QCDgrid

A Grid for UKQCD

National collaboration for lattice QCD

Chris Maynard

Contents

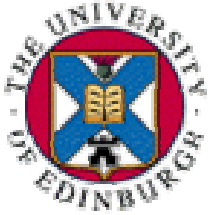

- UKQCD science
  - Bluffers guide to lattice QCD
  - Why UKQCD needs a grid
- QCDgrid
  - QCDgrid architecture
  - Metadata, metadata, metadata
    - XML schema
- Sharing data 

- Putting it all together

Lattice QCD

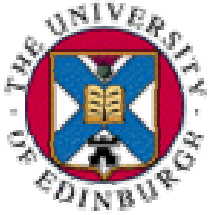

- Quantum Chromodynamics is …
  - Theory of strong interaction
    - Properties of hadrons (proton, pion etc)
    - from quarks and gluons
  - Strongly coupled
    - Perturbation theory only works in a specific regime
- Replace space-time with 4d lattice
- Compute numerically
- QFT → path integral
  - Integrate operators over all quark and gluon fields

Monte Carlo integration

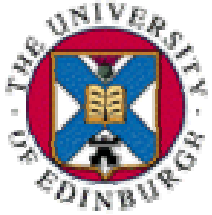

- Generate finite number of configurations of quark and gluon fields
- Probability given by the integrand
  - Importance sampling
    - Only generate configs which contribute
- Integral → finite sum
  - Compute operators on each configuration
  - Average over ensemble
  - Statistical estimate of expectation value

The hard bit

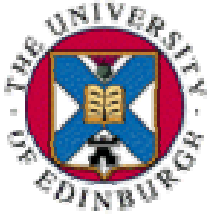

- The fermion matrix (quark-gluon coupling)
  - Proportional to the volume
  - Large and sparse and badly conditioned

$$C_N = \frac{\lambda_{\max}}{\lambda_{\min}} \propto \frac{1}{m_q}$$

- Up and down quark mass nearly zero!
  - Calculate determinant and inverse M
    - Generate Markov Chain
- Requires biggest computer you can find

QCDOC – system on a chip

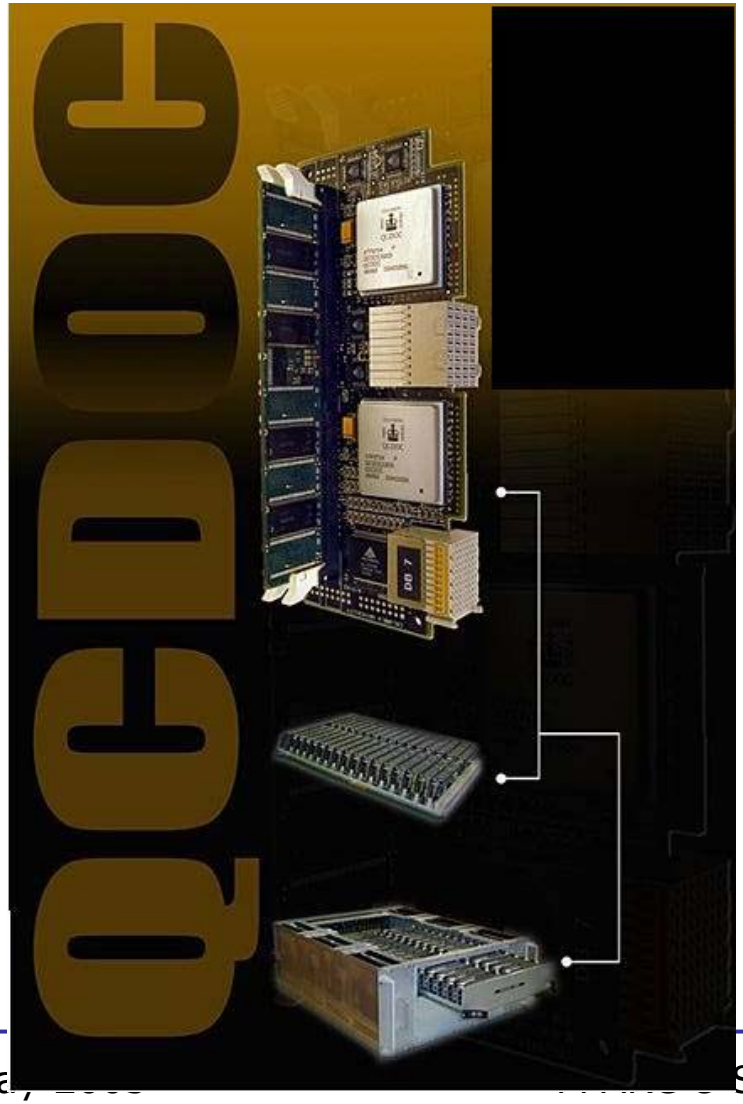

Data and processing

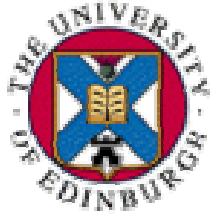

- Many data sets ensembles of hundreds of configurations
- Depend on lattice spacing, quark mass etc
- Very costly to produce
- UKQCD members distributed around UK
- Processing jobs on
  - Workstation, Linux cluster, QCDOC
- → data grid

What do I want from the grid

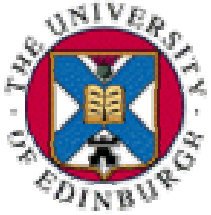

- Data security
  - Replication - Multiple copies of data
  - Who has read/write permission
- Data access
  - Metadata - discover what is stored
  - Get data without having to know location
  - Logistics of data moving taken care of

Machines – LHC terminology

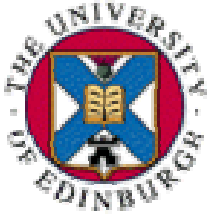

- Tier 1 system
  - 50TByte SAN system
    - QCDOC front end directly connected
  - Hold one copy of all UKQCD data
    - Machine exists, not yet part of grid
- Tier 2 systems
  - Edinburgh, Liverpool, Southampton, Swansea
    - Linux server + RAID array disk
    - Now: 5 TByte. Upgrade to12.5 TByte
- 50 TByte capacity + 50 TByte replicated

Software

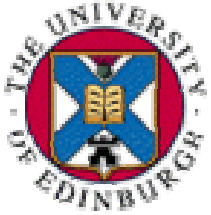

- Middleware: Globus 2.4
- RC: Globus RC
- MC: Native XML database eXist
- Bespoke Control Thread (CT)
- Command line client tools
- GUI client tools
  - Metadata browser
    - Based on OGSA-DIA browser
  - Job submission tool

