This document is a proposal to the National Science Foundation in response to the 1998 Knowledge and Distributed Intelligence in the Information Age (KDI) solicitation, targeted on the Knowledge Networking focus of the solicitation.

Project Summary

Objectives and Methods. The primary objective is a theory of effective, sustainable, geographically distributed teamwork. In an economy of large national and global organizations and virtual enterprises, collaboration among geographically distributed team members is essential. Increasingly, people are able to work together across distances by using technologies to collaborate in a shared virtual space. They create artifacts while conversing about the artifacts and the processes used to create them. Taken together, these artifacts and conversations constitute a knowledge network that must sustain the team's collaboration. However, virtual collaboration is also very challenging. Too often, the technology of the knowledge network becomes a hindrance, impeding the free flow of collective intelligence and shared understanding.

The theory will be based on empirical studies of distributed teams in the World Bank, Boeing, and Bell Atlantic. We will observe teams at work, interview team members, and analyze both their conversations and artifacts to determine the salient causal factors in the long-term sustainability of distributed collaborations. These observations are expected to yield a sort of ecological theory consisting of principles that differentiate environments that sustain teamwork from those that cripple it. Some aspects of the theory will be tested in experiments conducted either in the field or in a university laboratory.

A second objective is to capture the theory and its implications in an online "Handbook." This Handbook will serve as the research team's collaboratory, coordinating our research processes and results. The Handbook will encourage communication among researchers in this area and provide recommendations to members of distributed teams.

A third objective is to develop a software-based "Collaboration Assistant," based on the theory, that provides elements of a sustainable collaboration environment. The Collaboration Assistant will be conceptualized, designed, and implemented using techniques of iterative team design and participatory design. Pilot teams will use the Collaboration Assistant, providing validation of the theory and its usefulness.

Potential Impact. Sweeping social changes are rapidly changing the demographics of work, including the disaggregation of knowledge workers into distributed and partially collocated teams. Currently, few resources exist that can assist these distributed teams in improving the quality and sustainability of their collaboration, such as diagnosing and "healing" problems in their shared virtual space. Achieving the objectives of this proposal will result in the creation of both tools and theory that will provide support for knowledge networks and avoid many of the pitfalls currently experienced by distributed teams. If virtual teams become a more reliable and productive way of conducting the business of knowledge work, it could enable us to work more efficiently and in more flexible, enjoyable ways.

Project Description

Towards an Ecological Theory of Sustainable Knowledge Networks

Section 1. Conceptual Framework

As the global economy heats up, there is more and more pressure on knowledge workers to collaborate, to innovate, to make decisions, and to do all of this faster than ever. At the same time, sweeping social changes are disaggregating many knowledge workers, forcing them to work together at a distance, with few opportunities for face-to-face interaction. Telecommuting, multi-organization projects, executives who are frequently on the road, branch offices across a large organization or government agency, and the use of remote consultants and scientific experts all conspire to place new demands on teams[1] of knowledge workers and on knowledge networks.

While the telephone, fax machine, e-mail, and a host of other technologies provide the necessary communication infrastructure for these distributed workers, these technologies do not address the specific obstacles of *creative* collaboration at a distance. Creative collaboration involves the creation and sharing of two kinds of information: the object being created itself, and interactions about that object and the process of creating it. These are the two aspects of "reflection in action" (Schon, 1984), the duality of object[2] and process (Robinson, 1991; Conklin, 1992). For collocated teams, creative collaboration is challenging enough, but there are a host of conventional tools and practices of face-to-face collaboration that provide a literal shared space to support their work. For distributed teams seeking virtual collaboration, however, technology must provide a replacement for the literal shared space – a virtual collaborative space (Schrage, 1995) or common information space (Bannon & Bødker, 1997; Schmidt & Bannon, 1992).

---

[1] Our focus is on distributed knowledge workers collaborating in some kind of creative activity. This may be a formal team, a working group, or an ad hoc and open collective. In all these cases, membership is dynamic and the group's boundaries fuzzy. For simplicity, we loosely use the blanket term "virtual team" to mean any distributed group focused on a shared creative task or goal.

[2] Such objects of creation can be as simple as a policy memo or a spreadsheet, and as complex as a piece of legislation, a multi-disciplinary research proposal, a multi-organization business plan, or a hardware system design.

The most common and familiar kinds of virtual shared spaces are on-line discussion databases, which are created around a topic such as "Sailing" or "Italian Opera" and in which a self-selected group of people hold an electronic discussion by posting comments and responses. USENET newsgroups and the PLATO project at University of Illinois, Champaign-Urbana were early examples of this application, and now there are thousands of such public discussions using a variety of technologies including Lotus Notes™ and email listserve and shared folder systems. These technologies have been put to use within organizations for corporate projects ranging from community building and best practices repositories to project-specific collaboration spaces. For example, there are many case studies of organizational applications of Lotus Notes discussion databases (e.g. Lloyd & Whitehead, 1996; Orlikowski, 1995 & 1992).

Discussion tools fall far short of supporting the collaborative creation of complex evolving artifacts. Conversely, technologies exist that allow groups to collectively add to and modify a complex artifact, such as source control systems in software engineering, but these systems don't generally offer significant tools for communication about these artifacts and the process of creating them. Ultimately, virtual collaboration environments will need to provide tightly integrated support for both aspects of creative collaboration. Moreover, robust knowledge networking by virtual teams will require systems that allow knowledge workers to move quickly and seamlessly between object-building activities and discussions about the object and the process of building it (Mark, Fuchs, & Sohlenkamp, 1997).

The challenge is much greater than integrating existing technologies – the current technologies provide inadequate support for sustained distributed teamwork. For example, one member of our project team was involved in a consulting contract to help a client understand what was in their discussion database. The client, a geographically distributed senior team in a large bank, had been engaged for over a year in a Notes™ discussion about how to deal with an upcoming merger. This database had over a megabyte of threaded textual material concerning a variety of issues. The senior team could no longer understand what was in their database: what the group had learned, what the most important issues were, or what they should do about the merger. The consultants spent weeks analyzing the database, and in the end were able to create a summary that illuminated several major issues that were hidden in a fragmented and incoherent mass of textual material[3].

