



DESY 06-152


Status Report on ILDG activities

Karl Jansen

*John von Neumann-Institute for Computing NIC,
DESY, Platanenallee 6, D-15738 Zeuthen, Germany
E-mail: Karl.Jansen@desy.de*

A status report about the International Lattice Data Grid (ILDG) is given. Different countries participating in the ILDG have created regional lattice data grid solutions that are implemented, working and used. The remaining task and the focus of present activities is the development of the interoperability of these regional grids. A first, successful step in this direction is a metadata catalogue service which is already working interoperable.

*XXIVth International Symposium on Lattice Field Theory
July 23-28, 2006
Tucson, Arizona, USA*



1. The ILDG idea and Participating Countries

The wish to have a convenient tool to exchange valuable configurations from large scale lattice simulations is certainly high. Such a tool should

- allow for a simple, semantic-based web search of configurations of interest;
- enable the download of such configurations;
- provide the configurations in a standardized format;
- realize all this internationally and across borders.

The realization of such a wish became in sight with the rapid development of Grid technologies as used for international experiments such as the LHC. Inspired by a suggestion of R. Kenway, Australia, Germany, Japan, UK and USA decided in an initial video conference organized by Edinburgh [1] to define a project of the *International Lattice Data Grid* (ILDG). Later, France and Italy joined and, since the ILDG is an open project, it is to be expected and desired that additional countries will participate in the future, too.

Progress on defining ILDG standards has been presented during several lattice conferences [2]. It is the purpose of this write-up to show that ILDG is becoming a real and working infrastructure which will support and hopefully ease further research on lattice gauge theories. Many of the questions raised in the first meeting have not only been answered, but working solutions have been found and implemented. ILDG issues were discussed in half-yearly virtual video meetings which also served to observe and follow the progress of the project, see the URL of ref. [1] for a list of the 8 ILDG meetings that have taken place so far. The structure of ILDG is that it consists of a board comprising representatives of each participating country*, a Metadata working group and a Middleware working group, the members of which will be listed below. Let me also recommend the poster presentation at this conference of ref. [3] for additional information and more technical details.

In order to set the frame, let us have a look at the (quite substantial) supercomputer resources that today can be used for lattice field theory (LFT) in the countries participating to ILDG:

- Australia: The CSSM collaboration has access to about 2 TFlops[†] of compute power on commercial machines installed in part at APAC (Canberra) and SAPAC (Adelaide).
- France: The main resource is apeNEXT with 1.2 Tflops, installed at the university Rome I “La Sapienza”.
- Germany: There are 6 Tflops apeNEXT systems, installed in Bielefeld and in Zeuthen. In addition, there is (peer reviewed) computer time available at the German national supercomputer centers, 45 Tflops of a BG/L system and a 10 Tflops IBM Regatta system at the Research center Jülich and a 26 TFlops Altix System at the LRZ in Munich. Lattice physicists typically have access to about 10–20% of this computer power.

*Present members are R. Brower (USA), K. Jansen (Germany, chairman), R. Kenway (UK), D. Leinweber (Australia), O. Pene (France), L. Tripiccione (Italy), A. Ukawa (Japan). The chair is rotated yearly.

[†]We give peak performances.

- Italy: Again, the working horse is a 7.2 Tflops apeNEXT system, installed at the university Rome I “La Sapienza”.
- Japan: Here the machines are a 14.3 Tflops PACS-CS cluster in Tsukuba and a 57.3 Tflops BG/L system at KEK. In addition, there are smaller O(1) Tflops installations in Hiroshima, KEK and Kyoto.
- United Kingdom: The major source of computer power is a 12 Tflops QCDOC system, installed in Edinburgh.
- USA: In the US, there is the 12 Tflops QCDOC machine in Columbia, two 10 TFlops QCDOC machines at BNL, and a total of about 10 Tflops cluster systems at Fermilab and JLAB. In addition, peer reviewed computer time is available at national supercomputer centers at NERSC, ORNL and Pittsburg.

In most of the countries there are ambitious plans for future increase of computer power with the aim to reach Petaflops computing soon.

