Published in the Proceedings of the ACM Conference on Office Information Systems (Cambridge, MA, April, 1990), pp. 31-39. The figures in this version have been redone.

Answer Garden: A Tool for Growing Organizational Memory

Mark S. Ackerman[1] *and Thomas W. Malone*[2]
Center for Coordination Science
Massachusetts Institute of Technology

Answer Garden allows organizations to develop databases of commonly asked questions that grow "organically" as new questions arise and are answered. It is designed to help in situations (such as field service organizations and customer "hot lines") where there is a continuing stream of questions, many of which occur over and over, but some of which the organization has never seen before. The system includes a branching network of diagnostic questions that helps users find the answers they want. If the answer is not present, the system automatically sends the question to the appropriate expert, and the answer is returned to the user as well as inserted into the branching network. Experts can also modify this network in response to users' problems.

Our initial Answer Garden database contains questions and answers about how to use the X Window System.

1. Introduction

Consider the following common problems in organizations: Customers get incomplete or inconsistent answers to their questions and complaints. Employees get stuck or waste effort when they can't find the specific information they need. Meanwhile, the people who could answer these questions may, in some cases, be unaware of the problems people are having, or, in other cases, they may be overwhelmed by the sheer mass of simple questions they have to answer.

The Answer Garden system helps an organization solve these problems by providing a database of answers to commonly asked questions that grows "organically" as new questions arise and are answered. It is designed to help in situations where there is a continuing stream of questions, many of which occur over and over, but some of which the organization has never seen before.

Such situations are pervasive in many kinds of organizations: field service technicians repairing a company's products, telephone "hot lines" for customer problems, customer contact personnel (such as airline ticket agents) responding to customer complaints, teaching assistants helping students in large classes, coworkers helping each other with "lore" about how to use complex computer systems, and even primary care physicians treating their patients in medical clinics.

In this paper, we will describe a cooperative work tool, called the Answer Garden, that we have designed to help improve an organization's "memory" in situations like these. Even though there is little consensus in the organizational literature about how to define organizational memory (e.g., Orr86, Levitt88, Walsh89), the concept of organizational memory is an evocative and attractive metaphor. For our purposes here, we will define organizational memory as an organization's ability to benefit from its past experience in responding more effectively (e.g., faster or more accurately) in the present. We will be particularly concerned with situations where there is a body of "authoritative" knowledge about how to do things (such as using devices or following organizational procedures).

[1] ackerman@athena.mit.edu or ackerman@xv.mit.edu
[2] malone@xv.mit.edu

We begin by describing the key ideas in the Answer Garden system and the incentives people would have for using a system like this. Then we present a more detailed "tour" of the system and describe its implementation. Finally, we present a few observations from our first use of the system (with a database of questions and answers about how to use the X Window System), and we discuss a variety of issues that will arise in large scale use of systems like these (such as how to represent conflicting points of view about a question).

2. Overview of the system

Even if someone in an organization already knows the answer to a question, unless that answer can be easily located, other people may have to recreate it, over and over again. The Answer Garden helps an organization "improve its memory" by capturing answers to questions that would otherwise be lost and then providing tools to retrieve the answers later. There are three key ideas in the system:

(1) *A branching network of diagnostic questions helps users find the answers they want.* The primary tool to help users find answers is a branching network of diagnostic questions like those an expert might ask to diagnose a user's problem (or like the questions people playing the Twenty Questions game have to ask). For instance, the first question in a database for hardware diagnosis might be "Does the power light come on?" with different subsequent questions depending on whether the answer to that question is "yes" or "no". At the end of each path through the branching network of diagnostic questions is a set of questions people have asked in that situation and the answers to these questions provided by experts.

(2) *New questions from users are automatically routed to appropriate experts and then inserted (along with their answers) into the network.* Whenever users comes to the end of a path in the branching network and don't find the answer to their question (or when they have trouble answering a diagnostic question along the way), they can enter a new question. This question will be sent automatically to an expert who is knowledgeable about questions in that part of the network. Then, the expert's answer will be returned directly to the user who asked the question, and it will also be entered into the network (along with the question) at the point where the question originated. Thus, the Answer Garden grows its information database over time.

