

Exploring Multimedia Histories of Casual Interactions

Charlotte Tang, Gregor McEwan and Saul Greenberg

Department of Computer Science
University of Calgary, Calgary, AB, T2N 1N4
[tangsh *or* mcewan *or* saul]@cpsc.ucalgary.ca

Abstract

Many groupware systems now allow people to converse and casually interact through their computers in quite rich ways—through text, images, video, artifact sharing and so on. If these interactions are logged, we can offer these multimedia histories to a person in a manner that makes them easy to review. This is potentially beneficial for group members wishing to find and reflect on their past interactions, and for researchers investigating the nuances of online communities. Yet because we have little knowledge of what people would actually do with these histories, designing an effective history review system is difficult. Consequently, we conducted a user study, where people explored real data from an online community. Our study identified a set of tasks that people would do if they could review these histories of casual interaction. It also produced a list of parameters pertinent to how we could visualize these historical records in a tool. With the increasing popularity of computer-mediated casual interaction tools, this study provides an important guide for developing tools to visualize and analyze past multimedia conversations.

Keywords: *Conversational histories, multimedia, casual interactions, visualizations.*

1 INTRODUCTION

Conversations over computers are now common. Using bulletin boards, email, and/or instant messaging, people communicate both asynchronously and in real time through text and file exchange [11]. Using video-based media spaces, people see and hear each other through audio and video links [5]. Using MUDs, MOOs, collaborative virtual environments, and other multi-media gathering places (including our own Notification Collage [7] illustrated in Figure 1), whole communities interact with one another in public venues, where they stay aware of what others are doing and use that awareness to capitalize on opportunities for interaction.

Most research has gone into the groupware technology behind conversation support. In essence, these technologies focus on the present: how people encounter one another, and how they find, receive and respond to messages. In contrast, our interest is on the past, where conversations are captured as a history and then offered to a person in a manner that makes them easy to review. These histories are valuable for several reasons.

- *Group member’s perspective.* The person can browse and review one’s own interaction history to remind oneself of what has happened, or to retrieve critical details. The person can also review missed conversations that are relevant to them, and new members can familiarize themselves with the project work and cultures of the workgroup.

- *Researcher’s perspective.* The researcher can analyze group activity within the history of casual interaction. This is crucial from both a sociology perspective (to understand how conversation and casual interaction are affected by new communication media), and from a usability perspective (to fine-tune the groupware offerings and its interface).

While there has been some work on conversational histories, most concentrate on either formal meeting capture (e.g., video/audio capture tools [2]), or asynchronous text exchanges (e.g., such as threaded chat [14,16]). These approaches, summarized in Section 6, do not cover the new genre of groupware that is now emerging i.e., groupware systems that support both awareness and long-term casual interactions in a rich multimedia setting.

Consequently, we set ourselves the research goal of capturing and presenting histories of multimedia-based casual interactions that people could review. To achieve this goal, we pursued three sub-goals:

1. Elicit tasks that people would do if they could review and analyze histories of casual interactions.
2. Elicit useful representations that people would create to help them visualize and manipulate these histories.
3. Based on these tasks and visualizations, design and evaluate a tool that let people review these histories effectively.

We originally began with sub-goal 3, where we built the VisStreams tool for visualizing multimedia conversations [15]. VisStreams is briefly described in Appendix 1, and is illustrated in Figure 5. While the VisStreams prototype showed good potential, we felt that we needed a more grounded user-centered approach to help us understand our user requirements. Consequently, we reconsidered all sub-goals within the context of real data collected from the Notification Collage (NC), a multimedia-based groupware system supporting casual interaction that has been in daily use for several years [7].

This paper reports our work to date. To set the scene, we first provide background into casual interaction and the Notification Collage. We then delve into the primary topic of this paper: a study of sophisticated NC users who analyzed paper-based records of real casual interactions to see what they would want to do with these histories (sub-goal 1), and how they would visualize these records to help them perform their tasks (sub-goal 2). We then discuss how these results can inform the design of a tool that presents multimedia histories of casual interactions (sub-goal 3). We close with a brief description of related work.

2 BACKGROUND

2.1 Supporting Casual Interaction

A wealth of research into *casual interaction*—the spontaneous and one-person initiated meetings that occur over the course of the day—has repeatedly shown that it is a vital component of effective collaboration [8,17]. The glue behind these interactions is *awareness*, where people track and maintain a general sense of who is around and what others are up to [8]. This is easy when people are co-located: they are aware of many visual cues, such as noticing a closed door or that others are engaged in a phone

call, and people use these cues to better identify both opportune and appropriate times to initiate conversations.

Today's workplaces, however, often contain teams where members are separated by geographic distances, severely curtailing opportunities for casual interactions. This is a problem, because studies have shown that informal communication constitutes a crucial part (31%) of office activity [17]. Removal of such interaction significantly decreases effective collaboration [8].

To mitigate this deficiency, many groupware systems have evolved. First, *instant messaging* let people casually interact and converse through computers [11] e.g. ICQ, MSN and Yahoo Messenger. For awareness, they estimate and display another person's presence as measured by keyboard activity. People communicate mainly by typing. While instant messaging tools offer only an impoverished sense of awareness and a low bandwidth communication channel, they are extremely successful because they make casual interaction between distant collaborators possible.

