

Augmenting Organizational Memory: A Field Study of Answer Garden

To appear in *ACM Transactions on Information Systems*

Contact:

Mark Ackerman
Department of Information and Computer Science
Computer Science 444
University of California, Irvine
Irvine, CA 92717

714-824-7355

ackerman@uci.edu

Classification: C.2.4. (distributed applications); H.1.2 (user/machine systems); H.3 (information storage and retrieval), H.3.3 (information storage and retrieval), H.3.4 (systems and software/information networks/question-answering systems); H.4.3 (communications applications/bulletin boards, electronic mail); H.5.1 (hypertext navigation); H.5.2 (user interfaces/training, help, and documentation -- evaluation/methodology; H.5.3 (group and organizational interfaces); I.7.2 (hypertext/hypermedia). K.4.3. (organizational impacts)

Additional keywords: organizational memory, computer-supported cooperative work, CSCW, field studies, group memory, collective memory, community memory

Augmenting Organizational Memory: A Field Study of Answer Garden

Mark S. Ackerman

University of California, Irvine

Abstract

A growing concern for organizations and groups has been to augment their knowledge and expertise. One such augmentation is to provide an organizational memory, some record of the organization's knowledge. However, relatively little is known about how computer systems might enhance organizational, group, or community memory.

This paper presents Answer Garden, a system for growing organizational memory. The paper describes the system and its underlying implementation. It then presents findings from a field study of Answer Garden. The paper discusses the usage data and qualitative evaluations from the field study, and then draws a set of lessons for next-generation organizational memory systems.

This work was done while the author was at the MIT Center for Coordination Science. This research was supported, in part, by research grants from the X Consortium, Digital Equipment Corporation, the National Science Foundation (IRI-8903034), and the MIT International Financial Services Research Center.

Author's address: Department of Information and Computer Science, University of California, Irvine, Irvine, CA 92717. (ackerman@ics.uci.edu; <http://www.ics.uci.edu/CORPS/ackerman.html>)

1. Introduction

A recent but growing interest for organizations and groups has been to augment and manage their knowledge and expertise. With an impetus from layoffs, down-sizing, and internationalization of personnel, the search for new ways to access, maintain, and promote the organization's intellectual assets has become vital for many organizations. One possibility is to provide an organizational memory, some record of the organization's knowledge. To do this, we would like to know how much computer systems can be used to augment or supplement existing learning and knowledge mechanisms within organizations, groups, and communities. Field studies, within the context of use, are an important part of understanding how computer systems can provide these organizational mechanisms.

This paper presents findings from a field study of an organizational memory system, the Answer Garden. In short, the study found that augmenting organizational memory was possible, but it also uncovered a number of interesting issues and problematic design assumptions. This paper describes these as the basis for future work on organizational and group memories.

To provide the necessary background for the field study, the paper first briefly discusses the topics of organizational memory and organizational information seeking. Second, the paper describes Answer Garden and the hypermedia system that underlies it. After these sections, the paper then presents the field study and its results. The final section of the paper draws some conclusions and lessons from the field study for Answer Garden and similar memory systems.

2. Augmenting the Expertise of an Organization

Stewart [1991] raised the possibility that the organizational issues of finding the right expert, growing knowledge and memory, and managing intellectual property may become increasingly important. Properly managing its intellectual resources might enable the organization to function

more effectively and to prosper. Simply put, the likelihood that a perceived problem will find an organizational solution increases as the organization has additional resources to find previously created solutions or to create new solutions [Cohen, March, and Olsen 1972].

2.1. *Organizational Memory*

From the macro perspective of the organization, one method of managing its intellectual resources is to augment its organizational memory. It is beyond the scope of this paper to review the literature on organizational memory. In general, however, an organization of people should retain some knowledge of its past efforts and environmental conditions.

Organizational memory is therefore intrinsically linked with organizational learning [Cohen and Sproull 1996; Bannon, Michelis, and Soergaard 1996; Bannon, Michelis, and Soergaard 1996]. If an organization learns, then the result should be available later [Levitt and March 1988; Duncan and Weiss 1979]. Further, the contents of an organizational memory can be made available for further individual and organizational learning.¹ In this view, "organizational memory" is organizational knowledge with persistence.

A standard connotation of organizational memory, i.e., that material which can be captured in a written record, is only one form of organizational memory. In Walsh and Ungson [1991], organizational memory can be retained in six places: individuals, organizational culture, organizational transformations, organizational structures, organizational ecology, and external archives. (External archives are data external to the organization, not its own archives.) It should be

¹This is not to say, however, that memory is uniformly positive in nature. Memory merely serves as a potential source of organizational learning. Organizations can be encumbered by their past, and organizational forgetting can be critical [Sandoe and Olfman 1992]. Furthermore, organizational memory requires people to bring an active recontextualization to the remembering [Ackerman 1996; Bannon and Kutti 1996], an effort that may fail.

noted that individuals are a prime location for retention of the organization's knowledge. This list, however, should be expanded to include internal information repositories such as corporate manuals, databases, filing systems, and even stories [Yates 1989; Morgan and Root 1979; Orr 1986].

What does it mean for knowledge to persist within an organizational setting? March and Simon [1958] argued that organizations are concerned with achieving their goals in a manner that minimizes the drain on limited resources. If this is true, the goal-driven nature suggests that an organizational memory mechanism that is immediately tied to the on-going processes and considerations of an organization will be most important and useful to that organization. Furthermore, one might expect that organizational memory mechanisms most often are employed for recent events and outcomes because the situational context does not need to be reconstructed and substantially reinterpreted to use the information. In other words, we want to minimize the upstream costs of an organizational memory mechanism and make the downstream payoffs clear [Grudin 1989].

Thus, information technology can support organizational memory in two ways, either by making recorded knowledge retrievable or by making individuals with knowledge accessible. To augment organizational memory in this manner, an information system would need to incorporate elements of information databases and communication systems. Furthermore, we might expect that such an organizational memory system would be most useful if it were centered around a current organizational activity.

2.2. *Information seeking in an organizational setting*

Viewed from the perspective of the organizational member, rather than the organization as a whole, organizational memory is invisible or at least muted and hazy. At a micro-level, the issue is one of information seeking by some organizational member within an organizational context. As such, information seeking can be considered as the process of finding the right "piece" of organizational memory. The goal, then, is to decrease the effort involved in information seeking in

cases where the required information is not known to the individual involved, or to eliminate the information seeking in cases where the organization has redundant efforts.

In his seminal study of engineers' information seeking behavior, Allen [1977] found that their major source of information was direct contact and communication with colleagues. He determined that performance was related to the presence of gatekeepers, engineers who maintained contacts outside the group. Much of an organization's communication flow (for R&D engineers) is channeled through these gatekeepers, who track both the literature and the organizational members who know about various subjects.

Possibilities for computer support are offered in Gerstberger and Allen [1968]. This study noted some of the reasons why engineers would *not* go to colleagues, and would instead use other information channels. In their study of 19 engineers, Gerstberger and Allen found that the engineers chose not to go to the channel of the highest quality for technical information, but rather to go to the channel of highest accessibility (i.e., lowest psychological cost). Allen [1977] argued that the psychological cost was in the potential lack of reciprocity between giving and obtaining information and in the status implications of admitting ignorance.

