

Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy¹

by Jeff Conklin, Ph.D.

Abstract

Knowledge management is an essential capability in the emerging knowledge economy. In particular, organizations have a valuable asset in the *informal knowledge* that is the daily currency of their knowledge workers, but this asset usually lives only in the collective human memory, and thus is poorly preserved and managed. There are significant technical and cultural barriers to capturing informal knowledge and making it explicit. Groupware tools such as E-mail and Lotus Notes™ tend to make informal knowledge explicit, but they generally fail to create an accessible organizational memory. On the other hand, attempts to build organizational memory systems have generally failed because they required additional documentation effort with no clear short term benefit, or, like groupware, they did not provide an effective index or structure to the mass of information collected in the system. This paper explores the design of a project memory system that overcomes the barriers to capturing informal knowledge. The key component of this design is the use of a *display system* that captures the key issues and ideas during meetings. The emphasis in this approach is on improving

¹This paper was originally written several years ago. Since then, there have been many important papers and books published including Dialogue Mapping: Building Shared Understanding of Wicked Problems, Jeff Conklin, Ph.D.
Link to Dialogue Mapping book.

Keywords: organizational memory, corporate memory, organizational learning, knowledge management, teamwork, decision rationale, groupware

communication during meetings by creating shared understanding. The paper briefly describes a commercially available display system which uses hypertext to capture the thinking and learning in large, complex projects. The paper ends with a few examples of this kind of organizational memory system in action.

Introduction

Those who cannot remember the past are condemned to repeat it.
--George Santayana

The basis of the new economy is knowledge work, and the workhorse of this economy is the knowledge worker.² He or she has a strong formal education, has learned how to learn, and has the habit of continuing to learn throughout his or her lifetime.

The knowledge worker, unlike the blue collar and traditional white-collar worker, is an expert or specialist, because to be effectively applied, knowledge must be specialized. As a consequence, knowledge workers (unlike their clerk forebears) must routinely come together to solve complex problems—they work in teams.

Thus, it is not enough for a knowledge worker to be a good expert, he or she must also have the skills of *collaboration* with other knowledge workers. For example, he or she must be able to understand and be understood by people who do not have the same knowledge base, and who thus frequently have different values and a different model of the world. The productivity of a knowledge team will depend on its being able to communicate and relate despite such obstacles. However, our educational institutions little prepare us for this kind of high collaboration. As Peter Drucker says, “The productivity of knowledge work—still abysmally low—will become the economic challenge of the knowledge society” (Drucker, 1994).

² The ideas expressed in this introduction are borrowed directly from “The Age of Social Transformation,” by Peter F. Drucker, in *The Atlantic Monthly*, November 1994.

A *knowledge organization* is one in which the key asset is knowledge. Its competitive advantage comes from having and effectively using knowledge. Examples include the law office, accounting firm, marketing firm, software company, most government agencies, universities, the military, and significant parts of most manufacturing companies, whether they make cookies or cars. Knowledge organizations are the container for knowledge workers, the vehicle through which they apply their knowledge.

Increasingly, knowledge workers are the scarce resource in the economy, moving freely within and between organizations. This fluidity presents a new challenge for knowledge organizations. In the industrial era, factory machines were the scarce capital resource, and the people to run them were regarded as interchangeable. Moreover, organizations acquired assets through capital, depreciated them, and finally sold or junked them.

The dominant asset of the knowledge organization, however, is knowledge. Intellectual assets belong inherently to people, and are the organization's assets only through their application, capture, and reuse. If the people are unhappy, unmotivated, or unskilled in the art of collaboration, their precious intellectual assets are, from the organization's perspective, wasted. When these people leave, a valuable asset leaves with them.

Knowledge is the key asset of the knowledge organization. Organizational memory extends and amplifies this asset by capturing, organizing, disseminating, and reusing the knowledge created by its employees.

There are good reasons to pursue creating organizational memory³. If a person had a memory like the average organization, we would think he was very stupid, or suffering from a neurological disorder. Organizations routinely "forget" what they have done in the past and why they have done it. These organizations have an impaired capacity to learn, due to an inability to represent critical aspects of what they know.

³ The term "organizational memory" is sometimes used to refer to whatever exists today in the way of social conventions, individuals' memories, etc. In this paper the term refers to a new capacity for organizations, an augmented memory that is based on information technology.

But organizational memory is not just a facility for accumulating and preserving but also for sharing knowledge. As knowledge is made explicit and managed it augments the organizational intellect, becoming a basis for communication and learning. It can be shared among individuals working alone, by teams needing a project memory, and by the organization as a whole for *between-team* coordination and communication. “Given the nature of organizations and the competitive environment within which they exist, organizational learning and the accumulation of knowledge will be a source of immediate health as well as long-term survival.” (McMaster, 1995, p. 113)

Project memory is simply organizational memory for a project team; it has a more limited scope than organizational memory, but it is easier to implement and its benefits are easier to measure. This paper explores the design of project memory as an evolutionary stepping stone to organizational memory⁴. The key contribution of this paper is the power of using a *display system* in meetings as a transparent way to capture knowledge and to enhance the team problem solving process. Knowledge teams suffer much less rehashing and repetition when the team uses a display system to make its thinking explicit during the problem solving process. Meetings become much more efficient. Being more explicit also increases rigor and improves coordination. New people can come up to speed on a project much faster by reviewing the store of project knowledge; and when a team member leaves, at least some of his or her informal knowledge stays with the team.

This paper diverges from much of the previous work on organizational memory by placing its emphasis on the knowledge worker and the knowledge team, not the organization. It explores the question:

**How can we design a project memory system to
serve the immediate needs of knowledge workers
first, with secondary emphasis on the
organization’s needs?**

⁴ Overall, I will use the broader term “organizational memory” for the general discussion in Part One, and the narrower term “project memory” in the Part Two.

If project memory is so clearly beneficial, why is it not commonplace? Part One of this paper lays some conceptual groundwork and then discusses four barriers to creating an effective project memory: (1) informal organizational knowledge, like a wild animal, resists capture; (2) the usual approach to project and organizational memory, preserving documents, fails to preserve *context*; (3) knowledge loses its relevance, and thus its value, over time; and (4) the current litigious environment may create an economic incentive for “*organizational amnesia*”.

The second half of the paper ventures into exploring the design of a project memory system which takes the above challenges into account. What are the components and features of an effective project memory system? Part Two introduces the concept of a *display system*, which provides a knowledge team with a simple, practical structure for creating shared meaning and shared understanding.

The emphasis in this paper is on capturing knowledge in a way that is both nearly transparent and immediately beneficial for knowledge workers. While memory includes both storing and retrieving information, the problems of retrieval only come up once you have successfully solved the problem of capture. If capture is not transparent—if it involves extra work—it is not sustainable in a practical way.

The paper ends with a few examples of an organizational memory system in action, from clients who have been using a commercial product which implements the design ideas laid out in this paper.

Part One: Barriers to Creating Organizational Memory

The world that we have made as a result of the level of thinking we have done thus far creates problems that we cannot solve at the same level that existed when we created them.

--Albert Einstein

Knowledge is the small part of ignorance that we arrange and classify.

--Ambrose Bierce

Formal and Informal Knowledge

In order to understand why knowledge is hard to capture we must first distinguish two kinds of knowledge: formal and informal. Formal knowledge is the stuff of books, manuals, documents, and training courses. It is the primary work product of the knowledge worker, in the form of reports, white papers, plans, spreadsheets, designs, memos, etc. Knowledge organizations easily and routinely capture formal knowledge; indeed, they rely on it—without much success—as their organizational memory.

But there is another kind of knowledge as well. It is the knowledge that is created and used in the *process* of creating the formal results. If formal knowledge is the foreground, this knowledge is the background. It includes ideas, facts, assumptions, meanings, questions, decisions, guesses, stories, and points of view. It is as important in the work of the knowledge worker as formal knowledge is, but it is more ephemeral and transitory. This kind of knowledge is “wild”—it is hard to capture and to keep. Let us call this process-oriented stuff “informal knowledge.”