Next, *media spaces* offer distant collaborators a much richer sense of awareness and communication. They do this by capturing contextual information visible in the everyday

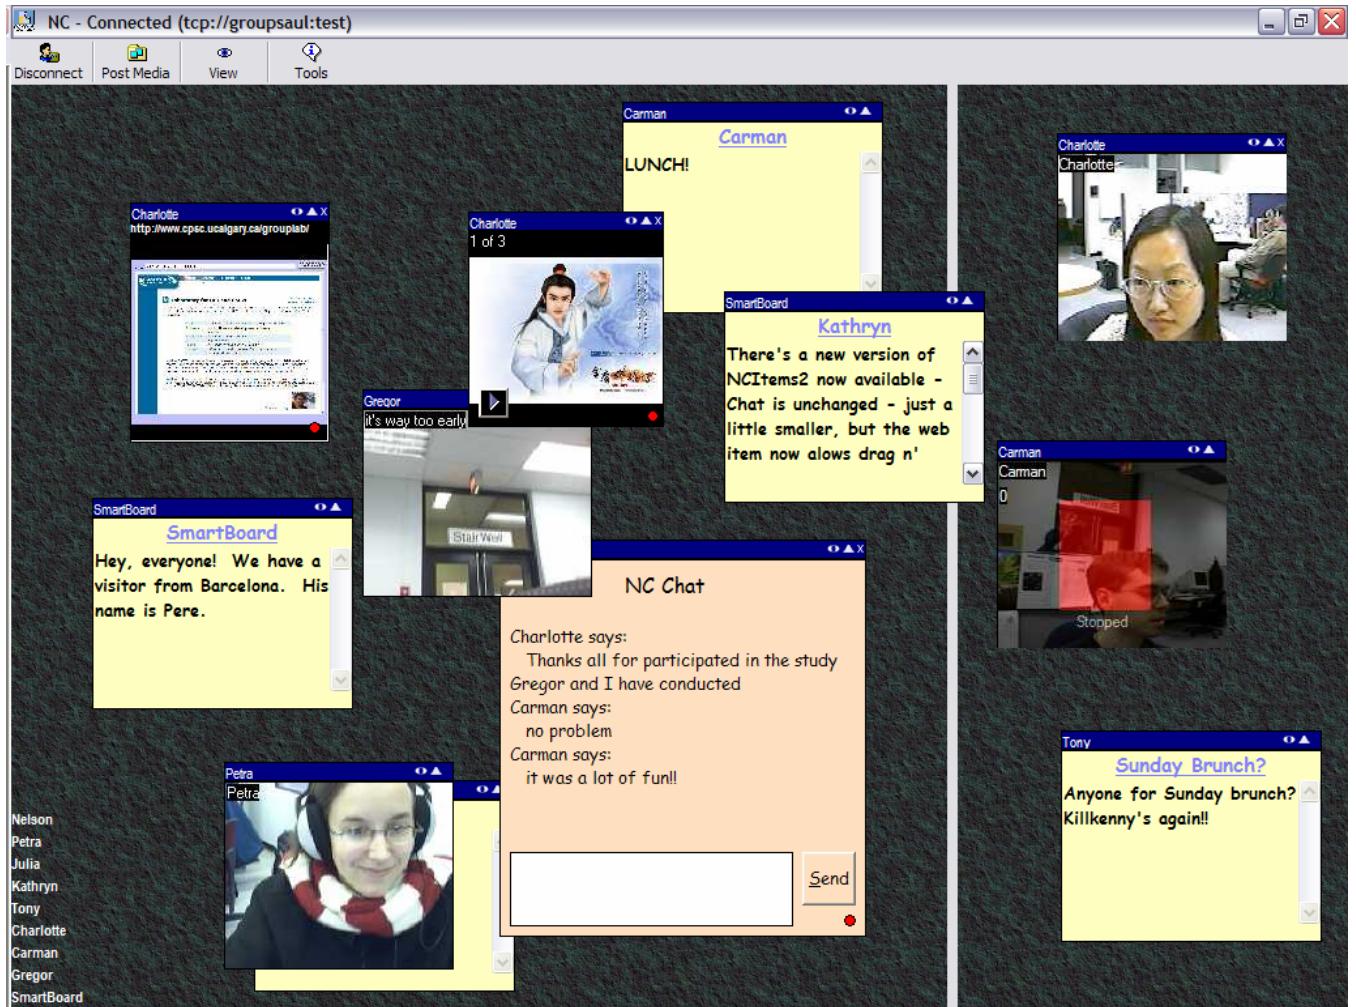


Figure 1. A typical Notification Collage screen taken at time of writing this paper

world as a multimedia stream, and by offering rich communication channels. Early media spaces typically created an always-on video and audio link between distant common areas and/or offices [17].

While instant messaging and media spaces generally support one-to-one conversations, *MUDs*, *MOOs* and collaborative virtual environments create public places for interactions. People see who is present in a virtual space, and can engage in public conversation with anyone there. Eavesdropping and joining into on-going conversations is commonplace. Depending on the system, conversations may be textual or may use real time video and audio. It is these public places that are our primary interest.

