Evolutionary Engagement in an Ongoing Collaborative Work Process: A Case Study

Thomas P. Moran, Patrick Chiu, Steve Harrison, Gordon Kurtenbach, Scott Minneman, William van Melle*

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

{moran,chiu,harrison,minneman,vanmelle}@parc.xerox.com, gkurtenbach@alias.com

ABSTRACT

We describe a case study in which experimental collaboration technologies were used for over two years in the real, ongoing work process of intellectual property management (IPM) at Xerox PARC. The technologies include LiveBoard-based meeting support tools, laptop notetaking tools, digital audio recording, and workstation tools to later access and replay the meeting activities. In cooperation with the IPM manager, both the work process and the tools were continuously evolved to improve the process. We supported and observed over 60 meetings, leading to a rich set of empirical observations of the meeting activities. We note some practical lessons for this research approach.

KEYWORDS: activity capture, audio recording, co-development, evolutionary engagement, LiveBoard, meeting support tools, notetaking, salvaging, work process support

1. INTRODUCTION

Our research goal is to create tools that can support the natural informal communicative activities of interaction and collaboration. Much of the work of groups, even in such orderly settings as structured meetings, takes the form of casual interaction—the give-and-take of conversational exchanges whereby they generate new ideas and develop shared understandings. This casual activity is poorly supported by most computational tools, which tend to focus on the outcomes of such activity, while ignoring much of how the group arrived at those outcomes. Furthermore, many attempts to gather such information end up formalizing the activity, making the participants conform to a way of working that suits the computer rather than supporting their natural work practices.

Exploring such tools is as much, if not more, a matter of understanding how and in what ways they are useful as it is a matter of inventing new tools. Thus, our strategy, which we call *evolutionary engagement*, has been to engage with a real work setting to ground and motivate the process of creating new tools. This paper is a case study report of an extended engagement with an ongoing work process. We describe the mutual evolution of the tools, the practices, and the work setting. From this experience we present some observations on the nature of collaborative activity with such tools, and we note some practical lessons of this research approach.

Note that we are talking about two levels of collaboration. The first level is the collaboration within the work setting, and the second level is the collaboration between us, the researchers, and the participants in the work setting. The two levels are somewhat intertwined, but this paper is principally about the first level.

1.1. Meeting Support, Capture, and Access Tools

We have been working on collaborative technologies for many years. One path of research involves meeting support tools: the *Colab* [18] multi-workstation meeting room environment; the wall-size *LiveBoard* [9] pen-based interactive display; the *Tivoli* [16] whiteboard-emulation software with implicit structures [15] on the freeform surface. (In this paper, we use the term LiveBoard to refer both to the hardware and to the Tivoli software running on it.) The second research path involved multimedia: the *Media Space* [5] shared awareness environment and the *WhereWereWe* [13] capture, indexing, and playback tools. These paths have converged into a "confederation" of tools for meeting support, capture, and access [14].

Computer Supported Cooperative Work '96 , Cambridge MA USA

We would like to acknowledge and thank our colleagues who worked closely with us during the project described in this paper: Sara Bly, Chuck Hebel, Bill Janssen, and Don Kimber.

*Gordon Kurtenbach is now with Alias Research, 110 Richmond Street East, Toronto, Canada M5C 1P1.

This confederation is described later in this paper and in additional detail in [14]. Here we want primarily to emphasize the interrelated goals that we are trying to achieve with these tools. The first goal is to ***support*** the natural communicative and interactive activities that people engage in during the course of collaborative work. We want to provide tools that are immediately useful in informal collaborative settings. The LiveBoard is our main vehicle for such tools. At minimum, the tools must not inhibit or distort the natural activities. The second goal is to ***capture*** records of the activities. We do this with both the support applications and with unobtrusive multimedia recording. The third goal is to provide ways to ***index*** the captured materials. This can be accomplished by intentional activities (e.g., notetaking), as a side-effect of the natural collaborative activities (e.g., switching a page on the LiveBoard), or by automatic derivation (e.g., speaker identification in the captured audio). The final goal is to provide ***access*** to the indexed, captured records so that the original activities can be revisited and used as effective resources for further work. We have come to call this activity "salvaging"; it involves the use of browsers, media players, and editing and authoring tools.

We are not alone in our goals of meeting support or multimedia capture, and we build on earlier efforts. Our work is distinguished by our attempt to *integrate technologies* to provide support for *natural, informal* collaborative activities. Early meeting room support, such as the EDS *Capture Lab* [12], was aimed at supporting meeting processes. The EDS facility provided tools, e.g. for voting and decision making, that tended to impose order on the meeting process more than working with the informal activity. In spite of its name, capture was not provided as a resource for the participants.

Recent work on multimedia capture facilities, such as Bellcore's *Streams* [7], has focused on making and accessing recordings of presentations. These tools assume a model of use (e.g., speaker/audience setting, largely broadcast media distribution, remote viewing and playback) where notetaking is a private activity outside the scope of the system. Carroll's *Raison d'Etre* [6] is aimed at multimedia dissemination of retrospective project narratives, which are consciously recorded and authored by a project "historian." Our work, in contrast, tries to make structure and reuse an outcome of the naturally occurring work activity.

There has been some work in utilizing scribbled notes to improve the recall of meeting activities. The *We-Met* system [21] provided a shared pen-based display for notetaking and for later replay of the sequence of notetaking activities. The *FiloChat* system [20] used a pen-based computer and digital audio recording to provide a single user with a means to take notes in a meeting and to use the notes to index the replay of the recording. Although it deals with many issues common to our effort, the FiloChat work emphasizes personal use and does not encounter many issues that arise when the same functionality is used collaboratively.

1.2. Evolutionary Engagement Approach

Given our commitment to this line of research, we wanted to deal up front with the issue of utility — *Under what circumstances is it useful to access captured activities?* — as well as with the issues of usability — *How can we design tools that support natural activity?* We could best do this by engaging[1] a real setting with a real, ongoing work process and by evolving the design of the tools along with the practices and processes for utilizing them.