As Allen pointed out, information seeking is not a straight-forward information transfer. For people with a career anchor in technical expertise [Schein 1978], each information seeking interaction is a double-edged sword. It provides the possibility for consorting with experts (or relative experts) and thus gaining prestige and status. On the other side, it also provides the possibility of being seen as incompetent, a possibility either so frightening or so likely that most engineers prefer using the documentation, other technical literature, or friends. In order to facilitate information seeking in the communication system, one should decrease the status implications and need for reciprocity. Any system attempting to augment information seeking (and use the organizational members as part of an organizational memory system) will need to consider this status interaction.

3. Answer Garden

Answer Garden [Ackerman and Malone 1990; Ackerman 1993] supports organizational memory in two ways: by making recorded knowledge retrievable and by making individuals with knowledge accessible. In the standard configuration of Answer Garden, users seek answers to commonly asked questions through a set of diagnostic questions or other information retrieval mechanisms.

A typical Answer Garden session is shown in Figures 1 through 4. Figure 1 shows the user's initial view of Answer Garden in its X Window System environment. This first node describes the particular information database being used; in this situation, the information database is about the X Window System.² The warning to the user about her being monitored is required by MIT rules.

There are several different ways to find information in Answer Garden. One way is to travel down a set of diagnostic questions that lead the user to the information sought (Figure 2). It is similar to the system playing the game "Twenty Questions." The user traverses the question nodes, selecting the appropriate button with his mouse. Figure 2 shows the screen after the user has selected several questions. At the end of these diagnostic questions, the answer is a particular piece of information. In this example, the user is looking for the name of an X graphics program.

In this particular information database, answers can either be specific pieces of information (as in Figure 2), or they can be more general groupings of opinions, tutorials, and code examples. However, there is no restriction on the type of information or display for the node. The end node could also be a picture of a person, a program segment to dynamically create or retrieve the

²The information database about X is only a sample database. Two computer science departments are building databases about their classes and requirements, and various organizations have built information databases about system administration. There are also databases on a hardware product and on a software system.

information, a set of questions and answers, an advertisement for a distributed service, a tutorial, or anything else. (Answer Garden is quite extensible by application programmers.)

Figure 3 shows a second way of searching for the information. More experienced users may wish to go directly to the information, and they can use the grapher view of the same diagnostic questions. (In some Answer Garden applications, the grapher view is a reference view, and therefore offers an alternative to the diagnostic questions.) The directed acyclic graph of diagnostic questions is projected into a tree to ease navigation for the user. The "..." indicator in some tree items indicates there is a subtree as well. Figure 3 shows a situation where the user has selected the Answer Garden tree, and then selected a specific answer. (It is the same answer he would have selected from the diagnostic question series.) Alternatively, from the "Other" menu on the beginning node, the user can select additional information retrieval engines. Currently, Answer Garden has a free text retrieval engine and a keyword engine; we have also experimented with an adaptive retrieval engine. The structure of Answer Garden makes it possible to include almost any information retrieval engine for use.

In many cases, however, the information will not be in the information database. If an answer is not found or is incomplete (or if the user becomes confused or lost), the user may ask the question through the system (Figure 4). The "I'm Unhappy" button, present on each node, pops up a mailer for the user.³ The user then asks his question, and Answer Garden routes the question to an appropriate human expert.⁴ The human expert then answers him directly. If the question is a

³Answer Garden also has synchronous communication capabilities using the Zephyr communications package [DellaFera et al. 1988]. In an MIT installation, Answer Garden attempted to contact the expert synchronously; if that failed, the expert was sent the same message via e-mail. Telephone and voice mail support are possible extensions, again blurring the line between asynchronous and synchronous communication in the system.

⁴The expert normally answers via electronic mail. Usually, the user is anonymous. In that case, the

common one, the expert can also insert the answer into the database. As well, the expert can add any diagnostic questions that she feels might be necessary. In the field study described below, the message was sent to the expert anonymously to reduce status implications [Sproull and Kiesler 1991].

question and the expert's answer go through an intermediate layer. If the user is not anonymous, the expert could also phone or even show up in person.

The header for the electronic mail message is replaced before being sent with the electronic mail address of the correct expert or set of experts, a user history so the expert can determine where the user has been and what she has seen, as well as other, miscellaneous information.