(3) *Experts can modify the diagnostic branching network in response to users' problems.* When experts conclude from a particular user's question (or from usage statistics for the branching network as a whole) that the branching network is misleading or inefficient for users, they can add to or change the diagnostic questions in the branching network. This means that heavily used and well-understood parts of the network will presumably have well-honed branching networks, while the branching networks in lightly used or "frontier" areas will be more ad hoc.

Only a small number of simple components are needed technically to implement this system. Yet, such a system serves as an excellent testbed for a number of organizational and social questions.

3. Incentives for using the system

Why would people want to use a system like this? As Grudin (Grudin88) and others have pointed out, the success of cooperative work systems may depend in very subtle ways on the incentives that different people have for using them.

3.1. Incentives for questioners

For most questioners, the primary alternatives to using a system like Answer Garden would probably be (a) finding the answer in some other written medium, or (b) finding a person who could answer the question. In some cases, these alternatives may be preferable. For instance, if the questioner knows exactly where to find the answer in a printed manual, that might be faster and easier to locate and read.

In other cases, questioners may prefer to ask people they know, or people whose job it is to answer questions. Dealing with live people has the advantage that the expert can be more responsive in helping to formulate the question.

For instance, the questioner can usually just describe a problem in English rather than having to answer a series of multiple choice questions. In some cases, dealing with a live person may simply provide a pleasurable social interaction.

In an important number of situations, however, we believe these advantages will be outweighed by the advantages of using a system like Answer Garden. With Answer Garden, for instance, users can quickly find answers to previously asked questions at any time, even when they can't find any human experts. Moreover, since the answers in the Answer Garden have presumably been provided by authoritative experts, users can usually rely on them much more confidently than on answers from a random colleague down the hall.

Even when the answer to a question is not yet in the database, systems like Answer Garden provide a very efficient means of locating appropriate experts since new questions will be automatically routed to them.

3.2. Incentives for experts

For experts, Answer Garden eliminates the need to answer many simple questions. The experts no longer have to answer the same questions over and over again, and can instead concentrate on more interesting problems. For example, one expert in our pilot study estimated that less than 10% of questions asked of him required him to find new information.

3.3. Incentives for organizations

For an organization, Answer Garden provides better, faster, easier to get answers. Thus the efficiency, consistency, and quality of the answers given to internal or external customers is higher than before. It also reduces the need to have large amounts of experts' time devoted to repeatedly answering the same questions. Finally, as the costs of information technology decline, it becomes feasible to codify and store electronically more of the knowledge that is now stored only informally in the minds of people.

4. A Tour of the Answer Garden

The user invokes Answer Garden from the prompt line or a suitable menu. Figure 1 shows the initial pane, the control panel. Because the Unix environment in which the system operates is a multi-tasking, multi-windowed system, other programs (like those shown in the background of Figure 1) can also be simultaneously displayed. From the control panel, the user can choose between the two major ways of locating answers: (1) multiple choice questions, and (2) a tree graph.

4.1. Helping users locate answers

The primary way users locate answers in Answer Garden is by answering a branching series of multiple choice questions, starting at the control panel. Each question is presented in a structured text node (like those shown in Figure 2). When the user selects one of the answers by clicking on the appropriate button, the next structured text node appears, and the user continues answering multiple choice branching questions until she finds the question and answer for which she is looking. There are a number of automatic placement options for the nodes.

In many cases, we expect this branching network to be a series of diagnostic questions like those an expert (or an expert system) might ask to diagnose a user's problem. It is important that these questions be phrased in a way that users can understand them. For example, the questions should usually refer to symptoms the user can observe (e.g., "Does the power light come on?"), not to underlying causes or subtle distinctions only an expert can recognize (e.g., "Does the disk whine sound normal?").

Some users, especially experts, may already know roughly where their question is answered and will not want to "click" their way down a long branching chain to find it. In these cases, users can select the "View tree" option in the control panel and see a tree (like the one shown in Figure 3) of the nodes in the branching network. A user can then jump immediately to any node in the network. The grapher pane may include links to other grapher panes; in the current version, these are marked by using the "..." convention to imply additional items.

