8

TOWARD PORTABLE IDEAS

Mark Stefik
John Seely Brown
Xerox Palo Alto Research Center

INTRODUCTION

We honor creativity in our culture, especially that of the individual genius, but creativity is as much a social as an individual affair. When people of different backgrounds come together, new ideas can arise from their conversations. Sometimes new ideas are built up incrementally from the fragments of different viewpoints. Ideas can be made more robust when they have been bounced around, critiqued, polished, and repackaged by a group.

We have all been in situations where our part in developing ideas has been but one of many contributions. We have all benefited from the wisdom of a second opinion, or been surprised on occasions when good ideas came from unexpected sources. Nonetheless, idea creation, like other aspects of intellectual work and even routine office work, is usually conceived in terms of the contributions of isolated individuals. Conventional wisdom has been slow to recognize the importance of collaboration and teamwork. The wisdom about teams that "many hands make light work" refers to hands not minds, and certainly not to committees.

Most approaches to office automation and computer-mediated work have focused on individuals rather than groups. The landscape is littered with failed computer systems that were supposed to make light work of various office tasks (Bikson, 1987). Studies of the acceptance and use of computer systems in offices have shown consistently that a major factor determining the success or failure of such systems is whether the designers took into account the habits, needs and activities of work groups.

During the past two or three years, however, there has been a burst of new thinking about computer-mediated work. In contrast with the terminology of personal computers and computers to empower *individuals* to do their best, we are starting to hear much more about *interpersonal computing*. This means different

things to different people, and new jargon has started to appear including "cooperative computing," "collaboration technology," and "work group computing." The somewhat awkward term "groupware" has been proposed for computer software that is specifically intended to aid in the work and coordination of activity of a work group.

This chapter presents a particular vision of possibilities that we have found intriguing and that we believe could have profound effects on the functioning of organizations. Crafting tools that actually help collaboration is a very subtle enterprise. There are two parts to our thesis. The first is that creative genius lies in the social substrate itself. Secondly, the interaction of ideas properly externalized and appreciated leads to wonderful combinations and results. Some of this synergy and exposure can be enhanced by tools in the social infrastructure. There is a streak of genius and creativity in each of us. That streak can be tapped by creating a medium in which ideas can rub productively against each other. We propose a medium of active and sharable workspaces for developing and explaining information.

In the next section we focus on the Colab project, one of several projects studying collaboration and supporting technology at Xerox PARC. We sketch out the basic premises of the project, describe one of the experimental systems that has been developed, and present some of the questions and issues that we have encountered in the work so far. In the following sections we reexamine and critique some of the basic assumptions of this work. This prepares us to ask more basic questions, and also to propose ways that the ideas and technology for group work could prove more significant and valuable than in our original vision, especially for research and engineering organizations.

The study of collaborative tools requires transcending what we call the "techno macho" syndrome, the fascination with technology or methodology for its own sake. We should not get carried away with the belief that technological artifacts or decision methods in themselves will help that much. Many obvious attempts to apply technology to the work setting have hindered more than they have helped. However, for the authors, there is no denying or escaping our role as technologists. We are not disinterested observers of technology and the social scene. We are, in fact, concerned with the limitations of the status quo and are actively trying to invent new and more productive ways of working. This chapter is intended to stimulate those who want to think beyond current technologies and work practices.

COLAB MEETINGS AND CONVERSATIONS

The Colab project (Stefik et al., 1987) was conceived as an experiment in computer support for meetings. We imagined that professional people in meetings should have the same kind of access to computers that they have in their offices

for private or isolated work. In support of this we created the Colab meeting room as shown in Figure 8.1. The meeting room provides a computer workstation for each participant in a face-to-face meeting. At the front of the room is a blackboard-sized touch sensitive screen (which we call a "liveboard") capable of displaying an image of approximately a million pixels.

Most of the meetings that we had in mind when we started the Colab project take place by a small group in front of a whiteboard or some other vertical and erasable writing surface. In these meetings, a creative group is engaged in discussion and work activities using notations on the whiteboard to formulate and explain their thoughts and to keep notes during the meetings. For us, the whiteboard is the dominant medium used in meetings; it is a medium that we all use constantly in our daily work. We could see many shortcomings of the whiteboard, several of which are discussed in the following. The whiteboard became the technology to beat in inventing a more powerful medium for meetings, and we decided to beat it by creating a computational medium that kept the best properties and brought in new capabilities as well.

Our new medium distributed a computational whiteboard to every participant in a meeting. To promote shared viewing and shared access to what is written during the meeting, the Colab software is oriented around a concept for multi-user interfaces that we call WYSIWIS (What You See Is What I See—pronounced "whizzy whiz"). In a WYSIWIS interface, all the meeting participants can see exactly the same information on their displays. Colab meeting tools support this illusion by maintaining synchronized views across workstations. In addition, each person can point to things on the display with a personalized "telepointer" that is made visible in real time to the other participants. Colab software also supports private windows. Private windows correspond to notepads; public WYSIWIS windows correspond to whiteboards.

But how could computers possibly help in meetings? One could approach this question by enumerating the aspects of meetings that are annoying and then investigating which of them might be ameliorated by computer technology. This would be an awesome task, requiring at the onset some substantial focus to limit and to identify the kinds of work to be investigated. However, we approached the issue from another direction, trying to understand the properties of a computer medium and then imagining the kinds of meeting situations in which computers could make a positive difference. Most of our intuitions are based on information processing concepts both for computers and the nature of work in meetings. These concepts also provide some insights about the kinds of work activities for which computers might make a difference.