The first step was to identify a setting that was well matched to our research goals. We sought a work process that might benefit from using recorded materials and participants who would be willing to engage with us for an extended period of time. The setting had to be supportable with our limited resources; our tools were not portable, which implied a local setting. We found that the intellectual property activities at PARC satisfied these requirements. We were fortunate that it had an enthusiastic champion who welcomed exploring the technological possibilities and introspecting about the process changes they implied. This was key to making the engagement itself a mutual collaboration.

We devoted considerable preparatory effort to understand the culture of these activities and to plan the engagement. Our strategy was to support the current work practices in straightforward ways and then to evolve a series of small improvements of the technology and practices, which implied an extended time commitment. A big part of this task was to develop the trust and confidence of the setting's participants in the reliability of our tools and service, while at the same time managing their expectations of what we could deliver and the scope of our commitments.

We want to emphasize that we are researchers, not developers. While committed to engage with the specific concerns of the chosen setting, we oriented to those aspects of the setting that matched our long-term research objectives. Also, we wanted to go beyond the specific; we saw that many of the issues we encountered in the setting are more general. During this engagement, we have been rebuilding the architecture underlying many of our tools. This redesign is grounded by a sense of reality based on our experience of the work practices that make the tools meaningful.

Our approach is in the tradition of "action research" in that we intervened in a real-world situation. It is similar to the approach of *participatory design* of work and technology in Scandinavia [2], not so much at the level of design methods, but at the higher level of aiming to improve the quality of work, communication, and worker participation. Our approach bears strong resemblance to the techniques of *work-oriented design* [4,8] and *product co-development* [1,10]. However, these approaches are primarily used for system development, not research. But we do share with all these design practitioners the experience of rapidly co-evolving our technologies and the user's work practices. See [3,17] for other case studies of these evolutionary phenomena.

---

1. The verb **engage** has several pertinent and interesting meanings: to participate with, to assume an obligation, to become enmeshed with, to require the use of, to enter into conflict with, etc. [*American Heritage Dictionary*]

2. THE WORK PROCESS: INTELLECTUAL PROPERTY MANAGEMENT (IPM)

2.1. The Overall Process

PARC instituted a process for managing its intellectual property in the 1970s: a complex, ongoing process involving perhaps a hundred people, plus the community of researchers who are the inventors. Researchers are encouraged to document their novel ideas by writing *Invention Proposals* (IPs), which are typically 4-12 page (and occasionally longer) documents describing their inventions. IPs are being created continually, and they have to be evaluated and prioritized, so that the best and most important IPs are turned into patents.

This evaluation involves a complex decision-making process that is modeled on the scientific peer review process. The main vehicle is the *Technology Assessment Panel* (TAP) of technical experts from the labs. There are several different TAPs, covering a wide range of technical domains, such as solid state devices, large area electronics, image processing, software architectures, software applications, user interfaces. Each TAP meets on a regular basis (about every 6-8 weeks) to review submitted IPs and give them ratings. The results of the TAP evaluations are fed back to the inventors and reported to their managers. The IPs are then ranked by the TAPs, with some adjustments by research managers. Attorneys prepare patents from the IPs according to the rankings.

The central figure in this process is the manager of IPM who coordinates all its activities (in this paper, we call him *C*). His job is to keep the process running smoothly, report to management on the status of intellectual property, and to look for ways to improve the quality and quantity of intellectual property by refining the process and adapting better technological support.

The scope of the IPM process is very broad. Thus we focused our engagement on the central processes involving the TAP meetings.

2.2. How the TAPs Were Run

Let us look at the processes of the TAPs more closely to see how they were run before our engagement with them. C's role was to assign IPs to the appropriate TAPs, schedule and chair the TAP meetings, report on the results of the meetings, and coordinate with managers and attorneys.

Before each TAP meeting, C sent a set of IPs to the members as pre-reading. They brought their annotated copies of IPs to the meeting. In the meeting they dealt with the IPs one at a time. For each, they discussed its value along many dimensions, such as novelty, utility, readiness, urgency. There was a lot of give-and-take in the discussion, and members' views about an IP could be changed considerably in the process of trying to reach a consensus. They gave each IP a summary rating code reflecting the integration of the various dimensions; but from the discussion there were always a lot of qualifications, caveats, suggestions, etc., beyond the rating code. C managed the agenda, facilitated the IP discussions, brought the group to consensus on the rating, and at the same time took extensive handwritten notes on the content of the discussions, so that he could report on them.

C's most critical task was outside the TAP meetings: creating reports summarizing each IP and its assessment. These reports served many purposes. They provided feedback to the inventors; they informed the managers and attorneys about new intellectual property; and they helped TAP members recall the status of the set of IPs when they later ranked them. Given the extremely wide range of deeply technical matter in the different TAPs, this was a challenging task.[2]

C created his IP reports after the TAP meeting, using the notes he took during the meeting. Because of his workload, he was often not able to work on a report until a month or more after the TAP, by which time his notes were difficult to interpret. That, plus the technicality and subtlety of the arguments, drove C to extensive consultations with some of the TAP members. Sometimes this was not possible and very sketchy reports had to suffice. C then circulated a draft of the IP reports to the TAP members for critique, and finally he issued the IP reports to the inventors, managers, and attorneys. Inventors sometimes felt that these reports were not accurate. C found this very frustrating.

Each TAP met twice a year to rank all the recent and still-active IPs by priority for patenting. This was not a simple matter of sorting them by rating, since other factors, such as business considerations, new technical developments, and the history of the IPs, came into consideration. The members had printouts of the current rank ordering, plus summary notes about all the IPs. The first step in a ranking meeting was to review all the IPs (in the current rank order) to remind everyone what they were about. During these reviews, members often revised their assessments of the IPs. In the process, they also noted which IPs should be ranked higher or lower and which IPs should be deactivated. The actual ranking technique was to sort the IPs into quartiles. Usually, this was first done roughly by quartile, followed by a refined sorting within quartile. Often, IPs were grouped with related IPs that could be treated as units. Sometimes C simply took notes on a printout of the current ranking as members made suggestions for ranking changes. Sometimes they wrote lists of IPs on a whiteboard as they suggested different placements of IPs in the new ordering. This process could get very confusing at times: when there were many IPs to consider, when the changes from the current ranking were great, or when there was a lot of contention about ranks. After the meeting, C prepared a ranking document that was circulated to the managers for further changes in the TAP's rankings.