When we add up the above listed resources, we find a total of about 150 Tflops for LFT today. The available computer time clearly allows for a significant push of the simulation parameters towards a physical situation with small enough pseudo scalar masses, large enough physical volumes and small lattice spacings. The configurations generated are very precious and it is one of the aims of the ILDG to make best use of these configurations. Already now, large collaborations – UKQCD, RBC, JLQCD, QCDSF, MILC, CSSM, CP-PACS/PACS-CS, ETMC, SESAM/T χ L/GRAL – are storing configurations on the grid, employing ILDG standard format, using ILDG infrastructures and making these configurations available to the corresponding members of the collaborations for further analysis.

The upload of configurations is proceeding rapidly and today we can already find more than 70.000 dynamical configurations residing in ILDG and waiting for download. In table 1 I summarize those physics plans of various collaborations from which configurations will be put on the grid. As can be seen, a great variety of actions is used and the set of configurations that will become available eventually is certainly very interesting. Besides the configurations listed in table 1, there exist configurations from older simulations, i.e. staggered $N_f = 2$ from MILC, Wilson $N_f = 2$ from SESAM/T χ L/GRAL and $N_f = 2$ tadpole improved Wilson from CP-PACS.

The different collaborations have setup their own corresponding policy for allowing to download and use these configurations. Possible rules are *immediate access; an acknowledgment in papers that use these configurations; waiting periods of six months between upload of configurations and giving access to them; draft of papers using the configurations in advance or a waiting period for the submission of key publications*. In addition, most collaborations want *citations of their key work for which these configurations were originally produced*. It appears therefore to be a wise idea to contact the collaborations before accessing the configurations and ask for their particular policy. Some of them are also thinking of collaborating on certain physics questions.

It should be emphasized that uploading configurations is non-trivial and needs quite some work. I believe therefore that there should be an applause to those collaborations who are willing to spend such an effort and an appeal to other collaborations to follow these examples.

Collaboration	fermion action	Flavors	a [fm]	am_q/m_{PS} [GeV]	aV_{\max}
RBC/UKQCD	domain wall	$N_f = 2 + 1$	0.12	$m_q = 0.0[2, 3, 4]$	$24^3 \cdot 48$
MILC	rooted staggered	$N_f = 2 + 1$	0.06-0.125	$\frac{m_l}{m_s} = 0.[1, 2, 3]$	$48^3 \cdot 144$
PACS-CS	NI Wilson	$N_f = 2 + 1$	0.07-0.12	—	$28^3 \cdot 56$
CSSM	FLIC	$N_f = 2 + 1$	0.12	$m_{PS} = 0.3$	$20^3 \cdot 40$
QCDSF	NI Wilson	$N_f = 2$	0.05-0.11	$m_{PS} = 0.25 - 1$	$32^3 \cdot 64$
ETMC	tm Wilson	$N_f = 2$	0.07-0.12	$m_{PS} = 0.25 - 0.5$	$32^3 \cdot 64$

Table 1: Configurations that are or will be in the near future on the grid. Quark masses and volumes are given in lattice units. am_l denotes the light, am_s the strange quark mass. NI stands for non-perturbatively improved and tm for (maximally) twisted mass fermions. aV_{\max} denotes the maximal lattice size that is planned for the simulations. In general, a sequence of also smaller lattice sizes is aimed for. Note that for the non-perturbatively improved Wilson fermion simulation the Iwasaki gauge action (PACS-CS) and the Wilson gauge action (QCDSF) will be employed.

1.1 Finding Configurations

In order to find the configurations listed in table 1, you have to query the various metadata catalogues of the regional grids. For those catalogues which are already ILDG compliant, one of the web interfaces which act as a portal and allow to query all catalogues can be used. Let us take the one operated by the German/French/Italian LatFor Data Grid (LDG) [4] as an example. Going to the site given in [4], you will find a *list of ensembles*, see fig. 1. From the list of ensembles you can select a *list of configurations*, see fig. 2. Finally, you can select a particular configuration to obtain the specific information for this configuration in a (hopefully) self-explanatory manner, see fig. 3. To actually retrieve configuration brings us to the next section.

