Proceedings of the UK e-Science All Hands Meeting 2004, © EPSRC Sept 2004, ISBN 1-904425-21-6

Collaborative tools in support of the *e*Minerals Virtual Organisation

Martin T Dove[1,2], Mark Calleja[1], Richard Bruin[1], Jon Wakelin[1], Martin Keegan[2], Stuart Ballard[2], Gareth Lewis[3], S Mehmood Hasan[3], Vassil Alexandrov[3], Rik Tyer[4], Ilian Todorov[4], Paul Wilson[5], Maria Alfredsson[5], G David Price[5], Clovis Chapman[6], Wolfgang Emmerich[6], Stephen Wells[7], Arnaud Marmier[8], Steve Parker[8], Zhimei Du[9]

1. Department of Earth Sciences, University of Cambridge, Downing Street, Cambridge CB2 3EQ
2. National Institute for Environmental *e*Science, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0EW
3. Department of Computer Science, University of Reading, Whiteknights, PO Box 225, Reading RG6 6AY
4. Daresbury Laboratory, Daresbury, Warrington, Cheshire WA4 4AD
5. Department of Earth Sciences, University College London, Gower Street, London WC1E 6BT
6. Department of Computer Science, University College London, Gower Street, London WC1E 6BT
7. Royal Institution, Albermale Street, London W1S 4BS
8. Department of Chemistry, University of Bath, Bath, BA2 7AY
9. School of Crystallography, Birkbeck College, Malet Street, London WC1E 7HX

Abstract

We outline the experience, from within, of working towards establishing the *e*Minerals project team as a functioning virtual organisation. We mostly describe our experience with the tools we have been using, with particular emphasis placed on the project-wide use of the personal version of the Access Grid.

1. Introduction

The *e*Minerals project is one of the NERC eScience testbed projects ("Environment from the Molecular Level"). The aim of the project is to develop a cross-institute collaborative infrastructure to facilitate a programme of computer simulations of environmental processes at a molecular level (Dove *et al*, 2003). The science applications concern issues such as nuclear waste encapsulation, adsorption of organic pollutants on soil particles, and weathering and precipitation of minerals. The simulations we run are particularly challenging when we seek to incorporate higher levels of realism into them, including complexity of the system, high-accuracy methods to properly model chemical bonds, and large samples in order to study processes that occur over length scales that are larger than the molecular length scale. The science outcomes so far from the *e*Minerals project are described elsewhere (Alfredsson *et al*, 2004). The grid infrastructure that is being established for this work involves integrating both computational and data grids, and various components of this work are discussed in a number of other papers in the 2004 All Hands collection (Calleja *et al*, 2004a; Tyer *et al*, 2004; Blanshard *et al*, 2004) and outlined in the Appendix. The third component within the *e*Minerals project is the support for the human interactions that are an essential part of the collaborations within the project; this is the subject of this paper.

The *e*Minerals project consists of teams from Cambridge, Bath and Reading Universities, Birkbeck College, University College London, the Royal Institution and the CCLRC Daresbury Laboratory. The work is managed under three teams, who focus on the grid developments (the 'grid team'), developments of the simulation codes and methodology (the 'application team'), and the science being carried out (the 'science team'). Members of the project work within any number of these teams.

As we will describe below, close collaboration between the members of the *e*Minerals project, both within and between the three teams, is essential to the success of the project. Accordingly we are our own experiment in how to manage a project in a way that maximises the amount of interaction that can be carried out between geographically distributed partners. This brings us to the notion of the "Virtual Organisation" (VO), a concept that is increasingly being cited within the eScience/grid community and which is certainly taken seriously by the international particle physics community. In this paper we will describe how the *e*Minerals project maps onto the concept of the VO, and we report our experience of using IT tools to support collaboration across the *e*Minerals VO. Since this is a report from within the project team, it is not possible to be objectively analytical about the human dimension, but we can report the extent to which various tools have helped foster collaboration.

2. The concept of the 'Virtual Organisation'

The concept of the Virtual Organisation arose around 10–15 years ago, but many of the practices of virtual organisations can be traced back at least four decades. For example, Sor (1999) has described how many of the features of virtual organisations can be discerned within the organisation of the housing construction industry in Western Australia in the early 1960's. Much has been written about the more specific concept of the 'Virtual Business', and many definitions of VO's are particularly pertinent to the industrial sector. The strong motivation for the formation of VO's in industry is the need to reduce costs – this was the driving factor that saw the drive towards collaborations in Western Australia. The economic idea is that costs can be saved if partners with complementary expertise work together towards some common objective. In particular, the infrastructure costs of a group of small units are likely to be much lower in total than the corresponding cost to a large organisation.