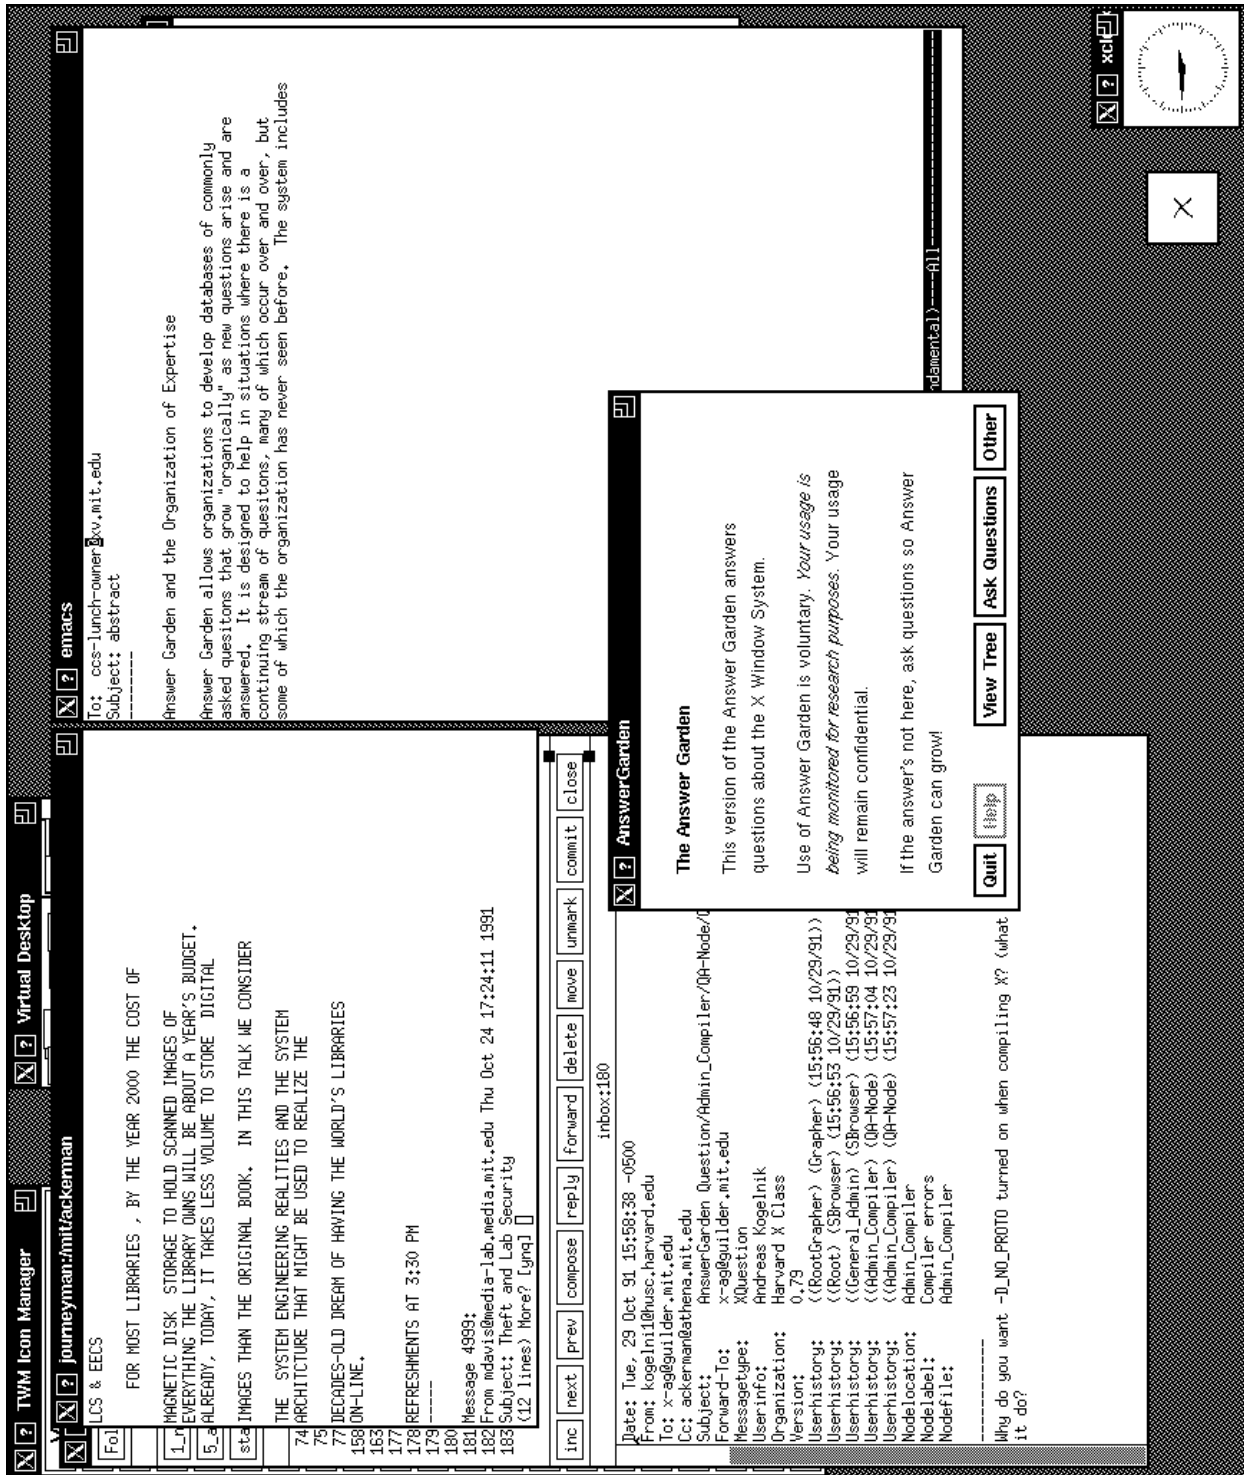


Figure 1: Initial screen for the Answer Garden application

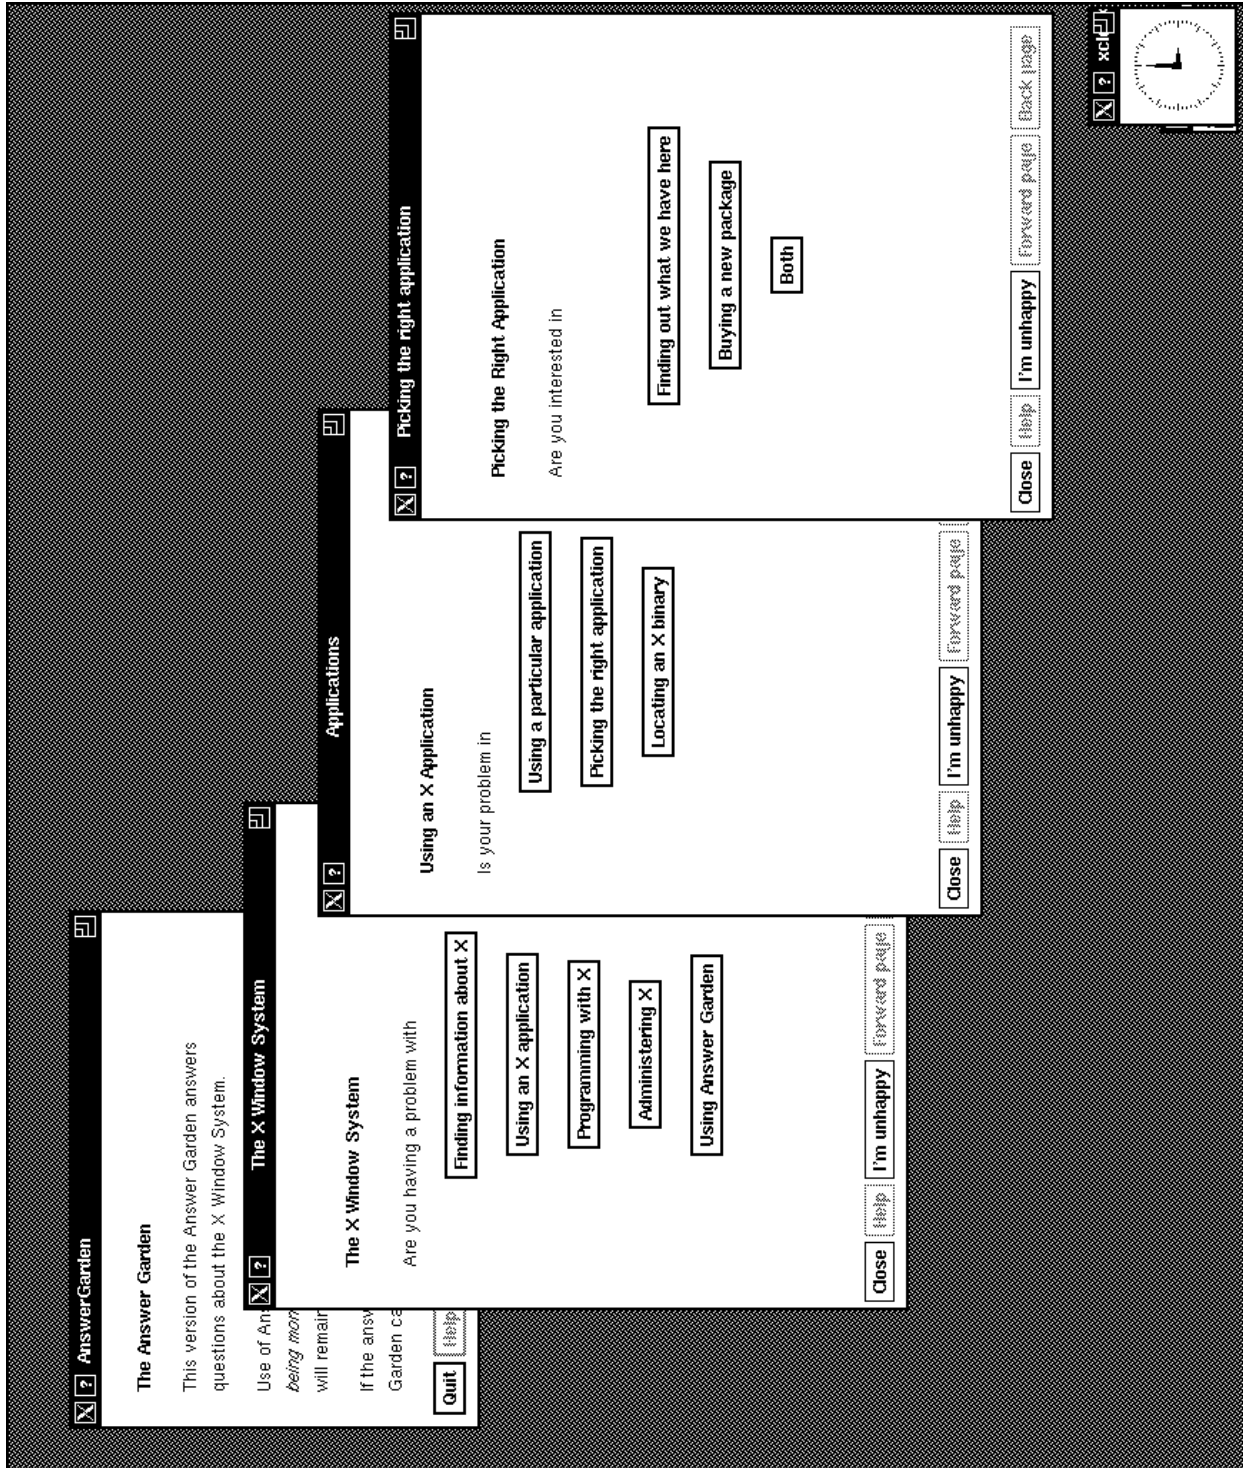


Figure 2: Diagnostic questions in an Answer Garden session

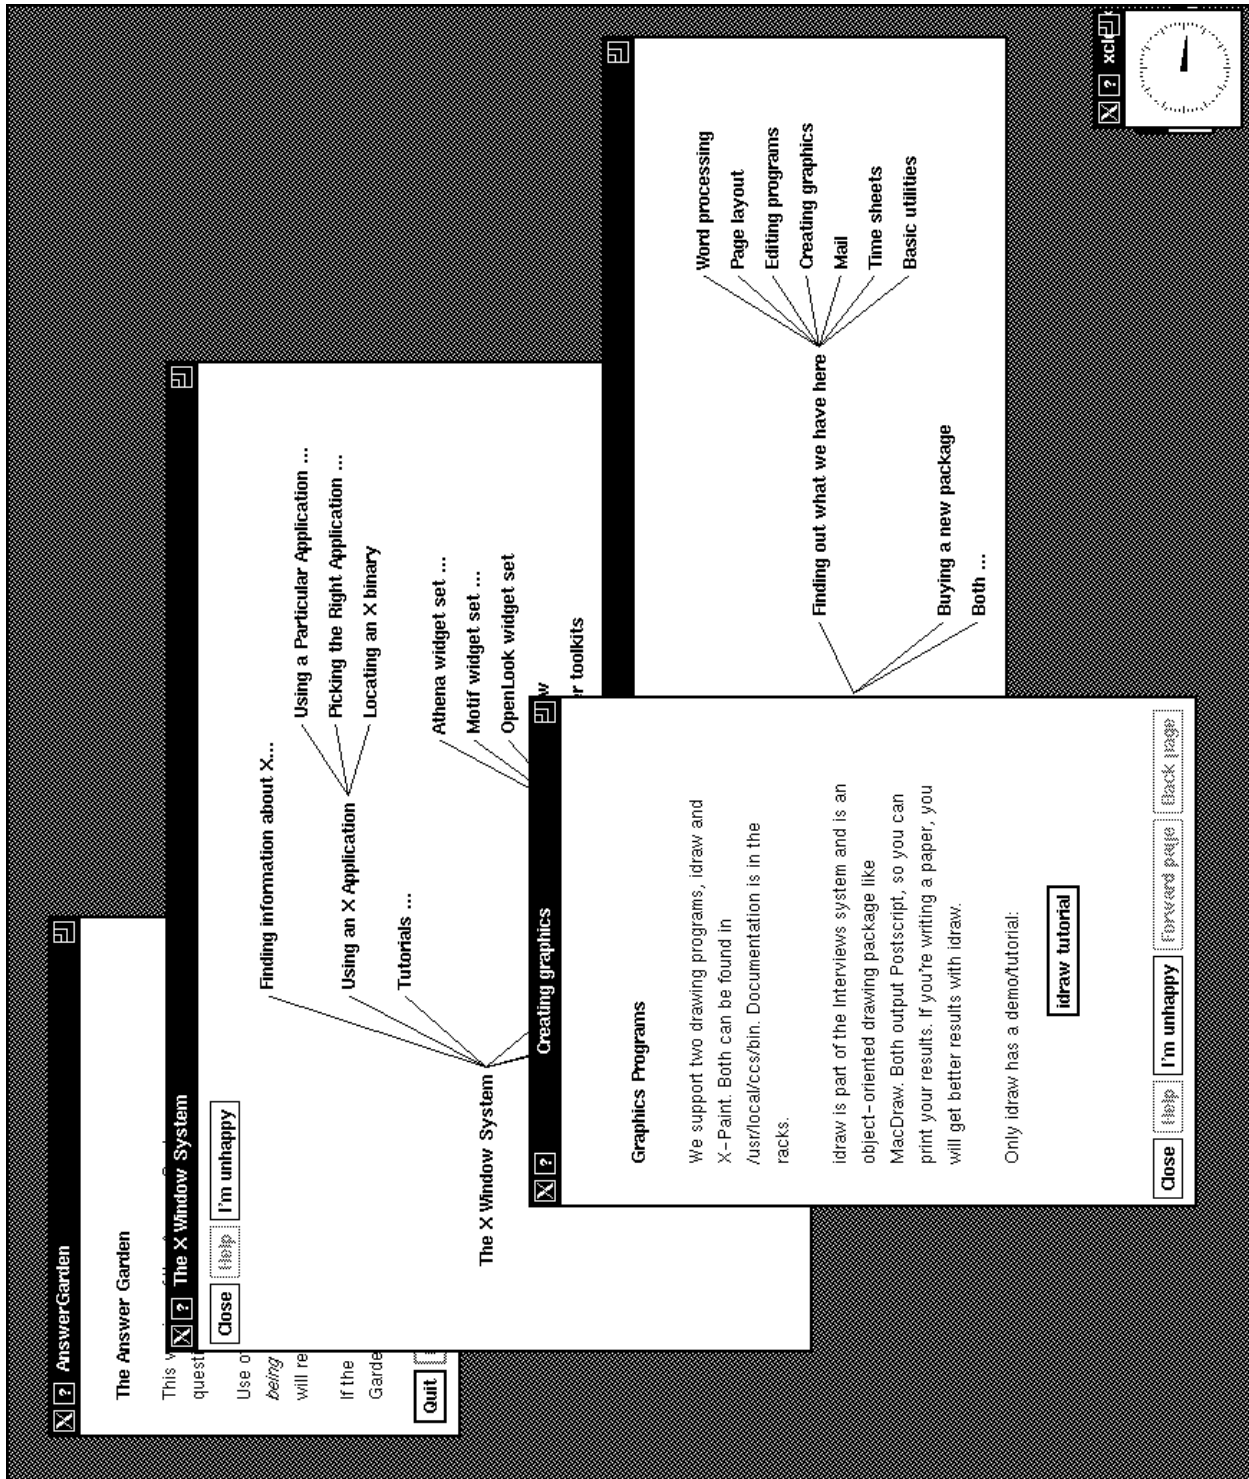


Figure 3: The diagnostic tree in Answer Garden

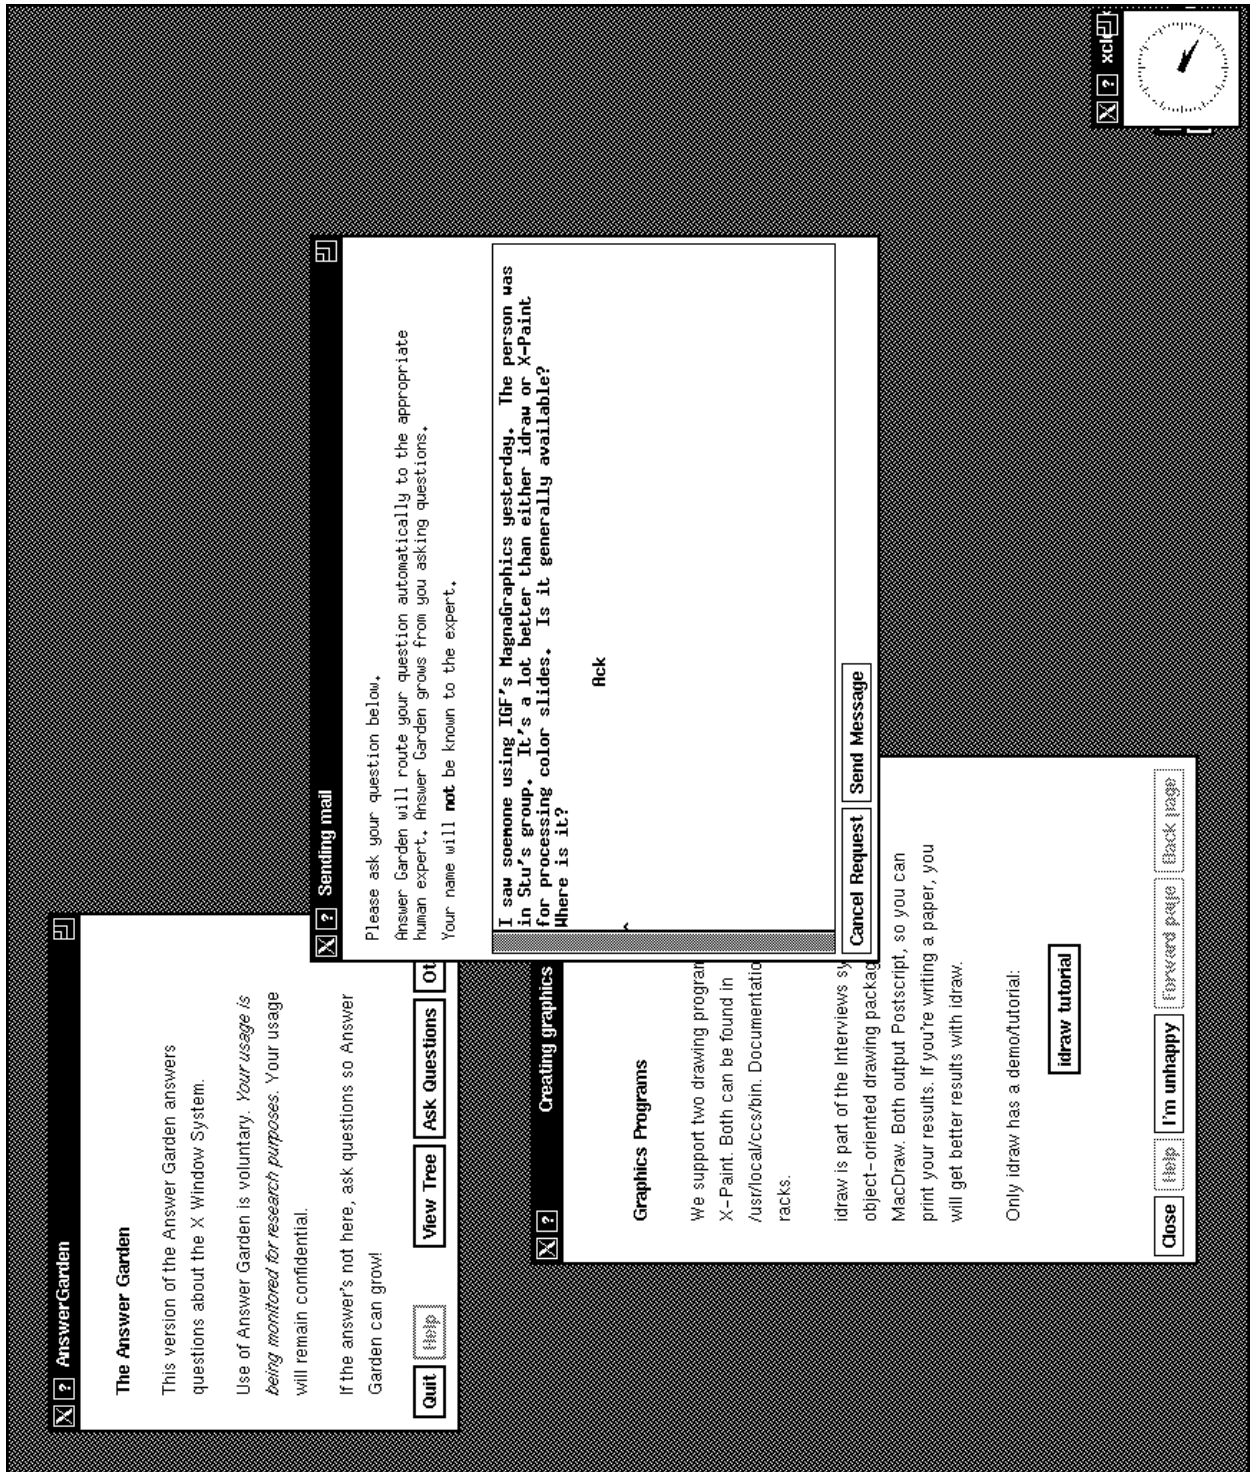


Figure 4: Mailing a question in an Answer Garden session

Answer Garden thus provides a mechanism for growing a body of information over time. (Hence the name.) Both experts and users perform their normal duties, and Answer Garden provides incentives for both the experts and the users. While users have to first browse the database to find the answer, they gain the ability to find the right person of whom to ask their question. This is often a considerable problem in many organizations. While the experts must structure the database, they gain the capability to rid themselves of the most commonly asked questions. As a whole, the organization, group, or community gains a corpus of information, an organizational memory.

