LQCD:
Particle physics from a supercomputer

October 12, 2022 | Eric B. Gregory | FZJ

OVERVIEW

- A little bit of physics
- LQCD basics
- LQCD data challenges
- IO-SEA solutions

SOME QUICK PHYSICS

LQCD: Lattice quantum-chromodynamics

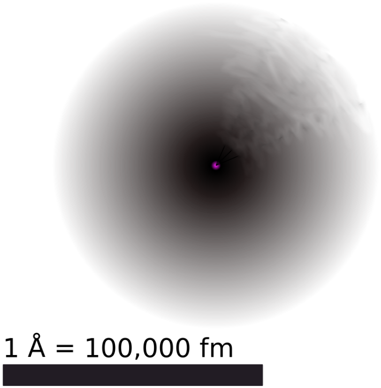

SOME QUICK PHYSICS

LQCD: Lattice quantum-chromodynamics

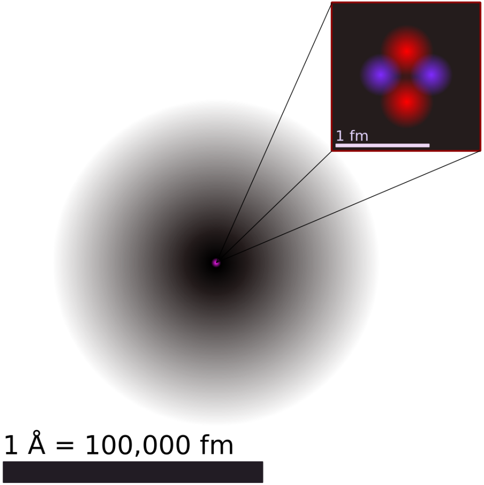

QCD

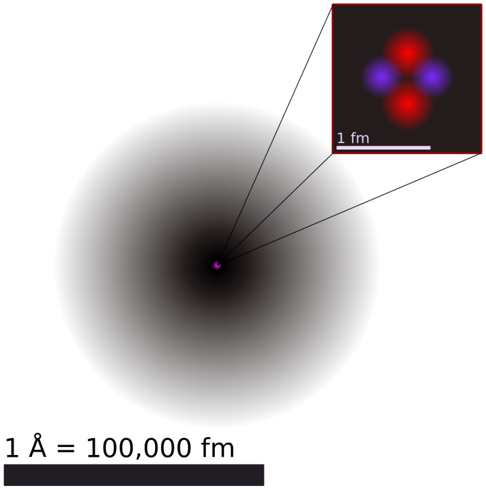

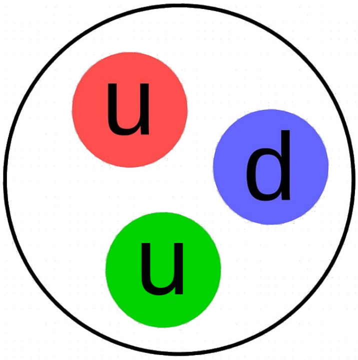

QCD

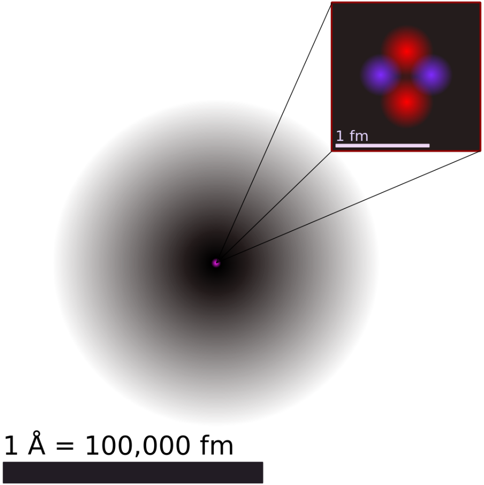

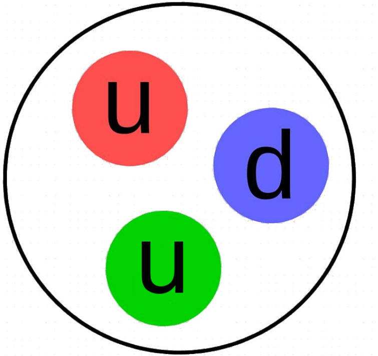

Proton is a *hadron*, a particle made of quarks bound together by the strong force.