As a buzz phrase, there appears to be no agreed definition of the term 'Virtual Organisation', but there are a number of key characteristics that can be said to be implicit in the idea. Our working understanding of the concept of the VO is that it is a collection of people working together within an organisational structure that is distinct from their formal allegiances (probably not relevant in an academic context, but some members of a VO may be freelance without any institutional allegiance). A VO will have a particular mission, and may be time limited. Its members will inevitably be geographically dispersed, and will have responsibilities on behalf of their employer institute as well as on behalf of the VO. Membership is also likely to be dynamic, with members joining and leaving when their roles begin and end, rather than remaining members for the while duration of the project around which the VO is established. Some aspects of the VO are similar to those of a more traditional working organisation, but there are other aspects, such as a flatter hierarchy and a voluntary commitment, that are more peculiar to the VO. It is being recognised that the dependence on IT tools is one of the main characteristics of a VO, but it has also been argued that since VO's existed in practice before the IT revolution, the reliance on IT is not a defining characteristic.

It is useful to distil out of the general idea those concepts that are pertinent to an academic project, taking account of the fact that we do not see the same constraints as would be central to the concept of a Virtual Business. In particular, there is not the equivalent of a fixed objective – research objectives have to be sufficiently flexible to be able to develop during the lifetime of a project in response to new discoveries by the project team and competing research groups, and to adapt to the common situation where proposed research meets unpredictable problems. Moreover, there is no corresponding cost-reduction motivation, and usually it is expected that member groups will stay together for the whole lifetime of a particular project (rather than joining for short periods). The key point we can distil is that of the joining together of dispersed research groups to work together on a topic of common interest with a commonly-agreed management structure (*e.g.* one nominated leader and a steering committee) that is selected by acclamation rather than imposition. One might ask how this differs from a standard collaboration of distributed groups? We argue here that there are two features of a VO that are qualitatively different from features of a looser collaboration. The first is that there is a sharing of resources that is more akin to the manner in which resources are shared within a formal organisation. For example, there may be semi-formal policies on access to some of the shared resources, and a commitment on the part of the donor to ensure that access is properly maintained. The second particular feature of a VO is an interdependence between member groups that is built into the VO from the outset. It is possible for members of a collaboration to gain benefit from other members but to not be dependent on each other. One of the points of this paper is to show how IT can be used to ensure that the interdependence between VO partners is exploited to the fullest extent possible.

3. Mapping the *e*Minerals project onto the concept of the Virtual Organisation

Consortium grants are increasingly common in the modern funding era, and eScience projects are very much in this mould. Often such consortia are formed by groups with similar skills and related interests. In such a case, there may be no in-built interdependence on the consortia members. Thus there will be a tendency towards working within the traditional model of collaboration that is built upon regular but not frequent face-to-face meetings where progress is reviewed (a good consortium will gain a lot from these meetings), irregular email contacts where help/advice is sought (and the telephone used when this help is urgent), and the reading of manuscripts sent between partners. This is often as much as the partners expect out of the collaboration.

The UK eScience testbed projects have the interdependence between partners built into them from the outset. The *e*Minerals project, along with many others, consists of scientists, code developers and computer scientists (and some people who straddle two or three of these specialisations). Collaboration is essential if the project is to achieve its goal of constructing an integrated grid structure that meets the needs of the scientists. The scientists need to inform the grid team of their needs, and the grid team need to develop something that the scientists will genuinely find useful. The scientists will also need a lot of help adapting their usual work practices to the new grid-based way of working. The code developers will need to select their priorities based on the needs of the users through working closely with the scientists (as distinct from the more usual case where groups who use a particular code formulate periodic wish lists). The code developers can also use the grid structure to their benefit, and will need to interact with the grid team to ensure they get as much from the grid structure as the science team does. Thus we have sought to build collaboration between all project partners right from the outset. We need to work within the constraint that our teams are based in geographically distinct locations, and yet we want interactions to be much more frequent than would be possible if restricted to face-to-face meetings. It is our aim to eventually be able to enlarge the team, including bringing in international collaborators, and the concept of the VO will need to be robust enough to accommodate enlargement.

It should be noted that the science community from which the *e*Minerals project is drawn is not used to working within large close collaborations. It is much more characterised by individuals working with their own resources and codes. Consortia may be formed in order to gain access to high-capacity computing facilities, and partnerships may be formed between groups working on common problems. But these are a long way from the concept of the VO outlined in the previous section. We recognise that other areas of science, particularly the particle physics and astronomy communities, have a much stronger track record of the need to work within interdependent collaborations. It is thus an interesting experience to see how the *e*Minerals project will adapt to the new possibilities afforded by eScience to develop as a functioning VO.