2. Getting Configurations

Let us assume that you have browsed the metadata catalogue on the web as explained above, and you found your favorite set of configurations that just fits to address your physics problem. Let us further assume that you have even contacted the corresponding collaboration and you got green light for downloading their configurations. Here are then the next steps to proceed:

- The first thing you have to do is to get a grid certificate. For this you have to identify a local Certificate Authority (CA), which is willing to provide you with a certificate. Before you receive a certificate, the CA or one of its Registration Authorities (RA) will check your identity, e.g. ask for your passport. It is foreseen that ILDG resource providers will trust certificates issued by any CA that is member of the International Grid Trust Federation (IGTF, [5]). Of course, the step of obtaining a grid certificate has to be done only once.
- As a next step you have to become member of the Virtual Organisation (VO) ILDG. There is a policy being worked out about who can become a member of the VO ILDG. Basically, this is open for all people doing lattice field theory. Each regional grid has to nominate two representatives who can decide whether a particular person shall become a member of the VO ILDG.

- The next thing you have to do is to install software that will allow you to actually download the configurations. This software will depend on the particular regional implementation of ILDG in your country (see also the discussion on regional grid setups below). Let me take, for some convenient reason, as an example again the LatFor Data Grid (LDG) [4]. At our site, you would retrieve a Grid User Interface and a package with user tools called `ltools` [6] which runs on basically all Linux platforms.
- After you have obtained your grid certificate and you have installed the software you are now ready to get your favorite set of configurations. Let us take as an example the set you find on the web as shown in fig. 3. What you will detect is a *Logical File Name* (LFN) for the configuration you want to obtain. This LFN represents the globally unique grid address of your configuration and issuing `lget lfn` will then get you the desired configuration directly on your disk (if you have made sure that you have enough space there). Naturally, if you want a whole set of configurations, you will write a small script to retrieve all desired configurations.

It could be asked whether this all is worth the effort. On the negative side, there is certainly the initial task of getting your grid certificate and installing the software to make use of ILDG on your regional grid. But then, you can download very precious configurations on which you can address your physics application (which should, of course, not be in conflict with the configuration owner's ideas). Moreover, you know exactly the format in which the configuration is written and you are sure that for future downloads this format will be the same. Therefore, your measurement code will run immediately on any new set of configurations. As an additional point, it should be stressed that once you have overcome the initial difficulties, the download of configurations will be routine. Thus, the ILDG mechanism is a new way to use configuration for all their potential of physics applications, avoiding duplication and loss of information - something that has been very common in the past - and therefore increasing the overall efficiency of the community.

Of course, it remains to be seen in practice, whether this optimistic point of view is indeed realized or whether we are left with some practical issues that will drive life complicated.

Uploading the configurations is more tricky and needs extra work. First of all, configurations have to be stored in the ILDG configuration format, see ref. [1], a Logical File Name (LFN) has to be added and the metadata have to be created. For doing all this, the freely available LIME library is needed which can be obtained from the URL in ref. [7]. The metadata consist of the *ensemble XML file*, the *configuration XML file*, a *glossary file* and a *configuration checksum*. Finally, you have to store the binary configuration into a storage element and the metadata into the metadata catalogue. This typically can be done using just a single command. For instance, in LDG this command is called `lput`, where you have to specify the storage element you actually want to store the file physically and the names of the metadata document as well as the file containing the configuration.

Although there are many help tools on the corresponding ILDG sites, and scripts exist to automatize the upload procedure, there is clearly some effort necessary to store the configurations.

