Computer Supported Cooperative Work (2006) 15:385–411
DOI 10.1007/s10606-006-9025-0

Walking the Tightrope: The Balancing Acts of a Large e-Research Project

KATHERINE A. LAWRENCE
Collaboratory for Research on Electronic Work, School of Information, University of Michigan, 1075 Beal Avenue, Ann Arbor, MI, 48109-2112, USA (E-mail: kathla@umich.edu)

Abstract. Although e-Research has received much attention and acclaim in recent years, the realities of distributed collaboration still challenge even the most well-planned endeavors. This case study of an e-Research project examines the 'balancing acts' associated with multidisciplinary, geographically distributed, large-scale research and development work. After briefly describing the history and organizational design of this information technology and atmospheric science research project, I identify five paradoxical challenges that cannot be resolved: research versus development, harmony versus conflict, consensus versus top-down decision making, frequency and modes of communication, and fast versus slow pacing. Although collaboration and communication technologies supported the project's management and organization, most of the complexities faced by the team were not technological in nature. From the five paradoxical challenges associated with the project, I distill three cross-cutting issues that could be relevant to other e-Research projects of this magnitude: satisfying the multiple needs of a multidisciplinary project, managing information, and engaging all participants. I identify the practical implications of these challenges and issues, specifically that organizational and low-tech solutions – not the introduction of more sophisticated technology tools – are needed to solve these challenges and to better streamline coordination.

Key words: collaboration technology, communication, distributed work, e-Research/cyber-infrastructure, geographic dispersion, information management, multidisciplinary scientific research, paradox, project management, technology development

1. Introduction

In recent years, scientists and funding sources alike have heralded the opportunities afforded by e-Research, also known as cyberinfrastructure (CI) (Olson et al., 1998; Atkins, 2003). However, the development of these e-Research technology platforms involves a complicated blend of skills and interests. In particular, to build a cyberinfrastructure that supports a specific domain of science requires collaboration between scientists in that domain and computing experts interested in the specific challenges of high-performance computing. In the United States, several such initiatives have been supported through the National Science Foundation's Information Technology Research (ITR) program, which awarded grants between 2000 and

2005. The goal of the ITR program was to provide long-term support for high-risk IT projects with the potential to have an innovative impact on other science and engineering fields.

This paper offers a case study of an e-Research initiative at a midpoint in its progress. This portrait is an effort to help alleviate the scarcity of workplace studies that portray the gritty realities of collaborative work (Schmidt, 2000). The project under examination brought together meteorologists and computer scientists to develop a cyberinfrastructure for numerically predicting weather phenomena. This project, called Linked Environments for Atmospheric Discovery (LEAD), is developing tools for data discovery, analysis, simulation, and visualization. LEAD started in 2003 and was funded by the ITR program for five years. Like many ITR grants, the project involved researchers at the leading edge of their area of expertise, working at a distance from locations around the United States.

For many years, research scientists have experienced the difficulties of collaboration across geographic distances (Finholt, 2003). Distance collaborations demand a substantial overhead of time to coordinate logistics and other day-to-day features of collaboration (Cummings and Kiesler, 2005). For example, as the number of participating sites grows, collaborators have narrower windows for finding suitable meeting times across time zones, and the cost of travel for face-to-face meetings increases. In addition, the potential for the serendipitous generation and exchange of ideas that comes with close proximity to colleagues is simply lost (Kraut et al., 1990; Allen, 1997).

Although some technologies such as e-mail and videoconferencing have made document sharing and meetings easier to accomplish, the use of collaboration technologies brings its own complications. Technology is often expensive, and a single set of equipment may be shared by many people or departments, requiring rigid or advance scheduling. Sound quality may be impaired, rendering unfamiliar regional accents unintelligible. Sometimes cutting-edge technologies are unstable and difficult to use, and technical support staff, when available, may be unfamiliar with the newer standards or systems that underlie first-generation collaboration technologies. Thus, distance collaboration requires extra time and patience.

On top of this, routine communication between disciplines is more complicated. For example, terminology used in one specialized field is difficult to translate to another or may not even be recognized as something that needs to be translated (Heath and Staudenmayer, 2000). In some cases, the same term is used in different fields yet carries different meanings with it, so people must establish a common understanding of what they mean (Clark and Brennan, 1991). Identifying these pitfalls may require a rare person who can speak both disciplinary languages and recognize where communication gaps might exist. These factors add to the existing complexity of distance collaboration.

Despite complications such as these, by all accounts LEAD was already a successful collaboration by the end of its second year, in that it was able to coordinate a widely distributed team to pursue a common goal that was pushing the boundaries of computer science and meteorology. Yet LEAD also faced challenges that could never be fully resolved. These difficulties illustrate the paradoxical demands that multidisciplinary e-Research projects may need to balance. I begin an exploration of this project with a description of my involvement with this project and the data I collected. Then, after offering a brief history of the project that focuses on the origins, organization, and communication methods of the project team, I identify the specific, paradoxical challenges associated with this large, distributed e-Research project, describing the tradeoffs associated with each of these challenges. Curiously, my findings suggest that issues of organization and management are far more pressing than technological issues. What makes this case interesting and unique is that the project team is prevailing against the common technological barriers associated with this kind of work, and this success makes it possible to identify the more subtle but intrinsic management challenges that might otherwise be obscured, particularly those that may be characteristic of e-Research projects. From this case study, I distill several issues that similar projects involving large, distributed research teams might consider, emphasizing the multidisciplinary, information management, and motivational hurdles they are likely to face. I conclude with a discussion of the practical implications of these findings, asserting that technology should take a back seat, supporting role in helping e-Research projects to succeed.

2. Data collection and analysis method

The analyses reported in this paper were based on six months of intensive observation of and interviews with the LEAD project team. My role on the project was to study their process for developing CI. From a social science perspective, I was to offer support for collaboration and implementation and observe how my interventions took effect. The LEAD project represents one portion of an iterative study designed to observe two or three CI projects and derive best practices and requirements analysis assessment tools that could support similar development efforts in the future. This paper portrays a period in this study prior to initiating any significant interventions.

I attended more than 60 h of meetings via Access Grid (AG)[1] videoconferencing or via teleconferencing. I monitored traffic on five e-mail listservs that had frequent exchanges, sometimes more than 20 messages in a single day. I visited six of the nine participating institutions in person, for a total of more than 15 days on site, touring facilities and interviewing participants. These visits also included two 'All-Hands Meetings,'[2] a formal project review by the National Science Foundation (NSF), and a conference on geosciences information technology that was organized by LEAD. I also collected documents from meetings, e-mail exchanges, and the project's website, as well as formal status reports and published articles about the project.

I conducted approximately 28 h of semi-structured interviews with 25 project participants, including principal investigators (PIs), research scientists, software programmers, graduate students, support staff, and affiliated 'teacher partners.' The length of the interviews ranged from 32 min to 2 h and 15 min, with a mean length of 67 min. These interviews covered the history of the project and each participant's involvement in the project, the project's collaboration and coordination requirements, approaches and technologies for accomplishing collaborative work, issues of multidisciplinary collaboration, and the project's challenges and successes. Most interviews were digitally audio recorded, after which each interview was reviewed and summarized in a format that captured the major topics of the interview protocol, including direct quotations for illustration. These summaries were supplemented by handwritten notes taken during the interviews.