We can understand these two kinds of knowledge better if we recall a similar distinction — between matter and energy. One of the most important experiments in quantum physics shows that elementary particles such as photons and electrons have a dual nature, as things (i.e., matter) and as waves (i.e., energy). In the “double-slit experiment” a beam of electrons passes through a pair of slits in one surface before falling on a back surface. If one of the slits is closed, the beam projects

the image of the open slit onto the back surface, reflecting the “thingish” nature of the stream of electrons. But when both slits are open, the beam projects a refraction pattern onto the back surface, reflecting the wave nature of the electrons.⁵

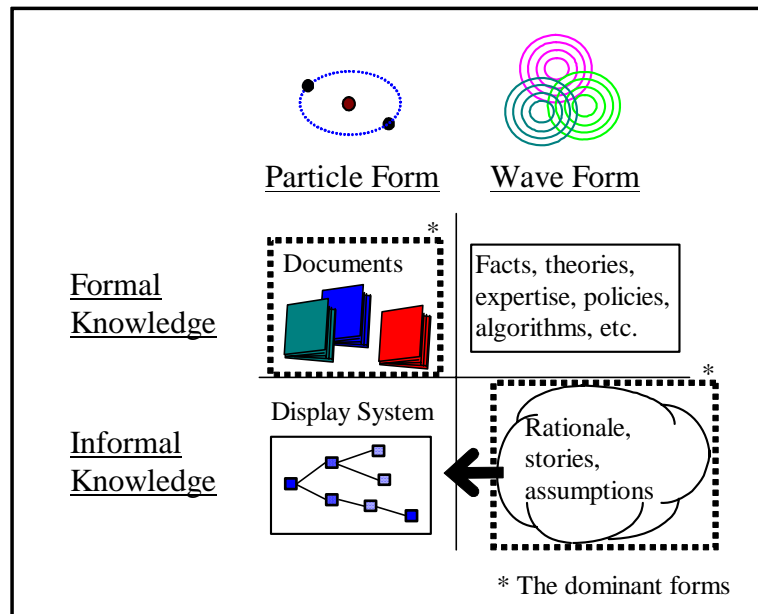


Figure 1: The Knowledge Form Matrix

Similarly, knowledge (both formal and informal) is particle-like when it is written down, such as in a report, or encoded in an information system. But knowledge has another form whose nature is much more like a dynamic force or energy: the interaction of a group of people in a context that leads to action. As with light, both forms of knowledge are valid, each form has contexts in which it makes the most sense, and in the course of knowledge work the action shifts back and forth between these two manifestations.

The distinction between particle and wave forms of knowledge illuminates the present reliance on formal knowledge for organizational memory (see Figure 1: The Knowledge Form Matrix). Formal knowledge usually

⁵ Even more intriguing is that, “on its way through the slit, the electron ‘knows’ whether or not the second slit is open. It knows what the scientist is testing for, and adjusts its behavior accordingly.” For more details, see (Wheatley, 1992, p. 62).

manifests in its slower, heavier particle form, typically in a document of some kind. Informal knowledge, however, rarely shows itself in its particle form; it usually manifests (usually in conversations and meetings) as fast moving, invisible wave energy. Informal knowledge can be described as a standing wave or refraction pattern created by the wave fields of the participants.

Embracing Process

Informal knowledge is as an organizational asset of immense value, because much of what is being created by and shared among knowledge workers never makes it into formal documents. The informal knowledge contains the background context for the organization's formal documents. For example, informal knowledge answers such questions as "Why did we do it that way?", "What would happen if we stopped doing such and such?", "Hasn't this problem been solved before?", "Did anyone consider trying this other approach?", and "What did we learn the last time this happened?"

**These are the questions that send project teams
scurrying to reinvent the wheel and repeating
discussions that have been "closed."**

For all of the value of informal knowledge, the current practices of knowledge work fail to capture, share, and reuse this asset. Because it is often invisible (like most energy forms), informal knowledge tends to be viewed as "just talk," "soft information," or as nothing at all.⁶

One reason for the widespread failure to capture informal knowledge is that Western culture has come to value results—the output of the work process—far above the process itself, and to emphasize things over relationships. In other words, we have an *artifact-oriented* culture, and

⁶ This is analogous to operating a lumber mill as if the wood chips and sawdust were a waste product, to be hauled off and burned. By changing the process, for example, by gearing up to produce particle board, the mill might create a new product, thus capturing the value inherent in the wood chips and sawdust.

this orientation is reflected in the way we work. Within the artifact-oriented value system, formal knowledge—knowledge in its particle form—is what counts. Informal knowledge, being wave-like and thus generally invisible, is devalued and ignored.⁷

If our organizational memory system is to include informal knowledge, we must find ways to render the informal wave energy such that it can be represented in information systems. The arrow in the lower half of Figure 1 suggests that certain kinds of informal knowledge, such as decisions, rationale, stories, and assumptions, can be captured in their particle form in a *display system*, as described in Part Two.

Another reason that informal knowledge is lost is that the tools of knowledge work—based on computer and communication technology—little recognize or support the process of knowledge work. Our tools for knowledge work reflect the artifact-oriented ontology of our culture, and create the illusion that finished knowledge products (in the form of documents) spring forth from the word processor polished and complete. These tools systematically ignore the expensive thinking and learning that underlie the formal work product. In the pre-computer office, for example, when paper documents were circulated for review, valuable informal knowledge accumulated on the document in the form of notes in the margins. Such marginalia are less likely to be placed in a document, or to be preserved with it, in a word processor.⁸

An organizational memory that consists only of formal knowledge is bare and lifeless. It is like describing the ball game by giving the statistics, or the mystery novel by simply relating the plot outline. It also lacks the history and context behind the formal documents, and as a result, the organizational “memory” is essentially an immense heap of disconnected things, a giant organizational attic. Documents that contain formal knowledge that the organization has paid dearly to create live somewhere

⁷ In general, wave phenomena, such as informal knowledge, compassion, and community, are regarded as being less real. This is beginning to change as we deepen our understanding that a quality product can only be produced by a quality process.

⁸ Some word processors allow electronic “annotations” to be made to a document, but these are more like “optional footnotes.” To this author’s knowledge, this capability is not widely used for teamwork.

on the corporate network with enlightening names like
“H:\org\finan\arc\drg\693plan.doc.”⁹

Because the organization’s formal knowledge repository—the default organizational memory—is dead and without context, it is often ignored as a resource. If, however, an organization embraces its informal knowledge, then the rationale behind decisions and documents becomes the glue¹⁰ that holds the formal knowledge documents together and preserves their meaning (Conklin, 1993).

Attempts at Organizational Memory

Error is discipline through which we advance.

--William Ellery Channing

There was a project at an aerospace company some years ago in which the team decided it would capture its project memory. They preserved official reports, design documents, presentations, memos, meeting minutes—virtually anything that they wrote down. At the end of the project they had indeed created a project memory: an office completely filled with stacks of paper, extending almost to the ceiling. Even if someone thought there might be valuable information stored in that room, no one ever wanted to go in there and try to find it.

Other project teams have recognized the importance of preserving the informal knowledge involved in the project. The team leader—keen on the importance of capturing informal knowledge—instructs the team members to write down important ideas, decisions, notes, and communications,

⁹ The current solution—document management systems—does a better job of organizing the formal documents, but still relies heavily on *search* as the mechanism for finding things. Without the help of context, brute force search become less effective as the number of stored documents grows. The web of contextual relationships to other documents is also lost, or mostly lost.

¹⁰ Decision rationale and other forms of informal knowledge must be tightly integrated with the artifacts of concern—each one relies on the other to make sense. Moreover, there must be a smooth transition in both directions. Because people do orient to the artifacts of their work, the relevant informal knowledge must be indexed and accessible directly from those artifacts.

either on paper or in email. Everything goes along fine until a deadline or milestone approaches and the pressure increases—then the team quietly drops these extra documentation duties in favor of doing their “real” work—producing formal artifacts. Because of this, the ideas created and the decisions made during this highly productive crunch phase of the project are not captured.