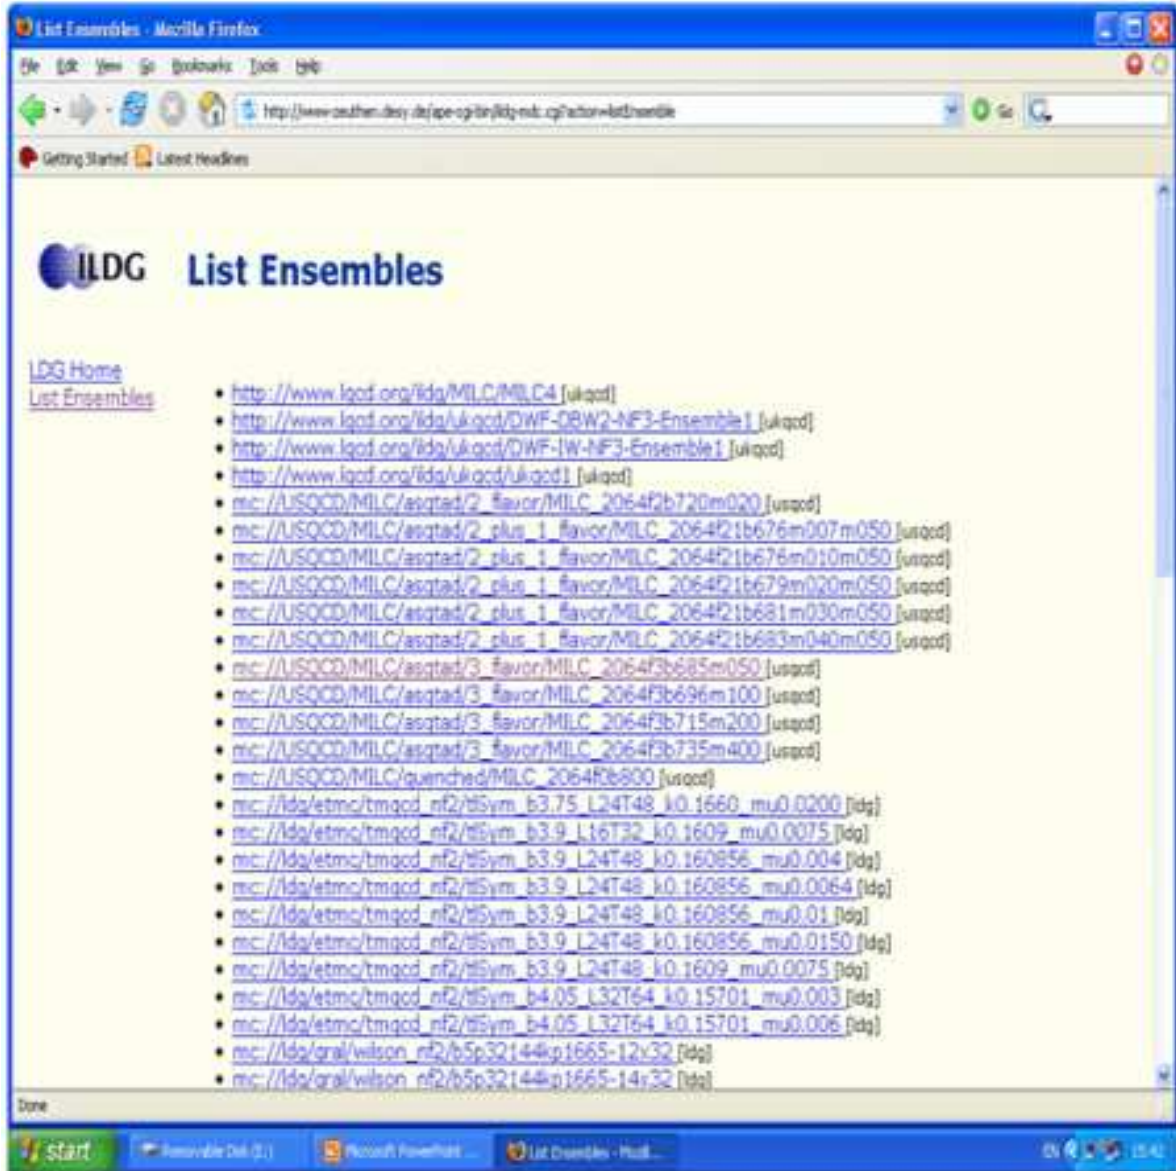


Figure 1: The list of uploaded ensembles of configurations as seen when using the LDG metadata catalogue.

3. Looking behind the scene

The fact that we have today solutions and working implementations of regional grids realizing the ILDG idea, is purely due to the hard work of people involved in the metadata and the middleware groups listed here ‡:

- Metadata Working Group: G. Andronico, P. Coddington, C. DeTar, R. Edwards, B. Joo, C.

‡In addition, there are many people working together with the Metadata and Middleware working groups. The names of these people can be found on the webpages of the corresponding collaborations.



