From INWA to INCA: an international collaboration in eSocial Science

A.D. Lloyd[1,2], V.Maxville[1], Y. Sun[1] and T.M. Sloan[3]

[1]Curtin Business School, Curtin University of Technology. GPO Box U1987, Perth, Western Australia 6845

[2]Management School, The University of Edinburgh, 50 George Square, Edinburgh EH8 9JY, UK

[3]EPCC, The University of Edinburgh, James Clerk Maxwell Building, Mayfield Road, Edinburgh EH9 3JZ, UK

Email address of corresponding author: Ashley@curtin.edu.au

Abstract. The INWA Grid went ‘live’ in December 2003, connecting Curtin Business School in Western Australia through the US to EPCC in Edinburgh, making it one of the longest distance and longest running Grid-based ‘collaboratories’ in the Social Sciences. Stable network connections between Australia and China were established in January 2005, allowing the extension of INWA to the network centre of the Chinese Internet at the Computer Network and Information Center of the Chinese Academy of Sciences (CNIC, CAS). The present collaboration with CNIC - Project INCA - focusses on innovation processes that can be supported by Grid technologies across a ‘virtual cluster’. This paper describes the current INWA infrastructure, which is migrating from 2nd to 3rd generation Grid technologies, and provides an increasingly homogeneous platform for collaboration across distinctively different socio-economic contexts. Here the focus of ‘innovation’ moves from enhancements in technical capability associated with enhanced computing and telecommunications, to their relationship with ‘innovation capacity’ – i.e. the ability to expand potential for innovation and extend an innovator’s reach. To explore this relationship, we outline a framework that is being used to track the capacity of this ‘virtual organisation’ as each advance in 2nd to 3rd generation technology is implemented.

Introduction

The INWA Grid project was initially funded by the Economic and Social Research Council as part of their “Pilot Projects in e-Social Science” and established a ‘virtual collaboratory’ connecting research groups located in three different continents: the UK (EPCC, Edinburgh), Australia (Curtin Business School, Perth) and China (Computer Network Information Center, Chinese Academy of Sciences, Beijing).

In establishing this infrastructure a number of social, legal and technical challenges were identified (Lloyd et al., 2005; Sloan & Lloyd, 2005, Lloyd, 2005): quality of international network connectivity, differences in domestic internet addressing and security policies, security of Grid middleware, and usability by Social Scientists. In some cases reconfiguration of existing Grid technologies was not sufficient to allow interoperation between sites and

Figure 1. The three nodes of the INWA Grid.

development of new software was required, including security modules for ‘transfer queue over Globus’ and an ‘open database connectivity’ driver for transparent access to Federated Grid databases (Jackson et al., 2005). The socio-legal challenges were more profound but are beyond the scope of the present paper.

This paper reports an Australian Research Council project in collaboration with Sun Microsystems and SingTel Optus (LP0454322) to apply the INWA (Innovation Node: Western Australia) infrastructure to support INCA (Innovation Node: Chinese Academy of Sciences). Here the focus of ‘innovation’ moves from enhancements in technical capability associated with distributed computing and enhanced telecommunications to their relationship with ‘innovation capacity’ – i.e. the ability to expand potential for innovation and extend an innovator’s reach (Carley, 1996).

The new forms of collaborating and organising implicit in the above statement aligns well with the UK research agenda for eSocial Science outlined by Procter (2004), with its focus on ‘facilitating collaboration both within and between research communities’ and emphasis on National Centre for eSocial Science’s (NCeSS) role in terms of ‘exploring and validating research innovations’ and ‘understanding and managing innovation’.

However, achieving the objective of forming ‘new research communities and networks’ within a sparse, globally distributed, Social Science community is at odds with recorded histories of innovation capacity development at an industry or national level (Marshall, 1890; Marshall, 1907; Schumpeter, 1934; Porter 1990; Mowery & Nelson, 1999) that emphasise scale and concentration at a local level. Such scale is, perhaps self-consistently, typically present in international Science collaborations (Forbes & Abrams, 2004) whose needs have been clearly expressed for the distributed computing community to develop supporting technologies. Though the resulting ‘Global Grid’ may be viewed as ‘general purpose’ technology (Bresnahan & Trajtenberg, 1995) and hence capable of supporting new eSocial Science communities, the expansion of existing small-scale research networks may increase the volume of communications at the cost of improved diffusion of ideas (Carley, 1996).

More recent reviews focussing on comparative analysis of eScience and eSocial Science practice have reported disciplinary differences evident in access to infrastructure by eSocial Scientists and the type of applications to emerge for eSocial Scientists (Borgman, 2005; Olsen, 2005), with few applications clearly identifying the benefits of operating within a Grid environment versus a Web environment.