4.2. Growing the information database

The scheme just described works well when the Answer Garden knowledge base contains correct answers to all the questions people might ask and a well-structured branching network to help people quickly locate the answer to their question. But what happens when the knowledge base of answers in the system is incomplete or out of date? What if a user doesn't understand a question? What if the answers to common questions are "buried" at the end of a long branching chain? Answer Garden provides two basic tools for dealing with these problems: (1) new questions, and (2) usage feedback.

New questions are answered by experts. Whenever users are not satisfied, they can click on the "I'm unhappy" button in any node. They might click this button, for instance, if the answer to their question isn't present at a "leaf" node, if they don't understand the branching question they are being asked, or if they are frustrated with the path down which the branching network is taking them.

When users click the "I'm unhappy" button, they get to describe their problem or ask their question directly to a human expert who is knowledgeable about the node at which the problem occurred. In the current version of the system, users communicate with the experts via electronic mail; future versions could allow synchronous communication with experts for parts of the database. Figure 4 shows an example of the kind of node through which users can ask their new questions.

These mail messages are mailed not only to the appropriate experts but optionally to an additional notification list. The list of experts and notification list can be changed on a node-by-node basis, allowing organizational customization. Thus, some experts may be within the boundaries of the group or organization, while others may be outside of it. When the expert answers, the answer is automatically inserted in the information database.

Experts use usage feedback to restructure the branching network. There are two kinds of feedback to experts about problems users are having with the knowledge base: (1) new questions, and (2) usage statistics (i.e., which nodes and paths are accessed most frequently). New questions and their answers are incorporated into the network immediately, thus correcting incomplete or inaccurate things as soon as a user discovers them. Note that this ensures that experts will spend their time on questions that are actually useful to people, not on esoteric things that no one ever wants to know about. New questions can also give feedback about where the branching network is ambiguous, confusing, or frustrating. Usage statistics provide another kind of indication about how the system is being used and where the problems are. For instance, if certain very long branching paths are traveled very often, this is an indication that the branching network should be revised to place the questions higher in the tree.

4.3. Answers

Information in the Answer Garden can include simple text, structured text, pictures, and active processes. As the database grows, mail messages can be included as simple text at first. Later, an expert can separate and structure the information, presenting it in a structured text node. A structured text node can include proportional text layout, a variety of fonts styles, and active buttons. Active nodes can trigger actions at run-time such as querying a remote database, and therefore build dynamic information. Thus Answer Garden can serve as the "front-end" for an entire set of separate, user-defined programs.

5. Implementation and sample database

The current implementation for Answer Garden is written in C for Unix (Berkeley BSD 4.3) platforms. It uses the X Window System extensively. Since Answer Garden is to be available as widely as possible, all work has been done with the Xt toolkit and the Athena widgets, which are publicly available (McCormack89, Swick88). Most of the functionality is in new widgets compatible with the Athena set.

The Answer Garden uses the standard BSD Unix mail and file systems. Answer Garden nodes are kept as Unix flat files; no additional database is required to run Answer Garden. This will degrade system performance when the number of nodes becomes large, but it has the sharp advantage of simplicity. Usage of the mail system is largely invisible to the end

user; Answer Garden invokes the mail system directly. Site-settable and user-settable options (resources) allow Answer Garden to be tailored to the various Unix mail handlers.

We are testing the initial Answer Garden system using an information database on the X Window System itself. Since X is a newly developed and widely used standard, many organizations do not have the range of expertise and support necessary for large-scale development efforts. Based on our discussions with potential users, we believe this database about X will be quite valuable to a very large number of people at many sites.

We are gathering questions and answers from the most heavily used X mailing lists and interviews with staff members of the X Consortium (the organization that develops and maintains the X Window System standards and software). We have asked the X Consortium staff to provide the most popular questions that they answer continuously, and they have agreed to answer the new questions that are asked using the system.

As the number of users grows and they ask new questions, the X database will grow. Thus, we are starting with a core of questions and answers, and expecting the users to expand our coverage.

