

Exploitation of Switched Lightpaths for e-Health: Constraints and Challenges

Lee Momtahan*

Oxford University Computing Laboratory

E-mail: Lee.Momtahan@comlab.ox.ac.uk

Andrew Simpson

Oxford University Computing Laboratory

E-mail: Andrew.Simpson@comlab.ox.ac.uk

The Exploitation of Switched Lightpaths for e-Science Applications (ESLEA) project is evaluating the feasibility of using switched lightpath networks to support various e-Science applications. This feasibility is being investigated via case studies, one of which pertains to the use of such networks to support distributed healthcare research and delivery. We consider some of the constraints and challenges associated with the use of switched lightpath networks in this context.

Lighting the Blue Touchpaper for UK e-Science - Closing Conference of ESLEA Project

March 26-28 2007

The George Hotel, Edinburgh, UK

*Speaker.

1. Introduction

The future of networking is moving towards switching in the optical domain, with the UK's effort being represented by the UKLight network. Within the UK, the national e-Science Programme [1] helped drive the development of infrastructures and applications to deliver 'large-scale' science. In this spirit, the Exploitation of Switched Lightpaths for e-Science Applications (ESLEA) project [2] aims to pioneer the use of UKLight to support a variety of e-Science applications. The concern of this paper is the e-Health 'mini-project' of ESLEA.

A two-pronged approach has been taken within the e-Health mini-project: the consideration of generic issues and the consideration of more specific issues pertaining to a particular application area have been split into two parallel tasks. Initial experiences with respect to the former were reported in [3]; initial experiences with respect to the latter were reported in [4].

The specific application associated with the e-Health mini-project pertains to typical use cases derived from the Integrative Biology project (IB) [5, 6], which is developing a customised virtual research environment to support researchers in the integration of biological information across multiple spatio-temporal scales. IB has established a community that is committed to sharing technologies and best practice with respect to heart and cancer modelling. The exploitation of networks such as UKLight has the potential to improve remote visualisation and to allow models to be run at more than one site—even when they are tightly coupled.

Current models of action-potential propagation within whole-heart models can involve finite element meshes with tens of millions of nodes, and up to 50 solution variables. Solutions in real-time for just a single heartbeat require (at least) millisecond temporal accuracy, resulting in data generation in the 10–100 TeraByte range from a single simulation. Establishing the feasibility of running a resource-intensive simulation model of electrical activity in the heart across a number of geographically distributed sites, using UKLight to provide the necessary low latency network, could offer significant benefits.

The structure of the remainder of this paper is as follows. In Section 2 we expand upon our use cases drawn from IB. In Section 3 we describe our experiences in trying to realise these use cases. Finally, Section 4 provides a brief summary of this contribution.

2. ESLEA and IB

IB is developing a customised virtual research environment that is capable of supporting research scientists in the integration of biological information across multiple spatiotemporal scales. Such integration could provide comprehensive descriptions at the system level, which can aid in the determination of biological function in both normal- and patho-physiology. It is intended that the tools developed by IB will support and accelerate the work of clinical and physiological researchers. The long-term beneficiaries of the work should be patients with heart disease and cancer, which together cause 60% of all deaths within the UK.

To this end, IB's end-users are drawn from the bio-mathematical modelling, cardiac physiology, and clinical oncology communities. These groups all wish to make use of mathematical modelling and HPC-enabled simulation to investigate multi-scale and multi-physics phenomena in the modelling of the heart or cancer. Much of this work is grounded in experimental data and involves

the development of complex mathematical models, usually involving the solution of non-linear systems of coupled ordinary and partial differential equations (or equivalent stochastic representations of the underlying processes on very small spatial scales) in complex, deformable three-dimensional geometries. In addition, very large quantities of simulation data are being generated which must be validated against experiments, curated, synthesised, analysed, and visualised.

IB is using, and has used, a variety of computational and storage resources distributed throughout the UK, including HPCx at Daresbury, CSAR at Manchester and NGS (National Grid Service) Data Clusters at the Rutherford Appleton Laboratory (RAL).