Computers provide more space for writing than whiteboards. The storage capacity of a whiteboard is quite limited and after a period of time, a group writing on a whiteboard must erase things in order to keep going. In a computer medium, the display space can be reused without discarding information because

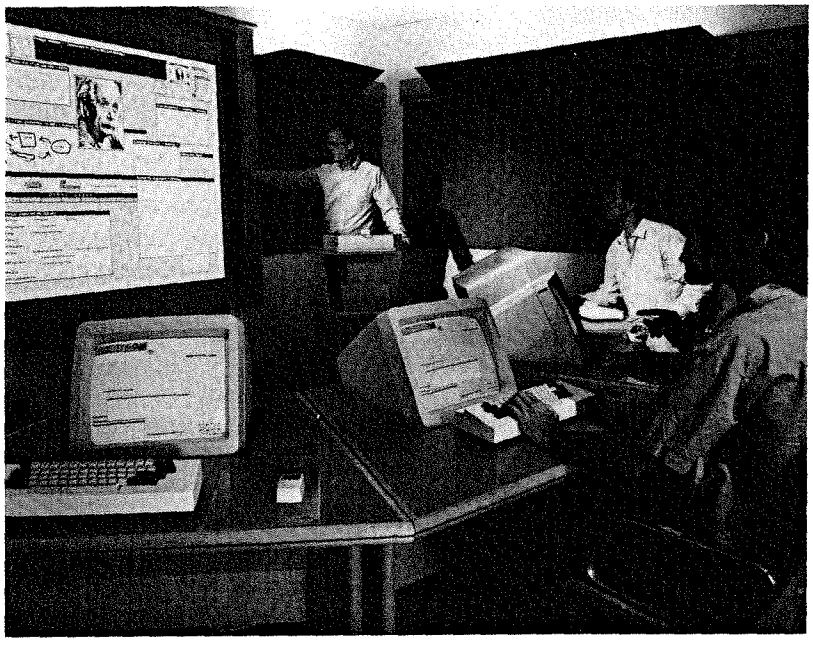

Figure 8.1. View of the Colab. The Colab is an experimental meeting room that provides computational support for collaboration in face-to-face meetings. It is designed for typical use by two to six persons. Each person has a workstation connected to a personal computer. The computers are linked together over a local area network (Ethernet) that supports a distributed database. In addition to the workstations, the room is equipped with a large touch-sensitive screen and a stand-up keyboard. (Photograph by Brian Tramontana)

symbols can be moved to and from file space. Furthermore, the file storage capacity of a computer is quite large and there are many techniques for organizing the display of information on a computer screen using windows, icons, and scrolling techniques.

Even with whiteboards, participants tend to build up large collections of written symbol structures that provide the common ground for reference. With more space for writing, participants can build up potentially larger sets of shared writings. This can be important for meetings that last for several sessions. Backup on a file system makes it possible to recall and display things even if they were developed in previous meetings.

The abundance of space in a computer is no excuse for neglecting to manage space as a resource, but here again the computer medium offers some advantages.

Using techniques from bitmapped user interfaces, items can be quickly and easily rearranged on a computer screen. In contrast, on a whiteboard one must manually copy and erase symbols in order to rearrange them. This flexibility makes it possible to organize a screen, reducing clutter. This enables a group to organize space more easily for the purposes of visualization, and accommodation of shifts in focus.

A computer medium provides computational leverage that can be used in many ways. For example, in a resource allocation meeting, the computer could provide visible spreadsheet capabilities. It could display information in alternative formats for easier manipulation or better understanding. It can also provide search services for finding information in large sets.

WYSIWIS interfaces can relax some constraints on communication and cognitive processing. By enabling participants to use a shared written medium, the bandwidth of communication is potentially increased since more than one person can add information at the same time. However, even if the apparent increase in bandwidth of communication in a Colab setting is not significant, freeing the constraints on parallelism and serial communication may improve the quality of deliberations by enabling meeting participants more freedom in scheduling their attention and cognitive activities.

These general capabilities of a computer medium suggest that the Colab would have the most advantages in meetings that include manipulation of substantial amounts of information, such as meetings by engineers in which complex designs are discussed and compared.

Meetings Tools: An Example

We use the term meeting tool to refer to computer software in support of groups in meetings. Just as users of personal computer software need different tools for different purposes (e.g., text editors, spreadsheets, mail systems), so too do meeting participants need different tools for different purposes (e.g., tools for agenda control, brainstorming, negotiation, and argumentation).

Cognoter is a Colab meeting tool used in our lab about once a week. It is used to organize ideas for presentations, reports, talks, and papers. Cognoter supports a meeting process in which participants come together, usually without having prepared any materials ahead of time. Meeting participants determine the audience and goals for their presentation, the topics to be included, and the overall organization. The output of Cognoter is an annotated outline. Figure 8.2 shows an example of the visual display created by Cognoter when it is being used.

The organization of Cognoter is described in the following. To convey a clearer sense not only of what it is but also how it is used, we have included some informal observations of use and meeting phenomena. However, at the time of this writing, adequate recording capabilities for observing meetings were just becom-

ing operational. Consequently, formal and quantitative studies of group behavior in the Colab are still ahead of us.

