How web community organisation can help grid computing[♦]

Miranda Mowbray
Enterprise Systems and Storage Laboratory
HP Laboratories Bristol
HPL-2006-34(R.2)
July 25, 2007*

community practices, grid, peer-to-peer

OurGrid is a web-based community whose members can use each others' spare computing power. When an OurGrid member is not using his own computer, it can be used by any other member. This paper describes how the community aspects of OurGrid have been crucial for its success. It argues that grid economies, which provide another method for sharing computing power, might also benefit from having a web-based community structure.

* Internal Accession Date Only
♦ International J. Web Based Communities, Vol.3, no. 1, 2007
Approved for External Publication

How web community organisation can help grid computing

Miranda Mowbray

Enterprise Systems and Software Lab
Hewlett Packard Laboratories, Bristol
Filton Road, Stoke Gifford
Bristol BS34 8QZ, UK
Fax: +44 (0)117 312 8003
E-mail: miranda.mowbray@hp.com

Abstract: OurGrid is a web-based community whose members can use each others' spare computing power. When an OurGrid member is not using his own computer, it can be used by any other member. This paper describes how the community aspects of OurGrid have been crucial for its success. It argues that grid economies, which provide another method for sharing computing power, might also benefit from having a web-based community structure.

Keywords: grid; peer-to-peer; community practices.

Reference to this paper should be made as follows: Mowbray, M. (2007) 'How web community organisation can help grid computing', *Int. J. Web Based Communities*, Vol. 3, No. 1, pp.44–54.

Biographical notes: Miranda Mowbray is a Technical Contributor at Hewlett Packard's European research laboratories, in Bristol, UK. She co-founded E-mint (www.emint.org.uk), the association for online community professionals, and co-wrote *Online Communities: Commerce, Community Action, and the Virtual University*, published by Prentice-Hall. Her PhD is in Algebra, from London University. Mowbray's research papers have appeared in conferences and journals in the areas of mathematics, philosophy, business studies, theoretical biology, law, sociology, and computer science. Her current research interests are grid computing and peer-to-peer systems. Her travel writing has appeared in three books in O'Reilly's *Travelers' Tales* series.

1 Introduction

This paper is concerned with OurGrid, a web-based community whose members can use each others' spare computing power. When an OurGrid member is not using his own computer, it can be used by any other member. OurGrid has been described in several previous publications[1] in particular, for an OurGrid primer, see Cirne *et al.* (2006). However, this paper (together with the much shorter conference paper that it is based on) is the first to focus on OurGrid's web-based community aspects. OurGrid's exploitation of the fact that it is a web-based community has been crucial to its success.

This paper may be of interest to members of web-based communities who would like to know about community practices that have contributed to OurGrid's success, and so might also be useful for their own communities. It also may be of interest if you would like to increase the computer power that you can use, in exchange for making your spare computing resources available to others.

1.1 Previous approaches to sharing computing power

For decades, distributed computing has allowed computing power to be shared within an organisation. A user of a networked computer can send applications to be carried out by other computers in the same organisation. More recently a more ambitious architectural vision has arisen, that of *grid computing* (Foster *et al.*, 2001; Berman *et al.*, 2003) in which computing power can be shared transparently between different organisations.

There has been considerable effort into realising grid computing, and several grids are now in operation: for instance, CERN's LHC Grid (CERN, 2005), TeraGrid (Beckman, 2005), ChinaGrid (Jin, 2004), and the UK e-Science Grid (Hey and Trefethen, 2004). However, the middleware stack that these grids rely on is large and complicated, and as a result is regarded with suspicion by security experts. Joining any of these grids is a complex operation, and it requires human negotiation, so its complexities cannot all be delegated to a computer.

Once it is technically possible for your computer to be used by someone who belongs to a different organisation, the question arises as to why you would want to let him do that. An answer which has been widely discussed in the computing industry is that he would pay you for the use of your computer, in a secure, automated fashion, using a *grid economy* (Buyya *et al.*, 2000; Buyya, 2002). However, implementing secure and reliable global accounting and billing mechanisms for computing power shared between different organisations over a grid is a hard problem. Although work on grid economies is progressing, none of the grids mentioned above have yet developed full grid economies, and the elements required to support a future grid economy have sometimes proved problematic. In particular, the e-banking and cryptography infrastructure required can provide a single point of failure. In one grid the certificate revocation server went down twice in six months, and each time the entire grid stopped working.