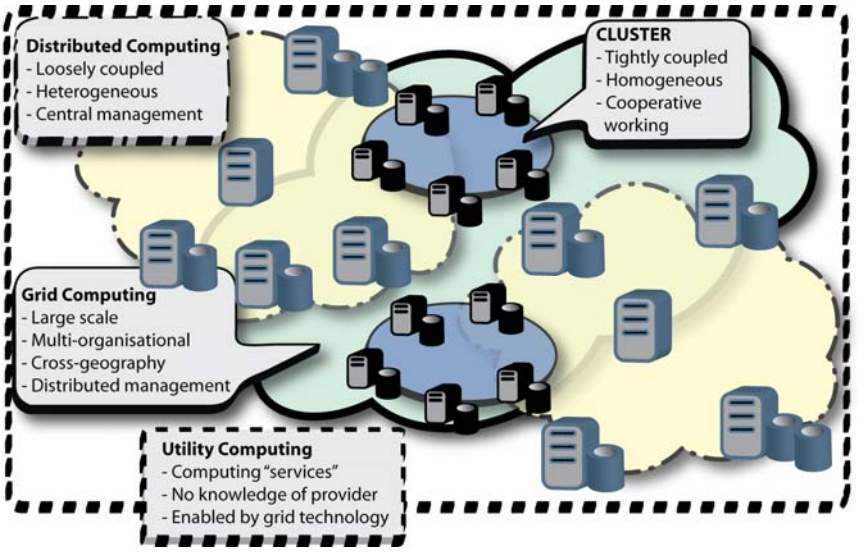

Figure 2. Grid and related computing paradigms, adapted from Grimshaw (2006).

This difficulty of establishing the benefit of technical advances such as a ‘Grid Service’ over a ‘Web Service’ (Lloyd, 2004) is echoed in the Enterprise Computing treatment of the Grid’s benefits to Business (see for example http://www-1.ibm.com/grid/about_grid/benefits.shtml), where collaboration is identified as both a technology and a business benefit, but in neither case is it distinguished from earlier technologies in terms of a unique technical advance that leads to a unique improvement in the ‘reach’ of a user.

Grimshaw (2006) clarifies the distinctiveness of Grid computing by charting the evolution of high performance computing from clusters in a local area, to loosely coupled distributed computers over a wider area, arguing that it is Grid computing that allows these different types of computing resource to be virtualised, managed and shared (Figure 2).

As these systems move from small numbers of identical computers tightly coupled into clusters and managed as a single resource in a single location, through to distributed computing across larger numbers of heterogeneous platforms, we still retain centralised management.

Moving to a Grid environment however offers the potential of coupling orders of magnitude more platforms and to make these interoperate effectively, management tasks have to be distributed. The trust relationships required in such networks are established through information exchange, typically with reference to a ‘trusted’ third party determined by local security policies, e.g. atypical use of Kerberos authentication in a Grid implemented for the US DoD Major Shared Resource Centers (Grimshaw and Natrajan, 2005).

In this typology, Grid Computing is presented as the ‘glue’ that makes the interoperation of different architectures of systems distributed in different ways possible, however, whilst Grid Computing is intended to enable the management of trusted and secure distributed resources, it is also intended to underpin ‘Utility Computing’ in which services are consumed on a commodity basis with no knowledge of the provider required by the client.

The consequently large scope and scale of interactions that can be described as ‘Grid Computing’ or supported by the ‘Grid’, with workflow and state requirements that can vary from minutes to months, based on as yet untested business models for ‘Commercial Data

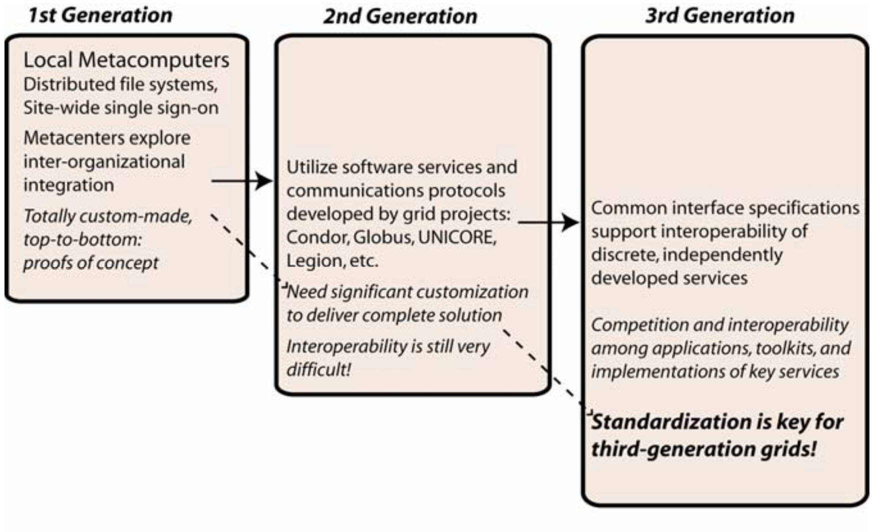

Figure 3. Three generations of Grid Computing, attributed to Catlett, C., adapted from Grimshaw (2006).