Several things can reduce the quality of a presentation. It could fail to include some important topics; it could dwell on irrelevant or unimportant topics; or it could address the topics in an incoherent order. To avoid these pitfalls Cognoter organizes the process into specific stages. Each stage incrementally increases the set of actions available to the user.

The stages in Cognoter are brainstorming, ordering and grouping, evaluating, and generating an outline. We originally adopted this structure from a similar one that proved useful for us in non-computational settings; however, there are significant differences in the uses and effects of the stages when computers are introduced.

Brainstorming Ideas. The brainstorming stage is intended to foster the free-flowing contribution of ideas. There is one basic operation: A participant selects an empty space in a public window and types in a word or phrase characterizing an idea.

Unlike usual brainstorming meetings, there is no waiting for turns in Cognoter; any participant can enter an item at any time. Often the inspiration for an item is triggered by another participant saying something or entering an item in a public window. Thus, communication (or loosely "the conversation") in Cognoter takes place both by voice and over the computer medium. All items appear on everyone's displays. Participants can annotate items with longer descriptions to clarify their meanings.

Organizing Ideas. The order of items for the presentation is established in Cognoter by incremental and local steps. There are two operations: linking items into presentation order, and putting items into delineated subgroups. If item A is linked to B (meaning A comes before B), and B is linked to C, then A comes before C. If item A is linked to a subgroup, then it comes before every item in the group. By these transitive and distributive operations, a small number of explicit links can tightly constrain the order of items in the outline.

The linking operation often takes place in conjunction with an oral justification. For example, if "expenses" and "bottom line results" were items, a participant might argue out loud "We have to talk about expenses before bottom line results because otherwise management won't understand the results." This relation is represented visually in Cognoter as an arrow linking the item labeled "expenses" to the one labeled "bottom line results." It is also possible to move related ideas to an idea subgroup in a separate window. Before moving items, it is common practice to put them in a spatially compact cluster. This allows comment on the coherence of the proposed grouping.

In Cognoter, the overall task is richer than in traditional brainstorming sessions. For one thing, the task is not finished when some ideas have been generated. Preparing a presentation requires organizing and evaluating ideas as well. Further-

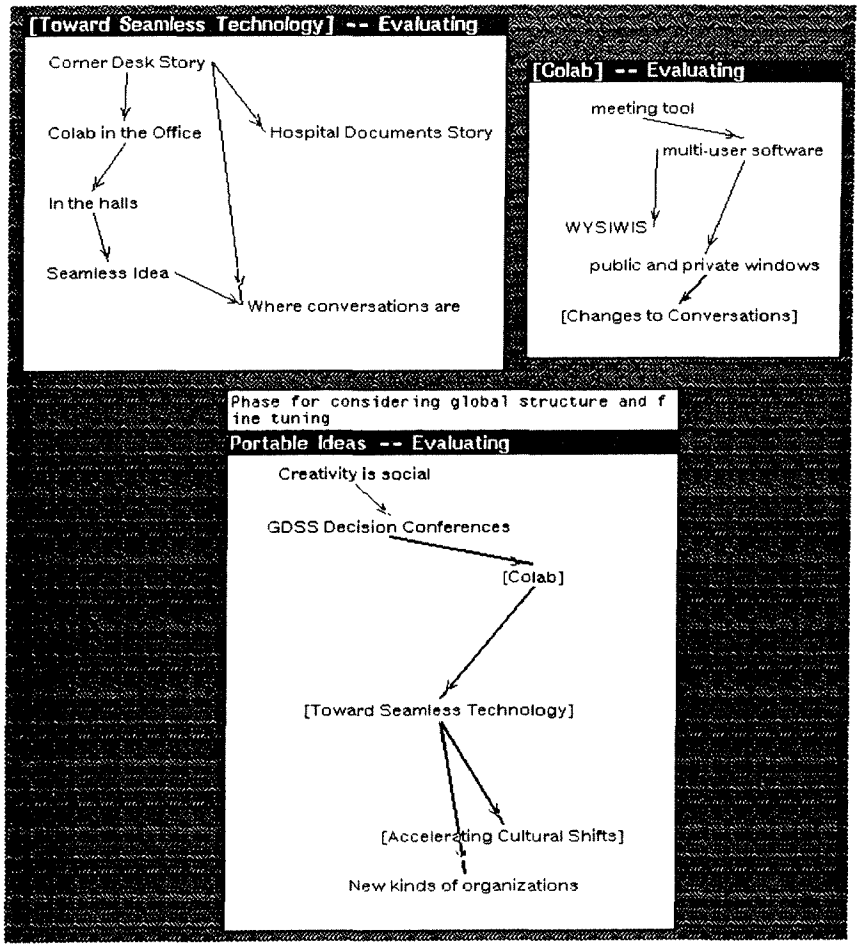

Figure 8.2. Interacting with Cognoter. Cognoter is a meeting tool for organizing a presentation. Cognoter guides this process through several stages: brainstorming, organization, evaluation, and outline generation. The items on the display are short expressions that refer to the ideas for the presentation. In this figure the items have been organized into three major groups. The arrows between the items indicate an ordering relationship; they imply constraints about items that must be presented before other items. Cognoter automatically generates presentation outlines that take into account these constraints.

more, our informal observations of the meetings indicate that people form subgroups that focus on the development of particular aspects of the subject matter. Since subgroups of ideas are usually put into separate windows, each subgroup of people can focus around one window or another. The frequency and significance of this behavior and the importance of supporting it with meeting tools will be a subject for systematic study when our observational facilities become operational.