The mechanics of running the TAPs was constrained by paper-based techniques. This was the case not only in the meetings themselves, but more importantly in the larger, more subtle communicative role that the TAPs played in PARC. As members of the technical community, TAP members serve to cross-fertilize ideas throughout the center across workgroup boundaries. One important reason for tar-

2. C was formerly a researcher and so was knowledgeable about some of the technologies, but he was a novice at others.

geting the TAPs for engagement was that we could work on the mechanics as a way to get at the larger issues.

2.3. Scope of This Case Study

The case study described in this paper covers over three years of our involvement with the IPM process: a "pre-engagement" period of a year and the engagement with the TAPs for two years.

We report here on the activities of five of the TAPs that we supported and observed. These are representative of the technical domains, size, breadth, and personality of all the TAPs. During the two-year engagement, the five TAPs had over 60 meetings, totalling over 100 hours. Each TAP met from 8 to 16 times. Four of the TAPs averaged 4 to 6 people at their meetings (including C), while the largest TAP averaged 9 people. The meetings ranged from 20 minutes to over 3.5 hours long, averaging 100 minutes. The TAPs considered anywhere from 1 to 14 IPs per meeting, averaging 6 per meeting. They spent an average of 13 minutes on each IP.

We provided meeting support and capture technology for all the TAP meetings during the engagement, plus a workstation for C to access the captured records and create reports. At each TAP meeting, we had an observer who took notes on the meeting process, debriefed C after the meeting, and emailed project notes to the rest of our group. (We draw quotes from these later in this paper.) Each TAP meeting was also videotaped, both as a backup and as an aid in creating process notes. We also had a technical support person available for each meeting and salvaging session.

3. PREPARATION FOR ENGAGEMENT

3.1. Preparation Prior to Engagement

We first considered the TAP as a target work setting in October 1992, and we began having discussions with C about how our technologies might contribute. For a year after that[3] we observed several TAP meetings, videotaping some. We also tried some experiments to get a feel for how our technologies might fit into the IPM processes. For example, we videotaped and indexed[4] two IP brainstorming meetings and had researchers use the videos to cull the interesting ideas that came up. In mid-1993 we proposed an experiment to one of the TAPs in which we would make a video of the TAP discussion and then use segments from it as a rich feedback to the inventors. The TAP members resisted this, because they did not want their uninhibited discussions, which they felt were valuable, leaving the room. We abandoned this experiment, but the experience raised our awareness of the sensitivity of making and using recordings in this setting.

We developed a strategy of minimizing changes in the TAP processes, at least initially. We focused on the LiveBoard to support the meeting process and on audio capture to help C with the reporting process. By Fall 1993 we committed to the TAP as a work setting, and we worked on shaping our technologies to fit into the ongoing TAP practices, targeting the first TAP meetings of 1994.

3. We were considering other work settings during this period.
4. We used the Marquee video logging tool [19].

Figure 1. Plan View of the Technology and the Physical Arrangement of TAP Meetings.

Twelve people can sit around the meeting table, the TAP members **M**, the coordinator **C**, and an observer **O**. In front of the table is the LiveBoard **L**, which the coordinator **C'** operates. On the table are two audio microphones **A**. The laptop notetaker **N** was introduced in 6/94.

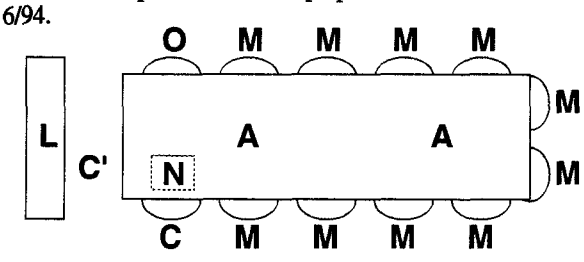

3.2. The Collaboration Technologies

We set up one of the PARC meeting rooms for the TAPs. Figure 1 shows how the technologies fit into the room. The goal of the setup was to keep the look and feel of an ordinary meeting room. There was a long table holding up to 12 people with two audio microphones. The LiveBoard at the head of the table was easy for everyone to see. C and the observer occupied the chairs near the LiveBoard. With this physical arrangement, the TAP members were in close proximity and could easily interact directly with each other.

The LiveBoard ran the Tivoli whiteboard program to display pages of materials to support each TAP meeting. The first page held the agenda of IPs to be reviewed. Following this was one review page for each IP on the agenda. A review page provided a form to record the rating and space for taking handwritten notes of the discussion of the IP; Figure 2 shows an example of a simple review page.

A user interface was provided in Tivoli to start up a digital audio recorder and to pause and restart it. A flashing red "Recording" sign was displayed during recording (see Figure 2). The added value provided by this technology was automatic indexing of the audio recording. Every action in Tivoli, such as switching a page (signalling that the TAP's attention was switching to another IP) or making a stroke with the LiveBoard pen, was timestamped by Tivoli, thus providing index points into the audio record.

C was provided with a workstation, which we called the "Salvage Station," that provided access to the materials created on the LiveBoard and the audio recording made during the TAP meeting. The Salvage Station interface, shown in Figure 3, presented the Tivoli display, a set of playback controls (play, stop, and forward or backward 10 seconds), and a word processor for creating the IP reports.