QCD

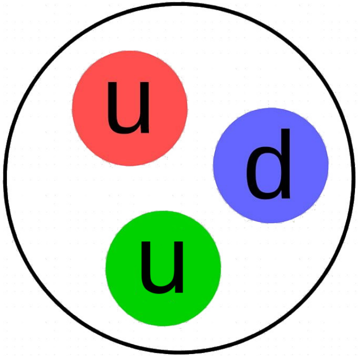

We say protons have the *quantum numbers* of three quarks:

QCD

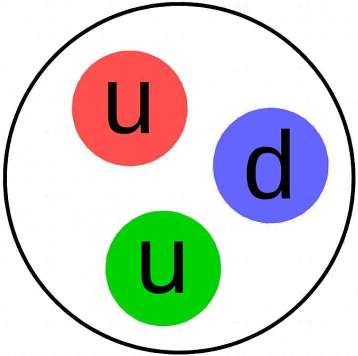

We say protons have the *quantum numbers* of three quarks:

Quarks have a color charge, so-called because:

$$\text{red} + \text{blue} + \text{green} = \text{neutral}$$

QCD

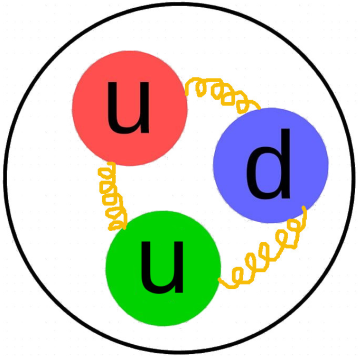

We say protons have the *quantum numbers* of three quarks....

Quarks have a color charge, so-called because:

red+blue+green= neutral

Hadrons also have gluons contributing to their properties.

QCD

We say protons have the *quantum numbers* of three quarks....

Quarks have a color charge, so-called because

red+blue+green= neutral

Hadrons also have gluons contributing to their properties.

For a complete understanding of a hadron, we must understand the quantum field fluctuations.

QUANTUM FLUCTUATIONS ARE IMPORTANT!

QUANTUM FLUCTUATIONS ARE IMPORTANT!

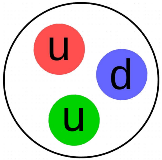

QUANTUM FLUCTUATIONS ARE IMPORTANT!

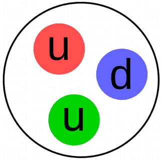

$2 \times M_{\mathrm{up}}$ $+M_{\mathrm{down}}$

QUANTUM FLUCTUATIONS ARE IMPORTANT!

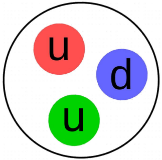

$$2 \times M_{\mathrm{up}} \quad + M_{\mathrm{down}}$$
$$2 \times (2.2\ \mathrm{MeV}) \quad + (4.7\ \mathrm{MeV}) \quad \approx 9\ \mathrm{MeV}$$

QUANTUM FLUCTUATIONS ARE IMPORTANT!

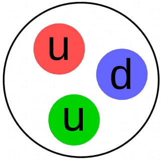

$$2 \times M_{\mathrm{up}} \quad + M_{\mathrm{down}}$$

$$2 \times (2.2\ \mathrm{MeV}) \quad + (4.7\ \mathrm{MeV}) \quad \approx 9\ \mathrm{MeV}$$

But ...

$$M_{\mathrm{proton}} = 938\ \mathrm{MeV}$$

QUANTUM FLUCTUATIONS ARE IMPORTANT!