This is not an isolated anecdote. We believe, based on direct experience in this industry over the past decade, that for every published case study of a successful

[3] The analysis was conducted using a rhetorical structure called Issue Based Information System (IBIS) which structures conversations in terms of questions, possible answers, and pros and cons (Kunz & Rittel, 1970). IBIS is the basis of the QuestMap™ software, which was used to map out the analysis for final report to the bank.

virtual collaboration there are many more unsuccessful attempts by virtual teams to use CSCW and groupware tools. Of course, the reports of these failures do not find public attention as readily as the successes, but there is nevertheless a small but significant literature on CSCW and groupware failures (Neilson 1997; Orlikowski, 1992; Grudin, 1994 & 1990; Markus & Connolly, 1990; Star & Ruhleder, 1996).

The lack of success may range from participant disappointment and frustration with the tools and interactions to cancellation of the project. As the size of the project team increases, and the duration of the project extends in time, there is inevitably a point where virtual collaboration spaces cost more – in human attention and energy, in time, and in frustration and disorientation – than the value they provide. Inevitably, this leads to various forms of resistance and non-participation (overt and covert) in the virtual dimension of the collaboration. At this point, we can declare that *the virtual collaboration system has failed*, irrespective of the ultimate success of the project.

Our experiences with virtual project teams, and the literature on the pitfalls and challenges of successful CSCW implementations, suggest that many causes of failure stem from a poor ecological fit between knowledge workers and the virtual environments they try to inhabit. Knowledge networks need to be adapted to human needs, preferences, and abilities (Turnage, 1990). They need to be flexible enough to respect the cognitive, cultural, and educational differences among the many participants in virtual collaborations. In addition, they need to reflect a theoretical understanding of the cognitive and social processes of creating, developing, maintaining, and dismantling knowledge networks.

We propose here to assemble systematically the elements of an ecological theory for knowledge networks, with a special emphasis on issues of viability and long-term sustainability in the human/technology ecosystem. This is an ambitious undertaking, but we believe that the chances of success are good. This research team has over 80 years of collective experience in research, design, development, and deployment of CSCW systems. Significantly, our experiences and studies have dealt with successful and *unsuccessful* CSCW implementations. Others on the team have staked out new territory in the field of virtual collaborations. Several members of this team have strong international reputations in both academic and commercial communities, and are known for their theoretical contributions to science. Finally, we believe that we have designed a research program, described below, which maximizes the likelihood of far-reaching scientific discovery.

Why sustainability is crucial

The incredible success and growth of the Internet and World Wide Web are a testament to the effectiveness of many aspects of the modern communication infrastructure. Email, for example, has evolved from an arcane engineering curiosity to a business necessity within a few decades. However, most of the innovation in communication and collaboration technology has been focused on the "low hanging fruit", the little pieces of technology that had large impacts and payoffs.

The remaining fruit, such as effective and sustainable knowledge networking, is not so easy to reach. A few pieces of this higher fruit may be readily "plucked" by new and emerging products. However, when it comes to large complex projects, such as the design of an entire modern passenger aircraft or the development of an effective and appropriate loan to a developing nation, no single technology will provide all of the elements necessary for a diverse group of experts to maintain a rewarding and sustainable virtual interaction.

In focusing on sustainability, we do not mean to imply that knowledge networks should last forever, or that any pause in growth implies an ecological problem. Virtual collaborations, like face-to-face collaborations, have a natural rhythm and pace to them; before they die they may go to sleep and even go dormant for periods of time. The practical challenge for knowledge networks is with illness and potentially early death – the failure of the ecology to sustain a virtual team for its chartered or intended life span.

Among the many impediments to successful virtual collaboration are broad issues of cognitive overload, lack of awareness of others' activities, lack of context about the purpose and history of shared objects, lack of appropriate feedback and acknowledgment of contributions to the shared space, disorientation, unbalanced participation, divergence without convergence, and problems with bringing new team members into a complex and established collaborative space.

While each of these breakdowns could be studied individually, this set of issues reflects the use of tools developed in the absence of a practical theory, a theory that would cover the cognitive, social and architectural issues concerning any sustainable virtual environment.

First Steps Towards an Ecological Theory

Ecology is a good model for the kind of theory that would be appropriate for this research. In place of hard laws the field of ecology offers principles[4]. It describes

[4] "Most field studies of living organisms now have an ecological orientation. Basic principles [include] the universality and nature of adaptation; the concept of populations as dynamic groups;

patterns of interactions among diverse populations, sometimes competing, sometimes cooperating, in a textured and changing environment. It is a field that, while rich in quantitative methods, creates theories about systems and wholes.

One possible ecological principle for sustainable virtual collaboration is borrowed from economic theory. If the "net interaction value" is defined as personal value less personal cost[5] of a virtual interaction, then a principle might be that *the more often an individual's net interaction value is negative, the less likely that individual is to continue to participate in the virtual collaboration.*

Another principle might deal with "critical mass" phenomena, e.g. *the larger the proportion of relevant stakeholders who participate in a virtual collaboration, the more likely the collaboration will sustain itself (all else being equal).*

A third principle of sustainability might describe the role of *mediation: there is a natural "information entropy" which, over time, creates incoherence and cognitive overload in collaborative spaces, and there is a complementary counter-entropic force, usually but not necessarily human, which adds order and coherence to these spaces.*

An ecological theory of virtual sustainability will offer, in addition to principles, specific "laws" or hypotheses describing the dynamics of virtual collaborative environments. Provisionally, these hypotheses fall into three broad categories: factors relating to *coordination and process*, factors relating to *shared objects* (e.g. context, usefulness), and factors relating to *orientation and cognitive overload* in the collaborative space itself[6]. Here are some testable working hypotheses dealing with these three factors, to provide a concrete example of the kind of theory the research aims to create.

Hypotheses about coordination and process:

1. One cause of unsustainability of knowledge networks is loss of awareness or incomplete awareness of activities of others in the system. Cues about what other people are doing are important in helping individuals to orient in the collaborative space and to pace their interactions. (Gutwin, Roseman, & Greenberg, 1996; Gutwin & Greenberg, 1998; Mark, Fuchs, & Sohlenkamp, 1997)

2. Teams routinely create and share coordination objects, such as action item lists, project timelines, and bug lists, to sustain their interactions. Making these coordination objects easier to share across distributed groups, both

---

[and] the principle that all organisms are parts of multispecific communities, which are organized and interacting units." (Simpson et al, 1957)

[5] Of course, these terms must be given precise measurable definitions.