A few project teams have attempted to capture their thinking and learning by audio or videotaping their meetings. Inevitably these teams end up with a staggering volume of tape. The key bits of knowledge they need later on are in there somewhere, but who has the time to watch or listen to it all to find them?¹¹

A few projects have had the luxury of a project historian or librarian—someone whose job is to capture and organize the knowledge created in the course of the project. Unfortunately, these jobs don’t seem to last very long. Without a way to capture the informal knowledge as well, and without an organizational commitment to accessing and using this knowledge, it is just an expensive way of filling the attic. Moreover, historians and librarians tend to be regarded by project teams as outsiders, which can further complicate their job.¹²

The concept of creating, archiving, and using learning histories, as part of the research and engineering function, is an old one. It is considered good engineering practice to create reports documenting “lessons learned” on a project. However, even in companies where this practice has become part of the “standard operating procedure,” it is very difficult to find instances of the resulting document actually being referenced in the next project, or being referred to by subsequent projects.

More recently, some organizations have attempted to use groupware tools such as Lotus Notes™ to create and manage organizational memory. Groupware is designed to be used for informal communications, and thus has the potential to become a repository for an organization’s informal knowledge. Unfortunately, groupware messages and documents tend to

¹¹ Some recent research suggests improved technology for capturing and indexing video clips of project history and rationale (Carroll et al, 1994; Minneman et al, 1995), but the fundamental challenge of indexing all of these snippets remains.

¹² For an account of one project historian’s experience, see (Shum et al, 1993).

lack any structure, so the repository that gets created is just an electronic version of the “attic full of stuff.” Most groupware applications include some search capability so that users can search the database for particular keywords. However, if the groupware application has caught hold, and is thus heavily used, its database is usually too large and too jumbled and incoherent for retrieval to be very satisfying or successful.

These experiments in capturing organizational memory paint a gloomy picture of the prospect of preserving the most precious asset of the knowledge organization. They illustrate that you can’t create a useful memory store just by capturing lots of information, you must somehow organize it in ways that create and preserve coherence and “searchability.” These experiments also illustrate a fundamental tension in the design of an organizational memory system. When does this information organizing, structuring, and indexing work get done?

Most current implementations of organizational memory postpone this organizing effort as long as possible, or try to do it automatically in the background (e.g., using artificial intelligence techniques). Neither of these approaches can capture the critical informal knowledge that gives the information context and has it make sense.

For all its potential, we have not yet found a way to tap the value in an organization’s informal knowledge. My conclusion is that the creation and use of organizational memory cannot be a by-product, an extra bit of work hanging on the side of the knowledge organization’s main production process.

If we are to find ways of preserving the asset of informal knowledge, we must look within the practices of everyday teamwork and change them.

Creating an effective organizational memory system entails creating new tools and new practices, making changes in technology as well as culture¹³.

The final sections of Part One outline the four chief barriers to effective organizational memory.

1. Making Informal Knowledge Explicit

A world community can exist only with world communication, which means something more than extensive [communication] facilities scattered about the globe. It means common understanding, a common tradition, common ideas, and common ideals.

--Robert M. Hutchins

One of the biggest challenges facing teams is effective communication. Almost all group dysfunctions come down to an inability or unwillingness to say what we really mean or to hear and *understand* what another says.

The Holy Grail of teamwork is shared understanding.

A knowledge worker, as a specialist, lives in a rich and complex world defined by his or her area of specialization. When knowledge workers come together in teams, each person's depth of experience and knowledge comes into play in the team's knowledge base. But these differences in expertise can also block shared understanding just as surely as if each person were speaking a different language. To overcome this hurdle the team must slow down and take the time to understand each other. Indeed, they must actively create a framework of shared understanding, especially about key concepts and terms, so that their

¹³ This paper focuses on the technology issues, not the issues of changing culture. But I believe that technology innovation and culture change efforts must be designed to synergize with each other. For a discussion of a culture change approach that is consistent with the ideas in this paper, see (Eppel & Conklin, 1995).

collaboration is not repeatedly undermined by subtle but disastrous miscommunication.

Lack of shared understanding can be debilitating. Too often, team members see the problem they are working on differently, then attribute the differences to incompetence or hostility in the others, and so shift into a “battlefield mentality” of protecting turf and taking new ground when possible. Some level of competition may be healthy, but a knowledge organization cannot afford the sheer waste of human energy and gumption that results from the prevailing level of skills and tools for group communication.

One element of creating shared understanding is making informal knowledge explicit. This means surfacing key ideas, facts, assumptions, meanings, questions, decisions, guesses, stories, and points of view. It means capturing and organizing this informal knowledge so that everyone has access to it. It means changing the process of knowledge work so that the focus is on creating and managing a shared display of the group’s informal thinking and learning.¹⁴ The shared display is the transparent vehicle for making informal knowledge explicit.

A remarkable thing happens when knowledge teams use a kind of shared display, a *display system* (described below, page 27), to treat informal knowledge as if it were valuable. Not only is a coherent record of the team’s thinking and learning created, but the team actually works better and more productively. There is less repetition in meetings, more rigor in decisions, and it is easier to bring others (e.g. new team members, management, and other stakeholders) up to speed on the team’s thinking and learning. In other words, when you take process-oriented knowledge seriously, the process itself immediately improves. More bluntly, a powerful way to avoid organizational stupidity is to take the process of knowledge work seriously enough to capture and share the informal knowledge involved.

¹⁴ Some readers may be concerned that formalizing the “soft stuff” like this can, in the process, destroy it. This is always a risk with language, but the approach I propose below respects the need for inconsistency, incompleteness, ambiguity, and all of the other “messy” qualities of the rich process of human communication. In particular, this approach does not seek to formalize or quantify the decision-making process.

2. Documents without Context

The second barrier to effective organizational memory is that the usual approach to organizational memory, preserving documents, fails to preserve the *context* which gives the documents meaning, the very thing that allows them to be useful in the future, when the context has changed.

Because current notions of organizational memory assume a repository of artifacts, they focus on preserving, organizing, indexing, and retrieving only the formal knowledge as it is stored in documents and databases. For some tasks, formal knowledge alone is sufficient; for example, when it is time to write the new annual report, you might start with last year's annual report as a template.

However, most knowledge work is performed in the quest for solutions to “wicked problems” (Kuntz & Rittel, 1972), problems for which there is no clear and agreed upon definition of the problem, and, indeed, in which the problem itself is apt to change over time. Wicked problem solving is characterized by making lots of assumptions, educated guesses, and decisions under conditions of uncertainty. Decisions must frequently be revised or even retracted. In contrast with the linear techniques¹⁵ that have been adequate for solving “tame” problems, wicked problems require both traditional linear techniques and a heavy dose of *social interactions*: conversations, meetings, presentations, phone calls, email, etc. When you are working on a wicked problem, your primary goal is not so much to find a “right answer” as to find a solution—and an understanding of the problem—which has broad ownership.

In this context, an “attic” of formal documents is simply not rich enough to support knowledge work. For example, a team may come together for many meetings in the course of resolving a wicked problem, but the practice of creating and circulating meeting minutes is a relatively blunt instrument for creating continuity and coherence among these meetings. Meeting minutes are sketchy, represent only one person's point of view,

¹⁵ Traditional wisdom dictates a linear, or serial, problem solving approach: (1) define the problem, (2) gather the data, (3) analyze the data, (4) formulate a solution, (5) implement the solution. Ideally, you visit each of these steps once, using the output of each step as the input to the next.

and usually lack the energy and context of the conversations they were meant to capture.

Because projects devoted to resolving wicked problems can often stretch into months and years, a project memory for informal knowledge becomes more than just a good idea. As the meetings spread over days and weeks one can soon sense the number of ideas that are getting repeated over and over, and the growing number of issues that are slipping through the cracks. An explicit project memory provides more continuity among these sessions, allowing the group to pick up where it left off, with a minimum of repetition and loss of important issues. As team membership changes over time, or the project is handed off to a completely new team, the project memory can in principle reduce the likelihood of false starts and duplication of previous work.

As the number of ideas and issues accumulates, it soon overwhelms the memory capacity of even the smartest team members. How often is progress in a meeting blocked by disagreement over what was discussed or concluded at a previous meeting? Is it not strange that we accept this state of affairs as somehow normal and inevitable? A shared memory for the group creates coherence within the mass of formal and informal project knowledge, allowing everyone to relax, focus on shared understanding, and think more creatively.