Centers' and 'Grid Resource Reseller' (Foster et al., 2004) represent a significant challenge in terms of building pervasive access to this infrastructure for users outside 'big science'.

Key to this process in a multi-vendor, heterogeneous world, are common interface specifications, an established requirement for collaborative systems (Navarro, Prinz and Rodden, 1993). However, even amongst 'Metacenters' (Figure 3) interoperability has traditionally been achieved by 'stove-piping'. This couples systems together to prove concepts rather than enable extensibility. Even at the end of the 2nd generation there was "little or no interoperability [and] I think no is a better phrase!" (Grimshaw, 2006).

Grid Computing is now at the beginning of its 3rd Generation (Figure 3). Since 2004, Grid Services have been treated as special instances of Web Services allowing Grid Services to be built on top of a relatively mature stack. This 'Specifications Landscape' is shown in Figure 4, where the Web Services Foundation is shown as either 'Standard' or 'Evolving'. In contrast the Core Grid Services underpinning the spectrum of computing that ranges from 'Grid' to 'Utility' has major gaps that need to be addressed before 'Virtual Organisations' detailed in the OGSA Use Cases (Foster et al., 2004) can be functionally supported.

In terms of the efficiency with which a 'virtual organisation' of highly distributed and generally small teams of eSocial Scientists (Forbes and Abrams, 2004) who want access to high-performance computing, the challenges are even more profound: "no matter how much money I spend it's still going to be 70 milliseconds from the East Coast of the US to the West Coast [and] 180-200 milliseconds from the East Coast to Japan – no amount of money is going to change that and so you have to plan for that [..] otherwise the performance is going to be really bad" (Grimshaw, 2006). Since the scope of Social Science is typically much greater than its scale, it seems like likely that the current inequity of access reported for eSocial Scientists by Borgman (2005) and Olsen (2005) will continue until business models emerge for Utility Computing that aggregate enough demand for a more pervasive infrastructure to become economic.

Success, as Olsen (2005) notes: "is not inevitable as technologies evolve [and arises from] a mix of social and technical factors", with Bergman (2005) expanding on a range of incentives to collaborate or not to collaborate that echo Carley's (1996) observations that 'innovative groups, research and development organizations, and intellectual communities may wish to

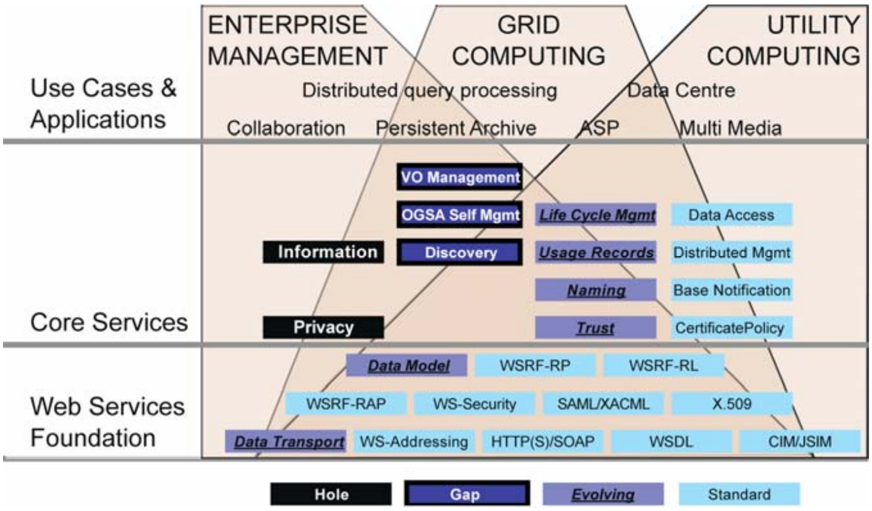

Figure 4. 'Specifications Landscape - April 2006', adapted from Grimshaw (2006). Note that this landscape is described by Grimshaw as 'highly volatile' and has been re-cast here only to illustrate the evolving nature and status of Grid standards.

restrict membership so that the new ideas they generate will diffuse first within their group', and further that geographical asymmetries may exist even where technology platforms are interoperable as 'if learning novel information first is advantageous economically or socially then cultural or knowledge restrictions will be advantageous'.

In the context of this evolving landscape, the INWA Grid reported here may be regarded as a special case where migration of tools and techniques from eScience to eSocial Science enforces a high degree of alignment between the planned needs of eScience and the emergent needs of eSocial Scientists, who in this case have had to establish only one node where no alternative existed at the project's inception: in Perth, Western Australia.

The INWA Grid: an infrastructure for innovation?

Just as 'third generation' Grids assume common interfaces to enable interoperation across functional, institutional and national boundaries, so too "third generation innovation policy assumes [the ability] to release the potential for innovation that is embedded in other sectors or policy domains [i.e.] that coherence may be achieved by ensuring cross-sectoral optimisation of the components of various sectors' innovation policy through co-ordination and integration." (OECD, 2005).