Behind the development of the *e*Minerals VO is the deployment of the *e*Minerals minigrid, as described in Calleja *et al* (2004a), Tyer *et al* (2004) and Blanshard *et al* (2004). This is an integration of both compute and data resources, wrapped within Globus security (based on GT 2.4, but also implemented within GT 3.2). The minigrid consists of project-owned and contributed compute resources; the former are a group of three 16-node Linux clusters running PBS, and the latter are condor pools (including a very large pool of 930 teaching PC's at UCL running Microsoft Windows, and a smaller 24 machine mixed-platform Condor pool in Cambridge) and parallel computing facilities (including a 24-node IBM pSeries machine at Reading). We will shortly add a second 40-node linux cluster and extend the Cambridge Condor pool, and anticipate adding access to a Cambridge-wide Condor grid as described by Calleja *et al* (2004b).

The data resources are based around the twin pillars of the Storage Resource Broker (SRB) from the San Diego Supercomputing Centre and the Daresbury data portal. The use of the SRB overcomes one of the limitations of Globus, namely difficulties in copying files through the Globus gatekeeper. The SRB allows data to be handled through a server, and allows location-independent access to the data for the users. The data portal handles the metadata associated with the various studies carried out within the project.

Access to the *e*Minerals minigrid is via eScience X.509 digital certificates. We have also established an *e*Minerals certificate authority to grant access to members of the VO who are unable to obtain eScience certificates (usually these members are international collaborators). The *e*Minerals compute resources for production tasks can only be accessed through remote Globus run commands, apart from one cluster for which we allow access via gsissh in order to allow code compilation and testing. We are developing command-line scripts and a compute portal to make access for the users as straightforward as possible.

Developing the *e*Minerals minigrid has required close collaboration between members of the grid team, and only through equally close collaboration has it been possible to help the scientists make use of the resources of the minigrid. This collaboration has made extensive use of the tools discussed in the following two sections. The development of the *e*Minerals minigrid illustrates both the resource sharing and team interdependence that, as we have argued, differentiate the VO from a loose collaboration.

We make one final remark in this section. The success of business VO's depends on the use of standards, so that all partners work to the same system and interoperability is built into the VO from the outset. In our case, one standard that we adopt is the use of the Chemical Markup Language (CML) to describe the simulation data. CML is an application of XML that is designed to handle the science that drives our project, and it enables data to be imported and exported into and out of many of the codes used in the *e*Minerals project. Aspects of CML are described in two other papers in this All Hands collection (Garcia *et al*, 2004; Zhang *et al*, 2004). Similarly, for interoperability at the grid level, we have made use of the benefits of the use of grid standards, as will be continued with the Web Service Resource Framework (WSRF). The *e*Minerals project has already started advancing in the WSRF direction, with a number of its members being responsible for the development of a WSRF-enabled version of Condor in collaboration with the University of Wisconsin Condor team, using the *e*Minerals minigrid as a testbed for its deployment.

4. IT/Grid tools

Based on the above discussion, it is obvious that the *e*Minerals project has had to pay some attention to providing the IT infrastructure in support of the operation of the VO. There are a number of tools that can be used, some of which are emerging as grid tools within the current eScience initiative. In this context we note that there is a long history of the development of IT tools for collaborative work (such as email). The formal beginning of such systems and their study began in mid-1980's with the emergence of Computer Supported Cooperative Work (CSCW; Grundin, 1994). Since then, significant advances have been made in understanding the ways in which remote collaboration is carried out and the issues relating to dynamics of cooperative work. CSCW is technology 'independent', which means technology is not the major driving force behind the discipline. CSCW is socially dependant and looks at the way people interact and collaborate with each other, and attempts to develop guidelines for developing technology to assist in the communication process. The goal of CSCW is to discover ways of using computer technology to further enhance the group work process through support in the time and place dimensions. Recently, attention has been focused on the manifestation of CSCW systems. This has been realised in form of various computer applications and tools which support and augment collaborative work. Some of these tools are more appropriate to short-term project work with hard deadlines and inflexible deliverables, such as shared calendars, workflow tools and project management tools, and we have not taken them on board in the *e*Minerals project.

4.1 Interactive collaboration tools

We now briefly review our experience of a range of tools that have the potential to enhance collaboration across the *e*Minerals VO. Each tool has its own advantages and disadvantages, which we consider under three broad category: the 'access cost' to the communication, that is whether it is easy or difficult to initiate a new communication; the 'potential for instantaneous interaction'; and the 'setup cost', whether financial or in terms of necessary initial effort or expertise. The overall spectrum of communications tools consists of the following:

eMail: Although relatively new, email is now pervasive in the scientific community, and can now be said to be the traditional IT method for communication. It is cheap (zero setup cost), easy, and automatically available for everyone (low access cost). It is virtually platform independent. The difficulties we face in the use of email to support collaboration within the *e*Minerals VO are that it does not support instant communication (the speed of response will depend on people's email receipt setup, and email does not demand attention), and that with the welter of email communication (genuine and unsolicited) it is gradually losing its effectiveness as a means of communication (in general, people are no longer able to read and respond to all emails they receive, particularly emails sent to a wide circulation list).