Moreover, knowledge work that bears on wicked problems requires tools and processes which allow the group to explore a variety of interpretations and points of view about the nature of the problem, the essential solution criteria, the key stakeholders, and the constraints. It must facilitate the group's evolution toward a shared understanding and shared goals. It must help the group create a *story* about what they are doing, thinking, and learning that will facilitate understanding and buy-in as the group's solution moves out into larger circles of stakeholders. The project memory thus becomes a living document that tells the story of the project. It preserves the *context* of the work as it evolves. This project memory is most naturally represented in the form of a web of information which includes facts, assumptions, constraints, decisions and their rationale, the meanings of key terms, and, of course, the formal documents themselves.

3. Relevance and Size

Many are always praising the by-gone time, for it is natural that the old should extol the days of their youth; the weak, the time of their strength; the sick, the season of their vigor; and the disappointed, the spring-tide of their hopes.

--Caleb Bingham

For all of its frailties, human memory has an extraordinary capacity for **relevance**. The third challenge for an effective organizational memory system—a system that includes informal knowledge—is that knowledge tends to lose its relevance, and thus its value, over time. Informal knowledge, being more contextual and wave-like, is even more dynamic in this way. An organizational memory system should therefore, like human memory, have the capacity to recall whatever is relevant and salient to the moment¹⁶. Closely related to this is the problem of the sheer size of organizational memory. There will be **inconceivable** volumes of corporate knowledge accessible on-line in the near future—if only you could find the specific bits of knowledge that are relevant to your immediate problem! This is already the case on the World Wide Web.

One popular view has it that the way people preserve relevance is by uncluttering their minds—forgetting most of what they have learned. But this is not true in human cognition. Human memory does not deliberately flush old information to make room for new, nor does it overwhelm you with too much information when you are trying to recall something—there is no problem (in everyday experience) with “remembering too much” or “knowing too much.” We often think highly of people with good memories, and pity those who are forgetful. Forgetting is an aspect of memory, but the mechanism is not a purging but a gradual fading over time through disuse.

On the other hand, there is also a risk that too much memory can make an organization stupid, as in the frustration of having an idea dismissed with “That won’t work—we tried it before.” It seems sometimes that the group would be smarter if it didn’t remember so much about the past.

¹⁶ The process of relevant recall may be quite *active*: more like reconstruction than retrieval. For more on this, see (Bannon & Kuutti, 1996).

Many a man fails to become a thinker only because his memory is too good.
--Nietzsche

The problem is not remembering too much but *misapplying* knowledge from the past. Simply being able to recall the past is not enough. In the face of the need for ever greater capacity for innovation, one must also be smart about how old knowledge should be applied.

The size of organizational memory can be a considerable obstacle, and relevance is a delicate thread. Current searching and filtering mechanisms still perform poorly in the face of a huge amount of textual information. The heart of the problem is that human memory has an extraordinary capacity for *meaning*, and thus for relevance. Unlike computers, biological memory appears to be holographic in nature: a given piece of knowledge is not stored in a single address or data field, it is distributed all over the brain. And unlike computers, human memory is associative in nature: a given piece of knowledge is accessed through a rich network of semantic associations. Biological memory gains an extraordinary degree of resiliency and flexibility through its holographic and associative neural network mechanisms.

How can we create analogous mechanisms for organizational memory, which allow it to preserve vast amounts of information while providing discerning retrieval of relevant knowledge when and where it is needed? This is an important issue for project memory (memory within a single project), but it is a critical one for the memory of a whole organization.

Here is an example. A team is working on developing a new kind of valve for an oil recovery device that their company manufactures. The valve must be able to operate at very cold temperatures, and not get stuck even if the oil flowing through it turns very viscous. The specialists on the team (in mechanics, fluid dynamics, cryogenics, etc.) have standard resources they use to get the basic formal knowledge they need for the design. But, generally, they proceed as if they are working in a kind of vacuum. How do they know if someone in the company has tried to create such a valve in the past? Perhaps an earlier effort was mounted but failed; perhaps it was technically successful but put on the shelf; perhaps someone showed that it was a practical impossibility and abandoned the effort. Perhaps there is, somewhere else in the company (or in a partner company) a

specialist who has thought about this problem, but not published any of her results.

For a computer system to augment this team's memory, it should have the following features:

- it makes it easy and natural to review similar cases and projects for information that would be helpful now;
- it avoids false hits by paying attention to relevance and meaning (e.g. On the valve project, doing a search for “oil” and “valve” does not return an heap of documents about cholesterol and heart disease);
- at the same time as avoiding the false hits, it finds items which are related in interesting ways (e.g. A search for “oil” and “valve” does return items dealing with “refrigeration plumbing”);
- it preserves context, by providing that retrieved “fact-oids” come with the context in which they originally made sense;
- it transparently captures whatever knowledge this team creates and automatically adds it to the organizational memory;
- as meaning evolves over time, the links and indexes in the system evolve correspondingly.

The oil valve project team is not alone: the same problems confront knowledge teams who need to know who else in their enterprise has worked with a given client, or has struggled to comply with a given law or regulation, or has used a given piece of equipment or software, or has bought from a given vendor.

Technically, there are exciting possibilities for the use of hypertext, groupware, intelligent agents, neural networks, advanced search techniques, genetic algorithms, and other computing technologies to provide “relevance retrieval” access in large databases—retrieval which respects the meaning relationships among the stored items. However, in

terms of providing the features listed above, these technologies are still experimental¹⁷.

For the near future at least, human intelligence and effort will remain a key component of the kind of intelligent retrieval that respects meaning and relevance. Some level of human expertise will be required, such as a librarian who can track subtleties of meaning and help with the indexing and structuring of the organizational memory. Moreover, as language and meaning evolve over time, some intellectual work must go into the re-indexing and re-structuring necessary to keep the organizational memory from becoming a historical curiosity.

Social mechanisms can also be used to assure relevance and meaning. In Japanese corporations, knowledge workers have many different jobs over time, moving around the organization so that they become part of a rich human network of experience and knowledge. Japanese corporations also have “gate keepers,” people whose whole job is technology transfer, i.e. cross-fertilization of knowledge among divisions and from outside the corporation. This practice seems to reflect a commitment to knowledge “retrieval” in the widest sense.

It is important to bear in mind that the problems of retrieval (size, meaning, relevance) only come up once you have successfully solved the problem of creating an organizational memory that contains knowledge that workers need and want. Creating such a memory store requires capturing the knowledge and organizing it effectively. Thus the senior challenge is transparent capture that preserves relevance and meaning, in other words, capture of informal as well as formal knowledge. Capture of informal knowledge can easily be, and generally has been, prohibitively expensive, as described in the previous sections (“Attempts at Organizational Memory,” page 10, and “2. Documents without Context,” page 15). However, if you shift the emphasis from “how to capture all this stuff?” to “how to improve the process of teamwork such that capture happens by itself?” you get a surprising answer. The thrust of this paper is that use of a display system to increase shared understanding—by

¹⁷ One particularly interesting experiment using an Email-based system for sharing expert’s answers to common questions is described in (Ackerman, 1994).

making informal knowledge explicit—improves the knowledge work process, and, as if by happy accident, solves the capture problem as well.

4. Litigation and Organizational Amnesia

If a man harbors any sort of fear, it percolates through all his thinking, damages his personality, makes him landlord to a ghost.
--Lloyd Douglas

O friend, never strike sail to a fear! Come into port greatly, or sail with God the seas.
--Ralph Waldo Emerson

There is a fourth barrier to organizational memory that should be mentioned. Spurred by their legal departments, a few American corporations are adopting a policy of “organizational amnesia”: the systematic destruction of all unneeded personal notes and documents at regular intervals. The thinking behind this policy is that, in the event of litigation or criminal prosecution, it is dangerous for anything to exist in writing that could be used against the corporation. Since the legal mechanism of “discovery” allows lawyers from the outside access to *any* documents that are not explicitly protected under “client attorney privilege,” the risk of expensive judgments against the corporation may have created an economic incentive for amnesia.