The Salvage Station provided C with random access into the audio at the index points. For example, we see in Figure 2 that C had scribbled a note "922261 may be similar to this IP." If he wanted to replay the discussion at the time that note

Figure 2. A Typical Tivoli IP Review Page.
[from the TAP meeting of 3/15/94]

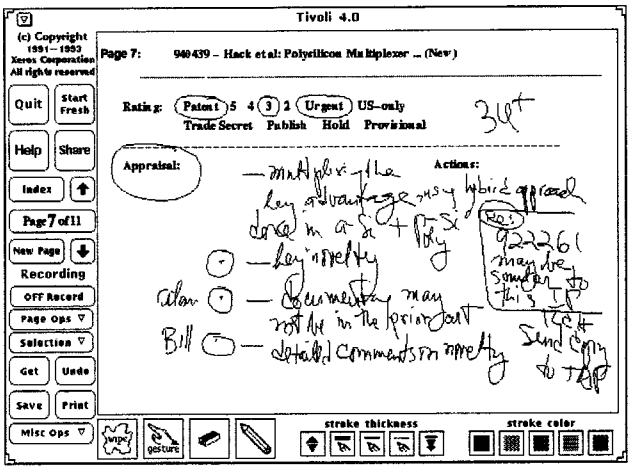

was made, all he had to do was touch one of the note's strokes, the "9" say, with Tivoli's "play tool" to hear what was said when he wrote the "9." Thus, by using the materials from the meeting, C could quickly access the discussion for a particular IP by using the Tivoli page structure, and he could access points within an IP discussion by using the notes he took on the LiveBoard during that discussion.

4. ENGAGEMENT AND EVOLUTION

4.1. Initial Engagements

> With the new year I am going to try to do a better job of capturing your comments and insights on IPs and providing better TAP [reports] and written feedback to inventors and the attorneys. [Email from C to TAP, 1/18/94]

The initial engagement with each TAP was critical to attaining credibility. There were two main concerns: They did not want the productivity of their meetings disrupted, and they were wary about being recorded.

We addressed the disruption concern by not changing the way they conducted their meetings. The only visible difference to them was the presence of the LiveBoard on which C took handwritten notes (facing away from them) and that C often jumped up from the table to the LiveBoard.

We addressed the concern about recording in two ways. First, we established an agreement that the audio was to be used only by C to help him write reports.[5] Second, we agreed to record only short "recap" segments of the meeting. Recaps did require a small change in meeting process, which seemed like a good thing to do in any case. The TAP members agreed that, after each IP discussion, they would take a few minutes to recap the discussion *for the audio record.* C had to manually pause and restart the audio recorder; thus, everyone was clearly aware of when the recorder was on.

5. They also agreed to our videotaping all the TAP meetings during the engagement. Only members of the research project had access to the videotapes, the audio record, and written materials.

Figure 3. Salvage Station Display.

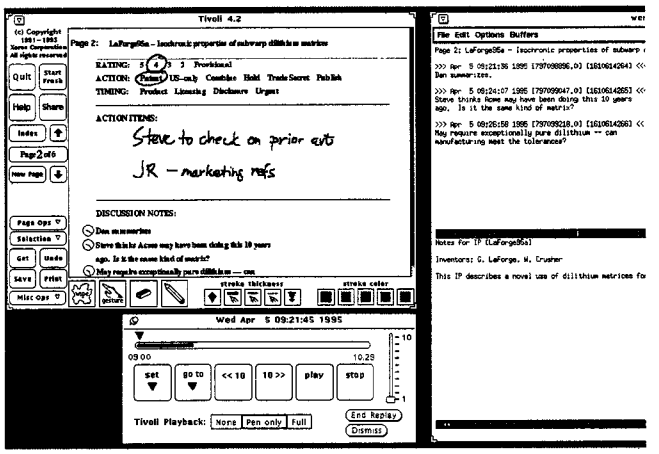

C briefed the TAPs about the changes beforehand via email and at the beginning of the meetings. There were a lot of new procedural steps required of C to make everything go smoothly. For example, part of the procedure we worked out was to circle the word "Appraisal" on the LiveBoard (see Figure 2) to index the start of the recap recording. We practiced the procedure with him in a rehearsal session.

The initial three TAP meetings proceeded smoothly, without incident or disruption. From the point of view of the TAP members very little was different.

The fourth TAP was a ranking meeting. For this the LiveBoard was used to manipulate lists of IPs as the TAP members discussed them and gradually decided on the priority ordering.[6] The "implicit structure" capability [15] of Tivoli made list manipulation easy, and TAP members were "wowed" as items on the LiveBoard seemed to slide magically into place while items were being moved around. Also, they were able to make scribbled annotations in the lists. The TAP members appreciated the ease with which the ranking task could now be done. The technology was more noticeable in the ranking meeting, because it provided an immediate service and improvement in the meeting process.

C used the Salvage Station to create IP reports. He was able to easily get to each recorded recap by using the circle around the "Appraisal" as an index. He would listen to an entire recap, and then sometimes listened to portions of it again. Interestingly, C often let the audio continue playing while he was composing the text of his reports.

4.2. Evolution of Technology and Practice

Many new features of the technology evolved from the needs of practice, and these new features in turn changed the practice as it adapted to the new features. The timeline in Figure 4 shows how these evolved in historical order. We illustrate the evolution with a few examples:

6. Ranking sessions, having no recaps, were not recorded.

Figure 4. Time Line of the Case Study.

Date	*Pre-Engagement*	*Sec.*[a]
10/92	first suggestion to focus on TAPs; first discussions with C	3.1
11/92 –5/93	observe TAP meetings	
2/93	experiment: video of IP brainstorm	
7/93	proposed experiment on video feedback to inventors rejected by TAP	
7/93	experiment: video of IP brainstorm	
11/93 –12/93	prepare tools for engagement (meeting room; LiveBoard; audio recorder; Salvage Station; procedures; training)	3.2
1994	***Engagement Begins***	
1/19	audio recording of discussion recaps only; scribbled LiveBoard notes	4.1
1/26	C uses "audio button" idiom	4.2 E
1/26	first "topic notes"	4.2 F
2/2	first "process notes"	
2/15	LiveBoard support of ranking process	4.1
3/2	TAP members compliment C on reports	4.2 A
4/6	double backup for audio recorder	5.2 B
4/13	start recording audio of whole meeting	4.2 B
	On/Off Record controls	4.2 C
5/25	clock objects	4.2 E
6/23	laptop notetaker (with manual beaming)	4.2 G
7/27	auto-beaming of notes from laptop	
9/26	poll of two TAPs re laptop notetaking	
10/14	experiment: salvaging TAP meeting for new ideas	5.2 F
11/10	notetaking during ranking (for indexing)	6.1 C
12/6	link objects (for jumping between pages)	4.2 H
1995		
2/1	incremental ranking started	4.2 I
2/8	new model LiveBoard installed	
2/13	experiment: managers meeting	5.2 F
3/22	data imported from Lotus Notes	4.2 H
9/21	"quick reports" of TAP meetings	6.1 B
10/3	take "bottom line" notes in TAP	
11/2	use "hear audio" notation	6.1 C
11/17	experiment: C salvaging TAP with speaker identification information	5.2 F

a. Section of this paper where item is discussed.