The analysis of the observations and interviews took place throughout the duration of the observation period described in this paper. As I took fieldnotes during meetings and as I read e-mail communications and project documentation, I kept track of recurring or interesting events and interactions. I also noted specific challenges that the project team seemed to face. Later, as I conducted interviews, I also listened for comments that confirmed or contradicted my observations, as well as new themes or issues raised by the participants. Once the interview recordings had been summarized, I reviewed them for common themes, looking for both the manifest and latent content expressed by the participants (Boyatzis, 1998). Some issues were very self-evident, such as the volume of e-mail communications and meetings and the load that these interactions created for the team members. As I moved between fieldnotes, interview notes, and ongoing meetings and interactions with the project team, I considered how newer themes and incidents might relate to earlier ones and incorporated those insights. For example, I found that people at multiple levels of the project – PIs, students, and staff – expressed concerned about the shifts in priorities that accompanied the different ways of describing the project tasks and the consequences that these changes had for their own work. Meanwhile, I had also witnessed the discussions and debates about how to describe and define tasks. From my longitudinal observations and the participants' own words about these topics, I identified both the research versus development and pacing issues. Finally, as I assembled this paper, I conducted a holistic analysis of the data, searching for the gestalt, or 'unifying nature' of the team members' experience of the project (Patton, 2002, p. 59). Ultimately, the themes I have identified are very much grounded in the participants' experience (Glaser and Strauss, 1967), but my interpretation of what they hold in common and how they are distinct is how I have made these experiences tell a coherent story. The three general issues for e-Research are abstracted further from the five original 'balancing acts,' identifying what is common among them. Based on these analyses, the following sections attempt to capture the overarching themes as well as the contradictions within them.

3. A brief history of the e-Research project

3.1. THE PROPOSAL WRITING STAGE

The precursor of the LEAD project was a two-year research expedition called the Modeling Environment for Atmospheric Discovery (MEAD). The purpose of that project was to enable the retrospective study of hurricanes and storms through simulation, data mining, and visualization on the TeraGrid[3], an enormous, integrated computational infrastructure for scientific research that is sponsored by the NSF. Members from the participating institutions in MEAD began discussing a subsequent project that would allow real-time, on-demand, and dynamically adaptive analysis of mesoscale[4] weather phenomena. Traditionally, weather analysis and forecasting is informed by data collected from static instruments, and no systems exist for redirecting instruments and computational resources toward potentially interesting or threatening weather patterns as they appear. The LEAD project was meant to be a paradigm shift for the field of mesoscale meteorology, using CI to conduct forecasts and analyses that respond to emerging weather phenomena. The LEAD CI would also allow the integration of large, disparate datasets to improve forecasts. Moreover, high-performance systems to do this kind of data-intensive modeling are typically available to very few researchers. LEAD would allow people to have access to cyberinfrastructure-supported technologies and data with a simplicity that had never before been available.

The originator of the ideas behind LEAD was reluctant to head up the project, recognizing that he had neither the computer science background nor the time available for running a new large project. Nevertheless, over time, as he found people eager to join the project, he became the de facto project director. His selection of participants was driven by what the project needed to do, both technically and in terms of domain content and educational outreach, not based on the geographical convenience of their institutional affiliations. In other words, he sought cutting-edge expertise and leadership in an area rather than the chance for collocation. Many of these participants had been part of the MEAD project (or other collaborations in the past) so had a long track record of working with each other. Other participants were invited for their expertise with undergraduate and graduate education and research in the atmospheric sciences. The PIs were excited by the opportunity to work on a project that was technically challenging and intellectually stimulating for both the computer scientists and the atmospheric scientists on the project. Altogether, the team held expertise in grid computing, data mining, high-performance computing, data provision, software development and support, meteorological modeling and analysis, undergraduate meteorology education, graduate atmospheric sciences education, and radar.

Their initial proposal to the NSF was rejected. The team revisited their proposal, and in their second application, they were funded as an ITR grant. Eight of the institutions were funded for the entire five years, and one was to join the project in its fourth year. With an ambitious charge ahead of them, the project officially started in October 2003.

3.2. ORGANIZATION OF THE PROJECT

From the beginning, the project faced a tremendous challenge of organizing and coordinating the nine geographically distributed institutions, which brought together more than 100 people. Three of the institutions concentrated on computer science, three offered primarily meteorological expertise, and three specialized in a blend of the two disciplines. (Figure 1 shows the geographic and disciplinary distribution of the participating institutions.) The team included principal investigators (PIs), research scientists, research programmers, software engineers, graduate students, post-doctoral fellows, and project coordinators. The core group was equally divided between those with a technology focus and those with a meteorology focus. Out of all these people, only one was funded full time on the project; the rest were committed to additional projects and thus had to balance their time commitments.

To manage and coordinate the multiple elements of the project, the team was divided into 'thrust groups.' Initially, the project included four thrust groups focused on specific areas of expertise – Data, Tools, Orchestration, and Meteorology – plus two cross-cutting test beds – Education & Outreach and LEAD Grid & Web Services. Each group included members from each discipline and each institution concerned with that part of the system, and each was led by PIs from two institutions. Over time, the Tools thrust group became less active, and a Portal thrust group formed as their needs became more apparent. Curiously, other than the PIs running each thrust group, membership in the groups was quite loose, and many team members participated in more than one. Some initially participated in several groups but reduced their involvement due to the extensive amount of time required for participation. Even though the thrust groups maintained multidisciplinary, cross-institution involvement, the localized expertise at each institution was

Total number of people includes core team members and less active participants.

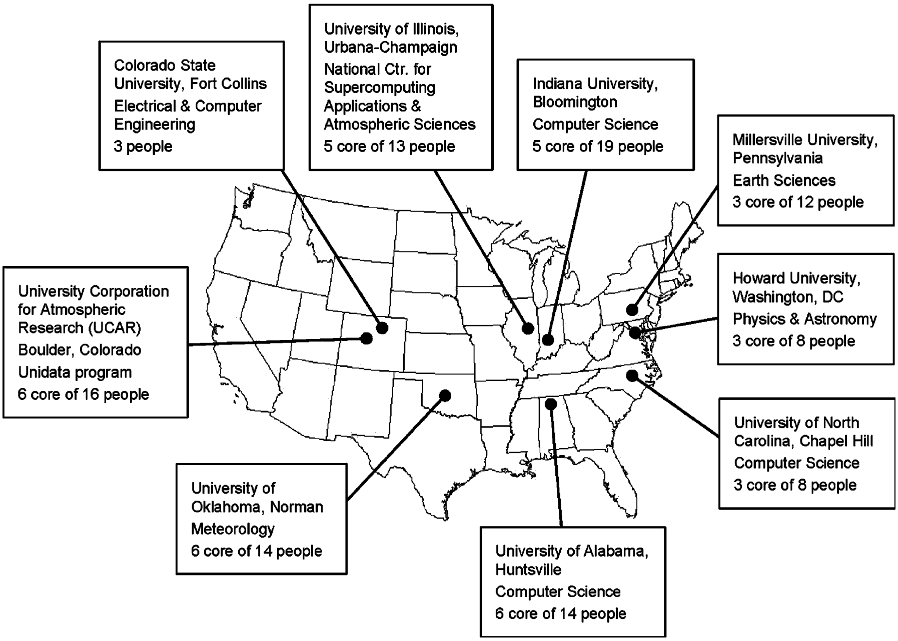

Figure 1. LEAD participating institutions, locations, disciplines, and team sizes.

such that particular aspects of technical research and development were often concentrated in a single institution.

