ISR Institute for Software Research
University of California, Irvine

Decentralized Virtual Activities and Technologies: A Socio-Technical Approach

NSF Award # 0808783

Walt Scacchi
Univ. of California, Irvine
wscacchi@ics.uci.edu

Bonnie Nardi
Univ. of California, Irvine
nardi@ics.uci.edu

Alfred Kobsa
Univ. of California, Irvine
kobsa@uci.edu

David Redmiles
Univ. of California, Irvine
redmiles@ics.uci.edu

Cristina V. Lopes
Univ. of California, Irvine
lopes@ics.uci.edu

Richard N. Taylor
Univ. of California, Irvine
taylor@ics.uci.edu

Gloria Mark
Univ. of California, Irvine
gmark@uci.edu

December 2008
ISR Technical Report # UCI-ISR-08-4

Institute for Software Research
ICS2 221
University of California, Irvine
Irvine, CA 92697-3455
www.isr.uci.edu

www.isr.uci.edu/tech-reports.html

1. Introduction and motivation

We see growing, widespread interest in the development and use of decentralized systems and virtual world environments as possible new places for engaging in collaborative work activities. The Gartner Group [1] declares that within five years 80% of Internet users and Fortune 500 companies will have an online presence in a virtual world of some kind. Furthermore, they estimate the collaborative and community-related aspects of these environments will dominate in the future. Elsewhere, there is widespread interest in stimulating new technological innovations that enable people to come together through social networking, file/media sharing, and massively multi-player online game play. This new generation of networked computing environments seems headed towards increased socialization, interaction, communication, and collaboration that span multiple organizational boundaries as its primary purpose. But how do we get there from here? Is it sufficient to just let the market of entrepreneurial vendors and technological innovators simply decide who needs what? The history of computing in organizational and work settings reveals that the adoption and integration of new technologies is rarely simply a matter of buying the lowest cost technical alternative. More often such history reveals a legacy of many failed or problematic efforts to develop and deploy computing systems that arise from a lack of understanding or recognition of the ways how people's work and social activities are situated in organizational and technological contexts. These contexts configure, constrain, or enable some types of activities to flourish while others are displaced, either unintentionally or intentionally. Our goal in this proposed study is to empirical investigate and comparatively analyze the development and use of decentralized virtual activities systems (DVAS) within a sample of five diverse organizational settings using a research approach that foregrounds a balanced socio-technical perspective.

Our five research partners serve multiple roles in this project. First, they *serve as a source of real-world problems* for us to tackle. Rather than speculating about the problems of the development and use of DVAS, or working with abstract theoretical models, we will study how DVAS arise in practice in real settings. Since the social and organizational context in which development takes place is a critical aspect of a successful development model, this sort of engagement is essential. Second, as organizations facing the daily problems of distributed development, our partners view these as practical problems to be tackled, and *will engage directly with us in developing strategies and solutions*. Third, they *serve as test-beds* for early evaluation of proposed new solutions. As organizations with real and pressing problems surrounding inter-organizational work, our partners provide ideal settings in which to evaluate research outputs *in situ*, providing more rigorous evaluation than would be possible in laboratory settings or simulations. In other words, by working with our research partners we ensure a continuous engagement with real world settings at all stages of the project.

Beyond this, we believe a large multi-investigator effort is required to conduct such a study. Why? First, a study of DVAS requires a broad interdisciplinary understanding of the problems and a broad and interdisciplinary approach to their solution. We are a unique interdisciplinary team comprised of researchers with years of scholarly and professional expertise that we have applied to investigating socio-technical issues of collaborative systems. We are ideally positioned to study the issues critical to DVAS. We are proposing a large multi-site, multi-partner research endeavor that provides the greatest practical opportunity for generalizable results from comparative analyses of both in-situ field studies and technology prototyping efforts. A small/mid-size team does not have the resources or expertise to realistically study and develop prototypes for DVAS in such varied organizational settings. Only a large team has the potential to realize all of the preceding opportunities, benefits, and constraints; we assert that no amalgamation of individual investigator studies can cover the same amount or diversity of research interests or research partners or provide such comprehensive expertise in ways that can more effectively achieve comparable results. Accordingly, we now turn to describe and explain our proposed research effort, starting first with some background that helps further identify the focus of our study.

2. Background

What is a decentralized virtual activity system? First, an *activity system* is a computer-based environment that encompasses a web of participants, processes, information artifacts, information technologies, products and services, customers/users, organizational setting(s), as well as human, technological, and resource infrastructures [50] [107] that are interrelated in ways that can enable workplace or other activities [48]. An activity system is a contemporary *socio-technical system* whose social and technical elements are interdependent and co-evolving, so that attempting to discount/ignore either the social/technical dimensions results in ineffective systems that are poorly understood and wasteful of scarce resources. An activity system is similar in concept to a work system [6], but we do not see it as limited to only routine business processes or other work activities in an enterprise setting, since other productive activities including building social relationships, playing games, or engaging in exploratory experiences are all within our view of an activity system.

Second, a *virtual activity system* (VAS) is one where some of the elements of an activity system either exist or are accessible only in a computer-mediated form, and may be distributed across multiple sites/computers. Text-based virtual realities (also called Multi-User Dungeons or MUDs [20], virtual worlds [98] [8], networked virtual environments [87], and persistent online games like *World of Warcraft* or *City of Heroes* are all kinds of VAS, as are computer-based models or simulations of enterprises as rendered within an Enterprise Requirements Planning (ERP) system, computer-aided manufacturing systems, multi-agent problem solving systems, and virtual organizations. The vast majority of open source software development projects primarily rely on text-based communications through "informalisms" like threaded email discussion lists, project wikis and instant messages that are supported by ad hoc arrays of personalized software or artifact development tools, to enable decentralized software systems and development processes [82] [83] [85]. In a sense, these OSS developers work in virtual worlds of online text, communication and discourse through informalisms that move in, out, and across both shared and private information repositories, as well as across organizational boundaries. The multi-party telephone conference call (with or without Web-based supplements and presentation materials) are common, widespread, and often ready-to-use technical systems that enable work to be decentralized and globally dispersed, are also considered (early) virtual worlds. On the other hand, the Web is generally not identified as a virtual world, nor is the use of email, nor is an integrated development environment with shared repositories for building software. Yet each of these can be very effective in supporting technical work activities, software system design, or system development processes that can be logically centralized and physically distributed. As such, a VAS is one where work activities occur via creation and update to information artifacts that move in and out of information repositories and the communications among participating contributors, so that these work activities can be physically distributed but logically centralized and coordinated [64]. Thus, VAS allow for development and use of virtual worlds based on text, graphic, electronic media, databases, or some combination thereof. Similarly, VAS allow for workplaces that exist only online and lack a physical workspace or specific setting.

Next, a *decentralized activity system* (DAS) is one where elements of the computer-based environment are not only physically distributed across multi sites, but also where administrative or social control, allocation of resources, and decision-making activities with regard to how the DAS is developed or used are performed locally in an autonomous or discretionary manner (at least to some degree). DAS both seek to engender and accommodate heterogeneous computing systems that can interoperate at arms length via information sharing protocols, processes, and data representations. Both the Internet and World-Wide Web were implemented as DAS, as are other computer-based environments for peer-to-peer applications like file sharing or resource-sharing. Virtual organizations, through communication, data-sharing and coordination technology support, have been able to respond to dynamic changes in the environment by adopting decentralized structures [5] [9] [23] [22] [49]. DAS must self-organize, and DAS developers and users must organized themselves into participatory roles, role-sets and role-migration paths so that the DAS can persist and thrive as participants join and leave the effort. Collective activities that span across the weak social network ties can nonetheless give rise to a global social movement [11] or computerization movement, like the Free Software Movement [30, 32].

Finally, a *decentralized virtual activity system* (DVAS) is an activity system whose elements can be both virtual and decentralized. Here we look for emerging computer-based environments that seek to create or articulate a virtual organization or virtual world that spans multiple interacting organizations with diverse stakeholders (who may be in conflict with one another, yet need to cooperate or coordinate their activities) that seek to collaboratively interoperate as a loosely coupled alliance, rather than as a hierarchically controlled network of administrative authorities or resource controllers. In conjunction with our research partners, described later, we have identified five different settings where DVAS are being developed and used, or where efforts are underway to create and employ a DVAS.

Decentralization is not inevitable, but it has the potential to be very effective in many situations, as do centralized systems and administrative regimes with distributed developers and end-users. Similarly, highly visual virtual worlds may or may not be the workplace of the future for everyone. Subsequently, we seek to understand where, when, how, what kind, and why decentralization works best, and whether or not decentralized virtual worlds are a necessary or useful part.

Overall, it remains unclear as to what principals, models, or guidelines people can employ to help determine whether/how their technical work activities, software system designs, or development processes might best take advantage of the kinds of decentralization that we have identified. Similarly, people need guidance to recognize when such efforts like the move to virtual worlds of one kind or another are a good/bad fit, unnecessarily complicated, or otherwise inappropriate choice that might otherwise be better served through alternative means. Our proposed research investigation seeks to discover and refine such principals, models and guidelines based on both empirical field studies and experimental system prototyping efforts to ground are findings. We similarly seek to demonstrate and iteratively assess how such results can be adopted, assimilated, and refined into work practices in the enterprises of our research partners.