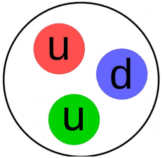

$$2 \times M_{\mathrm{up}} \quad + M_{\mathrm{down}}$$

$$2 \times (2.2\ \mathrm{MeV}) \quad + (4.7\ \mathrm{MeV}) \quad \approx 9\ \mathrm{MeV}$$

But ...

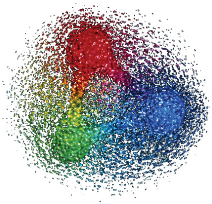

$$M_{\mathrm{proton}} = 938\ \mathrm{MeV}$$

QUANTUM FLUCTUATIONS ARE IMPORTANT!

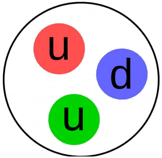

$$2 \times M_{\mathrm{up}} \quad + M_{\mathrm{down}}$$
$$2 \times (2.2\ \mathrm{MeV}) \quad + (4.7\ \mathrm{MeV}) \quad \approx 9\ \mathrm{MeV}$$

But ...

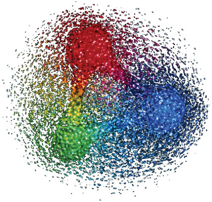

$$M_{\mathrm{proton}} = 938\ \mathrm{MeV}$$

To understand properties of hadrons, we must take quantum fluctuations into effect.

WHAT CAN WE HOPE TO CALCULATE?

WHAT CAN WE HOPE TO CALCULATE?

- Properties of hadrons
  - mass
  - internal structure
  - decay probabilities
  - ...

WHAT CAN WE HOPE TO CALCULATE?

- Properties of hadrons
  - mass
  - internal structure
  - decay probabilities
  - ...
- Existence of unobserved states

WHAT CAN WE HOPE TO CALCULATE?

- **Properties of hadrons**
  - **mass**
  - **internal structure**
  - **decay probabilities**
  - **...**
- **Existence of unobserved states**
- **BIG QUESTION:**

Does

$$\{\text{experiment}\} - \{\text{theory}\} \stackrel{?}{=} 0$$

WHAT CAN WE HOPE TO CALCULATE?

- **Properties of hadrons**
  - **mass**
  - **internal structure**
  - **decay probabilities**
  - **...**
- **Existence of unobserved states**
- **BIG QUESTION:**

Does

$$\{\text{experiment}\} - \{\text{theory}\} \stackrel{?}{=} 0$$

Physics beyond the Standard Model?

QCD ON THE LATTICE

Model is discretized 4-D box of (Euclidean) space-time:

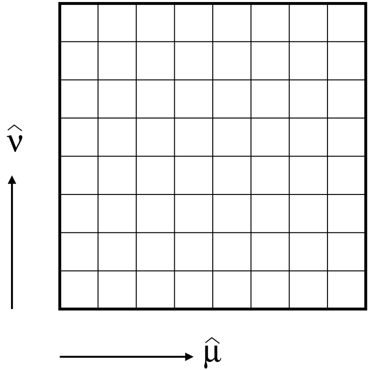

QCD ON THE LATTICE

Model is discretized 4-D box of (Euclidean) space-time:

- Quark fields $\phi(x)$ live on lattice sites
  3 (or 3×4)-component, complex:

$$\phi(x) = \begin{pmatrix} \phi_0 \\ \phi_1 \\ \phi_2 \end{pmatrix}$$

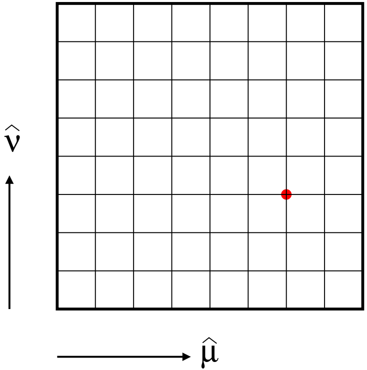

QCD ON THE LATTICE

Model is discretized 4-D box of (Euclidean) space-time:

- Quark fields $\phi(x)$ live on lattice sites
  3 (or 3×4)-component, complex:

$$\phi(x) = \begin{pmatrix} \phi_0 \\ \phi_1 \\ \phi_2 \end{pmatrix}$$

- Gauge fields $U_\mu(x) = \exp(iagA(x))$
  live on links: (3×3)-component, complex

$$U_\mu(x) = \begin{pmatrix} U_{00} & U_{01} & U_{02} \\ U_{10} & U_{11} & U_{12} \\ U_{20} & U_{21} & U_{22} \end{pmatrix}$$

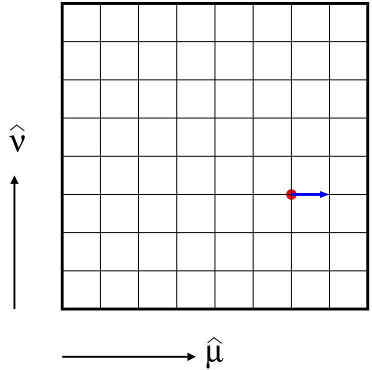

LQCD WORKFLOW

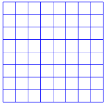

- Generate Markov chain of lattice gauge field configurations

LQCD WORKFLOW

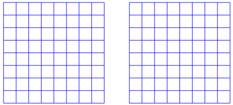

- Generate Markov chain of lattice gauge field configurations

LQCD WORKFLOW

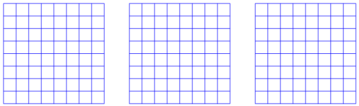

- Generate Markov chain of lattice gauge field configurations

LQCD WORKFLOW

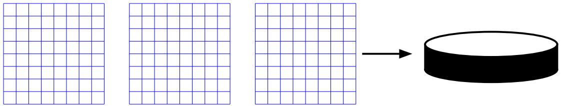

- Generate Markov chain of lattice gauge field configurations
- Save each to disk