Such thinking, where it exists, creates a major obstacle for the creation of organizational memory. It insists that only the most formal and sanitized forms of knowledge may be allowed to persist. It puts everything that is written down or stored in a computer under the lens of “can this information possibly be used against us.” Most adults know that you learn the most if, when you make a mistake, you acknowledge it and reflect on what you have learned from it. But in an organizational amnesia environment, mistakes must be avoided at all costs, and denied if they occur. How can organizational learning possibly take place in this environment?

Time will tell whether this kind of anti-memory policy is cost effective in the long run, and whether companies that pursue it can compete as knowledge organizations. Of course, following the guidelines in this paper, one could still use a display system to build shared understanding in knowledge teams, but then simply erase that display (and all records and copies of it) at regular intervals. This temporary “memory” would at least promote shared understanding on a project, but not long term nor between-team memory.

To summarize, the obstacles to an effective organizational memory system fall into two categories, cultural and technical. The cultural barriers include the following: a cultural emphasis on artifacts and results to the exclusion of process; resistance to knowledge capture because of the effort required, the fear of litigation, and the fear of loss of job security; and resistance to knowledge reuse because of the effort required, and the low likelihood of finding relevant knowledge. The technical barriers include: how to make the knowledge capture process easy or even transparent, how to make retrieval and reuse easy or even transparent, and how to ensure relevance and intelligibility (i.e., through sufficient context) of retrieved knowledge.

Part Two: Steps Toward Organizational Memory

So far we have considered the obstacles for creating an effective organizational memory. Part Two offers a solution to the problems outlined above, provides a brief theoretical justification for it, and gives some examples of its application. The solution offered focuses on creating project memory, but also expands the design space for organizational memory systems.

Memory Mediated by Short Term Store

In our search for how to design an effective project memory system, it is worthwhile to consider, as metaphors, the operation of three other information processing systems which possess long term memory:

computer architecture, human cognition, and living cells. These metaphors suggest a theory about the missing link in project memory.




Figure 2: Simple model of computer memory

Computer memory basically has two components, in addition to the processor (see Figure 2). The processor works directly with RAM (Random Access Memory), and the disk (or network) is used for long term storage of data.¹⁸ The important feature of this model is that *RAM mediates memory*. The processor only works with data stored in RAM. All data must be read from the disk into RAM, where it is processed, changed, or displayed to the user. Similarly, any data that must be kept between “processing sessions” must be stored on the disk, and must go through RAM to get there. Nothing gets into or out of the disk without going through RAM.

Current theories about human memory tell a similar story. In the simplest model, human memory has two components: Short Term Memory¹⁹ (STM) and Long Term Memory (LTM). Of course, there is the component of cognition that corresponds to the computer processor: Processing (see Figure 3). (I call it “Processing” to sidestep such questions as “Is there a *processor?*”, “Where in the brain is the processing taking place?”, etc.) STM participates directly in cognitive processing (e.g. thinking, reflecting, acting, and so on), but its size is very small²⁰. LTM is used for permanent “storage” of information and experience.

¹⁸ Strictly speaking, the processor in this case is a *program*, the short term store is the program’s *data structures*, and the long term store a *file or database* on the disk or network.

¹⁹ Sometimes called Working Memory.

²⁰ Cognitive scientists believe the size of human STM is about seven “chunks,” plus or minus two, depending on expertise and other factors. The contents of an STM memory chunk can be quite complex, but it must be a single thing—it must have enough internal coherence that it only takes up one chunk. The concept of “tigers” is complex, but in a list of animals it only takes up one STM chunk.




Figure 3: Simple model of human memory

Again, the point is that *STM mediates memory*. To “commit something to memory” is to focus processing on it in STM with enough repetition that a permanent trace of it is created in LTM. To “recall something” is to bring it from LTM back into STM. Nothing gets into or out of LTM without going through STM.

Life has a long term memory, as well: genetic information stored in the DNA in the nucleus of the cell. Here, again, the mechanism by which the genetic information is accessed involves an intermediate store. In this




Figure 4: Simple model of cellular memory

case, *RNA mediates memory* (see Figure 4). All of the cell’s metabolic functions (the “processing” in this system) are performed by proteins. The blueprint for the structure of each protein, i.e. its sequence of amino acids, is stored in the DNA of the cell, but is not retrieved directly from the DNA. Instead, the DNA is “transcribed” into a complementary molecule called RNA. The RNA is then used as the template from which the proteins are manufactured. And again, nothing gets into or out of DNA without going through RNA²¹.

²¹ Life also uses RNA as an intermediate in “storing” new information, but the story is a bit more complex. Briefly, organisms “remember” by successfully creating offspring; at the moment of conception, RNA is used to bind the DNA strands from each of the parents into a new DNA pattern, the blueprint for the offspring.




Figure 5: Pattern of mediated memory in information systems

Thus, there seems to be a pattern for memory in which there is some kind of fast, flexible intermediate store²² between the processing component of these information systems and their stable long term memory (see Figure 5).




Figure 6: How project memory is currently implemented

However, virtually all implementations of project memory are structured without the short term store—we attempt to move information directly from the immediate work process directly into storage (be it bookshelf, file cabinet, or computer database), and directly from storage into the work process (see Figure 6).

The problem with this approach is that, in the pattern, the short term store performs at least two important functions: it provides the processing element with a well-defined information focus (a small and discrete subset of long-term storage), and it provides a staging ground which structures and indexes the information for long-term storage.




Figure 7: Proposed model of project memory

²² In computer science terms, memory that serves this function is called a “cache.”

One could object that, in the prevailing model, project memory has an *operational* short-term store that is the desktop or computer screen of whoever is accessing the organization's long-term memory. That person may have taken a document off their shelf or out of the corporate library, or they may be looking at a document on-line, or they may be reviewing the results of a search they have performed in some database. But these examples really just illustrate information moving directly from long term storage into the work process, without any intermediate structure or store. Also, these examples only illustrate retrieval of formal knowledge items by a lone individual. In the terms of the pattern for memory systems, there is currently no *structure* that mediates between group knowledge work and the organization's vast memory (see Figure 7). The central proposal of this paper is that the missing ingredient for project memory is a structure that functions as an "project short term store," and that this mediating structure can also greatly facilitate the group problem solving process.

Criteria for Project Short Term Memory

It is beyond the scope of this paper to explore why it is that memory seems so often to use an intermediate store, but this pattern may be helpful in our design efforts. What do we need, by metaphorical extrapolation, to create and use an "project short term store" that mediates project memory? This structure would have similar criteria to those of STM, RAM, and RNA:

Criterion	Description of Project Short Term Memory
A	It would provide a small, high-speed store, the structure of which supports and enhances both individual and workgroup processes.
B	It would structure and index the information held in it to be moved easily into, and back from, long term storage.
C	It would provide just the right subset of long term storage that is needed for the task at hand.

D	It may operate by a “constrained copy” mechanism, in which processing never makes changes directly to the long-term item, but to a local copy of it, which, when processing is done, gets committed back to long-term store. (This mimics how computer memory, at least, works.)
---	--

There may be other criteria for a mediating short term store as well. The essential point is that our efforts to design a project memory system should be informed by the design of other kinds of memory systems.

Display Systems

What sort of structure might act as a project short term store? I propose that a special kind of shared display for team meetings called a *display system* fits the criteria. A display system has three components: *capture* of information into the system, a *structure* by which the information is organized, and a representation and *display* of that information, usually to a group. A simple example is a checkbook register: by recording checks you capture the information in the system; the structure includes date and amount of the check, and who it is written to, but not, for example, where you were when you wrote the check or how you felt about it; the display in this case is the register itself. When used to facilitate a meeting, a display system acts as a *shared space* in which the thinking and learning of the group is recorded and displayed. It is beyond the scope of this paper to explore in any depth the impact of such a shared display on group dynamics, but here are a few immediate benefits:

1. As a person's comments are captured in the display, they feel heard;
2. As the display grows, it is easier to relax and reflect on the problem, because you don't have to be trying to keep track of all the points being made in your head;
3. If the topic is contentious, or there is internal conflict in the group, the shared display tends to draw people into a spirit of collaborating to create the display, rather than arguing to see who is right or who will win the debate.