Research objective: Our research objective is to understand the constraints, properties, opportunities, and affordances associated with different DVASs, as well as understanding what practices, guidelines, principles, models, and theories can best shape the development and use of a DVAS in different types of settings. For example, a new startup venture like Unimodal Inc. chooses to have a research team at UCI prototype a new activity system for simulating a next-generation personal rapid transit (PRT) system using a virtual world platform like *Second Life*. Furthermore, they both want to allow or encourage other users in *Second Life* to try out the new simulated PRT system to see what users learn from their interaction with each other when exploring the new system. In so doing, Unimodal and the UCI researchers subsequently seek to link activities within the simulated virtual world to digital objects (PRT passenger pods, PRT network rail switches, embedded sensors, etc.) or events distributed in the physical world of an actual PRT system prototype (soon to be built), as well as how these users and the system's developers make sense of their individual and shared experiences with the physical/virtual elements.

Figure 1: Simulated PRT system modeled in *Second Life* (from UCI Prof. C. Lopes and colleagues) on left, PRT system prototype envisioned in physical location (from Unimodal Inc.) on right.

In our proposed study, we will engage a series of field studies and prototyping experiments with five research partners including Unimodal Inc. that will provide empirical observations, support comparative analyses, and enable formulation of principles, methods and theoretical models that generalize across our study samples so as to inform future developers and users about how best to build and employ new DVASs in socio-technically effective ways.

We now turn to describe some of the related research that informs our proposed research study and variables.

3. Research variables

Our proposed research effort is focused understanding how systems supporting decentralized virtual activities and technologies (DVAS) are developed, used, and adapted in diverse settings. From this, we see an array of variables that span a socio-technical landscape of alternative choices for action and system design. These variables (or sets of related variables) range from (a) how to bring together diverse stakeholders in immersive worlds in ways that enable mutual sense-making and dialectical evaluation of new DVAS applications through alternative *representations that span physical and virtual realities*; (b) *how people who collaborate through DVAS create new policies and practices that engender or mitigate conflicts* in the resulting work practices, which in turn shape and guide their primary work activities; (c) *how people engage the work of creating, sustaining, and adapting relationships in DVAS environments* that most effectively meet their needs; (d) how can geographically dispersed people articulate development processes, work practices, and project communities for DVAS; (e) how people can *manage or control their privacy concerns as well as how such concerns are exposed to others* when interacting through artifacts or avatars; and (f) how to define and provide *technical foundations for security and trust in DVAS*.

Each of these issues can be briefly described in turn as follows.

Representations and realities: A primary attraction of DVAS is the ability to blend real artifacts and people with virtual ones to accomplish, for example, a design task. An instance of this is seen in our background work in the design of a rapid transit system: engineering specifications of the real artifacts are used in a virtual environment to explore design choices. Of course the real artifacts are not actually used; representations of them are. The goal of this aspect of our research is to understand the issues and trade-offs involved in designing representations of non-virtual entities for use in a DVAS, yielding principles for guiding developers in creating new ones.

Research questions that arise from this goal include the following. First, what are appropriate ways to represent people, especially people playing specific roles? For instance, in a decentralized software development project, an individual may play a key role as the configuration manager. Should the representation focus solely on this role, or should characteristics perhaps not directly related to the official role be included? How does the social web of developers in a project affect appropriate choices for representation? Second, when should representations be used? With what fidelity? What kind of flexible boundary can be used between the two? For instance, in the software development project, when is it more effective to manipulate representations of the software artifacts as opposed to the artifacts themselves? Or should all meetings have to take place in a virtual setting? For all participants? Third, when are "representations" appropriate which do not have any natural counterpart in reality? For instance, composite or intangible entities, such as "this software project" could have an explicit representation in a DVAS.

In short, we want to find out to what extent the different stakeholders of a complex infrastructure change can come together in an immersive DVAS world like those enabled in *Second Life*, and whether such common "live" representations can support mutual understanding of the system and its broader consequences.

Conflicting policies and practices: Our goal is to study how a socio-technical approach can enable people who are collaborating in a highly decentralized work structure to create new policies and practices that relate to their new work practices, in order to achieve benefits for the decentralized team and organization.

The challenge for reconciling different practices is that there is no standard approach in a decentralized setting, as there would be in a traditional hierarchical, centralized structure. People must find a way to articulate differences and self-organize to set new standards and conventions of practice. To achieve this, different perspectives must first be made explicit [88] [86] [54]. Focusing on the shared objects used in decentralized work practices is a way to understand the relationship of work routines and practices, and work products. Bergman et al., [12] discovered that these boundary objects can be successfully used as active facilitators to help distributed workers develop shared and congruent perspectives and policies by serving to promote shared representations, transform knowledge, mobilize for action, and legitimize new knowledge. We will investigate how technology can enable perspectives to be linked to such boundary objects in decentralized work forms to assist people in transforming their practices to achieve congruency across distance.

Relationship work: We seek to understand how people create and sustain relationships in decentralized activities in order to design technologies to more effectively support relationship work. We will examine the shape and dynamics of social networks, the means by which people contact and exchange information with one another, formal and informal mechanisms by which relationships are created and sustained, and what is unique about relationships in decentralized activities. The goal of this research is to provide broad understandings that will enable people in DVAS to examine and interpret their own practices surrounding relationship work as well as to inform the design of new technologies that will be undertaken as part of the proposed research. This is consistent with a socio-technical approach in which both practices and technologies are seen to mutually and dynamically affect and transform one another. Recent research suggests that social networks in distributed organizations are held together by small cliques that work across the boundaries of the groups within the networks [109]. We will examine network dynamics to understand how specific relationships, such as those in cliques, influence the function of social networks in DVAS. We are also interested in the whole suite of technologies people use in relationship work [61] and the ways in which new technologies, such as those we develop, will alter such practices. It is also important to remember the possible role of face-to-face communication in DVAS. Although they are primarily distributed, some critical face to face interaction may take place and, in line with our socio-technical approach, we will be alert to understanding its possible role in activities that are primarily conducted in virtual space [62].

Discovering DVAS development processes, work practices, and project community dynamics: We seek to develop empirically grounded models that account for the conditions, circumstances, and events people working in formal organizations or in ad hoc communities when they address whether to: (a) embrace free/open source software (FOSS) or proprietary development techniques as a major mode of system development of DVASs? (b) enable the creation and deployment of decentralized virtual environments that are built from FOSS versus proprietary components, systems, or applications? (c) enable end-user developed DVAS application systems or systems for sharing user-created open content? and (d) enable socio-technical networking of DVAS across enterprise boundaries to form alliances with external partners or competitors? The resulting models will help identify fundamental principles, properties, and theories accounting for DVAS development and use.

Managing privacy and maintaining awareness of impressions: Our research on impression management for improving privacy management will address the following research questions: How can the impression(s) one conveys through DVAS systems be made more visible to oneself? What DVAS mechanisms can empower users to manage their impression(s) appropriately? How is appropriateness of impression(s) interpreted and evaluated? How can DVAS mechanisms be made to integrate seamlessly with user practices to avoid undue burden on users? To what extent do DVAS systems with improved impression visibility and management succeed in improving privacy management?

Security and trust: The absence of a central, trusted authority in DVASs (a consequence of decentralization) implies that they are potentially subject to attack or inappropriate use. Our goal is to understand how to make DVASs useable in a world where malicious behavior is to be expected. In absence of such understanding, the best DVAS technology will never be used in organization-critical applications. Three research questions follow from this goal. (1) What are the risks most critical to DVASs? Scoped more narrowly, what are the risks critical to the DVAS of our research partners? (2) Can previously identified mechanisms for mitigating trust and security risks satisfy the identified needs of DVASs? If not, how not, and why? (3) What new security and trust mechanisms can augment the prior work and be effectively incorporated into DVASs, to satisfy any newly identified needs?

4. Insights from Prior Research

Each of our six research variable sets builds from a legacy of prior research that our project investigators have performed, as well as related research from others. These insights help to further explain why we have selected each of the research variable sets introduced in the previous section. Accordingly, we now review what we already have learned about each of these issues for further investigation in this proposed effort.

Representations and realities: Activities around deploying new technologies always involve a diverse collection of stakeholders to come together in a fairly decentralized manner. One good example is Personal Rapid Transit (PRT). Because of concerns with the environment PRT has recently got considerable attention from the part of local and central governments all over the world. PRT is has the potential to drastically reduce air pollution and redefine public transportation. Although the interest in PRT has been growing, its deployment in real settings involves a convergence between technologists, politicians, service providers, and, especially, it has to have a broad support from the general public. All these stakeholders have their own set of concerns and political power, they are by nature decentralized, and their coordination for a common goal is a major hurdle.