Control Thread

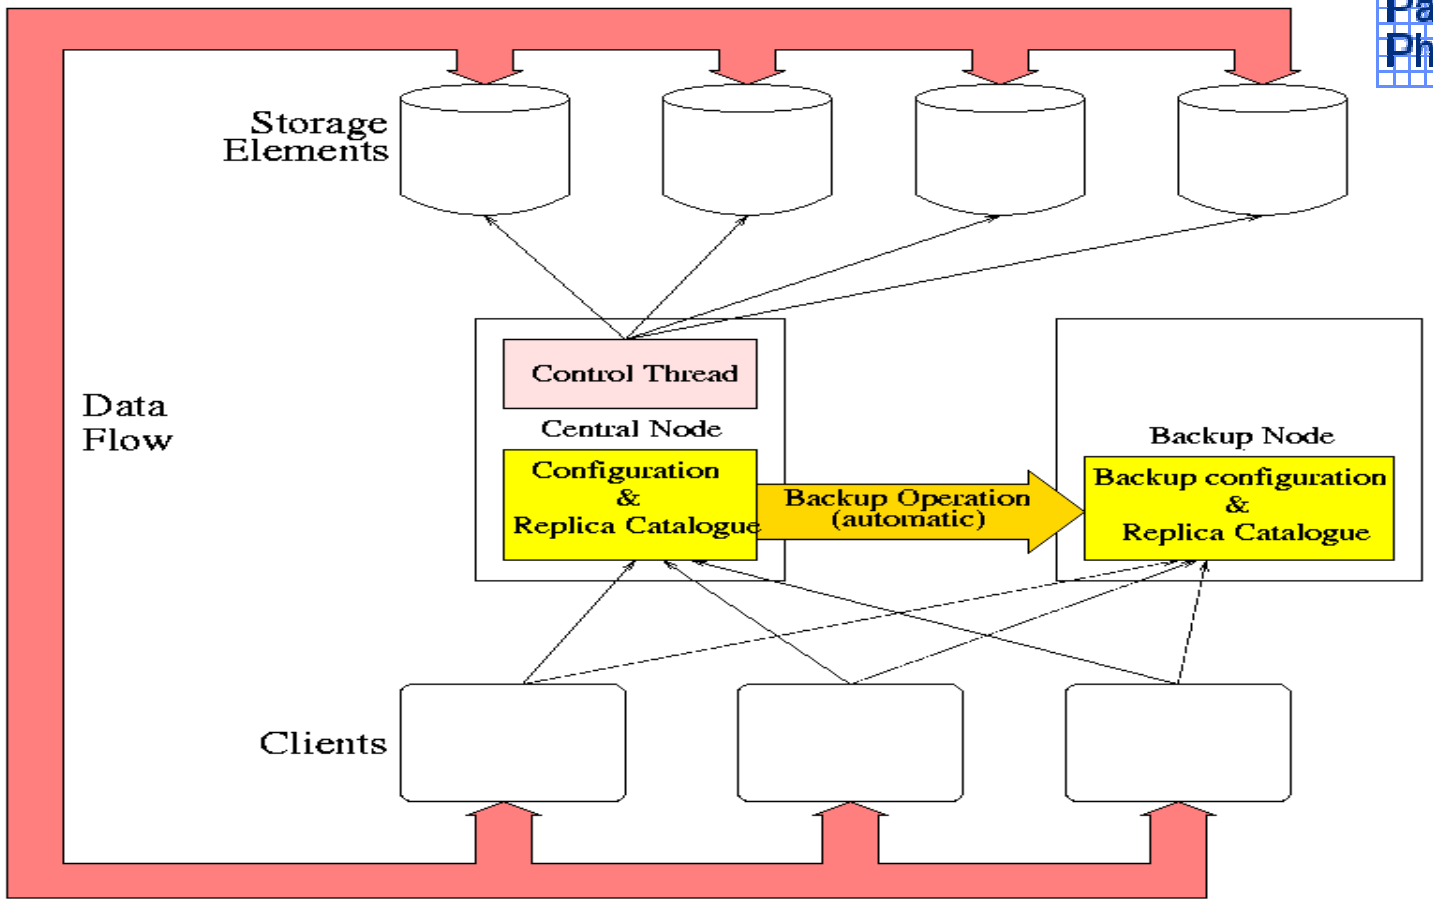

Command line tools

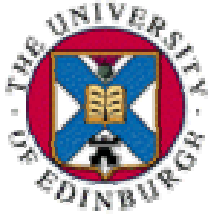

- For grid administrator
  - Add/remove machine
  - Disable/enable machine
  - Add/remove users (see later)
  - Remove Data
- For user
  - Add, retrieve, store locally, data
  - Job submission post-production on grid data
    - Remote machine not required to be QCDgrid

QCDgrid GUI

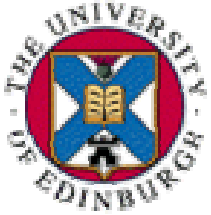

- Same functionality as command line
- Metadata browser
  - Search and read metadata
- Based on OGSA-DIA browser
  - Freely available from `qcdgrid.forge.nesc.ac.uk`
- Find data via metadata
- Can read MC without authentificaton

Authenticated access

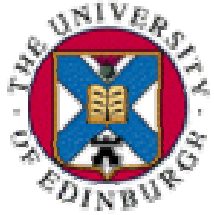

- QCDgrid data access
  - X509 certificate from trusted CA
    - Currently only UK e-Science CA
    - Will trust other CAs! For data sharing
  - QCDgrid is not closed
    - Not anonymous, but authenticated access.
  - ILDG data sharing
    - Technically could move to anonymous access
    - Require data sharing policy from ILDG board

Job submission tools

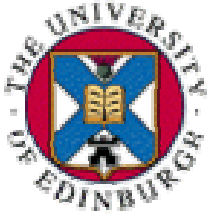

- Can talk to any system running globus
  - You need authenticated access
- Submit job and which data on QCDgrid
- Get back results
  - Data can be stored on QCDgrid
- For instance NGS can be used to process QCDgrid data

Grid concepts - Namespace

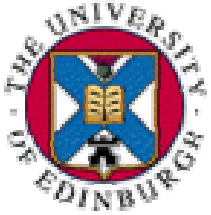

- Conventional namespace is a mathematical set
  - Used in definition of Logical filename
- XML Namespace defined by W3.org as
- *A collection of names identified by a URI reference*
- XML namespace has internal structure
- Can amalgamate namespaces

