A Sociotechnical Exploration of Infrastructural Middleware Development

Charlotte P. Lee[1], Matthew J. Bietz[2], Katie Derthick[1], Drew Paine[1]
[1]University of Washington, 423 Sieg Hall, Seattle, WA 98195, USA
[2]University of California, Irvine, Donald Bren Hall 5042, Irvine, CA 92697, USA
{cplee, derthick, pained}@uw.edu, mbietz@uci.edu

ABSTRACT

While previous CSCW research has noted that computer scientists have their own research interests pertaining to cyberinfrastructure development projects, most have focused on the research imperatives of scientists. This qualitative, interview-based study investigates the perspective of computer scientists developing middleware software for cyberinfrastructures at two supercomputing centers. This paper examines how technologists develop and sustain middleware applications over time by leveraging expertise and partnering with different research domains in order to achieve long-term infrastructural goals.

Author Keywords
Cyberinfrastructure, sustaining development, software, infrastructure, qualitative methods

ACM Classification Keywords
H.5.3 Computer-supported cooperative work

General Terms
Human Factors, Design

INTRODUCTION

As part of ongoing explorations in CSCW investigating how complex, large scientific cyberinfrastructures are created and sustained over time, this qualitative, interview-based study of computer scientists at two supercomputing centers investigates *how middleware developers sustain long-term research agendas over time*. Scientific cyberinfrastructure development necessarily requires collaboration between domain scientists and computer scientists (who we here crudely group with developers and engineers). A common finding is that the collaboration between scientists and computer scientists (also referred to as technologists or developers) is contentious or difficult [8,11]. However, more problematic is the trope to also paint the already over simplified domain vs. computer scientist division as embodying a split in long vs. short-term concerns. Previous CSCW research has explored the sometimes difficult negotiations of cyberinfrastructure collaborators that work according to different temporal scales [9], for example the fast pace and shorter temporal frames of developers vs. the slower pace and longer temporal frames of information managers [5]. We find that similar to information managers and domain scientists, middleware developers are also constrained by the short time frames provided by government project funding which inevitably are not long enough for more than initial infrastructure development.

Cyberinfrastructures (CI) and eScience refer to the computational and communication infrastructures that support scientific work. Large-scale, distributed cyberinfrastructures are sociotechnical systems that offer new opportunities to conduct increasingly complex and data-intensive research that will answer some of society's most pressing environmental, health, and energy problems. The field of CSCW has increasingly considered the challenges of these complex systems [4,10] because a great deal of sociotechnical innovation occurs first in the scientific realm, making the study of scientific work a harbinger of new ways of working.

The importance of middleware, software that lies between the network and the application, is well accepted in the world of CI and is frequently considered by its developers to itself constitute infrastructure [3]. For our paper, we look at middleware as software systems that may include some end-user applications, but that primarily functions at a lower, back-end layer. Understanding how to support middleware development will shed light on how to support other large-scale, long-term cyberinfrastructures designed to support innovation, both inside and outside of the scientific milieu.

LITERATURE REVIEW

Infrastructures are inherently relational - they are embedded in networks of connections to other systems and infrastructures [11]. Similarly, rather than being systems cyberinfrastructure can be thought of as "networks or webs that enable locally controlled and maintained systems to interoperate more or less seamlessly" [2]. At the same time, the concept of cyberinfrastructure has expanded to include the notion of virtual organizations. The variety of terms used to describe CI illustrate the difficulty of talking about a complex, multi-faceted phenomena.

In line with the historical accounts provided by Edwards et al. [2], Lee et al. [7] found that cyberinfrastructure-in-action

CSCW'12, February 11–15, 2012, Seattle, Washington, USA.

requires a hybrid of old and new organizational forms. When describing the on-the-ground business of creating and maintaining cyberinfrastructure, the notion of *human infrastructure* continues to be a useful lens; *human infrastructure* is the collaborative partnerships of various researchers, technologists, and others who develop and maintain the software, hardware, tools and applications of CI. *Human infrastructure is a more complex and heterogeneous form of organization that encompasses a multitude of collaborative forms including organizational structures, distributed teams, and personal networks.*

Synergizing is "a broad concept that includes strategic collaborative undertakings in pursuit of greater combined effects than individuals, groups, or organizations could effect on their own" [1] and is one way in which human infrastructure functions to enact CI. Bietz et al. [1] find that much of the work of developing CI involves creating and managing these relationships. The synergizing lens provides a conceptual link between global infrastructure creation and the local day-to-day work of systems development. The notion of synergizing highlights the process by which CIs come to exist and that CI requires the active coordination and collaboration of both the technical and social aspects of creating infrastructure. We use the notions of human infrastructure and synergizing in order to understand middleware as part of larger sociotechnical arrangements.

