

# A socio-technical framework for cyberinfrastructure design

Ann Zimmerman

University of Michigan, School of Information

Ann Arbor, Michigan 48109-2112, USA

asz@umich.edu

**Abstract.** As the concept of cyberinfrastructure (CI) matures people question whether it is something that can be designed or built. These doubts may stem from considering design as a time-limited process or the lack of a framework in which to view, integrate, and compare knowledge relevant to e-Science applications. This paper addresses these issues by identifying the components of a preliminary framework for CI design where *design* is viewed as an ongoing and long-term process. The framework is necessarily a socio-technical one because CIs interact with people, technologies, and other infrastructures. While there has been a commendable focus on science drivers and user needs in many CI efforts, more than this is needed if CIs are to adequately balance the needs of all users and other stakeholders. I argue that a framework for CI design should consider the following elements: maturity and complexity of the CI; users' capabilities, expectations, and perceived need for CI; research problems and their relation to computing; hardware, software, and associated technologies produced by others; the role and influence of all stakeholders; and mechanisms for coordination and feedback. This paper has two purposes: to introduce the framework and to communicate a new way to consider CI design.

## Introduction

Cyberinfrastructure consists of the advanced information technologies that facilitate access to distributed resources such as computational tools and services, instruments, data, and people, and e-Science is about new ways of doing science that draws on the capabilities that CI provides (Lynch, 2006). Both are popular topics due to the focus on CI in the United States and on e-Science in the United Kingdom and elsewhere. A specific aim of CI-enabled science is to transform scientific practice in order to address larger and more complex problems.

Social scientists are now studying CI design, deployment, and use as a means to better understand, for example, modern scientific organization and practice or the way that social factors shape technology (e.g., Hine, 2006; Lee, Dourish, & Mark, 2006; Ribes & Bowker, in press). It is also increasingly common for social scientists to be involved as participants in initiatives to help develop or evaluate CI. This is particularly the case as those who are leading or funding CI projects realize that technology adoption depends on a wide range of factors, some of which concern aspects of the technology, but others that are social, cultural, or organizational in nature. This participation provides social scientists with access to rich sites of study and is yielding valuable insights and findings that have both applied and

theoretical implications (e.g., Lawrence, 2006; Ribes & Baker, 2007; Ribes & Finholt, 2007; Zimmerman & Nardi, 2006; Zimmerman & Finholt, 2007).

In spite of this growing base of knowledge, social scientists and others with an interest in CI have yet to develop a coherent framework that identifies the key elements for cyberinfrastructure design or a robust set of methods and tools to guide that design. In fact, people increasingly question whether CI can be designed or built (Freeman, 2007). These doubts may stem from a view of design as a time-limited process or from the lack of a conceptual container in which to place, relate, and integrate the growing body of new knowledge as well as findings from prior studies that are relevant to CI design. This paper addresses these issues. Specifically, I introduce the major components of a preliminary CI design framework where *design* is viewed as an ongoing and long-term process. The framework is necessarily a socio-technical one because CIs interact with people, technologies, and other infrastructures, including other CIs.

The framework was developed based on observations gained from an ongoing investigation into the growth and evolution of an existing large-scale, computing infrastructure and from results based on studies of other CI projects. I argue that a framework for CI design should consider the following elements:

- the maturity and complexity of the cyberinfrastructure
- users' capabilities, expectations, and perceived need for CI
- research problems and their relation to computing
- the hardware, software, and associated technologies used by the CI but produced and maintained by others
- the role and influence of all key stakeholders
- mechanisms for coordination and feedback

I introduce the elements here, but the framework is necessarily preliminary, and I hope others will critique and expand upon it. In addition, due to space limitations, I do not address the framework's implications; these are the subject of future work.

The remainder of the paper is organized as follows. I begin with a brief discussion of CI design. Second, I introduce TeraGrid, a centerpiece of U.S. cyberinfrastructure, and two studies of TeraGrid that form the primary basis for the observations reported in this paper. Third, I discuss the components of the CI design framework. I focus on the first three parts of the framework for several reasons. First, I have studied these in the most depth. Second, space limitations make it difficult to fully discuss each component. Finally, and most importantly, an additional goal of this paper is to communicate a new way to consider CI design, and this intent can be fulfilled through a partial examination of the framework elements.