Since PRT is a radical departure from the current transportation infrastructure, including self-driving vehicles that carry people, there is an uphill battle for its adoption. The different stakeholders often express concerns such as safety (fail-safe control of the vehicles in the guideway network), visual pollution from having to build guideways, and whether it's practical at all to embed such systems within cities, and use them on a daily basis.

One of the PIs has been working with a company in Southern California, Unimodal Inc. that is trying to deploy one such PRT system. As part of this collaboration, the PI developed a technically-accurate simulation of their PRT system using the virtual world *Second Life* [59] [53]. The guideway and control of the cars was developed according to the actual specifications given to us by the company's engineers.

We found numerous advantages in doing this simulation in an immersive 3D environment like *Second Life*. Not only were technical design problems uncovered, but several issues pertaining to usability and adoption emerged [53]. The fact that the simulated system is publicly available makes it easy for anyone – politicians, the media, and the general public – to experience it to some extent before it is deployed in the real world. This immersive experience can be done in a decentralized and relatively chaotic manner, at the stakeholders' own paces. But the virtual deployment site is also a common point, a common representation, which the stakeholders can refer to, and experience.

The insight we gained is that publicly accessible virtual worlds can be powerful tools not just for engineering design, but also for supporting the natural decentralization of stakeholders in complex techno-social situations.

Conflicting policies and practices: The interoperability of systems to support collaboration requires moving beyond purely technical issues; it also concerns the means and practices that users adopt to carry out their cooperative activities [86]. Decentralization compounds the interoperability problem as it introduces a new kind of interaction order. People who work in decentralized configurations have developed unique practices in their local work settings, and must follow locally established governmental or organizational policies. At the same time they interact in decentralized work structures, networked electronically across distance. This duality of settings creates a challenge for decentralization due to the different local practices and organizational policies that conflict [47]. Prior research has shown that people whose home base is distributed from their teammates, have different and even conflicting work practices influenced by organizational policies [66] [57], different reference frames of the technology (i.e. assumptions and expectations of use and purpose) [67] [75] [108], conflicting conventions for handling shared objects and groupware [37] [54] [89], different resources [89], and even use of different models and methodologies applied to the same problem though collaborating partners are of the same discipline and nominally of the same organization [56]. These differences are difficult to overcome and impact successful collaboration [65]. To our knowledge, no one has investigated how different perspectives, practices, and policies interact and impact decentralized work.

Relationship work: Decentralized activities distribute power throughout a social network. They are dynamic, flexible, and responsive to changing conditions. A key means by which people collaborate is through creating and sustaining human relationships [60]. Such "relationship work" in the contemporary context requires both technical capacity and social knowledge about the multiple stakeholders in the activity. Possible challenges for relationship work in decentralized activities are distributed relationships, rapid change, opacity of networks, cultural differences, varying motivations for participating in the shared activity, and time needed to build relationships.

Processes and coordination: Free/open source software development (FOSSD) is a widely practiced decentralized approach to building and sustaining large distributed software systems [11] [85]. It focuses on practicing the open access, examination, modification/creation, redistribution, and replication of shared, decentralized knowledge artifacts, development processes, and related socio-technical work practices. The artifacts include online chat transcripts, annotated source code, bug reports, etc. that act as boundary objects that can span multiple Web-based FOSSD projects [40]. The development processes include decentralized development or modification of a system's source code modules, and the collective composition, configuration, building, and testing of these modules into candidate or formal software system releases. The practices include self-organizing and continuously emerging FOSSD contributor roles and role migration paths [41], virtual project management (affecting FOSSD activity allocation, performance, oversight, and coordination) without formal manager roles or administrative authority [83], and others. Further, the OSSD project community relies on the decentralization of knowledge of the requirements and design of FOSS in order to insure commitment and socio-technical advancement with a project community [41] [85].

Privacy and awareness: When collaborative work is geographically and temporally distributed, collaborators find it challenging to be aware of each other's activities, routines, tasks, and availability. Yet, such awareness is crucial for increased efficiency and effectiveness of collaborative work [25] [39] [76] [63]. Collaborative software therefore increasingly provides means to disseminate awareness information to facilitate collaboration, such as in instant messaging systems, word processors [16], calendars [68], and programming environments [19]. Infrastructures that seamlessly and automatically capture, store, process and disseminate awareness have been implemented in a variety of domains such as workplaces [72], hospitals [10] [13], and conference centers [24].

However, collaborative needs for awareness are often at odds with individuals' desires for privacy. Prior studies [39] [69] [70] [71] [74] indicate that the inability to achieve a balance between awareness and privacy can lead to underuse of collaborative technology. Also, inadequate attention to privacy aspects may evoke strong user backlash, as was illustrated recently when the popular social networking site Facebook introduced new privacy-invasive awareness features. In such a case, organizations stand to lose their investment in collaborative technology, and face the prospect of longer-term damage due to the undermining of trust and credibility. It is therefore critical to consider privacy aspects when designing awareness mechanisms in collaborative systems. Patil & Kobsa [70] also found evidence that the motivation behind privacy concerns in collaborative settings is impression management, i.e. the desire of an individual to convey an impression of oneself appropriate for the context at hand.

Security and trust: Understanding the nature of decentralized applications and virtual worlds is critical before solutions can be designed for them. Our previous work [91] in building secure decentralized applications has revealed that developing appropriate threat and risk models is the most critical step in this process; it is only when threats can be characterized and are well-understood that effective countermeasures can be appropriately designed and deployed. Decentralization induces risk regarding the perceived multiple points of attack on information access or integrity from potentially insecure system interfaces. Any attempt to support coordination in an open decentralized system – one in which the set of participants may change and no central authority prevails to guarantee that all participants are non-malicious– must address, from the outset, the risks presented. The alternative is to create systems that will inevitably fall prey to malicious behavior and hence become insecure and vulnerable to exploits, surreptitious or remote take-over, or denial of service attacks.

Our previous work has revealed basic principles and effective techniques for designing specific types of decentralized systems that can (1) support security (information access) constraints at the architectural level [79], and (2) incorporate basic secure design and reputation-based trust models [3] [104] [77] [43] [21] in the architectures of participating members [93]. Since the participant architecture must protect itself against threats, appropriate countermeasures need to be incorporated within the architecture. Our experience with developing the PACE architectural style [94] for decentralized trust management gives us a good handle on how to identify countermeasures for threats.

Domain constraints can influence the choice of components in the participant architecture. For example, if the application requires exchange of rapidly changing critical data every few seconds, there may be a need for "estimator" components within the participant architecture to predict new data in case of communication delays [45]. Similarly, our previous work on software architecture-based decentralized trust management identified the need for a special component that encapsulates domain-specific conditions that determine trustworthiness [91]. If a participant architecture is composed of components belonging to different stake-holders, a connector-based security approach can be used to regulate information access within the participant architecture [79].

5. Research Partners

We have five research partners with whom we will be conducting our proposed study. The five are (a) The Aerospace Corporation, a federally funded research and development center (FFRDC), (b) Avaya Labs, a multi-site corporate laboratory focusing on research and development of telecommunications systems and software for business applications, (c) Discovery Science Center, a regional science center focusing on informal science education, (d) Northrop-Grumman Cyber Warfare Integration Center, a corporate laboratory focusing on development organizational knowledge management tools and techniques, and (e) Unimodal Inc., a startup venture focusing on research, development, and commercialization of personal rapid transit systems. Each of our five partners is engaged in the development and use of DVAS in their respective areas of interest. And each of these partners has agreed to provide our research team with access to their DVAS development or usage efforts, as well as inviting insights we gather that might further contribute to and improve the effectiveness of their development and usage efforts.

In the Computer Systems Research Department at *The Aerospace Corporation*, interest focuses on the development of an online environment for coordinating and providing oversight for a loosely coupled network of defense contractors who are building and integrating a system of systems application for the U.S. Air Force. These contractors are required to provide sensitive proprietary data and information regarding the performance of the systems they are developing or integrating to The Aerospace Corporation on behalf of the Air Force, yet these contractors do not want to have their proprietary data or information disclosed to their competitors. Thus, The Aerospace Corporation is interested in a new multi-contractor project oversight environment where they can securely receive and transmit proprietary data and information to/from the defense contractors they are coordinating, yet at the same time be able to partially reveal selected data or information about the performance of one or more artifacts to others in order to maximize the overall likelihood of success in the development. Additional interests within the scope of this management task include issue tracking and management, meetings, and software configuration management. Such a multi-contractor environment is envisioned as a DVAS.

In the Collaborative Applications Research group at *Avaya Labs*, our partner is focus on problems of how best to manage and coordinate distributed, multi-site, and multi-national software development projects through collaborative technologies. Though their software development projects are normally within their corporate boundaries, each site's development team is located within a corporate profit center that may be operating in a different time zone, country, and work culture. They seek to develop a DVAS to help them visualize the location, availability, and technical skills/capabilities of software developers who are part of a multi-site software development projects, yet who might be reassigned or made unavailable to the project team during their development efforts. Participating software developers in turn seek ways to control how information about their status, expertise, and availability is shared or distributed with others on the team in different corporate locations.