A display system for knowledge teams may use flip charts, white boards, overhead projectors, or computer projection systems for the display, and their structure may be as simple as a list of brainstormed items or a calendar, or as complex as a causal loop diagram or process model.

The specific display system presented in this paper has the following components:

- Capture: a facilitator (or “technographer”) typing at a keyboard or writing on a flipchart or whiteboard;
- Structure: the IBIS (Issue Based Information System) conversational model;
- Display: a software system that supports the structure, in conjunction with a computer display panel.

Here are the properties of this kind of display system:

1. Like a flip chart, it serves as a shared work surface for the ideas and issues being discussed by the group, and, like a flip chart, it is fast and convenient (this fulfills Criterion A, page 26, by providing a small, high-speed store);
2. It serves to increase continuity within a meeting and between related meetings, especially if some group members are not present for all sessions (this aspect also fulfills Criterion A, by supporting and enhancing the group’s process over time);
3. It helps focus the group’s attention away from playing out dysfunctional group dynamics and onto sharing and creating informal knowledge together (this aspect fulfills Criterion A, through the power of the structure to focus the group’s attention on particular aspects of the problem space);
4. It provides a constraining structure for the key pieces of informal knowledge in the conversation (this aspect fulfills Criterion B, through the particular linguistic distinctions made by the IBIS structure, e.g. questions and ideas);

5. It captures the context of any specific ideas, decisions, and actions that the group creates (this aspect fulfills Criterion B, through the hypertext webs that record any decisions and their rationale as an intrinsic part of the conversation maps);

These are the properties of any display system, as defined here. The main purpose of a display system for knowledge teams is to facilitate the thinking and learning of the knowledge workers as they meet over time to discuss and solve wicked problems. The secondary purpose of such a display system is to serve as the interface to the project memory, so that their informal (and formal) knowledge is transparently captured, and previously captured knowledge is easily recalled. To fulfill this secondary purpose, the display system must have one more property:

6. Optionally, it connects to a database of previous discussion maps, and supports search and navigation within this informal knowledge base (this aspect fulfills Criterion C, by making storage and retrieval of previous conversation maps easy).

Creating a project memory system requires creating tools and practices for transparently mediating between knowledge work and the project knowledge base. Since this is a large innovation, an evolutionary approach is called for.

First, give knowledge workers a display system (i.e., the Short Term Store in Figure 8) that immediately improves the quality of the knowledge work process and removes some of the recurrent frustrations of meetings, such as lack of clarity and rigor in decisions, rehashing the same ideas, reopening closed decisions, hidden agendas, lack of clear results, politics, and so on. The exact nature of this display system will depend on many local factors, including the group's process maturity, the group and corporate culture, their technological sophistication, the nature of their project, and the time frame of the project.




Figure 8: First step in evolutionary design approach

Second, connect the display system to the organizational computing network (the organization's Long Term Store) so that display maps can be preserved and shared among different people, teams, and sites (see Figure 7).

Third, develop the practices and advanced technologies that (i) transparently index and store informal knowledge captured in the display system, and (ii) retrieve relevant chunks of informal knowledge as they are needed by users of the display system.

This third step is necessary to complete the organizational memory design, and it depends on having effective storage and retrieval mechanisms that connect the short-term store to the organization's long term memory on the corporate computer network. But this is merely a technical problem, and these emerging technologies are not a prerequisite to gaining the many short-term benefits of a display system. Moreover, the biggest challenge in introducing an innovation as sweeping as an organizational memory system is the general slow speed of organizational culture change. By starting with using a display system for the short term store in a project memory system, we begin with a focus on tools that augment whatever culture change efforts are underway in the organization, but also tools which deliver immediate relief from the "organizational pain" of trying to solve wicked problems using only traditional linear methods.

The notion of display system for a team is not new, it just hasn't been used much for the kind of creative knowledge work that I am focusing on in this paper.

Air traffic controllers have their radar screens,
stock brokerages have the "big board," a doctor's
office has a shared appointment book, and a
football team has the downmarkers, the clock,
and the scoreboard.

Each of these is a display of key information that is changing frequently and must be shared by everyone on the team.

A central assumption of this paper is that most knowledge work happens in groups, and that group work is largely conversations. If this is true, then it would seem natural, and indeed compelling, for workgroups to adopt display systems that support and enhance their conversations.

Generally, the closest knowledge teams come to using a display system is to use a flip chart or white board during a meeting, or to write up and distribute the minutes of their meetings. While these are a start, the flip chart and white board have neither the detail nor structure to create much shared understanding, nor do they serve the cause of memory very well. Meeting minutes create a memory trace, but do little to create shared understanding.

The innovation I am proposing here is to treat the *process* of knowledge work—especially the informal knowledge that comes out in conversations and meetings—as a critical and valuable asset, and to capture it in a way that facilitates learning and creativity. For example, it means respecting what each person is saying by capturing the essence of his or her statement in a display system for everyone to understand and appreciate. By doing this, the facilitator brings the group to a practice of listening carefully to each person's point. In a strong sense, the use of a display system institutionalizes *listening* as a vital and honored part of the group work process.

Although such a practice may seem a far stretch from the way people interact with each other in most corporate meetings, the short-term payoffs can make it self-reinforcing. One common objection to this added level of rigor in meetings is that it takes too long; but, once teams have experienced it, they realize how much time they have been spending in repetitive and tangential discussions—and the true economies become clear. Since meetings constitute much of the work process in a knowledge organization, meeting process improvement is a high leverage opportunity.

In addition to better shared understanding, there is a shift in dynamics that happens when a group commits to really listening to each other. The

sense of collegial mutual respect increases, and power plays and petty politics diminish.

The use of a display system creates a shared sense of thinking and learning together, as opposed to determining who has the right answer.

Debate and conflict can still occur, but the chance of getting stuck in that mode is much smaller. The playing field becomes more level, since the strength of a proposal depends more on the strength of the case for it than on who is making it. Creative thinking and learning becomes more important, worth attending to and capturing.

To summarize, the use of a display system for the informal knowledge in meetings and teamwork makes the knowledge explicit, improves shared understanding in the team, and shifts the dynamics from adversarial to inquiry. The display system also functions as a project short term store, and is thus an essential stepping stone to designing an effective organizational memory system.

A Structure for Creative Conversations

Display systems consist of capture, structure, and display. The choice of the structure depends on what aspect of knowledge needs to be made explicit. For example, process modeling (Rummler, 1990) makes the inputs, steps, decision points, and outputs of an organization's processes explicit, while causal loop diagrams (Senge, 1990) make causal relations and feedback loops explicit. Since we are designing organizational memory, a good starting point for what to make explicit is: What do organizations and projects routinely forget that they most need to remember? One answer might be documents and data. There is already a very strong interest in, and some tools for, finding (or not losing) this kind of formal knowledge. But remembering formal knowledge alone is not enough, and won't have much impact on organizational learning.

My answer to the question ‘what do organizations need to remember’ is: *decisions*²³ and the knowledge surrounding them. That is, we must learn to capture the decision, the rationale behind it, the open questions related to it, the assumptions behind it, and any related supporting information. Admittedly, these things are only a part of the informal knowledge being created in a knowledge organization, but they are a critically important part, especially in any organization that is engaged in any sort of planning or design, or where innovation is essential.²⁴

It turns out that these aspects of decisions and decision making are all neatly handled by a conversational model developed in the early 1970's called IBIS (Kuntz & Rittel, 1972; Conklin & Begeman, 1989). IBIS (short for Issue-Based Information System) classifies all of the points in any creative conversation into four simple elements: questions, ideas, pros, and cons²⁵. Although this structure is simple enough to be easy to learn and use, it is surprisingly powerful. All of the major components of informal knowledge can be concisely expressed in IBIS, including ideas, facts, assumptions, definitions, questions, decisions, tradeoffs, guesses, inferences, and points of view. Creative conversations can be captured and rendered in terms of a web of IBIS questions, ideas, and pros and cons, yielding a vivid map that makes the structure of the conversation explicit. There are other models that serve for other aspects of informal knowledge (e.g. action workflows, Toulmin diagrams), but IBIS is the best we have found so far for creative conversations.

