Collaborative Knowledge Management Supporting Mars Mission Scientists

Irene Tollinger
NASA Ames Research Center
MS 262-4
Moffett Field, CA 94035
+1(650)604-5740
Irene.V.Tollinger@nasa.gov

Michael McCurdy
NASA Ames Research Center
MS 262-4
Moffett Field, CA 94035
+1(650)604-3291
Michael.McCurdy@nasa.gov

Alonso H. Vera
Carnegie Mellon University
NASA Ames Research Center
MS 262-4
Moffett Field, CA 94035
+1(650)604-6294
Alonso.H.Vera@nasa.gov

Preston Tollinger
San Jose State University
NASA Ames Research Center
MS 262-4
Moffett Field, CA 94035
+1(650)604-1216
ptollinger@arc.nasa.gov

ABSTRACT

This paper describes the design and deployment of a collaborative software tool, designed for and presently in use on the Mars Exploration Rovers (MER) 2003 mission. Two central questions are addressed. Does collaborative content like that created on easels and whiteboards have persistent value? Can groups of people jointly manage collaboratively created content? Based on substantial quantitative and qualitative data collected during mission operations, it remains difficult to conclusively answer the first question while there is some positive support for the second question. The MER mission provides a uniquely rich data set on the use of collaborative tools.

Categories and Subject Descriptors

H.5.3 [**Group and Organization Interfaces**] – Computer-supported cooperative work, Collaborative computing, Evaluation/methodology

General Terms: Design, Human Factors, Measurement

Keywords: Knowledge Management, Collaboration

1. INTRODUCTION

The MERBoard software provides a collaborative workspace for collocated mission scientists for the Mars Exploration Rovers (MER) 2003 mission. MERBoard was initially proposed as a digital whiteboard to support informal scientific collaboration. The advantage of MERBoard over traditional paper media like flip charts and physical whiteboards was that data would be persistent, retrievable, and distributable in a primarily collocated fashion. Although there has been significant work on collaborative knowledge management [4,5,19,22,24], collaborative content creation [1,3,8,10,13,15,18], and electronic whiteboards/large scale displays [6,7,9,11,12,17,20,23], there is relatively little that looks at the intersection of the three, and in particular that examines the problem of knowledge management on large screen displays with an emphasis on real world deployment and use. This paper presents data collected during the deployment of such a system, focusing on knowledge management.

The purpose of the MER mission is to further Mars exploration through the deployment of twin robotic rovers outfitted with a payload of scientific instruments. Launched in June of 2003, the rovers landed in January 2004 and surface operations continued through April 2004 successfully completing the 90 sols (Martian days) of the nominal mission. Both rovers have continued to perform well and are in an extended phase of operations as of the writing of this paper (August 2004). Each rover can perform several hours of activity per sol (e.g. taking photographs or driving) contingent on the availability of limited resources such as battery power. At the beginning of each sol, the rover receives a new set of commands to execute and transmits the data gathered during the previous sol. The mission science and engineering teams then have 18 earth hours to process the data received, analyze the data, and use it to decide what science activities to request for the following day.

The Athena Science Team, consisting of more than 120 scientists, is responsible for analyzing data products on a daily basis and working with mission engineers to plan rover operations. During the daily science assessment period, the team is divided into five Science Theme Groups: Atmospheric Science, Geology, Mineralogy and Geochemistry, Soil and Rock Physical Properties, and Long Term Planning. It is critical for the Science Theme Groups to work collaboratively within and across groups in the Science Assessment Room (Figure 1). Scientists work individually and together to generate hypotheses and devise ways to test those hypotheses. Within each group and across groups they must come to an agreement on a prioritized set of activities to be delivered to the engineering team. The prioritized set of activities represents the consensus of the entire team called the Science Operations Working Group (SOWG) .

CSCW'04, November 6–10, 2004, Chicago, Illinois, USA.

Figure 1: Science Assessment Room (one per rover) containing the majority of the scientists on shift, workstations, projectors, and six MERBoards.

On any given rover and sol there are nearly twenty staffed science positions, including a Science Theme Group lead and documentarian for each theme group, and an SOWG Chair and documentarian. In addition, there are usually between two and eight supporting Science Theme Group members who collaborate with the leads throughout the sol. These scientists are the primary target users of MERBoard.

2. MERBOARD DESIGN PROCESS

At the beginning of the MERBoard project in August of 2000, two team members, an ethnographer and a computer scientist, set out to find gaps in the science process that could be addressed by the incorporation of appropriate technologies. They conducted user research in the form of ethnographic analyses of an early mission simulation in August of 2001. They also conducted open-ended interviews with scientists to better understand their work practices. They saw scientists working on flip charts which were difficult to share with those not in the room and difficult to modify, as ink on paper is by nature. Often these flip charts were mislaid. Screen real estate was another gap in the process. Scientists held up laptops for viewing by groups, making comments like, “if you could see my analysis it would show that…” The MER mission planners also participated in the interviews. One high-ranking planner said he wanted to save “everything.” Simultaneously, the computer scientist explored technologies that could support the science assessment process. Based on the data, they identified the concept of a workspace built around a large plasma screen with a touch-screen overlay, providing an electronic whiteboard to support the informal collaboration among small groups of scientists [23].

Approximately one year later, the project brought in HCI researchers to participate in design, development and deployment of the tool through a series of mission simulations between August 2001 and the final pre-landing release in December of 2003.

3. THE MERBOARD PLATFORM

The goal of the MERBoard project was to allow shoulder-to-shoulder collaboration with a sizable view space such that groups of scientists and engineers can annotate images, do free form drawing, and do strategic planning. The goal was to support these collaborative activities with storage, retrieval and sharing capabilities. Given that the MERBoard is an *ad hoc* tool and that the science users vary greatly in their comfort with technology, the original directive for MERBoard was “Palm Pilot” simplicity: appliance-like usability and reliability.