The expert then answers the user through electronic mail, and if the question is a common one, the expert can insert the question and its answer back into the database. Thus, users are not limited to the information in the system; if the information is not present, they can tap the social network in a natural way. As a result, the construction of the information is iterative, and the corpus of information grows over time.

Several important considerations behind Answer Garden include:

- Answer Garden extends the standard information retrieval model [Lancaster and Fayen 1973] to include the social network of the organization. It does so by including elements of asynchronous or synchronous communication systems within the system.
- The incentives for building and maintaining an Answer Garden do not necessarily result from each person deciding to be cooperative per se. Each user and each expert has incentives to work separately towards the construction of such a Garden. The users get to find answers, and experts can rid themselves of commonly asked questions.
- Answer Garden assumes that the information database will be iteratively designed and built. The design inherently allows user feedback for the

correction of mistakes, requests for adding materials, and changes in the indexing structure.

- The design of Answer Garden allows for the production of information on demand. Answer Garden grows where users need information, and the information producers need to create information only where it is needed.

3.1 *The Answer Garden Substrate*

Most of the system facilities described were provided by an underlying hypermedia system called the Answer Garden Substrate (AGS). There have been a number of Answer Garden databases built using AGS; additionally, other implementations of Answer Garden exist in Lotus Notes and the World Wide Web. AGS, which was designed to support organizational memory applications, has several advantages over other systems:

- All node types and general services in AGS are extensible and open. Thus, it is relatively easy to add functionality to existing node types or to add node types. This is important in developing new organizational memory applications.
- AGS provides a number of typed information nodes suitable for supporting organizational memory. This allows for varying types of information manipulation, processing, and display. To the user, AGS attempts to provide uniform display and navigation.