This convergence appears encouraging as it offers a broader base on which to value investments in Grid technologies and infrastructure. However: "Innovation policy in OECD countries has mostly been seen as an extension of R&D policy" (OECD, 2005). This promotes a linear view of the relationship between R&D expenditure and innovation capacity, illustrated in Figure 5.

Whilst a linear view is accepted to have weaknesses (Gibbons et al., 1994) the relative importance of R&D intensity over the shape of its (regional) distribution means that pathways for international innovation are often 'black-boxed' when conceptualising national innovation systems. This is true of the UK Science and Innovation Investment Framework 2004-2014 (HMSO, 2004; Storey, 2005), and consequently any interfaces with international innovation systems are (a) harder to isolate, and (b) harder to assess in terms of their coherence. In turn

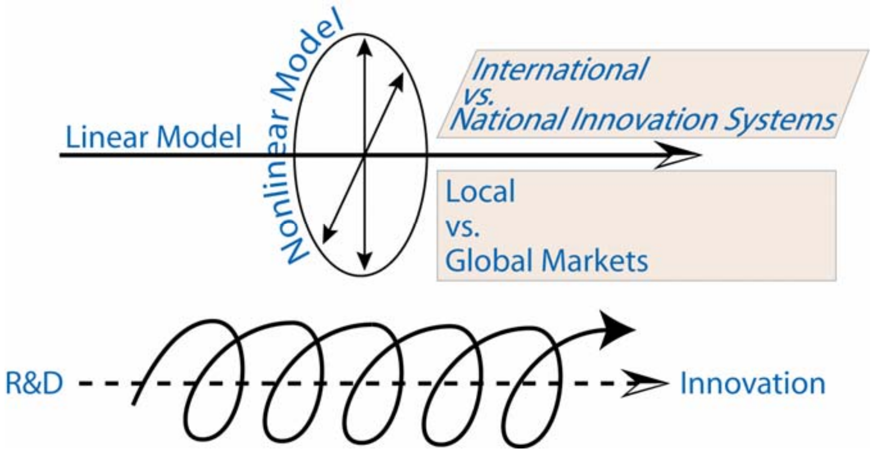

Figure 5. Linear versus nonlinear views of innovation systems. As markets globalise and innovation systems are coupled, the relationship between R&D inputs and outputs recognised ex post as Innovations becomes more complex.

this makes the outcomes of investment in such systems, or underlying technologies such as Grid Computing harder to predict and hence to promote

Nevertheless, as part of the project we have engaged in tracking and analysing the evolution of the infrastructure, reflecting on decisions and their rationale, as well as recording 'lessons learned' as part of the process. For example, in the INWA Grid, it was necessary to establish a node in Perth, Western Australia, where no alternative existed, essentially to transfer technology from EPCC (Edinburgh). However, in the extension of INWA (Perth) to INCA (Beijing) a decision was made to install a separate system in the Computer Network and Information Center at the Chinese Academy of Sciences to ensure that the Grid could be managed as a single entity, with common policies and configurations, allowing the focus of the collaboration to be on complementary skill sets supported by a suitably scaled high availability system, rather than efficient access to existing resources per se.

The degree of complementarity was explored as part of the baseline survey of each node's characteristics and socio-economic context. To assess the relative distribution of technical skills of the academic members of the teams physically located in the EPCC and INWA nodes, each staff member was asked to rate their experience and expertise in each of the following areas:

- [B] Business
- [DM] Data Mining
- [KM] Knowledge Management
- [CS] Computer Science
- [CP] Computer Programming
- [G] Grid Computing
- [HPC] High Performance Computing
- [Com] Communications

The summed results (Figure 6) indicate that the groups in EPCC and INWA have quite different skill orientation, though a strong correlation is evident in Computer Programming.

A second level of interactions occurs with the social networks that exist between the team members. The respondents to our survey were asked to identify five key team members locally and non-locally. In Figure 7, the resulting list of members of the INWA project is provided. Members from the UK are indicated as triangles, those from China as circles, and Australian members as squares. Within each of these groups, subgroups are circled, showing which members are employed by the respective institutions: EPCC (Edinburgh), CNIC/CAS (INCA, Beijing) and Curtin University of Technology (INWA, Perth).