Figure 2: The list of configurations when one of the ensembles is selected.

Maynard, D. Pleiter, J. Simone, T. Yoshie

- Middleware Working Group: G. Beckett, D. Byrne, M.Ernst, B. Joo, D. Pleiter, M. Sato, C. Watson

In order to establish even a regional grid infrastructure is highly non-trivial. Many components have to work together as sketched in fig. 4. It is therefore a very big step forward that at basically all participating sites regional grid infrastructures have been developed. Let me list the characteristics of the different regional grids that exist today.



Figure 3: The information on a selected configuration.

- Germany/France/Italy

These countries use the LatFor data grid (LDG). LDG has a metadata catalogue [8] to query, download, upload and manage metadata. The user software tool is `ltools` with commands such as `lget`, `lput`, `lls`, `lupdate` to download, upload, list and update configurations, respectively. The regional grid infrastructure is based on LCG-2 compliant components. For further details see ref. [4]. To store configurations, there are 50 Terabytes storage space in Germany and 5 Terabytes in France available.

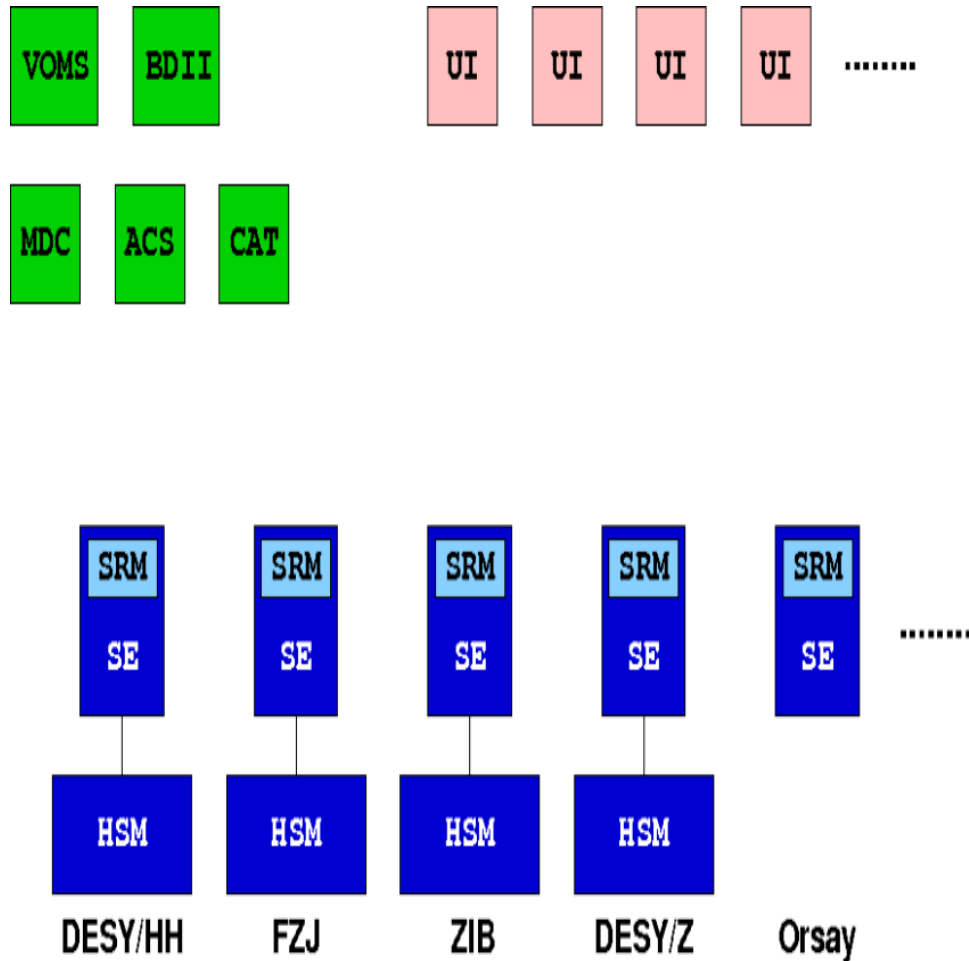


Figure 4: The major components for a grid infrastructure. The abbreviations are: UI: User Interface, SE: Storage Element, SRM: Storage Resource Manager, VOMS: Virtual Organization Management System, MDC: Metadata Catalogue, BDII: Berkeley Database Information Index, HSM: Hierarchal Storage Management, ACS: Access Control Service, CAT: File Catalogue. All of these components have to work together smoothly in order to have a functioning regional grid infrastructure.