2.2 The Notification Collage

The Notification Collage (NC) [7], pictured in Figure 1, is a groupware system that combines features of instant messengers, media spaces, and MUDS/MOOs. Its basic goal is to provide a modest-sized group of intimate collaborators with a public electronic meeting place. Ideally, the NC emulates characteristics of how people work together in open rooms such as team rooms and research laboratories.

Through the NC, people create and broadcast *media elements*. At the time of writing this paper, media elements include text notes, slide shows, web pages, video snapshots, chat tools, and so on. NC then arranges these media elements as a publicly viewable collage, visible on people's desktop computers and on large wall-sized displays situated in a public area (Figure 1).

The media elements act as rich information sources. They provide the group with awareness not only of each other's interpersonal state, but of interesting artifacts. People also communicate through these media elements. For example, Figure 1 reveals (through the video snapshot) which members of this group are present or absent as well as what they are doing, and some photos and web links that people found interesting. It also reveals several announcements that people can read asynchronously at their leisure, and an on-going real time conversation in a chat box. In practice, NC has proven effective at providing opportunities for casual interaction by attracting people's attention and thus arousing their interest [7]. People notice what is going on, they communicate through waves or gestures through the video snapshot and by writing notes, they post artifacts that may be of interest to others, and they use it to signal events.

3 THE STUDY

3.1 Introduction and motivation

When we began this project, we articulated several distinguishing characteristics that must be considered in a conversational capture tool. From these, we developed the VisStreams system—summarized in Appendix 1—that would visualize these characteristics in what we thought would be an effective manner [15]. The problem was that the more we worked on VisStreams, the more we realized that the range of possible features we could include within

it was boundless. We also realized that while the visualizations presented within VisStreams let people pursue some tasks, it hindered the way they could do other tasks. For example, while we could dynamically filter the view, it was very difficult to view conversational units.

Hence, we decided that we needed a more grounded approach to help us uncover user-centered requirements for a visualization tool displaying multimedia histories of casual interactions. To do this, we conducted a study to answer two questions that correspond to sub-goals 1 and 2 of our research.

- 1) What questions would people ask of a history of multimedia data of casual interactions?
- 2) If the raw data were made available to them, how would people re-arrange it to answer these questions?

To answer these questions, we gathered sophisticated NC users as study participants, and had them analyze paper-based records of real raw data captured from the Notification Collage. The study roughly fell into two phases matching these questions: a brainstorming phase where we elicited tasks that they would do over these records, and a data re-organization activity centered on how they would visualize these records to help them perform their tasks.

3.2 Participants

Study participants were 13 computer science students (12 graduate / 1 undergraduate), all of whom were doing research in Human Computer Interaction. All subjects had personal experiences using the Notification Collage, and were members of the NC community whose data we had captured. This selection was deliberate, as we wanted 'expert' subjects who could respond from a group member's perspective (i.e., how one would personally use a visualization tool to review interactions of their own community), and from a research perspective (i.e., how one would use this tool as a scientist trying to understand these interactions). Because participants were members of the NC community whose data we had captured, we found them highly motivated. They wanted to discover things about their community as they analyzed the data. Because subjects had used the Notification Collage for real purposes, they were all familiar with its basic concepts.

3.3 Method and Materials

Each study session (about 1 to 1.5 hours long) involved a single participant working through four stages.

Stage 1. Pre-test questionnaire. Participants completed a questionnaire asking about their familiarity with the Notification Collage, VisStreams, CSCW and groupware, and data analysis techniques.

Stage 2. Brainstorming tasks. To answer question 1, we asked the participants to brainstorm tasks they would perform with a tool that let them review multimedia histories of casual interactions. As they brainstormed ideas, we gave only positive feedback to encourage idea

generation. All ideas were noted and used to guide the third stage.

Stage 3. Data manipulation and visualization. Participants were provided with a booklet containing five days of raw log data captured from the NC (a page from this booklet is shown in Figure 3). Raw events were presented as time-ordered rows annotated by the time and date of the event. The event itself was printed on a detachable PostIt™ Note ~3.5cm x 5cm in size.

Raw events appearing on these PostIts included: a person's connection/disconnection on the NC, a media item's appearance/disappearance, and the detailed contents of media items as they changed e.g., text posted to sticky notes, pictures posted to slideshows, and video frames appearing in video snapshots. As is visible in Figure 3, raw events displayed the data in human-readable form i.e., image data were presented as images (the 5th and 7th row), sticky notes as its current text contents (1st and 2nd row) and other events as explanations.

As seen in Figure 4a, we also provided participants with large tables, several whiteboards, whiteboard markers, blank PostIt notes, pens and masking tape.

To answer question 3, we then asked participants to perform the tasks they had brainstormed in Stage 2. In particular, we asked them to create visualizations by restructuring these raw materials in any way that made sense of the data. Because events were on detachable PostIts, participants could move them around freely on the whiteboard, and annotate them as needed (Figure 4b gives an example of what they did). While we also had a list of tasks prepared *a priori* to offer participants if they were stuck, we only had to use these twice. These data visualizations were recorded in detail by both investigators as well as a video tape.