Each MERBoard consists of “controller” software running on off-the-shelf hardware. The hardware consists of a high-powered PC running Microsoft Windows, a 50 inch plasma display, a touch-sensitive screen overlay, and a custom stand (see Figure 2). There is also a supporting network infrastructure that connects the boards to one another and also to a pair of servers (for centralized file storage and configuration management). There are 18 MERBoards installed in the three-floor mission support area at JPL. Each Science Theme Group has a dedicated MERBoard and there are MERBoards in the engineering areas as well as the room where the SOWG meetings are held.

Figure 2: A MERBoard in use by mission scientists.

The controller software is an application written in Java that occupies the full screen and, for the most part, completely masks the underlying operating system. The platform consists of three “core” applications: a whiteboard, a web browser, an application called “remote” that allows view and control between boards. The software is designed to be extensible. For MER, the boards are running several internal and third party “plug-in” applications. Internally, the project developed a decision-tree tool for the Long-Term Planning Group called “Sol Tree Tool” (see Figure 3). The project also incorporated other mission tools as third party applications such as a data navigator and schedule viewer. There are supporting “global tools” such as screen capture and email, which allow users to perform these functions in any application.

4. MERBOARD USAGE SCENARIOS

In order to place the design decisions and usage data described later in context, the following scenarios are provided to frame MERboard usage. They are based on observational data and are representative of typical science tasks.

> ***Scenario 1:*** Several scientists (three to six) discuss two possible plans for the next sol in front of the MERBoard Sol Tree Tool (see Figure 3). The first option is to stay at the current location, “Stonehenge” and collect more science data. The second option is to drive toward a possibly more interesting target. Each of these options is represented as a node on the decision tree. During the course of the discussion, as other options are suggested scientists encode them in the Sol Tree. The tree also

tracks progress toward mission success metrics for different branches of activity. Often, the Sol Tree created in this context is presented to the larger group in order to facilitate discussion about long-term planning.

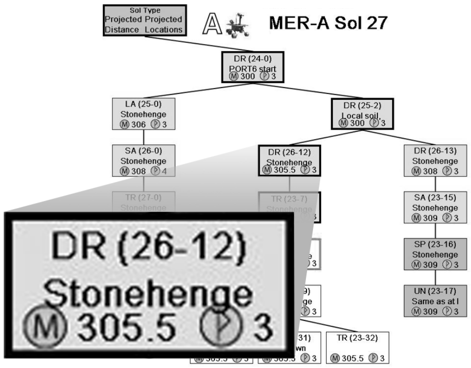

Figure 3: Sol tree representing the Long Term Planning Team's decision tree covering 7 Sols of operations. (Expanded view of single node inset.)

> ***Scenario 2:*** Three Geology Theme Group members are looking at a large image. They draw arrows to targets in the distance they may like to visit in the coming sols. The MERBoard is a convenient option because they are able to email their annotated image to collaborators at remote institutions.

> ***Scenario 3:*** A Soils Theme Group member has done an analysis of a rover traverse and overlaid it onto an image in Adobe Photoshop (see Figure 4). He uploads this image to the MERBoard in order to take advantage of the large display format and discusses it with several colleagues all standing around the board.

5. USER RESEARCH LIMITATIONS

When we set out to design a tool to fill the gaps identified by an ethnographic analysis conducted during the first major mission simulation in 2001, we were faced with many questions regarding the eventual work practices on the MER mission. There was no currently running Mars surface mission to provide data for user research. While the MER mission planners were running simulations three years before the mission, those simulations were very different from the actual mission in scale, which affected all other aspects of the exercises: duration, schedule, location, facilities, number of people, rover hardware, science support tools, and, of course, scientists' own awareness that these were "sandbox" simulations.

The purpose of the mission simulations was to improve the various aspects of mission design, from processes to facilities. We were able to observe and collect data during many of these simulations. Nevertheless, the traditional HCI iterate-and-test model was less effective because the current iteration of MERBoard was always being developed based on data from the previous iteration of the mission design. The net effect of these differences between the simulations and the mission's final design was that our users' tasks and operational roles were not defined in practice until shortly before the landing of the first rover.

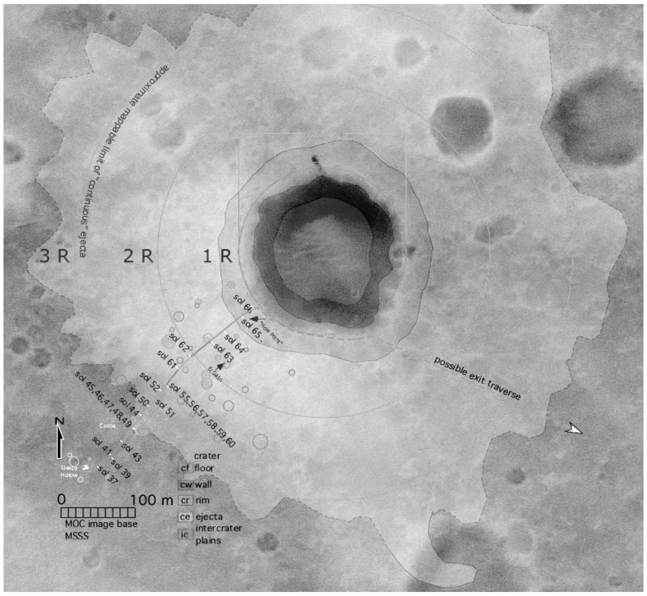

Figure 4: A sample composite image created by a scientist in Photoshop of the type that might be displayed on the MERBoard for group discussions.

From the point of view of designing tools to allow collaborative knowledge management, duration was one of the more challenging limitations of the simulations. The longest simulations were ten days compared to the 90-day nominal duration for each MER rover. Most simulations lasted less than five days. This format did not begin to test the scalability of our knowledge management design. In fact, it did not even allow us to test whether knowledge management tools were needed at all. There was no reasonable way to conduct a traditional user test that would serve as validation because without the user being involved in the creation of the data, she would lack the contextual cues upon which knowledge management systems depend. On this issue and many others, we made design choices based on our understanding of the evolving mission design, but with little actual usage data.