LQCD WORKFLOW

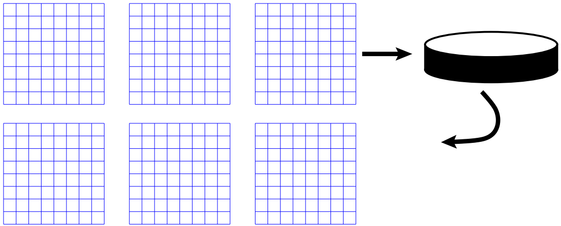

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk

LQCD WORKFLOW

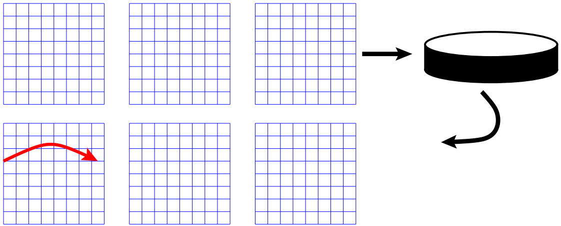

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$

LQCD WORKFLOW

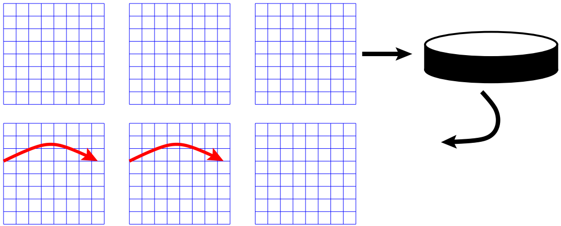

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$

LQCD WORKFLOW

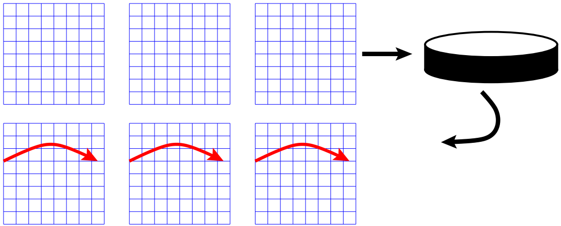

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$

LQCD WORKFLOW

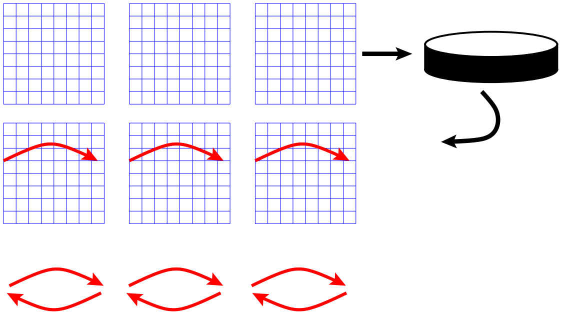

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$
- Contract propagators to generate hadron correlators $C(t)$

LQCD WORKFLOW

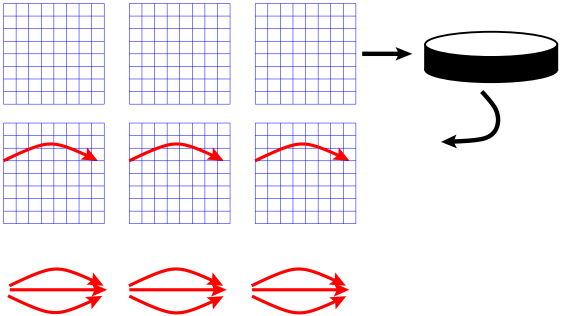

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$
- Contract propagators to generate hadron correlators $C(t)$

LQCD WORKFLOW

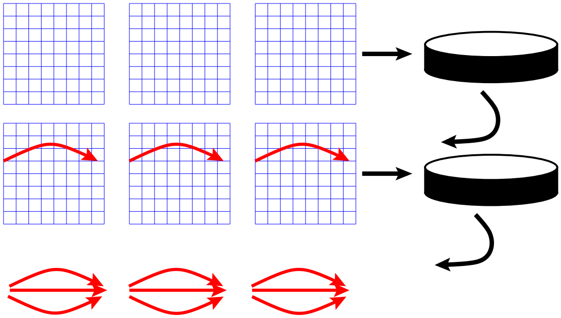

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$
- Contract propagators to generate hadron correlators $C(t)$

LQCD WORKFLOW

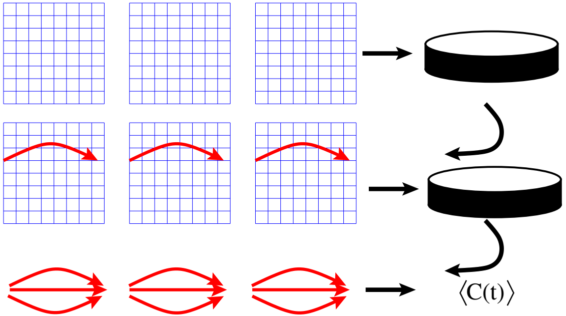

- Generate Markov chain of lattice gauge field configurations
- Save each to disk
- Load configuration files from disk
- On each, calculate quark propagators: $p(x) = M^{-1}(x, y)q(y)$
- Contract propagators to generate hadron correlators $C(t)$
- Average over ensemble gives expectation value
- Fit correlators to extract physical quantities, e.g., hadron masses

LQCD I/O AND DATA CHALLENGES

- Bigger lattice size: $V = N_x^3 \times N_t$ $\longrightarrow$ smaller systematic uncertainties
- More lattice gauge field configurations N_{conf} $\longrightarrow$ smaller statistical uncertianties
- More points in parameter space (m_q, β) $\longrightarrow$ smaller systematic uncertanties

- LQCD is a very homogeneous problem with many available levels of concurrencies.
- Algorithms are highly scalable.