Stage 4. Post-test questionnaire. A short questionnaire gave participants opportunity to comment on the limitations of both the data representation and provided materials. The participants were able to express where the materials prevented or hindered them in their representations.

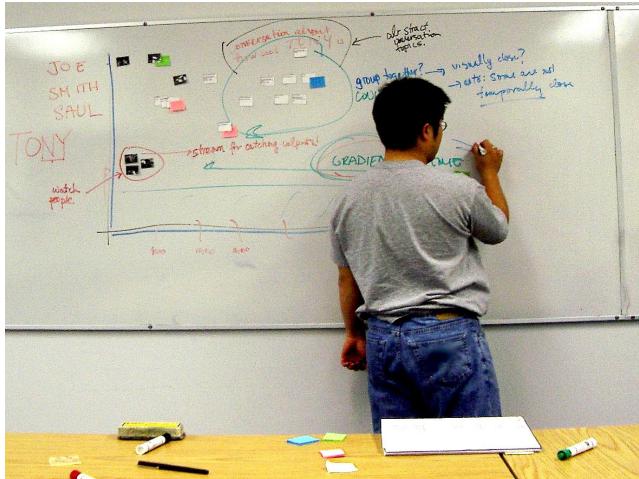


Figure 4. (a) A participant restructuring the data, and

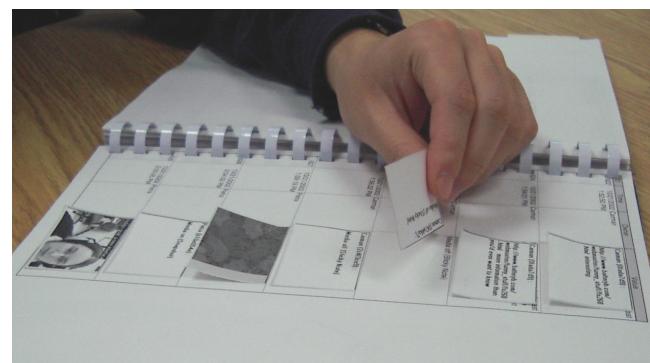


Figure 3. A participant manipulating the history log

4 RESULTS

The next sub-sections describe the results from each of the stages of the trials.

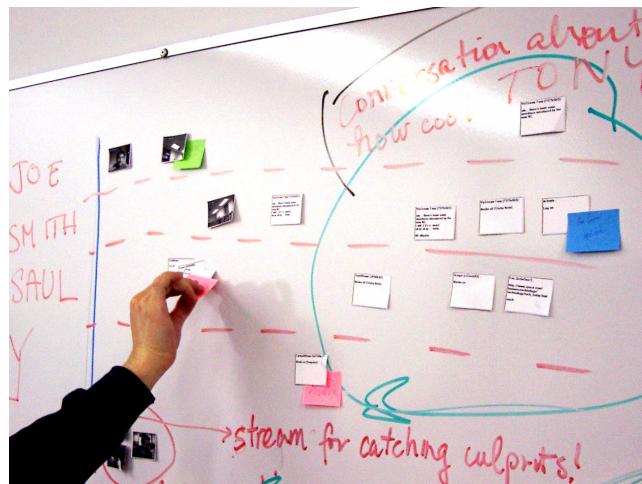
Stage 1: Pre-Test Questionnaire

The questionnaire confirmed that our participants fit the demographics we wanted. That is, they were very familiar with the Notification Collage, and they frequently used other groupware (typically an Instant Messaging system). The majority had some experience with data analysis, although this expertise varied considerably across both depth and domain. While all had some prior exposure to our VisStreams prototype; none had extensive experience with it or had used it for real purposes.

Stage 2: Brainstormed Tasks

Participants generated a large list of potential tasks that they would perform over a history of multimedia casual interactions. We augmented this list with other tasks we saw people perform in Stage 3 i.e., tasks that emerged as people worked with their visualizations. We analyzed these lists by categorizing them and looking for patterns.

All tasks fell naturally into five categories, as shown in Tables 1 and 2, and as detailed below. We will explain how these tasks serve the perspective of a group member interested in personal exploration and of a researcher interested in patterns and social networks.



(b) a close-up of the visualization he is creating

Task categories				
Interaction primitives	Conversations	Rhythms & trends	Descriptive statistics	Community & social networks
← Perspective →			of group member	of researcher

Table 1. Task categories and perspectives

Interaction primitives. Many participants had tasks centered on the interaction primitive held by the contents of a media item i.e., the events that occur, the conversational utterance, the actual media that appears. Almost all tasks involving interaction primitives concern *finding* specific media items or login data. As listed in Table 2a, these include finding past comments about what someone said, who was present at a moment in time, locating particular information (such as a URL held in a sticky note), locating a picture in a slide show or a video snapshot of a person. Participants formed these tasks mostly from their perspective as a group member i.e., they wanted to know about things that happened that were personally relevant to their membership in the NC community.