Grid concepts – logical filename

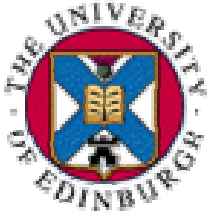

- Logical filename (LFN) is a name in a namespace which identifies a file
- Often it is a URI
  - Not `machine.domain:/path`
  - Data grid LFN references a file which can have several copies
  - Replica catalogue maps LFN to file instances

Grid concepts – replica catalogue

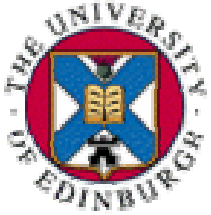

- Replica catalogue maps LFN to actual file instance
- Data grid - several copies of file
- Replica catalogue tracks number and location of file instances
- Data access is via the LFN and the replica catalogue

Metadata

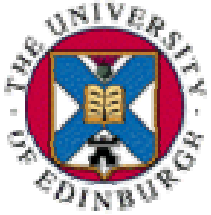

- Data about data
- "meaningful" filenames not enough
- Require a scheme for organising metadata
- Scheme has to be extensible
  - New things not previously thought of
- Only know full metadata when data created
- Metadata capture is very important

A broken scheme

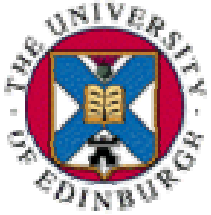

`D52C202K3500U010010_LL3450X_FL3400X_CMesonT00T31`

- UKQCD filename
  - What does X stand for?
- Wilson, Rotated, Clover
  - Many different clover. Scheme broken
  - X means none of the above!

Dynamical c_{SW}=2.02

NP determined – no information

XML

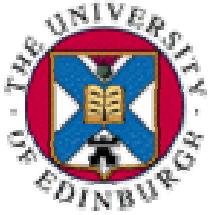

- XML eXtensible Markup Language
- www.w3.org/XML
  - XML is for structuring data
  - XML looks a bit like HTML
  - XML is text, but isn't meant to be read
  - XML is verbose by design
  - XML is a family of technologies
  - XML is license-free, platform-independent and well-supported

FAQs about XML schema

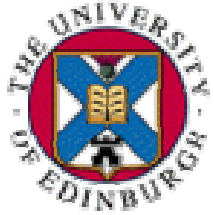

- What is XML schema?
  - Collection of rules for XML documents
  - An XML schema is itself an XML document
- Why do we need an XML schema?
  - Computers can read and understand XML IDs
  - `<length>16</length>`
  - Meaning of length is context dependent

QCDml1.1

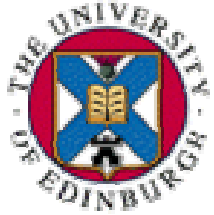

- Metadata split into two schemata
  - Ensemble XML `<markovChain/>`
  - Config XML `<gaugeConfiguration/>`
    - N.B. use lowerCamelConvention
- ILDG website for XML schema files
  - http://www.lqcd.org/ildg
  - Go to Metadata and follow links
  - Version 1.1 online and ready to use

Namespaces

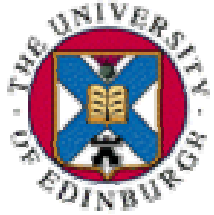

- Example XML ID for UKQCD data

```
<?xml version="1.0" encoding="UTF-8" ?>
- <markovChain xmlns="http://www.lqcd.org/ildg/QCDml/ensemble1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/ensemble1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Ensemble.xs
    <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
      Ensemble1</markovChainURI>
  + <management>
  + <physics>
  + <algorithm>
  </markovChain>
```

- XML Namespace defined by W3.org as
- *A collection of names identified by a URI reference*

First namespace

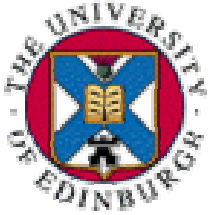

- URI defines namespace for QCDml

```
<?xml version="1.0" encoding="UTF-8" ?>
- <markovChain xmlns="http://www.lqcd.org/ildg/QCDml/ensemble1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/ensemble1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Ensemble.x
    <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
      Ensemble1</markovChainURI>
  + <management>
  + <physics>
  + <algorithm>
  </markovChain>
```

  - This is the default namespace
  - All elements of QCDml belong to this namespace

Second namespace

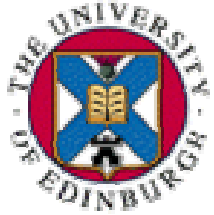

- Namespace of XML schema itself

```
<?xml version="1.0" encoding="UTF-8" ?>
- <markovChain xmlns="http://www.lqcd.org/ildg/QCDml/ensemble1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/ensemble1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Ensemble.xsd"
    <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
      Ensemble1</markovChainURI>
  + <management>
  + <physics>
  + <algorithm>
  </markovChain>
```

- Prefix **<xsi:>** for elements of XML schema
- XML ID is valid against WC3 XML schema

SchemaLocation

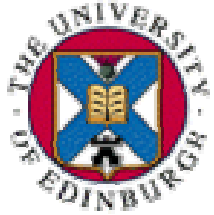

- The namespace of the schema

```
<?xml version="1.0" encoding="UTF-8" ?>
- <markovChain xmlns="http://www.lqcd.org/ildg/QCDml/ensemble1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/ensemble1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Ensemble.xsd"
    <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
      Ensemble1</markovChainURI>
  + <management>
  + <physics>
  + <algorithm>
  </markovChain>
```

- The file which contains the schema
- URI namespace *can* be URL of the schema instance – *not* compulsory

Logical filename

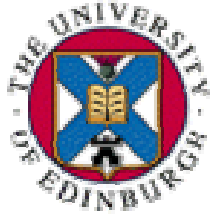

- Unique URI for a file in a namespace

```
<?xml version="1.0" encoding="UTF-8" ?>
- <markovChain xmlns="http://www.lqcd.org/ildg/QCDml/ensemble1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/ensemble1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Ensemble.xsd"
    <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
      Ensemble1</markovChainURI>
  + <management>
  + <physics>
  + <algorithm>
  </markovChain>
```

- Uniquely identifies this ensemble in ILDG namespace

Validation

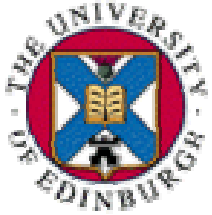

- Verify XML ID is valid against a schema
  - Schema aware applications can use XML ID
- Can write XML in `vi, emacs` etc
- CMM uses XMLSpy for schema and ID manipulation
  - built in validator, create XML ID from schema
- http://www.w3.org/XML/Schema
  - Many different tools