An alternative approach to sharing computing power between different organisations is that of *volunteer computing* (Sarmenta, 2001; Anderson and Fedak, 2005). There have been quite a few examples of volunteer computing, but the most famous is probably SETI@home (Anderson *et al.*, 2002). Volunteer computing uses the idle time of volunteers' computers. Computer users typically have some periods of time during which they need to use their computer, and some periods of time during which they have no applications to run and leave their computer idle.

Suppose a project team wants to carry a large computation which is *embarrassingly-parallel*, that is, it consists of a large number of tasks that can be performed in parallel without any communication between the different tasks. (Such computations are called embarrassingly-parallel because they are so parallel that it is embarrassing.) Embarrassingly-parallel applications exist in many fields, including data mining, simulations, parameter sweeps, computer imaging, film animation, pharmaceutical development, computational biology, and meteorology. These applications typically have the property that the limit on the amount of computing power that they can profitably use is extremely high; for example, additional computing power may allow the computation to run faster or to produce more detailed (and hence more useful) output. To use volunteer computing, the team makes some special screen-saver code available for volunteer members of the public to download. When a volunteer's computer is not being used by the volunteer, and so would otherwise be idle, the screensaver code uses the volunteer's computer to perform some of the tasks, and sends the results back to the project team.

The screensaver code is lightweight, joining the system as a new volunteer is easy, and volunteer computing scales well up to millions of volunteers. The shared volunteer computing infrastructure BOINC (Anderson and Fedak, 2005) allows scientific projects to use volunteer computing without having to create their own infrastructure. The volunteers lose nothing by volunteering, since their computers would otherwise be idle. However, they also do not gain anything concrete by volunteering – except a good feeling from having contributed, for example, to improving malaria control in Africa (AFRICA@home, 2006).

Only projects which can persuade large numbers of people to volunteer can harvest much voluntary computing power. As pointed out by the climate*prediction*.net volunteer computing project team, getting the word out to the public is a challenge, and the best hope is to have an application of widespread public appeal combined with marketing and public relations to generate media interest (Christensen *et al.*, 2005). Even with these advantages, climate*prediction*.net suffered from a high attrition rate of volunteers. It is not practical for an individual user to use volunteer computing to increase the computer power available for his routine computer use.

1.2 OurGrid

OurGrid is a web-based community for sharing computing power. It has been in production since December 2004. It began as a research project at the Federal University of Campina Grande, and later gained some financial support from HP Brazil. The principal reason for HP's interest in the OurGrid project is that the project is carrying out cutting-edge innovation in a technology area of strategic interest to the company. However, it is also seen as part of HP's global citizenship efforts (Hewlett-Packard Development Company, 2006). In March 2006, OurGrid comprised 323 machines on 17 sites, up from around 150 machines in March 2005 and around 220 in August 2005 (Cirne *et al.*, 2006). For a snapshot of OurGrid's current size and status, visit http://status.ourgrid.org.

As in volunteer computing, members join OurGrid by downloading a lightweight client which runs tasks on their computer when it would otherwise be idle. However, the tasks do not all belong to a single application run for a project team; instead, the tasks may be part of an application submitted by any OurGrid member. That is, the organisational model for OurGrid is a community, in which all members can both contribute and benefit, as opposed to a service (where some only perform the service and others only receive it, as in volunteer computing) or a market (in which benefits are bought and sold, as in a grid economy).

The OurGrid code incorporates an innovative resource allocation algorithm which ensures that the more spare computing power that a member donates to other OurGrid members, the more he can expect to receive from them. Thus there is an incentive for an OurGrid member to let other members to use his computer, which is that in return he can use their computers (when they would otherwise be idle) to perform any embarrassingly-parallel application that he pleases, for free. OurGrid members do not choose which application to donate their spare computing power; the resource allocation depends only on the relative past generosity of the members who are currently requesting donations. As a result, in contrast to volunteer computing, members can obtain computing power from OurGrid in order to run applications that are not of wide public interest.