## Cyberinfrastructure design

Peter Freeman (2007) noted that "'design' implies having a set of specifications to which the resulting design responds" and that you can describe or define what you are designing (n.p.). Like many new ideas, the notion of cyberinfrastructure has and will continue to evolve. Most conceptualizations of CI represent it as a middle layer consisting of the systems, services, and

technologies shown in the highlighted portion of Figure 1 (Atkins et al., 2003). Additionally, hardware, software, personnel, and organizations are present in each of these parts.

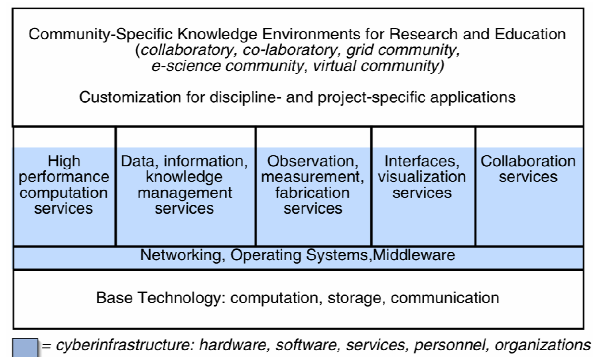


Figure 1: Illustration of Cyberinfrastructure and Surrounding Layers.

Below the CI layer are base technologies and above it are applications and collaboration environments tailored to specific user communities. This figure has been useful in helping to communicate the once new concept of CI, and it remains valuable. For example, funding agencies need a definition to help them plan programs that meet specific goals (Freeman, 2007). There are three limitations, however, to this conceptualization in terms of design. First, the distinctions between CI and other layers are hindrances for CI design because the "'design' process must constructively involve the technical, policy, legal, and usage communities..." (Freeman, 2007, n.p.); some of these communities, as well as the communities' interactions with each other, are missing from this figure. Second, there are multiple, overlapping CIs that must be accounted for in design, and cyberinfrastructures are embedded within and dependent on other infrastructures (Jackson, Edwards, Bowker, & Knobel, 2007). This is shown to some degree in Figure 1, but it is not explicit. Finally, users interact directly or indirectly with all layers and may not appreciate or care where one starts and another one stops. Thus, designing CI to meet the specifications presented in Figure 1 would miss some key features and many important interactions.

This paper argues that we must cast a wider net and consider all aspects of the CI landscape and the diverse stakeholders that are part of the scene in order to ensure that future tools, technologies, and policies adequately meet the needs of all users and other stakeholders and help to evolve and sustain the CI itself. Observations from a study of an existing large-scale computing infrastructure and from investigations of other CI projects are used to support this argument and to help identify the broader set of elements that must be considered.

## Learning from TeraGrid

TeraGrid is a collaborative, NSF-funded project that provides integrated access to computational resources, primarily in the form of supercomputers, large amounts of storage space, visualization services, fast networks, and software.<sup>1</sup> TeraGrid also provides support in the use of its resources and services. While some of the nine resource providers, including the National Center for Supercomputing Applications and the San Diego Supercomputer Center have been in existence for more than 20 years, TeraGrid is fairly new and has only been in operation since late 2004. The motivation in bringing these individual sites together was to

<sup>1</sup> See <http://www.teragrid.org>.

create a grid-based infrastructure to enable scientific discoveries that would otherwise not be possible.

TeraGrid is an excellent umbrella for studying the larger space of CI because, among other reasons, it is a key part of the U.S. cyberinfrastructure that serves diverse users, that must consider, interact, or partner with other infrastructures, that develops software and incorporates software produced by others, and that is working toward scaling up both in terms of the number and types of users. These points are best illustrated through a description of TeraGrid's three main objectives. In short, these objectives are to *enable deep science*, *have wide impact*, and *provide open infrastructure*. The first is about exploiting the cutting edge of TeraGrid systems and services to address the most advanced computational science problems in various domains. The second objective relates to bringing TeraGrid resources and services to communities that have not traditionally used them. This is primarily being accomplished through the TeraGrid Science Gateways program. Early in its history, TeraGrid recognized that many disciplinary communities were building elements of their own cyberinfrastructure, and so TeraGrid set out to form partnerships that would provide TeraGrid resources and services to user communities through tools and environments they were already using (Catlett, Beckman, Skow, & Foster, 2006; Zimmerman & Finholt, 2007). In this strategy, TeraGrid's role is as a back end service provider with the gateways serving as the front end to the user. Finally, the open aspect refers to the fact that TeraGrid architecture is service-oriented and emphasizes open standards. The word open also reflects TeraGrid's willingness to explore partnerships with other service and resource providers, including campuses, peer grids, etc.