Although subgroups of collaborators can work mostly independently, they can also communicate. For example, one group may decide that some of the items in its windows don't fit with the others, and may put them back into the general pool or offer them to another subgroup. Further communication is then required when the items are reconsidered, perhaps by the whole group. When subgroups rejoin, participants can recap the changes made in the subgroups.

Evaluating Ideas. During this stage the subgroup boundaries tend to dissolve and the meeting participants function again as a single group. Participants try to understand the organization of the presentation as a whole. Items that seem irrelevant or less important than others can be deleted. Outlines are generated by Cognoter upon request, and ambiguities in the ordering can be highlighted. Participants can argue whether particular items are irrelevant or unimportant when compared with others.

Expanded Dimensions for Conversations

Tools like Cognoter embody more than a shift from a whiteboard to a computer. Effective collaboration has some tacitly held rules, such as taking turns in conversation. There are multiple roles such as inventing, critiquing, reformulating, scribing, and summarizing. People switch roles during a conversation, and the switching itself follows certain rules bearing on rhythm, momentum, and topical focus. For example, in an effective collaboration one party will hold back and not interrupt a second party who is obviously "on a roll" generating a stream of related ideas.

Computer media change some of the basic parameters of conversation and enable profound changes in the shape of the conversations. When we move from personal to interpersonal, the requirement for personal intelligibility of the subject matter shifts to a requirement for *mutual intelligibility;* the meaning of conversational terms shifts from being internalized and fixed to being externalized and negotiated. Expressed in a computer medium, communications persist in a form that is tangible, external and manipulable. We conjecture that this substantially increases the amount of information for which there is a lasting awareness and shared understanding by participants in a meeting.

The coordination of intellectual work around manipulable icons draws on familiar skills for the coordination of physical work. Sorting can be done with icon manipulation. One moves items between buckets until they are in the right places.

In both cases the multi-user interfaces indicate when objects are being manipulated (e.g., edited) by someone, providing visual clues for team coordination.

The same kind of manipulative action can be used for indicating when one participant wants another to work on an item. Thus, it is possible to pick up an item and drop it into the workspace of a second participant. This is very much like the physical act of picking up a physical object (e.g., a piece of paper) and handing it off to someone else for attention. A single action removes the object from one's own inventory and adds it to someone else's inventory for attention.

The possibility of simultaneous communications relaxes many constraints. It increases the possible bandwidth of communication in meetings. Changes like this raise many questions for which we do not yet have answers. In ordinary conversation, the meaning of what is said often crucially depends on the context of what was said just before. In Cognoter, multiple things can be "said" at once in the computational workspace. It can be argued that such capabilities introduce confusion into the meeting process. On the other hand, this capability can be enormously freeing in the context of a fast-moving brainstorming session. Reading is faster than listening so it is possible to scan the items being created by several others and occasionally to respond to them. When something puzzling comes along on the screen, however, it is not necessary to tend to it at once. Unlike oral communication, there is no need to remember a confusing item because it remains in the workspace inventory for later processing. Any systematic process for going through the items will encounter the item again for later consideration. A written workspace is amenable to scheduled and systematic processing of comunications. Thus, even if the potential increase in bandwidth does not result in an overall increase in communication, it may be important for other reasons.

During the brainstorming phase of Cognoter, the parallel action in proposing ideas reduces the usual verbal communication to coordinate turn-taking and synchronization. A participant can enter an idea whenever it comes to mind. Oral conversation tends to drop off radically, since so much of the communication load shifts from ears to eyes. The speech resource becomes more available for questions that clarify points.

It is important to note that parallel action is not altogether absent in non-computational meetings. Videotapes of design meetings show that groups of designers working on large sheets of paper engage in much parallel sketching activity. Furthermore, shared workspaces can be created in other media such as video. Architects working together through such a medium have reported an intensity of engagement and productivity (System Concepts Laboratory, 1987) similar to the informal reports by users of shared workspaces in the Colab.

Another profound change to the dimensions of conversation is the possibility of equal access to public data. In Colab, the conversational acts that enable a participant to modify the public display or to assume the role of chairman can take place in a fraction of a second. In contrast, with an ordinary meeting room with

table and blackboard one must negotiate the transition, rise from a chair, walk to the board and so on. By lowering the hurdles of transition the technology creates a potential for broader participation and for more flexibility in roles.

Whether and when these changes in conversations are beneficial is still to be determined. For example, accelerating the pace of a meeting as in the brainstorming phase may give participants less time to think. On the other hand, freedom from serial turn-taking may alleviate some of the problems of production blocking (due to the limitation that only one group member can talk at a time) reported by other studies of brainstorming in more conventional media (Diehl & Stroebe, 1987). In some cases, the context surrounding the generation of an item may be lost. Important effects also happen at the transitions between meeting processes. Among the transitions are the formation and dissolution of subgroups, the shifting of participants from one topical focus to another, and the transitions of the conversational patterns of the whole group as it shifts from having a single focus of activity to having multiple conversations and subgroups and then back again. However, our purpose now is to begin to understand some of the differences, not to evaluate them.

Seamlessness at Work

As we move on to the next phases of our research, we are also ready to challenge the assumptions of the Colab design with an issue that seems to dwarf all the others:

Meetings do not take place (exclusively) in conference rooms.