A. Initial Feedback From the New Reporting Capability. An important aspect of the initial engagements was building credibility. C had long had the most difficulty reporting on the largest TAP. While its first meeting with audio recording proceeded with little disruption, the benefits of the audio had to be taken on faith. By the time of its second meeting, C had used the audio to create reports, which he had circulated to the TAP members for their review. In the second meeting, the TAP members immediately complimented C on his "better notes" and said that they "appreciated the greater thoroughness." There was, from this time on, a greater appreciation of the audio. They felt comfortable just making remarks, knowing it would be captured and picked up later by C.

B. Extent of Audio Capture. Recording only the recaps proved to be unworkable. First, a strict boundary between discussion and recap was difficult to maintain. Second, pertinent remarks made during discussion were not captured:

> During a lengthy discussion phase (not recorded) of an IP, R apparently said something that clarified a lot of the previous discussion; C said "that's great, R, remember to say it again during the [recap]." 10 or 15 minutes later during the [recap] (being recorded), C did remember to ask R to state his insight again, but of course R couldn't remember what it was he had said or even what the point was he was making at the time. [Project notes, 2/4/94]

> At one point, someone said "oops, we're not recording." [Project notes, 3/4/94]

Third, there was the hassle of turning the recorder on and off, and often there were errors in doing this (both too much recorded and failure to record a recap). After three months (during which most TAPs had met twice), the TAP members were ready to accept having the whole meeting recorded.

C. On/Off Record Controls. The TAP members wanted a feeling of being in control of the full recording; so an *On/Off Record* facility was provided. It not only allowed forward control (i.e., pause and restart), but also backward control (i.e., it allowed the last *n* minutes of the meeting to be either put On or taken Off the Record).[7] While the TAP members were pleased to have this control, they actually made little use of it. Occasionally, they would pause the recording for a time when they anticipated a sensitive discussion coming up. They never took anything already recorded Off the Record, though the possibility did arise:

> At one point someone made a remark that offended one of the TAP members. There was an awkward moment, as this startled everyone. There seemed to be concern for the confidentiality of the meeting (the offended person looked at [the observer]). However, asking to take this Off the Record would only have emphasized the awkward moment that the group wanted to just pass by. [Project notes, 4/14/94]

D. Notetaking Length. C's style of taking notes during an IP discussion evolved considerably in response to the technology, the effect on the TAP meeting, and the usefulness in reporting. Before the engagement, C took extensive notes on his hardcopies of IPs. At the beginning of the engagement, C took handwritten notes on the LiveBoard during the recaps,

7. This was implemented by recording the whole meeting, with segments of the meeting being marked as On or Off the Record. At the end of the meeting, all the Off-Record segments were purged.

and also on hardcopy during the discussions. His LiveBoard notes were long. Since he was the focus of attention during recaps, the members had to wait while he wrote them, which noticeably slowed down the meeting. Then C started writing very short notes on the LiveBoard, taking fuller notes on the hardcopy IPs during the discussions. Then, when the whole meeting started being recorded, he again wrote long notes on the LiveBoard. Since these notes were taken during the IP discussions between the TAP members, they didn't pay much attention to whether he was keeping up with them. C again shortened his notes so he could keep up with the discussion. When the laptop became available, C began taking longer notes (roughly 50% longer), because he could type them much faster.

E. Notetaking Idioms. Early in his LiveBoard use, C drew little circles-with-dots that he called "audio buttons" (see Figure 2). These graphically served as bullets marking each note, plus they represented entry points into the audio that he could use during salvaging. We eventually reified these as *clock* objects on the LiveBoard. A clock object looks like an analog clock with the hands set to the time it represents. Clocks were used to represent the time of beamed notes from the laptop (see Figure 3). They could also be created manually by gestures. Clocks had the feature that they visually portrayed the course of time over the static spatial image of the LiveBoard page. A related idiom was an offset time notation. Because there is always a delay in notetaking (it takes time to recognize something as noteworthy), C would write e.g. "–30" to indicate that the noted event started 30 seconds before the start of the note. During salvaging, he would touch the audio button or clock and then immediately reset the audio play point back 30 seconds.

F. Notetaking As Indexing. Before the engagement, C's notes had to carry the content and meaning of the discussion, for this is all he had to rely on in creating the reports. But with audio as a resource, the notes only had to serve as *indices* into the audio. We can crudely classify C's notes into three types (from verbose to terse): *content notes* (e.g., "This idea seems not patently distinguishable, but could be covered as a claim in..."), *topic notes* (e.g., "discussion of novelty"), and *process notes* (e.g., "Robert's comment"). Topic and process notes are indexical and need the audio (or one's memory) to fill out the content. C quickly began to use topic and process notes. C used mostly topic notes, but he also seemed to find person-centered process notes a meaningful way to parse discussions. There were very few content notes when he had to scribble them on the LiveBoard; but, after he got the laptop, more content notes appeared (roughly 1/4 of the notes). The ability to use indexical notes relieved C from having to distill the full meaning of the discussion into his notes and thus permitted him to pay more complete attention to the discussion.