3.3. COMMUNICATION

Appropriately, the team made use of available e-Research technologies to support the development of their own e-Research technology, but they also relied heavily on more 'old-fashioned' methods of communication. Initially the project primarily conducted regular meetings by videoconference using the AG. In every two-week period, six thrust group meetings were held on the AG, plus every month included a two-hour All-Hands Meeting and a one-hour Education meeting on AG. In addition, the PIs were scheduled to meet every two weeks on a teleconference. The project was budgeted to support two face-to-face All-Hands Meetings each year.

Over time, the load of such frequent AG meetings became onerous, and the AG meetings were reduced to a weekly, two-hour, All-Hands session and a monthly Education meeting, with occasional, as-needed meetings for other thrust groups. Occasionally the AG was used to offer tutorials on topics and technologies that were particular to one discipline yet unfamiliar to the other.

One person said about the tutorials: 'That helped a lot, I think, in connecting people about the project and crossing these disciplinary boundaries.' Alongside the AG technology, the team sometimes conducted demonstrations of prototype LEAD technologies using VNC (Virtual Network Computing), which allows people at multiple remote computers to view or control another computer's display via the Internet. Overall, the team's synchronous interaction shifted toward decreasing specialized thrust group AG meetings and increasing All-Hands, PI, and face-to-face meetings, particularly on specific topics or pressing issues.

Layered on top of these modes of communication were 10 e-mail listservs serving all thrust groups and test beds, several subsets of administrators and PIs, and the project team as a whole. These listservs had heavy traffic. For example, in the month preceding the project's initial, mid-term site visit[5] from the NSF, the list going to all the team members circulated 187 messages. As a complement (or perhaps a remedy) to the volumes of e-mail traffic, the team briefly experimented with Instant Messaging (IM), but that technology did not catch on particularly well. No one seemed to have an explanation as to why IM did not spread. If prior research offers any indication, perhaps they were so busy that asynchronous communication was easier and more effective than introducing a new mode of communication; possibly the equivalent of 'chat' took place mostly among those who were collocated (Handel and Herbsleb, 2002; Herbsleb et al., 2002). Although most documents were circulated by e-mail, a password-protected web site stored copies of many of the documents that were relevant to the entire group, such as reports, project plans, and slide presentations. The internal website was set up as a traditional web server so that documents had to be posted by the project coordinator for download by others, which unfortunately meant that spur-of-the-moment document sharing was always conducted by e-mail, and documents could not readily be shared through a common repository.

4. Findings: the balancing acts

The LEAD project was, by most standards, well managed. The formal structures for organizing were thoroughly conceived, yet they remained flexible enough for the inevitably evolving needs of the project. Communication was extensive. The team members were very collegial with each other – both within and across disciplines. Many interview participants remarked how friendly and pleasant the entire experience had been. One remarked, 'We have a bunch of mature people who know how to interact and work through problems. Lots and lots of dialogues through e-mails, AGs, telecoms, and visits.' Some attributed this harmony to the charisma and enthusiasm of the project director, while others felt it was due to the precedent of trust and respect that had been built by working on prior projects together.

Many circumstances had converged to make the project an operational success. When asked what was their greatest success to date, team members raised the same themes again and again. They recognized that to have identified a common goal and to move toward that goal with strong commitment was a healthy early indicator of success. One PI suggested, 'I think we're unique from other ITRs in the cohesiveness that the group has, our enthusiasm for a common goal, our tight integration of the education group. So again, I am extremely optimistic that we can succeed where others have failed. So I think, in the last couple years, in building that trust, in building those relationships, I think this has been our success.' The team members were proud of their ability to educate each other, and out of that understanding they had already made 'baby steps' of research and development progress by the end of the second year. Not only was the technology coming together, but outside organizations were aware of the importance of their work and were eager to have some association with the project.

Even so, the project faced challenges of balancing between equally compelling alternative directions, *both* of which needed to be satisfied for the project to fully succeed. First, the team had to allocate their energies to a blend of research and development. Second, the team members needed to encourage harmony while welcoming the benefits of conflict. Third, decisions sometimes were made by consensus and other times required a top-down mandate. Fourth, the project team had to cope with multiple modes of communication among the large number of individuals involved, finding a way to share sufficient information while not overwhelming themselves. Finally, the project had to move forward at a pace that was fast enough to get the job done but slow enough to fit the realities of the project. These paradoxes were not a matter of choosing between one option or another, but rather finding a way of comfortably accommodating both realities. I describe each of these 'balancing acts' in the sections that follow.

4.1. RESEARCH VERSUS DEVELOPMENT

The most apparent tension in the project was the necessity of balancing research versus development priorities. Officially, the project was a computer science *research* grant, but at the same time it was expected to *develop* technologies for use in a disciplinary context.[6] Ultimately, meteorology researchers and educators needed a product with which they could do their own research and experimentation. Thus, as the project advanced, the participants faced an ongoing debate as to how precisely they needed to define the project's components. On the one hand, for the purposes of developing an appropriate set of features, ensuring those features fit together, and estimating the time it would take to complete their development, the team needed to define specifications and architecture. These specifications and architecture would also smooth communication and agreement about what they were building. However, for those computer scientists engaged in pure research, too much specificity would possibly constrain the 'discovery' intrinsic to research and the option to pursue new pathways as they emerged. Moreover, given the unknown elements in the process, research could not necessarily happen in predictable time intervals. The researchers' need for unrestricted flexibility was as essential as the developers' need for clarity.

The needs of the atmospheric scientists on the project hovered somewhere between these two perspectives. Although they could be engaged with and offer opinions about the priorities, content, and design of the final product while the computer science researchers conducted their research, they also needed a viable product before they could use it for research or education in their own areas of interest (Atkins, 2003). Well before the end of the project, the atmospheric scientists would need an opportunity to 'test drive' the software, see how it was meeting their needs, and provide feedback to guide later versions of the product. Furthermore, as a consequence of the 'paradigm shift' they expected to generate, conducting a requirements analysis was difficult to do. To date, not many atmospheric scientists had ever had easy access to the computing power and data that LEAD would make available, so future users could not readily imagine how they might incorporate such a resource into their existing work. Even atmospheric scientists using the most advanced analysis and modeling technologies had never used a *dynamically adaptive* system, so they too had to imagine how they might use it. For these reasons, the team could just as well make an educated guess, develop prototypes, and then get feedback from potential users of the product (Boehm, 1988). Again, to do this, they needed software in hand.

Ironically, the tension between research and development was less of an issue across disciplines than within the computer science discipline. One of the institutions involved in the project was invited to join the project for its skill at deploying and supporting applications and data resources to the atmospheric science community. This group's priority was developing a product that its constituents could reliably use whereas the computer scientists conducting research were more accustomed to demonstrating a proof-of-concept – sometimes unstable or highly customized software which often required the people who created it to run it. Again, the atmospheric scientists were not partisan in this discussion; they sympathized with the computer scientists' need to do research while desiring to have a working product to do their own research and education. As one person summed up the issue, 'If we go back to our original proposal, it says that we will make available to the community a robust and reliable LEAD capability which the community can use, grow, and develop as an education or research tool. That says to me that the end user has to have something that is appropriate to the level of that end user and is deployable.'

In sum, to fulfill research opportunities for not just the computer scientists but also the meteorologists, some degree of development needed to be done. But in doing so, the computer scientists would be constrained, and in *not* doing so, the meteorologists would face delays in their own research and education efforts. Ultimately, the most typical way of balancing the two priorities appeared to be episodic bursts of development, in anticipation of impending demonstrations, conferences, or reviews by external parties, underscored by steady research. By the end of the second year, however, the need to develop a product that could be released to a friendly testing audience became much more immediate, and the team began more systematic efforts to specify and build an initial, limited version of some of the more basic features.