Here is an example of the power of IBIS for knowledge work. In a one-year field study, a five-person software team working on a commercial product used IBIS as the structure of their design meeting minutes. They also kept track of the effort that they put into capturing the IBIS minutes, typing them into the computer, and keeping them organized. In the process of reviewing the IBIS record, they found 11 errors in the software and its specification. They were able to calculate that by finding these

²³ A decision is a final conclusion or choice, a point of commitment to a certain idea or path.

²⁴ The entire engineering field of *design rationale* is based on the importance of focusing on and supporting the decisional aspect of informal knowledge. An excellent overview is Carroll & Moran (1995).

²⁵ Rittel termed these “issues,” “positions,” and “arguments.” We have given them slightly more intuitive names.

errors when they did, they saved between *three and six times* the cost of documenting their design thinking in IBIS (Conklin & Burgess Yakemovic, 1991). In other words, the documentation effort more than paid for itself in process improvement and attendant cost savings. It is difficult to get empirical measures like this in a commercial setting, but this result clearly shows that statements like “capture is too expensive” or “it takes too long” are simply wrong, at least in the case of capturing the process of creating large knowledge products.

IBIS is an excellent structure for display systems, because it is easy to learn and use, it can represent nearly everything that gets expressed²⁶, it does not require the meeting participants to change their behavior, it exposes assumptions, rhetoric, and hand waving but allows for a full range of creative expression, including equivocality, ambiguity, incompleteness, and inconsistency. The display system allows the group to focus their attention on a shared space, and the IBIS structure allows the contents of that display to remain well structured and to create shared understanding by letting each participant see how their ideas relate to the group’s thinking. Most important, the structure does not require group consensus for the display to be coherent: the display can hold widely different, even opposing, points of view, contradictory facts, and equivocal statements without losing its power as a shared space for thinking and learning.

A Tool for all Reasons

Of course, even a single meeting can produce a large number of IBIS elements. Knowledge groups need tools to support the use of display systems, and to transition that knowledge into project memory. QuestMapTM²⁷ is a computer tool for capturing and managing any size of IBIS map (see Figure 9), and any number of interlinked maps, among large numbers of users. QuestMap is thus one of the first in a new generation of computer tools that support the *process*—not just the

²⁶ The primary exception is action items (e.g., promises, requests, etc.).

²⁷ While QuestMap is no longer available, a more powerful version with the same capabilities is available for free in Compendium. Download the software at <http://compendiuminstitute.org>.

products—of the knowledge team's daily work. As conversations are captured and stored in QuestMap, the organization begins to accrete a memory of informal knowledge, linked within itself (using hypertext links) as well as with any related documents that are on the corporate network. Moreover, it is not necessary to have a face-to-face meeting to create IBIS maps. Knowledge workers sitting at computers in their offices can conduct “virtual meetings” in which issues are raised, discussed, and resolved. To our knowledge there are no other commercially available hypertext tools which support capturing meeting conversations in IBIS.

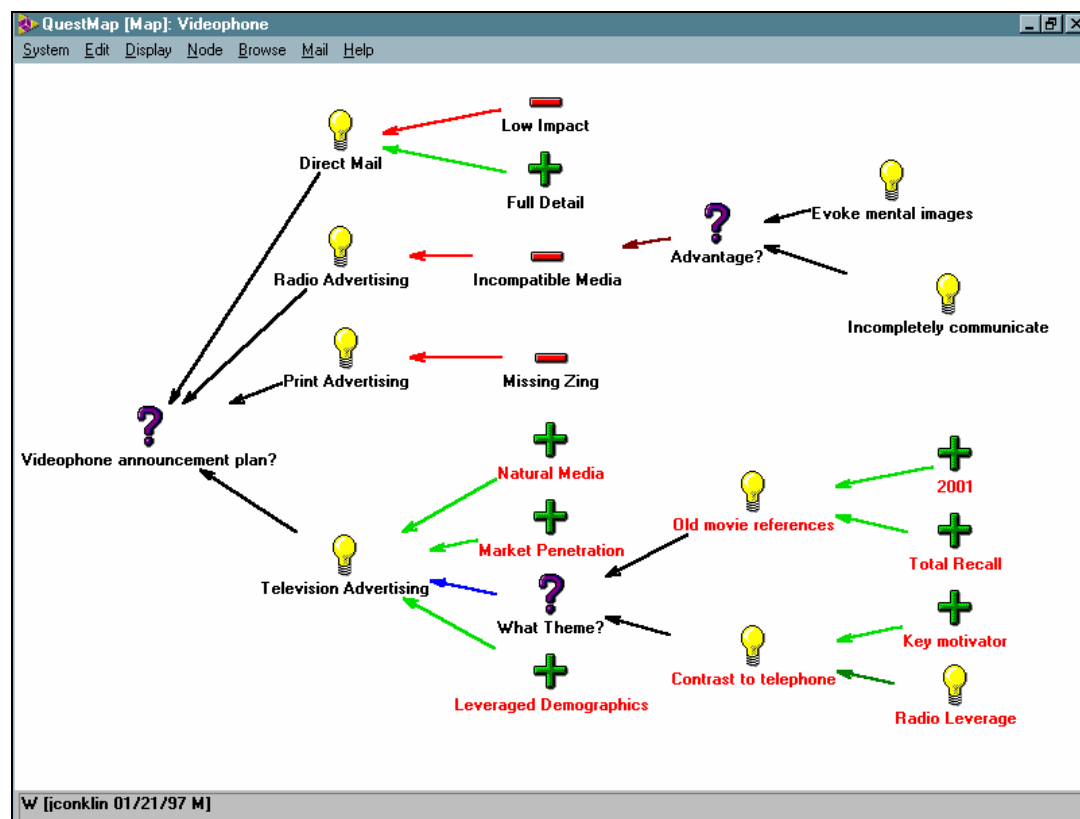


Figure 9: A Sample QuestMap Screen

Examples of Organizational Memory in Action

Continuity from Meeting to Meeting

The author facilitated a series of planning meetings for a utility industry group. We used QuestMap to capture the key points in the meeting conversation and a display projector to display the QuestMap screen on the wall so that everyone could see it. At each meeting the maps we created corresponded to the main agenda topics for that meeting, such as “By-law changes,” “Budget,” “Dues,” and “Who is the customer?” At one meeting one of the specific issues in the dues discussion was “What are the criteria for the dues formula?” The group explored this question and came up with three criteria. At the next meeting several months later, during the discussion of dues, the question about criteria came up again. The group started to rehash the question, but within a few seconds the QuestMap facilitator had navigated to and pulled up the map from the previous meeting onto the shared display.

After reviewing what they had said before, the group added two new criteria to the list. Then they were able to return to their discussion of dues, confident that they had not wasted any time and that they were back up to speed with their earlier thinking. In this case, the retrieval process was transparent to the group—except for their delight at the power of effortlessly reusing informal knowledge. The time between meetings of a workgroup can be as little as a few hours and as long as a year or more; organizational memory must handle the smaller “remembering events” at least as well as the larger and longer-term sort.

Virtual Meetings Add Rigor, Save Time

An environmental planning division at an electric utility company explored and resolved a complex problem with only two formal meetings, one at the beginning and one at the end, during the two month project span. In this group every workstation had QuestMap installed on it, and the staff used it for “virtual meetings” in which the conversations took place electronically, as with an electronic bulletin board. At the beginning of this project to remove a contaminant from a site, the group manager was faced with

choosing between an existing cleanup effort that was projected to go on for many years and had cost \$15 million so far, and a new experimental treatment that would remove the contaminant much faster but would cost \$8 million to implement, and was not guaranteed to succeed. After the initial meeting, the project team used QuestMap to explore the pros and cons between these two options.