We envision the system as having three levels of information: personal, group, and general. The group may be a particular site in an inter-organizational system. Most of the information will come from the members of the X Consortium, and will be generally true for all X users. For example, X Consortium staff may add nodes about new applications that run on specific hardware platforms, and all sites will get this information. Group or site information may include information true only at a given location, such as who the site administrators for X are.

6. Usage experience and issues

In the first phase of this study, the Answer Garden system has been used irregularly by about 8 people in 4 groups at MIT for a period of 4 months. In the next phase of the study we will include several groups in external organizations as well. Our early experience has already highlighted several issues that we expect to be important as the usage and size of the database increases.

Authentication of information. One of the problems of bulletin board systems is that it is difficult to determine - let alone control - the quality of the information. Answer Garden provides some technical mechanisms for access control, but much of the control will be social. In the sample X database, for example, only local experts and X Consortium staff members can change the information database. The X Consortium staff members are, de facto, the ultimate experts, and therefore authentication of information is not an issue.

However, this locking mechanism can be bypassed either at the Answer Garden level or at the Unix file level. Local site administrators could, of course, adulterate the database. Moreover, we have no control over site-specific nodes; these will be under the control of local experts. We surmise that an editing function, perhaps within each site, may become necessary to control the quality of the information and maintain the database.

A future enhancement almost universally requested is to allow users to annotate the entries. We also think it useful to allow users to form their own access trees for quicker retrieval of frequently used information.

Scaling. We expect the X information database to ultimately grow to thousands of nodes. As the number of nodes grows, ease of access for users will become more critical. Currently, only the branching questions and tree graphs exist; however, Answer Garden has been designed to allow a multiplicity of retrieval engines. For example, we are currently investigating semi-structured retrievals and dynamic trees to provide greater flexibility across ranges of expertise, machine characteristics, and technical interests of users.

In addition, potentially thousands of users, software engineers and others, will be using this database. These users are spread across diverse organizations and geographic areas, and access to a centralized site is not practical. Many of the information subtrees will be organization and site specific. For example, the availability and location of X binary files is always specific to the site. Therefore, the X information database is fully replicated at each site. Updates, however, should

occur continuously, each day or more often. To do this will require a steady flow of mail traffic to each site, bearing the changes to the information databases.

Representing multiple points of view. In the scenario shown in the figures, an expert has arranged the branching questions in such a way that the answers are specific and small. In other cases, a topic may be too recent or general for the expert to determine how the information database should be structured. Or there may be differing opinions, even among experts, on the correct answer. In this case, a "moderator" might help organize the branching network, but not provide authoritative answers.

For example, a user might want to know about the advantages of using X terminals over workstations. After answering the appropriate questions - whether the problem is a question of X usage, of platforms, and of hardware - he can browse through a large amount of mail traffic on the various types of terminals and their vendors. From time to time, the moderator might select the most interesting messages or modify the branching network to reflect the differences in terminals and vendors.

7. Related Work

Answer Garden has points of commonality with a number of previous systems. For example, videotext systems provide a structured branching network to help users find information. They do not usually, however, allow users to easily ask new questions and have the answers automatically added into the database via electronic mail.

Electronic mailing lists, computer conferences, and bulletin boards allow ad hoc sharing of information among widely distributed groups in response to new questions, but they usually provide only the most rudimentary kinds of structure. For instance, it is common for computer conferences (e.g., Hiltz81) to include a few levels of topics and subtopics, but we are not aware of any system that treats this structure as a diagnostic branching network.

At the opposite extreme of structure, expert systems often provide elaborate diagnostic branching networks, but they usually require significant effort to construct and they are not ordinarily thought of as being automatically extensible in response to new user questions and answers from experts via electronic mail.

The kinds of questions and answers used in Answer Garden could certainly be provided in a multi-user hypertext system (Conklin87, Aksyn87, Trigg88). In fact, one can think of Answer Garden as an example of a multi-user hypertext system. We are not aware, however, of any other hypertext applications that exploit this kind of diagnostic branching network and are automatically extensible through electronic mail.