4.2. HARMONY VERSUS CONFLICT

A second issue was how to strike the right balance between agreement and discord. One of the hallmark features of the project was the collegiality of the team and how much they genuinely enjoyed working with each other. On the other hand, so many experts with different bases of experience inevitably endorsed different solutions to the same problem. Similarly, the team members' roles on the project and professional goals tended to promote different priorities that did not always mesh with those of some of the other members of the team.

When members of a group have different functional backgrounds, conflict about how to perform a task is to be expected (Pelled, 1999). This type of conflict over task-related goals and processes can be productive, promoting more original and complex solutions through the juxtaposition and resolution of these opposing perspectives (Hirshberg, 1998; Kurtzberg and Amabile, 2000). The project director recognized the value of these differences, and he welcomed the occasions when team members disagreed. At the same time, some project members expressed uneasiness about the different decision paths that people were taking and the time it would take to reconcile those differences later in the project. Also, the team members needed to maintain a certain level of comfort with each other and feel that their hard work would not be for naught. The unspoken challenge was how to balance the benefits of diverse thinking and conflict with the equally valid advantages of harmony.

The distributed nature of this e-Research initiative complicated the resolution to this issue. Unlike people working on projects that are collocated, many of the team members on this project did not have the day-to-day opportunities to know each other's dispositions and perspectives on work, and the technologies available for communication did not necessarily enhance their understanding. One person explained the frustrations of the

Access Grid this way: 'No personality development and no relationship development happens on Access Grid. That is the main point of interaction. You really don't know the people who you are dealing with. So there is no progress over time... It has to get better over time, but that doesn't happen on Access Grid.' Likewise, e-mails lack most of the non-verbal cues that convey the overall tone of the communication (Friedman and Currall, 2003). Of course, team members indicated that they communicated more frequently with those who were collocated, further solidifying their connections to those at their own institutions.

This complex social structure, however, did produce at least one distinct benefit. Most of the PIs on the project had worked together on prior projects, forging cross-disciplinary connections long before LEAD began. Indeed, their familiarity with each other and the friendships between many others on the project, now at different locations, traced back to time spent at the same university or university department or through advisor-student connections. They had prior experience of working well together, and their rapport appeared to spread to the rest of the group. One PI explained his theory of why their interactions were so successful:

> One reason I think by and large this project works really well is that...almost everybody in the project has known each other a long, long time. So they satisfy what I call the 'Principle of Least Surprise,' which is: If I have a certain set of data, I can predict what the other person will do, and they can predict what I will do. Therefore they will act on prediction. And the 'Principle of Least Surprise' says that you are almost always right.

This kind of familiarity and knowledge encouraged trust. Members of the team also indicated that they felt that their particular strengths were valued and respected. One person explained that their significance was established this way from the start:

> I think a lot of it is maybe because of the way that we were asked to participate...I think everyone that was brought into the project was sort of introduced as ... 'Okay, this person is brought in because they have this expertise or because we know – they've done *something*.' ... So I think people sort of came into it feeling special, in a sense, and then at the same time the other people were aware that, 'Hey, this person is here because of *this*.' That they're just not somebody that you grabbed – and I think that had something to do with it. A lot of those things are – The way they start out is the way they stay.

Based on past experience with NSF site visits, both insiders and outsiders to the project believed that this goodwill was not very common among large, multidisciplinary projects and felt that this project was noticeably harmonious compared to others they had observed or had participated in.

Thus, the LEAD project team was productive and satisfied despite – or perhaps because of – the differences between them. Nevertheless, the communication media used to bridge distances did not readily help team members grow more comfortable with each other. However, the shared history between many of the team members generated an overall tone of goodwill that made differences of opinion less consequential.

4.3. CONSENSUS VERSUS TOP-DOWN DECISION MAKING

A third balancing act concerned decision making. On a project as large as LEAD, many decisions needed to be made at all levels and degrees of complexity. Overall, standards and interfaces between components needed to be set. Interdisciplinary subsets of the team needed to establish unprecedented aspects of the software's functionality that combined cutting-edge technology with sophisticated elements of meteorology. Faced with evolving and uncertain technologies that supported their own project, the team sometimes needed to second guess the trajectory of a particular product or standard. The selection of priorities and goals could thereby affect activities at all the institutions involved. On a daily basis, smaller decisions about design, capabilities, and sub-components needed to be made.

These decisions often required input from both disciplines and thus involved multiple institutions. Inevitably, their opinions or preferences differed some of the time. Although discussion of the options was polite and participatory, a clear path to a decision was not always self-evident. The project director encouraged a consensus-making approach, partly motivated by his opinion that he did not have sufficient expertise about computer science issues to dictate the solution to a technical dispute. But arriving at consensus in these circumstances could take a long time and often could not be fully resolved without someone making the final decision. Such decisions could be awkward. For instance, at one face-to-face meeting at a crucial junction in the project, some of the PIs pushed for dividing those at the meeting into subgroups in order to move the planning process forward. Others were not so sure that breaking into groups would be effective because of the interdependencies of the components they needed to discuss. Finally, the project director said, 'I sense that some people are kind of uncomfortable right now, and I'd like to hear from anyone to sort that out.' In response, people raised concerns about various facets of the project, and by default, the group ended up working as a whole throughout the remainder of the meeting. In cases like this, consensus could neither be found nor forced, and arriving at a satisfying decision required a bit of both tactics.

Another decision-making complication that may be fairly unique to e-Research projects spread across multiple academic institutions is that academia does not have an authority structure that is comparable to most other types of organizations such as commercial businesses, non-profits, or governmental agencies. Once tenured, professors have a great deal of freedom. Very little can be leveraged to influence them or force them to comply with an externally made decision or procedure. Because the PIs had final accountability for the use of the LEAD project's funds, they had to do what they felt was wisest; however, some institutions were best served by pursuing technology formats or protocols that were not necessarily compatible with those at other institutions. The challenge was how to integrate the differences or force a choice between them when neither had to comply.

A related decision-making issue was who beyond the PIs had the authority to make decisions. The project managers or coordinators without PhDs did not consistently have the clout to make recommendations that those with PhDs would follow. Although the project had administrative and management guidance by two very able individuals with undergraduate degrees, they themselves and others expressed concern that their decisions or requests were not readily followed by those who did not see them as authorities. In addition, the PIs – who were often busy with other responsibilities – could become decision-making bottlenecks when they did not endow their students or staff with the authority to make day-to-day decisions, and because the project was so vast and complicated, it was hard for those students and staff to necessarily know if the size of the decision merited upper-level oversight.

In these ways, the project had a constant need to balance where and how decision-making authority emerged. Although consensus appeared to be the best approach, it was not always feasible. Moreover, valuable expertise existed below the responsibility level of the PIs. Consequently, the team faced an unresolved tension over the optimal decision-making strategy for many situations.

4.4. WHAT IS THE 'RIGHT' AMOUNT OF COMMUNICATION?

Another challenge was how to have sufficient communication among the team members without inundating everyone. Frequent and inclusive communication appeared to maintain energy and momentum on the project, so the LEAD team members tended to err on the side of keeping others informed. Also, due to the multidisciplinary nature of most facets of the project, team members could not always be sure of who would want or need to know about something. Consequently, meetings and messages were frequent and voluminous. But finding the right balance between dissemination and restraint was not easy.