G. Public and Private Notetaking. Before the engagement, C's paper-based notes were private. With the advent of the LiveBoard, part of C's notes (those taken during the recaps) became public. With full audio recording, all his scribbled notes became public. At C's request we provided him with a laptop so he could type his notes at the table. We provided the capability to "beam" his typed notes to the LiveBoard, so that they could be integrated with the ratings and notes on the LiveBoard. We faced a question of how public to make C's typed notes. We did not want the TAP members to be diverted from attending to each other during the discussions, which argued for the notes to remain private. At first, we allowed C to "manually" beam all his notes at once at the end of a discussion, at which point the TAP members did not want to take the time to review them. Some members requested that the notes be made more public, because the notes helped them see whether C was understanding them. We did not want to beam them continuously (word at a time, say), because we felt the resulting dynamics on the LiveBoard would be distracting. A compromise was to beam a note at a time. (A typed note was taken every 2-3 minutes and consisted of 1-5 lines of text.) This allowed TAP members to glance at the notes as they came up. Some members did this rather consistently, while others ignored them. Occasionally, there were comments on the accuracy of a note, and scribbled annotations were made on the LiveBoard. Members of two of the TAPs were polled on this facility, and all were satisfied with this beaming scheme.

H. Import IP data from Notes. A Lotus Notes database of IPM information was independently being created during this period by C. Once it was available, information was exported from Notes to automatically create LiveBoard pages for each TAP (which was previously done manually). Ranking, summary, and status information about all the active IPs for a TAP was also included. Links between pages allowed easy navigation between agenda, reviewing, ranking, and summary/status pages. Thus, a "working set" of information for a TAP was available at every meeting and easily accessible on the LiveBoard. This replaced much paper that used to be brought to the meetings. This capability was used extensively, especially during ranking sessions.

I. Incremental ranking. Because a ranking page was being provided at each TAP meeting, C began a practice of ranking IPs incrementally at each meeting, rather than waiting to do it semi-annually. Usually, the new IPs reviewed at a meeting were slotted into the current rank ordering, although this step was sometimes skipped (members often ran out of steam after reviewing a lot of IPs). In 1994, only 40% of the meetings included ranking activity; in 1995 some ranking was done in over 70% of the meetings. The incremental rankings were partial rankings, so there was still a need for a thorough reconsideration of the IP rankings at the end in 1995.

5. EMPIRICAL RESULTS

Having supported, designed for, and observed so much collaborative activity during this engagement, we observe some of features of the collaborative work with these technologies and reflect on this style of research.

5.1. Features of Collaborative Activities

A. Focus Between People vs. Focus on Artifacts. The members of one TAP reacted strongly when, on an occasion in which a member could not be present, it was suggested that

they have a discussion on-line; they wanted to preserve what they called the "interaction chemistry" that a face-to-face meeting provided. We have tried to preserve this essential feature of the TAP meetings by not diverting them from each other during discussions. That is why notetaking is done "on the side" where it can be ignored. Sometimes an artifact-centered mode of discussion is appropriate, such as when considering a complex ranking or when writing down a consensus rating. Focus of attention is a very subtle issue to be carefully considered in designing collaborative tools [12].

B. Rapid Shifting Between Topics in a Meeting. A productive meeting is a balance between staying focused and following the flow of topics. Any technology that depends on clearly categorizing segments of a meeting can expect trouble. An example of this is our early attempt to record only the recap portions of the TAPs. At first this seemed to work cleanly. But gradually the boundary between discussion and recap eroded, and it became very difficult to manage the recorder. Often, when they were recapping, a new idea arose, and they were immediately back into discussion. Another example is the shifts between IPs being discussed. Sometimes, when a discussion brought up a point that applied to an earlier IP, the group was able to shift their attention effortlessly to the other IP for a moment and then back. It is difficult for any technology to track these rapid, subtle shifts.

C. Leaping Beyond Structural Categories. A related phenomenon is that the activities in a meeting seldom stay strictly within the confines of a particular structure or agenda. This is a problem for collaborative tools. An example of this is the ranking activity. We provided a representation that allowed the TAP members to move IPs around in various regions of the LiveBoard to rank and categorize them. This structured manipulation capability was provided on top of a freeform scribbling capability [15]. This was fortunate, for some kind of scribbled annotations were always being made that would have been impossible to anticipate (e.g., see Figure 5). Without this capability, the meeting activities would have been overly confined.

D. Awareness of Audio as a Resource. One might guess that the TAP members would become unaware of the audio recording capability once they were used to it. To a great extent that seems to be the case. There is little evidence that it inhibited discussion. (But remember that they trusted us to keep the audio records confined to C's use.) On the other hand, they were constantly aware of it as a resource. C would remark "I'll get that off the audio" a few times each meeting. Often C would do something to create an index point for the audio. For example, he would ask a member to hold a remark a second so he could start a new note indexing the anticipated remark; a couple of times he had someone repeat a short remark so it could be more accurately indexed. Another example was the "prompted index," where C would ask someone a question or prompt someone to make a comment while he created a note anticipating a valuable response. In general, TAP members felt that their talk at the meeting was being heard and used when appropriate. In that sense, the audio allowed the meeting to flow freely without worrying that something they said might be missed.

Figure 5. Representation for Ranking IPs.
[from the TAP meeting of 12/15/95]

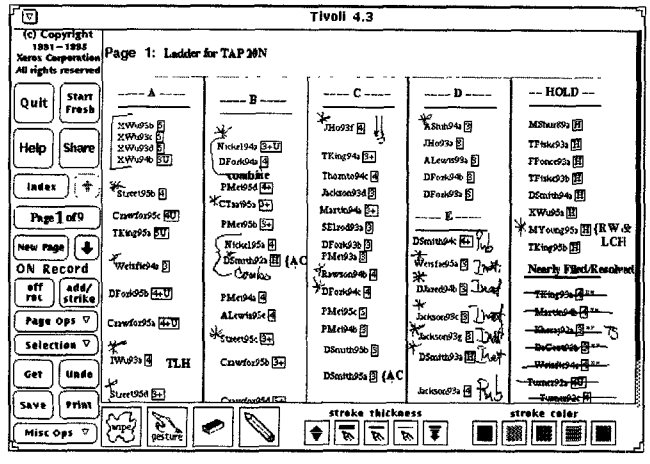

E. Coordination and Competition Between Paper and Electronic Materials. The TAP meetings have a lot of materials to deal with. Much was in paper form: the IPs being reviewed, the agenda, private notes, old IPs and reports on them, ranking and status documents, etc. To this we added materials on the LiveBoard. One issue was redundancy. For example, there was an agenda on the LiveBoard in addition to a paper agenda sent to each member before the meeting. In the latter half of the engagement, we provided the capability that all the ratings from the review pages automatically appeared on the agenda page, thus providing a summary of the "results" of the meeting. However, C had a hard time shaking the habit of writing the ratings on his hardcopy agenda, even though he always made a printout of the LiveBoard materials. The issue of coordinating different materials arose especially during the period when part of the material was on paper and part on the LiveBoard. It was difficult to look at an item on the LiveBoard and then find the corresponding information on paper, thus causing a lot of wasted time shuffling through papers. This is not a hard problem to solve; but it is an easy one to overlook.