At the *Discovery Science Center*, they are in the business of providing hands-on informal science education experiences to more than 120,000 K-8 grade students, and upwards of another 300,000 members of the public who visit DSC year round. DSC in collaboration with the UCI Game Culture and Technology Laboratory recently completed the development and deployment of a single-player and multi-player online science learning game environment called, *DinoQuest Online* [2]. DSC is now launching a new phase of development and internationalization of DQO that will provide support for collaborative problem solving among DQO game players (school age children) via social networking and communication services, as well as expanded multi-player game content, and multi-site networked game play across regional science centers in the U.S. and abroad. DQO was designed to conform to California Science Education Standards for the life sciences in grades K-6. As a result, DSC is providing Web-based access to DQO to students in more than 40 Southern California school districts, as well as developing teacher training and curricular materials that are designed to help teachers and students get the most of the DQO science learning game experience. But these teachers, students, parents, and school administrators are all interested in seeing how well such a game environment works for assisting or enabling young students to learn the basics of life science, and perhaps even plant the seed within them for a career in science. Members of each of these groups have expressed interest in being able to add to, extend, or modify the DQO science game content or software. This in turn has encouraged DSC to look for ways to accommodate open source software (and content) development methods as a way of further engaging its customer and partner base, but at the same time seeking to maintain or improve the quality and educational value of the DQO science learning games. Thus, the ongoing development and use of DQO is to make it more of a DVAS that both continues to support immersive science learning game play, as well as decentralized open source software development activities and local tailoring of DQO to regional science centers in the U.S. and elsewhere.

The Cyber Warfare Integration Laboratory at *Northrop Grumman* has been investigating potential applications of virtual world technologies to support the prototyping and development of advanced military systems and training applications. Their attention has been focused on articulating ways and means for system developers, designed, project managers, customers, and sub-contractors to better understand both what the envisioned system under development looks like, or how it is to be used/experienced, in order to better capture the various kinds of organizational knowledge and expertise that must be mobilized to insure the system's feasibility and subsequent production. They have been experimenting with the prototyping of new facilities and systems for military applications using virtual world technologies like *Second Life*, as well as a variety of computer game technologies. Thus, what makes their efforts relevant to our study is through how they employ, prototype, and evaluate different VAS technologies to engage and elicit project specific knowledge and expertise in ways they can capture, represent, and access across stakeholders working in different enterprises.

Unimodal Inc. is a start-up venture focusing on the development and commercialization of personal rapid transit (PRT) systems for regional and venue-specific deployments. PRT systems are still mostly experimental, yet there is great interest in their potential to relieve transportation congestion and reduce the carbon footprint of current transit options. But what do such systems look like, how do they operate, where/how will they be installed and configured to run, are among the multitude of questions that consumers, system developers, financial investors, and regional transportation and governmental authorities are asking. To help answer such questions, Unimodal has engaged a research team lead by Prof. Crista Lopes to create PRT system mock-ups within the virtual world of Second Life to better help articulate the answers to these questions, but also to help identify and answer new questions that emerge along the way. In this regard, the PRT system research and development effort underway at Unimodal is employing virtual world technologies to help a diverse community of stakeholders to see/visualize and interactively experience a virtual PRT system operating in a virtual world, in order to help these stakeholders make sense of what they see and experience, as well as to help elicit their concerns, interests, and questions regarding the emerging design and commercialization of a proposed PRT system at different installation sites. As such, it is yet another variation of a unique DVAS development and usage effort for us to study and potentially influence in the future.

6. Research approach and methods

Our research approach is empirical. Our methods are primarily focused on qualitative observation and inquiry through field study, and experimental in that they focus on prototyping and evaluating new DVAS concepts, techniques, and tools in the context of different types of organizations and applications of specific DVAS development and usage efforts.

It is always a challenge to model the array of factors that affect real world decentralized work activities, e.g. work pressures, career trajectories, local influences, routines, experience and expertise, etc. Adding another layer of variability arises through a consideration of how such factors are mediated through interaction or collaborative activities that take place in virtual settings. Therefore, we find the best methodology for understanding the complex interplay of organizational, social, and technological factors involved in reconciling alternative activity system configurations (whether decentralized, virtualized, or both together) is through *in situ* studies. However, this alone cannot produce generalizable results unless effort is made to structure and situate these studies across a set of diverse approaches to developing and using DVASs that are found in a diverse set of organizational settings. Therefore, our field studies must also produce data, activity patterns, work/play practices, development processes, and usage scenarios/experiences that can be analytically coded and compared. It is through comparative analysis across cases arising in different settings, with different DVAS configurations or usage scenarios, or with new concept demonstration or prototyped DVAS mechanisms that we can articulate more generalizable results. Here's how we make this work.

First, each of our six analytical variables (described in Section 3) is associated with one project investigator who is already vested in prior research, and therefore brings both prior research expertise as well as an analytical eye with which to observe conditions and events that impinge on their variable under study.

Second, we have identified five research partners with who we will engage our field studies. However, each of our partners presents a different set of research variables of interest. Table 1 represents our current understanding as to which research variables are central to the DVAS development or use efforts underway at each of our five research partners.

Table 1: Initial variable-partner matches

	Aerospace Corporation	**Avaya Labs**	**Discovery Science Center**	**Northrop Grumman**	**Unimodal Inc.**
Representations and realities			XXX	XXX	XXX
Conflicting policies and practices	XXX	XXX		XXX	
Relationship work		XXX	XXX	XXX	XXX
Processes and coordination	XXX	XXX	XXX	XXX	XXX
Privacy and awareness	XXX	XXX			
Security and trust	XXX		XXX		XXX

This distribution of research variable-partner matches signifies that we have a sufficient diversity of variable mixes across the five research partner sites. This provides us a basis for conducting comparative analyses across (a) individual variable-partner matches (comparing results from studies of each cell in the table), (b) range of variables found in each organization (comparing across a column), (c) range of organizations articulating each variable (comparing across each row), and (d) all research variables across all organizations (comparing across columns and rows together). This gives us a maximum comparative analytical capability that can inform or suggest generalizations that account for data that articulate the variables of our proposed study.

Last, our research method must insure a balanced study of socio-technical processes, practices, constraints, opportunities, and affordances that help characterize how different DVASs are developed and used in their particular multi-organizational setting. Accordingly, three of our research variables (security and trust, privacy and awareness, and representations and realities in an immersive world) are associated with experimental development and usage of new DVAS technologies--tool or functional mechanism prototypes that may be incorporated into their respective research partner DVAS project effort. The other three research variables (processes and coordination, relationship work, and conflict practices and policies) focus on identifying, articulating, and modeling/codifying recurring social practices, activity patterns, and processes that characterize the development and/or use of situated DVAS. Though we will begin focusing on one enterprise partner for each analytic variable, we will continue studying these variables with other partners as well as shown in Table 1.

For our study of social practices and processes we will apply rigorous qualitative methodology as outlined by [33] [102] [103] [28] [27]. Methods will consist of intensive observation of work activities, repeated semi-structured interviews among participants (by telephone or in-person), and document analysis, when documents are available. The use of these methods provides a rich corpus of data for understanding the complex processes and relationships of interest to us.

Representations and realities: We will study prototyping of immersive models and simulations of personal rapid transit (PRT) systems within a virtual world of *Second Life*. This will include ongoing work with our research partner Unimodal Inc. and several county governments in Southern California, in order to expand the scope and diversity of issues and public policy concerns that are being addressed through early use of the current PRT simulation in *Second Life*. We will make the PRT simulation world open and publicly available for experimental use and in-world evaluation by the diverse participants, and will engage with the local community in having people experience the system.

Conflicting practices and policies: Ethnographic techniques will be employed with Northrup Grumman to 1) understand the array of different perspectives and practices that exist in DVAS found in and 2) to investigate the impacts that new DVAS development and usage that emerge. Ethnographic methods have been widely deployed in the study of collaborative distributed organizational settings yielding insightful findings (e.g., [66] [80] [4] [65] [31] [31] [54] [61]). In some cases, it will be beneficial to isolate particular variables of interest in a laboratory setting for usability testing or to test the effectiveness of particular features of a system. For these investigations, experimental studies in a laboratory will be conducted.

Relationship work: Building on previous research on relationship work [60] [61] [62], we will examine relationship work in decentralized activities that arise in the development and use of virtual worlds for organizational knowledge management at Northrop Grumman. Qualitative work including field observations, in-depth interviews, and document analysis will inform the design of focused surveys to gain broader insights and identify quantitative patterns about how people create and sustain relationships in decentralized activities. These insights will be used to articulate the requirements and design features of DVAS technologies intended to support relationship work in decentralized settings.

Processes and coordination: The primary approach involves empirical study of DVAS development and deployment projects via multi-modal ethnographic discovery and modeling of observed development processes, work practices, and project community dynamics [84] found at our research partners, such as the Discovery Science Center. Current research studies [85] employ this approach in the study of open source software development projects focusing on the development of (a) Internet information infrastructure systems; (b) networked computer games, (c) scientific research in astrophysics and bioinformatics, and (d) administrative computing applications. In the proposed effort, we see DVAS development projects underway or under consideration in each of our enterprise partners, including in the Discovery Science Center in its efforts to continue its development and expansion of interactive, online science learning games and collaborative problem-solving environments publicly accessible over the Web.