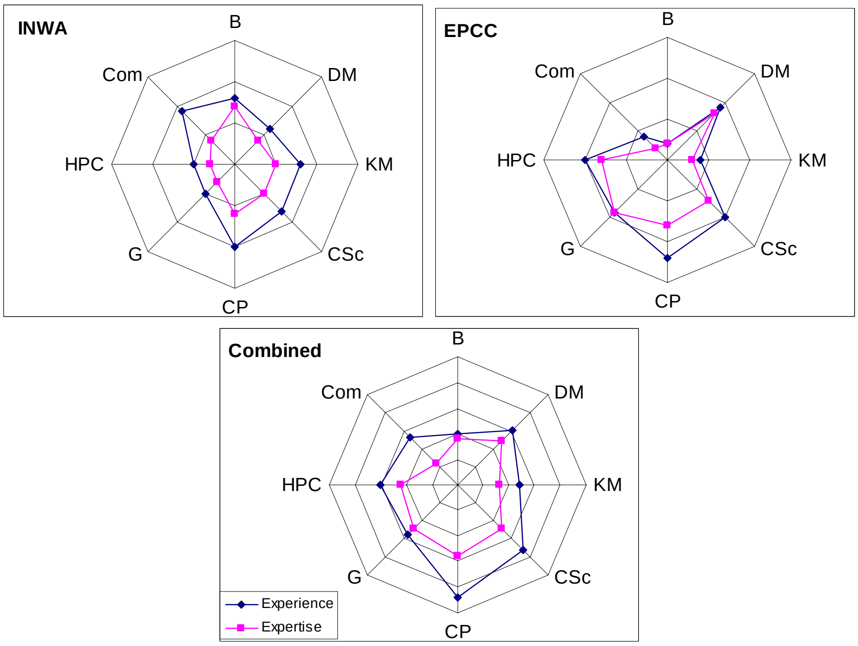

Figure 6: Experience and Expertise in the INWA and EPCC nodes. Cross-section: first quarter 2006. Note that sample sizes are identical in both sites and hence the results have not been normalised in this instance.

Figure 7 charts the responses from six team members (outlined in bold) and provides an indication of the difference in structure in the three groups. The members with the most references are the three in the centre who are the principle connections in each node. The INCA node in China has much fewer members and links as the network is still forming. The people identified in the UK group are entirely within EPCC and each member of the node has relatively few links. A combination of Figures 6 and 7 would suggest that the EPCC node has been the primary source of technical expertise, requiring a smaller number of external links, whilst the INWA node in Perth is managing the collaboration. However, the principle value of this representation is in tracking the network that evolves as part of the collaboration, rather than determining the net contributions by each partner as the relative sizes of the organisations collaborating are significantly different, as is the breadth of expertise available at each site.

It is this broader context to the INWA infrastructure that provides a link to the literature on industry cluster formation and competitiveness, from ‘local clusters’ of the type observed by Marshall (1890) to ‘national innovation systems’ (Mowery and Nelson, 1999). Furman et al. (2002) argue that these multiple perspectives are consistent with Porter’s (1990) framework describing the innovation orientation of national clusters (Figure 8).

The emphasis in Figure 8 is Porter’s and used to highlight the importance of ‘information infrastructure’ on competitiveness, yet this work pre-dates the rapid expansion of the Internet and adoption of the World-Wide Web - raising questions over its sensitivity to developments in information and communication technologies and their convergence, which as Figures 2-4 indicate, are developing rapidly in terms of complexity.

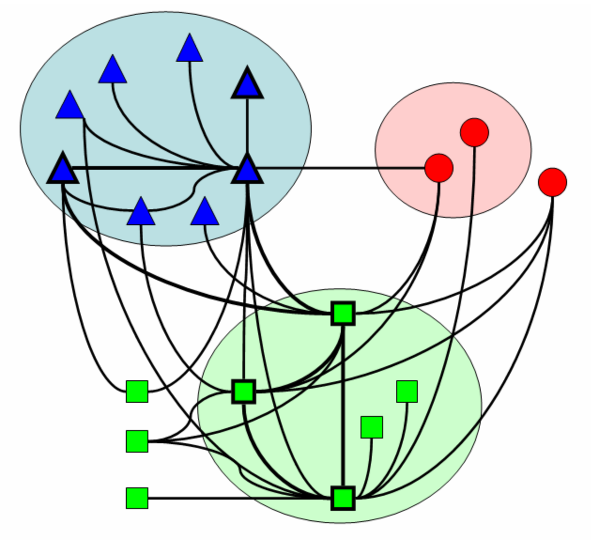

Figure 7: INWA project social network. Cross-section: first quarter 2006.

Furman et al. (2002) also appear to embed an assumption in their argument that the locus of innovation arising from clustering is in product innovations that are developed and commercialised, and not in the commercialisation process itself. This assumption is emerging as contentious, with a study reported in the Sloan Management Review indicating that clustered firms generate fewer new products for their R&D expenditure than non-clustered firms, and hence their success in revenue terms arises not from technological innovations; rather '*they get an edge in finding markets and customers and tailoring products for them*' (Yu 2002).