E-mail alone was overwhelming. As described in Section 3.3, the amount of e-mail on the LEAD project was atypical. As one PI remarked in an interview, 'This project generates more e-mail than any two projects I've worked on.' Team members tended to be inclusive, recognizing the benefits of maximizing others' input, yet people could not keep up with the flow. Many team members indicated that they already had too much other e-mail, so the volume of LEAD messages was yet another drain on their attention. The unfortunate consequence was that several people expressed concerns that when they sent out messages to one of the listservs, they failed to get the response that they would have expected, as if their messages had fallen through the cracks. Consequently, they had to follow up with specific people in order to resolve whatever issue they wanted to address. Using e-mail for scheduling, too, was complicated because some people were unreliable about responding to requests, due to frequent travel and other commitments that created e-mail backlogs.

Meetings, too, were more difficult to facilitate across institutions and time zones. Initially, the biweekly thrust group meetings on AG seemed useful, but the frequency and time commitment became too burdensome, which led the team to decrease the number of meetings to occasional, as-needed thrust group meetings plus a weekly All-Hands AG meeting. Moreover, people had definite preferences for ideal modes of communication. Some thought that AG was great, since it offered the visual cues that teleconferences could not. Others expressed frustration at the variable sound quality and inherent technical instability of AG. Some felt that teleconferencing – despite lacking the visual cues – felt more conversationally intimate. Consistent with prior research, some suggested that nothing was better than a face-to-face meeting for making substantial progress on an issue, and the electronic modes of communication were merely effective as a follow-up medium (Finholt et al., 1990). One PI explained their process:

> When we do get together, we do have very fixed sets of issues you have to deal with... You can very quickly isolate... the hard problems [that] we sit down face-to-face to solve... And then you can get something solved, and then you can use electronic mediation afterwards to work out the final details. You develop the final plan together face-to-face and then the rest of it, instead of dropping into a bucket, you have a continuation from that meeting.

However, face-to-face meetings were costly in terms of both time and financial expense; inevitably someone had to travel a distance or change time zones for any such meeting, so choosing a venue was always a compromise. Finding a time when everyone was available to invest one or two days was not easy either. After hosting meetings at different participating institutions during the first two years, the group eventually decided to minimize travel time and cost by scheduling their All-Hands Meeting at a centrally-located, major airport hotel, within driving distance of two of the institutions involved.

Other issues emerged in the tradeoffs between communication specificity and efficiency. For example, in the course of reducing the volume of meetings and condensing most AG meetings into a regular 'meeting for all,' more of the meeting time appeared to be less relevant for more of the people involved. In other words, the meeting topics were not relevant to all of the people all of the time, plus many topics could only be discussed at the most general of levels due to time constraints. Some people were concerned whether it was even appropriate to bring up detailed questions in such a meeting, but then when was the right time to raise them? Another problem was that people realized that much of the meeting would not be relevant to them, so they would multitask, replying to e-mail or doing other work during the meeting. Unfortunately, this meant that people were not fully 'present' in meetings, and questions posed to the entire group would often be met with complete silence or just a handful of people seemed involved enough to respond.

Certainly, the challenge of using these different types of communication in an optimal way could have no simple resolution. However, the particular sticky points faced by the team suggest that some difficulties could have been related and reinforcing each other. For example, perhaps the reduced quality of particular communication media tended to also reduce the overall quantity of *effectively* transmitted communications, thus requiring more volume to make sure that the communication found its target. In other words, when e-mails were too numerous to process or meetings were too time consuming to give undivided attention, perhaps the lack of response from others prompted more e-mails and additional meetings to happen. Another possibility was that the communications needed to be partitioned into more discrete units, either by dividing group membership into more specific sub-topics or by aggregating and filtering some communications through a single, less frequent, regular communication medium such as a project blog.

4.5. FINDING THE RIGHT PACE

The fifth issue of balance was how fast or slow the group should move forward with elements of their plan. Part of this issue was driven by the concern (described in Section 4.1) of not rushing the research agenda too much and cutting off future possibilities for exploration. But overall, the thorniest issue was that the extraordinary magnitude and complexity of the project made grasping the requirements and interactions of all the parts very difficult.

Even while thrust groups were experimenting with aspects of the technology and determining standards and interfaces, the entire group was developing ways of understanding the project as a whole. At the start of the project, the LEAD components were described as 'building blocks' and later called 'atomic tasks,' but these concepts did not appear to make sense to many people on the project. At the All-Hands Meeting in the summer of 2004, the team mapped out a series of what they called 'prototypes' that would help guide their research and development, but the framework of prototypes never gained wide acceptance, perhaps because they still had substantial research to do before they could implement them. One year later, in August 2005, someone coined the term 'mini-LEADs' in an effort to capture the idea of small, functional units that could be used and tested by a limited number of friendly meteorologists. Corresponding changes to the graphic representations of the project and execution plans emerged alongside the name changes. Unfortunately, by the time the project had run for nearly two years, team members were expressing confusion and concern at the latest new way of describing the project. They felt that using so many different labels over time to describe the same thing was problematic. One person explained why this was a problem:

> I think there is way too much emphasis on documentation... of plans, projections, and definitions. That has changed over the two years that LEAD has been ongoing. We used to call [what is now] mini-LEAD 'prototypes,' and now we call [them] 'mini-LEAD' ... Mini-LEAD is nothing but prototypes...So everyone who was already struggling trying to collaborate through communication between these various groups now has to go out and become familiarized with new jargon that all really means the same thing.

Some believed that the new labels had prevented the team from making concrete progress by keeping them focused on higher-level discussions, whereas others speculated that the delay in development progress had more to do with the need for additional research before they could be ready to implement a stable product.

When I began to observe the project, midway through their second year, a contentious issue at an All-Hands Meeting was whether it was appropriate to define more specific requirements for the upcoming phases of research and development. The concerns of both the development-oriented members and the research-minded members were duly acknowledged, but at the time, they achieved no resolution. By default, the project team continued with their research-focused activities. In parallel, the team made some efforts to cement some requirements. In anticipation of the site visit, the group co-authored a system architecture and requirements document that outlined in great detail how each component would work, and each thrust group and test bed itemized their goals for the next 12 months. The evolving architecture and requirements document grew quite large – more than any one person could comfortably read or comprehend – yet it was a touchstone that, in theory, anyone could consult. Once the site visit panel expressed concern over the lack of an operational system for educational use, the project's priorities shifted toward more systematic development. Later, some PIs indicated that they believed that the team had not been ready for a conversation about development until that time; efforts to introduce it earlier had been unsuccessful.

Thus, although the project had a vision and an outlined plan, the reality of how long each stage might take was unclear. Like many teams that take stock at a midpoint of a project, the group decided to shift priorities after the site visit (Gersick, 1988). Even so, the questions remained: When are people ready for a transition to a new stage in a large e-Research and CI development project? How long is it appropriate to discuss the big picture before focusing on the finer details?

4.6. SUMMARY

In this section, I have described a set of challenges that, despite good intentions and thoughtful execution, were inherently irresolvable. Instead the team found itself 'walking the tightrope,' trying carefully not to lean too far in either direction. Although some of these issues may be specific to this project, finding the right balance – between research and development, harmony and conflict, consensus and top-down decision making, frequency and modes of communication, and a fast or slow pace – suggests challenges that might generalize to other large, multidisciplinary e-Research projects. I discuss these in the next section.