[6] This list corresponds to the process/object duality mentioned above, but adds a third factor dealing with the coherence of the structure of the network of links between the first two.

asynchronously and as a resource in distributed meetings, contributes to sustainability (Farkas & Poltrock, 1996; Sharples, 1993).

3. An explicit representation of the team/project lifecycle, assessment as to where you stand within it, and a guide showing to how to move through it would help maintain process orientation in a virtual project team. (The process of team formation and evolution is complex. Leaders have tacit knowledge about how to move a team through a process, and they access that knowledge in face-to-face meetings. When in virtual collaborations, they don't have that, e.g. they may not recognize that they don't have alignment about team goals.)

4. A team's processes are supported by creating shared objects with affordances that guide team members to the next step in the process. For example, a document ready for review and approval may have editable sections for comment and signature while the content of the document remains locked. The object guides users to the appropriate tasks. Such objects engage tacit knowledge and support teamwork more effectively than explicit representations of the team processes.

5. Individuals who receive little or no acknowledgment of or responses to their contributions in a shared space will tend to withdraw from participation in the on-line collaboration. This becomes especially critical the more the group's purpose is abstract or strategic knowledge work (as opposed to group object-building activity).

6. Greater levels of member participation and wider distribution of participation (i.e., no one is dominating) will be associated with higher levels of perceived user satisfaction with the system and greater sustainability.

Hypotheses about creating and using shared objects:

7. People will only put attention/energy into a shared object, and thus make effective use of it, if they have some sense of its context: "Why am I looking at this? What am I expected/allowed to do with it?" (Context includes *who* and *why* and *what next*. *Who* covers author, intended audience, and readers. *Why* could be answered by an annotation, or by what else the object is linked to. *What next* addresses where the object is in the flow.)

8. Corollary to #7: How others have used objects in shared space helps people orient to those objects. (Example scenario: one is trying to make sense of a PowerPoint presentation that has been placed ambiguously in the shared space; using Collaboration Assistant one finds out that X used this presentation on Tuesday to a management group as part of the pitch for

Project Y and there was a lot of interest and enthusiasm; this important context makes one more likely to remember where it is and to make effective use of that object later.)

9. Sustainability will be a function of usefulness of shared objects to the individuals in the system. (If most of the information in the collaboration space is not personally relevant or salient to an individual, it will take much higher levels of extrinsic motivation to secure their creative input over time; and, it is more likely that they will withdraw or that their contributions will be suboptimal.)

Hypotheses about orientation and cognitive overload:

10. Individuals who get saturated or disoriented by information spaces will tend to drop out; activities/services which reduce cognitive load will increase sustained use of the shared space (Conklin, 1987). (As formal and informal information accumulates it becomes harder and harder to orient in the shared space, and to know what the salient information is; this causes a kind of cognitive overload that discourages people from participating in the shared space.)

11. Corollary to #10: New participants in the virtual collaboration are especially susceptible to this kind of overload, and require special kinds of summary and introduction to be "apprenticed into" the virtual collaboration (Poltrock & Englebeck, 1998 & 1997).

12. Corollary to #10: Mediation in a shared space (e.g. providing summaries, archiving irrelevant or old information) can sharply reduce cognitive overload.

13. Systems that support different modes of interaction (e.g. from independent to loosely coupled to tightly coupled) and allow users to move easily among them will be more satisfying to users and inherently more sustainable (Haake & Wilson, 1992; Streitz et al., 1992). (Systems that fragment activities into different applications and media, with little or no sharing among them, will be less sustainable.)

These hypotheses, if validated, have various implications. Some suggest the need for *mediation* that would serve to prevent or ameliorate factors that otherwise decrease sustainability (Okamura et al, 1994) (mediation in this sense could be the role of a human or the activities of a computer program such as the Collaboration Assistant). Other hypotheses suggest features of virtual collaboration systems that, if implemented, would make such systems more habitable and collaborations in them more sustainable. Thus, the emerging ecological theory has both *social* and *architectural* implications.

The goal of this research project is to use empirical studies to develop this initial small set of tentative hypotheses into a practical theory of knowledge network sustainability. This theory will be expressed in the form of two deliverables: a Handbook of virtual collaboration, and a computer program, a Collaboration Assistant, which provides practical guidance to virtual teams.

Section 2. Outline of the Work Plan

The primary deliverables from this research are a theoretical framework, represented in published papers and project documentation, as well as a computer program that embodies elements of the theory and mediates virtual collaborations in knowledge networks.

The plan for this undertaking consists of five research activities (each is described more fully in its own section below):

1. Activity 1: *Observations* – Observing and analyzing existing on-line collaborative spaces and virtual teams (those that have died as well as those that continue to survive) for evidence of the causal factors of sustained collaborative health.
2. Activity 2: *Handbook* – Building a theory and assembling a Handbook consisting of the theory and principles, with qualifications and examples. This will be a living document, accessible on the WWW.
3. Activity 3: *Experiments* – Conducting controlled experiments designed to test hypotheses like those listed above, and to qualify and quantify specific parameters of essential sustainability principles.
4. Activity 4: *CA* – a prototype software "Collaboration Assistant" which implements the principles of the emerging ecological theory, and which mediates potential problems in the health of a collaborative space, sometimes automatically, sometimes by advising a human mediator.
5. Activity 5: *Testing* – Evaluating and testing both the theory and the Collaboration Assistant with real world distributed teams who are using collaborative technology to solve live problems.

Each of these activities has a leader and a primary center of activity. These activities are not sequential – *Observations* and *Handbook* dominate the first year, *Experiments* and *CA* dominate the second year, and *Testing* dominates the third year, but there is also considerable overlap among them (see attached Project Management Plan).

Section 3. Activity 1: Observing and analyzing virtual collaborations

This activity will provide the empirical foundation for development of an ecological theory. We will conduct a systematic examination of organizationally-based virtual teams in their natural environments to determine the key factors relating to sustainability and productive functioning of their shared spaces. Where practical we will focus on project teams engaged in a finite creative collaboration, as opposed to operational or functional teams. This research will extend and integrate ongoing studies at Boeing (Poltrock & Engelbeck, 1998 & 1997), Bell Atlantic (Selvin, 1996), and the World Bank[7]. The studies underway now at these organizations are all concerned with support for effective distributed or partially distributed teams, but they differ greatly in approach, methodology, theoretical framework, and focus. This research will establish a common framework and methodology that will enable comparing results obtained in different organizations.