LQCD I/O AND DATA CHALLENGES

- Bigger lattice size: $V = N_x^3 \times N_t$ $\longrightarrow$ smaller systematic uncertainties
- More lattice gauge field configurations N_{conf} $\longrightarrow$ smaller statistical uncertainties
- More points in parameter space (m_q, β) $\longrightarrow$ smaller systematic uncertanties

- LQCD is a very homogeneous problem with many available levels of concurrencies.
- Algorithms are highly scalable.

If given the opportunity, LQCD practitioners will (easily) fill up your machine in search of more precise/accurate results.

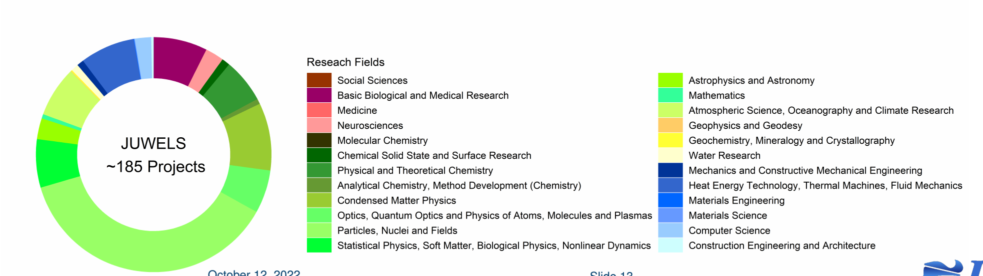

LQCD I/O AND DATA CHALLENGES

Lattice size $V = N_x^3 \times N_t$

- **HMC step**

  A lattice gauge configuration file has:

  $$V \times \underset{\text{dimensions}}{\underset{\uparrow}{4}} \times (\underset{\text{color}}{\underset{\uparrow}{3 \times 3}}) \times \underset{\text{complex}}{\underset{\uparrow}{2}} \text{ real numbers (72 reals /site).}$$

  For $V = 64^3 \times 128$, double precision: 2.4×10^9 doubles; $\sim$ 20 GB/file

  Few $\times 10^3$ gauge configuration files per ensemble

- **Spectrum step**

  A propagator (or eigenvector) file has

  $$V \times \underset{\text{spin}}{\underset{\uparrow}{4}} \times \underset{\text{color}}{\underset{\uparrow}{3}} \times \underset{\text{spin}}{\underset{\uparrow}{4}} \times \underset{\text{color}}{\underset{\uparrow}{3}} \times \underset{\text{complex}}{\underset{\uparrow}{2}} \text{ real numbers (288 reals /site)}$$

  For $V = 64^3 \times 128$, double precision: 9.7×10^9 doubles; $\sim$ 80 GB/file.

  $\sim$ 10 — 100 per gauge configuration.