AGS has also been used to build two additional organizational memory systems: LiveDoc and the ASSIST. LiveDoc consists of some formal documentation, such as a computer manual. LiveDoc uses "side notes" to mark and answer any additional issues or questions that may arise in use. This additional information can include full discussions by organizational members, or questions and answers about the documentation or underlying process. If the user has a question or problem with

a particular section of the document, the question and answer series can be placed by location in the document. Documentation writers can later harvest these side notes and modify the original documentation.

The ASSIST is an application, developed by the Harvard-Smithsonian Astrophysical Observatory, that allows astrophysicists to easily find documents, data files, and software components. In essence, the application writers embedded over a thousand software modules within a surrounding memory and help system. The ASSIST is in world-wide use by the astrophysics community. The ASSIST is more fully described in [Mandel et al. 1992; Ackerman and Mandel 1995].

3.2. *Related systems*

Answer Garden combines many different kinds of systems in a novel manner. Interestingly, the addition of various capabilities that may be found in differing computer applications leads to a very different type of application in practice.

Answer Garden differs from standard information retrieval systems (e.g., Bartschi 1985) in its addition of communication capabilities. Answer Garden varies from computer-mediated communications systems such as enhanced mail Borenstein 1992 or bulletin board systems (e.g., [Hiltz and Turoff 1981; Horton and Adams 1987]) in its capability for retrieval and its emphasis on growing a memory. Answer Garden differs from consulting systems using expert system technology (e.g., [Gwei and Foxley 1990; Ram, Hayne, and Carlson 1992; Terveen, Selfridge, and Long 1995]) because of its reliance on human experts, its resulting non-brittleness and its emphasis on iterative design without knowledge engineers. Answer Garden also differs from other forms of organizational memory, such as design rationale (e.g., [Conklin and Begeman 1988; Lee 1990; Terveen, Selfridge, and Long 1995]), in its emphasis on informal information and communication flows.

4. Field Study of Answer Garden

Answer Garden offers an alluring application, providing the possibility of capturing portions of an organization's memory at a relatively low cost to that organization. But, would it work in practice? I undertook a field study of Answer Garden using the sample information database about the X Window System to determine whether, in fact, it would.

I chose to study software engineers using the X Window System. I did this for a variety of reasons including the iterative nature of user interface development, the complex nature of the X Window System and its toolkits, and the presence of electronic mail and network connections in a workstation environment. Software engineering is of special interest for an examination of expertise and information seeking because there is no body of commonly accepted knowledge (aside from low-level data abstractions and algorithms which all programmers are expected to know). Furthermore, the software world is undergoing constant revolution, preventing an individual engineer from acquiring lasting expertise. Software engineers using the X Window System, then, were ideal candidates for information seeking through electronic means such as Answer Garden.

4.1. Field sites, participants, and research methods

Two field sites provided almost all of the data in the field study. These two sites were a research group at MIT and a class in the Harvard extension program with a total of 59 potential users. Additional field sites included a software group at a federally-funded laboratory, a product team in a software company, an advanced development group in a large computer company, and a support group in another computer company. Two of these supplemental sites were local; therefore, I was able to obtain interview data.

The Harvard site was a Harvard extension class in two-dimensional graphics programming using the X Window System. The Harvard class had 44 students and 3 instructors. Its students were largely working software engineers, and as such, used Answer Garden in their workplaces, potentially

substituting it for asking questions and seeking information from organizational colleagues. The MIT research group, with 12 members, was engaged in sponsored research, and their work flow appeared to be similar to that in a software R&D company.

There was, additionally, another set of participants in the field study: the experts who answered questions. I wished to provide the study participants with authoritative information, so that they would not need to judge whether it was accurate or not. Two colleagues and I served as the first-line of answerers. The seven experts included staff members of the X Consortium as well as consultants with more than three years worth of experience in Motif or OpenLook (popular varieties of X). Other experts volunteered, but were not needed.

The study used multiple data collection procedures including an initial questionnaire, usage data at the "mouse-stroke" level, critical incident interviews, final questionnaires, and field observations and interviews (with both users and experts). The primary methods were the usage logs and the critical incident interviews. Most critical-incident interviews, held with users shortly after their use of the system, ran between 15 and 20 minutes in duration, including time spent obtaining qualitative data concerning the interviewee's responses. For example, users often provided (and were encouraged to provide) additional information on why they rated the system as they did.⁵ Because the critical incident interviews were not systematically randomized, only qualitative data from the 49 critical incident interviews will be described.

The observational data was used to supplement and bolster the qualitative evaluations and quantitative usage data. In general, a concerted effort was made to use the qualitative and quantitative data together.

⁵Users were reassured in the interview that negative responses were as valued (or more so) than positive ones.

4.2. Field study assumptions

In addition to the design considerations behind Answer Garden in general, several other assumptions and simplifications were made in the field study. An implicit assumption in the field study was that there were separate groups of experts and of information seekers. This is clearly not the case in organizational life, where people are arrayed along many continua of expertise over many subject domains. Other versions of Answer Garden could allow for more of a community interaction.

Additional simplifications in the field study were made about the type of information to be captured in the information database. An implicit assumption behind Answer Garden is that the information within an information database should be easily classifiable and easily broken into node-sized "chunks." Furthermore, this information database contained only static information, and did not include connections to people, software modules, or dynamic information. (The Answer Garden Substrate does support these types.)

The version that was released to the field sites was essentially the same as the version distributed on the MIT X11r5 contributed software tape. Because it was designed to be readily available, it is quite robust, and only a handful of minor bugs have been reported. However, this version did not include information retrieval engines (other than the standard, naive retrieval engine) or window stacks (to tidy up the user's screen).

5. Patterns of Use

...we wish to present some trends and observations of the uses, strengths, and weaknesses of this ... tool. ...we have tried to be as candid as possible about the weaknesses and research problems.... We hope this candor does not create an overly negative impression about what we feel is a very positive research effort. (Conklin and Begeman discussing gIBIS [1988], p. 147)

Before discussing some of the more interesting findings from the field study, it is important to describe how much and in what ways the system was used. This section presents material as a background for the subsequent Evaluations and Lessons sections.

5.1. Answer Garden usage

During the study, Answer Garden was used 194 times by 35 users in the two major sites (mean number of uses per user = 5.54, s.d.= 7.01). Since Answer Garden was designed to be a system that would be used intermittently, I expected usage would occur only when the user had a difficult question, perhaps twice a week or even once a month for consistent users. In fact, this is what was observed. The maximum that any user used Answer Garden was 28 times over three months.

This pattern of usage did not appear to be from lack of interest. Only five of the users appear to have been just looking from curiosity; the others self-reported serious usage. Moreover, these 35 users were 59% of the 59 potential users; additionally, several potential participants reported being unable to connect to the Harvard Answer Garden for technical reasons.

Answer Garden's use was not uniform across the test period. Figure 5 shows the aggregate usage. Usage was higher initially as users experimented with the application. Usage dropped in week 11 since this was the Thanksgiving holiday. Additionally, usage dropped after week 14 when the MIT group ended their use with the end of MIT's semester.