We will study how distributed and partially distributed teams perform their work, focusing on factors relating to coordination and process, factors relating to sharing objects, and factors relating to orientation and cognitive overload. Generally, these teams work together both synchronously and asynchronously, and we will observe both kinds of interactions and attempt to understand the role of synchronous interactions and how these could be supported in an asynchronous collaboration environment. Factors we plan to analyze for include:

- level of awareness of other's participation in the collaborative space (Hypothesis #1),
- level and accuracy of awareness of the status of the team with respect to it's milestones and lifecycle (#3),
- level of shared understanding among group members about their respective roles in the project,
- individual's overall satisfaction in the collaboration and their satisfaction with specific aspects of it,
- individual's sense of being heard and respected within the group and it's impact on their participation (#5),
- level and accuracy of awareness of the structure of the shared space (#10),
- level of cognitive overload in the shared space (#10),
- ease of access to shared coordination objects (#2), and
- level of usefulness of the shared objects in the collaborative space (#9).

Both qualitative and quantitative analyses will be employed in this activity. Assessments will include: observation of teams at work; examination of written

[7] Nothing yet published.

team products; process, and history (e.g., computer logs, e-mail, online discussions, product iterations, stored materials); detailed participant interviews; surveys examining both team process, attitudes toward the process and people, perceived team success/failure; and organizational or supervisor evaluations of the teams' efforts. Team purpose within the organization will also be examined as a potentially relevant factor (e.g., software development teams vs. strategic document production teams) as will potential individual differences in functioning in virtual collaboration environments.

Data collection sites will include Bell Atlantic, Boeing, the World Bank, and Corporate Memory Systems[8]; letters of support from appropriate organizational sponsors within these organizations are attached.

A matched-pairs design will be used to select for assessment pairs of teams within a given organization of the same team type/purpose, where one is currently active and has functioned successfully for at least one year, and another has failed to meet expectations set for it (either disbanding prematurely, prior to initial expectations, or continuing to exist, but with limited productivity). In both cases, "success" of the virtual team will be assessed based on evaluations of both participants and supervisors. Teams expected to have an extended life span (i.e., where sustainability is a vital factor) will be selected. A matrix illustrating sampling and proposed comparisons appears in Table 1.

	Team Type	
Organization	Active Success	Limited Success
A	1	2
B	3	4

Table 1: Four types of teams to be studied

The desired comparisons are:

More Successful vs. Less Successful Virtual Teams	(1 and 3 vs. 2 and 4)
Organizational Impact on Sustainability	(1 and 2 vs. 3 and 4)

In addition, we plan to examine at least one database constructed within a state-of-the-art hypertext collaboration environment[9]. The software project team which created this database over a two year period made heavy use of several kinds of hypertext links not available in other collaboration systems (Selvin, 1998 & 1996), and thus encountered different kinds of opportunities and barriers than teams using more conventional systems (e.g. email, Lotus Notes™) would encounter. This analysis should provide a revealing counterpoint to the other

[8] One of the co-PI's, Conklin, was a founder of Corporate Memory Systems, where extensive virtual collaboration (though collocated) occurred over a 5 year period, and has full access to the defunct corporation's QuestMap database.

[9] An augmented version of QuestMap™ developed at Bell Atlantic.

observational sites. This work is being contributed by Bell Atlantic – see attached letter of intent.

These assessments are extremely time and labor intensive, but are required to get a full picture of how these teams have made use of the technology available to them, and how that technology has supported or failed to support their efforts. Results of these assessments will be analyzed with an eye toward further developing and refining our hypotheses about virtual team success. The observations and analysis from this activity serve as an empirical foundation for all of the other activities; thus great care will be taken to maximize the opportunities for these empirical results to inform the *Handbook, Experiments,* and *Collaboration Assistant* activities. In particular, analyses will reveal better definitions of sustainability and optimal success criteria that can be used in *Experiments* (Activity 3).

Section 4. Activity 2: Assembling a Handbook of Knowledge Ecology

Overview. The research team will use the ongoing development of an on-line Handbook as a way of coordinating the research process and results during the entire three-year project, constituting a "collaboratory" for the distributed research team. This activity will include explicit theory building efforts, in both electronic and annual face-to-face workshops, among the research staff, building upon existing conceptual models (e.g. Ellis & Wainer, 1994). The Handbook is also intended as a vehicle for distributing the emerging theory to, and engaging in on-line dialogue with, other researchers, partner organizations (e.g. Boeing, World Bank, and Bell Atlantic), and the public.

Contents. The Handbook will contain many different kinds of information, all interlinked:

- Methods and instruments used in Activity 1 (*Observations*) and Activity 3 (*Experiments*);
- Issues, notes, ideas, questions, hypotheses, analyses, and theories about the data in those activities;
- On-line discussions (both asynchronous and synchronous) among the project researchers;
- Where appropriate[10], research reports and papers, both published and in-preparation;

[10] Appropriate to the politics of Web publication – some journals don't accept papers previously available on the Web, other journals publish only on the Web.

- Traditional handbook content, such as tips for virtual team leaders, pitfalls of prolonged collaborations, and how to facilitate asynchronous design sessions; and
- Information on the design, applications, and case studies of the Collaboration Assistant.

Some materials in the Handbook will be "backstage" information (Giddens, 1990), intended only for communication and coordination among members of this research team. In particular, it is important that researchers engaged in observational and empirical studies (*Observations* and *Experiments*) connect frequently with those working on the design of the Collaboration Assistant (*CA*), and the Handbook will provide an important medium for this cross-disciplinary collaboration. Other materials, in the "frontstage," will contain information intended for distribution beyond the research team.

In particular, the frontstage of the Handbook will be designed for two audiences: people who are on virtual teams, and people who are designing and developing virtual collaboration environments. For example, if the empirical work supported Hypothesis #11, the Handbook might suggest ways of orienting new project participants to an existing knowledge network, such as having a senior member of the project conduct an introductory tour of the space using an application sharing tool. For developers, the Handbook might recommend specific tools and practices that would facilitate new users getting oriented quickly and with confidence. If the hypothesis on the importance of acknowledgment is validated (# 4), the Handbook might suggest a feature for electronic "nodding" in the communication space; not necessarily agreement with the "speaker," but simple acknowledgment that what he or she said has been heard.