Configuration XML

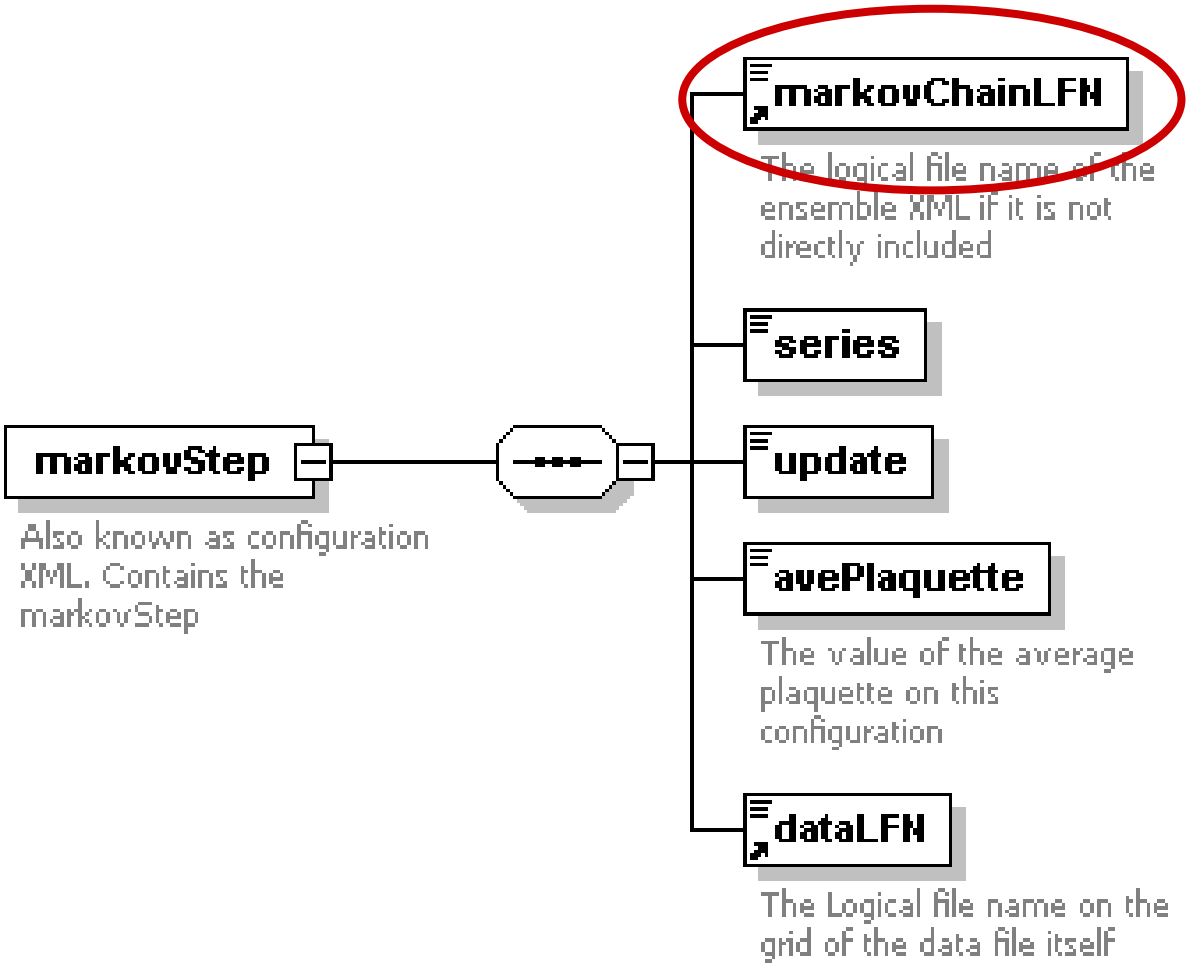

Logical File name of the ensemble in the ILDG namespace

dataLFN

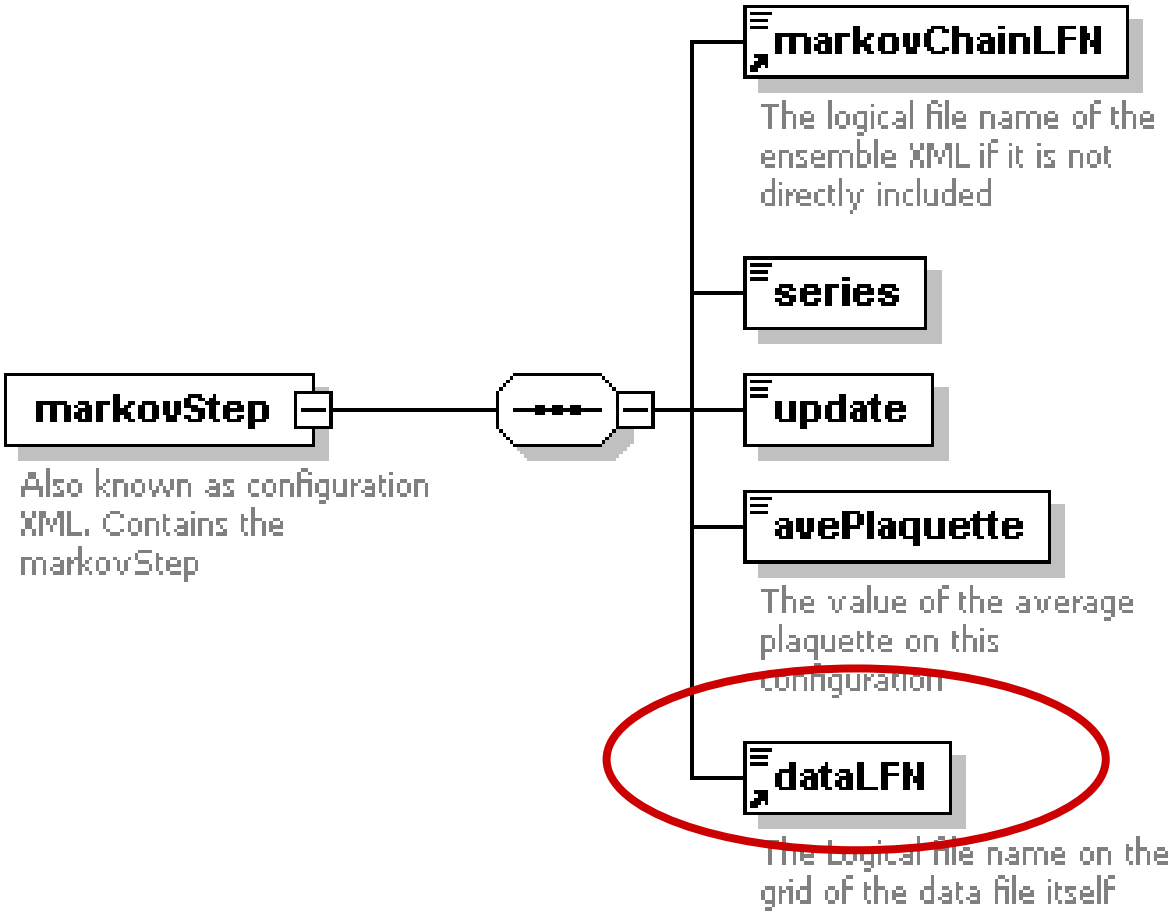

Logical File name of the configuration in the ILDG namespace

The markov chain

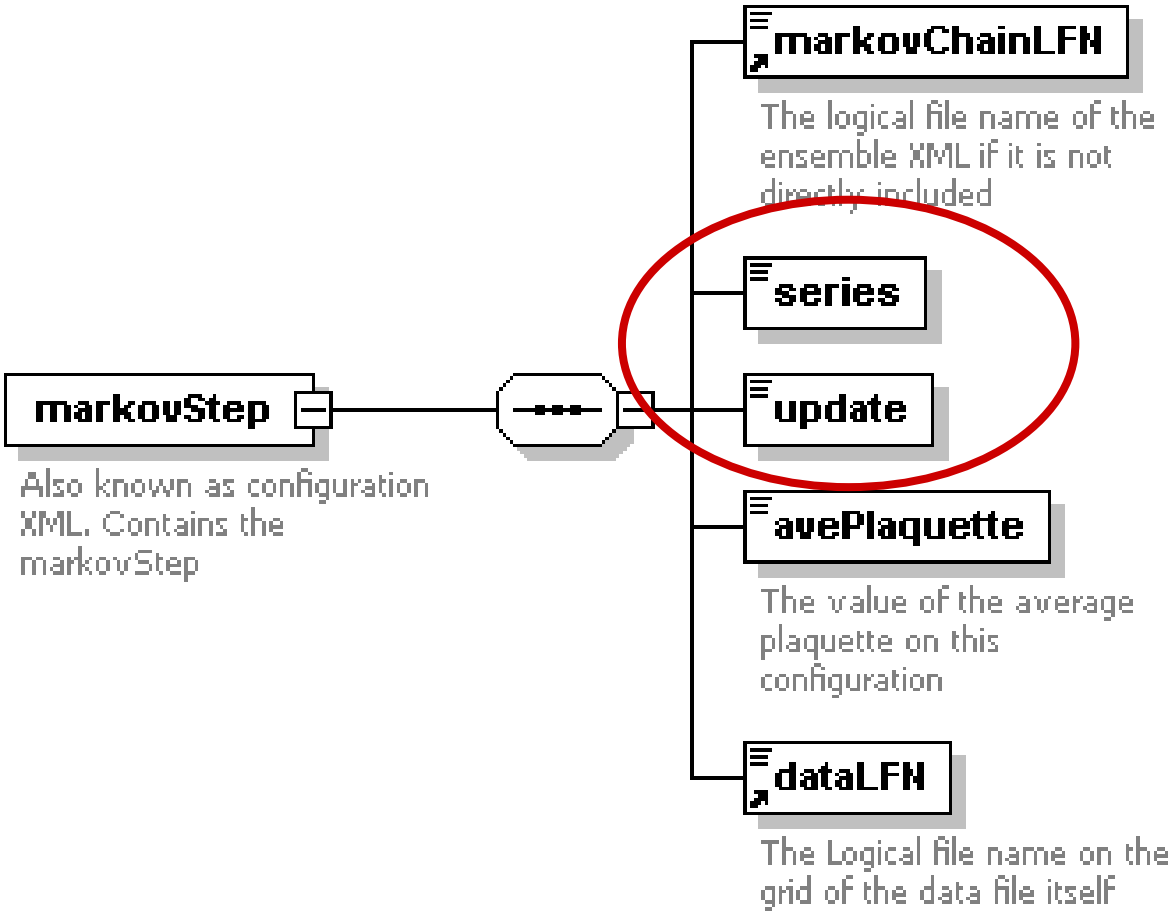

Where the configuration is in the trajectory of markov chain

avePlaquette

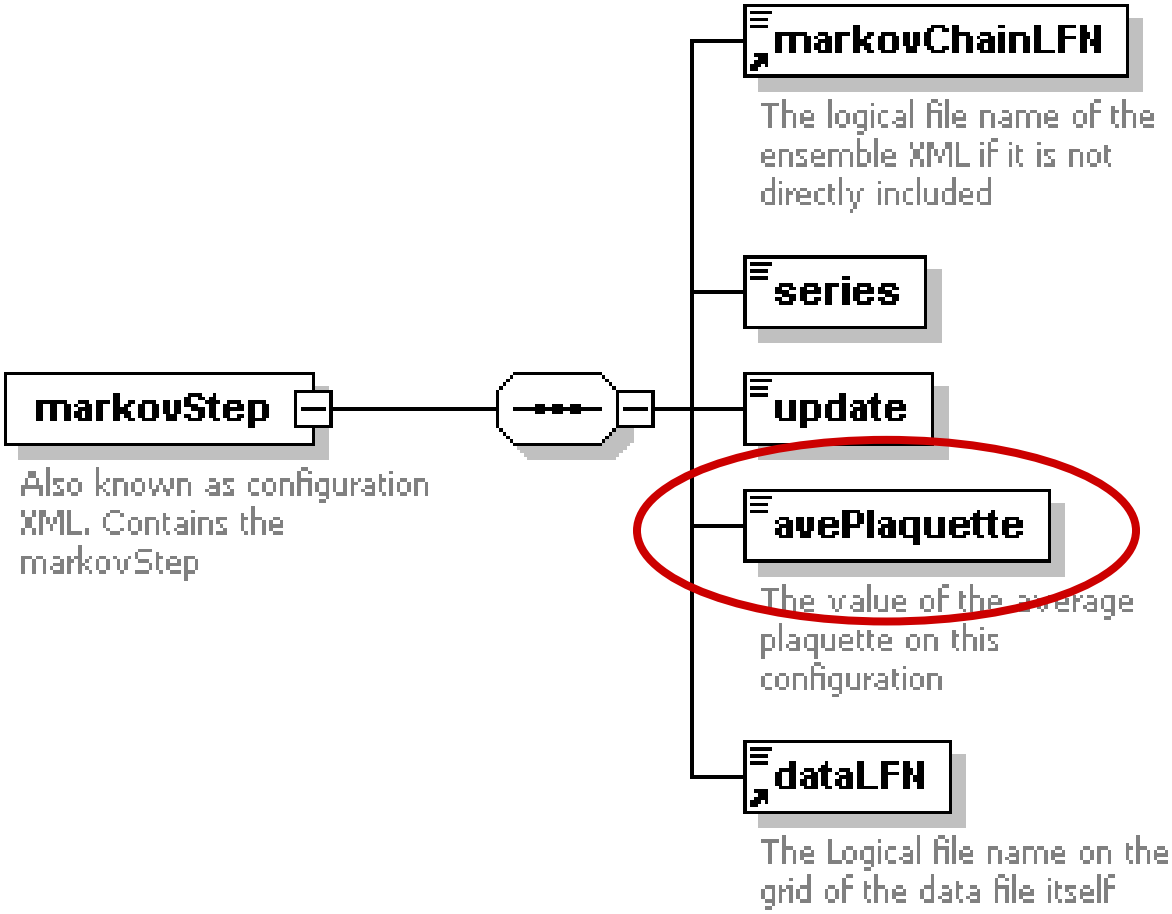

Very useful metadata, can be used to check data transformations are correct

Config: UKQCD example

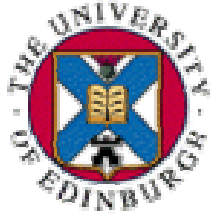


```
<?xml version="1.0" encoding="UTF-8" ?>
- <gaugeConfiguration xmlns="http://www.lqcd.org/ildg/QCDml/config1.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.lqcd.org/ildg/QCDml/config1.1
    http://www.ph.ed.ac.uk/ukqcd/community/the_grid/QCDml1.1/QCDml1.1Config.xsd">
  + <management>
  + <implementation>
    <algorithm />
    <precision>double</precision>
  - <markovStep>
      <markovChainURI>http://www.lqcd.org/ildg/ukqcd/DWF-IW-NF3-
        Ensemble1</markovChainURI>
      <series>1</series>
      <update>485</update>
      <avePlaquette>0.608425760908252</avePlaquette>
      <dataLFN>DWF/NF3/IWB2.2/M0.02/V16X32X8/RHMC/ukqcd/DWF-IW-NF3-
        Ensemble1.485</dataLFN>
    </markovStep>
  </gaugeConfiguration>
```