Figure 5: Aggregate usage by week

The two sites, however, showed different usage patterns. The MIT group started two weeks earlier than the Harvard group and displayed gradual increase in the number of sessions per week through week 6. The MIT group then fell into a steady usage of two to four sessions per week, with use from nine of eleven potential users. A spike during week 8 suggests that usage can increase by a factor of 100% during "normal" periods. Interviews with these MIT participants revealed no systematic pattern since usage was driven by a combination of local deadlines and steady-state work. This is the same as the demands on experts' time now; people do not ask questions of experts on a specific schedule. Such high variations in usage increase the difficulty of scheduling experts for answering questions through Answer Garden, and suggests that a workflow component is required for similar systems. As might be expected, the usage from the Harvard class was driven by class deadlines, leading to a more irregular pattern.

5.2. User groups

Even taking into account the user's location, usage varied considerably by individual. A simple diagram of the number of sessions by user (Figure 6) shows that six users were "heavy" users, accounting for 61% of the total sessions. Heavy users often kept Answer Garden running in an iconified state between uses. One heavy user reported:

You know that I find it valuable when I've always got it running. It's just one of my normal suite [of tools].

The other participants could be further divided into two groups, a finding obscured in Figure 6. An analysis of usage by person, broken down by the date of their session, shows that there were a group of users that used Answer Garden more infrequently, but over a long time period. These 12 participants were classified as "intermittents." Some of these intermittents used Answer Garden for as long, but not as intensively, as the heavy users. That this group of intermittent users existed was extremely satisfying since it indicated that some users found Answer Garden sufficiently useful to return after several weeks. These users used Answer Garden as was originally predicted, as an occasional tool for information seeking. Additionally, there was a group of seventeen people, or 49%

of the participants in the study, who used the system for a short period of time; they were categorized as "tire-kickers."⁶




Figure 6: Usage by user

The qualitative responses indicated that the standard problems in introducing a new software system were present in the field study. For example, some non-users had some difficulty in connecting, while others did not want to spend the time to learn the system. However, there were some new difficulties as well. Even a heavy user spoke of the issues of finding information within Answer Garden:

There was a learning curve issue in finding out where the good information was, you know? ...The people who persevered got good at it.

Users also spoke of more subtle issues. One important issue was an ambiguity in the users' minds about what questions were appropriate:

I was often unclear as to what were legitimate Answer Garden questions.

This may have been one of the factors that separated intermittents and tire-kickers. This issue will be discussed further below.

⁶I attempted to find determining demographics for these user groupings. However, the statistical power for this sample size (n=35) was not sufficient to test for anything but large effects [Cohen 1977].

5.3. Mail usage

Users asked many questions in the field study. In 53% of the 194 sessions, mail was sent for a total of 121 mail messages. (In 16 sessions, users sent multiple mail messages.) There were an additional 43 times that users opened mail windows, but did not send any message. The reasons were very obvious in the critical incident interviews. A typical user response during an interview was: "I could have used more info [in the database]."

Largely this was the result of an incongruity between the coverage in the database and the size of the topic domain. The X information database was small (98 nodes at the start of the study with approximately 225 questions and answers) compared to the number of possible questions and answers about X (at least thousands of potential questions). This suggests two possibilities: future Gardens should be restricted to smaller domains so users will be more likely to find answers to their questions already in the information database, or the initial period of activity will be largely devoted to building the domain coverage.

While most field study activity built the answer database, there were indications that users also found the existing database useful. Eighty-three percent of the questions were mailed after they examined the database. (The remaining were posted from the root or a top-level node by eight users. Three users were gently admonished during the field study to stop posting from the root, and two did so. The other five were infrequent or normally posted correctly.)

The subjective evaluations (described below) were consistent with this. More Harvard users found the existing answers useful than did those at MIT. This is perhaps unsurprising since everyone in the Harvard group was going through similar class exercises, similar to the repetitive organizational activity that Answer Garden was designed to support.

6. Evaluations by Users and Experts

Because of Answer Garden's design, two groups participated in this field study, the users and the experts. In short, the qualitative evaluations from both groups showed a fairly complex set of nuanced activities. Generally, people found the system satisfactory for their use, but there were places where the system needed to provide additional support.

6.1. *Users' evaluations*

The 49 critical incident interviews, conducted with the participants about their use of the system on a session basis, included 30 negative statements and 23 positive statements about the system. (Several participants gave mixed responses.) Overall, the majority of numeric scores in these same critical incident interviews were favorable.

In general, the overwhelming emphasis in the performance evaluations for the sessions was that the users were satisfied as long as they got an accurate solution quickly. The source of the answer did not matter. Users rated the application more highly when an answer was already in the information database or when they got a quick response by electronic mail to their question. A typical evaluation was: "It was very good. Whoever answered it was exactly right, and very quick too." Participants gave an unfavorable rating when they got no answer or a slow answer. In fact, almost all (83%) of the unfavorable responses were due to not being able to get the information quickly, either in the information database or by electronic mail.

In addition, users did not like it when they had to wade through a considerable amount of text to find that the answer was not present. If the answer was not in the information database, they wanted to know quickly so they could ask their question through electronic mail. In fact, being able to quickly determine that the information was not in the database was often seen as a positive attribute. A majority (60%) of the favorable responses in the interviews mentioned being quickly able to

determine whether the information was in the database, getting a quick response, or finding the information in the database.

Users appeared to also be comfortable with information seeking through a combination of standard information retrieval system and the social network. This conclusion is based on indirect evidence, primarily the differences in comments between the first weeks of the study and later. During the first weeks of the field study, users had to be prompted to accept the use of the social network as part of the information retrieval process. Participants consistently reported not wanting to bother the experts or expressed surprise at the idea that they could ask the question of a human if the answer was not present. (They had been told how the application worked repeatedly in introductory talks and handouts, implying that they were not used to the idea of Answer Garden at first.) Several steps were taken to reinforce their understanding, including a sentence on the initial node reminding them to ask questions.

After the first month, the performance evaluations suggested that their continued usage of the application was natural. Users stopped evaluating the information in the database separately from the electronic mail response; the system evaluation was based on both together. Users did not answer that they did not ask a question through the system because it seemed strange or novel even though they were prompted for such responses. By the middle of the field study, tying the social network into the information retrieval system appeared to be normal to the users. Said one user:

I always went straight to AG at this point [after looking at the manuals and asking people nearby], instead of trying to guess someone who would likely know the answer, because I trusted that there was a mechanism in AG that would guarantee a satisfactory answer.

In addition, Answer Garden seemed to ameliorate the status implications of information seeking for many users. No one mentioned negative status implications in the critical incident interviews. This was true for both the people who sent a question and for those who did not. In fact, session evaluations were on the whole appreciative of reducing the status interactions in information seeking. When it was mentioned, the users liked being able to ask their questions anonymously. Said one

interviewee: "[It's] a vehicle where you're not intimidated; you're asking anonymously and through text." In a quarter of the interviews, participants said that the application helped them ask for information that anyone would be expected to know. Another half (56%) of the interviewees said that their questions were intermediate in difficulty.

However, when asked about the information itself, users had a mixed reaction to the information gathered through Answer Garden. In nearly 20% of the interviews, participants pointed to Answer Garden as providing information they could not have received elsewhere. However, about a third of the users reported problems with the specificity of the material (either too high or too low) and the level of the explanation provided by the expert (either too high or too low). This was an indicator of another issue, the problematic dichotomy of expert and novice, as will be discussed in the Lessons section.