6. RELATED WORK

Design for the MERBoard drew on previous CSCW work in three primary areas: collaborative content creation, collaborative version management, and interaction with whiteboards and large touch-screen displays.

There is a great deal of research on collaborative content creation both synchronous [1,13,10] and asynchronous [8,15,18] that describes various tools in support of collaboration for tasks such as writing or drawing. However, adoption of these tools has historically been relatively low [2,25]. There have been several theories presented that attempt to explain this low adoption rate [14], but most of them associate adoption strongly with the extent to which tools match sometimes widely varying types of collaboration.

Therefore, we made it a priority to design the MERBoard for the specific type of collaboration that we and other system designers anticipated for MER. Since the majority of content creation on the MERBoard was likely to be done "shoulder-to-shoulder," that is, multiple people standing around a single physical board, fine grained coordination around content creation is not explicitly supported in our tool. Instead, we relied on social protocols to dictate this. We also did not anticipate a great deal of explicit structure around the content creation process (i.e. no explicit delineation of stages such as "brainstorming," "planning," "writing," "editing," etc.) [14,15] given the relatively unstructured style of collaboration observed around flipcharts during field tests.

During the mission, both physical space and teaming were arranged in such a way that they took advantage of many known benefits of physical collocation during collaborative work [3]. We expected a mode of interpersonal collaboration similar to that of TeamX [16] in which power was distributed fairly evenly across collaborators, contributors were highly collocated, and there was a strong incentive and little down-side to collaborating (in contrast to some environments that are potentially hostile to collaboration such as that observed by Orlikowski [19]).

In light of the type of asynchronous collaboration that we expected, we identified version control as a key requirement for the MERBoard. Version and configuration management has been studied in several domains, in addition to writing and drawing, such as software development [22, 24]. Many existing versioning schemes present fairly sophisticated models that are robust to pathological situations including branching, parallel editing by multiple authors, dynamic locking of content, and eventual re-merging into a single master copy [25]. However, each added feature brings complexity. Since one of our overarching goals for the MERBoard was walk-up-and-use learnability, we opted to draw primarily upon versioning schemes used in traditional writing and drawing tasks and make many aspects of the process transparent to the user. For example, in contrast to the Shared Books' [15] content locking model which requires explicit check-out and check-in of documents (and is typical of many versioning systems used today), we designed a system whereby version management and branching intent is derived by the system from familiar user actions. For example, if a user performs a "Save As…" the system will automatically create a branch and preserve the version history.

The original inspiration for the MERBoard's large touch-screen form factor was derived from work done on the IBM BlueBoard [23] and work on some promising applications of large-screen technology to traditional tasks [12]. In particular, touch-screens seem well suited to supporting collaboration by groups of physically collocated users [6,9,11,12,20,21], the primary mode of collaboration we were trying to support. In addition, there has been some work on knowledge management for large shared displays [7,17]. The MERBoard shares many design goals with Flatland [17] in that it is intended to have a low threshold for initial use, provide a look-and-feel consistent with informal collaboration, and support a range of usage modes from public to small groups.

The work on MERBoard extends existing work in two primary ways. First by implementing a tool that is novel in its combination of functionality, and second by providing some analysis of its use in an operations environment. We implemented a tool at the intersection of collaborative content creation, collaborative version management and interaction with whiteboards and large touch-screen displays – a design space that is currently sparse. While we drew upon previous work in each area, it remains difficult to make direct comparisons to the MERBoard. For example, looking to previous work for version management design was of limited usefulness because the introduction of the collaborative content creation and large touch-screen components significantly changed many design parameters. Second, we have collected a rich set of both qualitative and quantitative data over the course of the MERBoard's deployment in mission operations, and this paper presents some analysis of that data.

7. PRIMARY RESEARCH QUESTIONS

At the outset, two research questions were posed, the answers to which should inform other groups developing tools for collocated collaboration. First, does collaboratively generated problem-solving content, of the kind that people usually create on whiteboards and flip charts, continue to have value after it is created? Second, can *group* knowledge management actually work, or does it always depend on management by one individual (or a small subset of the group)? We use the term *group knowledge management* to mean the process of jointly creating, modifying, storing, retrieving, and sharing content by groups and individuals. If it is the case that content created on-the-fly during collaborative problem-solving (e.g., on whiteboards and easels) has no persistent value and if true group knowledge management will not work, then no tool or design will change this.

8. KNOWLEDGE MANAGEMENT CONVENTIONS

Single user desktop knowledge management conventions, while not without problems, are well established. The lack of collaborative knowledge management conventions manifests itself as four specific areas relevant to the knowledge management design of MERBoard. These areas, described below, are: document organization, content ownership, document modification and versioning, and data privacy and security.

Desktops have hierarchical file systems with platform specific metaphors like the Mac's Desktop or Windows' "My Computer" that offer frameworks for organizing documents. Applications developed for those platforms conform to standards like default locations for saved files. In collaborative contexts, there is no default save location let alone established metaphors or frameworks for organizing documents.

Content ownership is not problematic in single user, desktop systems, which have one user logged in at a time. A document is owned by the user who was logged in at creation time. Collaborative systems, in contrast, are in use by multiple users at a time and, unless ownership is explicitly declared, it is impossible to establish. Although several designs have been implemented in specific contexts [8,15,18], collaborative systems have yet to adopt a standard for allowing a particular user to flag content as their own.

Desktop users have full authority over the modification and versioning of content on their personal machines. Collaborative users, on the other hand, maintain shared content. For example, two colleagues may create a document together in one session, but in a separate session one user may change the shared content such that the document will appear different when opened next. A colleague's changes may or may not be what the group or group member is looking for. While there are complex, domain specific tools that provide collaborative versioning (e.g. CVS for software development) there is no usable, familiar collaborative versioning convention.