- United Kingdom

In the UK the QCDgrid is used. QCDgrid also has a metadata catalogue based on the native XML-database “eXist” [10] complete and deployed on the UKQCD development system. It is now running on the production system providing access to the real metadata. Access is realized through the ILDG sample clients metadata. User tools come, as within LDG, as command line tools such as `put-file-on-qcdgrid`, `get-file-from-qcdgrid`. In addition, the UK QCDgrid is working on developing web-based tools that will allow for graphical user interfaces to handle the metadata. The storage space in the UK is about 80 Terabytes across the UK but mainly in Edinburgh, see ref. [9].

- Japan

The Japanese Lattice Data Grid (JLDG) has also implemented an eXist-based metadata cat-

alogue. The user software consists of commands such as `Gftp` on a specially developed Gfarm file system. Security is realized by the Grid Security Infrastructure (GSI) [11]. Also in Japan 50 Terabytes of storage space is available for the ILDG. More information is provided at ref. [12].

- USA

The USQCD has a metadata catalogue up and running. The metadata catalogue is constructed as a web service. The user software to download and upload configurations is not yet completely finished. USQCD has 50 Terabytes of storage elements at NERSC, BNL and FNAL. Further information can be found at ref. [13].

- Australia

In Australia we again see the metadata catalogue to be complete. User software for handling configurations is realized through a web portal. In Australia we find 25 Terabytes for storage. Additional information is given at the ref. [14].

The above list demonstrates that indeed substantial progress has been achieved towards the implementation of the ILDG idea:

- We have five working regional ILDG infrastructures ready, LDG (Germany/France/Italy), JLDG (Japan), QCDgrid (UK), USQCD (USA) and Australia.
- There are about 250 Terabytes of storage space available at these sites. This would allow to store, roughly, 200.000 configurations on a $32^3 \cdot 64$ lattice which would need some time to be generated.
- Actually, already now 70 000 configurations are available on the ILDG (many of them on smaller lattices than $32^3 \cdot 64$).

4. Interoperability

With the regional grids now functioning, ILDG has to face the next challenge, namely the interoperability of these regional grids. One big progress in this context is the operation of interoperable metadata catalogue services. This means that it is possible for the sites to browse all each other's metadata catalogue, see fig. 5 for examples of the appearance of the web-browser at the participating sites. In the development of this service, it was necessary to define the interface in terms of the web service description language (WSDL) and to agree on a behavioral specification.[§] In addition, a set of tests has been defined to verify ILDG compliance of particular services.

However, there are still most significant efforts required to achieve the interoperability of remaining components which are: *Security, File Catalogues and Storage Elements*. It is worth stressing that also here progress has been made already. It could be demonstrated successfully that file transfers between LDG (at DESY), QCDgrid (at EPCC) and USQCD (at Fermilab and JLAB) are

[§]WSDL essentially defines the name of the services as well as the name and type of the input arguments and the return values. The behavioral specification defines, e.g., what should happen in case of errors, which status codes are returned, etc.