Privacy and awareness: We plan to conduct a field study of distributed collaborators at one of our industry partners, Avaya Labs, to uncover attitudes and practices regarding impression management and how these relate to privacy considerations. Data gathering will be performed through surveys, interviews, non-participant observation, focus groups, and experiments. Multiple methodologies are needed partly to gain a holistic understanding of impression management practices, and partly to offset the methodological challenges in privacy research [73].

Results of the field study will be used to generate requirements for prototype implementations of new impression management and privacy controls available at the user interface to Avaya Labs' DVAS. To allow early exploration, the initial prototype will be restricted to an electronic communication system in use at Avaya Labs, such as an instant messaging system. Based on the results, we will extend it to provide a generic and extensible solution for other collaborative systems (e.g., shared calendars, source code repositories, email, blogs). It is likely that such a solution will take the form of a middleware layer that provides impression visibility and management as a "service" via an Application Programming Interface (API). Alternatively, it may take the form of a suite of plug-ins for individual collaborative systems.

Security and trust: Our approach uses threat and risk models that apply to DVAS to identify/develop (a) appropriate trust models to counter them, *and* (b) secure technologies that can be encapsulated within the participant architectures to support the trust models and hence counter the threats. Our initial specific focus will be to work with our partner at The Aerospace Corporation to investigate the threats, risks, and requirements for trust that arise in the context of their effort to develop and project oversight environment for managing system of systems development projects.

Identifying the threats and risks is our first task. Our previous work [95] has already identified a number of critical threats of decentralized systems. But virtual worlds and interactions between them will surely impose some additional constraints and risks that need to be understood. Therefore, we plan to pursue an in-depth examination of the experiences and the findings of the virtual world research community, the CSCW community [42] [97] [90] [7] and our own prior experience to pinpoint the dominant characteristics and vulnerabilities of these systems.

Second, we will identify appropriate trust models. Numerous trust models exist in the research literature [34] [106]. Depending upon the nature of the collaboration involved, a suitable trust model that can counter the identified threats needs to be identified. For example, if the entities are mostly concerned with regulating access to critical resources, credential and policy-based trust models [14] [15] [52] [101] may be adopted. Or, if entities need to determine whether reported information can be trusted, reputation-based trust models [105] [44] [38] [81] [99] [58] [92] may be adopted. Both kinds of models may be required in some circumstances. We plan to develop techniques and tools to support the selection of appropriate trust models. Towards this goal, we have already developed an initial suite of tools that help an application developer to choose a suitable reputation-based trust model. These include the TREF framework [95] that helps identify an initial set of reputation models and the SIFT simulator [96] that simulates this set of reputation models under varying threat scenarios.

Third, we will identify design principles that will provide specific guidance on how to design secure participants. For this, we will draw upon our prior experiences [91] [79] as well as leverage contemporary literature on various types of trust and security frameworks [35] [17] [51] [18] [26] [36]. The final step in our approach is to leverage our extensive experience with software architectural styles to choose an appropriate architectural style to incorporate these design guidelines.

7. Assessment methods and metrics

In order to assess the impact of the technologies we develop, we will first conduct a baseline assessment of the six variables of interest at each research site. This work can be conducted by graduate students guided by the project faculty investigators and will involve interviewing and observing our research partners to understand current practices and to understand what is important to them in terms of enhanced socio-technical support. Thus our assessment measures will emerge from the research and our interactions with our partners. This means that we will have a set of measures that include both measures common to all partners as well as some that are specific. As our sites are carefully chosen to represent diverse DVAS, we believe this approach will be fruitful in leading to scientific understanding of DVAS as a single class of systems as well as understandings about specific kinds of DVAS. Once the baseline study has been conducted we will be in a position to measure the impact of the technologies and practices we design on the six variables, to ascertain how they change activities in DVAS.

Beyond this, how will we or others come to know whether the proposed research undertaking has produced meaningful and usable results from our study? To help determine this, we presented above a representative set of research questions aligned with each of the six research variables that are the focus of our study. The answers, explanations, interpretations, models, or metrics we provide in response to the systematic observational data we collect, and the multi-layered comparative analysis we undertake, ultimately determine whether we have realized our goals. As such, we now describe in greater detail how we plan to assess our research goals in each of the six research variable areas.

Representations and realities: Building on our efforts to date in prototyping an immersive virtual world for evaluating a new PRT system with our research partner, Unimodal Inc., we will collect and interpret data addressing the following detailed variables of interest. These include determining how many people visit the SL site, what stakeholder communities they represent, and the degree to which the PRT immersive simulation is used/referenced in stakeholders meetings. The particular questions of blending real and virtual elements will focus on identifying the critical circumstances and attributes when melding the elements is successful (such as when an element is itself virtual, such as control software). Issues of representation and reality will similarly be addressed in the context of our other partners.

Conflicting policies and practices: Following the approach used by Mark [54], we will focus on identifying what new conventions of practice and policy have been formed, how they address decentralized work, and importantly, whether they are *followed* by collaborating partners. Perspectives and practices, and policies are dynamic, and constantly evolve and we will also track the extent to which these change over time to adapt to new decentralized configurations. Ethnographic techniques will enable us to identify and evaluate the relevance of new practices and policies to conditions.

Relationship work: The research will examine the entire ecology of tools and techniques for relationship work used by participants, beginning at our research partner, Northrop Grumman. We will investigate how their tools and techniques are used in conjunction with existing collaboration tools such as those for social networking tools, email, instant messaging, video conferencing, and so on. We will identify ways in which the tools enhance one another as well as remaining gaps that need to be filled for more effective relationship work. We will investigate how new mixes of tools emerge, some replacing, enhancing, or complementing others. We will assess how work practices change, and in what ways, using the new tools, discovering how work practices evolve to meet the demands of relationship work in decentralized activities. Subsequent studies will pursue these questions in the context of the other research partners, several of whom have similar situations.

Processes and coordination: Our approach to the use of multi-modal ethnographic and modeling of DVAS processes, practices, and project dynamics that we observe in our research partners is founded on principles of grounded theory development via ongoing comparative analysis of multiple cases of situated practice. Our data collection and analysis methods for process discovery are geared towards both formative and summative assessments of what we can learn from our studies. Our formative assessments rely on a transparent, reconstructable method for data collection, coding and cross-coding, and comparative analysis of multiple cases (both similar and dissimilar). Our summative assessments result from our multi-modal analysis and modeling of the data we collect, and the alternative interpretations we use to present such summary findings and models [84] [85].

Privacy and awareness: We will employ several evaluation techniques throughout the prototype building activities, specifically rapid prototyping and heuristic evaluation in the context of our study at Avaya Labs. We will also conduct usability studies at intermediate stages to iteratively improve the prototype. We will deploy the final prototype versions at UC Irvine as well as industry partners, and collect usage logs as well as conduct interviews with users to understand how well the prototypes meet their goal of supporting impression visibility and management (and in turn of improving privacy management). We will also gather feedback for further improvements. Cross-site comparisons will be made to illuminate how the system should adjust to different organizational and work contexts.

Security and trust: We will evaluate our approach in the context of our research partners, beginning with The Aerospace Corporation's system of systems project oversight and coordination situation. Specifically, we will begin with study of their problem domain, identifying the threats and risks present. We will then use previously developed tools and techniques [93-95] as the basis for identifying a suitable trust model (for those threats so amenable). Such trust models will then be integrated within each participant's architecture using an architectural style that leverages the past PACE work and incorporates connector-based security policies [79] [78]. Next, the application comprising of these participant architectures will be subjected to a variety of threat scenarios and the behavior of each participant studied. These experiments will help us evaluate whether and to what extent these participant architectures in concert with the trust model allow DVAS participants to establish useful trust relationships with each other and control correctly the security of information being exchanged. The threats relevant to this situation are primarily those of inappropriate data access, rather than overtly malicious behavior. Iteration on this process is expected, as shortcomings in the prior work are identified relative to the new circumstances of DVAS, and as differences between our different partner sites are discovered.

8. Intellectual merit and broader impacts

Our focus on DVASs provides the opportunity to study people's experiences with simulated versions of new technological systems while these systems are being conceived and designed, prior to their eventual deployment and use in a broader social setting. We believe that research on DVASs will have a vital impact on society. As development and use of DVASs becomes more common practice and as organizations continue to become more decentralized, new methods and policies will need to be identified and tested to enable people to collaborate successfully. We also believe that our study has economic value as it will help organizations to carry out decentralized work effectively with smoother coordination, so that they can better compete in the global market. DVASs will be resilient to environmental disruptions as collaboration will be able to be conducted from anywhere, anytime, using representations of people, artifacts, and activities. Our results will also have an important impact on higher education, as new people entering the workforce will have to gain skills in developing systems, and in conducting work, in a decentralized setting. Our management plan (below) provides more details on how our results will be integrated into different education settings.