Meetings take place wherever people get together and have conversations. Recently, one of us spent time observing the use of documents by nurses at the Pacific Medical Center in San Francisco. What was impressive and interesting was how much clarification, co-ordination, and negotiation among nurses took place over their clipboards. If a narrow-minded computersmith wanted to bring information processing to the hospital situation, the first bad idea might be to make the documents available on a workstation located somewhere down the hospital corridor. That would completely ignore the conversations and interactions of the nurses where they meet. Like other intelligent human beings, nurses could probably cope with a poorly designed computer system. However, to be most helpful, we suspect that the technology of record keeping and conversation should be as familiar and easy to use as a blackboard and as readily available and portable as the clipboard, paper, and pen.

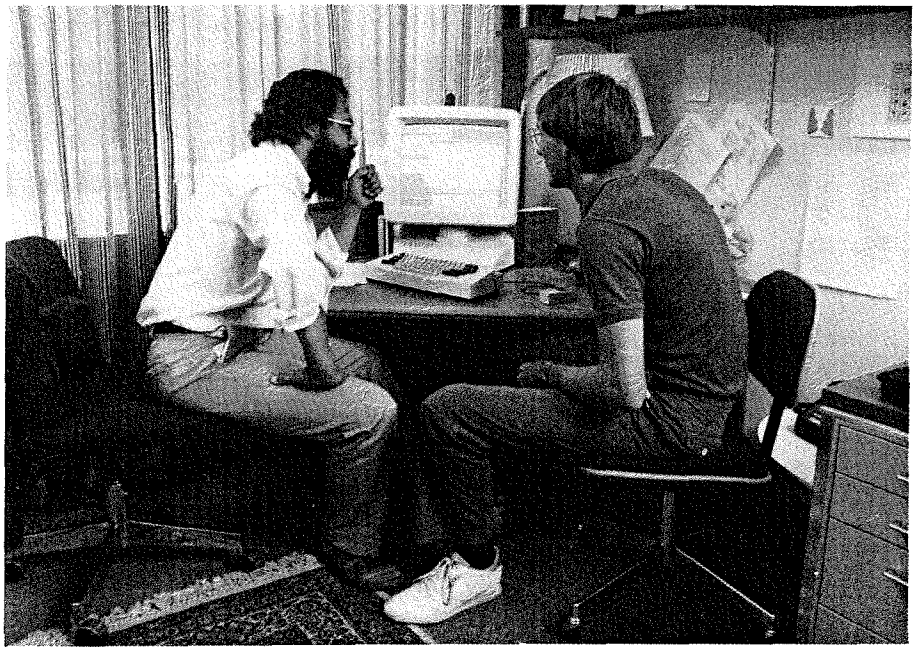

Figure 8.3. Team programming at a corner desk. Wedge-shaped desks were developed for use by our group in team programming. A critical factor bearing on the successful use of these desks was locating them in offices. Programming activities start in offices, and teamwork begins there. When we first placed the desks in conference rooms, we observed that programming that started in an office never moved out to the desks. Instead, collaborating programmers would just squeeze into the office situation. The problem was presumably that the overhead of moving the computing environment and debugging context to a second location was greater than the benefit of increased elbow room. (Photograph by Brian Tramontana.)

To illustrate this issue of the use of technology and the location of meetings, we present a story about some very low technology that we misapplied at PARC. The technology was a corner desk, a wedge-shaped desk for holding a computer workstation that could be located in the corner of a room as shown in Figure 8.3. We wanted to promote team programming on our research projects and these new desks offered ample elbow room for two people sitting together. We placed the

desks in our regular conference rooms, equipped them with computer workstations, and anticipated that they would also be useful for demonstrations and for visitors when office space was tight.

After the corner desks were installed, we noticed that they were never used. Team programming was happening occasionally, but it always started in somebody's office. Someone might start with a system debgging question, or have a programming puzzle or idea. In every case people jammed their chairs together in the office and squeezed around the workstation. At no point did they move to the corner desks in the conference room.

Concluding that the desks were a failure, we decided to have them stored in a warehouse. To salvage something from them, one of the authors who liked the aesthetics of the corner desks and had an office shape that could use one easily, decided to discard a table and regular desk from his office and to use a corner desk instead. Shortly thereafter, team programming was observed to occur in his office, and furthermore, the idea of putting the desks in offices was legitimized. The corner desks are now mostly located in our offices, where programming occurs. Now they are serving their intended purpose, team programming occurs on a regular basis, and several more corner desks have been built.

Returning to the assumption behind the Colab meeting room, we note again that meetings take place regularly in offices, not just meeting rooms. This raises the question of how offices should be equipped. Here we believe that the Colab experience is relevant. In taking a prescriptive stance, we can predict what could be possible.

We believe that one of the most useful additions to the infrastructure of an office would be a large touch-sensitive display: an office liveboard as sketched in Figure 8.4. A large display creates a focus of attention for a team working together. Furthermore, with other workstations based on CRT displays or flat panels around the room, we could presumably get some of the WYSIWIS meeting phenomena that we have observed in the Colab.