The activity pertaining to IB and ESLEA was split into three phases.

The intention of the first phase was to set up a link between Daresbury and RAL, with a view to transferring TeraBytes of simulation data between the main computation site at Daresbury and the main storage site at RAL.

The plan for the second phase was to utilise the link established in the first stage for the purposes of visualisation—with data being streamed for storage also being passed to visualisation servers based at RAL. The motivation for this was that such processing had the potential to reduce the high bandwidth stream of raw simulation results into a relatively low bandwidth stream of geometry information—which could then be shipped to the end-user for rendering on their desktop over normal low bandwidth networks.

The third phase was more ambitious, and involved the execution of a heart model across heterogeneous compute resources. In modelling heart mechanics, the equations that model the mechanical pumping action are coupled to the electrical activity generating the muscular contractions via a large system of ordinary differential equations which model the ion flows through the cell membrane. A potentially attractive approach is to solve the ODE systems on a massively parallel processing (MPP) architecture whilst simultaneously solving the mechanical problem on a symmetric multi processing (SMP) architecture. To make this approach feasible would require very rapid communications between the HPC resources, with minimal latency, jitter and packet loss. The intention was to use UKLight to couple the MPP and SMP resources (respectively the HPCx and CSAR facilities), to provide an optimal architecture on which to execute the heart model.

3. Experiences

First, a large number of people external to the ESLEA and IB projects (such as systems administrators of the end systems and local networks at the relevant sites) were required to facilitate the first, relatively simple, use case. There was initially an understandable reluctance to connect the storage facility to UKLight in case the potentially very large workload from this connection caused a denial of service to other users. Eventually, however, the relevant parties agreed that the use case was feasible and would work on it.

To execute the first use case, it took on the order of a month to install the necessary Globus client software and obtain grid certificates to connect to the NGS Data Clusters at RAL. One of the factors with the software installation was that there were multiple sources of information, which were slightly different. For example, there is a generic installation guide from Globus, but this contains none of the specifics of the UK NGS. There was also a step-by-step instruction guide produced by the IB project, but this was in part out-of-date as the UK NGS was in a ‘state of

flux’ at the time. In addition, advice involved reinstalling the Globus software as root—as non-root installs were not recommended; this, obviously, involved a significant amount of rework.

Eventually, these problems were overcome and it was possible to make some baseline measurements for the data transfer use case. The next task was to make the same measurements with the data routed over the UKLight network.

Even when UKLight was correctly provisioned, tests revealed that the end-to-end data transfer was an order of magnitude below what was expected. Diagnosing this problem proved very difficult. The ESLEA e-Health project did not initially have login access to the relevant nodes of the SRB storage facility, a prerequisite to performing the most basic of diagnostics. The security policy associated with the SRB prohibited external login, but after a period of time we successfully negotiated the login access we required.

Diagnostic tests could then be performed. The end-to-end network was tested using the iPerf program. iPerf needed to be compiled on HPCx which runs the AIX operating system, but this did not work out-of-the-box. The HPCx help desk could not resolve the issue, although it was at least responsive. Members of the wider ESLEA community were drawn upon for support with compiling iPerf and this proved extremely fruitful. The diagnostic tests revealed that the network was the bottleneck, rather than the end systems.

Unfortunately, drilling down into the network issue was not possible for the ESLEA e-Health project. The ESLEA project is not aware of the topology of the network between the end systems and the UKLight point of presence, and in any case has no administrative rights to the relevant network devices, which would normally be needed to run diagnostics. Therefore, the problem was passed to the local area network support. This problem was never resolved.

The second use case required modifications to the visualisation software being developed for IB. Instead of post-processing simulation results from a data file, it was required to work with a stream of data being received from the network. Understandably, those responsible for the visualisation software were not willing to make these changes, and it would have been difficult for anyone not already working on the software to make such changes; as such, the use case was dropped.