RESEARCH SITES AND METHODS

Distinct from typical approaches that study particular domain-specific projects, this research focuses on how resources transcend and cross cyberinfrastructure project collaborations. We investigated two of the United States' premier supercomputing organizations, the National Center for Super Computing Applications (NCSA) at the University of Illinois Urbana-Champaign and the San Diego Supercomputer Center (SDSC) at the University of California, San Diego, in order to discover commonalities across these two institutions that might be generalizable to similar computing organizations. These research sites afford a perspective that focuses on software that can support myriad cyberinfrastructures.

Both SDSC and NCSA work with multiple scientific communities coming from various academic disciplines. Each project is funded individually, directly from the NSF or another funding body, with group leaders responsible for writing grants. CI developers at NCSA and SDSC typically are involved with multiple projects at a given point in time. These projects can range in size from two or three people up to the hundreds of people involved in projects like TeraGrid (www.teragrid.org).

We interviewed 20 participants at NCSA in spring 2009 and 12 participants at SDSC in summer 2009, for a total of 32 interviews. Interviews were semi-structured to ensure both consistency across interviews, but to also allow for deeper questioning to elicit new insights. Participants in this study included upper-level management, principal investigators on grants, staff computer scientists and programmers, and domain scientists who work with NCSA or SDSC on CI development. Most of our participants work on software rather than hardware development. Interviews were semi-structured, were scheduled for an hour, and lasted between 15 and 86 minutes (median: 56 min.). Interviews were recorded and transcribed for analysis. Data were analyzed using a grounded approach. We began with a closed coding scheme based on the interview protocol and then iteratively created a codebook including open codes based on new discoveries. Finally a series of memos were generated to further identify and develop thematic relationships.

FINDINGS: SUSTAINING MIDDLEWARE DEVELOPMENT

Cyberinfrastructures tend to be funded from either the base level of hardware and network development or from the top level of funding from the particular domain science or sciences. This has created an interesting tension regarding the creation of middleware in cyberinfrastructure. Identified by some as the cyberinfrastructure layer [3], middleware receives comparatively little direct funding. Given the transformational aspirations of NSF research funding, iterative and somewhat incremental improvements of middleware are not considered to be innovative enough to warrant funding.

Most of the work of the centers is funded from "soft" money, that is, funding for relatively short-term projects, often directed toward a specific application domain. Many of our participants find that a significant amount of their time is spent developing new proposals to stay funded. The gap between the need for long-term cyberinfrastructure and the reality of short-term funding for developers is recognized [2]. The contribution of this research is showing how computer scientists address the challenge of developing middleware and researching middleware development beyond short time horizons.

Our research finds that participants are very strategic in terms of how they: synergistically seek collaborators (individuals, groups, and organizations), including hiring individuals who bring both code and the expertise for building and using that code; use legacy code as a stepping stone to move towards more advanced middleware software; re-scope and rename middleware as funding sources and projects change; and insert into and sometimes customize middleware for different domain-specific cyberinfrastructure efforts. How these different activities contribute to sustainable middleware development is discussed below using representative examples from our data.

Leveraging Sociotechnical Resources

As noted earlier, human infrastructure is a hybrid of organizational forms. These organizational forms are also constantly in flux, changing shape in accordance to demands on the virtual organization [7]. In the soft money environment at our research sites, attempts to slow or control the flux often involved trying to keep development teams together despite projects beginning and ending and despite projects differing in their scope and scale. Having experience and prior investment in one area can lead to continued funding to pursue similar topics. Often one project is an explicit continuation of one that came before. Some projects receive continuing funding or structure the project to be funded in phases.