5. General issues for e-Research

As unique as the LEAD project may seem, this case study reveals several crosscutting issues that are likely to complicate even the most successful and harmonious e-Research endeavors. These issues fall into three main areas: satisfying the multiple needs of a multidisciplinary project, managing information, and engaging all participants. I explore each of these below.

5.1. SATISFYING THE MULTIPLE NEEDS OF A MULTIDISCIPLINARY PROJECT

Any large project will include people with different roles and objectives, but an e-Research project that involves multiple institutions and different disciplines in each of those institutions is even more likely to encounter agendas that are somewhat independent of each other. The institutions have come together for more than just the overarching needs of the project – the individuals involved also want to advance their own research streams and secure additional funding for their immediate group or laboratory – and the project at hand is just one episode in a stream of projects.

The LEAD project indicates some of the ways that these multiple needs might be manifested. For example, some institutions may prioritize research while others prioritize development, and these goals may work at cross purposes. Underlying these motives might be differences in career stages; an academic who has tenure may be considerably more relaxed about when or where to publish than a more newly-minted researcher. In addition to the need to publish at a certain pace, an untenured researcher may not benefit from publishing in a journal that is multidisciplinary, interdisciplinary, or outside his or her field. Other people on a project may be more acutely aware of the development challenges they face and be eager to apply resources toward resolving open issues that might close off certain research opportunities for others. Educators might likewise want the project to speed up so that they can begin using the products of the research in their classrooms.

The decision-making process may also affect participants on a multidisciplinary project. Whether the process is top-down or by consensus, if the interests of a particular discipline are not championed at the right time, the corresponding part of the project may be delayed or forced to move too quickly. Some voices may be stronger than others, rendering them more influential than they should be. Even if the disciplines need not work closely together, their work may be highly interdependent such that certain choices by one disciplinary group or subgroup will affect the choices of another. These discrepancies may spur productive conflict leading to a solution that is better for all, or they may result in a compromise that fully satisfies no one. Taking stock of the multiple interests on a project can at least reduce the possibility of surprises later. Whether these interests are made explicit through conversation or simply by identifying the stakeholder characteristics, knowledge of these different needs might smooth the decisions, transitions, and differences that surface throughout the project.

5.2. MANAGING INFORMATION

As seen in this case study, a large e-Research project inevitably faces greater challenges for managing information. Collaboration between multiple institutions and disciplines complicates logistics, interaction, communication, and content sharing. Even with effective tools and systems, project participants can find themselves overwhelmed.

For example, when a project is collocated, the task of keeping everyone posted on developments, changes, or urgent issues may simply require a walk down the hall or a note left on a desk. Serendipitous interactions during the day can surface emerging problems and solve them in short order. In contrast, the members of a multi-institutional project must make special plans to meet in person, and all other communications require some form of intermediate technology. Face-to-face meetings must be arranged around multiple academic calendars, and if people on the project represent different disciplines, they may not have the benefit of touching base at a national or regional conference.

With intermediating communication technologies as a given, the members of a large e-Research project must work out – often through trial and error – what works best for them. For instance, as straightforward and common as e-mail has become, projects need to establish their own norms and expectations. Who needs to be informed about which events? What issues merit discussion or input from others? How quickly must someone respond to a message? Do all listserv messages require a response? What topics are better discussed in a particular format or venue? In addition, when a topic needs to be discussed in real time, a project group must decide what technologies are most effective for their circumstances, in terms of quality (video or voice only?), convenience (who has access?), and cost (what technologies must be upgraded or purchased?). Likewise, such meetings prompt questions of participation, frequency, and timing, particularly when spanning time zones.

Finally, the task of sharing documentation and data is rendered more difficult as well. Beyond basic compatibility issues such as computing platforms and applications, most institutions require special permissions for outsiders to gain access to servers and other technology resources. Shared documents can be posted on a Web or FTP site, but the security of the files and who will maintain the site are considerations. For ‘living’ documents or databases that are regularly updated, the group must be concerned with tracking changes and establishing the most recent version. Sharing documents by e-mail is more expedient but complicates version tracking. Newer technologies for collaborative editing, such as wikis and blogs, are beginning to be integrated into the mainstream for such purposes but are often slow to be accepted in that role (Tepper, 2003; Lamb, 2004). On top of this, there must be a framework for helping members understand the structure and relationship of each document to others so that information in context can be readily located (Schmidt, 1991; Mariani and Rodden, 1996; Grudin, 2006). Ultimately, shared repositories are useless if not accompanied by the social support and communication required to navigate them (Pipek et al., 2003).

Thus, managing information may require careful thought in advance and considerable flexibility over time. Technology has made information sharing and communication easier in some ways but has introduced new, unsolved complications. As seen in other industries, electronic information can be shared effectively among heterogeneous disciplines working apart from each other (Reddy et al., 2001). In the future, e-Research project teams need to be alert for new processes or technologies that are solving similar problems in other areas, whether research-focused or not.

5.3. ENGAGING ALL PARTICIPANTS

Finally, a large e-Research project at multiple institutions faces the challenge of keeping all the participants engaged with the project. Those who are involved in multiple projects are likely to be unavailable at least some of the time when their presence is desired. But engagement may be more than physical participation. Full engagement may involve intellectual or emotional attention to the project. As described in the LEAD case, there are several aspects of a large e-Research project that can make engagement more difficult to achieve.

First of all, various aspects of the project plan and execution may decrease engagement. As described earlier, priorities may be set toward an aspect of research or development or some other set of choices that may make participants feel that they have nothing to contribute to the current tasks. Also, differences of opinion that create tensions between different subgroups on the project may cause some to disengage, either because their preferences have not been represented or the strain is more than they want to feel. This raises the question of when the benefit of 'creative tension' is not worth the emotional cost to those involved. Furthermore, aspects of the project pace may leave some people feeling as though the project has moved forward without them while others may feel so overloaded that they cannot contribute as much as they need to. These issues are not unique to large, multidisciplinary, geographically distributed projects, but the circumstances of such a project may introduce factors that make disengagement more likely.

Another risk to full engagement is the multiple ways that people can be connected to the project. A project may come together because the principal investigators have personal or professional bonds, but their collocated colleagues may feel no such attachment to the project. The history between those who know each other helps them trust each other (and their associates) much more readily than those who are thrown together by an artificial association (Kramer, 1999). Also, the hierarchies inherent in academic research may make it difficult for the project to harness the available types of expertise. Leaders and subordinates monitor different aspects of a task, such that those at lower levels may have hands-on expertise that is unknown to those above (Yanow, 2001). Newer or loosely affiliated participants may lack a sense of the big picture and how their work fits into the project's overall goals. They may mistakenly perceive that they have little that they can add to other aspects of the project. Also, decision-making structures may imply (fairly or unfairly) that some people have little authority for making or influencing decisions, so they may not offer their opinions or become invested in the outcome of a decision to which they did not contribute.

Finally, communication challenges may introduce other means by which people become disengaged. People have preferred modalities of learning and for interacting with others (Gardner, 1983). For example, teleconferences may work well for the aurally inclined, but not for those who are tactile or visual. Fidelity of the communication technologies may cause people to be misunderstood, or at the very least, create an unnatural barrier that does not exist in face-to-face meetings. Moreover, these technologies allow people to 'attend' a meeting while their attention is directed to e-mail or other tasks, compromising their involvement. Those who cannot keep up with e-mail volume may not follow the thread of a subsequent discussion. In addition, language similarities and differences across disciplines can prompt misunderstandings. The same word may have very different meanings to people in different fields, yet they may not realize it. For example, on the LEAD project, some people thought that a software 'release' implied an out-of-the-box readiness for installation on any computer whereas to others it simply meant a somewhat stable version of the product that could be shared with interested colleagues. Specialists may not realize that they are using jargon that others do not comprehend. Thus, as people join a project, bringing different expectations and expertise, the challenge is to involve each of them to the degree that they are optimally challenged without being overloaded or aggravated.