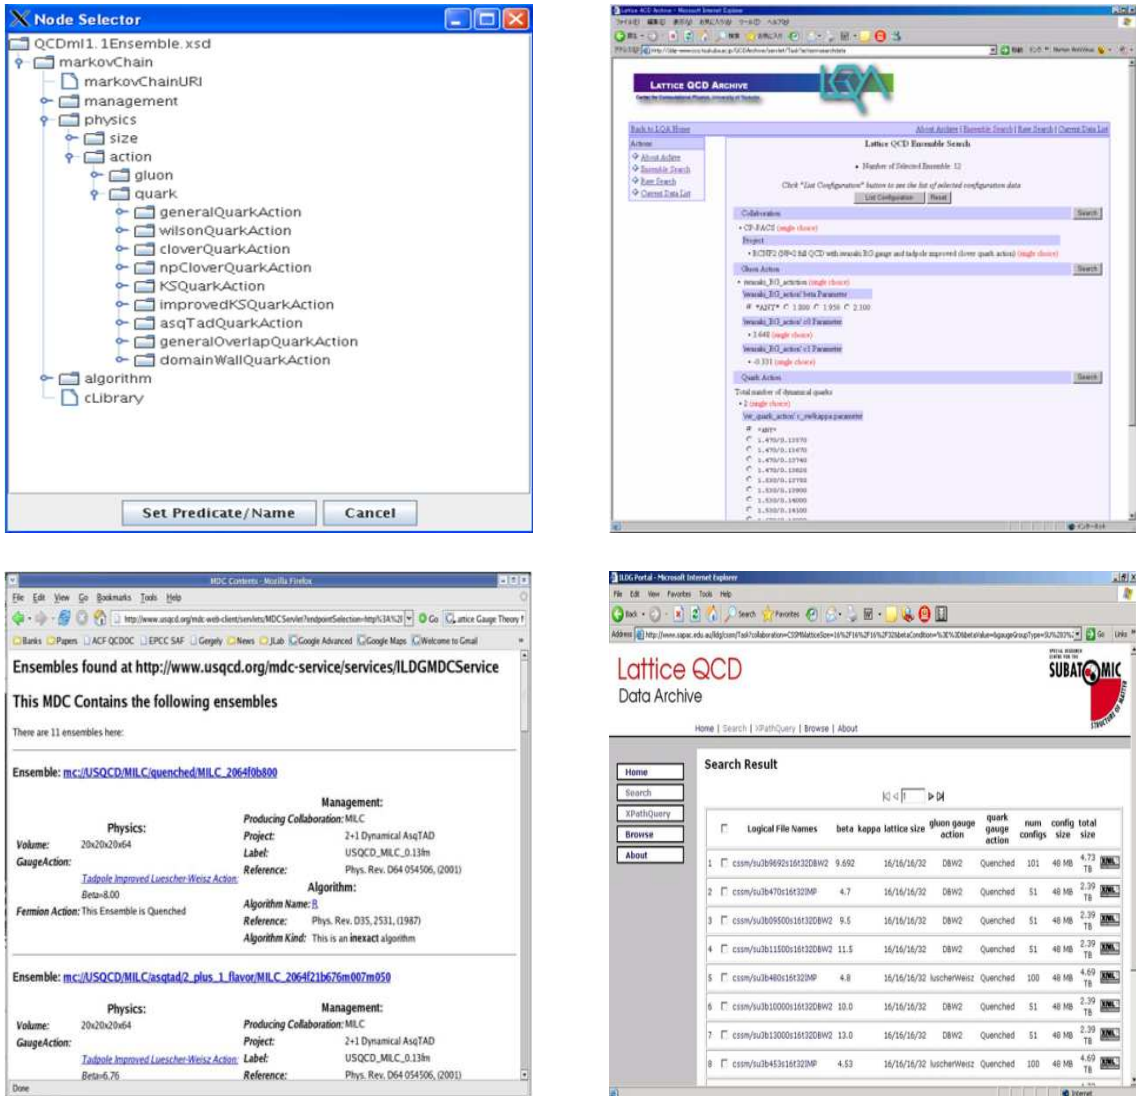


Figure 5: The appearance of different web browsers for listing configurations: UK QCDgrid (upper left), Japan JLDG (upper right), USA USQCD (lower left) and Australia (lower right). For examples of the German/France/Italy LatFor Data Grid (LDG) see figs. 1-3.

possible. Another point is that through a virtual organization membership service (VOMS) provided by the global grid community the VO ILDG can be managed. A draft for policies on who can become a member of the VO ILDG exists and will be completed soon.

5. Summary

In discussions with many lattice practitioners, it was clearly expressed that the ILDG idea is extremely useful and valuable. Nevertheless, there was always a number of questions, typical ones (with some answers) are:

Can I determine access rights myself?

This very desirable feature is currently only supported at the regional grid level which allow to

define groups, rights for each group and ensembles of configurations. In particular, it is not only possible to manage access rights for configurations, but to also delete and replace them.