Media. The primary medium for the Handbook will be the WWW. Some more polished aspects, such as principles for sustainable virtual teams, will be publicly accessible. Connected to these pages will be interactive spaces designed primarily for the immediate "customers" of the research, such as employees at Boeing, the World Bank, and Bell Atlantic, where they will be able to ask questions about Handbook materials and provide interactive feedback about their experiences using those materials.

By the second and third years, there will be other ways to distribute and transfer the theory: published papers, a printed version of the Handbook, training courses, pamphlets, tip sheets, and so on, as appropriate.

Novel Aspects. One of the most intriguing aspects of the Handbook activity is that its goal is to successfully create and maintain a sustainable collaborative space for this research project[11]. As we identify potential principles of healthy virtual collaboration, we will build them into the Handbook (to the extent that current technology allows) and directly test them on ourselves. We plan initially to construct the Handbook using Microsoft's FrontPage® and The Soft Bicycle Company's Consensus @nyWARE® systems. As resources allow we anticipate extending the Handbook using a variety of other collaborative tools as well, to further our learning and theory building. An additional benefit of using a virtual collaboratory approach is that it will make it relatively easy to expand our research team, to invite specific research colleagues to review a theory, for example, or assist with the interpretation of data.

Section 5. Activity 3: Conducting Empirical Studies

Based on the results of the observations and analysis of the virtual teams examined in Activity 1, *Observations* (p. 10), systematic hypotheses will be derived for empirical testing with experimental or quasi-experimental designs. We expect that Activity 1 will expand our provisional list of hypotheses beyond the empirical scope of a single grant. Grounded, theoretically supportable hypotheses around the three primary focus areas of sustainability – coordination and process, object sharing, and orientation and cognitive overload – will be given priority for testing under this proposal. Ongoing organizational teams will be studied when practical, with laboratory research examining basic underlying processes, as well as to test underlying principles with greater precision. Greatest attention will be given to hypotheses showing greatest evidence of validity in Activity 1, *Observations*, and those with the greatest potential for feeding useful information into the design of the Collaboration Assistant.

These studies will involve better understanding the dysfunctionalities of virtual teams, as uncovered in Activity 1, as well as putting those in the context of the advantages face-to-face groups enjoy that are problematic for virtual collaborators. For example, recent work by Offermann and Eller (1998) suggests that face-to-face decision making teams outperform computer-mediated teams both in terms of decision quality and decision time. Yet some computer-mediated groups performed as well as the face-to-face groups, indicating that some virtual teams are able to overcome difficulties posed by their communication medium. Understanding the specific reasons for success and failure of computer-mediated teams, especially as revealed and articulated within a shared space, will help insure that more teams achieve success.

---

[11] Indeed, there have been many times in the creation of this proposal across three time zones when such a collaborative space would have been extremely useful!

For example, our experiences with virtual teams strongly suggest that the distribution of participation across the team may be related to team functioning (Hypothesis #5). A study by Barry and Stewart (1997) found that the proportion of extroverts in a team related to team effectiveness, with either too many or too few extroverts in a team being associated with poorer performance. Practical experiences with collaborative spaces suggest that, for some teams, participation dominance by a small proportion of the team may be dysfunctional, and that a more even spread of participation is preferable. These issues can be tested, either by examining the existing distributions of personality profiles in existing organizational teams in relation to their participation, process, and output, and/or systematically varying those compositions in laboratory-created teams. Consideration of the social roles played by team members can also be considered in projecting appropriate participation rates, with expertise in a particular area of the group's mission escalating expected inputs for a given person, and inexperience decreasing it.

Social psychologists have long known that the distribution of participation in groups is highly amenable to individual and group feedback, with automated feedback able to completely reverse initial levels of participation (Smith, 1972; Short et al., 1976). Appropriate research on thresholds of participation differences and the responses of members can lay the groundwork for the development of Collaboration Assistant functionality that would monitor individual member participation and give feedback as to relative participation both to individuals (i.e., one's own contributions vs. other members) or to the group or moderator. Such systems could also perform a leadership role by encouraging the involvement of more peripheral participants, or alerting the leader/facilitator of the group that such action may be appropriate.

Similar research on the potential advantages of programmed computer assistance in managing overly large data resources, cueing facilitators on the need for summarizing and focusing attention, or providing succinct group histories and current project status that allow new members to enter into the group more quickly are all reasonable. In all of these respects, we propose that systems assistance capacities can serve as "substitutes for leadership" (Kerr & Jermier, 1978; Howell et al., 1990), and will endeavor to demonstrate this through controlled experimentation.

Section 6. Activity 4: Constructing a prototype Collaboration Assistant

Overview. The Collaboration Assistant activity (abbreviated CA) will conceptualize, design, and implement a software prototype assistant suitable for use by pilot test groups. It will actively provide information, analyses, and advice to moderators and collaborators within a distributed, asynchronous virtual collaboration. The design of this assistant will be based upon, and represent a pragmatic embodiment of, our working hypotheses of sustainable knowledge networks. Thus, use and evaluation of this CA will constitute a strong form of hypothesis testing.

We conjecture that most of the sustainability factors exposed in the empirical studies (both *Observations* and *Experiments)* will be concerned with *collaboration-specific* phenomena, such as the distribution of knowledge in the group, the various roles the members take on, the conventions and practices of sharing knowledge, the rewards for being explicit, etc. These phenomena generally remain concealed by the current document centric desktop applications. Thus, we expect the CA will be predominantly concerned with collaboration-specific parameters such as amount of access to shared objects, awareness of other's presence and activity in the shared space, and the other factors suggested in the description of Activity 1, *Observations* (p. 6).

CA Functionality. The functionality of the Collaboration Assistant will be partitioned into two categories. The first category is "Shared Object Coordination," which is concerned with gathering, manipulating, and disseminating artifact information. As much as possible we will avoid creating application-specific features; rather, we will treat a complex shared object within a virtual collaborative space as a "black box" with links going into and out of it. For example, the MS Windows environment now makes it relatively easy to create web and OLE links between documents (artifacts), and the CA might track the creation and evolution of those links. The web of links could also be analyzed by the CA for patterns of traversal and use by different members of the team, perhaps leading to "usefulness" metrics on shared objects (Hypotheses #8 and #9).