In sum, even for the most compelling project, circumstances may conspire to disengage some participants. Planning, priority setting, and differences of opinion can leave out some people. Familiarity with each other can vary across the team. Communication technology and disciplinary language can get in the way of effective interaction. Whether through the project's structure or sequence of development, through the functional expertise of the team, or through some other heuristic, a system for involving people at multiple sites can be successful if tailored to the primary coordination needs of the group (Grinter et al., 1999). Not all needs can be satisfied, but recognizing differences may pave the way to greater involvement.

6. Conclusion

In this paper, I have documented the 'balancing acts' required by a large, e-Research project and have extrapolated several general issues that may be relevant to similar projects. The development of generalizable theory is limited by the nature of this case study, namely the idiosyncratic aspects of a single project. The participants in this study emphasized that the LEAD project was unusual among collaborations they had observed, particularly because the team had experience working together in the past, and these prior relationships had added a degree of tolerance and trust that smoothed over some difficulties. Even so, the broader management challenges they have faced are not new to distributed software development projects, but the features of e-Research appear to magnify them in some ways. For example, unlike commercial enterprises, universities do not have strong management controls that help motivate or punish employees, and a multi-university project has even less ability to regulate those involved. Hierarchies that would simplify decision-making or enforcement of deadlines are almost nonexistent in academic settings. Moreover, whereas corporate software projects have making a profit as the ultimate goal, funding for research and development projects such as ITRs attract participants for different reasons and rewards that may not be compatible. These prominent issues are clearer here perhaps because the project team has put so much energy into communicating as much as they can, thereby illuminating fundamental management issues rather than the more technologically-based issues identified in so much research on computer supported cooperative work.

Despite the unique aspects of the LEAD project, this study has practical implications for the types of solutions that might support distance collaboration on such projects. Specifically, such solutions may not be technological, and those solutions that require technology may not need to be particularly sophisticated. In fact, more benefit may come from less 'glamorous' solutions that simply streamline collaboration by acting as a multi-platform broker between scheduling or project management tools that are already in use at collaborating institutions. I briefly describe these ideas below.

First of all, this case highlights the reality that some of the significant challenges to successful e-Research projects are not necessarily solved with technology. In fact, issues of establishing priorities, creating development plans, handling disagreements, making decisions, and finding a comfortable pace of progress are not technology dependent except to the degree that technology can facilitate meetings and document exchange at a distance. Instead, the issues raised by this study require both awareness of the potential pitfalls and regular discussion among team members as to how well they are managing each other's expectations. Such conversation could accommodate the evolving nature of a project and routinize what might otherwise be an awkward or uncomfortable self-evaluation. If team members recognize that change and introspection are necessary and inevitable, they may feel more comfortable adjusting over time. At a higher level, funding sources could potentially provide guidelines, based on aggregate data from successful e-Research projects, for recommended frequencies of face-to-face meetings and how to budget accordingly. Likewise, the extra costs associated with multi-site coordination, in particular for the hours required to manage large projects, should be an assumed part of any budget.

Second, for e-Research projects that can be helped by tools and technology, the most sophisticated tools may not be the best solution. Given that e-mail was originally designed simply to expedite asynchronous communication, it has moved into a much more central role in day-to-day work than expected. Unfortunately, with a more prominent role, e-mail complicates personal information management and takes up large amounts of time (Whittaker and Sidner, 1996). In interviews, team members indicated that they did not want to introduce another piece of technology or another information source that they would have to monitor. This suggests that one possible solution is to find ways of using existing tools more effectively. For example, some e-mail applications allow the sender of a message to indicate the priority of the note, but these settings do not necessarily transfer between different e-mail applications. However, if users who require a prompt reply were to put an agreed-upon indicator in the subject line (such as an asterisk or a reply deadline), their teammates would know to respond as soon as possible. Such a solution is easy, convenient, and requires no sophisticated technology. Other studies of software management techniques have shown that low-tech solutions can be more useful and more readily accepted than computer-based tools that are meant to accomplish the same function (Pagonis and Cruikshank, 1992; Whittaker and Schwarz, 1999; Beck, 2000).

Finally, large gains may be made by streamlining the coordination needs of a large project. Along with time spent in meetings, e-mail seemed to demand the largest amount of time related to collaboration or coordination. Tools that could facilitate the task of scheduling meetings or providing status reports without clogging e-mail in-boxes would reduce the overhead associated with distance collaboration. For example, a tool that could seamlessly interface with project members' existing calendar applications, acting as a broker for scheduling meetings across institutions, would reduce one type of hassle. In general, any kind of tool that can simplify technology incompatibilities between sites – rather than introducing yet another standard – would offer a huge advantage.

As e-Research becomes more prevalent and more multidisciplinary projects are funded to develop the infrastructure for e-Research, additional case studies would be useful for elaborating the observations offered in this paper. New issues and new solutions may appear when the technology grows more powerful and more prevalent. In the meantime, the most pressing concerns appear to be those that simply reflect the basic challenges of collaboration and management. Such challenges have no easy answers. Perhaps as e-Research blossoms, producing successes and failures, the repository of lessons learned over time will serve to help future large projects like this one safely walk the tightrope.

Acknowledgements

I want to thank the members of the LEAD research and development team for their time and insights that informed this paper. I also want to thank Il-hwan Kim for his assistance with data collection and analysis, and Tom Finholt and Ann Zimmerman for their intellectual support of this paper. The research on which this article is based was funded by the National Science Foundation through the National Center for Supercomputing Applications.

Notes

1. The Access Grid (http://www.accessgrid.org) is a grid-enabled means of group-to-group videoconferencing that allows multiple locations to join the same virtual venue. Each location can display video from all parties on a large-format screen as well as view shared presentations. A text chat portion of the AG environment allowed those who were responsible for monitoring the connection software to consult with those at other sites without interrupting the ongoing meeting.
2. An 'All-Hands Meeting' is literally a meeting of all the people with hands-on involvement on the project. In reality, budgetary and time constraints often meant that some subordinate team members could not travel to a face-to-face meeting or a busy principal investigator might send another participant from his or her institution in his or her stead.
3. Since 2001, the TeraGrid project has received more than $98 million to connect eight partner institutions. As of September 2004, they provided more than 40 teraflops of computing power and nearly two petabytes of rotating storage over a dedicated national network operating at 10–30 gigabits/s.
4. Mesoscale meteorology refers to atmospheric phenomena on spatial scales as large as hundreds of kilometers, but often smaller. Mesoscale weather includes thunderstorms, terrain-induced windstorms, coastal storms, and hurricanes.
5. A site visit is a review of a project by a panel of experts that have been selected by the funding agency. Typically, site visits happen at a midpoint in a project to make sure that the project is on track and the funds are being spent appropriately. The project team is expected to prepare a report and summary presentations to convey the project status to the panel.
6. An interesting contrast is the approach taken by the Joint Information Systems Committee (JISC) in the United Kingdom. This organization explicitly finances and supports *only* the implementation and development of cyberinfrastructure and information technologies. Although these technologies are to be used in the context of teaching, learning, research, and administration, the JISC does not expect research to be a product of its funding.