## Methods

I have been studying TeraGrid since 2006 as part of two NSF-funded projects. The first is an evaluation research study that is investigating the needs of current and emerging users of TeraGrid, the impact of TeraGrid on scientific practices and outcomes, and the partnership among the nine sites. To date, I, along with colleagues in the University of Michigan's School of Information (UM-SI), have conducted a large survey of individual TeraGrid users, interviewed more than 85 individuals, including TeraGrid personnel, Science Gateway developers, individual users, and CI experts, observed and participated in internal TeraGrid meetings, and read many documents written by and about TeraGrid and about high-performance computing (HPC). The second activity is a project to plan for the future of TeraGrid after 2010 when the current resource provider grants expire.<sup>2</sup> The planning process is being led by a steering committee drawn from the user and other key stakeholder communities and facilitated by myself and colleagues at the UM-SI.

My participation in two other projects has also informed the observations reported here. First, from late 2003 through 2005, I was a member of the Science of Collaboratories project (SOC) (Olson, Zimmerman, & Bos, in press). Among the many distributed scientific collaborations that SOC studied were collaborations between computer scientists and domain scientists to develop CI for particular user communities. Second, I participated as a member of the cyberinfrastructure committee for the National Ecological Observatory Network (NEON) Design Consortium (Zimmerman & Nardi, 2006). I have also drawn on the work of others studying CI (e.g., Lawrence, 2006; Lee, Dourish, & Mark, 2006; Ribes & Finholt, 2007).

---

<sup>2</sup> See <http://teragridfuture.org>.

# A socio-technical framework for cyberinfrastructure design

In the sections that follow, I introduce each element of the CI design framework. As noted in the Introduction, I focus on the first three elements of the framework. Where appropriate, I also identify the endpoints of an axis associated with components or sub-components and illustrate key points along the axis using examples drawn from TeraGrid and studies of other CI projects. The ends of each axis are not intended as measures of advantage, success, or any other form of evaluation; they are simply illustrative of the range of cases that could exist and will be useful in future work to consider the implications of the various elements and sub-elements. In addition, the ways in which framework components interact with each other are discussed only minimally here, but they are important and deserve greater attention in future analyses.

## Characteristics of the cyberinfrastructure: Maturity and complexity

There are potentially many characteristics of cyberinfrastructures that are relevant to design. Here, I identify two aspects: maturity and complexity. While it may seem paradoxical to consider CI characteristics in CI design, this section shows that they are important.

### Maturity

The maturity of the CI relates to its stage of development. At the upper end of the maturity axis are full production cyberinfrastructures. TeraGrid is an example of a CI at this end of the scale, although like many CIs, TeraGrid is both a development and a research project. Lawrence (2006) has discussed the tensions this dual purpose can cause, especially in the early stages of a CI. Since CIs are inherently about technology, this tension is likely to be part of their nature. At the other end of the maturity scale are cyberinfrastructures that exist only as lists of requirements and/or diagrams in a design document. Although some of the technologies to be used in the CI may exist, there is yet no technological instantiation of the CI as a whole, and there are no active users. An example of an immature CI is the proposed CI for NEON. NEON is one of several U.S. environmental observing systems being planned to help predict and solve environmental grand challenges. If NEON is funded, cyberinfrastructure will be an integral part of it—connecting people to each other, to remote instrument and sensor networks, to data, and to software. At this point, NEON CI is present only in a report.<sup>3</sup> Finally, in the middle of the axis are prototype systems, which to varying degrees embody the essential functions and form of the production CI, but which often have few, if any, actual users. GEON is an example of a prototype CI (Ribes & Bowker, in press).

### Complexity

A second important characteristic of a CI is its complexity. This is an evolving concept that I currently define as a combination of 1) the number and intricacies of related parts (internal and external) of which the CI is comprised, on which it relies, and to which it must respond, and 2) the scope of the CI in terms of the diversity of the users' characteristics (see below). The first aspect of complexity is affected by other elements of the framework including the hardware, software, and associated technologies used by the CI but produced and maintained by others, the stakeholders other than users, and the need for coordination and feedback. It is

---

<sup>3</sup> Networking and Informatics Baseline Design for the National Ecological Observatory Network [http://www.neoninc.org/documents/NIBD\\_2006Oct23.pdf](http://www.neoninc.org/documents/NIBD_2006Oct23.pdf)

important to note that complexity cannot be viewed outwardly. For instance, something that looks complicated is not necessarily complex and a well-designed interface can hide complexity. In addition, to describe a CI as complex is not to say that the complexity is well managed.