The client code runs the tasks in a sandbox based on Cambridge University's Xen virtual machine (Barham *et al.*, 2003), which isolates the tasks from the rest of the computer, and has no access to the network. This helps to prevent poorly coded or malicious tasks from causing security problems. It also carries out an integrity check after task execution, to prevent malicious tasks from installing Trojan horses. OurGrid is designed to work for up to 10 000 computers. In comparison to grids such as the LHC grid it is simple to join, and joining does not require a human negotiation. Its code is lightweight, and it does not require secure e-banking or a network-wide cryptographic infrastructure.

OurGrid's architecture for the sharing of computing power is that of a peer-to-peer community (Oram, 2001). That is, it is completely decentralised, and hence so has no single point of failure. There have been several other proposals for peer-to-peer sharing of computing power (Oliveira *et al.*, 2002; Butt *et al.*, 2003; Taylor *et al.*, 2003; Lo *et al.*, 2004), but unlike OurGrid these proposals do not address the problem of providing an incentive for peers to share. The experience of peer-to-peer file-sharing systems in which there is no incentive to share has been that in these systems a significant percentage – in some cases, a majority – of participating peers *free-ride*, that is they only consume resources, never contributing anything back to the system (Adar and Huberman, 2000; Saroiu *et al.*, 2002; Hughes *et al.*, 2005).

Evaluation results summarised in Figure 5 of Cirne *et al.* (2006) indicate that almost all the time there are computers being donated over OurGrid. This means that almost all the time there are both some OurGrid members whose computers would otherwise be idle, and some other OurGrid members who have a current demand for computing power that their local resources cannot satisfy.

In the rest of the paper, I will first outline the role of web-based tools and communication in OurGrid, and the effects of the build-up of social capital. I will suggest that similar communication might play a role in grid economies. I will then describe an application for which the community aspects of OurGrid are especially relevant. I will end by indicating how readers may benefit from, or contribute to, OurGrid themselves.

2 Web-based community support for grids

In this section, I will describe some ways in which OurGrid successfully exploits the fact that it is a web-based community. Other web-based community projects might find some of these strategies useful in their own contexts. I will then suggest, based on the experience of OurGrid and some research into unsuccessful and successful online commodity markets, that grid economies might also benefit from having web-based community infrastructure.

2.1 *OurGrid seen as a web-based community*

A crucial aspect of OurGrid's code is that it was developed by and for a web-based community. Web-based communities allow distributed collaboration, opportunities for testing ideas, and easy feedback from members on system functionality. All of these have been essential in the development of the code. OurGrid's code is open source, allowing any member (or indeed any non-member) to contribute to its development. Different functionalities were developed by different members of the community, in different locations, working together over the internet. For instance, OurGrid's resource allocation algorithm (Andrade *et al.*, 2004) was developed by a collaborative team distributed between Brazil and the UK. OurGrid's own technology has been used to run simulations and tests of new technical ideas for OurGrid itself, trying them out in a secure way on members' computers (Duarte, 2005). Several changes in technical direction for OurGrid came from suggestions by community members or arose from their experiences with using the OurGrid code (Cirne *et al.*, 2006, Section 7); this experience is normal in the development of open source projects.

There is an OurGrid users' mailing list, with an archive linked from the OurGrid website, which is used to report and discuss problems and features, and to share experiences. Members of the development team take the responsibility to answer questions on the users' mailing list, if they are not adequately answered by other users. There is a separate developers' mailing list, also web-archived, for people developing OurGrid or discussing its architecture; and an announcement list to report new code releases and bug-fixes. The use of separate mailing lists means that members can choose for themselves the level of information that they receive, and that meta-level discussions on how to improve OurGrid can be separated from questions on how to use the current version.

The OurGrid development team has found wikis particularly useful for coordinating distributed tasks. For example, there is a wiki of contact information for the development team. Whenever a team member makes contact with an organisation that expresses interest in OurGrid, the team member can add this contact to the wiki and provide updates on the progress of that relationship. That way, any other team member who comes into contact with the organisation can simply and quickly find out the state of the organisation's involvement. There is an introduction page used to introduce members of the development team; this is more important than it might first appear, because the team is large. Another wiki page records a summary of current research activities. This makes it easier for the members of a large and geographically distributed research team to be aware that another team member's work might be relevant to their own. Wiki technology is especially well suited for this purpose, since it makes it easy for each researcher to update the summary of her own research. There is also a users' wiki, which allows members to contribute useful information, hints, and pieces of code. For instance, the users' wiki contains a script to manage OurGrid user agents in the machines of a site containing many machines, and code and instructions to run OurGrid user agents as a Windows service. If a user writes a piece of code for her own convenience in using OurGrid, the ease of sharing code through the wiki means that with very little additional effort she can potentially spare many other users from repeating her work.