Perhaps the closest precursors to Answer Garden are systems designed explicitly to aid in providing immediate answers from technical support personnel to questions about hardware and software (Coppeto89, Dyson89). With the exception of OLC, these systems appear to be designed for use only by skilled technical intermediaries, not by the end users who actually have the questions. Furthermore, these systems do not appear to include a diagnostic branching structure to help users locate answers.

8. Conclusion

The Answer Garden is not a radically new kind of system. We believe it shows, however, how a relatively simple combination of well-known concepts can provide a surprisingly powerful platform for a new kind of cooperative work application. As information technology becomes more pervasive, we believe that tools like this--for capturing and exploiting organizational knowledge that was previously stored only informally or not at all--will become increasingly important.

Acknowedgements

This work has been funded in part by research grants from the X Consortium, Digital Equipment Corporation, and the National Science Foundation. We would like to thank Stuart Madnick, Bob Goldberg, Brian Pentland, Ralph Swick, Win Treese, Wendy Mackay, and Michael Schwartz for their help in conversations leading to design insights.

References

1. Akscyn, Robert; McCracken, Donald; and Elise Yoder. "KMS: A Distributed Hypermedia System for Managing Knowledge in Organizations." *Proceedings of Hypertext '87*, pp. 1-20.
2 Conklin, Jeff. "Hypertext: An Introduction and Survey." *Computer*, September 1987, pp. 17-41.
3. Coppeto, Thomas J.; Anderson, Beth L.; Geer, Daniel E., Jr.; and G. Winfield Treese. "OLC: An On-Line Consulting System for UNIX." *Proceedings of the Usenix Summer 1989 Conference*, 1989, pp. 83-94.
4. Dyson, Esther. "Computers, Customers and Hand-Holding." *Forbes*, August 7, 1989, p. 128.
5. Grudin, Jonathan. "Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces." *Proceedings of the Conference on Computer-Supported Cooperative Work*, 1988, pp. 85-93.
6. Hiltz, Starr Roxanne, and Murray Turoff. "The Evolution of User Behavior in a Computerized Conferencing System." *Communications of the ACM*, November, 1981, pp. 739-762.
7. Levitt, Barbara, and James G. March. "Organizational Learning." *Annual Review of Sociology*, 1988, pp. 319-40.
8. McCormack, Joel; Asente, Paul; and Ralph R. Swick. *X Toolkit Intrinsics - C Language Interface.* Cambridge, MA: Massachusetts Institute of Technology, 1989.
9. Orr, Julian E. "Narratives at Work: Story Telling As Cooperative Diagnostic Activity." *Proceedings of the Conference on Computer Supported Cooperative Work*, 1986, pp. 62-72.
10. Swick, Ralph R., and Mark S. Ackerman. "The X Toolkit: More Bricks for Building User Interfaces, or Widgets for Hire." *Proceedings of the Usenix Winter 1988 Conference*, 1988, pp. 221-228.
11. Trigg, Randall H. "Guided Tours and Tabletops: Tools for Communicating in a Hypertext Environment." *Proceedings of the Conference on Computer-Supported Cooperative Work*, 1988, pp. 216-226.
12. Walsh, James P., and Gerardo R. Ungson. "Organizational Memory: Structure, Functions, and Applications." Unpublished manuscript, 1989.