Data privacy in the desktop context is addressed by a password-protected login and, often, by single user machines. In a collaborative context there is an expectation that at least some data is public. The concept of designating public versus private data may be familiar, based on interaction with document management systems, but those systems do not provide a common set of conventions or defaults. Overall, desktop document management is based on the preferences of one individual. In a collaborative space there are multiple users' preferences to consider, and while organizational conventions are necessary it is unclear how they should be structured.

Given the dramatic difference in the needs and usage patterns between the single user and collaborative contexts we did not seek to reuse existing single user conventions. We sought instead to deal with these problems as they arose by consulting previous work on collaborative systems and designing new conventions where appropriate.

9. MERBOARD KNOWLEDGE MANAGEMENT DESIGN

This section describes four primary aspects of knowledge management on the MERBoard.

9.1 Content Ownership Design

MERBoard is intended to be fundamentally "placeless." That is, no matter which board a user walks up to she will be able to orient herself and access her data. In order to accomplish this, the design separates the notion of the data from that of the physical board through Personal Object Models (POMs).

The POM is an iconic representation of an individual, group or archive (usually a thumbnail-sized picture), which is based on the Pcon (personal icon) convention used by IBM's BlueBoard [23]. A primary design goal for MERBoard was to encourage informal use by eliminating the need to login using a traditional username/password scheme. The POM provides the user or group a consistent way to identify themselves to the system in collaborative scenarios where the user's identity is important. These scenarios include: which user or group is this document associated with, who is this email from, and who is requesting to remotely view this MERboard. The user identifies himself to the board with a maximum of three taps by invoking a POM chooser and selecting his likeness (see Figure 5). Users access their group's POM similarly to store and retrieve group content.

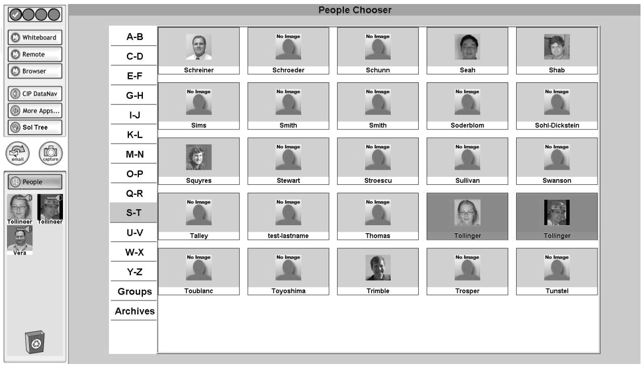

Figure 5: The People Chooser allows users to identify themselves to the MERBoard by selecting their POM.

We decided that associations between people and documents were necessary to limit the amount of data for users to search or browse in order to find a particular document. In our design, a user expresses ownership through an explicit interface-level association between a person and a document. Pressing the "owners…" button in an application opens a sliding, non-modal dialog referred to as a "drawer" that contains all the owners of the document (see Figure 6). There are two primary ways of adding a POM as an owner. The first is to manually drag a POM to the document's "owners drawer." The second is to add owners via the dialogue that results from tapping the "Save As…" button. The "Save As…" dialog displays the owners in a manner consistent with the "owner's drawer" in each application. When a new document is created, its default owner is the archive for the MERBoard it is created on. When a document is closed without an explicit save, the document is automatically saved to the associated POM which, in the case of a document that has not had an owner added, is the board archive. This reduces the chances that users will inadvertently lose work.

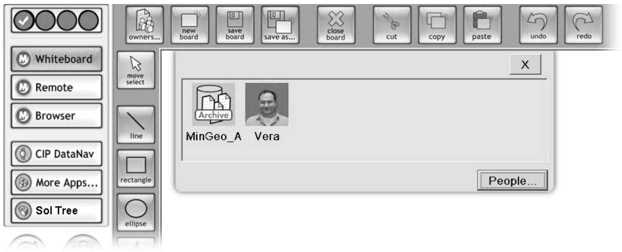

Figure 6: The "owners drawer" used to add owners to a document via dragging a POM into it or selecting the "People…" button.

9.2 Document Organization Design

Collaborative Knowledge Management is handled in part through a centrally stored document repository dubbed "MERspace." MERspace (see Figure 7) is accessed on any board through individual POMs which, when tapped, invoke a view containing that individual or group's subset of the shared repository. Each person or group's MERspace is a one-deep hierarchical representation of files that a user has either saved to her own space or received from another user. In order to give a file to another user, the colleague's POM must be available on the screen. The user must drag the file from their MERspace and drop it on the colleague's POM. The file is then available in the colleague's "inbox" folder.

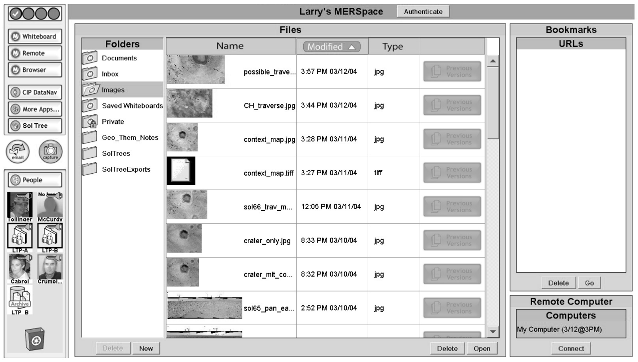

Name	Modified	Type	
possible_trave...	3:57 PM 03/12/04	jpg	Previous Versions
CH_traverse.jpg	3:44 PM 03/12/04	jpg	Previous Versions
context_map.jpg	3:28 PM 03/11/04	jpg	Previous Versions
context_map.tiff	3:27 PM 03/11/04	tiff	Previous Versions
sol66_trav_m...	12:05 PM 03/11/04	jpg	Previous Versions
crater_only.jpg	8:33 PM 03/10/04	jpg	Previous Versions
crater_mit_co...	8:32 PM 03/10/04	jpg	Previous Versions
sol65_pan_ea...	2:52 PM 03/10/04	jpg	Previous Versions

Figure 7: A user's personal MERspace where they can access and organize their files and folders.