The OurGrid website[2] is also used to spotlight interesting applications and development activities, thus rewarding high-quality application or development work with publicity and diffusing information about this work to people who might benefit from knowing about it; and to make available all OurGrid's documentation, including the quick-start guide, manual, list of frequently asked questions, and related research papers. Moreover, web technology is essential for the ease of joining OurGrid, because the method of joining OurGrid is to download the OurGrid client from the web.

The OurGrid community has been fostered through face-to-face communication as well as web communication. Regular face-to-face meetings have been organised for OurGrid community members. This has been productive, both technically and socially. It appears to be a general rule that members of successful web-based communities wish to meet each other face-to-face, and that face-to-face meetings between some of the members, where practical, can enhance the web-based community, with online and offline communication complementing each other. Calvi (2006) calls the resulting online-and-offline communities 'semi-virtual'. For an example in a different context of the beneficial impact of face-to-face meetings between some members of a web-based community, see Burgos *et al.* (2006).

The social capital developed within OurGrid provides an extra disincentive for malicious uses, in addition to technological security protections such as the sandbox mentioned earlier. It provides an incentive for members to make useful contributions. Perhaps most importantly, it has had the effect of increasing and enhancing the interaction between OurGrid members. Several OurGrid member organisations are now collaborating with each other in deeper ways than simply using each others' spare computing power.

2.2 Grid economies may need web-based communities

The importance of person-to-person web-based communication in the development, testing, use and growth of OurGrid may be relevant to organisations developing other models of sharing computing power, in particular grid economies. The grid economy model is that of an online commodity trading market, where the commodity is computing power. Kollock and Braziel's (2006) study of failed online commodity trading markets shows that one of the reasons for failure was an incorrect assumption that identity and the network of social relationships were unimportant to the functioning of an online market. In Altra Online Exchange, a rare successful business-to-business online market that Kollock and Braziel examine, social relationships between traders were still important and not precluded by the online system. Developers of Grid technology have so far focussed on building correctly-functioning and reliable software to carry out automated trading of computing power between computers. This focus is understandable, because it is a complex and challenging technical problem. However, it may well turn out that grid economies also need a network of human social interactions to function well, just as it has turned out to be the case for OurGrid and for earlier online commodity markets.

This issue was foreseen to a certain extent by Foster (2003), who stated that the success of the Grid concept "demands effective community mechanisms" for coordinated research and development of core technologies, testing, packaging, documentation, and support and training of user communities. However, this list appears not to include the communication between users that Kollock and Braziel found important for online commodity trading, or the feedback from users to developers that was valuable for OurGrid.

Interestingly, social features appear in volunteer computing projects, although the depth of social interaction in these projects is not generally as strong as in OurGrid. The BOINC infrastructure supports volunteer profiles, message boards (which are used to sort out technical problems, but also sometimes for purely social conversations), and links to project-related news and information. It also produces statistics showing which teams of volunteers have donated the most computing power, in order to encourage donations through competition driven by team identification; this idea goes back to much earlier volunteer comuputing projects for factorising large numbers (Lenstra *et al.*, 1996). Since there is no economic motivation for volunteering, it seems reasonable to try to use social factors to motivate volunteers. In a poll of SETI@home (2002) volunteers, 58% said that their main reason for running SETI@home was for the good of humanity, 17% to "keep my computer productive", and slightly over 5% were mainly motivated by the chance of becoming famous or of getting their name on a top 100 list on the website. It may be that developers of grid technology have so far not put much consideration into social factors because they believe that these are unnecessary if there is a concrete incentive to share computing power. The experiences of OurGrid and of online commodity markets suggest that such a belief would be mistaken.