Although some members of the staff favored going ahead with the experimental treatment, the existing field data was inconclusive, and the manager proposed (in QuestMap) that they do a pilot (i.e., reduced scale) project to reduce the risk. The staff studied the proposal and reported back that the experimental process would not work if scaled back to a pilot level. The manager probed a little deeper. Over several iterations, all in QuestMap, the staff went off and did more research, but each time came back with reasons that the pilot would not work. And each time the manager would push back on some piece of their reasoning, exposed in the QuestMap map, and ask “Yes, but did you consider this?” or “Is this really true?”

At the end of two months the staff did come up with a very clever solution: it involved building a slightly reduced scale version of the treatment system, but using plumbing that would have to be installed anyway if the experiment failed and the existing cleanup system had to be continued into its next phase. A final face-to-face meeting was held, and the decision was made to proceed with the pilot project. The team had conducted an extended and rigorous analysis of the problem with very few face-to-face meetings. As a bonus, they had a complete record of the research and rationale that led to the decision.

A Creative Solution plus Project Memory

The environmental affairs group at a large electric utility had been working hard on a wicked problem involving a new substation. The company already owned the site for the substation, but a species of bird had been found on the site that was about to be listed as endangered by the Department of the Interior. There were also some very messy politics involving apparent dealings between the county and a local developer. All of the options of what to do so far had some major drawbacks. The team

sat looking at the QuestMap map of the problem for a while, then someone quietly said, “I have a crazy idea: why not give the property to the U.S. government?” After some exploration, the team agreed this was a reasonable solution to pursue—definitely an “outside the box” idea. The group felt that the use of the display system had contributed to the creativity of this solution.

After a flurry of activity the environmental group came up with a recommendation for management. Things died down, but six months later a new development caused the whole issue to come back up to “Condition Red”. The group was able to go back to their QuestMap maps of the project and quickly come back up to speed on all the options and information. By adding some new information, and updating some existing items, they could now clearly see what the best option was for the current circumstances. A new course of action was adopted, and this, too, was captured in the group’s memory.

Summary

Workgroup computing (“groupware”) tools take an important step in the direction of facilitating knowledge work, and their databases inherently create some degree of organizational memory. But such tools also can—and do—create mountains of incoherent rubbish. The problem is that, to avoid the attic-full-of-stuff syndrome, knowledge must be organized and indexed as it is being captured, without creating a burden to the people who create it.

The theory I am advancing in this paper can be summarized as follows:

1. The concept of organizational memory, and the possibility of an effective organizational memory system, has growing importance in the global knowledge economy, but many organizations are letting their most valuable asset—their informal knowledge—go “up the smokestack.”
2. Current implementations of organizational memory fail for a variety of reasons, including: (a) a broad cultural focus on work products over

- process, (b) lack of tools which make capture and reuse of knowledge transparent.
3. The challenge is to design an organizational memory system which offers sufficient short term payoffs to knowledge workers that they will use the system, both to capture knowledge as they are creating it and to look for and reuse existing knowledge.
 4. The next step in the evolution of organizational memory is the use of a display system to (a) focus knowledge workers on improving shared understanding and coherence in their project meetings, and (b) capture the group's informal knowledge—in context—and link it with the project's formal products in an easy and natural way.
 5. The display system mediates access to the organizational memory. It should (a) be readily available for all team meetings, (b) be linked into the organization's computing network so that information can be easily stored and retrieved, and (c) structure informal knowledge in a way that enhances the process of creative teamwork. The IBIS structure is generally a good starting point.
 6. Such display systems are not mere theory—they exist and work (an example is QuestMap). Display system tools do not solve the whole organizational memory problem, but they reflect essential principles that must be embodied in the design of an effective organizational memory system.

Once a team or organization has recognized the value in its informal knowledge, and has begun to capture and manage it appropriately, the group has the key raw ingredients of project memory, and ultimately of organizational memory. GDSS has direct experience, through our clients, of the value of this memory for a team when they come back later and need to take up where they left off. Of course, as the size of the organization—and its memory—increases, new problems of scale emerge that are both technical and cultural in nature. The good news is that the short term payoffs from using display systems generally pay for the cost of implementing them, thus easing the evolution toward a complete organizational memory system.

Acknowledgments

I am very grateful to Michael Ayers, Tobin Quereau, Bill Weil, Al Selvin, Simon Buckingham Shum, Robert Conklin (my dad), Mike McMaster, Kai Brown, Marie Case, Michael Begeman, and Lauren Powers for their comments on earlier drafts of this paper. Conversations with Joe Griffith and Faith Flores were essential in fleshing out the STM/LTM metaphor.

References

- Ackerman, Mark S. (1994) "Augmenting the Organizational Memory: A Field Study of Answer Garden," in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW)*, ACM, N.Y.
- Bannon, Liam J., and Kari Kuutti (1996). "Shifting Perspectives on Organizational Memory: From Storage to Active Remembering," in *Proceedings of the 29th Hawaii International Conference on System Sciences (HICSS-29)*, Vol. 3, I.E.E.E. Press.
- Carroll, J. M., Alpert, S. R., Karat, J., Deussen, M. S. V., & Rosson, M. B. (1994). "Raison d'Etre: Capturing Design History and Rationale in Multimedia Narratives," *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems* (pp. 192-197), ACM: New York.
- Conklin, Jeff, and Michael Begeman (1989). "gIBIS: A Tool for All Reasons," *Journal of the American Society for Information Science*, 40, pp. 200-213.
- Conklin, Jeff and KC Burgess Yakemovic (1991). "A Process-Oriented Approach to Design Rationale," *Human-Computer Interaction*, 6, pp. 357-391. Also in *Design Rationale: Concepts, Techniques, and Use* (1995) Lawrence Erlbaum Associates, NY.
- Conklin, E. Jeffrey (1993). "Capturing Organizational Memory," in *Proceedings of GroupWare '92*, D. Coleman (Ed.), Morgan Kaufmann, San Mateo, CA, 133-137, 1992. Also in *Groupware and Computer-Supported Cooperative Work*, R. M. Baecher (Ed.), Morgan Kaufmann, 561-565.
- Drucker, Peter F. (1994). "The Age of Social Transformation," *The Atlantic Monthly*, November 1994.
- Eppel, Richard, and E. Jeffrey Conklin (1995). "Blending Cultural Transformation and Groupware to Create a Learning Organization," in

GroupWare '95 San Jose Conference Proceedings, D. Coleman (Ed.), The Conference Group, Scottsdale, Arizona, pp. 265-276.

Kuntz, Werner, and Horst Rittel (1972). "Issues as elements of information systems," Working Paper No. 131, Institute of Urban and Regional Development, University of California at Berkeley.

McMaster, Michael D. (1995). *The Intelligence Advantage: Organising for Complexity*. Knowledge Based Development Co. Ltd., Douglas, Isle of Man, UK.

Minneman, S., Harrison, S., Janssen, B., Kurtenbach, G., Moran, T. P., Smith, L., & van Melle, W. (1995). "A Confederation of Tools for Capturing and Accessing Collaborative Activity," *Proceedings of the ACM Multimedia-95 Conference*, ACM: New York.

Moran, Thomas P. & John M. Carroll, Eds. (1996). *Design Rationale: Concepts, Techniques, and Use*, Lawrence Erlbaum Associates, Mahwah, New Jersey.

Oinas-Kukkonen, Harri (1996). "Debate Browser - an Argumentation Tool for MetaEdit+ Environment," *Proceedings of The Seventh European Workshop on Next Generation of CASE Tools (NGCT 96)*, Crete, Greece, May 1996.

Rummler, Geary (1990). *Improving Performance: How to Manage the White Space on the Organization Chart*. Jossey-Bass Publishers, San Francisco.

Senge, Peter (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*, Doubleday Currency, N.Y.

Shum, S., MacLean, A., Forder, J., & Hammond, N. V. (1993). "Summarising the Evolution of Design Concepts Within a Design Rationale Framework," *Adjunct Proceedings of InterCHI'93: IFIP/ACM Conference on Human Factors in Computing Systems*, Addison Wesley, Amsterdam, pp. 43-44. Also available at <ftp://kmi.open.ac.uk/pub/simonb/interchi93.ps>.

Wheatley, Margaret J. (1992) *Leadership and the New Science*, Berrett-Koehler, San Francisco.