Instant messaging (IM): Instant messaging gives much better instantaneous chat facilities than email, and is being used within the *e*Minerals project for discussions between small groups of people. It has a very low immediate access cost and nearly zero setup cost (simple installation and initial registration), and is available for all platforms in the project. It is less useful for larger numbers (> 3) of participants for two reasons: difficulty in maintaining social control (who should speak next, particularly since it takes time to type in a message), and because communication easily becomes tangled when participants follow several lines of thread in an IM discussion. Some organisations formally block the use of IM tools.

Web tools: At the lowest and slowest level, information can be transmitted to members of the *e*Minerals project via the *e*Minerals web site (www.eminerals.org). The value of a project website is that it is owned by the project members and can be shaped to meet the needs of the project (whether for dissemination or access to information). However, not all team members have access to the directory structure that supports the project web site, and they have no option but to rely on those who do have access to deposit information.

The development of the 'wiki' concept (from the Hawaii 'wiki wiki' meaning 'superfast') removes the problem of members of the VO not having direct write access to the project website. A wiki is a web site that can be freely edited by the community (document pages altered, new documents added, links created) using a simple markup language. Changes are implemented instantaneously (without review). We have set up a private wiki site for exchange of information within the *e*Minerals project, to which access is limited to members of the *e*Minerals VO through the use of X.509 certificates. People can deposit news information, edit documents (this paper was written on the wiki, and only transformed into a document format at the final stage), deposit information (for example, in order to collate information for a project review), and post questions and answers. Unlike normal web sites, the wiki has a relatively low access cost for members of the VO to post information to. As with all web-based tools, the wiki does not support instant communication. Its main problem is that it relies upon other team members regularly looking at it.

Access Grid based on Access Grid suites: The widespread use of the Access Grid is one of the successes of the UK eScience programme, and many projects (including ours) have made considerable use of the Access Grid for project meetings. The Access Grid network suites based in the UK eScience Centres provide an excellent support for remote meetings (there are around 30 suites in UK academic institutes). Videoconferencing via the Access Grid provides the best method for instantaneous communication. However, there are a number of difficulties associated with the use of Access Grid suites, including having all sites required for a particular team to be simultaneously available at suitable times, the fact that some team members have to travel considerable distances to get to their nearest Access Grid suite, and the need to negotiate support for a trained operator at each site. This corresponds to a very high access cost. The setup cost is prohibitive for most projects.

Personal Access Grid: The Access Grid software is packaged for use on a single PC as well as for use in a dedicated suite. Once running, it is easy for several team members to initiate an Access Grid session from their desktops, gaining many of the advantages of the Access Grid without the various disadvantages outlined above. Initial setup costs may not be not light (see below), but access costs are not high. We will discuss implementation and usage issues in the next section of this paper.

4.2 Other support tools

In addition to the communication tools reviewed above, we are using two other tools to support the *e*Minerals VO:

Tools for shared applications: In addition to verbal discussion on the Access Grid, it is sometimes useful for participants in a discussion be able to look together at representations of data, such as viewing the atomic arrangements in a molecular simulation. The VNC (Virtual Network Computing; Richardson *et al,* 1998) tool provides some level of support for this in that is allows users to share a single desktop and it is traditionally used within the Access Grid to share applications. A crystal structure displayed on one desktop can be seen by all participants in a discussion, and all participants are able to manipulate that structure. However, it does have several features that make it less suitable for this role. These include the following: *a*) the user must share the entire desktop rather than a single application, which may not be desirable; *b*) one of the participants runs the VNC server to which all other participants connect, and to