There are five default folders in MERspace: documents, images, inbox, saved whiteboards, and private (which requires authentication to view its contents). At save time, users are not given the option to choose where their document will be stored to minimize overhead. Storage location is determined by document type (e.g. whiteboards are saved in the saved whiteboards folder). Users have the option to create and name new folders and designate them as public or private. Previously saved files can be moved into these new folders. The files in MERspace can be sorted by date-modified, name, and type. The default sort is date-modified.

9.3 Document Modification and Versioning Design

Versioning, the questions revolving around who has the right to modify a document, how changes are propagated, how to limit the possibility of one user objecting to or clobbering the changes of another, is a particularly thorny issue inherent in the context of collaborative knowledge management. The science team, our target user group, is composed of senior researchers selected based on research proposals submitted for review. The context is a secure building on the campus of the Jet Propulsion Laboratory that requires passing through several layers of security and using a key card to open a locked door. Therefore, by default the right to modify a document is open to all, even those who are not owners of a document, based on an understanding of the target user group and context. Users have the option to mark data as private which is covered in the next section.

Automatic change propagation is likely to confuse users of desktop computers who do not expect their file spaces to look different unless they themselves have made changes. In order to address this concern, collaborative versioning systems typically adopt a "branch and merge" strategy to handle asynchronous changes made by peers [4,15,22,24]. The fundamental premise is that a users' space should appear just as it was left the last time it was used. However, given the fast paced analysis and decision making required by the daily schedule, we decided the danger of providing out of date data was worse than any possible confusion caused by change propagation. The current design has one central document that all owners have a link to in their MERspaces. If user A modifies the document, user B's space is updated and when user B opens the document in question, he or she sees the new version of the document (see Figure 8). If a user wants to explicitly separate their version they can do a "save as" which maintains the version history but the file will no longer be changed when the original document is updated.

The MERBoard architecture has a versioning scheme designed to support the user should he or she wish to go back to an earlier version of a document. Each time a user or the system saves a document, a new version is created. However, given our hypothesis that access to previous versions would be infrequent, only the latest version of a document is visible in the user's MERspace. There is a "previous versions" button associated with documents that have previous versions, which brings up a dialog listing the previous versions in a format consistent with MERspace.

The versioning design was also intended to reduce the number of files created by offering access to older versions without requiring the user to "Save As…".This design provided previous versions beneath fewer files in order to limit the amount of clutter in users' MERspaces. Reducing clutter makes sorting and organizing information easier.

9.4 Data Privacy and Security Design

Some of our design decisions were driven more by mission mandate than by perceived user needs. For example, it is possible to access, display, store, and even share ITAR (International Trade in Arms Restriction) classified information on the MERBoard. Therefore, it was necessary to provide the ability to designate folders as private spaces that are password protected. The user must authenticate before accessing private folder content and their authentication expires after 10 minutes of inactivity. A folder labeled "Private" (with a distinguishing locked icon) exists as a default folder in every users' MERspace to demonstrate this functionality. However, the burden of keeping information private and protected necessarily falls largely to the user in that users must explicitly move data into the private folder as the mission used no software level flag to identify ITAR data.

Besides the mission mandate for secure storage spaces on the MERBoard, the scientists had concerns about security. One of the lead scientists realized early on that email functionality requires the user to identify who an email is "from" and became worried that others would send email "from" him. The authentication method described above was required to add a user's POM to the "from" field in order to send an email. We worked to provide a consistent design to address the disparate issues, by making authentication board-wide, allowing access to all protected areas and actions with a single password until authentication expires.

The design team felt that such security measures, although necessary, were somewhat contrary to the collaborative design of the MERBoard and that a 50 inch display was not likely to be used for truly private data. As a result, we were careful to isolate usage modes in which sensitive information could be accessed, and keep interaction around security features to a minimum.

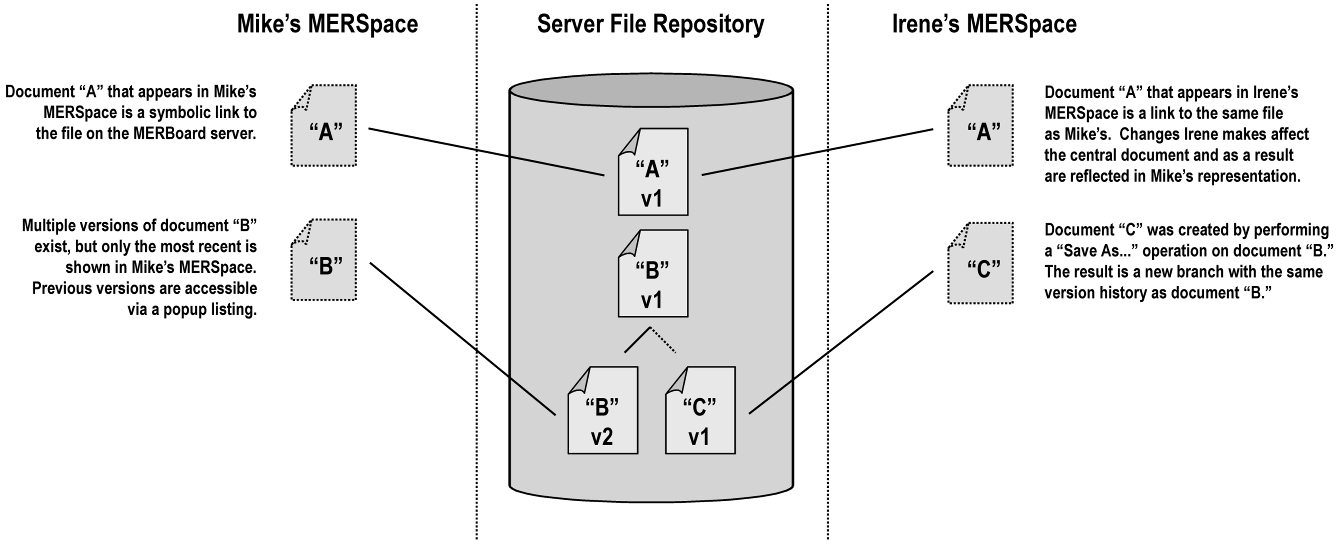

Document "A" that appears in Mike's MERSpace is a symbolic link to the file on the MERBoard server.