The second category of functionality is "Process Coordination," which is concerned with gathering, manipulating, and disseminating information about activities, participants, interactions, and formal/informal communications. For example, the CA might have features for orienting virtual team members in the knowledge network, such as a graphic view of all of the objects and interactions in the space color-coded for how recently each object was created or changed. Previous studies on the utility of graphical views of communication spaces (Begeman 1990) have shown that such a view would allow team members to tell at a glance where the "hot spots" were in the shared space (Hypothesis #10). The scheduling and managing of work tasks (e.g. workflow capabilities) is another one of its important functions concerning process (#2). We understand that much of the team knowledge is passed along by informal interactions. Thus, it is useful for the CA to be cognizant of which participants frequently chat with whom. We thus intend that the CA will have built-in dimensions of awareness, including some degree of task awareness, organizational awareness, and social awareness.

When groups are not co-located, the issues of team cohesion and shared context tend to be particularly aggravating. We believe that the Collaboration Assistant can include features that may significantly help to address these issues. For example, within a large team, it is useful simply to know who is currently working on the artifacts, and who has been contributing or non-contributing. This is one specific example of our working hypothesis that "cues about the activities of others will make the system more sustainable" (Hypothesis #1). This information can be presented to participants specifically (by name) or statistically by the CA. The CA can maintain certain barometers of team viability. It is thus possible that the CA can help team cohesion by noting team deficiencies, monitoring group member satisfaction, or suggesting helpful actions to appropriate team members or to those playing a mediator role.

Design Process. The CA will be conceptualized, designed, and implemented using techniques of iterative team design (Norman and Draper, 1986) and participatory design (Schuler & Namioka, 1992; Kyng, 1991). It will encode and utilize aspects of our theory of sustainable knowledge networks. The design team will perform conceptualization work, as well as producing a specific implementation; this work will start in the first year.

We envision that the CA will connect to and augment existing virtual collaboration systems (e.g. Lotus Notes™, QuestMap™, and Consensus @nyWARE®). Within the conceptualization work, we propose to first explore taxonomies of collaboration primitives that are common among such collaborative systems. We will investigate possibilities of a collaboration language or protocol (e.g. Lai et al, 1988) that enables our prototype, through a standard interface (see Figure 1), to communicate with different groupware systems.

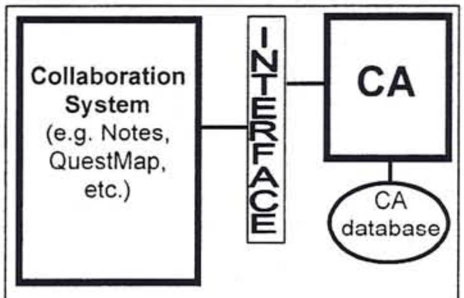

Figure 1: Architecture of the CA

Within the implementation work, there will be an initial (potentially throwaway) prototype designed and programmed by a team of implementers and users at the University of Colorado, followed by one or more enhanced prototypes developed at GDSS. The initial prototype will be constructed within a three-month coding period, and will be useful to show functionality and interface possibilities. Comments from experts and potential user communities will be elicited, and then one or more redesign and re-implementation efforts will take place during the second year of the research work. This will produce a prototype that is suitable for use by external pilot test groups. Testing, evaluation, and some continued development will proceed during the third year of the research.

Section 7. Activity 5: Evaluating and Testing

The proposed research will yield two technologies to be evaluated in Activity 5, a Handbook representing an ecological theory of sustainable knowledge networks including guidance for distributed teams, and a prototype Collaboration Assistant that advises participants in distributed teams.

The evaluation strategies for these technologies will differ because their purposes differ. The Handbook is intended to serve as a resource for researchers, developers, and end users, and it will evolve as research and development unfold. Its value will be measured by how it is used.

The theory in the Handbook will be tested in Activity 3. The Handbook itself will be accessible on the web where we will monitor its use. We will track the frequency of access to different parts of the Handbook, thereby measuring the degree of interest in those facets of the theory and guidance represented in those parts. We will also solicit input from visitors to the Handbook. They may comment on parts of the theory, provide anecdotes or data supporting or refuting the theory, and evaluate the Handbook's utility.

The Collaboration Assistant will be evaluated in field experiments. Because of its prototype status, we will not deploy it to teams engaged in mission-critical work. Instead, we will recruit distributed teams that can accommodate the uncertainties and difficulties experienced when using prototypes. Examples of potential evaluation teams include a conference organizing committee or a team concerned with computing support. These are both examples of distributed teams composed of people with computing expertise. If the Collaboration Assistant proves successful with such teams, we will consider teams that are less receptive to new technologies.

Dissemination of Results.

Results generated in the course of this research will be disseminated through the traditional channels of reports, white papers, conference papers, and journal articles. In addition, we plan to use the Handbook (Activity 3) as a major interactive forum for the on-going dissemination of research ideas, theories, and results. Finally, the Collaboration Assistant prototype will be available via Internet download to the public.

Institutional Space and Equipment Commitments.

The University of Colorado Department of Computer Science is an active research department with 25 full time faculty members. The department is closely aligned with the Institute for Cognitive Science, and provides all of its faculty with up-to-date computing equipment in their offices and homes. Since 1985, the department has been constantly supported by NSF Computing Infrastructure grants (NSF CER grants followed by NSF RI grants), along with being the recipient of numerous industrial grants and gifts. Department equipment includes modern computing equipment on every employee's desk, labs full of PCs, MACs, Sun workstations, mainframes, etc. The most recent RI program grant to the department focussed upon providing networking and multimedia distributed computing facilities using high speed network technology, fiber optics, and ATM switches. The University of Colorado is committed to excellence in research. Thus, computing equipment, infrastructure, and other resources of the University will be made available to perform this research.

Group Decision Support Systems, Inc (GDSS) is an consulting and facilitation firm specializing in the use of state-of-the-art synchronous and asynchronous communication tools for group planning and decision making. GDSS's facilities in Georgetown in Washington, DC, offer a modern high-tech environment for research into collaboration. These facilities will be completely available to Dr. Jeff Conklin in support of this research project. In addition, the face-to-face conferences held by the research team will be hosted by GDSS in its meeting room facilities.