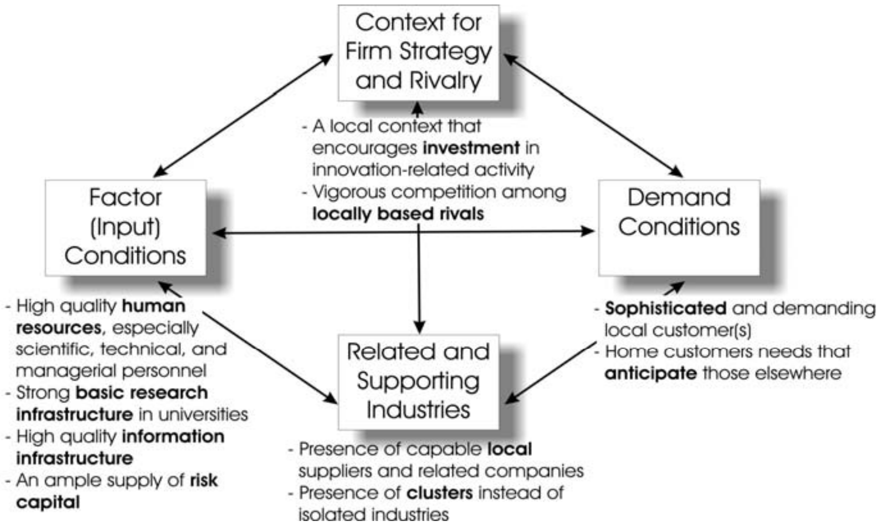

Figure 8. The innovation orientation of national industry clusters (Furman, et al., 2002).

This accords with Phillips (2002) view of clusters that are recognised as global centres of excellence, such as the Saskatoon-based biotechnology-centred innovation cluster, where '*it* *is clear from the evidence, however, that the innovation cluster is not in any way independent or self sufficient [..] with the result that less than half the value added to the product is added locally'* (Phillips 2002).

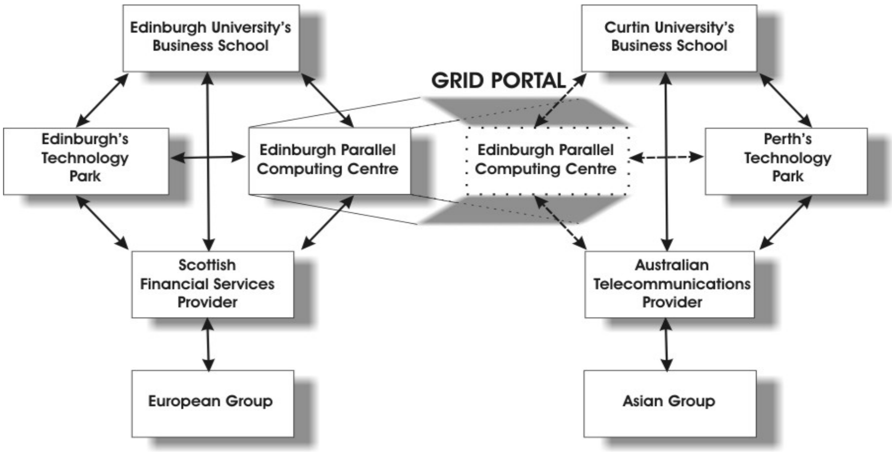

Figure 9. Distributing high-performance computing capabilities and expertise (dashed lines) between clusters in Scotland and Western Australia.

Yu (2002) and Phillips (2002) may be viewed as consequences of globalisation of markets and supply chain raising the relative contribution to competitiveness of distributed organisational forms operating as 'virtual clusters' where efficiency gains are realised through minimising (though not eliminating) redundancy. This complementarity is illustrated in the representation of the INWA Grid in Figure 9, where the Grid is used to export a capability required to establish a 'virtual' complementary cluster and thereby enable exchange of expertise.

A review of Figures 1-9 illustrates the scale and scope of interactions that are anticipated for Grid technologies, some of the complexity (and flexibility) of the technologies being developed and the social networks that evolve. In outline these describe the framework with which we are tracking current developments and their impact on the operation of the 'virtual' cluster.

These should inform an understanding of the relationship between the underlying technologies, capacity for innovation and innovation outputs, in similar technological contexts but distinctive socio-economic contexts, which include the network centre of the fastest growing economy in the world and the world's most isolated (by geography) capital city.

Since the INWA Grid is managed as a single entity, the common-mode events that we are able to track and contrast focus on the evolving Grid computing and communications technologies and their standards that comprise the architecture of the INWA Grid (Figure 10).

In terms of the Grid technologies and application that comprise the INWA Grid in Figure 10, a number of enhancements are currently possible that may be described as 3rd generation technology:

OGSA-DAI/OGSA-DQP – Upgrade to the latest version of OGSA-DAI (Open Grid Services – Data Access & Integration) and implement OGSA-DQP to allow Distributed

Query Processing (DQP). This upgrade supports queries over Grid Data Services (GDS)s allowing data access and analysis to be combined, improved efficiency of dynamic accesses to federated databases, including parallel queries, as well as consistent access to metadata.