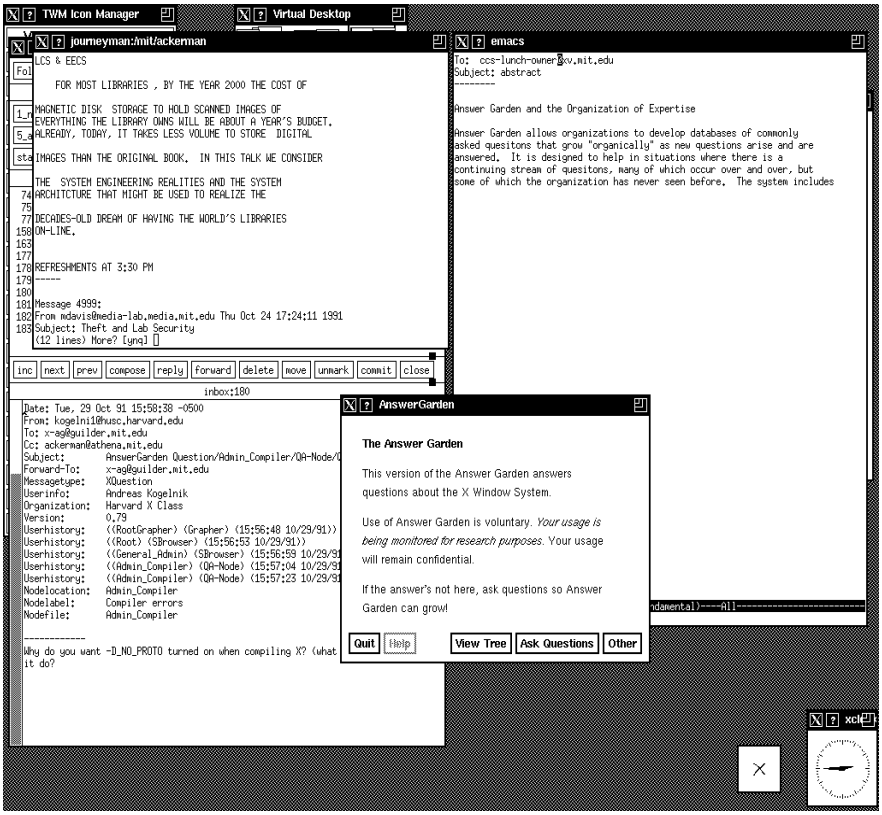

Figure 1

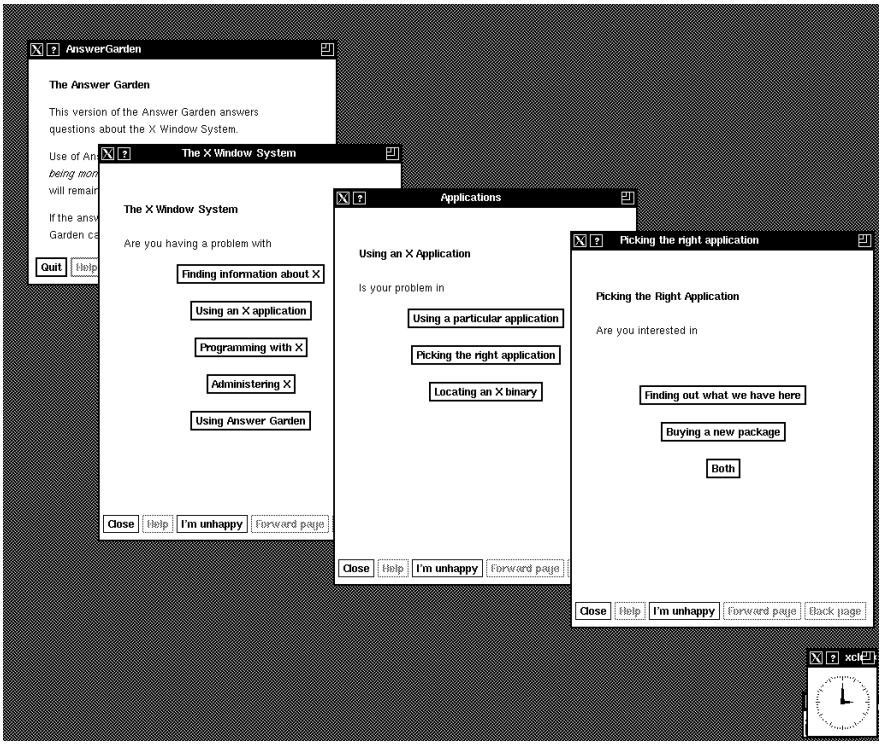

Figure 2