George Washington University currently provides Lynn R. Offermann with a private office complete with a Dell Pentium PC with ISN line to the GW mainframe, non-graphics access to the Web through Lynx, and a DeskJet printer. Ethernet connection to the Web is available down the hall, with current installation plans for full desktop internet access for her building scheduled for Fall 1998. She also shares a research room with PC and printer on the fourth floor, currently housing one of her doctoral students. Research space on the first floor is available to all faculty on an as-needed, scheduled basis; this semester she has had use of 5 rooms for a pilot study on distributed team decision making. Conference rooms are also available on a reservations basis.

Boeing's Applied Research and Technology (AR&T) organization is a center for advanced applied research in computing application development and has facilities to support that mission. This includes a variety of UNIX servers, including HP, IBM, SGI, and Sun that support over 600 users and over 550 workstations. The workstation mix includes over 250 UNIX machines (Sun, SGI, HP, Dec, Alpha, IBM), roughly 220 PCs and an assortment of miscellaneous machines. An ATM backbone is also part of the environment which includes Internet access and access to Boeing's mainframe computers. AR&T is located on a 90 acre site in Bellevue , Washington that houses the Boeing Shared Services Group Headquarters and Boeing's largest data center. The entire complex consists of over 1,000,000 square feet of space devoted to computing support and applied research activities.

In support of its commitment to this proposal as discussed in the Letter of Intent, Bell Atlantic will provide sufficient work space and computing equipment to support Bell Atlantic personnel (specifically Albert M. Selvin) working on this project, as well as conference rooms in Bell Atlantic work locations as needed, high-speed Internet access, local- and wide-area-networking, remote access, speakerphones and telephone service to support teleconferencing, and other technologies as necessary. Bell Atlantic is currently both researching and developing distributed shared space technology and methods, so its participation and commitment to space and equipment for this proposal will result in greater leverage for Bell Atlantic's investment in this area.

Performance Goals

The primary deliverable of this research project is a **theory of sustainable knowledge networks**. The deliverable comes in two forms. The first is an written "theory handbook," a textual and web-based compilation of data, hypotheses, principles, and annotations; it documents the research findings and emerging theory at many levels, from hand-on practical tips for virtual team leaders, to published papers on social and architectural principles of sustainability, to interview transcripts and other essentially raw data forms. The second form is a prototype implementation of the theory as a "Collaboration Assistant," a system that provides warnings, data views, analysis, and advice about counter-sustainability conditions emerging in a virtual team's knowledge network.

The program to create this theory has five activities that span three years (see the Project Description for details of the activities):

1.	*Observations*	Observational studies of knowledge networks.	Year 01
2.	*Handbook*	The written vehicle for the theory.	
3.	*Experiments*	Controlled experiments to explore selected issues in depth.	Year 02
4.	*Collaboration Assistant (CA)*	The executable vehicle for the theory: a program.	
5.	*Evaluation*	Tests of the Handbook and Collaboration Assistant in pilot groups.	Year 03

Performance goals for the three years of the program are as follows.

Year 01. Our initial focus will be on the *Observations* and *Handbook* activities. The goal of *Observations* in Year 01 is to identify at least 4 virtual teams (both "live" and "deceased") within our partner corporations (World Bank and Boeing), gather data about their collaborations, analyze the data in terms of sustainability principles and factors, and document the analysis. The goal of *Handbook* in Year 01 is to construct a web-based "research collaboratory" using off-the-shelf tools (e.g. Consensus @nyWARE or Lotus Notes) and to assemble within it (as much as feasible) the linked network of data and observations from *Observations*. *Handbook* will include a continuous theory-building effort that includes all research staff in this program. Finally, an initial design effort will happen in Year 01 on the *Experiments,* and a first "throwaway" prototype of the *Collaboration Assistant (CA)* produced.

Year 02. In Year 02 our emphasis will be on *Experiments* and *CA*, but reduced levels of work will also continue in *Observations* and *Handbook.* In particular, we anticipate that the elements of an ecological theory – hypotheses and principles – will come into clear focus this second year. The goal of *Experiments* in Year 02 is to test specific hypotheses created in *Observations* and *Handbook.* The goal of *CA* in Year 02 is to take the rapid prototype from Year 01 and "harden it" into a working prototype, suitable for use by pilot test groups. The goal of *Handbook* in Year 02 is to develop a short course (and associated training materials) in moderating virtual teams. We anticipate presenting several papers at various conferences during this year (e.g. CSCW).

Year 03. The final year of the project focuses on *Evaluation,* both of the *Handbook* materials as well as of the prototype *CA.* The goal of *Evaluation* is, using pilot groups from various organizations (not restricted to Boeing and World Bank), to gain clear feedback about advantages and breakdowns encountered in real work situations using both *Handbook* and the *CA* prototype. This evaluative material and its analysis, as well as the required NSF reports, will be the major work products from Year 3.

Project Management Plan

A distributed research team such as one proposed here will require that special attention be given to creating and maximizing good coordination of studies and experiments, a high level of shared understanding about goals and issues, and exceptional sharing and collaboration about interim research results. Dr. Conklin will augment Dr. Ellis's PI role by serving as the overall Project Manager. Drs. Ellis and Conklin worked together closely in the Software Technology Program at MCC and have a very high level of trust and respect for each other and an excellent working relationship.

In addition, each of the five research activities will have its own Lead Researcher (who is also either PI, co-PI, or a Senior Researcher), as shown in Table 2.

	Activity	Leader	Center
1	*Observations*	Conklin	GDSS
2	*Handbook*	Conklin	GDSS
3	*Experiments*	Offermann	George Washington University
4	*CA*	Ellis	University of Colorado, Boulder
5	*Testing*	Poltrock	Boeing

Table 2: Leaders of the five research Activities

Each Lead Researcher will oversee the coordination of research efforts and the sharing of information, electronic and otherwise. Bell Atlantic is donating 25% of Al Selvin's time to this research effort, and he will be participating in both *Observations* (analysis of Bell Atlantic software project data) and *CA* (participating in design and testing).

In order to maximize cross-fertilization of ideas, hypotheses, and results among the Centers, we will use web-based collaboration tools, monthly conference calls, and a mid-year face-to-face conference (on the East Coast) to facilitate the research collaboration. For four months during Year 01 the *Observations* teams will have an additional weekly one-hour conference call for active sharing of ideas, issues, and hypotheses. Drs. Ellis and Conklin will arrange for additional face-to-face meetings twice a year for administration and management oversight coordination.