Multiple versions of document "B" exist, but only the most recent is shown in Mike's MERSpace. Previous versions are accessible via a popup listing.

Document "A" that appears in Irene's MERSpace is a link to the same file as Mike's. Changes Irene makes affect the central document and as a result are reflected in Mike's representation.

Document "C" was created by performing a "Save As..." operation on document "B." The result is a new branch with the same version history as document "B."

Figure 8: MERBoard versioning scheme

10. DATA COLLECTION

Through the duration of nominal surface operations on Mars for both rovers, data was collected to assess MERBoard's use. Log files captured all user interaction and all user created files were saved. In addition, 60 sols of audio/video of users, synchronized with MERBoard screen capture were collected. In situ activity coding capturing the distribution of tool-use by scientists, ethnographic observation and structured interviews were collected following a counter-balanced sampling design for 60 sols across both missions.

The results reported here cover use of the MERBoard through mid-March 2004 on both rover missions. Because the Spirit Rover landed 20 days before the Opportunity Rover, the data analysis presented represents MERBoard use for 68 sols on Spirit and 48 sols on Opportunity. Overall, the use of MERBoard has been fairly low. Initial analysis of in situ activity coding indicates that users physically interacted with the boards 2% of the time they were on shift. Nevertheless, our data indicate that 22% of users had documents associated with them or otherwise identified themselves to the system in the course of interaction, and created a total of 799 unique files.

A systematic analysis of log file data was used to identify usage patterns relevant to our high level research questions. After presenting this data, possible interpretations are considered based on structured interviews and ethnographic observation gathered during the sixty sols (over 500 hours) of in situ data collection.

11. MERBOARD KNOWLEDGE MANAGEMENT RESULTS

11.1 Content Ownership Results

In our analysis of the POM-based ownership model, we first sought to gauge the level of use. Out of 217 users and groups, 22% of users have associated themselves with at least one document. We found that documents stored in MERspaces had on average 2.2 users associated with them (max=13). Of all documents, 25% had multiple owners. Of documents with multiple owners, the average number of owners was 5.7 (as compared to the overall average of 2.2). Of documents opened more than once, 25% were accessed from different MERSpaces.

Our observations showed that people infrequently accessed another user's personal space (even without security restrictions) and would tend to stick to their own personal space and their group's space. This seems to support the hypothesis that these documents were being used by different people who opened the same document but from different locations.

Since our design associates one "owner" (the Archive) by default with every document created, it is important to note the distinction between these default owners and owners expressly added by users. While the percentages above include both types of owner (those added by the system and those added by the user), we found that 67% of documents saved included at least one owner that was explicitly added by a user.

11.2 Document Organization Results

One of our key questions about MERspace usage was to what extent people used the ability to create new folders and move documents to organize their spaces beyond the default organization. Current data indicates that only 5% of users have created new folders in their MERspace. Only 4 documents have been moved from their default save location to another folder in their MERspace. There were only 3 instances of documents being delivered to others by dragging from the file view onto another user's POM (note that it is also possible to share a document by adding a user as an "owner"), and only 2 instances of users changing the default sort, both times the sort was changed from the default date ordering to name ordering.

The above log data indicates that it was relatively uncommon for users to change the way the information was organized in their MERspace. Based on structured interview data, users generally seemed satisfied with the default organization. The following interview excerpt is a representative example.

> **Interviewer:** How often do you access files from previous sols versus the current one?
> **User:** Quite a lot I mean on any given – since I've been here for this mission – every day I've gone back to previous sols to look for images that I recall that we had to compare with what I have now?
> **Interviewer:** You mentioned you have a folder on the MERBoard. How do you prefer to organize your files personally? Once you have them on the MERBoard or have them in your personal-
> **User:** Can't say there's been a lot of organization yet I just sort of dump stuff in there.
> **Interviewer:** And does that sort of date ordering work for you now?
> **User:** Yeah, yeah pretty much. Eventually it won't work but there hasn't been enough happening yet.

The incidence of people changing the way their MERspace is organized remained low perhaps because the average number of documents in group and personal spaces was 9.01 (median=4, max=117). While the default organizational scheme scaled to support this usage, it is unclear what would have happened if MERBoard usage had been higher.

11.3 Document Modification and Versioning Results

We have some evidence that versioning was understood and used in some cases. Of the documents created on the MERBoard, 22% have previous versions associated with them. Since a version is created every time a document is saved, this percentage is not necessarily indicative of a user's intent to create an explicit version. However, a few pieces of data suggest that the versioning functionality was used to recover documents.

We found that out of 471 instances of a user opening a saved document, there were 20 instances of users invoking the "Previous Versions" dialogue and opening a document that was not the most recent version (4% of file opens). There were 19 instances of users invoking the "Previous Versions" dialogue without actually opening a previous version which seems to imply that users could not find a version that matched their expectations, were exploring the interface, or pressed the button by accident.

Since a new version of a document is created each time it is saved, it is also useful to look at saving behavior in the context of versioning. There is a distinction between the behavior of "Save" and "Save As…" in that "Save As…" creates an explicit fork in the version tree (although previous versions of the document are preserved). In our analysis of the log data early in the mission (sols 1-19 of Spirit), we found that users chose "Save" 39% of the time and "Save As…" 61% of the time. However, analysis of data later in the mission (sols 1-68 of Spirit and 1-48 of Opportunity) revealed a reversal in favor of "Save," with users choosing that option 57% of the time as opposed to 43% "Save As…". This reversal may be explained in a number of ways, but one interpretation is that an increased understanding of the MERBoard versioning scheme resulted in a decreased desire to create explicit version forks with "Save As…". Instead, users may have been more willing to rely on their ability to recover previous versions of the same file. It is also possible that users realized they did not tend to go back to previous versions and simply stopped using the higher overhead "Save As…" operation. There is presently insufficient data to come to a conclusion, but future interview data may shed light on this.