Conversations. Another set of tasks are characterized by grouping media items together into blocks that roughly define a ‘conversation’. All tasks involve identifying conversations and conversational properties. As Table 2b illustrates, example tasks include identifying and reconstructing past conversations, locating past conversations based on topic, identifying participants in the conversation, and relating conversations over time. These tasks predominantly help group members recall past conversations and details about it. However, some tasks are also useful to the research perspective, e.g., what defines conversations and their memberships.

Rhythms and trends. Some tasks ask questions about rhythms and trends i.e., summaries of the history that uncover patterns. These questions are typically phrased as ‘when does this usually happen’, ‘how long does it happen for’, ‘what is usually...’ and ‘how often does it happen’. Table 2c lists the specific questions, such as when do conversations occur, when are people usually present, how often do media items change, what topics a person usually talks about, etc. While these tasks appear more research oriented, they maintain a strong element of personal usefulness. For example, Begole et al describe how visualizations of work rhythms can enhance a group’s awareness of one another [1].

Descriptive statistics. Other tasks are questions that can be answered through descriptive statistics derived from analyzing the history of all media items. As listed in Table 2d, these tend to ask ‘who’ and ‘how many’ questions about individuals, such as who participates the most, how many times does one post, and so on. These tasks tend to be more research oriented.

Community and social network. The final set of tasks concern questions about community and social structure.

a) Tasks centered on interaction primitives
<ul style="list-style-type: none"> • see group member presence at a particular time • find answers to common questions • find past comments • find URLs • find slideshow pictures • find past video snapshots
b) Tasks centered on conversations
<ul style="list-style-type: none"> • identify and reconstruct conversations • find past conversations • discover participants of conversations • compare and relate conversations
c) Rhythms and trend Tasks
<ul style="list-style-type: none"> • what times to conversations occur? • what different purposes are sticky notes used for? • when are the people usually present? • which people tend to be lurkers vs interactors? • what are the patterns of numbers of people logged in? • how long do people usually stay logged on? • are there patterns in the number of media items posted? • are there patterns of posting different types of media items? • how many and how long are the breaks in daily activity? • how stable is the NC software? • how often do media items change? • what types of posts do I mostly make? • what is usually the busiest time period? • what are the patterns of activity for an individual? • what topics does an individual discuss? • what content are communicated on different media item types?
d) Descriptive statistics tasks
<ul style="list-style-type: none"> • who participates in discussions the most? • what is the maximum number of people that participate in a single conversation? • how many items do individuals post? • how many times did an individual log on in a time period? • who posts most often? • who responds to posts most often? • how often is the “clear” function used on sticky notes?
e) Community and social structure tasks
<ul style="list-style-type: none"> • can we identify social networks by analyzing which people are around at the same time and how they interact? • can we discover credibility and reputation? Who supplies good answers? How is reputation gained in the community? • what are the social norms of the community? • what are the taboos of the community? • what form does ‘social policing’ take? • what makes people decide to use the NC, IM, or email at different times? • can we evaluate the role of NC in collaboration? • how does co-located vs distributed use of the NC compare? • what are the synchronous vs asynchronous behaviors on the NC? • how does NC interaction differ from that on a bulletin board?

Table 2. Task categories and specific questions

These questions typically focused on discovering social norms and relationship networks. Table 2e includes examples of how people try to reconstruct the social networks in the user group, such as who has strong social ties to whom, questions about reputation and social norms

and so on. It also asks questions that compare communal use of the NC as a collaborative resource e.g., how people use it for co-located vs. distributed communication. The questions in this category are complex and at a high level of abstraction, and are almost entirely research oriented.

To summarize, the analysis revealed many questions that we grouped into five categories. We believe these categories can be organized linearly into a continuum that reflects data granularity and it fits to particular user perspectives (Table 1). Categories on the left side tend towards details about particular interactions, and are of most interest to group members involved in the community. Categories on the right side tend to gather abstractions about the interaction history, which are likely most relevant to a researcher.

Stage 3: Data manipulation and visualization

We gave participants the raw data, and asked them to answer the questions they had posed by restructuring and manipulating this data in any way that helped them make sense of it (as one person did in the example illustrated in Figure 4b). We then observed and recorded how they manipulated and visualized this data.

From these collective observations, we saw that most people created visualizations generally including one or more of these five organizational parameters: time scale, conversational granularity, level of detail, media item streams, and data perspective (Table 3). Each parameter is further layered into factors that describe the visualizations in finer detail. The parameters and how they interact with one another are discussed below.

Time scale. Time played heavily as an organizing principle in many visualizations, where people considered interactions over varying time scales i.e., over an hour, over days, over months, the entire logging period, and so on. However, a few visualizations did not use time at all, e.g., one showed a social network portraying frequent interactions with no indication of time.

Conversational granularity. Most participants expressed the need to group related media items (primarily sticky notes) into conversational units. Consequently many visualizations were organized around conversations. Some displayed only single conversations, which visualized a single conversation stream between a subgroup of participants. Others displayed multiple conversations containing several conversations; these included concurrent, overlapping, and temporally separate conversations. Of course, many other views did not discriminate by conversations and just showed events by some other organizational principle.