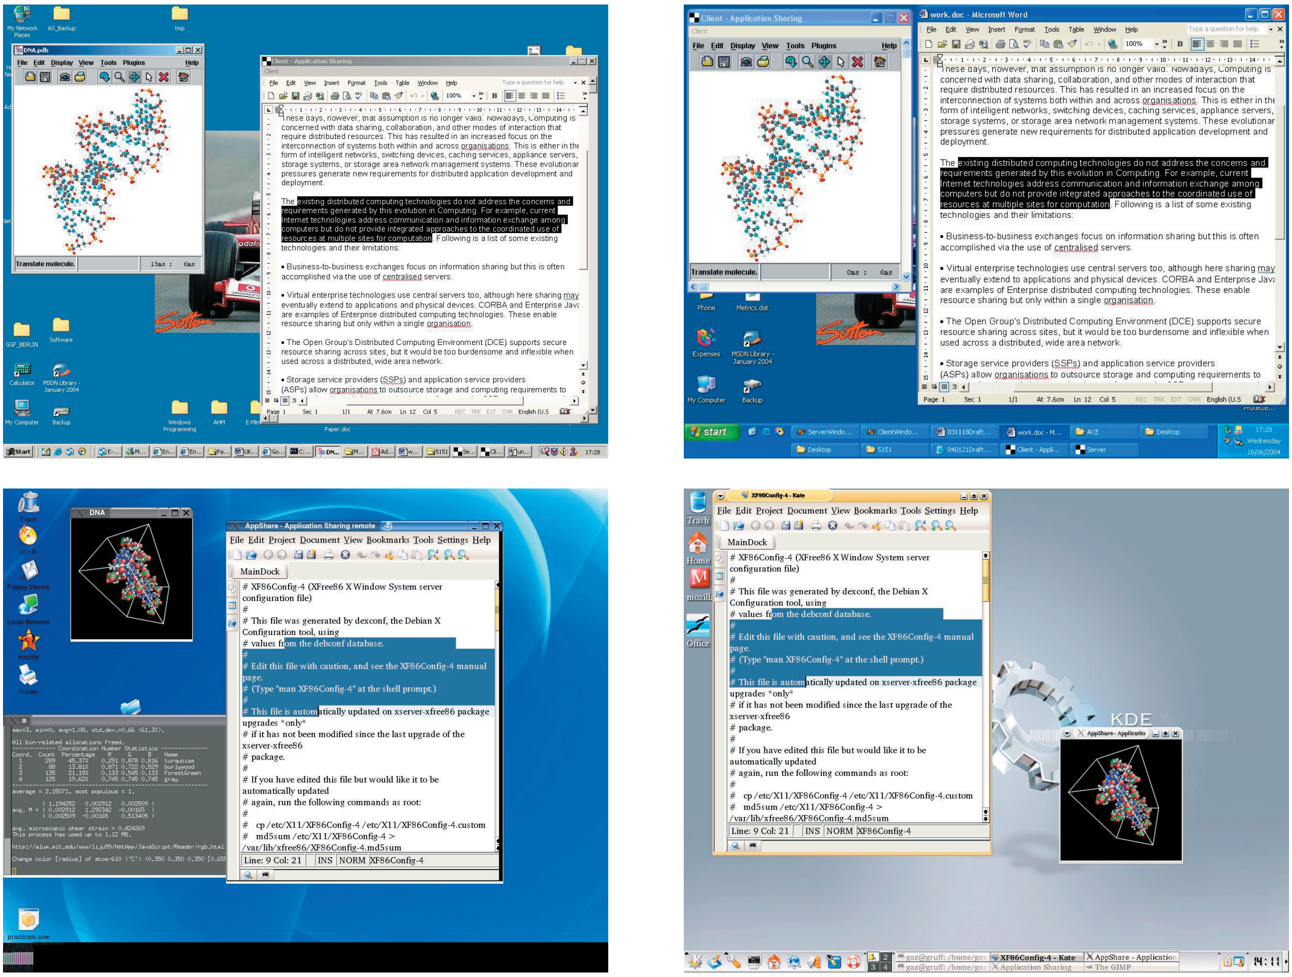

Figure 1. Examples of the application sharing tool being developed within the eMinerals project. The top images show application sharing with Microsoft Windows, and the bottom screens show application showing with Linux. In both cases, each window shows both server and client windows.

achieve this each participant must know the IP address of the machine running the server and a password set on the server; *c*) VNC has a high latency; *d*) VNC uses unicast for data transmission rather than multicast. To pick up the last point, the main problem with using unicast in a group environment is that the server machine needs to create and send multiple copies of the same data. This increases the load on the server as well as flooding the network with duplicate packets, leading to scalability issues. There is a plugin to VNC, which acts as a Java proxy to use multicast to direct the unicast traffic to other members of the group, but this is not part of the core VNC product. There are several problems associated with this approach; only the application owner has control over the application leaving other participants to be mere viewers, it leads to performance degradation and it is complicated to setup and run. Using the multicast proxy is not ideal and is comparable to relying on a multicast-unicast bridge to use the Access Grid (see below).

We are developing a new application sharing tool, eventually to be based on multicast communication, that will be faster through specific optimisation strategies (including not needing to share the whole desktop), and which will support multiple platforms. Currently, we have two separate versions of the tool: one for Linux and other for Microsoft Windows. We are working to integrate these two components into a single tool. The Graphical User Interface is being designed to be similar to the video tool used in the Access Grid: different application streams will be displayed as small windows that can be selected to be displayed as a larger full size version. This will allow many participants to share applications and viewers to have a convenient (and familiar) view of all the available applications. Examples screen shots of this tool are shown in Figure 1.