Grid Data Miner – Installation of Data-Mining software written expressly for a Grid environment allows comparison with current ‘Grid-migrated’ Data-Mining software

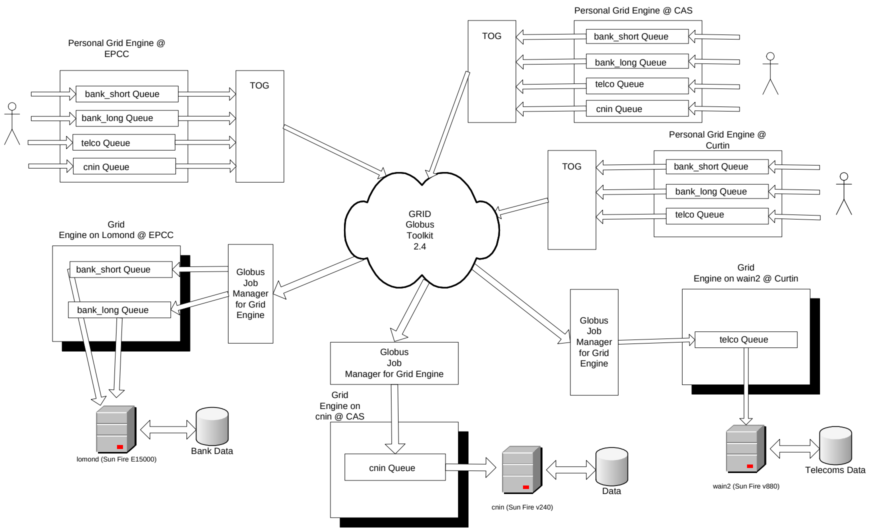

Figure 10. Technology schematic of the INWA Grid infrastructure connecting EPCC (Edinburgh, UK), Curtin Business School (Perth and the Computer Network Information Centre at the Chinese Academy of Sciences, Beijing).

Globus Toolkit GT4 – Upgrade to Globus Toolkit v4, open source Grid middleware that in this release is being re-targeted to the Enterprise community, with improved documentation, installation, administration and security. An explicit focus is ‘Web services that represent stateful resources on the Grid’ allowing differentiated applications to be developed within a Grid environment whilst allowing legacy applications to be supported. Critically, for global collaborative analysis of large data sets GridFTP has been re-written and demonstrated at data transfer rates of up 90% of the theoretical limit (www.globusconsortium.org).

Enhanced global network - Significant enhancements in access to network infrastructure between Perth and Singapore (4x155Mbps Perth-Singapore, 622 Mbps Singapore to Frankfurt)and hence onwards to Edinburgh) provide an alternative pathway for the INWA infrastructure that became available in 2006.

IPv6 – the migration to an ‘all IPv6’ infrastructure is expected to eliminate some of the heterogeneity in the INWA Grid, and specifically the addressing and security differences previously reported for the INCA node (Sloan & Lloyd, 2005)

Each of these technology enhancements of the INWA Grid might be expected to have an impact (longitudinally) on the innovation capacity of virtual organisation supported, providing an event study that may be tracked using the tools described above.

Conclusion

The INWA Grid represents one of the longest distance and longest running Grid-based 'collaboratories' in the Social Sciences. The present collaboration with the Computer Network Information Center of the Chinese Academy of Sciences, Project INCA, focusses on Innovation processes that can be supported by Grid technologies across a 'virtual cluster' with other INWA nodes.

The objective of enhancing 'the ability to expand potential for innovation and extend an innovator's reach' (Carley, 1996) is common to policies for both national and international innovation systems. However the artificial separation of these policies and tendency to view innovation policy as an extension of R&D policy makes it harder to "release the potential for innovation that is embedded in other sectors or policy domains" that is required of 3rd generation innovation policies (OECD, 2005). Further, it makes it harder to value the enhancement of capacity for innovation attending 3rd generation grid technologies.

The potential impact of this on investment in infrastructure for eSocial Science communities is particularly significant give the generally small scale but wide scope of eSocial Science collaborations and the 'gaps' in the stack (Figure 4) that relate to the analysis of social science data. Whilst 4th generation Grid technologies that support Utility Computing may broaden access to 'HPC on demand', the loosely coupled, highly distributed management environment anticipated may make this type of access to HPC even less compatible with the generally more stringent security and privacy constraints applied to the analysis of socio-economic rather than scientific data.

To help clarify the potential contribution of anticipated Grid technologies to enhancing the 'reach' of eSocial Scientists, we have placed the INWA Grid within the current landscape of grid technologies and standards. In combination with regional network enhancements in East Asia, this has allowed us to identify '3rd generation' grid technologies that can be adopted.

To explore how these enhancements may be related to the innovation capacity of the 'virtual cluster' supported by INWA, and thereby to help to inform the valuation of national investments in international innovation infrastructures, we have also outlined a framework that is being used to capture characteristics of the cluster as it evolves in relation to those technology enhancements.

Acknowledgments

The authors would like to thank A. Hume and M. Jackson of EPCC, B. Yan and K. Nan at the Computer Network and Information Center of the Chinese Academy of Sciences (CNIC, CAS), and D. Qian, Beihang University. We gratefully acknowledge the support of the UK Economic and Social Research Council (award RES-149-25-0005) for the initial phase of the INWA Grid; the Australian Research Council in partnership with SingTel Optus for its current support (award LP0454322) and Sun Microsystems and AARNet for continued support since the INWA Grid became operational in 2003.

