DataLake Applications: Lattice QCD

Gaurav Ray
(IFCA-CSIC)

Outline

- Lattice QCD at a glance
- Lattice Data Management
- Lattice and the interTwin DataLake
  - Experience
  - Rucio
  - Plans
- The DataLake and the ILDG
- Wrap-up

This talk recounts our efforts over the last year or so and is my attempt at thinking through the DataLake's potential effect on scientific data management from the scientist's perspective.

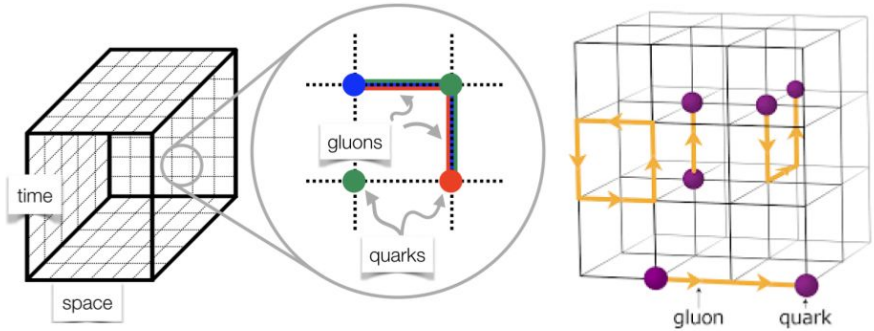

$$\langle 0|\mathcal{O}|0\rangle = \frac{\int [d\psi]\,[d\bar{\psi}]\,[dA_\mu]\,\mathcal{O}[\psi,\overline{\psi},A]e^{-S}}{\int [d\psi]\,[d\bar{\psi}]\,[dA_\mu]e^{-S_{QCD}}}.$$

- Probe the Standard Model of physics by numerically solving the field equations.
- Well-established and precise/well-controlled results.
- Only game in town in many cases.
- Lots of groups working on many different aspects all over the world.

- Lattice QCD simulations are executed at scale on HPC systems that are controlled by a batch system.
- A typical workflow involves the generation of up to O(1000) lattice field configurations,
  - the measurement of an observable of interest over those configurations,
  - and the statistical analysis of those measurements.

All of these steps, **especially the generation of configurations**, can be highly computationally intensive.

We almost always want to **keep these lattice ensembles**.

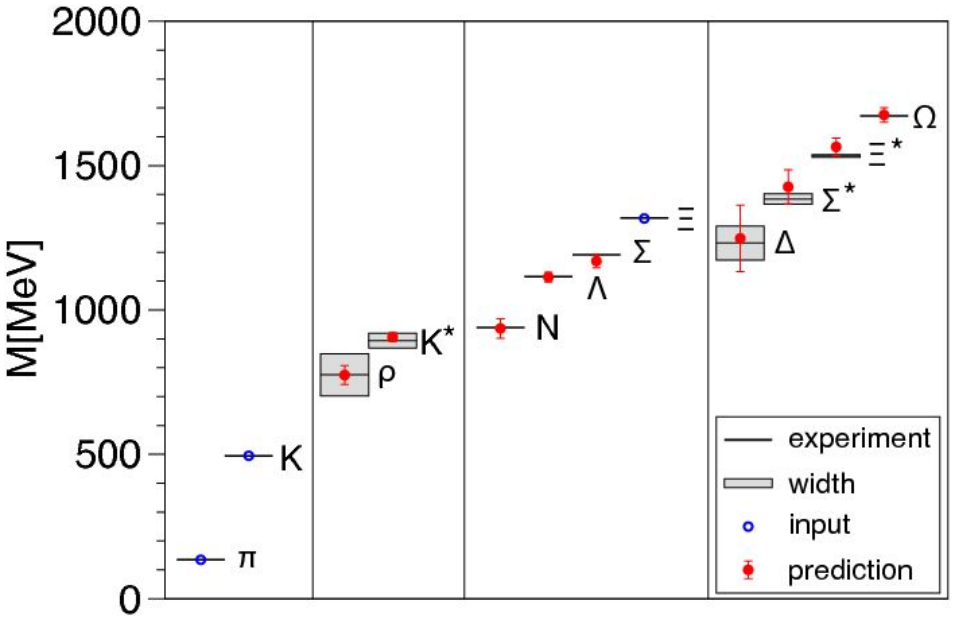

Bigger Lattices and Bigger Collaborations

- Since the advent of the field in the 80s typical lattice sizes have steadily grown.
- Individual lattices can now be as large as 100GB. Ensembles O(100TB).
- As in other fields, there has been a concerted push for the integration of FAIR principles into Lattice workflows.