Helpdesk software: Once the *e*Minerals minigrid was operational, the science users were then helped by the grid team to port their applications across to the minigrid and start running in earnest. Recall that access to the resources of the *e*Minerals minigrid is only via Globus tools, and users needed support to switch from the more traditional login methods. At an early stage it was realised that it was essential to implement a structured support system rather than have users guess who to ask for help. As the number and complexity of the problems encountered by the users increased, and considering the varying areas of expertise of the members, it became necessary to find means to allow groups of users to cooperate towards solving problems. It is in order to achieve this that we adopted the Open Ticket Resource System (OTRS), an open-source email support system (Edenhofer et al, 2003; see http://otrs.org/). Though email is traditionally a one-to-one form of communication, the OTRS system provides means for multiple users to manage help requests sent to a common address (*i.e.* helpdesk@eminerals.org). Incoming requests can be matched to the users most apt to respond, and several users can view, respond and add comments to existing help requests via an online Web interface. The OTRS is a relatively simple tool to use, that can be tailored to best match the capabilities and structure of the group. An example of this is the use of request queues and notifications. Several queues, which contain waiting help requests, have been set up to correspond to various areas of expertise (e.g. data transfer issues, grid tools, etc.). For each of these, a set of users most capable of responding was defined. Upon receipt, incoming emails are sorted into the right queues, and the corresponding helpdesk users are instantly notified of the help request by email. They can then log into the system to view, respond, comment, or move the request to a different queue.

Electronic team newsletter: Newsletters are a more traditional method to facilitate the flow of information within a collaboration, but nevertheless we have found that the use of a regular project newsletter has played a significant role in supporting the development of the *e*Minerals VO. Modern desktop publishing packages make it easy to put together a high-quality document, and the web and email allow for easy dissemination. The key to the successful use of the newsletter is the editor, and his/her ability to extract contributions from the project members against a set of fixed regular deadlines. Past copies of the *e*Minerals newsletter are available from http://www.eminerals.org/.

5. Setting up and using the Personal Access Grid

The tool that we have felt to be most useful in developing the *e*Minerals VO out of the initial collaboration has been the personal version of the Access Grid, and in this section we describe our collective "user's experiences' in terms of the problems we faced in rolling this out across the project, and in the way we have learned to use it. We note from the outset that in the general case the access cost (as defined above) is low, and the potential for instantaneous communication is high, but that there may be high setup costs in terms of staff effort (the financial costs are minimal; all that is needed is a webcam and microphone headset per participant). It should be noted, however, that the setup cost is only needed once per institute, and thereafter the access is easily expanded.

5.1 Setting up the Personal Access Grid

In this section we discuss the issues we faced in setting up the use of the Personal Access Grid across the *e*Minerals project. Specific points included:

Firewalls: Firewalls need to be configured, and you therefore need a sympathetic computer support person. Ideally the firewall should be set up to allow UDP multicast traffic, or five specified TCP ports if it proves necessary to use a multicast bridge arrangement (see below).

Multicast traffic: Some work is likely to be needed to allow multicast traffic through institute networks. Although multicast is not a new technology, until recently (particularly until the development of the use of the Access Grid) it has not been widely needed. As a result, it is often the case the network components (switches, routers *etc*.) have multicast forwarding disabled, and systems administrators will need to act on your behalf to ensure that multicast traffic can successfully be routed between users's computers and the outside world. Across the *e*Minerals team we have noted that different institutes have different experience of multicast technologies, and some institutes have been able to enable the use of multicast somewhat faster than other institutes.

Troubleshooting multicast: The Access Grid community has provided a software tool called the Multicast Beacon. This involves both server and client tools. The clients broadcast multicast signals, and these are routed through the server to other clients. Thus each client can monitor whether it can send and receive multicast traffic. The UK JANET service runs a beacon server on behalf of the UK Access Grid community. The beacon tool is an excellent diagnostic tool for troubleshooting multicast problems across a network, although in practice the beacon will demonstrate the existence of a multicast routing problem but will not necessarily demonstrate the complete lack of any problem!

Multicast-unicast bridging: In practice, it will not be possible to set up all collaborating sites to use multicast quickly, and in some cases possibly not in the lifetime of a project. For example, within the *e*Minerals project we found that there are severe bandwidth/networking limitations for one of our sites, which only accesses JANET via a slow link to a neighbouring university. In our case, we have needed to set up a multicast-unicast bridge, using the tool developed by Mike Daw in Manchester. Although the Access Grid Support Centre provides bridge facilities, a collaboration with longer term requirement will need to set up their own bridge. We used a low-cost PC with linux installed just outside the firewall in the lead author's department. This was relatively easy to set up and has worked well for the project. When people use unicast methods, they simply run the video (vic) and audio (rat) tools from command line rather than via the web interface to the Access Grid.

Multicast reliability: Multicast is not proving to be fully reliable. Our experience is that service can be seen to go on and off, as monitored by the beacon tool.

Operating systems: We mostly use the version of the Access Grid for Microsoft Windows. Our experience with the Linux version is mixed. There are two main

Figure 2. Linux setup for the personal Access Grid, showing how, if desired, it can be run with more than one monitor in order to give more space for the participant windows. The photographer's window can be identified by the flash!