When participants focused on conversation(s) rather than individual media space events, exact time scales seemed unimportant to them (although relative order and sequencing was still highly relevant). They worked with conversation blocks regardless of the time interval between

Time Scale	Level of detail
▪ Entire logging period	▪ statistical summary
▪ Months to a year	▪ conversation
▪ Weeks	▪ individual items
▪ 1 to 7 days	▪ individual details
▪ hours	
▪ not important	
Conversational granularity	Media item streams
▪ single conversation	▪ collaborator appearance / disappearance
▪ multiple conversations	▪ lifespan
▪ not important	▪ single stream
	▪ interacting multiple streams
	Data perspective
	▪ details
	▪ overviews

Table 3: Parameters of visualizations

them. The only relevant time information was the sequencing of media item events within each conversation.

Level of detail. Different visualizations revealed different levels of detail about the data. These ranged from abstract overviews to raw data streams: statistical summaries, conversational threads, crude yet identifiable individual media items, or media items in full detail.

When examining conversations, most participants visualized the full details of all media elements constituting a conversation. One person, however, considered each conversation as a building block without revealing any underlying details, where these blocks were used to show relations with other blocks.

Media item streams. Another way participants organized their data was to look at events that happened within one or more media item streams. Some participants concentrated on the appearance or disappearance of collaborators within the video snapshot stream. Others were interested only in the lifespan of streams i.e., how long a stream persisted on the NC. Some focused on changes within a single stream, such as the evolving text in a sticky note. Others would look at the interaction between streams such as the relationship between multiple sticky notes and photo elements. In most of these cases, people were interested in which streams and events were visible (and thus of interest) to other collaborators.

Of all the streams, participants tended to be most interested in the sticky notes because these held the actual conversations. Slide shows were second in popularity. Participants generally used video snapshots to discover the presence of NC group members; most compacted this video stream into a representative snapshot that served as an icon, where they could expand it later or play it as a movie.

Data perspective. Likely because we gave participants raw primitive data, all participants' initial visualizations were detail-oriented. What was surprising was that detailed visualizations persisted over time. Even when we gave people gentle hints to try overview visualizations containing larger time-scales, some participants remained uninterested, while others created short-lived overviews as a way to locate details they wanted. However, a few participants shifted on their own accord to create

overviews; these were people more interested in conceptual explorations, such as understanding the community and its social structure.

In general, when people created visualizations looking at a very short period of time, they typically displayed full details and all media streams i.e., they wanted to see exactly what was going on during that interval. When their visualizations considered longer time periods, they tended towards much more abstract overviews e.g., conversational units *vs* detail; stream relations *vs* stream details.

Stage 4: Post-test Questionnaire

Overall, people were satisfied with the raw materials we gave them to create their visualizations. Participants felt they could express their ideas well through the combination of sticky notes for the media events, the whiteboard, and the annotation tools. They thought the whiteboard useful because they could place items on them and draw/erase graphics and notes. Still, a few participants felt constrained by the physical nature of their materials e.g., that they were unable to overlay data or to show animations effectively.

Participant's reaction to the raw data we gave them in the booklet was generally good. The majority were extremely positive about working from real data, and praised the richness of the data since it was complete. However, a few people complained that the booklet contained much irrelevant data; for the most part these were the events signaling media item appearance and disappearance.

5 DISCUSSION

The main goal of this particular study was to motivate requirements for a visualization tool that captured multimedia histories of casual interactions. Our results provide the tasks that become our user-centered requirements, and the visualization parameters imply directions for interface solutions.

From these results, we recommend several task-centered views that the visualization tool should provide, centered around the primary questions that motivate users of such a tool.

Detailed views of short time segments

Many tasks and visualizations concern people asking '*What, in detail, is happening around this moment of time?*' To answer this question, the history visualization system should provide a detailed view that displays all activities in full detail over a short time segment (minutes to an hour or two). This could perhaps take the form of an animated playback tool (e.g., showing a replay of NC activities), or a tool based around a timeline.

Conversation view

Other tasks and visualizations are centered around '*What conversations are occurring and how do they relate to one another?*' This implies that a person must be able to group relevant streams of media items (such as sticky notes and related conversational artifacts) into conversational units, and that they should be able to view these units as

conversational threads. These threads should somehow summarize or hint at their content, with details on demand. While the visualization needs to obey time ordering of conversations, it does not need to reflect real time i.e., time gaps could be compressed. These visualizations also need to show who is actually present, as people may be interested in knowing who can overhear the conversation even though they are not participating in it directly.

Rhythms and trends view.

Another common question that is reflected in both the tasks and the visualizations of our participants is '*What rhythms and trends occur on the NC over time?*' To answer this, the history visualization system should portray abstract views of patterns and trends over long time periods, e.g., as done by Begole et al [1]. These would, in general, be statistical summaries or graphs of one or more variables of interest, such as the daily rhythm of people's presence and activities, active people *vs* lurkers, and so on. Time is usually important in detecting rhythms, but may be of lesser importance for uncovering other trends. An interface could perhaps list and have the user select from the factors suggested in Table 2c+d, which would generate a view displaying the relationships between these factors.

Community and Social Structures View.