References

Borgman, C. (2005): 'Disciplinary Differences in e-Research: An Information Perspective', First International Conference on e-Social Science, Manchester, June 2005.

Bresnahan, T. F and Trajtenberg, M. (1995): ‘General purpose technologies ‘Engines of growth’?’, *Journal of Econometrics*, vol. 65, issue 1, January 1995, pp. 83-108.

Carley, K. M. (1996): ‘Communicating New Ideas: The Potential Impact of Information and Telecommunications Technology’, *Technology in Society*, vol. 18, no. 2, pp. 219-230.

Catlett, C. (2004): ‘Making Grids Real: Standards and Sociology’, Grid and Cooperative Computing, GCC 2004: Third International Conference, Wuhan, China, Oct. 21-24, 2004

Forbes, I. and Abrams, D. (2004): ‘International social science research: Craft industry or baby behemoth?’, *International Social Science Journal*, 180, 227-244.

Foster, I., Gannon, D., Kishimoto, H. and Von Reich, J. J. (2004): ‘Open Grid Services Architecture Use Cases’, *GFD-I.029*, Global Grid Forum, October 2004. Furman, J. L., Porter, M. E. and Stern, S. (2002): ‘The determinants of national innovative capacity’, *Research Policy*, vol. 31, 2002, pp. 899-933.

Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P. and Trow, M. (1994): *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*, Sage, ISBN: 0803977948

Grimshaw, A. (2006): ‘Defining the Grid: An Introduction to the Open Grid Services Architecture’, *Global Grid Forum OGSA Working Group*, Webcast, April 26, 2006. Available: http://www.presentationselect.com/hpinvent/archive.asp?eventid=1188

Grimshaw, A. S. and Natrajan, A. (2005): ‘Legion: Lessons Learned Building a Grid Operating System’, *Proceedings of the IEEE*, vol. 93, no. 3, pp. 589-603.

HMSO (2004): *Science & innovation investment framework 2004 – 2014*, HMSO, July 2004, ISBN: 184532031X

Jackson, M. J., Lloyd, A. D. and Sloan, T. M. (2005): ‘Enabling Access to Federated Grid Databases: An OGSA-DAI ODBC Driver’, AHM2005: Innovating through e-Science, Nottingham, September 2005.

Lloyd, A. D. (2004): ‘The Grid, Virtual Innovation and the Tyranny of Proximity’, in J. Cha and M. Sobolewski (eds.): *The Worldwide Engineering Grid*, Balkema Publishers 2004.

Lloyd, A.D. (2005): ‘The Grid and CRM: from ‘If’ to ‘When’’, *Telecommunications Policy*, vol. 29, pp.153-172.

Lloyd, A. D., Nan, K., Qian, D., Sloan, T. M., Sun, Y. and Yan, B. (2005): ‘Linking e-Science capabilities for e-Social Science communities: extending the UK-Australia INWA project to the Chinese Academy of Sciences’, First International Conference on e-Social Science, Manchester, June 2005.

Marshall, A. (1890): *Principles of Economics: an introductory volume*, Macmillan, London.

Marshall, A. (1907): *Elements of economics of industry*, Macmillan, London.

Mowery, D. and Nelson, R.R. (1999): *Sources of Industrial Leadership: Studies of Seven Industries*, 1999, Cambridge University Press, Cambridge, MA.

Navarro, L., Prinz, W. and Rodden, T. A. (1993): ‘CSCW requires open systems’, *Computer Communications*, vol. 15, no. 5, May 1993, pp. 288-297.

OECD (2005): *Governance of Innovation Systems: Synthesis Report*, OECD, ISBN: 9264011021.

Olson, G. (2005): ‘Long-Distance Collaborations in Science: Challenges and Opportunities’, First International Conference on e-Social Science, Manchester, June 2005.

Phillips, P. (2002): ‘Regional systems of innovation as a modern R&D entrepot: the case of the Saskatoon biotechnology cluster’, in J. Chrisman et al. (eds.): *Innovation, Entrepreneurship, Family Business and Economic Development: A Western Canadian Perspective*, 2002, University of Calgary Press.

Porter, M.E. (1990): *The Competitive Advantage of Nations*, Free Press: New York.

Procter, R. (2004): ‘Developing the NCeSS Research Agenda’, All-Hands Meeting, Manchester, 6th July 2004.

Schumpeter, J. (1934): *The theory of economic development: an inquiry into profits, capital, credit, interest, and the business cycle*, 1934, Harvard University Press, Cambridge, Mass.

Sloan, T. M. and Lloyd, A. D. (2005): 'Extending the INWA Grid', First International Conference on e-Social Science, Manchester, June 2005.

Storey, D. (2005): 'Science & innovation investment framework 2004-2014', Institute of Physics Seminar, 9 February 2005.

Yu, L. (2002): 'How location clusters affect innovation', *MIT Sloan Management Review*, vol. 44, no. 1, 2002, pp.16.