Just as common, however, are informal threads that connect projects over time. In most cases, new projects are proposed specifically because they allow an individual or group to continue working in the same or a similar vein. One PI was able to leverage his expertise to help get a new grant using the same technology in a different way: *"Because of my track record there... we had the expertise, we had the hardware, we knew how to [set up the systems]...we got another grant"* (Paul, Project Lead). We see here the human infrastructure working to maintain some organizational stability by leveraging the principal investigator's expertise. This representative example shows how expertise is leveraged with and through middleware development. Expertise is used to synergistically control flux in human infrastructure—keeping a team together in the face of the end of project funding.

However, human infrastructure not only works to maintain stability, but also to support innovation, by bringing in new sociotechnical resources. One participant told us about recruiting a specific computer scientist to work at the center. *"I was a very strong supporter of bringing Jerry here… And then I invited him to be part of our team so that we could get access to his [Application X]."* Hiring Jerry was a way to bring both his expertise and his technology into the organization.

Jerry recounted the same story in his own words. He talks about developing middleware for his previous job: *"So [Middleware Y] had a life there as part of the Collaboratory. It was inside our [Application X] that we had built there, so there was about a year left of funding on that project going when I came here, so I as the PI brought some of the funding with me."* When Jerry arrived at his new job, he realized that his new project, while in a very different domain, had some of the same requirements as his earlier project. Even though the state of the art was moving away from the approach used in Middleware Y, the new technologies were *"not ready for prime time."* Middleware Y would do what they needed and was stable, and even though it didn't use the latest technologies, it would set the project up for an easier transition to the new technologies when the time came: they would *"be aligned with where we think things are going next."* In other words, Jerry was hired with the understanding that he would bring useful expertise and usable middleware with him. Furthermore, there was an understanding that his software would be a functional stop-gap to fulfill urgent workflow needs.

In the above example we see planning for technological obsolescence of Jerry's Middleware Y, with the understanding that newer underlying technologies would be necessary to sustain the larger system. At the same time, there is more at work than simple retrieval of knowledge or trading on specialized knowledge. Instead, human infrastructure synergistically sustains middleware development by leveraging sociotechnical networks [6] comprised of people and technology—in other words, how people enlist human resources and convenient, appropriate technologies in order to do more with less.

Advancing Research to Accomplish Infrastructure

From an institutional perspective, cyberinfrastructure can be scoped in different ways; for example, modeling and analytical tools and even computing resources may or may not be considered part of the CI, depending on disciplinary needs [12]. The middleware that our study participants develop, however, is generally considered to be foundational for cyberinfrastructure. For our participants, the goal of cyberinfrastructure is to provide an infrastructure in the traditional sense of a plug-and-play technology (at least from the user's standpoint). For the developers in our study, the infrastructural vision — the loftiest goals of cyberinfrastructure — resides on the continued iterative development of middleware.

As middleware receives comparatively little direct funding, it may at first seem that the push to appeal broadly to different application domains is purely an economic decision, but there is also a larger goal to create software that would be broadly useful. As domain scientists would like other scientists to build on their discoveries, so too would computer scientists: *"So we made a decision early on that we wouldn't focus in a discipline like astronomy or earth sciences or business, but we would try to build an environment, a cyberinfrastructure or infrastructure that could support a breadth of applications"* (Mick, Project Lead). The ultimate goal is to build an infrastructure that is broadly useful, or generalizable, across domains.

Developers and managers find ways to drive middleware development—perpetuating the advancement of their software—by attaching them to new, non-computer science, domain-specific scientific research questions. This entails finding different application domains in which their middleware can be applied.

So I used to work a lot on bioinformatics. I worked on it for two years, and then suddenly it disappeared from the radar of the funding agencies. And then, I couldn't raise funding, so then, I had to move to other domains. But the technology was the same, so I could reuse it. Just having, if I had focused only on bioinformatics, right now I would really struggle, right? But having that flexibility of moving from project to project, and building enough preliminary results *to propose to new things is to our advantage.* (Paul, Project Lead)

The group lead found that he would have to attach his research to different application domains in order to get funded. Even though he was continuing his own research in the same area, working on multiple projects gave him the flexibility to target different grant areas and therefore different funding sources.