Understanding the unique challenges of relationship work in decentralized activities and designing DVAS technologies to meet the challenges will have both theoretical and practical impact. Theoretically we want to determine the contours of decentralized activity, allowing comparison to other social forms such as traditional hierarchies, communities of practice [100] and the rational bureaucracies described by Max Weber. We believe this theoretical work will be broadly useful in the fields of human-computer interaction, computer-supported collaborative work, organizational studies, and design research and practice that incorporate the wider social context in which technologies are developed and used. Practically, our DVAS studies and prototype tools, techniques, and concepts will enable more effective relationship work in decentralized activities. Through study of their use we will be able to understand future directions needed for continued development of tools for decentralized activity.

Our effort to develop empirically grounded models and theories of decentralized virtual environment development processes, work practices, and project community dynamics both builds on and complements our current studies of open source software development projects. Our studies of the development and deployment of DVASs with our enterprise partners will serve as reference models for how organizations can learn practices, guidelines, models, and theories to follow. For example, our partnership with the Discovery Science Center will provide valuable results to other similar learning centers when they seek to develop and deploy online interactive exhibits and learning games that foster education practices. Similarly, results of our studies with our other enterprise partners can be applied to similar organizations.

The long-term impact of our research on privacy and awareness in the development and use of DVAS is likely to be substantial. Achieving our outcomes will lead privacy-sensitivity to become a standard design requirement in the development of DVASs and similar collaborative systems, and mechanisms for the user to manage their impression on others will become regular to-do items in software design specifications. This is especially important in the light of the fact that most collaborative systems so far focus on the awareness benefits since those are the primary purpose behind building the system, while privacy management often gets secondary attention. We expect some impacts may even go beyond collaborative systems and affect all systems that involve privacy and impression management by computer users.

Our proposed approach, techniques, and tools to the study of how best to provide effective security and trust mechanisms when developing a DVAS will contribute towards the development of secure collaborative open decentralized systems that will be acceptable to and usable by the target audiences. The importance of this should not be underestimated: if appealing DVASs are created, and if they remain defenseless in the face of malicious activities, they will be unused. Provision for security and trust is essential, not an option. Our approach, moreover, will not just identify abstract policies and practices to follow, but provide specific techniques for incorporating security and trust technologies in DVAS software.

9. Results from Prior NSF Funded Research

Project investigators Scacchi, Kobsa, Lopes, Mark, Redmiles, and Taylor all have prior NSF support from at least three different NSF funded research grants.

Award IIS-0205724, amount: $1,800,000, 7/2002-9/2008 (Completed). "ITR: An Integrated Social and Technical Approach to the Development of Distributed, Inter-Organizational Applications." PI/Co-PIs: Taylor, Richardson, Kobsa, Redmiles, Dourish, Mark, van der Hoek. Senior Personnel: Scacchi. This project was a broad empirical study examining the relationship between distributed, inter-organizational management structure and the structure of software built by and for those inter-organizations. The investigation included (a) examining decentralized software development projects, both open source and more standard approaches, (b) examining the role of privacy in decentralized systems, (c) examining security and trust concerns in decentralized systems, and (d) creating technologies to support decentralized development. Over 130 publications resulted from this project, including books, journals, conference papers, and workshops. Over 20 Ph.D students were supported over the course of the project.

Issues of decentralization and consensus that occur in open-source communities were explored, where diverse agencies with independent, often conflicting goals are yet able to come together to produce a high-quality product. Interactions between different teams at NASA were studied, whose communication and interaction patterns exhibit characteristics of decentralization despite ostensibly being part of the same agency. Investigators examined how workers negotiate and manage their membership in multiple communities within a large-scale organization. Investigations delved into how people manage to negotiate and switch their identity/membership among these multiple communities, and why people need these different communities to function effectively in the organization.

To support awareness in distributed configuration management and to aid coordination and collaboration activities among geographically distributed developers, several software tools and visualization techniques were developed as part of this project. At the same time, this need of awareness is frequently at odds with an individual's desire to keep private some of this information. Effectively balancing awareness and privacy needs has proven to be a significant challenge for designers of awareness systems and related infrastructures. Investigation also focused on identifying concerns arising out of decentralization that affect security and trust management solutions for decentralized systems, addressing those concerns through the development of a suitable architectural style, providing design guidelines towards constructing decentralized systems, constructing actual systems using that style, evaluating these systems against threat scenarios, and using experimental results to refine architectural style and design guidelines for future use.

Award no. IIS-0534771, Amount: $335,000, 11/01/05 to 10/31/08, *Discovering the Processes, Practices, Community Dynamics and Principles for Developing Open Source Software Systems*, PI: Walt Scacchi. This research seeks to account for prior results and recent research findings in the area of free/open source software development processes, work practices, and project dynamics. This effort address topics such as (a) the identification of software informalisms as different types of artifacts and communication media that are used to facilitate and coordinate FOSS development projects, as well as serve as decentralized sources of developer knowledge [29] [85], and (b) the role migration and socio-cultural mobilization of socio-technical resources, mitigating conflict, and facilitating career/occupational development in FOSS projects [30] [41].

Award no. 0724806, Amount: $726,455, 09/01/07 to 08/31/10, *SDCI Data New: Trust Management for Open Collaborative Information Repositories: The CalSWIM Cyberinfrastructure*, PI: Cristina Lopes. This project will study and support the California Sustainable Watershed Information Manager. The work just started aims at creating a publicly updatable, Wiki-based online encyclopedia of *all things watershed* that includes all watersheds in California. The CalSWIM Wiki will only succeed if issues of trust are taken into consideration in its design. Professor Lopes also has prior results from NSF supported projects including a CAREER award, in which she continued her original innovative work on Aspect-Oriented Programming [46] with an emphasis towards assessment of new approaches to software design and development.

References

[1] Gartner Group. <http://www.gartner.com/it/page.jsp?id=503861>, 2007.

[2] DinoQuest Online. <http://www.dinoquestonline.org>, 2007.

[3] Abdul-Rahman, A. and Hailes, S. A Distributed Trust Model. In *Proceedings of the New Security Paradigms Workshop.* Langdale, Cumbria UK, 1997.

[4] Ackerman, M.S., Hindus, D., et al. Hanging on the Wire: A Field Study of an Audio-Only Media Space. *ACM Transactions on Computer-Human Interaction.* 4(1), p. 39-66, 1997.

[5] Ahuja, M.K. and Carley, K. Network Structure in Virtual Organizations. *Organization Science.* 10(6), p. 741-757, 1999.

[6] Alter, S. The Work System Method: Connecting People, Processes, and IT for Business Results. Work Systems Press, 2006.

[7] Appelt, W., Pajo, S., et al. Supporting Trust Building in Distributed Groups by Appropriate Security Technology. In *Proceedings of GROUP '07*. Sanibel Island, Florida, USA, November 4-7, 2007.

[8] Bainbridge, W.S. Review: The Scientific Research Potential of Virtual Worlds. *Science.* 317(5837), p. 472-476, July, 2007.

[9] Baker, W. The Network Organization in Theory and Practice. In *Networks and Organizations*, Nohria, N. and Eccles, R. eds. Harvard Business School Press: Boston, MA, 1992.

[10] Bardram, J.E. Applications of Context-Aware Computing in Hospital Work: Examples and Design Principles. In *Proceedings of the 2004 ACM Symposium on Applied Computing.* ACM Press, New York, NY, USA. Nicosia, Cyprus, 2004.

[11] Benkler, Y. The Wealth of Networks: How Social Production Transformation Markets and Freedom. Yale University Press, 2006.

[12] Bergman, M., Lyytinen, K., et al. Boundary Objects in Design: An Ecological View of Design Artifacts. *Journal of the Association for Information Systems, Special Issue on Systems Analysis and Design*. Forthcoming.

[13] Black, J.P., Segmuller, W., et al. Pervasive Computing in Health Care: Smart Spaces and Enterprise Information Systems. In *Proceedings of the Second International Conference on Mobile Systems, Applications, and Services.* Boston, MA, 2004.

[14] Blaze, M., Feigenbaum, J., et al. Decentralized Trust Management. In *Proceedings of the IEEE Symposium on Security and Privacy.* p. 164-173, May, 1996.

[15] Blaze, M., Feigenbaum, J., et al. RFC 2704 - The KeyNote Trust-Management System version 2. <http://www.faqs.org/rfcs/rfc2704.html>, 1999.

[16] Cadiz, J.J., Gupta, A., et al. Using Web Annotations for Asynchronous Collaboration Around Documents. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2000).* ACM Press, New York, NY. Philadelphia, PA, USA, 2000.

[17] Cahill, V., Gray, E., et al. Using Trust for Secure Collaboration in Uncertain Environments. *IEEE Pervasive Computing Mobile and Ubiquitous Computing.* 2(3), p. 52-61, August, 2003.

[18] Capra, L. Engineering Human Trust in Mobile System Collaborations. In *Proceedings of the 12th International Symposium on the Foundations of Software Engineering (SIGSOFT 2004/FSE-12).* Newport Beach, California, USA, November, 2004.