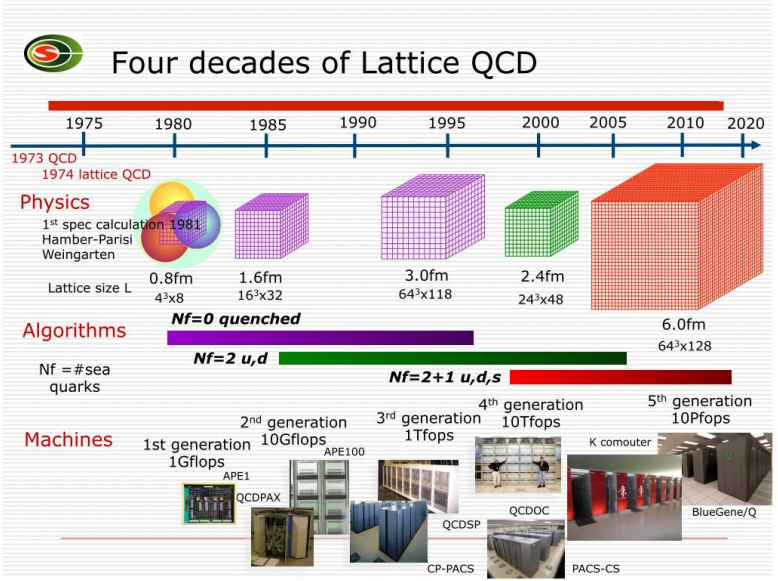

Credit: Akira Ukawa (2013)

How and where the increasingly large amount of lattice data will be stored is a related but separate problem. For now let's assume we have enough space.

A simplified typical Lattice workflow

1. Data is generated at HPC facility
2. Data is transferred out of HPC facility
3. Analysis of data takes place
4. Data is copied to long-term storage

Data Management within the Lattice community

- Each lattice collaboration forms its own policies wrt its own data and is bound by the terms of its funding.
- Wide variation: Some collaborations have developed sophisticated data tracking policies with info-rich websites. Others rely on individual members for access to data.
- As collaborations grow/shrink and personnel leave/are replaced data is liable to be forgotten. Sometimes by the time someone remembers it is too late!

Enter: The interTwin DataLake

- The interTwin DataLake is managed by the Rucio software, which ensures scalable and efficient data transfer and storage.
- Based off the ESCAPE DataLake which was developed at CERN by and for the experimental physics community.
- Each DataLake is formed from a network of storage endpoints (RSE) orchestrated via a Rucio Server.
- This work is being carried out as part of the interTwin project as it should help researchers create Digital Twins.

For more technical information about the DataLake please see Dijana's presentation

- Start small: We transferred sample lattice data and toy data between the VEGA and DESY RSEs. O(10GB)
- Needed to familiarise ourselves with the Rucio client. The Rucio CLI is fairly straightforward.
- Rucio Replication Rules are an alternative way of thinking about data. Takes a bit of getting used to.

```
$ rucio whoami

$ rucio list-rses

$ rucio upload --rse DESY-DCACHE --scope gsinharay test_lattice

$ rucio add-rule gsinharay:test_lattice 1 VEGA-DCACHE

$ rucio get gsinharay:test_lattice

$  rucio attach gsinharay:test_ensemble gsinharay:test_lattice
```

Mapping Rucio permissions on to Lattice collaborations

"Rucio assigns permissions to accounts. Permissions are boolean flags designating whether an account may perform a certain action (read, write, delete) on a resource (RSE, account, replica, etc.)." - Rucio docs

- Some collaborations may not want to give every member of the group read **and** write permission.
- Write permission restricted to a smaller group.
- Embargoed and blinded data is increasingly more common.

A ‘Lattice’ DataLake

- DESY has set up a ‘LatticeQCD’ data lake with
  - the latest version of Rucio,
  - using the public FTS server,
  - and ILDG's Indigo IAM for identity management.
- Indigo IAM as authentication is a welcome improvement over the prior EGI Check-In.
- At the moment, the data lake doesn't have any RSEs, users, or data.
- Future: Someone from the lattice community will need to be trained to administer a lattice data lake.

External Catalogues - The ILDG

- The **International Lattice Data Grid** is an organisation, founded in 2002 and based at DESY that hosts Lattice data and Metadata for the community.
- It is split into autonomous regional grids (LDG, JDG,...) (but with an ILDG-wide User Registration).
- Each Regional Grid provides Metadata and File catalogue and storage services. These exist to autonomously manage the data in that region, independently of other Regional Grids.
- It has defined a standard for lattice metadata and is responsible for updating it.

ILDG-DataLake integration

The ILDG can support multiple locations for each file. Each location is recorded as a URL. DataLake support could be included by letting ILDG accept URLs that look like so:

➢ rucio://rucio.example.org:8443/LatticeQCD/RAY:my_test_file

where

- rucio.example.org identifies the Rucio endpoint,
- 8443 is the port on which Rucio is listening,
- LatticeQCD is the name of the Virtual Organisation, and
- RAY:my_test_file is the file's Rucio Data Identifier (RAY is the file's scope and my_test_file is the filename).

Adding (Rucio) Storage Endpoints

- The lattice community is highly international.
- More RSEs with more storage space can only be a good thing (from the selfish POV of a lattice researcher). Network effects important.
- System Administrators should be able to set up the required storage partitions and permissions easily.
- We hope an endpoint will be setup at CESGA soon.

Example use case: CESGA-Zeuthen transfer

- We have a bunch of data at Zeuthen that we want to move to CESGA. O(30TB).
- 1st step - Open a connection via FTS
- 2nd step - Open a connection with Rucio
- 3rd step - Transfer data via the DataLake

Summary

- DataLakes can help the lattice community in our quest for FAIR data. They could be especially useful for researchers who wish to make their data readily available.

- Our experience in interTwin working with the DataLake Developers at DESY has been promising thus far.

- We now have a lattice DataLake though it's not populated with anything. A more involved technical demonstration is just around the corner.