The final question asked is '*What is the social structure of the community and what are the relationships between community members?*' This visualization would be highly abstract. Time is of little concern, and the challenge would be to either manually or automatically distill the raw data into a form that answers a specific question e.g., a socio-grid showing the patterns of communication between particular members.

Supportive Features

Our study results, including participant comments, also suggest other features that, while not central to a particular visualization parameter, are listed because they help people important supportive tasks.

- *Search tool.* Many people asked for a search tool to help them find particular information e.g., conversational fragments, particular collaborators, etc.
- *Presence markers.* Most participants would like to use the history to seek presence information of a particular collaborator. Because some group members always remain logged on, simple presence on the system does not suffice. Yet manually analyzing the video stream or other personal activity is too much work. The system needs to supply an automated way to determine presence and place it in the visualization as markers.
- *Time-axis.* Time plays a strong role in several views, and thus we expect one axis would represent time (most visualizations used the horizontal axis). Because people are interested in different time scales, this axis should be scrollable and/or zoomable. A side effect of zooming into this axis could be to increase resolution of the data visible within it.

- *Grouping tools.* Some views suggest the ability to group and collapse data. Conversations are the main example, where people need to indicate regions within one or more streams that reflect a conversation.
- *Filtering tools.* Because there is an overwhelming amount data, people wanted the ability to filter their views to remove extraneous materials. Techniques such as dynamic queries can help here [6].

We realize that these are just hints for what the visualization tool should include, and that many details need to be worked out. For example, instead of providing discrete views, perhaps the system could ‘morph’ one view into another, thus revealing information about the history that answers questions from different perspectives. We will leave this for future work.

6 RELATED WORK

While many other systems visualize temporal data or conversations, we are not aware of any that tackles the difficult problem of visualization multimedia conversations.

There are, of course, many multimedia editors that create time-based sequences. While these seem far away from visualizing histories, we could learn much from the techniques they offer to manage large amounts of multimedia data by time. *Adobe Premiere*, for example, is a typical video editor that lets people combine many media types (audio, video, images, titles) into a linear and playable sequence. *Macromedia Flash* is an example of an animation editor that lets people create sequences with alternate branches.

Various systems organize historical data by time. *LifeLines* [13] visualizes histories of medical records by means of timelines, while *Dynamic Timelines* [9] is an elegant time-based history of photography. Somewhat more ambitiously, *Lifestreams* [4] is a replacement of the desktop metaphor, where all user documents are organized by time. People can navigate through this stream, search it, or filter it so that it displays particular categories of information.

A variety of systems use time to help people analyse observational data. These include *Timelines* [12], and *The Observer* (Noldus Inc.), which allow people to collect, scan, codify, categorize and map the order of observational data as they occur over time. These include video data.

Other systems visualize text messages, such as those occurring through chat and instant messaging systems. Some concentrate on conversational threads found in chat systems, in asynchronous email, or on bulletin boards [14,16]. A different tact is taken by *Conversation Landscape* [3], where participants are arrayed along the x-axis and time is represented by the y-axis. Vertical lines along the y-axis represent a participant’s conversational thread. Postings (i.e. conversations) are shown as horizontal segments intersecting each vertical conversational thread; the width of a segment reflects the

length of the message. Even though detail of the textual message could be displayed upon demand, the main goal of Conversation Landscape is to reveal the interaction pattern of the conversation at a glance. We also see many point examples of how people have visualized conversational data manually e.g., Begole et al’s many graphs displaying work rhythms of distributed groups [1].

To examine topic-specific communications, the Communication-Garden System [18] used floral representation to visualize the activities occurred in a text-based chat tool. It provides different views for examination. It shows the liveliness and duration of discussion about a specific topic. It also describes how the dynamism of the discussions changes through time. Finally, it depicts the involvement of each participant so that experts in specific topics can be easily identified.

Conversations are usually logged as time-ordered streams of data so that the temporal orientation should be preserved. Yet, due to the massive amount of data captured for informal communication, this temporality can easily be lost when fine detail needs to be scrutinized. The *Hierarchical Video Magnifier* [10] lets users work with large quantities of time-based data at different levels of granularity. The top level of the magnifier presents an extremely time-compressed overview. A video marker analogous to an adjustable sliding window lets users narrow the region to be displayed. The next level then displays the selected portion in less compressed view. Similarly, the user can recursively apply the magnifier to successive levels until optimal detail is attained. Since all previous levels remain visible to the users, the temporal context is not lost with the benefit of fine-grained visualization as successive magnification is employed.

CONCLUSIONS

Our research goal has been to explore the issues involved in presenting and analyzing histories of multimedia-based casual interactions. This is a challenging task and not one that has been addressed before.

Using data collected from the Notification Collage, the VisStream prototype was a first attempt at a solution for the problem. However, initial prototyping experiences demonstrated a need for a more grounded approach to the design.

We conducted an exploratory user study to determine firstly the tasks that were required for the tool, and secondly how the data should be presented to perform the tasks.

From the study we obtained a clear categorization of the types of tasks to be supported and characterizations of the most useful styles of visualizations.

The results of the study provide some clear directions for future prototyping of the VisStreams tool. We hope that they also provide some insight for other designers of similar tools.