3 An OurGrid application: water management

OurGrid members have used the OurGrid technology to run many different applications, including molecular dynamics (Veronez *et al.*, 2003), simulations, computer imaging, and data mining (Silva *et al.*, 2004). An example of a project for which the community aspects of OurGrid are particularly relevant is the SegHidro collaborative project (Araújo, 2005; Araújo *et al.*, 2005). 'SegHidro' is an abbreviation of the Portuguese for 'hydrological security'. The project's aim is to support water management in North-East Brazil, an area where repeated droughts and poor planning have made a huge number of people suffer unnecessarily. There is a recurrent theme in the local music and literature of the misery and migration that occur in times of drought (Guerra, 2002; Swarnakar, 2003).

SegHidro is funded by the Brazilian Ministry of Science and Technology, via the funding body FINEP[3]. Within SegHidro, several government organisations are using the computation power that they can gain from OurGrid to improve weather forecasting for this area. Water reservoir managers are using these more accurate forecasts, plus computation and simulation power from OurGrid, to improve the risk management of their resources. Other organisations are using computing power from OurGrid, and information from the previous two types of SegHidro project members, for the management of alluvial aquifers – underground water sources that were previously used in an unsystematic fashion. Agricultural management centres are interested in SegHidro's forecasts and scenario simulation capabilities in order to perform better agricultural planning for sites in their geographical area of responsibility.

For this application, volunteer computing could certainly be used. However, the use of OurGrid means that the SegHidro project can exploit the communication channels and social capital already available within OurGrid. Clearly, the more closely that the different organisations work together, the more effective they can be, and the better the potential impact on water availability for inhabitants of this area. By sharing each others' spare computing power over OurGrid they are already cooperating with each other, and the aim is that their joint membership of OurGrid will assist them to collaborate in other ways as well, both over the web and face-to-face, sharing data and human expertise in addition to computing power.

The SegHidro project is a fine example of an application of web-based communities that is more interesting than the common use of such technology to support e-commerce. I would encourage researchers and developers of web-based community technology to look for humanitarian applications for their technological expertise, if they are not already doing so. It might not be obvious that cutting-edge research in networked computing could contribute to improving the lives of small-scale farmers in Brazil who may not own computers themselves, but that is precisely the potential result of SegHidro's use of OurGrid technology.

4 What next?

At the moment only embarrassingly-parallel applications can run on OurGrid. Restricting OurGrid's use to this type of application makes it easier to ensure that the system is efficient and secure. However, the OurGrid team is investigating how to extend OurGrid's code to also support workflow applications.

There are also plans to extend OurGrid's resource allocation algorithm to provide an incentive for the sharing of data between members in addition to the sharing of computing power. Several computations – including, for example, some of the computations using meteorological data that are mentioned in the previous section – require a relatively small amount of computing power but access to a very large amount of data. Currently, the only aspect of a computation carried out on behalf of another peer that is rewarded is the computing power used in carrying out the computation. The provision, storage and transfer of data relating to a computation are services that a donating peer may provide for another, and the resource allocation algorithm might be adapted to reward these services as well as the provision of computing power, repaying these data-related services either with similar data-related services or with additional computing power.

5 Conclusion

This paper has described some ways in which OurGrid exploits the fact that it is a web-based community. This fact has been crucial in its design, development, testing, adoption, growth, and in improving its usefulness to its members.

In contrast to other approaches to sharing computing power between different organisations, on the one hand OurGrid provides an incentive for members to share their computers, and can be used by individuals to increase the computing power available to them for running applications which are not of general public interest: and on the other hand it is lightweight, easy to join, and does not have a single point of failure.

Although OurGrid is designed to work smoothly with 10 000 computers, it currently has many fewer. Communities of 10 000 users require different social and technical infrastructures from communities of 300 users, so changes to the way the community is organised will be necessary as OurGrid grows. Another current limitation of OurGrid is that it can only be used for embarrassingly-parallel applications. However embarrassingly-parallel applications are useful in many fields, and there are plans to extend the range of applications further.

Acknowledgements

Thanks to all the OurGrid team. Special thanks to Walfredo Cirne, Eliane Araújo, and Carlos O. Galvão. This paper was written to the music of Luiz Gonzaga.

References

Adar, E. and Huberman, B. (2000) 'Free riding on Guntella', *First Monday*, October, Vol. 5, No. 10, http://www.firstmonday.dk/issues/issue5_10/adar/index.html.

AFRICA@home (2006) 'What is MalariaControl.net?', http://africa-home4.cern.ch/malariacontrol .html.