References

Allen, Thomas J. (1997): Communication Networks in R&D Laboratories. In R. Katz (ed.): *The Human Side of Managing Technological Innovation*. New York: Oxford University Press, pp. 320–330.

Atkins, Daniel E., et al. (2003): Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure. Arlington, VA: National Science Foundation. Available at http://www.nsf.gov/cise/sci/reports/atkins.pdf.

Beck, Kent (2000): *Extreme Programming Explained: Embrace Change*. Boston: Addison-Wesley.

Boehm, Barry W. (1988): A Spiral Model of Software Development and Enhancement. *Computer*, vol. 21 no. 5, pp. 61–72.

Boyatzis, Richard E. (1998): *Transforming Qualitative Information: Thematic Analysis and Code Development*. Thousand Oaks, CA: Sage.

Clark, Herbert H. and Susan E. Brennan (1991): Grounding in Communication. In L.B. Resnick, J.M. Levine and S.D. Teasley (eds.): *Perspectives on Socially Shared Cognition*. Washington, DC: American Psychological Association, pp. 127–149.

Cummings, Jonathon N. and Sara Kiesler (2005): Collaborative Research across Disciplinary and Organizational Boundaries. *Social Studies of Science*, vol. 35 no. 5, pp. 703–722.

Finholt, Thomas A. (2003): Collaboratories as a New Form of Scientific Organization. *Economics of Innovation and New Technology*, vol. 12 no. 1, pp. 5–25.

Finholt, Tom, Lee Sproull and Sara Kiesler (1990): Communication and Performance in Ad Hoc Task Groups. In J. Galegher, R.E. Kraut and C. Egido (eds.): *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 291–325.

Friedman, Raymond A. and Steven C. Currall (2003): Conflict Escalation: Dispute Exacerbating Elements of E-Mail Communication. *Human Relations*, vol. 56 no. 11, pp. 1325–1347.

Gardner, Howard (1983): *Frames of Mind*. New York: Basic Books.

Gersick, Connie J.G. (1988): Time and Transition in Work Teams: Toward a New Model of Group Development. *Academy of Management Journal*, vol. 31 no. 1, pp. 9–41.

Glaser, Barney G. and Anselm L. Strauss (1967): *The Discovery of Grounded Theory: Strategies for Qualitative Research*. New York: Aldine de Gruyter.

Grinter, Rebecca E., James D. Herbsleb and Dewayne E. Perry (1999): The Geography of Coordination: Dealing with Distance in R&D Work. *Proceedings of International ACM SIGGROUP Conference on Supporting Group Work (GROUP), Phoenix, AZ, November 14–17, 1999*. ACM Press, pp. 306–315.

Grudin, Jonathan (2006): *Enterprise Knowledge Management and Emerging Technologies*. Winter Park, CO: Proceedings of the Human Computer Interaction Consortium.

Handel, Mark and James D. Herbsleb (2002): What Is Chat Doing in the Workplace? *Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW), New Orleans, LA, November 16–20, 2002*. ACM Press, pp. 1–10.

Heath, Chip and Nancy Staudenmayer (2000): Coordination Neglect: How Lay Theories of Organizing Complicate Coordination in Organizations. In B.M. Staw and R.I. Sutton (eds.): *Research in Organizational Behaviour*. Oxford, UK: Elsevier Science Inc., pp. 153–191.

Herbsleb, James D., David L. Atkins, David G. Boyer, Mark Handel and Thomas A. Finholt (2002): Introducing Instant Messaging and Chat in the Workplace. *Proceedings of the ACM Conference on Computer-Human Interaction, Minneapolis, Minnesota, April 20–25, 2002*. ACM Press, pp. 171–178.

Hirshberg, Jerry (1998): *The Creative Priority: Putting Innovation to Work in Your Business*. New York: HarperBusiness.

Kramer, Roderick M. (1999): Trust and Distrust in Organizations: Emerging Perspectives, Enduring Questions. *Annual Review of Psychology*, vol. 50, 569–598.

Kraut, Robert E., Carmen Egido and Jolene Galegher (1990): Patterns of Contact and Communication in Scientific Research Collaboration. In J. Galegher, R.E. Kraut and C. Egido (eds.): *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 149–171.

Kurtzberg, Terri R. and Teresa M. Amabile (2000): From Guilford to Creative Synergy: Opening the Black Box of Team-Level Creativity. *Creativity Research Journal*, vol. 13 no. 3, pp. 285–294.

Lamb, Brian (2004): Wide Open Spaces: Wikis, Ready or Not. *Educause Review*, vol. 39 no. 5, pp. 36–48.

Mariani, John A. and Tom Rodden (1996): Exposing the Collaborative Nature of Data: Working Together in the PIT. 7th International Workshop on Database and Expert Systems Applications (DEXA), 1996. IEEE, pp. 445–450.

Olson, Gary M., Daniel E. Atkins, Robert Clauer, Thomas A. Finholt, Farnam Jahanian, Timothy L. Killeen, Atul Prakash and Terry Weymouth (1998): The Upper Atmospheric Research Collaboratory. *ACM Interactions*, vol. 5 no. 4, pp. 48–55.

Pagonis, William G. and Jeffrey L. Cruikshank (1992): *Moving Mountains: Lessons in Leadership and Logistics from the Gulf War*. Boston: Harvard Business School Press.

Patton, Michael Quinn (2002): *Qualitative Research and Evaluation Methods*. Thousand Oaks, CA: Sage.

Pelled, Lisa Hope (1999): Exploring the Black Box: An Analysis of Work Group Diversity, Conflict, and Performance. *Administrative Science Quarterly*, vol. 44, pp. 1–28.

Pipek, Volkmar, Joachim Hinrichs and Volker Wulf (2003): Sharing Expertise: Challenges for Technical Support. In M.S. Ackerman, V. Pipek and V. Wulf (eds.): *Sharing Expertise: Beyond Knowledge Management*. Cambridge, MA: MIT Press, pp. 111–136.

Reddy, Madhu C., Dourish Paul and Pratt Wanda (2001): Coordinating Heterogeneous Work: Information and Representation in Medical Care. *Proceedings of the European Conference on Computer-Supported Cooperative Work ECSCW, Bonn, Germany*, pp. 239–258.

Schmidt, Kjeld (1991): Riding a Tiger, or Computer Supported Cooperative Work. In L. Bannon, M. Robinson and K. Schmidt (eds): Second European Conference on Computer Supported Cooperative Work, Amsterdam, September 25–27, 1991. Springer, pp. 1–16.

Schmidt, Kjeld (2000): The Critical Role of Workplace Studies in CSCW. In P. Luff, J. Hindmarsh and C. Heath (eds.): *Workplace Studies: Recovering Work Practice and Informing Design*. Cambridge, England: Cambridge University Press, pp. 141–149.

Tepper, Michele (2003): The Rise of Social Software. *NetWorker*, vol. 7 no. 3, pp. 18–23.

Whittaker, Steve and Heinrich Schwarz (1999): Meetings of the Board: The Impact of Scheduling Medium on Long Term Group Coordination in Software Development. *Computer Supported Cooperative Work*, vol. 8, pp. 175–205.

Whittaker, Steve and Candace Sidner (1996): Email Overload: Exploring Personal Information Management of Email. *Proceedings of CHI 1996, Vancouver, BC, Canada, April 13–18, 1996*. ACM Press, pp. 276–283.

Yanow, Dvora (2001): Translating Local Knowledge at Organizational Peripheries. Interdisciplinary Committee on Organizational Studies, University of Michigan, Ann Arbor, MI, March 30, 2001.