The web-based *Handbook* will serve not only as a medium of distribution of information and knowledge to a wider audience, but will also serve as a virtual collaboratory, a collaboration and coordination mechanism among all project researchers across the four centers, plus consultants and other members of our extended network. We plan to use the web-based Consensus @nyWARE® system for this interactive dimension of the research (Consensus @nyWARE® is being supplied by GDSS as part of its cost sharing).

The distribution of funds among the four Centers (the University of Colorado at Boulder and the three subcontractors) is described in the budgets. The chart at right describes the approximate distribution of resources among the five activities. As the chart shows, all activities are funded in each year (except Testing, which is only funded in Year 03), but different years emphasize different activities. The first major activity of the research project will be a series of planning meetings involving the entire research staff, via video teleconference, to plan the detailed milestones and accountabilities of the first year.

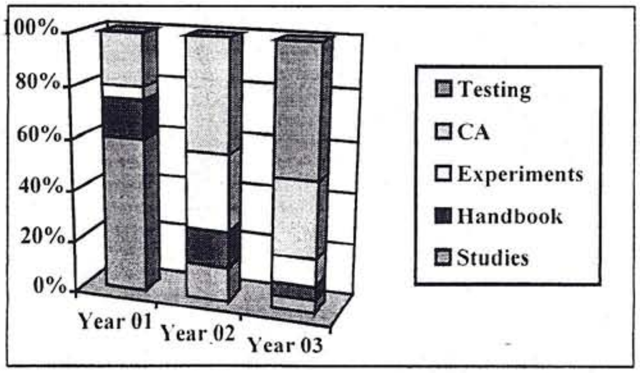

Letters of Commitment

Hard copies of letters of commitment from the World Bank, Boeing, and Bell Atlantic are being sent along with the Cover Page.

Facilities, Equipment, and other Resources

Laboratory

The Psychology department at George Washington University has several observation rooms on the first floor that are currently available for research on a reservations basis. If the proposal is funded, on of these would need to be dedicated to the grant, with requested equipment installed, for the life of the grant. In addition, networked PC stations would need to be available. This could be accomplished through sharing use of the Psychology laboratory space (around class times) or using multiple research rooms (as is being done now during pilot testing).

Computer

The University of Colorado Computer Science Department equipment includes modern computing equipment on every employee's desk, labs full of PCs, MACs, Sun workstations, mainframes, etc. The most recent RI program grant to the department focussed upon providing networking and multimedia distributed computing facilities using high speed network technology, fiber optics, and ATM switches.

Group Decision Support Systems, Boeing, and Bell Atlantic all have similar modern computing facilities. In the Psychology department at George Washington University Dr. Offermann has a Dell Pentium computer with an ISN line.

Office

All the research staff listed on this proposal have their own offices and full access to modern office facilities, e.g. telephones with speaker phone capability, voice mail and email, copy and fax machines, administrative assistance, and meeting rooms. In addition to her own office, Dr. Offermann shares a research room on the 4th floor with another professor, used largely to house two research assistants, store data and materials, and occasionally to run subjects completing paper-and-pencil measures. This research space would be needed to house Research Assistants under this request. If full funding is granted, some additional space would be required to house a total of 4 students in year 01, and 5 students plus one research coordinator in years 02-03.

Other

The World Bank has a newly installed observation laboratory that will allow researchers to observe real working groups engaged in collaborative work. The laboratory includes a two-way mirror to an observation room, whiteboards, a VCR and TV monitor, a computer, and tables and chairs that can be rearranged depending on the group's work needs. In the near future, it will also have audio and video taping equipment. One of the co-PI's (Offermann) is already engaged in activities which use this state-of-the-art research facility, and fully expects this facility to be available in support of the proposed research activity.

Major Equipment

The Psychology department at George Washington University has one video-observation room set up and available for use (currently being used in pilot testing) on the first floor. This room has full capacity for videotaping, with intercom capability. The control room also supports video editing. A similar facility to be available on a dedicated basis would be set up with funds requested under this proposal.

Other Resources

At GDSS, Dr. Conklin has a strong working relationship with a staff of 25 first-rate organizational consultants and group facilitators, as well as professional computing and graphics services. A sister software company, The Soft Bicycle Company, is in the same building and has skilled programmers and successful products that overlap heavily with the proposed research. The programming effort for the Collaboration Assistant would be performed by subcontracted programming staff housed in and supported by The Soft Bicycle Company, thus assuring timely and professional code production.

At Boeing Dr. Poltrock is currently the director of the Workgroup Integration Technology project within Boeing's Applied Research and Technology department, where he is currently supervising three research projects that overlap strongly with the proposed research. Dr. Grudin is currently a member of this research staff. Thus, Dr. Poltrock's staff, laboratory, and research momentum would be directly and immediately applicable to the sustainability research proposed here.

At the University of Colorado at Boulder, Dr. Ellis is a Professor in the computer science department, faculty member of the Institute for Cognitive Science, and Co-Director of the Collaboration Technology Research Group. Dr. Ellis thus has an extraordinary resource of talented and experienced colleagues, already engaged in related research (e.g. Dr. Gerhard Fischer's NSF-supported work on collaborative technology and organizational memory), as well as a pool of high quality graduate students to draw from in performing this research.

Dr. Offermann has been conducting on-going research on distributed team processes, using observation laboratories available at GW. The project she has been running has involved three research assistants. Like Dr. Ellis, Dr. Offermann is access to high quality doctoral students, some of whom would doubtless find rich opportunities in the proposed research.

Drs. Ellis and Conklin worked together for seven years in the Software Technology Program at Microelectronics and Computer Technology Corporation (MCC). They have a strong working relationship and a high level of trust and respect for each other, and this relationship is a strong asset for the management of this research project. Similarly, Drs. Grudin and Poltrock worked together for many years in the Human Interface Program at MCC, and have continued to work closely together since; for example, they are co-chairing CSCW '98, the ACM-sponsored major conference in the field. Also, there was a strong connection between the Software Technology and Human Interface Programs at MCC, thus Drs. Ellis, Conklin, Grudin, and Poltrock have associated professionally for over 10 years. Dr. Conklin also has worked with Mr. Selvin (at Bell Atlantic) for nearly a decade. Thus, in addition to their individual credentials, many of the lead scientists on this team have a strong background in collaborative research together in the CSCW field. During the proposal writing process Dr. Offermann became a strong member of the team, adding her expertise in experimental psychology to the mix.