However, there are two main features that intrigue us. One is the power of the computational medium for flexibly organizing space. Whether whiteboards are in conference rooms or offices, they never have enough space. All of the arguments that we advanced for Colab workstations for managing display and file space apply equally well to meetings in an office. The second feature that intrigues us is the possibility of an accurate, large-scale pointing device adequate for quickly sketching diagrams on a high resolution liveboard. The whiteboards of PARC are always filled with informal diagrams and symbols of great variety. The liveboard in the Colab, however, has a resolution of somewhat greater than one pixel per tenth of an inch. Although this is good enough to present a large image of one of our computer screens, it is much too coarse for smooth sketching at the liveboard.

This brings us back to our suggestion that the technology of conversation should be as familiar and easy to use as a whiteboard and as readily available as

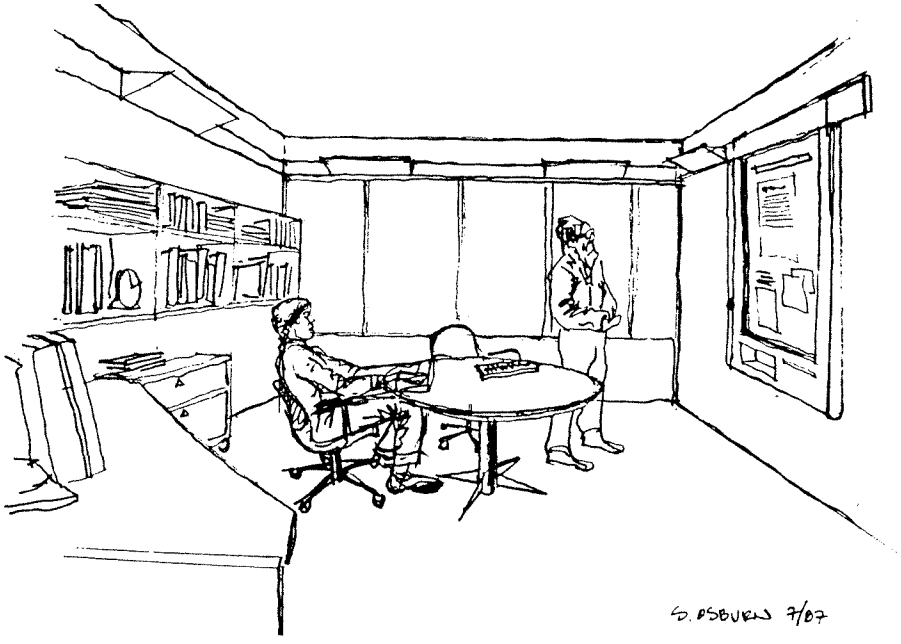

Figure 8.4. Redesigning the office. Colab-like devices could lead to a radical redesign of the office environment, in recognition of its use to support small meetings. For example, the environment could integrate personal electronic notepads with public liveboards on the walls. (Line drawing by Steve Osburn.)

paper and pen. Paper, pen, and whiteboards seem natural and easy to use, in part, because we are exposed to them and trained in their use at a very young age. They are also supremely flexible. You can write on them in almost any way that you desire, making text and figures slant up, slant down, surrounded by wiggly globs, and so on.

Many other meeting devices would be possible in an office and would work well with a liveboard, such as multiple keyboards, remote pointing devices, and small flat panel devices used as networked sketchpads. As in the Colab, pointing devices (e.g., mice or styluses) could enable participants to point to something on the public display without leaving their seats. All these devices should be small so that they could be stored out of sight, and they should be cordless. Another useful device would be a digital tape loop and audio gear so that one could recover a particularly apt turn of expression by playing back the last part of a conversation.

The well-equipped office should provide an information environment that connects seamlessly larger and more formal meeting rooms like the Colab. The same software and hardware should support conversations in both settings. In an office one should be able to prepare materials for larger meetings, and to continue small-scale follow-up meetings afterward.

Explaining an idea to a colleague is simplified when the context-setting sketches on a whiteboard are available. With active liveboards in offices and user interfaces based on new remote window systems, moving the contents of a liveboard from one office to another will be possible. To this end, means for moving the contents of a liveboard should be direct and simple. One could forward the contents to an office liveboard ("Send this to my office") or file it in a database. Similarly, one could retrieve something to show to a colleague in the coffee room ("Get the big-idea window from yesterday's conversation.")

This brings us to a most important concept about new capabilities: the portable meeting. A meeting that starts at one office on one day could be resumed by any of the participants in their own office or introduced to another colleague at yet another time or location. In conventional meeting situations, different kinds of records are kept for different purposes. Some things are written during the meeting for the purposes of explaining or developing a point; other notes are written by the participants for later use by themselves or others, and yet other things may be written such as minutes for explaining to parties that were not present at the meetings. With portable meetings the explanatory scribbles created during the course of a meeting become available for reuse at a second location without the need for manual copying. Furthermore, if one person explains a set of ideas using figures and symbols from a liveboard, the second person could gain access to this information and extend the script for explaining it to a third person.

Many visitors to Xerox's System Sciences Laboratory are surprised by the number of floor-to-ceiling whiteboards, each in its own corner with comfortable seats or couches around them. This is not an accident or just a sign of opulence. These areas were explicitly designed to foster small collaborative teams working in semi-private areas. Whiteboards enable people to create a large sharable context. By having so much of the discussion visibly displayed, it is very easy for someone walking by to gain a sense of what is going on and to decide whether to contribute. Further, one can come up to speed more quickly.