[19] Cheng, L.-T., Hupfer, S., et al. Jazzing up Eclipse with Collaborative Tools. In *Proceedings of the 2003 OOPSLA Workshop on Eclipse Technology eXchange.* ACM Press, New York, NY, USA. Anaheim, CA, USA, 2003.

[20] Curtis, P. MUDing: Social Phenomena in Text-Based Virtual Realities. In *Proceedings of the Conference on Directions and Implications of Advanced Computing.* 1992.

[21] Damiani, E., di Vimercati, S.D.C., et al. A Reputation-Based Approach for Choosing Reliable Resources in Peer-to-Peer Networks. In *Proceedings of the 9th ACM Conference on Computer and Communications Security.* Washington DC, November, 2002.

[22] DeSanctis, G. and Jackson, B. Coordination of Information Technology Management: Team-based Structures and Communication Systems. *Management Information Systems.* 10(4), p. 85-110, 1994.

[23] DeSanctis, G. and Monge, P. Introduction to the special issue: Communication Processes for Virtual Organizations. *Organization Science.* 10(6), p. 693-703, 1999.

[24] Dey, A.K., Salber, D., et al. The Conference Assistant: Combining Context-Awareness with Wearable Computing. In *Proceedings of the 3rd IEEE International Symposium on Wearable Computers.* IEEE Computer Society, Washington, DC, USA. 1999.

[25] Dourish, P. and Bellotti, V. Awareness and Coordination in Shared Workspaces. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96).* ACM. Toronto, Ontario, 1996.

[26] Dragovic, B., Hand, S., et al. Managing Trust and Reputation in the XenoServer Open Platform. In *Proceedings of the First International Conference on Trust Management.* Crete, Greece, May 28-30, 2003.

[27] Dubé, L. and Paré, G. Rigor in Information Systems Positivist Case Research: Current practices, trends, and recommendations. *MIS Quarterly.* 27(3), p. 597-635, 2003.

[28] Eisenhardt, K.M. Building Theories from Case Study Research. *Academy of Management Review.* 14(4), p. 532-550, 1989.

[29] Elliott, M., Ackerman, M.S., et al. Knowledge Work Artifacts: Kernel Cousins for Free/Open Source Software Development. In *Proceedings of the ACM Conference on Support of Group Work (GROUP07).* p. 177-186, Sanibel Island, FL, 2007.

[30] Elliott, M. and Scacchi, W. Mobilization of Software Developers: The Free Software Movement. *Information, Technology and People.* To appear.

[31] Feldman, M.S. Resources in Emerging Structures and Processes of Change. *Organization Science.* 15(3), p. 295-309, May-June, 2004.

[32] Gay, J. ed. Free Software Free Society: Selected Essays of Richard M. Stallman. GNU Press, Free Software Foundation: Boston, MA, 2002.

[33] Glaser, B.a.S., A. The Discovery of Grounded Theory: Strategies for Qualitative Research. Aldine de Gruyter: New York, 1967.

[34] Grandison, T. and Sloman, M. A Survey Of Trust in Internet Applications. *IEEE Communications Surveys.* 3(4), December, 2000.

[35] Grandison, T. and Sloman, M. Trust Management Tools for Internet Applications. In *Proceedings of the 1st International Conference on Trust Management.* Crete, Greece, May, 2003.

[36] Gray, E., O'Connell, P., et al. Towards a Framework for Assessing Trust-Based Admission Control in Collaborative Ad Hoc Applications. Distributed Systems Group, Department of Computer Science, Trinity College, Report TCD-CS-2002-66, 2002.

[37] Grudin, J. Why CSCW Applications Fail: Problems in the Design and Evaluation of Organization of Organizational Interfaces. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '88).* p. 85-93, ACM Press. Portland, Oregon, 1988.

[38] Gupta, M., Judge, P., et al. A Reputation System for Peer-to-Peer Networks. In *Proceedings of the Thirteenth ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video.* Monterey, California, June 1-3, 2003.

[39] Hudson, S.E. and Smith, I. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '96).* p. 248-257, ACM Press, New York, NY, USA. Boston, MA, 1996.

[40] Jensen, C. and Scacchi, W. Process Modeling Across the Web Information Infrastructure. *Software Process--Improvement and Practice.* 10(3), p. 255-272, July-September, 2005.

[41] Jensen, C. and Scacchi, W. Role Migration and Advancement Processes in OSSD Projects: A Comparative Case Study. In *Proceedings of the 29th International Conference on Software Engineering (ICSE 2007).* p. 364-374, Minneapolis, MN, May, 2007.

[42] Jones, S. and Marsh, S. Human-Computer-Human Interaction: Trust in CSCW. *SIGCHI Bulletin.*

29(3), 1997.

[43] Josang, A. and Ismail, R. The Beta Reputation System. In *Proceedings of the 15th Bled Electronic Commerce Conference.* Bled, Slovenia, June 17-19, 2002.

[44] Kamvar, S., Schlosser, M., et al. The EigenTrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the Twelfth International WWW Conference.* Budapest, Hungary, May 20-24, 2003.

[45] Khare, R. and Taylor, R.N. Extending the REpresentational State Transfer Architectural Style for Decentralized Systems. In *Proceedings of the 26th International Conference on Software Engineering (ICSE 2004).* Edinburgh, Scotland, UK, May, 2004.

[46] Kiczales, G., Lamping, J., et al. Aspect-Oriented Programming. In *Proceedings of the 11th European Conference on Object-Oriented Programming.* Volume 1241 of LNCS, p. 220-42, Springer-Verlag. 1997.

[47] Kling, R. Social Analyses of Computing: Theoretical Perspectives in Recent Empirical Research. *ACM Computing Surveys.* 12(1), p. 61-110, 1980.

[48] Kling, R. and Scacchi, W. The Web of Computing: Computing Technology as Social Organization. In *Advances in Computers*, Yovits, M. ed. 21, p. 3-85, 1982.

[49] Kock, N. Benefits from Virtual Organizations from Distributed Groups. *Communications of the ACM.* 43(11), p. 107-112, 2000.

[50] Lee, C., Dourish, P., et al. The Human Infrastructure of Cyberinfrastructure. In *Proceedings of the ACM Conference on CSCW (CSCW'06).* p. 483-492, Banff, Alberta, Canada, 2006. Nominated for Best Paper.

[51] Lee, S., Sherwood, R., et al. Cooperative Peer Groups in NICE. In *Proceedings of the IEEE Infocom.* San Francisco, USA, April 1-3, 2003.

[52] Li, N., Mitchell, J.C., et al. Design of a Role-Based Trust-Management Framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy.* p. 104-120, 2002.

[53] Lopes, C., Kan, L., et al. PRT Simulation in an Immersive Virtual World. In *Proceedings of the 1st International Conference on Simulation Tools and Techniques for Communications, Networks, and Systems (SIMUTools).* In submission.

[54] Mark, G. Conventions and Commitments in CSCW Distributed Groups. *Computer-Supported Cooperative Work: The Journal of Collaborative Computing.* 11(3-4), p. 349-387, 2002.

[55] Mark, G., Gonzalez, V., et al. Reconciling Different Perspectives: An Experiment on Technology Support for Articulation. In *Proceedings of the Fifth International Conference on the Design of Cooperative Systems (COOP 2002).* p. 23-37, IOS Press. Saint-Raphaël, France, June 4-7, 2002.

[56] Mark, G., Abrams, S., et al. Group-to-Group Distance Collaboration: Examining the "Space Between". In *Proceedings of the 8th European Conference of Computer-Supported Cooperative Work (ECSCW '03).* p. 99-118, Helsinki, Finland, September 14-18, 2003.

[57] Markus, L. Power, Politics, and MIS Implementation. *Communications of the ACM.* 26(6), p. 430-444, 1983.

[58] Marti, S. and Garcia-Molina, H. Limited Reputation Sharing in P2P Systems. In *Proceedings of the ACM Conference on Electronic Commerce.* New York, USA, May 17-20, 2004.

[59] Morla, R. and Lopes, C. Blurring the Line between Simulation and Real-world Prototypes using Second Life™. In 9th International Conference on Ubiquitous Computing (UbiComp 2007) - poster: Tyrol, Austria, 2007.

[60] Nardi, B., Whittaker, S., et al. ContactMap: Integrating Communication and Information through Visualizing Personal Social Networks. *Communications of the Association for Computing Machinery.* April, p. 89-95, 2002.

[61] Nardi, B., Whittaker, S., et al. NetWORKers and Their Activity in Intensional Networks. *The Journal of Computer-Supported Cooperative Work.* 11(1-2), p. 205-242, 2002.

[62] Nardi, B.A. Beyond Bandwidth: Dimensions of Connection in Interpersonal Communication. *Journal of Computer Supported Cooperative Work.* 14(2), p. 91-130, 2005.

[63] Neale, D.C., Carroll, J.M., et al. Evaluating Computer-Supported Cooperative Work: Models and

Frameworks. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work.* ACM Press, New York, NY. Chicago, IL, 2004.