Anderson, D.P. and Fedak, G. (2005) 'The computational and storage potential of volunteer computing', http://boinc.berkeley.edu/boinc_papers/internet/paper.pdf.

Anderson, D.P., Cobb, J., Korpela, E., Lebofsky, M. and Werthimer, D. (2002) 'SETI@home: an experiment in public-resource computing', *Communications of the ACM*, Vol. 45, No. 11, pp.56–61.

Andrade, N., Brasileiro, F., Cirne, W. and Mowbray, M. (2004) 'Discouraging free riding in a peer-to-peer CPU-sharing grid', *Proceedings of 13th IEEE International Symposium on High Peformance Distributed Computing*, Honolulu, Hawaii, pp.129–137.

Araújo, E. (2005) 'SegHidro web site', http://seghidro.lsd.ufcg.edu.br.

Araújo, E., Cirne, W., Wagner, G., Oliveira, N., Souza, E.P., Galvão, C.O. and Martins, E.S. (2005) 'The SegHidro experience: using the grid to empower a hydro-meterological scientific network', *Proceedings of 1st IEEE Conference on e-Science and Grid Computing*, Melbourne, Australia, pp.64–71.

Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I. and Warfield, A. (2003) 'Xen and the art of virtualization', *Proceedings of ACM Symposium on Operating Systems Principles (SOSP) 2003*, Bolton Landing, NY, pp.164–177.

Beckman, P. (2005) 'Building the TeraGrid', *Philosophical Transactions of the Royal Society, Series A: Mathematical, Physical and Engineering Sciences*, Vol. 363, No. 1833, pp.1715–1728.

Berman, F., Hey, A.J.G. and Fox, G. (2003) *Grid Computing: Making the Global Infrastructure a Reality*, Hoboken, NJ: John Wiley & Sons.

Burgos, D., Hummel, H.G.K., Tattersall, C., Brouns, F., Kurvers, H. and Koper, E.J.R. (2006) 'Influence of face-to-face meetings on virtual community activity: the case of learning network for learning design', *Proceedings of the 2006 IADIS International Conference on Web Based Communities (WBC 2006)*, San Sebastian, Spain, pp.105–112.

Butt, A.R., Zhang, R. and Hu, Y.C. (2003) 'A self-organizing flock of condors', *Proceedings of ACM/IEEE Supercomputing 2003 (SC'03*), November, p.42.

Buyya, R. (2002) 'Economic-based distributed resource management and scheduling for grid computing', PhD thesis, Monash University, Melbourne, Australia.

Buyya, R., Abramson, D. and Giddy, J. (2000) 'An economy driven resource management architecture for global computational power grids', *Proceedings of the 2000 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000)*, Las Vegas: CSREA Press, pp.517–525.

Calvi, L. (2006) 'Social issues in on-line community content sharing', *Proceedings of the 2006 IADIS International Conference on Web Based Communities (WBC 2006)*, San Sebastian, Spain, pp.340–346.

CERN (2005) 'Worldwide LHC computing grid', http://lcg.web.cern.ch/lcg/.

Christensen, C., Aina, T. and Stainforth, D. (2005) 'The challenge of volunteer computing with lengthy climate modeling simulations', *Proceedings of the First International Conference on e-Science and Grid Computing (e-Science'05)*, Melbourne, Australia, 5–8 December, pp.8–15.

Cirne, W., Brasileiro, F., Andrade, N., Santos, R., Andrade, A., Novaes, R. and Mowbray, M. (2006) 'Labs of the world, unite!!!', Special Issue on Grid and P2P Computing, *Journal of Grid Computing*, http://dx.doi.org/10.1007/s10723-006-9040-x.

Duarte, A. (2005) 'About GridUnit', http://gridunit.sourceforge.net.

Foster, I. (2003) 'Common middleware and SciDAC collaborations', Talk at *PPDG Review*, Fermilab, 28 April, http://www.ppdg.net/mtgs/phone/030411/PPDG-Site-Visit-April-2003.ppt.

Foster, I., Kesselman, C. and Tuecke, S. (2001) 'The anatomy of the grid: enabling scalable virtual organizations', *International Journal of Supercomputer Applications*, Vol. 15, No. 3, pp.200–222.