By all measures, TeraGrid is an example of a CI located at the high end of the complexity axis. For instance, TeraGrid coordinates the development, implementation, and support of common software components across its nine sites; it makes available a wide range of applications developed by others; and it requires a large degree of coordination both internally and externally. Further, TeraGrid's complexity is increasing, and this is likely to continue. One example of this is the tension between TeraGrid's *deep* and *wide* goals where the existing users, hardware and software, and policies and procedures may conflict with the needs of new communities of users (Zimmerman & Finholt, 2007). There is much more work to do in analyzing the issue of CI complexity, although by nature, CIs are likely to be complex, so this is a sub-component that may not exhibit wide variation at the extremes.

## Characteristics of the users: Capabilities, expectations, and perceived need

As with the CI itself, there are potentially many characteristics of users to be considered in CI design. In this section I discuss three inter-related topics: the users' capabilities, expectations, and perceived need for the CI. Together, these may be viewed as aspects of user needs. I give an example of each characteristic from TeraGrid.

### Capabilities

Capabilities refer to the skills and knowledge of the users in relation to those that are required in order to use the CI. As noted above, TeraGrid is a complex system and *use* of TeraGrid requires multiple steps and specialized knowledge. For example, users must write a successful proposal to receive an allocation on a TeraGrid resource or work with someone who has an allocation. Further, users must often select the most appropriate hardware on which to run their application, know what software and applications reside on particular machines, understand how to run jobs and retrieve data, and be familiar with disk storage policies or risk losing valuable data. Use of TeraGrid can also be viewed from a usage mode perspective. For example, some users run batch jobs on individual resources, others employ multiple distributed resources at once, and others access the resources through a Science Gateway and may not realize they are interacting with TeraGrid. Each usage mode requires varying capabilities on the part of the user. For instance, because there is significant heterogeneity in policies and procedures across the nine sites, the capabilities a user needs to run at more than one place can be quite high.

The multiple aspects of use make it hard to classify users by capability along only one axis. For instance, some users may be experienced in writing successful proposals to gain allocations on the TeraGrid machines while other users may be highly competent in making codes run more efficiently. Further, users may be able to capitalize on what others have learned before them, allowing them to bypass some of the capabilities they might otherwise need. For example, several interviewees, particularly graduate students, told me that they did not have to select the machine on which to run their jobs since there was a history in their project of using a particular resource. This has advantages and limitations. It reduces the startup time of new users, but it could also hinder the discovery of more effective ways to conduct work.

## Expectations

Expectations are the users' sometimes known, but often hidden, beliefs about factors that inhibit or encourage their use of new technologies. This characteristic does not conform well to one axis, but it is important to expose users' expectations as a means to help predict the likelihood that the CI will match one part of users' needs. A common example from our interviews of different expectations in relation to TeraGrid concerns the amount of time that experimentalists and computational scientists expect and, therefore, are willing to wait for the results of a submitted job to be returned. In general, experimentalists want a quick turnaround; if it is longer than a couple of hours or a day, the results may not be relevant when they are finally returned. As one interviewee stated regarding the goals of experimentalists: "It's generating ideas, generating 'what if' questions, getting answers quick, and narrowing down the search space so you don't have to do it in the lab." Computational scientists, however, as another interviewee noted, "are used to waiting two weeks to see their job run because there's no other way for them to do the science." The difference in expectations relates partly to the purpose for the computation. Experimentalists may be forming a hypothesis and want to quickly check an idea before going further, whereas for computational scientists, the results of a simulation are the basis of their science.

## Perceived need

Perceived need is defined as the degree to which users believe the CI will be helpful to achieve their goals. This is similar to the idea of *performance expectancy* (Venkatesh et al., 2003). Relative to this characteristic, at one end of the spectrum are scientists and engineers to whom HPC resources and services are essential—they could not do their work without them. The other end of the axis consists of people who have not yet begun to conceive a need for a particular CI. It is difficult to deal with this end of the scale when designing CI.