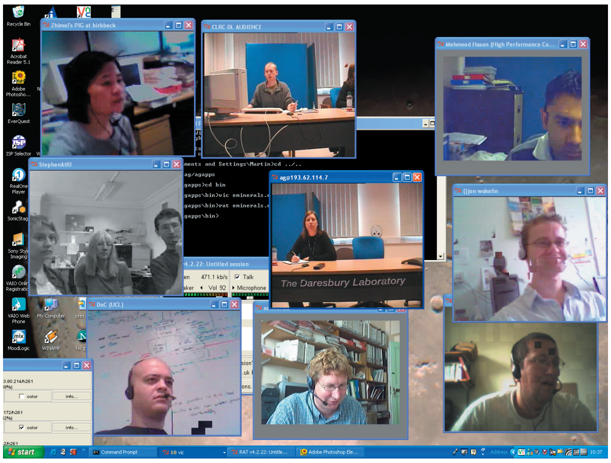

Figure 3. Screen shot from one of the eMinerals coffee morning sessions on the Access Grid, showing 9 simultaneous video streams. This screen shot was obtained on a laptop with 1400 x 1050 pixel resolution.

areas that cause problems. The first is in having drivers for USB webcams – there are drivers available for only a few USB webcams. The second is drivers for sound. This has been more of a problem. With Suse Linux we found we had no problems, but with other breeds of Linux sound proved to be a difficult problem. Some of these problems can be traced to conflicts between Open Sound System (OSS) drivers and the Advanced Linux Sound Architecture (ALSA) modules. The latter are required for use with the Access Grid, and methods to avoid problems are described elsewhere (Lewis, 2004).

At the moment there appears to be no options for Apple's Mac OS X. The vic program appears to work under OS X (from private correspondence), but there is not a version of rat for this operating system.

Screen area: One of the main advantages of running the full Access Grid within a purpose-designed suite is that 2 or 3 video projectors can tile a large number of participant windows across a wide area of a blank wall. Thus a major criticism of the use of the personal version of the Access Grid is that a wide screen area must be compressed within a single desktop screen. In practice, we have not found this to be a large problem. On the one hand, with additional video cards in the PC it is possible to use more than one monitor in order to increase screen area, as shown in Figure 2. On the other and, modern screens can have very high pixel resolution. Figure 3 shows a screen shot on a laptop with pixel resolution of 1400 × 1050. This screen shot shows that a single desktop can be capable of sustaining of order of 10 simultaneous participants.

Software version: At the time of writing, we are currently using version 1.2 of the Access Grid software. We plan to move over the version 2 when there is a wider community migration. It should be noted that version 1.2 is no longer easily available.

Virtual venues: We chose to set up our own *e*Minerals virtual venue for our collaboration. This needed to be done for us by the Access Grid support team. Without our own virtual venue we would have needed to work within someone else's virtual venue, the main problem being that we would not have unlimited on-demand access to the use of the Access Grid.

5.2 User experience of the personal Access Grid

The Access Grid is designed primarily to support large-scale collaborative interactions, rather than the more traditional videoconferencing between 2 individuals. We are using the Access Grid at an intermediate level: clearly at each end there are individuals, but we are making significant use of the ability of the Access Grid to support many simultaneous users. For project meetings, we may well often have around 10 participants, and we often mix in the Daresbury Access Grid node. As noted above, we do not find that the limited screen area (compared to the large display area of an Access Grid suite) is an issue for this scale of interaction.

We use the Access Grid to support our collaborations in a number of ways:

Project 'coffee mornings': We meet once a week for an informal meeting at which we exchange news and other information, and at which people can pose questions for discussion (Figure 3). We have specifically aimed to give this the feel of an informal coffee morning were people can pop in and out. We have found that this works well, with team members engaging naturally with each other.

Team project support: Our grid team has a number of large and small projects on the go (*e.g*. Tyer *et al*, 2004). The team has found it to be very advantageous to be able to meet weekly to monitor and develop progress.

Support work and ad hoc discussions: it has been very useful to use the personal Access Grid in "on demand" mode to facilitate discussions on specific issues as they arise, such as troubleshooting other work in the project, helping people install software, planning tests of tools *etc*. Although this could be done with other videoconferencing tools (e.g. tools that use

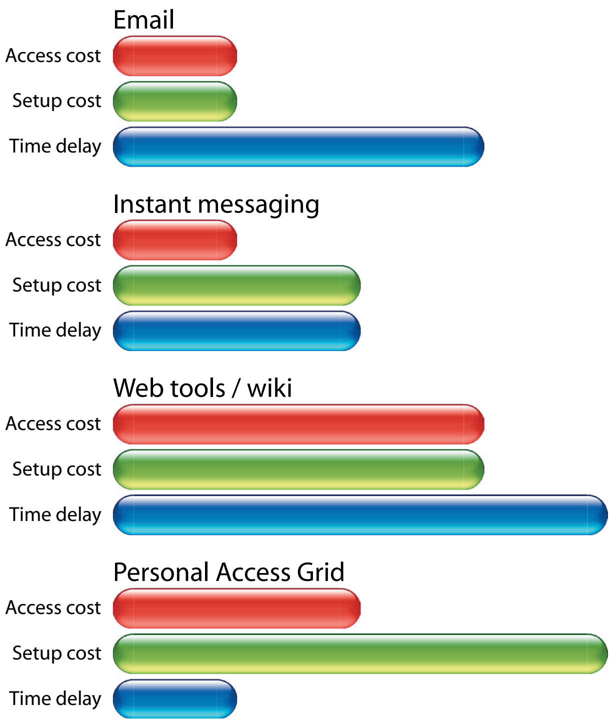

Figure 4. Rough comparison of the four main communication tools discussed in this paper. Shorter bars are better. Our argument is that the higher setup costs of the personal version of the Access Grid are well matched by the lower time delays for real communications.