6.2. Experts

As mentioned, there were seven external experts in the field study, all of whom had more than two years of experience. As mentioned, many of the experts answering questions were directly responsible for parts of the X Window System or its derivative vendor releases. The design of the field study tried to keep the workload on any of these external expert to a minimum. The most questions given to an external expert was two per week.

Five of the experts expressed initial fears about being overwhelmed by the task of answering questions. Accordingly, the experts demanded and were given the right to refuse to handle questions based on their workload. Two experts invoked this rule for almost half of the field study, suggesting that redundancy in expertise will be required for similar systems. The external experts also requested the right to not answer any question that required excessive time or was too specialized; however, this rule was invoked only once during the field study.

No expert complained about the time required for answering questions. However, the style of answering was an issue in the study. Four of the experts showed a marked formality in their

responses. Their answers were longer, containing substantially more detail and more cases than would a quick electronic mail response. When interviewed, these experts reported wanting to provide a more complete answer for two reasons. First, they tried to immediately generalize from the specific situation to a general answer. Moreover, they wished this general answer to be suitable for a large range of questions and special cases instead of growing the answer iteratively over time. Second, several reported that their response served as a public "badge" for them within the organization or community. As such, they wanted to make sure that their answer was complete and accurate - not doing so might reduce their status with people who did not know them well.

This push towards formality of response goes against the users' desires to have short, easily readable answers and the system goal of capturing informal flows of information. Not publicly identifying the expert might reduce the formality and further reduce the status implications in information seeking, but users use this identification to judge the potential accuracy of the information. This dilemma may argue for the necessity of an information "editor" or "moderator" to ensure authoritativeness and consistency in the information database.

Nonetheless, three of the external experts did provide informal responses. This suggests the possibility of keeping the responses informal. Again, an editor or moderator might be valuable in some situations; he could work to keep the tone informal.

7. Lessons from the Field Study

How successful were the design assumptions behind Answer Garden? The field study resulted in the following:

- ❑ The field study suggested that the type of organizational memory provided by Answer Garden could be built, although the final coverage of the information domain was not extensive. Users employed the system to find information by accessing the questions and answers provided as well as by asking new questions through electronic mail. The answers to these new questions were inserted into the

information database, thus growing the corpus of questions and answers. The combination of information retrieval and communications system was key to the successful operation of Answer Garden. Moreover, users appeared to find the combination usable and comfortable.

- ❑ Such systems should provide suitable incentives for use. The incentives appeared to work in the field study. Users found the system effective when they received correct and timely answers to their problems, and they were willing to ask questions through the system. However, the incentives for the experts were not tested since the system was not available for an extended period.

These findings demonstrated that Answer Garden could work in principle.

It was also hoped that such an organizational memory system as Answer Garden would result in the reduction of status implications and the need for reciprocity. The field study resulted in both a mixed achievement and some new insights about this design consideration:

- ❑ The study found that Answer Garden could reduce the status implications for many information seekers. The ability to ask questions anonymously and to ask questions of an appropriate expert were found to be beneficial by a number of users.

However, some users still had issues concerning the status implications in information seeking through the system. Whether these users would find their concerns ameliorated over an extended period of use is unknown. It may also be possible to reduce the status implications by providing access to lower-status help desk personnel.

- ❑ For the experts, the field study uncovered the need for experts to maintain their organizational "face." I had not expected, before the field study, that experts would also have status implications in their information providing role. Further work will

need to concentrate on finding mechanisms to reduce any status implications for the experts.

- ❑ The need for reciprocity appeared to have diminished, if not disappeared, in the two field sites. Users asked questions through the application even though they were not providing information back to the experts.
- ❑ A number of users continued to fret over their "bothering the experts." For example, one user said: "The one inhibition I felt using Answer Garden [was] knowing that the experts were typically busy and working on projects more important than my little application programs...."

This reaction could indicate continued concern over the status implications in information seeking. However, the users' responses suggested that software engineers may have an additional social consideration when seeking information, a consideration that had been previously masked by the concern over status implications.

One possibility for these comments might be a respect for the experts' position. The words the participants in the field study used could indicate a respect for the tasks and duties of people above them in a technically-based hierarchy. Because software engineers recognize technical expertise as a meritorious achievement and because their authority is based on technical expertise, they may not wish to "bother" the experts.⁷ This situation would then show characteristics of Blau and Meyer's example of a medical bureaucracy with its respect for professional authority [1987]. It is possible that the status implications were masking this additional concern for the experts' role and position.

⁷This is not the same as a status interaction (or concern over one). A status interaction involves a

If such a concern does exist, Answer Garden does little to ameliorate it. Indeed, Answer Garden may exacerbate it. At least in the field study version, Answer Garden identifies some organizational members as experts, placing them automatically above the user. This would imply that the clear-cut separation of experts and users is not only artificial, it may lead to operational difficulties.

An alternative possibility is that the concern resulted from the experts' being volunteers, rather than being organizationally responsible for these answers. In other words, users felt that answering questions required labor that was not organizationally rewarded. Shifting incentives might change the behavior and comments of the users. While I provided users with knowledge of the experts' incentives, it may not have meant as much to the users as clear organizationally based rewards.

Further research is required to differentiate whether either respect for professional author or concern over organizational rewards causes such comments.

- An interesting, and provocative, finding from the field study was that a large proportion of the users did not get answers that were at the right level or length of explanation. In reflection, the assumption that users should always have their questions answered by experts may have been false.

The design of Answer Garden in the field study assumed that the status implications and need for reciprocity were only negative and should be removed. It may be, however, that they also provide positive organizational benefits. For example, the status implications in information seeking and the need for reciprocity may serve to channel information seekers towards others at their expertise level, thus providing the seekers with answers at the right level and length of explanation.

two-way exchange or the potential of one; this is a deference to authority.

If this is correct, then it is not the status implications and need for reciprocity that are problematic *per se*. The organizational problem arises when the information seeking can be considered dysfunctional [Blau 1955]. For example, an organizational dysfunctionality occurs when the people at the same expertise level cannot answer a question and there is no person with greater expertise available.

In general, these latter findings showed that the social issues drove the specifics of use. Future organizational memory efforts, or computer-mediated communication systems with similar purposes, will need to ameliorate these additional social issues as well as achieve the original goals of Answer Garden.

8. Summary

Answer Garden presents a potential method of building an organizational memory. In its emphasis on helping people find answers to their questions and in saving the frequently asked questions, Answer Garden offers an alluring possibility of building such a memory without substantial cost. In general this was supported by a field study of its use.

The field study resulted in two intriguing findings that will need to be considered for future systems. First, some experts' perceptions of needing a public "face" in an organization or community pushed towards a formalization of their responses. Second, there are positive effects from reciprocity and status implications for the expertise network of an organization, and these effects route organizational members more effectively to the right level of expertise. Both of these findings will require some adjustment of the design assumptions behind Answer Garden and similar systems.

Some important caveats remain to these findings. The field study examined only the use among software engineers, and generalizations to groups with career anchors other than technical expertise are unclear. For example, there is a strong possibility that doctors or teachers would have different usage patterns. Furthermore, this study could not follow long-term results from Answer Garden usage, and many studies have suggested that the long-term effects of an information system cannot be

adequately determined from observing the short-term effects (e.g., [Hiltz and Turoff 1981]). For example, organizations could choose to use Answer Garden to prevent experts from being "interrupted." If users were prevented from this, it might lead to loss of social activity by both users and experts. Poking one's head in a office door has other purposes and benefits than information seeking.

Nonetheless, within the confines of the field study, Answer Garden "worked". After a short period of adjustment, users found