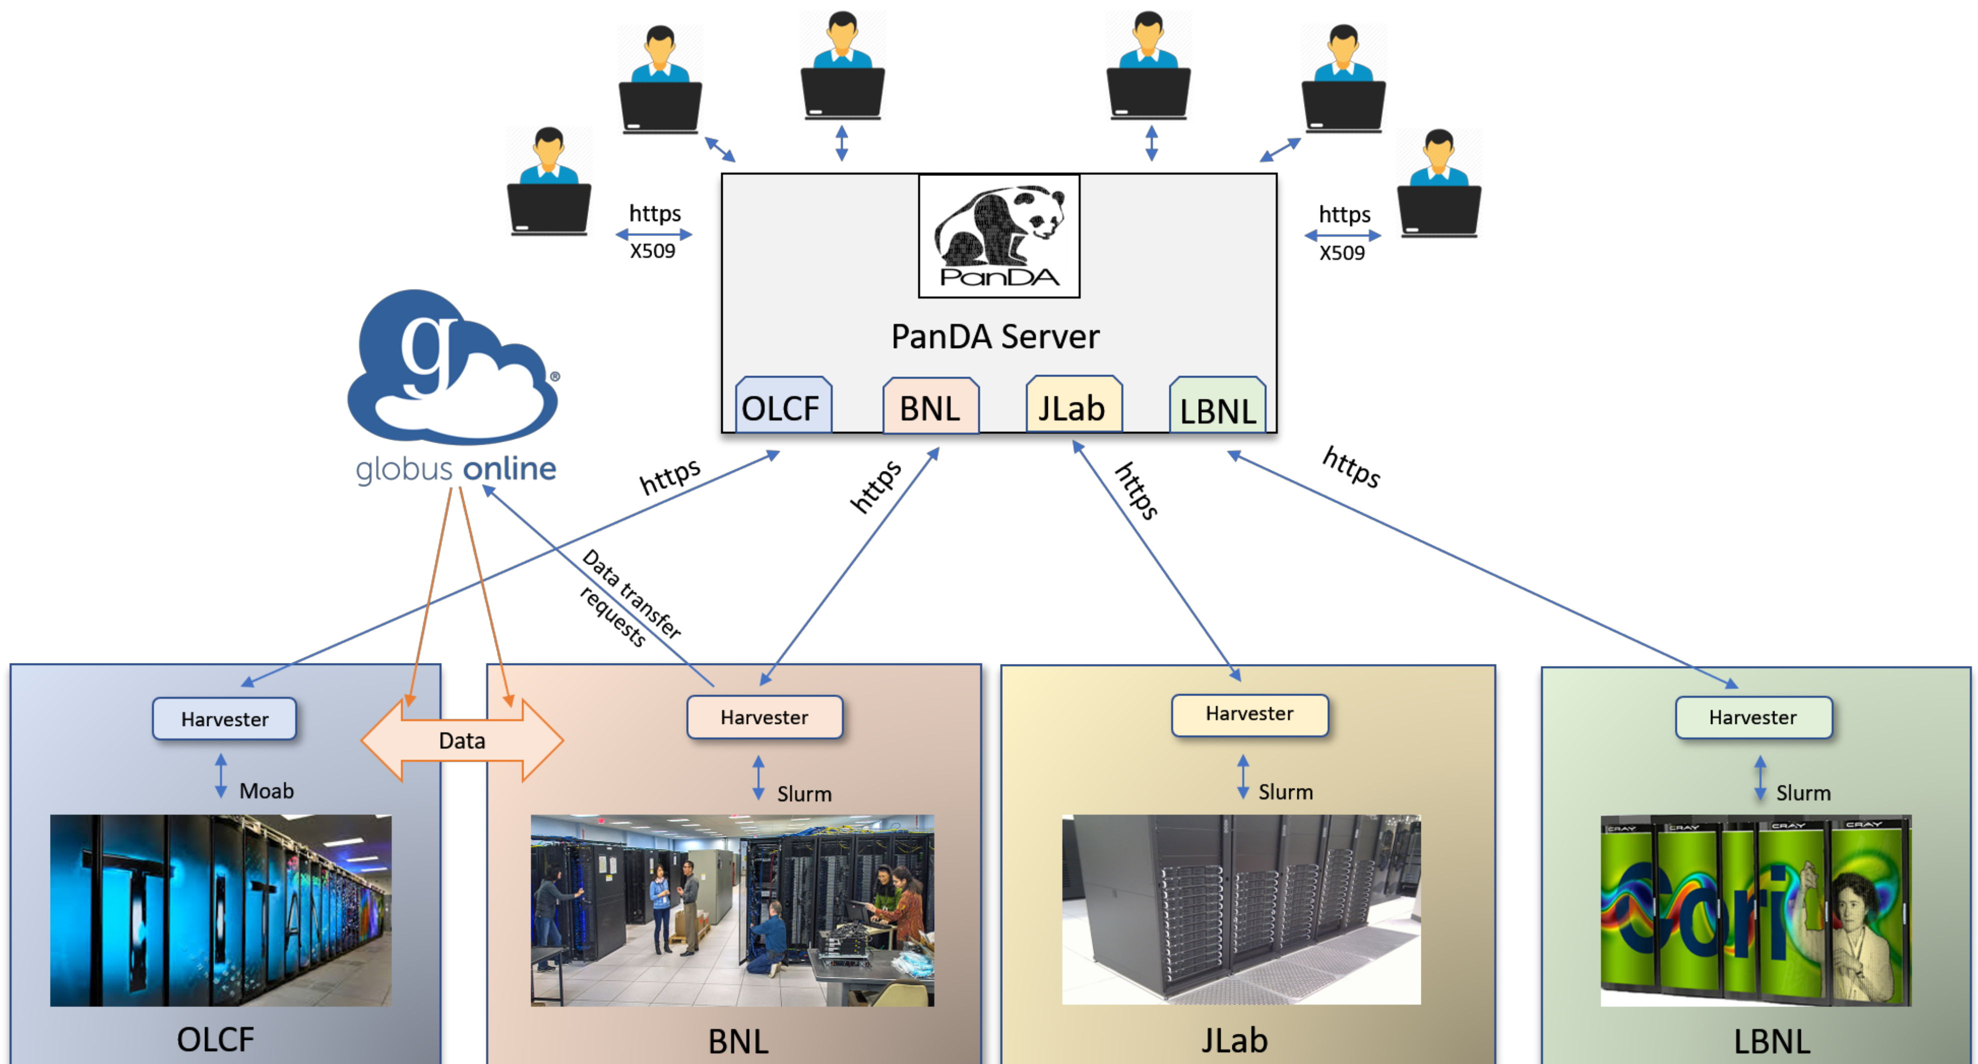

Reducing contraction/correlator costs

Gauge generation
- Leadership level
- Strong scaling

Strong scaling limited
- Comms/compute less balanced in recent machines
- Latency in comms important
- Require comms reduced linear solvers
- 1-flavor solvers

Sympletic PDE integrators

Data-parallel code gen. with comms

Analysis
- Throughput level
- O(1M) RHS-s/cfg

Throughput challenge
- Solvers still important
- Large problems / node -> minimize memory traffic
- Job coordination -> Grid based

Contractions

Correlators
- O(1M)/cfg

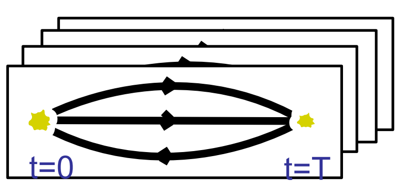

"Workflow" challenge
- Combinatorics -> improved graph theoretic methods
- Code generation for (sparse?) tensor contractions
- Job coordination -> Grid based

QCD for nuclei

- Quarks need to be tied together in all possible way

$N_{contractions} = N_u! N_d! N_s!$ $\quad (\sim 10^{1500}$ for $^{208}Pb)$

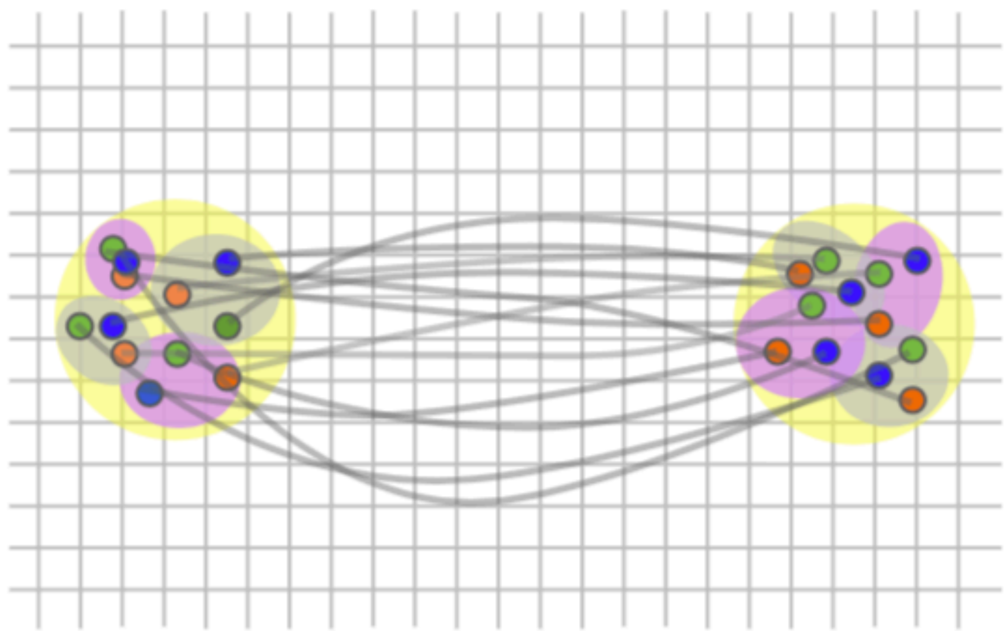

- Manage using algorithmic trickery - still significant graph contractions
- TACO @ MIT (cf TCE, Cyclops) http://tensor-compiler.org
  - general purpose tensor contraction compiler framework
  - target to QCD specific problems (sparsity patterns, …)
  - Allow efficient exploration of methodology

Other activities

- Kokkos and KokkosKernels
  - Kokkos team & Siva Rajamanickam (Kernels)

- Autotuning
  - Useful in several physics stages: Petabricks

- Opentuner
  - USQCD software stack

SciDAC-4 is a partnership

- Strong existing partnerships with ASCR & HEP community members
- Developing stronger ties with SciDAC institutes

- Already a big boost for LQCD & NP
  - fresh perspective on entire program

- Is significantly impacting our calculations
  - huge advance in gauge generation - accelerated our science campaigns
  - workflows optimally use available resources - lower bar for new researchers
  - better tools for (more easily) improving performance

SciDAC
Scientific Discovery through Advanced Computing