## Research problems and their relation to computing

This component is one of the most difficult to define. Although I describe it only preliminarily in this paper, I believe it is an important element in terms of taking the results of user needs gathered through planning processes and other methods and translating them into technical CI specifications such as requirements for machine architecture, network infrastructure, and data transfer and storage capabilities. I have identified two inter-related aspects to the relationships between research problems and computing, although there are likely to be others. The first concerns the ease of applying computational methods to the science. The second concerns how this first aspect relates to the technical specifications mentioned above.

The ease of applying computation to topics in a particular domain or sub-domain appears to be a mixture of the nature of the problems to be solved (e.g. linear or non-linear, scale addressed, etc.) the cost of employing the computational solution, the history of the application of computing in the domain or sub-domain, including the availability of trained personnel, and the existence and characteristics of the algorithms to solve user problems. How these factors relate to technical specifications is complex, but below I provide a simple example in relation to the first topic—the nature of the research problems to be solved.

A supercomputer is composed of processors, memory, I/O system, and an interconnect. Different configurations of these components are more or less effective for particular types of problems (National Research Council, 2005). For example, scientists who are trying to compare very large datasets prefer machines with large memory and high I/O. Other

scientists may not start with existing data but with a fundamental theory, and problems in this arena may rely more heavily on compute power. These comparisons were given by users who attended a recent TeraGrid planning workshop. An interview with a Science Gateway developer provides a somewhat different perspective on this issue.

The new communities that would benefit from high-performance computing find a real impedance mismatch between the way the machines are fielded, the way the staff are there to help them understand how you would use this in your science project, and also the way the machines are actually architected—the amount of memory, the portability of codes to that particular architecture.

Even when this information is known, using it to design—as TeraGrid must do—a complement of resources to serve an ever diverse set of user needs is an extreme challenge. There are numerous reasons for this, but one is certainly a lack of effective methods to collect, organize, categorize, and track user requirements and to synchronize it with hardware acquisitions and software development. Another problem is that technology changes quickly and hardware has a limited life-span.

## Other framework elements

In addition to the three elements discussed above, I have identified three other components of a CI design framework. Each of these is introduced briefly below.

### Hardware, software, and associated technologies used by the CI but produced and maintained by others

All cyberinfrastructures are dependent upon hardware, software, and associated technologies that are produced and maintained by others. These "others" may be commercial providers, nonprofit developers, or individual programmers. Thus, interfacing with the hardware, software, and other technologies also involves interactions with people and organizations.

### Role and influence of other stakeholders

So far, the framework has covered the following stakeholders: the individuals that are part of the CI, the users, and the individuals or organizations that provide hardware, software, or other technologies used by the CI. There are other stakeholders. In the case of TeraGrid, some of the other stakeholders include NSF, peer organizations such as other grid projects, other CIs, and current and potential partners, such as campuses. Defining the relative role and influence of various stakeholders in relation to design is an open and challenging problem.

### Mechanisms for coordination and feedback

Continual communication, coordination, interaction, and feedback between the CI and stakeholders is necessary in order to design CI that meets the needs of all users and other stakeholders and that helps to sustain the CI. Robust and scalable approaches to accomplish this are lacking. This should be a high priority area for future research and for CI funding initiatives. The growing body of knowledge on methods to develop and sustain online communities and the ready availability of social networking technologies may offer one means to address these challenges.

## Discussion

The framework presented in this paper is an ideal. Conceptual foundations, methods and tools are needed to help connect the dots between knowledge of scientific practice and the development of policies, systems, or technologies that move science in new directions. Currently, we lack the full complement of methods and theories to achieve a true "science of CI design".<sup>4</sup> For some issues and questions, existing concepts and methods can continue to be employed fruitfully, but in others these approaches need to be adapted or used more expertly and productively. By identifying the components of a design framework, it will be easier to discern where existing methods will suffice and where new approaches are called for. In addition, we need to understand the framework components in greater detail—and from multidisciplinary and interdisciplinary perspectives. Along the way, social scientists and others interested in CI design may find that they require more advanced technologies to address interesting new problems.