The phenomenon of portable meetings may enable additional possibilities if liveboards were introduced into public areas. One of the most heavily and productively used whiteboards in our laboratory is the one near the coffee service in the lounge area. Perhaps coffee centers foster creativity both because people encounter each other here, and because of the informality and relaxation that the centers suggest. Could the creative power of a center be tapped better if the whiteboard were replaced by a liveboard? After productive conversation one would not need to remember the ideas or to copy the contents to paper (see Figure 8.5).

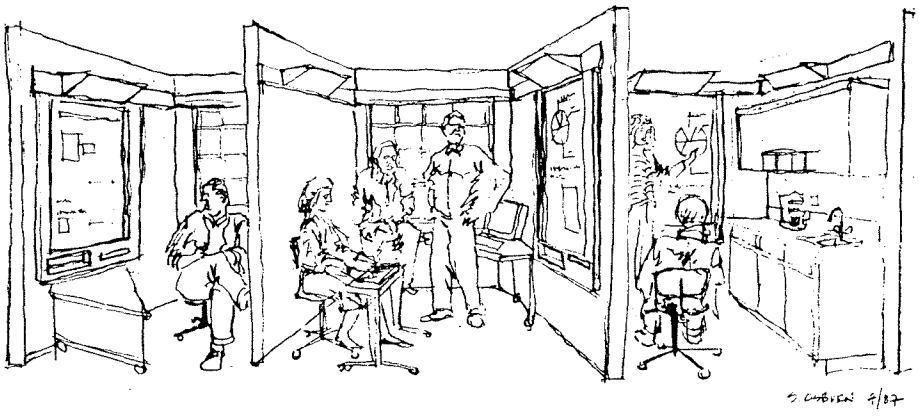

Figure 8.5. A well-equipped work group. Colab-like devices could lead to a radical redesign of all the areas where informal meetings take place. In a setting around the coffee pot, one could go up to the liveboard and retrieve a board from a previous conversation, or forward the board contents so that a conversation could be continued in an office. (Line drawing by Steve Osburn.)

This, finally, is what we have in mind for infrastructure for group work: a seamless environment of tools for conversation that extends from offices to the coffee room to the formal meeting room. One might even consider surfacing the tables in the cafeteria with interactive flat displays, providing an electronic version of the proverbial coffeehouse napkins on which so many important inventions reportedly have been born. In a seamless environment, the ideas of conversation can become not only external and directly manipulable as in the Colab, but also portable.

SEAMLESS TECHNOLOGY AND THE RESPONSIVENESS OF ORGANIZATIONS

In this section we want to step back from our general development to compare our thesis with a contrasting vision of computer support for meetings, and also to offer some thoughts on the effects that this kind of technology might have on organizations.

We began our discussion of computer support for meetings with a discussion of the Colab meeting room and the various assumptions that led to its design. The Colab is not the first example of computer support for meetings. Doug Engelbart's

early demonstrations (Engelbart, 1984b) during the 1960s and 1970s have become almost legendary. One area of earlier and continuing activity reviewed by Kraemer and King (Kraemer & King, 1986) is the use of computer support for a group decision conference. A decision conference is a conference organized by corporate executives facing a major and strategic decision. Often a specialized consulting team is brought in to conduct the meeting. The facilitation team introduces a formal decision method such as multivariate analysis to the executives and runs a conference around the use of the method. Sometimes, computer tools that support the decision method are provided, coupled with large screen televisions or video projectors to make information publicly visible. In the ideal case, the executive team generates its alternatives, identifies its assumptions, and achieves a consensus for a decision after about two days of intense work. At this point, the conference is over, everyone "goes back to work" to implement the decisions, and the model is discarded.

Similar as this concept of "computer support for meetings" may seem to the concepts that we are pursuing, it is actually a study in contrasts. One immediate difference is that decision conferences are aimed at a specific kind of executive decision meeting, usually involving resource allocation. We have focused on small work groups, especially design teams and research groups, rather than executive teams. A decision conference presupposes that there is a "decision" to be made and that a particular formal approach will provide a rational basis for making the decision. During an executive decision conference, much of the work is in determining and supplying parameters for the formal decision model. Our focus has ranged from formal to informal work processes for which an analytical and mathematical model would be inappropriate. Although tools like Cognoter have semiformal methods and other Colab tools have even more formal methods, there is an emphasis on flexible and informal notations especially in the portable meeting scenarios. Finally, the decision conference presupposes that a single meeting occurs, taking somewhere from a few hours to a couple of days. In contrast, we are concerned with a wider range of time intervals: from a few seconds corresponding to conversational acts as meeting events, to an hour or so corresponding to a single session of a meeting, to a few months corresponding to the continuing and portable meetings of a group project.

Seamlessness has several dimensions relevant to computer support for work. Seamlessness for groups means that tools scale gracefully between individual and group work. Seamlessness across locations means that it should be easy to move information workspaces from one setting to another. Seamlessness across tools means that information used in one tool can be moved easily to another. Seamlessness across time means that there are tools that enable one to easily browse information that was developed over time. Seamlessness across media would enable information from one medium (such as speech) to be referenced or copied to

another medium (such as text). Seamlessness in general refers to the ability to manage and move information fluidly.

Software and computer systems for groups will not develop if isolated from other computer systems. Indeed, there is a great incentive to be able to exchange data and to be able to use the same software for meetings and for individual work. We believe that software that extends gracefully from individual work to group work and back again will have advantages. We also believe that appropriate architectural approaches will make this a natural direction for software evolution (but that is a subject for another paper).