The software for the third use case was developed by the ESLEA e-Health project working closely with members of the IB project. The software took longer than expected to be developed. The scope of the software was reduced in that just an electrical model was developed. Nevertheless the software was capable of being run in parallel over the UKLight network, using MPICH-G2: a grid-enabled implementation of the Message Passing Interface for parallel computation.

CSAR went out of service before the end of the project, and therefore an attempt was made to run the software between the TeraGrid and HPCx. However, due to issues associated with the compilation of the large software stack, this was not achieved within the time-frame of the project.

Leveraging the IB experience, we see that collaboration within the heart modelling community stems from the development of a trust between scientists—and often these collaborations require the users to develop ‘experiments’ where disparate groups are required to run simulations, visualise results and contribute to joint papers. Scientists invite colleagues to participate in an experiment and require deployed systems and networks to ensure that these experiments, meshes, generated data and visualisation results are kept secure from competitors. The utilisation of high performance computing facilities requires users to manage accounts on multiple facilities and transport information across the world to secure data stores for real time visualisation or later analysis.

To realise the potential of utilising UKLight or similar networks for applications drawn from projects such as IB requires significant effort from many different parties—some of whom will never benefit (directly or indirectly) as a result of the activity. As often seems the case with large-scale projects, the political and social challenges have the potential to outweigh the technical ones. And if our experiences from ESLEA tell us one thing it is that if the potential of using switched lightpath networks for distributed heart modelling is to be realised fully, the political and social challenges will have to be tackled before any meaningful attempt can be made at tackling the technical ones.

4. Discussion

The e-Health ‘mini-project’ was an ambitious attempt to demonstrate the potential of switched lightpaths for e-Health applications. The chosen application was, in the authors’ opinion, the correct one—the potential benefits for the heart modelling community are significant, especially if the resources provided by the IB infrastructure are being leveraged. Unfortunately, the hurdles faced—coupled with a lack of ‘buy-in’—meant that progress was difficult. Our work has demonstrated that while other application areas have a pressing need for switched lightpaths *now*, the urgency of the need for them within the e-Health context is not quite as crucial—but we would argue that it is on its way.

References

- [1] A. J. G. Hey and A. E. Trefethen. The UK e-Science programme and the grid. *Proceedings of Computational Science (ICCS 2002), Part I* 3–21, Springer-Verlag Lecture Notes in Computer Science volume 2329, 2002.
- [2] C. Greenwood, V. Bartsch, P. Clarke, P. Coveney, C. Davenhall, B. Davies, M. Dunmore, B. Garrett, M. Handley, M. Harvey, R. Hughes-Jones, R. Jones, M. Lancaster, L. Momtahan, N. Pezzi, S. Pickles, R. Pinning, A. C. Simpson, R. Spencer, and R. Tasker. Exploitation of switched lightpaths for e-Science applications (ESLEA). *Proceedings of the 2005 UK e-Science All Hands Meeting*.
- [3] L. Momtahan, S. Lloyd, and A. C. Simpson. Switched lightpaths for e-health applications: issues and challenges. *Proceedings of the IEEE Symposium on Computer Based Medical Systems 2007*.
- [4] L. Momtahan and A. C. Simpson. Switched lightpaths for e-health applications: a feasibility study. *Proceedings of the IEEE Symposium on Computer Based Medical Systems 2006*.
- [5] D. J. Gavaghan, A. C. Simpson, S. Lloyd, D. F. Mac Randal, and D. R. S. Boyd. Towards a grid infrastructure to support integrative approaches to biological research. *Philosophical Transactions of the Royal Society: Mathematical, Physical and Engineering Science* 363(1883):1829–1841, August 2005.
- [6] S. Lloyd, D. J. Gavaghan, A. C. Simpson, M. Mascord, C. Sieunarine, G. Williams, J. Pitt-Francis, D. R. S. Boyd, D. Mac Randal, L. Sastry, S. Nagella, K. Weeks, R. Fowler, D. Hanlon, J. Handley, and G. de Fabritis. Integrative Biology: the challenges of developing a collaborative research environment for heart and cancer modelling. *Future Generation Computer Systems* 23(3):457–465, March 2007.