5.2. Pragmatics of Our Approach

Committing to an engagement with a real work setting for an extended period of time with a research prototype had considerable consequence for our day-to-day activities. Our experience was that you don't have to have a production system to do this kind of research, but you have to be prepared to compensate by showing that you are willing to make it work. You need to have, in return, the good faith and cooperation of the people you are engaging. Some lessons:

A. Be There. Tools don't run themselves. We found that we *always* had to be on hand when our tools were being used. At least one of us was always at each TAP meeting, to supply user advice, technical support, and moral support:

> Do the extra people in the TAP meetings (observer and technology supporter) disturb the meetings? -- I think it is the opposite; technology supporters make the occasional breakdowns not seem frustrating since we know it will be taken care of. [Email response from a TAP member, 9/26/94]

Occasionally, one of us would operate the tools for C in a TAP meeting to test if they were what was needed; then we refined them so that C could use them himself. Another benefit of being there is being able to observe.

B. Expect Disasters. Murphy's Law prevails. Needed equipment was "borrowed"; plugs were pulled; servers crashed revealing hidden system dependencies:

> The biggest glitch was that we were unable to connect Tivoli to WWW to make the recording. This is because the NMS on Zosma was not propagating to Liveboard01. S went down to CSL to investigate. Seems that since CSL is no longer making much use of NMS, nobody noticed when the master (and all other servers) stopped functioning.[project notes, 7/20/94]

We added a second level of audio backup, because C had become dependent on the audio. Analog backup tapes saved the day more than once for us.

C. Manual Effort. A great deal of manual effort preceded automated tools. For example, we manually set up the LiveBoard materials for the TAP meeting for many months before we automated this task. If we had not been willing to do this, we would have had to delay the engagement.

D. Tools Don't Have To Be Perfect. We put tools into use that were not ready for a general release. For example, the gestural manipulation tools for ranking on the LiveBoard were a bit quirky. Rehearsing and being there is what made this work. C's patience and sense of humor (e.g., his "rule is that it always works the second time you try it") and the understanding of the TAP members also helped. What was important was not so much the quirks, but that they were able to get their job done in spite of them. That is, the tools do have to be adequate.

E. Continuously Reflect on Process. We had a post-mortem with C after each TAP meeting about everything from the nitty-gritty of operating the tools to how the meeting went to the larger issues of the IPM process. Identified problems became goals to improve the tools. Easy fixes were done right away; larger fixes were prioritized with other development tasks. We had occasional discussions with TAP members to get their views. At the end of the first year, we had discussions about tools in each TAP. We found that discussions with individuals provided the most useful feedback.[8]

F. Side Experiments. Having become a part of the work setting, we were able to use the setting for several one-off experiments. For example, we extended the LiveBoard tools to support a lab managers' review of IPs. We explored the notion of salvaging the TAP recordings for new ideas (which often came up during discussions). We also experimented with automatically identifying speakers in the captured audio (see [11]) in two of the TAPs. This gave C the ability to see who was speaking when on a timeline on the Salvage Station. We ran a test of this with C, though we have not yet put this feature into regular use.

8. We did not focus as much on observing the salvaging activities as we did on the TAP meeting activities. This was because we felt the TAP setting was more sensitive and more important to get right at first. We did work closely with C in making sure the Salvage Station tools worked for him. In the latter part of the engagement, we made some closer observations of C's salvaging activities.

6. CONCLUSIONS AND REFLECTIONS

> [The TAP reports] seem much more comprehensive than those in the past. Much more useful in reconstructing what happened in the meeting, I think, than those of a year ago. ... It basically has no impact on my behavior at the meeting, aside from perhaps a tad of self-editing—but I'm not even sure about that. [email from a TAP member to C, 10/10/94]

We have installed a suite of novel collaborative tools into the IPM process, providing an ongoing meeting support, capture, and access service—the first reported demonstration of the utility of meeting capture tools in a real business process. The evolutionary process of creating and adopting the tools allowed them to sink deeply into the fabric of the process.

6.1. Impact

The impact of these tools has been substantial. Before our engagement, the IPM process was near breakdown. The complexity, breadth, volume, and pace of C's job was extremely frustrating. The requirement to assimilate and accurately report on a wealth of diverse technical discussions were exceeding C's ability to cope. The new tools turned this around, enabling C to be more effective and to feel more effective. He is able to produce qualitatively more complete, accurate, and timely reports of TAP discussions, according to TAP members, who are able to interact as freely as ever in the TAP meetings. C now has a greater range of job responsibilities and less support help, yet he is able to manage the IPM process more efficiently.

6.2. Reflections

A. Meeting Support. We have presented a wealth of results on the evolution and features of the practices and behaviors in meetings and their interactions with our meeting support and capture tools. The subtleties of how tools fit in to the meeting process is extremely important, for meeting interactions have a subtle fabric that is easily destroyed by insensitively-designed tools. We have succeeded in our goal that the TAP meetings retain their highly interactive, discussion-intensive character; the discussions flow quickly with the confidence that the important remarks will be picked up. The display of C's notes on the LiveBoard provides an additional check on C's understanding. The LiveBoard tools provide a good focus for making clear the consensus on ratings. The shared representation we have provided on the LiveBoard has qualitatively improved the ranking process. The facility to interactively manipulate the rankings makes the process of determining the ranking efficient, iterative, and truly collaborative. The linked-in working set of active IP information provides effective support for relating current IPs to the larger set.