Will the data be replicated?

Currently, replication is possible only within the regional grids. It is, however, planned to make replication beyond grid boundaries possible.

How can I check that I got the right configuration?

For each configuration a plaquette value and a checksum is provided.

Is the schema to describe data extensible?

Yes, more information can always be built in.

Can I put in algorithm information?

It is possible to add a name space for algorithmic information. However, ILDG demands only a limited information by default since algorithmic parameters are plenty and it will become too complicated to manage them.

What about propagators?

The ILDG is presently evaluating the possibility to have a simplified version, e.g. standardized format, for propagator storing. In general, it appears to be too complicated to repeat the example of configurations but discussions are ongoing.

Who is paying for all this?

In most countries, ILDG is embedded in larger grid projects that are mainly oriented towards LHC or other large scale experiments. ILDG profits from these developments and makes use of the software and hardware infrastructure that becomes available through these projects. For example, EPCC at Edinburgh and DESY at Hamburg/Zeuthen are Tier 2 centers within the LHC grid (LCG). In addition, ILDG is involved in a number of grid projects financed on a national or European level.

In conclusion, much progress has been made to develop regional grids for the ILDG idea. These infrastructures are presently successfully used by large collaborations, which have members at many different sites, to exchange their configurations. Again it should be stressed that the realization of these regional grids do not come for free but through the hard work of people involved in building-up these infrastructures.

The next milestone and target has to be the interoperability of these regional grid solutions. Here, the possibility of browsing each other's metadata catalogue has been successfully demonstrated, already. Even file transfer between several sites (DESY, EPCC, Fermilab/JLAB) has been achieved. Both of these accomplishments are important steps towards interoperability.

Thus, it remains to emphasize that it is time now to get a grid certificate at your local authorization site and become a member of the virtual organization ILDG.

6. Acknowledgments

This presentation at the 24th International Symposium of Lattice Field Theory in Tucson would not have been possible without the help of people sending me the material of their corresponding country. I am very much indebted to R. Brower, D. Byrne, B. Coddington, B. Joo, R. Kenway, D. Leinweber, C. Maynard, O. Pene, D. Pleiter, M. Sato, L. Tripiccion, A. Ukawa, A. Williams, T. Yoshie. A special thanks goes to D. Pleiter for a critical reading and very useful

comments to this manuscript. Finally, in the name of all ILDG members, I would like to thank the organizers of the conference to give the room to present the status of the ILDG.

References

- [1] <http://www.lqcd.org/ildg>
- [2] C. T. H. Davies, A. C. Irving, R. D. Kenway and C. M. Maynard [UKQCD collaboration], “International lattice data grid,” Nucl. Phys. Proc. Suppl. **119** (2003) 225 [arXiv:hep-lat/0209121]; A. C. Irving, R. D. Kenway, C. M. Maynard and T. Yoshie, “Progress in building an International Lattice Data Grid,” Nucl. Phys. Proc. Suppl. **129** (2004) 159 [arXiv:hep-lat/0309029]; A. Ukawa, Nucl. Phys. Proc. Suppl. **140** (2005) 207 [arXiv:hep-lat/0409084].
- [3] B. Joo, C. Maynard and D. Pleiter, these proceedings.
- [4] <http://www-zeuthen.desy.de/latfor/ldg>; M. Ernst *et al.*, “Using Grid Technologies for Lattice QCD,” Proceedings of CHEP06 (2006).
- [5] <http://www.gridpma.org>
- [6] <http://www.zib.de/i3hp/itools/>
- [7] The LIME library can be found at <http://www.physics.utah.edu/detar/scidac/>
- [8] O. Büchner *et al.*, “Datagrids for lattice QCD,” Nucl. Inst. and Methods in Physics Research, Vol. 559, Issue 1 (2006) 57-61.
- [9] <http://www.gridpp.ac.uk/qcdgrid/>
- [10] <http://exist.sourceforge.net/>
- [11] <http://www.globus.org/toolkit/docs/3.2/gsi/key/index.html>
- [12] <http://www.jldg.org/>
- [13] <http://www.usqcd.org>
- [14] <http://www.sapac.edu.au/ildg/cssm>