## Acknowledgments

Several colleagues at the University of Michigan's School of Information (UM-SI) deserve recognition for their part in the studies that form the basis for this paper, although I take sole responsibility for its content. Thomas Finholt is the Principal Investigator (PI) of the TeraGrid evaluation study and is Co-PI on the facilitation grant for the planning process. Magia Krause, UM-SI doctoral student, collaborated on the TeraGrid user survey. Katherine Lawrence co-organized and led the summer 2007 TeraGrid planning workshops. I also thank TeraGrid, especially Charlie Catlett, former Director of the TeraGrid Grid Infrastructure Group, for facilitating this research. I acknowledge the many individuals who contributed to this research through the survey and interviews and without whom the work reported here would not have been possible. I appreciate comments from three anonymous reviewers, which helped to focus and improve the manuscript. This paper is based upon work supported by the National Science Foundation under grants OCI-0602525 and OCI-0724300 to the University of Michigan. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

## References

- Atkins, D., Droegemeier, K., Feldman, S., Garcia-Molina, H., Messerschmitt, D., Messina, P., Ostriker, J. and Wright, H. (2003): *Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Panel on Cyberinfrastructure*, National Science Foundation, Washington, DC.
- Catlett, C., Beckman, P., Skow, D., and Foster, I. (2006): Creating and operating national-scale cyberinfrastructure services. *CTWatch Quarterly*, vol. 2, no. 2, pp. 2-10.
- Freeman, P. A. (2007): 'Is 'designing' cyberinfrastructure – or, even, 'defining' it – possible?' *First Monday*, vol. 12, no. 6 (June 2007). Retrieved September 7, 2007, from [http://www.firstmonday.org/issues/issue12\\_6/freeman/index.html](http://www.firstmonday.org/issues/issue12_6/freeman/index.html)

---

<sup>4</sup> I borrow this phrasing from the concluding chapter in Olson, Zimmerman, and Bos (in press), who ask whether there is a "science of collaboratories?" Credit for the phrase belongs to Bos.

- Hine, C. (ed.) (2006): *New Infrastructures for Knowledge Production: Understanding E-science*, Information Science Publishing, Hershey, PA.
- Jackson, S. J., Edwards, P. N., Bowker, G. C. and Knobel, C. P. (2007): 'Understanding infrastructure: History, heuristics, and cyberinfrastructure policy', *First Monday*, vol. 12, no. 6. Retrieved September 7, 2007 from [http://www.firstmonday.org/issues/issue12\\_6/](http://www.firstmonday.org/issues/issue12_6/)
- Lawrence, K. A. (2006): 'Walking the tightrope: The balancing acts of a large e-Research project', *Computer Supported Cooperative Work*, vol. 15, pp. 385-411.
- Lee, C., Dourish, P. and Mark, G. (2006): 'The Human Infrastructure of Cyberinfrastructure', In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '06)*, ACM Press, New York, NY, pp. 483-492.
- Lynch, C. (2006): 'Research libraries engage the digital world: A US-UK comparative examination of recent history and future prospects', *Ariadne*, vol. 46. Retrieved September 7, 2007 from <http://www.ariadne.ac.uk/issue46/lynch/>
- National Research Council. (2005): *Getting up to Speed: The Future of Supercomputing*, Washington, DC, National Academies Press.
- Olson, G. M., Zimmerman, A. and Bos, N. (eds.). (in press): *Science on the Internet*, Cambridge, MA, MIT Press.
- Ribes, D. and Baker, K. S. (2007): 'Modes of Social Science Engagement in Community Infrastructure Design', in Steinfield, C., Pentland, B., Ackerman, M. and Contractor, N. (eds.): *Communities and Technologies 2007*, Springer, Dordrecht, pp. 107-130.
- Ribes, D. and Bowker, G. C. (in press). 'Organizing for Multidisciplinary Collaboration: The Case of GEON', in Olson, G. M., Zimmerman, A. and Bos, N. (eds.): *Science on the Internet*, Cambridge, MA, MIT Press.
- Ribes, D. and Finholt, T. A. (2007): 'Tensions Across the Scales: Planning Infrastructure for the Long-Term', in *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work (GROUP '07)*, ACM Press, New York, NY.
- Venkatesh, V., Morris, M. G., Davis, G. B. and Davis, F. D. (2003): 'User acceptance of information technology: Toward a unified view', *MIS Quarterly*, vol. 27, no. 3, pp. 425-478.
- Zimmerman, A. and Finholt, T. A. (2007): 'Growing an Infrastructure: The Role of Gateway Organizations in Cultivating New Communities of Users,' in *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work (GROUP '07)*, ACM Press, New York, NY.
- Zimmerman, A. and Nardi, B. A. (2006): 'Whither or Whether HCI: Requirements Analysis for Multi-Sited, Multi-User Cyberinfrastructures', In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, ACM Press, New York, NY, pp. 1601-1606.