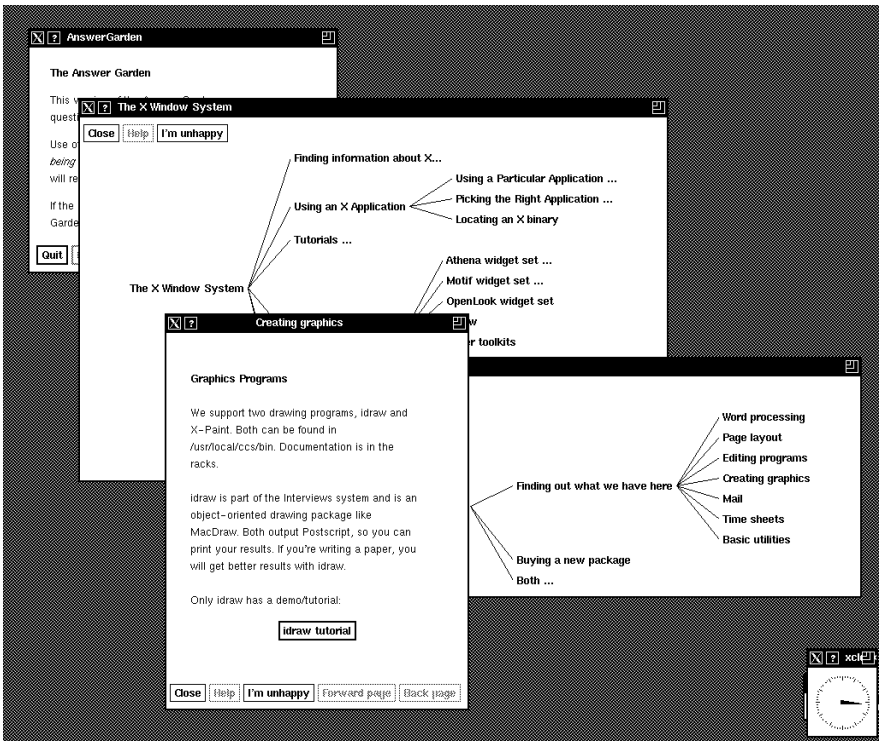

Figure 3

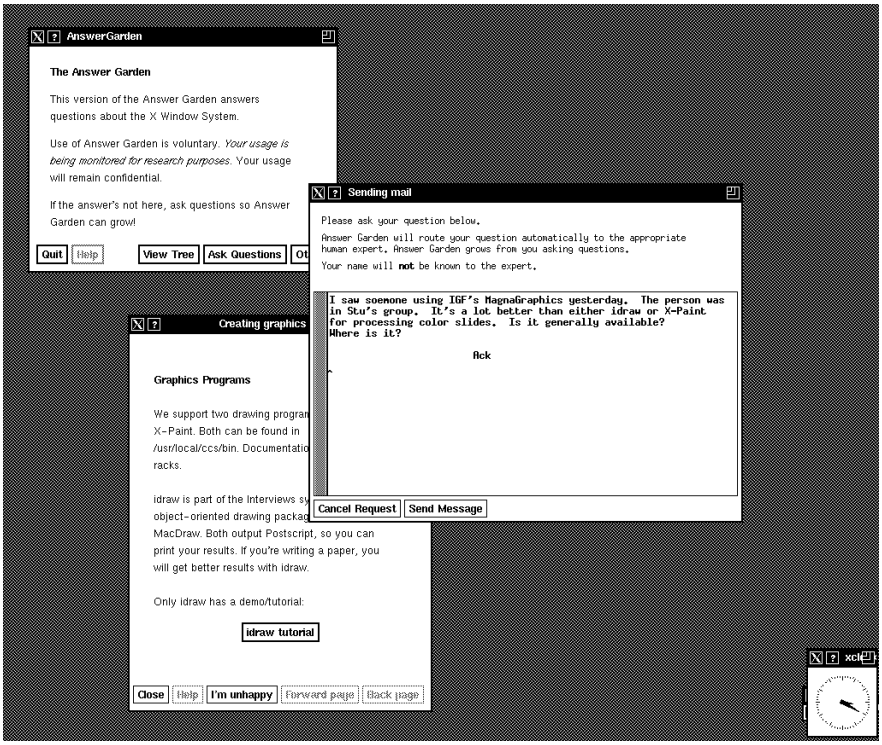

Figure 4