LQCD IN THE IO-SEA ENVIRONMENT

Data workflow on a traditional storage hierarchy

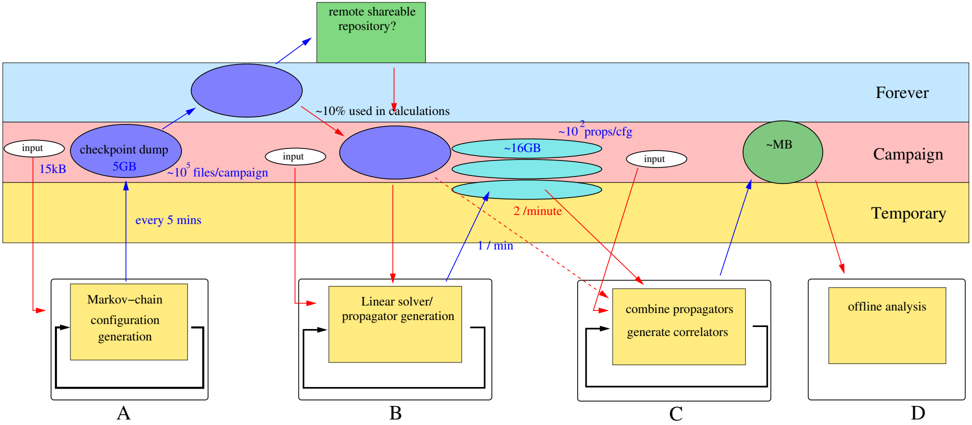

LQCD IO-SEA DATASETS

Group data in data sets (one possible solution):

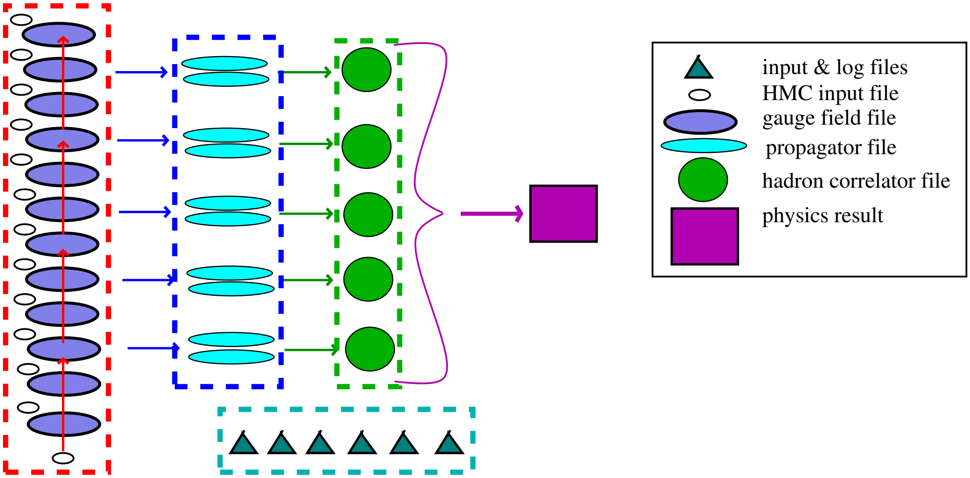

LQCD EPHEMERAL SERVICES

POSIX NFS service

- Requires no re-coding of the LQCD I/O library
- Primary service to be used with LQCD application

BB-NFS

- May provide benefit; will test on prototype

DASI interface

- Novel system of organizing data
- Possibly useful in organizing a sharable archive of lattice gauge field configurations
- Semantic keys include coupling constant β, quark masses m_q, lattice volume,...
- Would only use a stand-alone application, rather than re-coding

LQCD IN THE IO-SEA ENVIRONMENT

The “Workflow Description File”:

```
# lqcd-BC-workflow.yaml
workflow:
        - name: LQCD_BC

services:
        - name:  nfs-gauge-files
          type: NFS
          attributes:
            namespace: "beta_{{␣BETA␣}}_gauge_cfg_files"
            mountpoint:"/mnt/USER/gauge-fields/B{{␣BETA␣}}/"
            flavor: medium

        - name: nfs_propagators
          type: NFS
          attributes:
            namespace: "propagators-b{{␣BETA␣}}"
            mountpoint: "/mnt/USER/propagators/B{{␣BETA␣}}/"
            flavor: medium

        - name: nfs_input-and-logs
          type: NFS
          attributes:
            namespace: "input-and-logs-b{{␣BETA␣}}"
            mountpoint: "/mnt/USER/input-and-logs/B{{␣BETA␣}}/"
            flavor: medium

        - name: nfs_hadron_correlators
          type: NFS
          attributes:
            namespace: "hadron-correlators-b{{␣BETA␣}}"
            mountpoint: "/mnt/USER/corrs/B{{␣BETA␣}}/"
            flavor: medium
```

```
steps:
        - name: "step_B_props-b{{␣BETA␣}}"
          location: gpu_module      # generate props on GPUs
          command: "sbatch␣␣{{␣␣PROP_BATCH_SCRIPT␣}}"
          services:
                - "nfs-gauge-files" RW
                - "nfs_propagators" RW
        - "nfs_input-and-logs" RW

        - name: "step_C_contractions-b{{␣BETA␣}}"
          location: cpu_module
# hadron correlator contractions on CPUs
          command: "sbatch␣make_had_corrs-b{{␣BETA␣}}.sh␣--export=id={{␣ID␣}}
          services:
                - "nfs_propagators" RO
                - "nfs_hadron_correlators" WO
                - "nfs_input-and-logs" WO
```