Another important aspect of MERBoard save behavior as compared to desktop save behavior is the frequency of save for a given session. Through observation, we have noticed that users save relatively infrequently. While time between saves varies dramatically by task and by individual, the scientists often worked for hours without saving. Determining exactly how often users save is difficult because the end of a work session cannot be determined from the log data and documents tend to be left open. One explanation for the infrequency of saving is the robust implementation of content restoration. All open files are constantly backed up to prevent loss of work. When the MERBoard crashes, which itself is relatively infrequent, all currently open documents are restored upon board restart. Other explanations for infrequent saving are the observed high incidence of interruptions which may have caused users to forget to save or the collaborative nature of the context in which it was perhaps unclear whose responsibility it was to save.

11.4 Data Privacy and Security Results

Since many of the privacy features added to the MERBoard were mandated by the mission, we were interested to see the level of use associated with those features. In general, private folders (arguably the most rigorous optional privacy feature on the MERBoard) received very little use. As mentioned above, users were fairly unlikely to modify their MERspace from the default configuration. Only 11 users created folders and of those only 2 users created private folders. This is consistent with our expectation that users would not be concerned about privacy to the extent that they feel the need to lock out others in the group or to secure ITAR data in practice. There was almost no use of the security features that had been required by the mission security rules.

Early in the mission, people would access other users' personal spaces, but we observed this happening only when they had been present for the creation of the original document but did not have or did not realize they had a link to the document in their MERSpace. As the science team became more practiced, we did not observe any more instances of accessing files through other people's personal spaces. While people did not use the explicit security features, over time, they seemed to avoid accessing others' personal spaces.

We can further examine the privacy data by looking at some more subtle patterns. We found that of the 67% of documents that had an explicitly added owner, in 85% of those cases the user had also taken the extra action of removing the archive from the list of owners. The archive is a public file space associated with a particular board intended to serve as a repository for files not associated with a user or group. One explanation for the explicit removal of the archive is a lack of comfort with the archive's public nature. Another explanation may be a lack of understanding of the archive's purpose and location.

12. FINDINGS ON PRIMARY RESEARCH QUESTIONS

The results specific to particular design elements described above go part way toward addressing the two primary research questions used to frame this work. First, is the information generated in an informal whiteboard-like style useful after creation time? Second, is it possible to maintain group knowledge management that is not guided by an individual or small subset?

We have already presented some data relevant to the first question on utility of information beyond creation time. The fact that users are associating themselves with content (67% percent of all documents have at least one explicitly added user) and the fact that 25% of all documents were saved at least twice. These data suggest that users believed that content would be valuable to them at a later date.

We can further address this in terms of the frequency with which people return to their content and the period when that content is useful. 13% of documents created were opened at least once. However, based on observation, this measure may be deceptive in that many documents are never closed and are left open between one day and the next. Even when there is no active editing in progress users leave documents open in a sort of display mode or switch applications leaving a document open so they can quickly access it. On average, each file, treating all versions of a file as one, is opened 2.0 times. The average is high because these 13% of documents are opened multiple times. This, in conjunction with observational data, seems to imply that some files created on MERBoard are useful after their creation.

The lifespan of a document, the period between creation and last save, varies dramatically from less than a minute to 52 days. The average document lifespan is 8 days and the median is 1.3 days. These numbers take into account the high likelihood of users performing a "Save As…" rather than a "Save." We considered document A and document B (created from document A via "Save As…") as a single document. The data implies that most MERBoard documents' usefulness is limited to a relatively short period of time.

The second question, whether group knowledge management is possible, can be addressed in terms of association between groups and documents. Of all documents, 48% have a group POM as an owner. This indicates an inclination towards group knowledge management on the part of users. In general, documents do tend to be opened from group spaces. This answer is qualitative based on scanning a list of documents opened as it is difficult to get a true quantitative answer when any given document is opened an average of 2.0 times. Further, as noted above, 25% of documents opened more than once were opened from different MERSpaces. Of documents opened more than once, 64% were opened from both a personal and a group space.

At the beginning of the mission we observed team leads define protocol, as they have done with other systems, stating where the latest or 'golden' copy of a file would be stored. However, as the users learned the MERBoard system, we saw them put the document in all the spaces of the relevant authors and their group space. Thus, there is some indication that the MERBoard contributed to a more group oriented knowledge management process as opposed to a process driven by the group leader or small subset of the group.

13. CONCLUSION

Much like individual work on desktop computers, it appears that collaborative work products in the MERBoard context are saved, retrieved, and shared. Some content is likely to become "living documents," that users continue to work on from day to day. For example, Sol Trees created by the Long-Term Planning theme group generally had specific people as owners (rather than the archive), multiple owners and multiple versions. These are emergent work practices that were not foreseen by the designers. In trying to frame the current data, we have become aware that there is little in the way of a baseline for content lifespan and versioning on regular desktop computers. Furthermore, as the mission progressed it became increasingly clear that the boards were not being used in the way the designers anticipated. The primary intent was to create workspaces where mission scientists could collaboratively create and annotate images and other content in the way they would otherwise do with whiteboards and easels. The kind of content people create on whiteboards and easels is, almost by definition, not of permanent value. If it does have longer-term value, people tend to transcribe the content into more permanent and accessible format. Up to this point in the mission, there has been less of this kind of use and more of either just looking at images and then storing them or creating and evolving long-term plans using the Sol-Tree Tool. The first central question of this paper, whether content of the sort created collaboratively on whiteboards and easels has longer-lasting value is therefore difficult to answer from the current context. The second major question, whether groups can successfully manage collaboratively created content seems to be getting positive support thus far.