Name hierarchy

Unique name in for ensemble in ILDG namespace

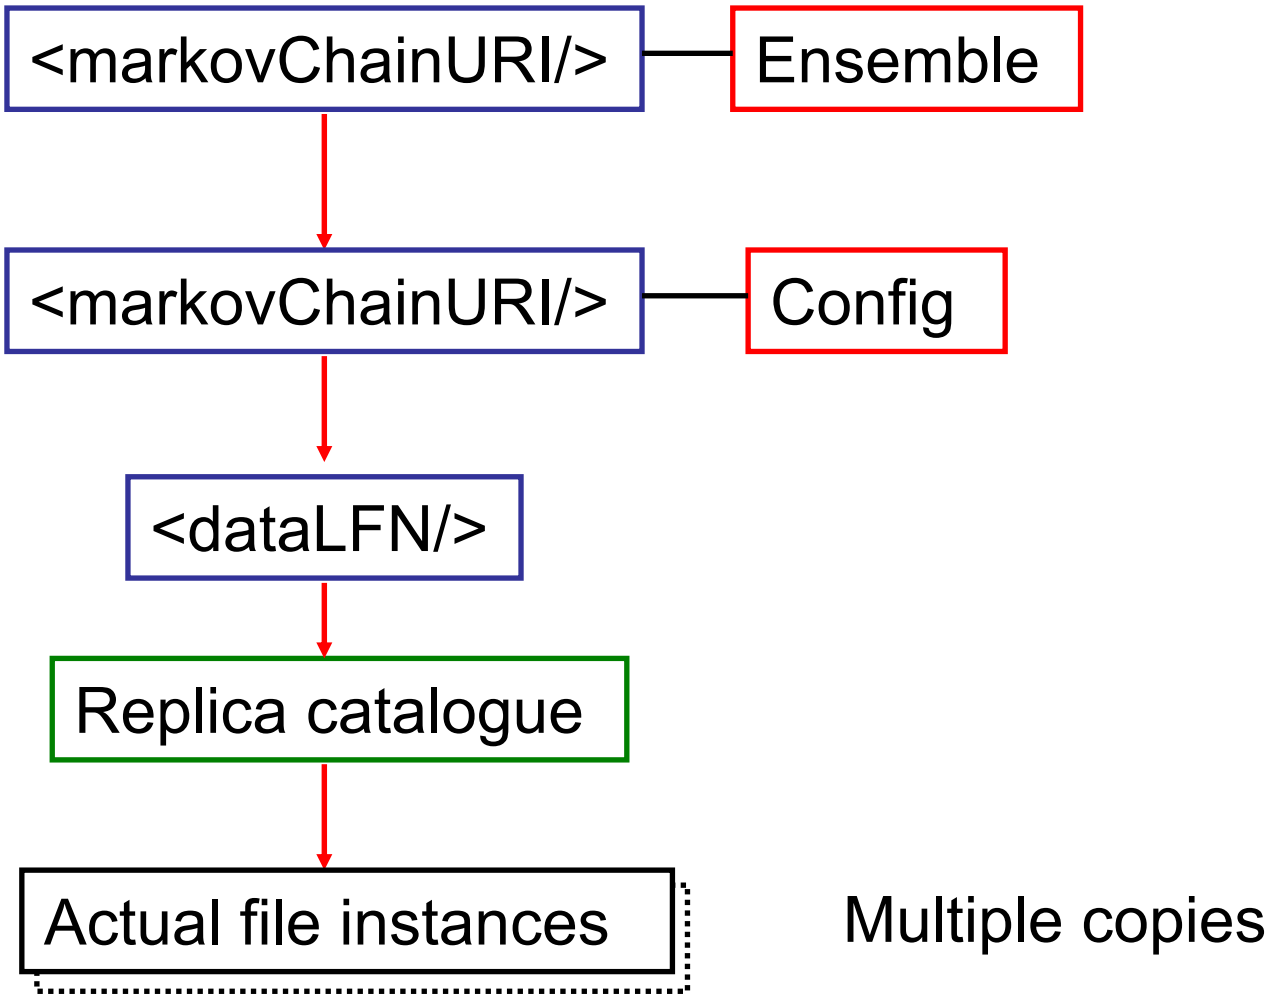

Metadata browser

Builds Tree from schema

Researcher walks tree

Selects relevant node

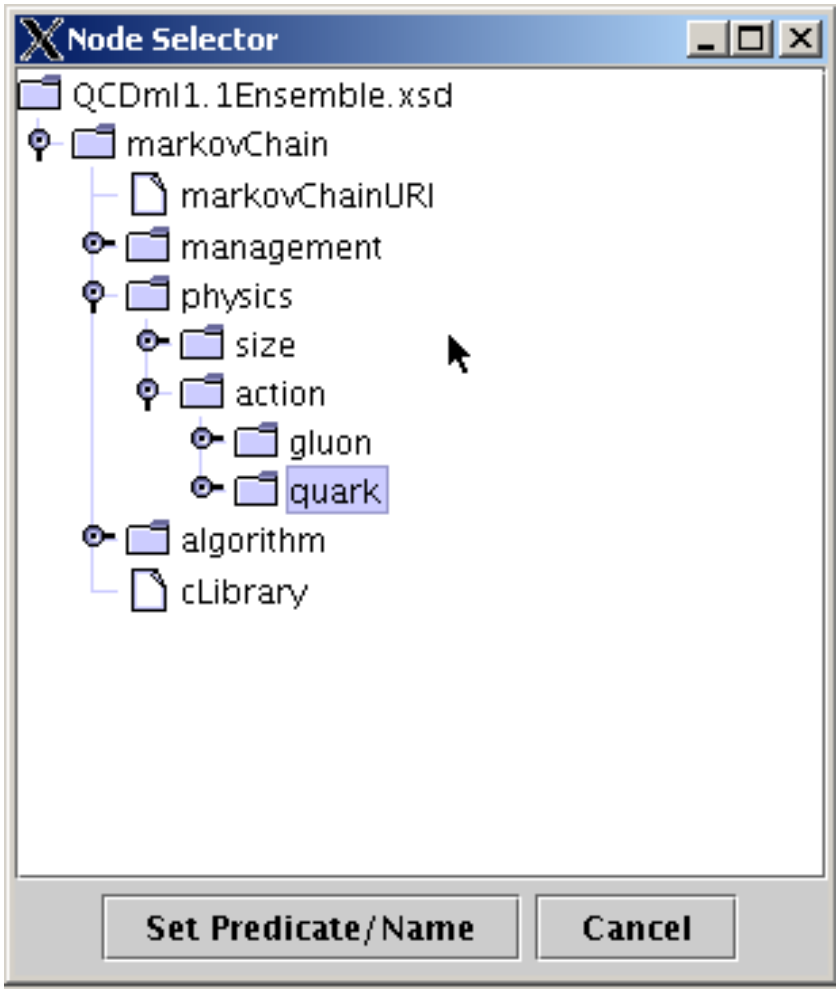

Graphical query

Construct XPath query via GUI

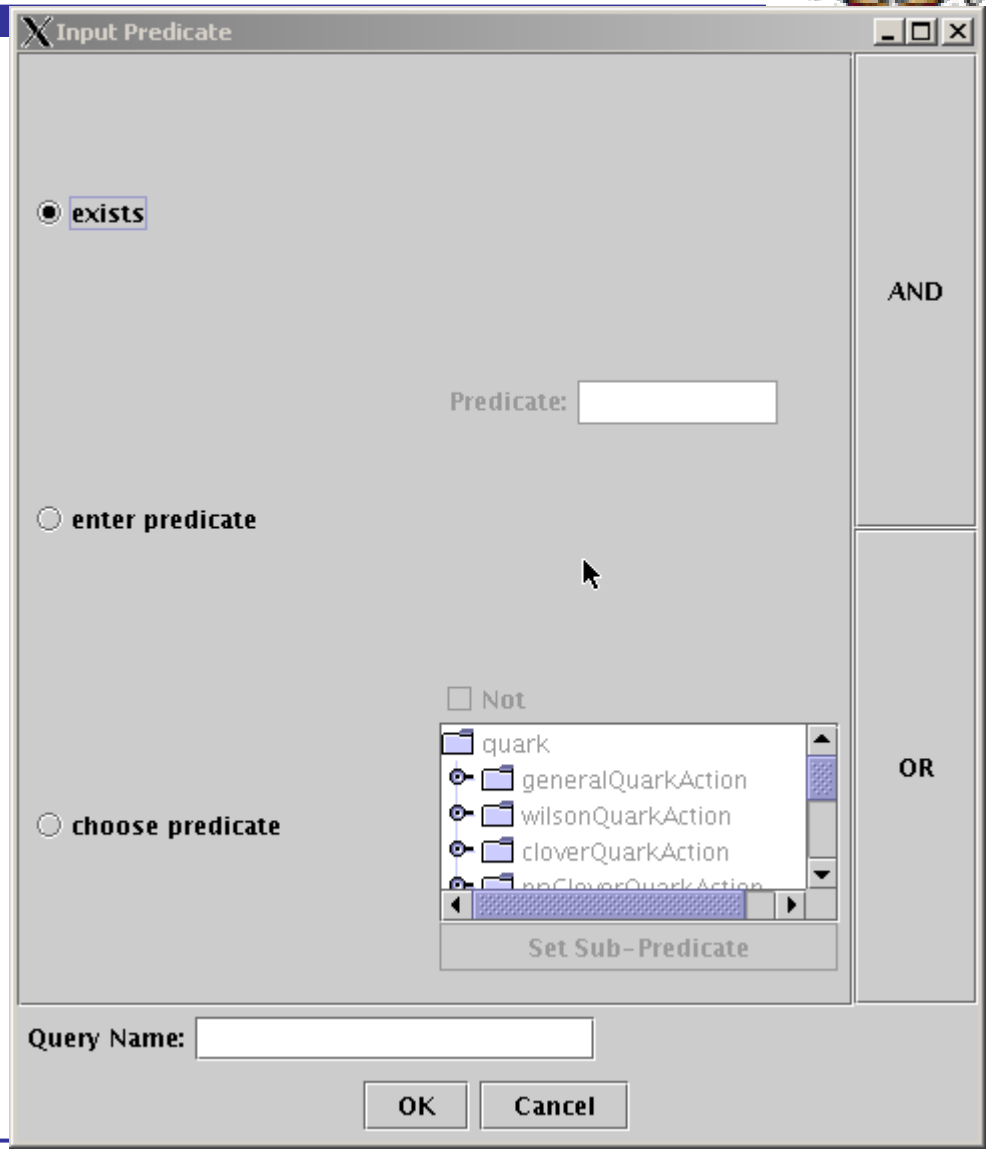

Returned results

Read the XML

Download data

Search the config XML for this ensemble

Extract LFN from config XML for command-line tool scripting