[64] Noll, J. and Scacchi, W. Supporting Software Development in Virtual Enterprises. *JoDI - Journal of Digital Information.* 1(4), 1999.

[65] Olson, G. and Olson, J. Distance Matters. *Human-Computer Interaction.* 15(2/3), p. 139-178, 2000.

[66] Orlikowski, W.J. Learning from Notes: Organizational Issues in Groupware Implementation. *The Information Society.* 9(3), p. 237-250, 1993.

[67] Orlikowski, W.J. and Gash, D.C. Technological Frames: Making Sense of Information Technology in Organizations. *ACM Transactions on Information Systems.* 12(2), p. 174-207, 1994.

[68] Palen, L. Social, Individual and Technological Issues for Groupware Calendar Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI).* ACM Press, New York, NY, USA. Pittsburgh, PA, USA, 1999.

[69] Patil, S. and Kobsa, A. Instant Messaging and Privacy. In *Proceedings of the HCI 2004 - The 18th British HCI Group Annual Conference.* p. 85-88, Leeds, England, September, 2004.

[70] Patil, S. and Kobsa, A. Privacy in Collaboration: Managing Impression. In *Proceedings of the First International Conference on Online Communities and Social Computing.* Las Vegas, NV, 2005.

[71] Patil, S. and Kobsa, A. Uncovering Privacy Attitudes in Instant Messaging. In *Proceedings of the 5th ACM Conference on Supporting Group Work (Group '05).* p. 101-104, ACM Press, New York, NY. Sanibel Island, FL, 2005.

[72] Patil, S. and Lai, J. Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI).* ACM Press, New York, NY. Portland, OR, 2005.

[73] Patil, S., Romero, N., et al. Privacy and HCI: Methodologies for Studying Privacy Issues. In *Proceedings of the ACM Conference on Human Factors in Computing Systems.* ACM Press, New York, NY. Montréal, Québec, Canada, 2006.

[74] Patil, S. and Kobsa, A. Privacy Considerations in Awareness Systems: Designing with Privacy in Mind. In *Awareness Systems: Advances in Theory, Methodology and Design,* Markopoulos, P., de Ruyter, B., and Mackay, W. eds. Springer Verlag: Berlin, Heidelberg, New York, Forthcoming.

[75] Pinch, T. and Bijker, W. The Social Construction of Facts and Artifacts. In *The Social Construction of Technological Systems.* p. 17-50, MIT Press: Cambridge, MA, 1987.

[76] Prinz, W., Mark, G., et al. Designing Groupware for Congruency in Use. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '98).* p. 373-382, ACM Press. Seattle, WA, November, 1998.

[77] Pujol, J., Sanguesa, R., et al. Extracting reputation in multi agent systems by means of social network topology. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems.* Bologna, Italy, July 15-19, 2002.

[78] Ren, J., Taylor, R., et al. Towards An Architectural Treatment of Software Security: A Connector-Centric Approach. In P*roceedings of the Workshop on Software Engineering for Secure Systems (SESS 05),* held in conjunction with ICSE 2005. St. Louis, MO, May, 2005.

[79] Ren, J. A Connector-Centric Approach to Architectural Access Control. Thesis. Information and Computer Science, University of California, Irvine, p. 222, 2006. <http://www.ics.uci.edu/~jie/Thesis.pdf>.

[80] Rogers, Y. Exploring Obstacles: Integrating CSCW in Evolving Organizations. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '94).* p. 67-78, 1994.

[81] Sabater, J. and Sierra, C. REGRET: A Reputation Model for Gregarious Societies. In *Proceedings of the 4th Workshop on Deception, Fraud and Trust in Agetn Societies.* Montreal, Canada, 2001.

[82] Scacchi, W. Understanding the Requirements for Developing Open Source Software Systems. *IEE Proceedings--Software.* 149(1), p. 24-39, February, 2002.

[83] Scacchi, W. Free/Open Source Software Development Practices in the Computer Game Community. *IEEE Software.* 21(1), p. 59-67, January/February, 2004.

[84] Scacchi, W., Jensen, C., et al. Multi-Modal Modeling, Analysis and Validation of Open Source

Software Development Processes. *Intern. J. Internet Technology and Web Engineering.* 1(3), p. 49-63, 2006.

[85] Scacchi, W. Free/Open Source Software Development: Recent Research Results and Empirical Methods. In *Advances in Computers*, Zelkowitz, M.V. ed. 69, p. 243-295, 2007.

[86] Simone, C., Mark, G., et al. Interoperability as a Means of Articulation Work. In *Proceedings of the ACM Conference on Work Activities Coordination and Collaboration (WACC '99).* p. 39-48, ACM Press. San Francisco, CA, February 22-25, 1999.

[87] Singhal, S. and Zyda, M. Networked Virtual Environments: Design and Implementation. Addison-Wesley: New York, NY, 1999.

[88] Star, S.L. and Griesemer, J.R. Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science.* 19, p. 387-420, 1989.

[89] Star, S.L. and Ruhleder, K. Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. *Information Systems Research.* 7(1), p. 111-134, 1996.

[90] Stiemerling, O. and Pacull, F. Secure Component-Based Groupware. In *Proceedings of the CSCW 2000 Workshop on Component-Based Groupware.* p. 43-48, Philadelphia, PA, December, 2000.

[91] Suryanarayana, G., Erenkrantz, J.R., et al. PACE: An Architectural Style for Trust Management in Decentralized Applications. In *Proceedings of the 4th Working IEEE/IFIP Conference on Software Architecture.* p. 221-230, Oslo, Norway, June, 2004.

[92] Suryanarayana, G. and Taylor, R.N. A Survey of Trust Management and Resource Discovery Technologies in Peer-to-Peer Applications. Institute for Software Research, University of California, Irvine, Technical Report UCI-ISR-04-6, July, 2004.

[93] Suryanarayana, G., Erenkrantz, J.R., et al. An Architectural Approach for Decentralized Trust Management. *IEEE Internet Computing.* 9(6), p. 16-23, November/December, 2005.

[94] Suryanarayana, G., Diallo, M.H., et al. Architectural Support for Trust Models in Decentralized Applications. In *Proceedings of the 28th International Conference on Software Engineering (ICSE 2006).* p. 52-61, Shanghai, China, May, 2006.

[95] Suryanarayana, G. and Taylor, R.N. TREF: A Threat-centric Comparison Framework for Decentralized Reputation Models. Institute for Software Research, University of California, Irvine, Technical Report UCI-ISR-06-2, January, 2006.

[96] Suryanarayana, G. and Taylor, R.N. SIFT: A Simulation Framework for Analyzing Decentralized Reputation-based Trust Models. Institute for Software Research, University of California, Irvine, Technical Report UCI-ISR-07-5, August, 2007.

[97] Van House, N.A., Butler, M.H., et al. Cooperative Knowledge Work and Practices of Trust: Sharing Environmental Planning Data Sets. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '98).* p. 335-343, ACM Press, New York, NY. Seattle, WA, November 14-18, 1998.

[98] Vince, J. and Earnshaw, R. Virtual Worlds on the Internet. IEEE Computer Society, IEEE Press: Los Alamitos, CA, 1998.

[99] Walsh, K. and Sirer, E.G. Thwarting P2P Pollution Using Object Reputation. Department of Computer Science, Cornell University, Report cul.cis/TR2005-1980, 2005.

[100] Wenger, E. Communities of Practice. Cambridge University Press, 1998.

[101] Yao, W.T.-M. Fidelis: A Policy-Driven Trust Management Framework. In *Proceedings of the 1st International Conference on Trust Management.* 2692, p. 301-317, Springer-Verlag Heidelberg. Heraklion, Crete, Greece, May 28-30, 2003, 2003.

[102] Yin, R.K. The Case Study as a Serious Research Strategy. *Knowledge: Creation, Diffusion, Utilization.* 3, p. 97-114, 1981.

[103] Yin, R.K. Case Study Research, Design and Methods. 2nd ed. ed. Sage Publications: Newbury Park, CA, 1994.

[104] Yu, B. and Singh, M.P. A Social Mechanism of Reputation Management in Electronic Communities. In *Proceedings of the Fourth International Workshop on Cooperative Information Agents.*

p. 154-165, 2000.

[105] Zacharia, G. and Maes, P. Collaborative Reputation Mechanisms in Electronic Marketplaces. In *Proceedings of the 32nd Hawaii International Conference on System Sciences.* Hawaii, 1999.

[106] Zacharia, G. and Maes, P. Trust Management Through Reputation Mechanisms. *Applied Artificial Intelligence.* 14, p. 881-907, 2000.

[107] Zimmerman, A. and Nardi, B. Whither or Whether HCI: Requirements Analysis for Multi-Sited, Multi-User Cyberinfrastructures. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'06).* p. 1601-1606, Montreal, Ontario, Canada, 2006.

[108] Zuboff, S. In the Age of the Smart Machine. Basic Books: New York, 1988.

[109] Julsrud, T. Core/periphery Structures and Trust in Distributed Work Groups: A comparative case study. *Structure and Dynamics: eJournal of Anthropological Related Sciences.* Forthcoming.