The other thing that I do is I go out and I look for new opportunities to fund applications that can drive our development because my model is kind of a unique one. ….To make [Software] viable and exciting to researchers in the application domains, I have collaborations with computer science departments, graduate schools, the library information sciences, bioinformatics groups that are developing new data management analysis and visualization techniques that we feed into this development environment. So there's big research and big technology going on here. We do big development, and we point that at big time applications… I'm looking for applications that can help fund this development and some of this research. (Mick, Project Lead)

As with other scientists, these computer scientists need to find and excite collaborators and funders about their research. A primary motivator for these computer scientists is furthering their own middleware development research. As noted earlier, the middleware layer is sometimes thought of as not just being part of the infrastructure but as being the infrastructure. The goal of advancing middleware, from the perspective of our participants, is the larger goal of creating infrastructure itself.

CONCLUSION

Unlike in other research [5], the developers in our study did not live natively or comfortably within short-term "project time." The difference may be attributed to particular organizational configurations that may or may not support developers who are truly computer *scientists* with a research orientation as opposed to engineers. While our participants' notion of "infrastructure time" was not on the scale of many decades as with ecologists, our research shows computer scientists engaging in elaborate and strenuous synergistic activities in order to sustain their software development research agenda well beyond the lifetime of projects.

The scale and scope of CIs is beyond the traditional core of CSCW and we are still struggling to find concepts that do justice to their complexity while still being comprehensible. The first challenge of being able to support scientific CIs is to understand how they are formed and maintained. Computer scientists that develop CI middleware (considered by some to be the most important technical layer) engage in simple, everyday activities to sustain future research and development including hiring people with particular technical expertise, using readily available and easily supportable appropriate technologies, and adapting software to be useful to a larger user base. Middleware developers strategically work towards expanding the scale of their middleware across domains in order to combat the short-term funding provided by funding agencies. Mobilizing human resources through synergistic activities is yet another way to maintain some semblance of stability in the face of constant flux. If middleware is indeed a key technical layer, cyberinfrastructures will benefit from additional research seeking to understand and support CI middleware development.

ACKNOWLEDGEMENTS

This research supported in part by NSF Awards OCI-0838601, IIS-0712994, and IIS-0954088. Thanks to our participants, Silvia Lindtner, and our anonymous reviewers for their assistance and feedback.

REFERENCES

1. Bietz, M. J., Baumer, E. P. S., & Lee, C. P. Synergizing in Cyberinfrastructure Development. Computer Supported Cooperative Work (CSCW), 19, 3-4 (2010), 245-281.

2. Edwards, P. N., Jackson, S. J., Bowker, G. C., et al. (2007). Understanding Infrastructure: Dynamics, Tensions, and Design (Workshop Report).

3. Hey, T., & Trefethen, A. E. Cyberinfrastructure for e-Science. Science, 308, 5723 (2005), 817-821.

4. Jirotka, M., Procter, R., Rodden, T., et al. Special Issue: Collaboration in e-Research. Computer Supported Cooperative Work (CSCW), 15, 4 (2006), 251-255.

5. Karasti, H., Baker, K. S., & Millerand, F. Infrastructure Time: Long-term Matters in Collaborative Development. Computer Supported Cooperative Work (CSCW), 19, 3-4 (2010), 377-415.

6. Latour, B. Science In Action. Harvard University Press, Cambridge, MA, 1987.

7. Lee, C. P., Dourish, P., & Mark, G. The Human Infrastructure of Cyberinfrastructure. In Proc. CSCW, ACM (2006), 483-492.

8. Olson, J. S., Hofer, E., Bos, N., et al. A theory of remote scientific collaboration. Scientific Collaboration on the internet (2008), 73-97.

9. Ribes, D., & Finholt, T. A. The Long Now of Technology Infrastructure: Articulating Tensions in Development. Journal of the Association for Information Systems, 10, 5 (2009), 375-398.

10. Ribes, D., & Lee, C. Sociotechnical Studies of Cyberinfrastructure and e-Research: Current Themes and Future Trajectories. Computer Supported Cooperative Work (CSCW), 19, 3 (2010), 231-244.

11. Star, S. L., & Ruhleder, K. Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. Information Systems Research, 7, 1 (1996), 24.

12. Wheeler, B. E-Research Is a Fad: Scholarship 2.0, Cyberinfrastructure, and IT Governance. In The Tower and The Cloud: Higher Education in the Age of Cloud Computing. R. N. Katz, Ed. Educause, 2008, 108-117.