H.323 protocols such as Microsoft Netmeeting), the personal Access Grid does just as well. The point here being that the ability to run ad hoc videoconferences is proving to be very useful in functioning as a VO.

Project management: When it comes to writing reports and new proposals, the Access Grid facilitates the interaction between team members that is essential. For gathering ideas, examples, summaries of work carried out, and lists of publications, the wiki tool described earlier is ideal to assist in the collation of information. However, in then working through the information, we have found that the use of the Access Grid has been invaluable.

Conclusion and summary

We have reported experience gained within the *e*Minerals project of using a range of IT/grid tools for the development of the infrastructure to support a virtual organisation. We have particularly emphasised the use of the personal Access Grid, providing some practical pointers to rolling this out across the VO and giving examples of how this has been a useful tool.

Some of our points are summarised in Figure 4, which compares the various factors we have taken into account in our discussion of the different tools available to support collaboration within a VO. Clearly there is a compromise. Some tools (e.g. email and instant messaging) are very easy to set up and use, but suffer from particularly difficulties as instant communication tools. Email as a tool is now overwhelmed by misuse and overuse, and IM supports only a small number of participants (the same criticism that can be levelled against the telephone). On the other hand, the personal version of the Access Grid has a high initial set up cost, but once setup it is easy to use and provides communication facilities that are not given by other tools. The setup costs should not be underestimated, but fortunately once it has been set up for one person at one site other users will immediately benefit.

Our experience with the personal version of the Access Grid as a collaborative tool has been very positive. We have not found that restricting the Access Grid to a computer desktop creates too many disadvantages compared to the full Access Grid, and has a number of important advantages over the full Access Grid in terms of easy and instantaneous access for all members of the VO. The low financial cost also ensures that roll-out across the whole consortium is affordable. We therefore are able to recommend the use of the personal Access Grid to other collaborations.

Acknowledgements

We are grateful to the Natural Environment Research Council for financial support and to the computer support personnel at our various institutes.

References

Alfredsson M *et al* (2004), "Science outcomes from the use of Grid tools in the *e*Minerals project", Proceedings of UK e-Science All Hands Meeting 2004

Blanshard L, Tyer R, Calleja M, Kleese K & Dove MT (2004) "Environmental Molecular Processes: Management of Simulation Data and Annotation", Proceedings of UK e-Science All Hands Meeting 2004

Calleja M *et al* (2004a) "Grid tool integration within the *e*Minerals project", Proceedings of UK e-Science All Hands Meeting 2004

Calleja M *et al* (2004b) "CamGrid: Experiences in constructing a university-wide, Condor-based grid at the University of Cambridge", Proceedings of UK e-Science All Hands Meeting 2004

Dove MT *et al* (2003) "Environment from the molecular level: an escience testbed project" Proceedings of UK e-Science All Hands Meeting 2003, pp 302–305

Edenhofer M, Wintermeyer S, Wormser S & Kehl R (2003) "OTRS 1.2 - Manual", http://doc.otrs.org/1.2/en/html/

Garcia A, Murray-Rust P & Wakelin J (2004) "The Use of XML and CML in Theoretical Chemistry and Physics Programs" Proceedings of UK e-Science All Hands Meeting 2004

Grudin J (1994) "CSCW: History and Focus". *IEEE Computer*, **27** (5), 19–27

Lewis G (2004) "Robust Audio Tool: Sound with ALSA howto", http://eminerals.esc.cam.ac.uk/eminerals/howto/accessgrid/piglet_rat.pdf

Richardson T, Stafford-Fraser Q, Wood KR & Hopper A (1998) "Virtual Network Computing", *IEEE Internet Computing* **2** (1), 33–38

Sor R (1999) "Virtual Organisations – Not Such a New Idea: A Case Study of the Housing Construction Industry in Western Australia". *Proc. 10th Australasian Conference on Information Systems*, pp 825-836

Tyer R *et al* (2004) "Portal Framework for Computation within the eMinerals Project", Proceedings of UK e-Science All Hands Meeting 2004

Zhang Y *et al* (2004) "JUMBO – An XML infrastructure for eScience", Proceedings of UK e-Science All Hands Meeting 2004