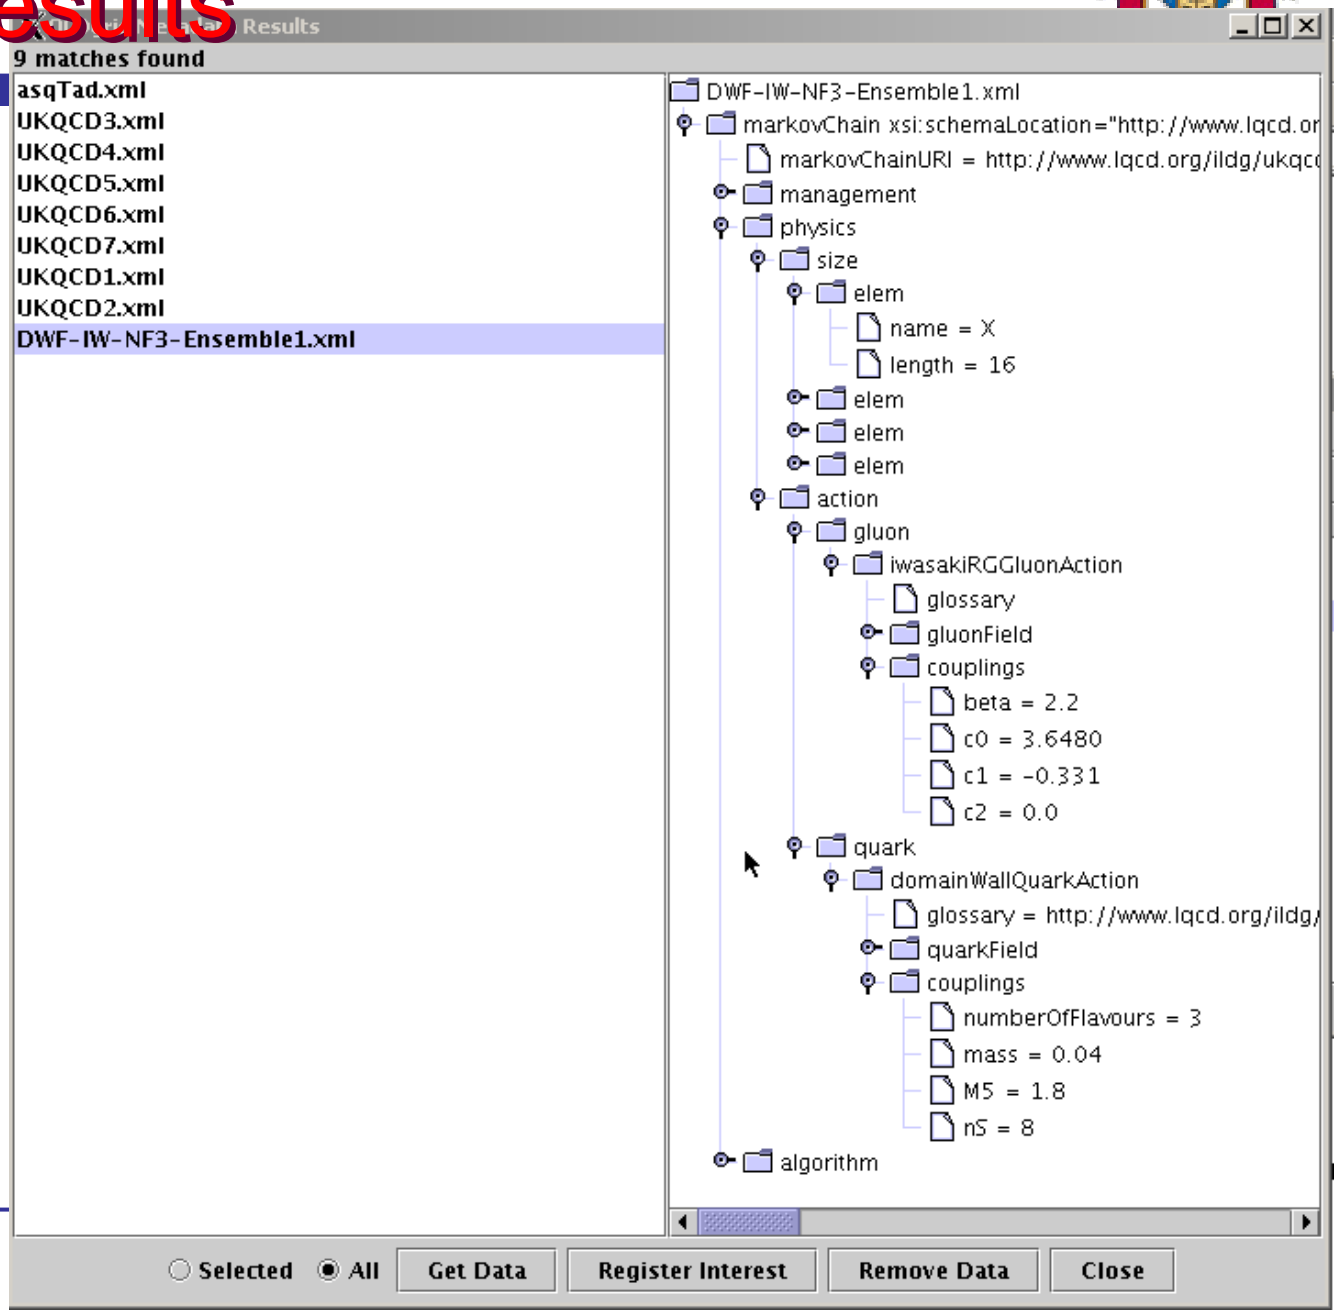

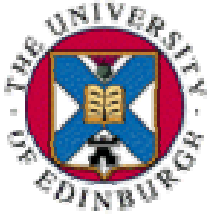

- Forum for data sharing amongst international collaborations
  - Share Code (open source)
  - Share Machines
    - QCDOC built in collaboration with Columbia, NY
  - Share data!
- Agree data format – done
- Agree metadata QCDml1.1
- Agree middleware – interfaces agreed

QCDgrid in ILDG

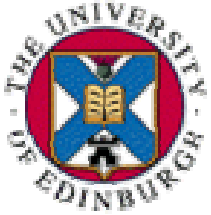

- ILDG is grid of grids
- Each member will offer data,
  - Not replicate each others
- Access other grids using own grid tools
  - Via common interfaces
- Technical problems are hard
- Sociological problems
  - E.g. Rules for sharing
  - Harder!

Summary

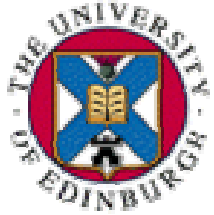

- QCDgrid is a secure, robust data grid
  - In operation
- Data access is via metadata
- Future work
  - Metadata mark-up tools
    - Currently rely on application codes and human intervention
  - Middleware tools which conform to agreed ILDG common interfaces