LQCQ IN THE IO-SEA ENVIRONMENT

```
#!/bin/bash
beta=3.6

#create namespaces for all hadron correlators:
iosea-ns create --auto-create-dataset hadron-correlators-b${beta}
#create namespace for input and log files for B & C
iosea-ns create --auto-create-dataset input-and-logs-b${beta}
# create empty namespace/dataset for propagators
iosea-ns create --auto-create-dataset propagators-b${beta}

start=500   # first 500 gauge files are not thermalized
skip=20        # we think it takes 20 updates for gauge fields to de-correlate
for (( num=start; num<=2000; num+=skip ))
do
    # also add input files for B and C to input-and-logs NS:
    iosea-ns put input-and-logs-b${beta}  input_propsolve_beta_b${beta}_${num}.xml
    iosea-ns put input-and-logs-b${beta}  input_had-corrs_beta_${beta}_${num}.xml
done

# start session
#session name: props-corrs-b3.6
iosea-wf start WORKFLOW=lqcd-BC-workflow-b${beta}.yaml SESSION=prop-corrs_b${beta}  BETA=$beta
```

LQCD IN THE IO-SEA ENVIRONMENT

```
#run B step
for (( num=start; num<=2000; num+=skip ))
do
    iosea-wf run  SESSION=prop-corrs_b${beta} STEP=step_B_props-b${beta} \
             PROP_BATCH_SCRIPT="solve_propagators-b${beta}-${num}.sh"
done

#check status of jobs:
iosea-wf status SESSION=prop-corrs_b${beta}

#run the C step
for (( num=500; num<=2000; num+=skip ))
do
    iosea-wf run SESSION=prop-corrs_b${beta} STEP=step_C_contractions-b${beta}} ID=$num
done

#check status of jobs
iosea-wf status SESSION=prop-corrs_b${beta}

#stop the session and release the datanode
iosea-wf stop SESSION=prop-corrs_b${beta}
```

SUMMARY & OUTLOOK

- I/O challenges will increase as faster HPC platforms & algorithms allow larger problems to be tackled.
- IO-SEA ephemeral services should improve LQCD I/O bottlenecks and allow us to get the most from compute time allocations.
- Testing on prototypes to begin in the coming months.

LQCD COMMUNITY SOFTWARE

$\longleftarrow$ different physics goals $\longrightarrow$

$\longleftarrow$ different simulation codes $\longrightarrow$

$\longleftarrow$ different architectures $\longrightarrow$

LQCD COMMUNITY SOFTWARE

⟵ different physics goals ⟶

⟵ different simulation codes ⟶

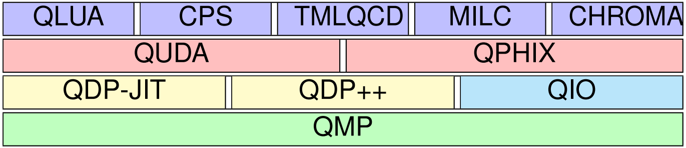

⟵ different architectures ⟶