B. Reporting. C is able to use the captured audio from his Salvage Station to produce noticeably more complete and accurate reports of IPs, because of the "riches" provided by

the audio record. Reports are written more quickly, because he seldom has to contact TAP members after the meeting to clear up points—a check of the audio is usually sufficient. It is difficult to measure the gains, because the reporting process is highly distributed and is not done just during explicit report-writing sessions. C has recently focused on more efficient reporting by disciplining himself to issue "quick reports" within a few days of a TAP meeting. To support this he has instituted a practice during the TAP meetings of making clear what the "bottom line" of each IP discussion is (similar to the earlier recap practice, but without worrying about the state of the audio recorder).

C. Indexing and Salvaging. One of our goals is to provide effortless indexing into captured records. We successfully provide indexing into meeting segments defined by IP discussions (at a 10-15 minute granularity). C can get to the correct IP discussion segment with no problem. We also provide indices with the notes taken on the laptop (at a 2-3 minute granularity). Recently, C has been marking a few of his notes with "HA" (hear audio) annotations, which he uses to focus his salvaging effort. We are also experimenting with speaker segmentation tools that provide indexing at a conversational interchange granularity. The task of salvaging is complex, and we will report on it in a later paper.

6.3. The Story is Not Over ...

It is difficult to come to "conclusions" because the evolution continues. We have become part of the work setting and the process. During the second year of the engagement we spent much of our time rebuilding the foundations of our tools to make them more architecturally solid and customizable. When they get installed in the work setting, they will enable a range of new possibilities for changes in the work process, such as multiple media documentation. But technological readiness is only half the story. The engagement has spawned reflections on the overall IPM process, and many of the participants are exploring new ways to improve the process. We are looking beyond the question of making reports better and reporting more efficient, to considering the issue of what role the reports play for all the parties involved and whether there could be a better form of documentation of TAP results and, indeed, of inventions. Researchers, managers, attorneys, and product managers each have different needs. The goal is not just to supply them with reports, but to enhance the communication between them. Especially, the processes of creating intellectual property needs to be made a more organic part of the natural research activities in the lab. Thus, we want our engagement in the process to go much deeper and to effect more fundamental changes.

ACKNOWLEDGMENTS

In addition to our close colleagues (mentioned on the front page), we thank the many TAP members for their supportive collaboration. We thank Polle Zellweger for many discussions, detailed critiques, and thoughtful suggestions on the paper. Leysia Palen worked with us on understanding the salvaging process. Karon Weber and Alex Poon carried out many of the early pre-engagement experiments with video indexing using Marquee [19].

REFERENCES

[1] Anderson, W., & Crocca, W. (1993). Engineering practice and codevelopment of product prototypes. *Communications of the ACM, 36*(6), 49-56.

[2] Bjerknes, G., Ehn, P., & Kyng, M. (eds.) (1987). *Computers and democracy: A Scandinavian challenge.* Aldershot, England: Gower Publishing.

[3] Blomberg, J. (1986). The variable impact of computer technologies on the organization of work activities. *Proceedings of CWCW'86*, 35-42.

[4] Blomberg, J., Suchman, L., & Trigg, R. (1994). Reflections on work-oriented design. *Conf. on Participatory Design*, 99-109.

[5] Bly, S., Harrison, S., & Irwin, S. (1993). Media spaces: Bringing people together in a video, audio, and computing environment. *Communications of the ACM, 36*(1), 28-47.

[6] Carroll, J., Alpert, S., Karat, J., Van Dusen, M., & Rosson, M. (1994). Raison d'Etre: Capturing design history and rationale in multimedia narratives. *Proceedings of CHI'94*, 192-197.

[7] Cruz, G., & Hill, R. (1994). Capturing and playing multimedia events with STREAMS. *Proceedings of Multimedia'94*, 193-200.

[8] Ehn, P. (1988). *Work-oriented design of computer artifacts.* Stockholm: Swedish Center for Working Life.

[9] Elrod, S., Bruce, R., et al. (1992). LiveBoard: A large interactive display supporting group meetings, presentations, and remote collaboration, *Proceedings of CHI'92.*

[10] Greenbaum, J., & Kyng, M. (Eds.) (1991). *Design at work: Cooperative design of computer systems.* Hillsdale, NJ: Erlbaum.

[11] Kimber, D., Wilcox, L., Chen, F., & Moran, T. P. (1995). Speaker segmentation for browsing recorded audio. *Proceedings of CHI'95*, 212-213.

[12] Mantei, M. (1989). Observation of executives using a computerized supported meeting environment. *International Journal of Decision Support Systems*, 153-166.

[13] Minneman, S., & Harrison, S. (1993). Where Were We: Making and using near-synchronous, pre-narrative video. *Proceedings of Multimedia'93*, 207-214.

[14] Minneman, S., Harrison, S., Janssen, B., Kurtenbach, G., Moran, T. P., Smith, I., & van Melle, W. (1995). A confederation of tools for capturing and accessing collaborative activity. *Proceedings of Multimedia'95.*

[15] Moran, T., Chiu, P., van Melle, W., & Kurtenbach, G. (1995). Implicit structures for pen-based systems within a freeform interaction paradigm. *Proceedings CHI'95*, 487-494.

[16] Pedersen, E., McCall, K., Moran, T. P., & Halasz, F. (1993). Tivoli: An electronic whiteboard for informal workgroup meetings. *Proceedings of INTERCHI'93*, 391-389.

[17] Rogers, Y. (1994). Exploring obstacles: Integrating CSCW in evolving organizations. *Proceedings of CWCW'94*, 67-77.

[18] Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S., & Suchman, L. (1987). Beyond the chalkboard: Computer support for collaboration and problem-solving in meetings. *Communications of the ACM, 30*, 32-47.

[19] Weber, K., & Poon, A. (1994). Marquee: A tool for real-time video logging, *Proceedings of CHI'94*, 58-64.

[20] Whittaker, S., Hyland, P., & Wiley, M. (1994). Filochat: Handwritten notes provide access to recorded conversations. *Proceedings of CHI'94*, 271-277.

[21] Wolf, C., Rhyne, J., & Briggs, L. (1992). Communication and information retrieval with a pen-based meeting support tool. *Proceedings of CSCW'92*, 322-329.