The MER mission provides a uniquely rich opportunity to study the role of collaborative tools in a process that demands a great deal of frequent and active collaboration. MERBoard and the other tools designed to support the mission have been successful in some ways and have failed in others. This paper presents an early analysis of a very small portion of the data being collected. Nevertheless, there have already been important insights. Although the overall use of the MERBoards has been relatively low, they are being used to create, store and retrieve collaboratively created content. The literature on collaborative content creation, knowledge management, and whiteboard/large-scale display use is populated with analyses based on researchers using the tools themselves or evaluating them with small test groups. Examples of large-scale deployments, designed using current methods by HCI researchers, and that can then be analyzed in detail are hard to come by. This deployment is providing many useful lessons along with a few of instances of positive reinforcement.

14. ACKNOWLEDGEMENTS

We would like to thank the MER science team for the time spent participating in user tests and offering feedback before the mission as well as their patience with our data collection personnel and equipment during operations. We thank Roxana Wales, the original ethnographer who studied the early work practice and mission design for her contribution to both the design of MERBoard and the research directions for data analysis. We also thank Pamela Hinds for her valuable comments that guided multiple drafts of this paper.

This work is funded by the NASA IS/HCC program.

15. REFERENCES

[1] Baecker, R.M., Nastos, D., Posner, I.R., and Mawby, K.L. (1993) The User-centred Iterative Design Of Collaborative Writing Software. In Proceedings, InterCHI '93, p. 399-405

[2] Beck, E., Bellotti, V. (1993). Informed Opportunism as Strategy: Supporting Coordination in Distributed Collaborative Writing. *European Conference on Computer Supported Cooperative Work*, 241-256.

[3] Covi et al, L., Olson, J., and Rocco, E. A room of your own: What do we learn about support of teamwork from assessing teams in dedicated project rooms? (1998) In Cooperative Buildings, N. Streitz, S. Konomi, and H. Burkhardt, Eds. Springer-Verlag, Amsterdam, The Netherlands. p.53–65

[4] Dix, A., Rodden, T., Sommerville, I. (1997). Modeling Versions in Collaborative Work. *IEE Proceedings - Software Engineering,* 144(4) pp. 195-205.

[5] Dourish, P., Edwards W. K., LaMarca A., Salisbury M. (1999) Presto: An Experimental Architecture for Fluid Interactive Document Space. ACM Transactions on Computer-Human Interaction, 6(2).

[6] Elrod S., Bruce R., Gold R. Goldberg D., Halasz F., Janssen W., Lee D., McCall K., Pedersen E., Pier K., Tang J., Welch B. (1992) Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration. Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems, p.599-607

[7] Fass A., Forlizzi J., Pausch R. (2002) MessyDesk and MessyBoard: two designs inspired by the goal of improving human memory. Proceedings of DIS'02: Designing Interactive Systems: Processes, Practices, Methods, & Techniques, p.303-311

[8] Fish, R., Kraut, R., and Leland, M. (1988) Quilt: a collaborative tool for cooperative writing. In Proceedings of Conference on Office Information Systems.

[9] Ganoe C. (2002) Supporting the collaborative meeting place ACM CHI 2002 Conference on Human Factors in Computing Systems, doctoral consortium, v.2 p.546-547

[10] Greenberg, S. and Marwood, D. (1994). Real time groupware as a distributed system: Concurrency control and its effect on the interface. *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work*, pp. 207-217.

[11] Guimbretiere F., Stone M., Winograd T. (2001) Fluid interaction with high-resolution wall-size displays Papers: Off the wall Proceedings of the ACM Symposium on User Interface Software and Technology, p.21-30

[12] Huang E. M., Mynatt E. D. (2003) Semi-public displays for small, co-located groups. Proceedings of ACM CHI 2003 Conference on Human Factors in Computing Systems, v.1 p.49-56

[13] Karsenty, A., Tronche, C. and Beaudouin-Lafon, M. (1993) GroupDesign: shared editing in a heterogeneous environment. USENIX Journal of Computing Systems, 6(2) p.167-195

[14] Koch, M. (1994). Design Issues and Model for a Distributed Multi-User Editor. *Special Issue on Computer-Supported Collaborative Writing,* Volume *3, p* 359 – 378.

[15] Lewis, B., Hodges, J. (1988). Shared Books: Collaborative Publication Management for an Office Information System. *Proceedings of the ACM Conference on Office Automation Systems*, p. 197-204.

[16] Mark, G. (2002). Extreme collaboration. *Communications of the ACM. Vol. 45(6), p.89-93*.

[17] Mynatt E. D., Edwards W. K, LaMarca A., Igarashi T. (1999) Flatland: New Dimensions in Office Whiteboards. Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems, v.1 p.346-353

[18] Neuwirth, C. M., Kaufer, D. S., Chandhok, R., & Morris, J. H. (1990). Issues in the design of computer-support for co-authoring and commenting. *Proceeding of Conference on Computer-Supported Cooperative Work*, p.183-195.

[19] Orlikowski, W. (1992). Organizational Issues in Groupware Implementation. *Proceeding of Conference on Computer-Supported Cooperative Work,* p.362-369.

[20] Pedersen E. R., McCall K., Moran T. P.,. Halasz F. G. (1993) Tivoli: An Electronic Whiteboard for Informal Workgroup Meetings Proceedings of ACM INTERCHI'93 Conference on Human Factors in Computing Systems, p.391-398

[21] Tang, J. and Minneman, S. (1991) Video Draw: A Video Interface for Collaborative Drawing, Proceedings of CHI, p. 315-322

[22] Rochkind, M.J. (1975) The Source Code Control System. IEEE Trans. on Software Engineering. SE-1(4): p. 255-265

[23] Russell D., Trimble J., Wales R. (2002) Two Paths From the Same Place. IBM Make It Easy Conference

[24] Sachweh, S. and W. Schäfer (1995). Version management for tightly integrated software engineering environments. in Proc. 7th Int. Conf. on Software Engineering Environments. IEEE Computer Society Press. p. 21-31

[25] Tatar, D., Foster, G. & Bobrow, D. (1991). Design for Conversation: lessons from Cognoter. *International Journal of Man-Machine Studies*, 34 p. 185-209.