Guerra, G.A.D. (2002) 'O Êxodo Rural no Cancioneiro Popolar. Triste Partida, de Patativa do Assaré', *Trilhas*, Publ. UNAMA, Universidade da Amazônia, June, Vol. 3, No. 1, pp.23–24, http://www.nead.unama.br/bibliotecavirtual/revista/trilhas/pdf/trilhas31a3.pdf.

Hewlett-Packard Development Company (2006) '2006 global citizenship report', http://www.hp.com/go/report.

Hey, T. and Trefethen, A.E. (2004) 'UK e-Science programme: next generation grid applications', *International Journal of High Performance Computing Applications*, Vol. 18, No. 3, pp.285–291.

Hughes, D., Coulson, G. and Walkerdine, J. (2005) 'Free riding on Gnutella revisited: the bell tolls?', *IEEE Distributed Systems Online*, June, Vol. 6, No. 6, pp.1541–4922.

Jin, H. (2004) 'ChinaGrid: making grid computing a reality', *Proceedings of ICADL'04*, Lecture Notes in Computer Science, Vol. 3334, pp.13–24.

Kollock, P. and Braziel, E.R. (2006) 'How not to build an online market: the sociology of market microstructure', S.R. Thye and E.J. Lawler (Eds.) *Advances in Group Processes: Social Psychology of the Workplace*, New York: Elsevier Science.

Lenstra, A.K., *et al*. (1996) 'RSA factoring-by-web: the world-wide status', USENIX Report, 11 April, http://www.npac.syr.edu/factoring/status.html.

Lo, V.M., *et al.* (2004) 'Cluster computing on the fly: P2P scheduling of idle cycles in the internet', *Proceedings of IPTPS 2004*, pp.227–236.

Oliveira, L., *et al.* (2002) 'P3 (Parallel Peer to Peer): an internet parallel programming environment', *Proceedings of Web Engineering and Peer-to-Peer Computing Workshop*, Lecture Notes in Computer Science 2376, Pisa, Italy, May, pp.274–288.

Oram, A. (2001) *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, Sebastopol: O'Reilly Media Inc.

Sarmenta, L.F.G. (2001) 'Volunteer computing', PhD thesis, Massachusetts Institute of Technology, Boston.

Saroiu, S., Gummadi, P.K. and Gribble, S.D. (2002) 'A measurement study of peer-to-peer file sharing systems', *Proceedings of Multimedia Computing and Networking (MMCN)*, San Jose, January.

SETI@home (2002) 'SETI@home poll results', http://seticlassic.ssl.berkeley.edu/polls.html.

Silva, F.A.B., Carvalho, S., Senger, H., Hruschka, E.R. and Farias, C.R.G. (2004) 'Running data mining applications on the grid: a bag-of-tasks approach', *Proceedings of 2004 International Conference on Computational Science and its Applications*, Lecture Notes in Computer Science, Berlin/Heidelberg: Springer-Verlag, Vol. 3044, pp.168–177.

Swarnakar, S. (2003) 'Drought, misery and migration: the fictional world of José Américo de Almeida's "A Bagaceira" and Jorge Amado's "Seara Vermelha" and "Gabriela, Cravo e Canela"', *Latin American Studies Association (LASA 2003*), Dallas, Texas, 27–29 March http://136.142.158.105/Lasa2003/SwarnakarSudha.pdf.

Taylor, I.J., Shields, M., Wang, I. and Rana, O. (2003) 'Triana applications within grid computing and peer to peer environments', *Journal of Grid Computing*, the Netherlands: Springer, June, Vol. 1, No. 2, pp.199–217.

Veronez, C., Osthoff, C. and Pascutti, P.G. (2003) 'HIV-I protease mutants molecular dynamics research on grid computing environment', *Proceedings of the Second Brazilian Workshop on Bioinformatics (WOB 2003)*, Rio de Janeiro, Brazil, pp.161–164.

Notes

1 http://www.ourgrid.org/twiki-public/pub/OG/OurPublications

2 http://www.ourgrid.org

3 http://www.finep.gov.br

Appendix How to join OurGrid

If you have a Linux computer connected to the internet, you too can join OurGrid. Download the client from http://www.ourgrid.org and follow the quick-start guide. The website also gives contact information if you would like assistance, or if you would like to contribute suggestions or help to the project. OurGrid is open source, and is free to join (as in free speech and also as in free beer).