When the technological advances necessary for making inexpensive, interactive liveboards is ready, we believe that a revolution in computer use by organizations will follow. Liveboards open up the possibility of computer use for purposes much less formal than is now typical. We believe that people should be able to use a computer as casually and simply as they use a whiteboard now, picking up the "chalk" and doodling.

Some people have referred to computers as a kind of information appliance. An appliance like a toaster is readily available and recognizable. It has a single main function which is recognizable from its presentation to a user. For a toaster, the recognizable display consists of the two slots on the top in which one places sliced muffins or toast. Different appliances present different interfaces according to their function and complexity. What is unusual about computers is that they perform multiple functions. It is as if we expected the same appliance to act either as a toaster, a refrigerator, or a microwave.

This brings us to the notion of the easy path from low functionality to high functionality. Starting from the basic whiteboard, one can imagine a gentle gradient toward increasing capability. The learning gradient should be such that a small investment has a large payoff in increased capabilities, quickly drawing one further into learning and using more. At the basic level, one uses recognizable chalk and an eraser in the usual ways. Simple ends need simple means. After learning about writing and erasing one learns to save and recall the contents of the liveboard. After storage and recall comes moving and copying symbols. Incrementally there is a range of functions that can be learned or not, but the entry point should be simple and familiar and the slope of the learning curve should be gentle. Through simple explorations a user should encounter new possibilities.

Putting liveboards in public places and in casual office use may lead to much greater use of computers in organizations. Much of human learning is by apprenticeship and imitation. A typical user of a computer workstation in a corporate organization is the accountant running a spreadsheet program off by himself in his office. A co-worker is unlikely to see him at work and may have little reason to encounter the arcane computer magic that the accountant uses. Learning by casual watching and imitation is both difficult and unlikely because one must go to the accountant's office, and, furthermore, one cannot easily watch both the computer

display (where the action is) and the keyboard (where the control is) at the same time. With liveboard software, it is possible to arrange it so the actions at the locus of activity (the chalk) will be much more visible than with the usual computer workstation. If they are more visible, they will probably be much more imitable so that computer skills could spread more rapidly through an organization.

Imitation of methods could occur even for specialized kinds of software. For example, tools like those used in the Colab can make methods and approaches (such as argumentation spreadsheets) more visible. Even in small organizations, some kinds of meetings are repeated so often that it would be worthwhile to create specialized methods and tools to support them. For example, financial planners could use liveboard-based tools for working through what-if scenarios with their clients. Attorneys could use special meeting tools to work through the terms of standard kinds of contracts. Thus, the methods and styles of argument can be rehearsed, shown, imitated, and reviewed in a seamless medium.

Work life in an organization is filled with conversations and to this end, seamlessness may turn out to be a key to the adoption of technology. Seamlessness facilitates imitation. We predict that conversation technology will spread when it is easy to use and is inexpensive.

We conjecture that seamless technology may lead to the more rapid propagation of ideas in organizations, affecting the resources and speed with which organizations can respond to new situations. Ideas that start in offices can spread to the coffee lounge. Similarly, ideas that arise from interactions in a coffee lounge can spread back to offices. Ideas can flow to or from meetings. In either case, an informal medium with memory can make it easier for people to explain ideas to each other and to combine and compare them. Thus, ideas become more portable between both locations and people.

For technology to be accepted, it must be perceived as filling an important need. For example, until a few years ago the common telephone was a form of office technology with a rude user interface. When a telephone rings, it is like someone is banging on your door, insisting that you drop whatever you are doing and give the caller your immediate attention. Nonetheless, ordinary rude telephones satisfied an important need and gained wide use. There is anecdotal evidence that Colab technology satisfies important needs. Faced with meeting in an unequipped office, a frequent Colab user may jokingly ask "how do I point?". Faced with a shortage of space on an ordinary whiteboard, a Colab user may ask "how do I shrink this?" or "how do I move this over?" Ordinary whiteboards are frustrating after experiencing superior media.

In closing, we return to our original theme of unlocking the creative genius inside of us. Genius takes many forms. It is not just the development of new ideas. There can be genius in negotiation, genius in management, genius in creating coherent plans, genius in all of the things that organizations do. Perhaps this genius can be unlocked by new tools for conversation that respect that we do not work in

isolation. In this view, the seamless tools for next generation meetings and conversations may shape our next generation organizations.

ACKNOWLEDGMENTS

This work was supported in part by contract N00140-87-C-9288 from the Defense Advanced Research Projects Agency. The chapter grew out of two talks: J. S. Brown's keynote address at the Conference on Computer Supported Cooperative Work at Austin, Texas in December 1986, and M. Stefik's address at the Symposium on Technological Support for Work Group Collaboration at New York University in May 1987.

The ideas in this chapter have evolved over an extended period of time in conversations with many others at Xerox PARC. Many of the ideas came from other members of the Colab project, including Dan Bobrow, Gregg Foster, Herb Jellinek, Stan Lanning, and Deborah Tatar. Special thanks also to Lucy Suchman and Stu Card for many conversations and insights. Thanks to Gregg Foster, Harley Davis, Giuliana Lavendel, Geoff Nunberg, Lee Sproull, Jeff Shrager, Stan Lanning, Deborah Tatar, and Eric Tribble who read earlier drafts of the chapter and offered helpful suggestions. Thanks also to Steve and Joan Osburn of Osburn Design, for contributing the sketches and for their work in designing the physical facilities of the Colab meeting room.