ACKNOWLEDGEMENTS

We thank members of the Interactions Laboratory at the University of Calgary who participated in our user study and who gave us invaluable feedback. This work was partially funded by NSERC and Microsoft Research.

REFERENCES

1. Begole, J.B., Tang, J.C., Smith R.B. and Yankelovich N. (2002), Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups, *Proc ACM CSCW 2002*.
2. Davis M. Media Streams: An Iconic Visual Language for Video Representation, In R. Baecker, J Grudin, W. Buxton and S. Greenberg (Eds) *Readings in HCI: Towards the Year 2000*, p.854-866, Morgan Kaufmann, 1995.
3. Donath, J., Karahalio, K. and Viegas, F. Visualizing Conversation, *Proc 32nd Hawaii International Conference on System Sciences*, 1999.
4. Fertig, S., Freeman, E. and Gelernter, D. Lifestreams: An Alternative to the Desktop Metaphor. *ACM CHI Video Program*, 1996.
5. Finn, K., Sellen, A. and Wilbur, S. (Eds.) *Video-Mediated Communication*. LEA Press. 1997.
6. Goldstein, J. and Roth, S.F. Using Aggregation and Dynamic Queries for Exploring Large Data Sets. *Proc ACM CHI 1994*.
7. Greenberg, S. and Rounding, M. Notification Collage: Posting Information to Public and Personal Displays, *Proc ACM CSCW 2000*.
8. Kraut, R., Egido, C. Patterns of Contact and Communication in Scientific Research Collaboration, *Proc ACM CSCW 1988*.
9. Kullberg, R. Dynamic Timelines: Visualizing the history of photography. *ACM CHI Video Program*, 1996.
10. Mills M., Cohen J. and Wong Y. A Magnifier Tool for Video Data, *Proc ACM CHI 1992*.
11. Nardi, B.A., Whittaker, S. and Bradner, E. Interaction and Outeraction: Instant Messaging in Action. *Proc ACM CSCW 2000*.
12. Owen, R.N., Baecker, R.M. and Harrison, B. Timelines, a tool for the gathering, coding and analysis of usability data, *Proc ACM CHI 1994*.
13. Plaisant, C., Heller, D., Li, J., Shneiderman, B., Mushlin, R. and Karat, J. Visualizing medical records with LifeLines, *Proc ACM CHI 1998*.
14. Smith M., Cadiz, J.J. and Burkhalter, B. Conversation Trees and Threaded Chats, *Proc ACM CSCW 2001*.
15. Tang, C. & Greenberg, S. (2002) VisStreams: Visualizing Temporal Multimedia Conversations. *Poster Proceedings of Graphics Interface 2002*.
16. Venolia, G. and Neustaedter, C. Understanding Sequence and Reply Relationships within Email Conversations: A Mixed-Model Visualization. Accepted *Proc ACM CHI 2003*.
17. Whittaker S., Frohlich D. and Daly-Jones O. Informal Workplace Communication: What is it like and how might we support it? *Proc ACM CHI 1994*.
18. Zhu B. and Chen H. (2001), Social Visualization for Computer-Mediated Communications: A Knowledge Management Perspective, *Proc Eleventh Workshop on Information Technologies And Systems 2001*.

APPENDIX 1. VISSTREAM

We developed VisStreams, a multimedia conversation capture and analysis prototype tool that allowed collaborators to *browse and review past conversations*, and researchers to *analyze the conversational history*. VisStreams is our first prototype, and was developed before we ran the study above. We describe it here for completeness, and illustrate it in Figure 5.

VisStreams has two major functions: *conversation capture*, and *conversation visualization*. For conversation capture, we instrumented the Notification Collage [7] so we could monitor all activities and conversations contained within its multimedia elements. These elements are then recorded as a temporal stream in a database.

For conversation visualization, VisStreams presents the history as a graphical timeline (Figure 5). We chose this approach because the natural ordering of communication occurrences by time also preserves the conversation's context. The catch is that there will be far too many media elements to display at full fidelity in this timeline.

Consequently, VisStreams offers several visualization, playback and filtering techniques to allow the evaluator to selectively examine the data.

In a nutshell, the data is viewable at three levels of detail. First, an overview displays individual media elements very compactly (a few pixels each; see top row in Figure 5). The evaluator would use this overview to visually discover and study communication patterns and trends. Second, evaluators can selectively zoom into selected portions of the data set, where individual elements or rows can be expanded at will. For example, the figure shows a case where several rows have been expanded by selecting the checkboxes at their far left. Third, evaluators can get details of one or more media elements on demand by selecting them; a full-sized version of selected media elements then appears on an actual Notification Collage display, shown in a somewhat squashed form at the right of the figure.

Secondly, VisStreams offers playback capabilities, where one can position a timeline (the vertical line in the figure) and 'play' the conversation back on the Notification Collage via the video-player style buttons located at the lower middle part of the figure.

Thirdly, media elements are dynamically filtered in several ways that encourage homogeneous decomposition—the process of repeatedly partitioning the same attribute to narrower ranges of the attribute values [6]. One can select a particular time period for examination by the date slider

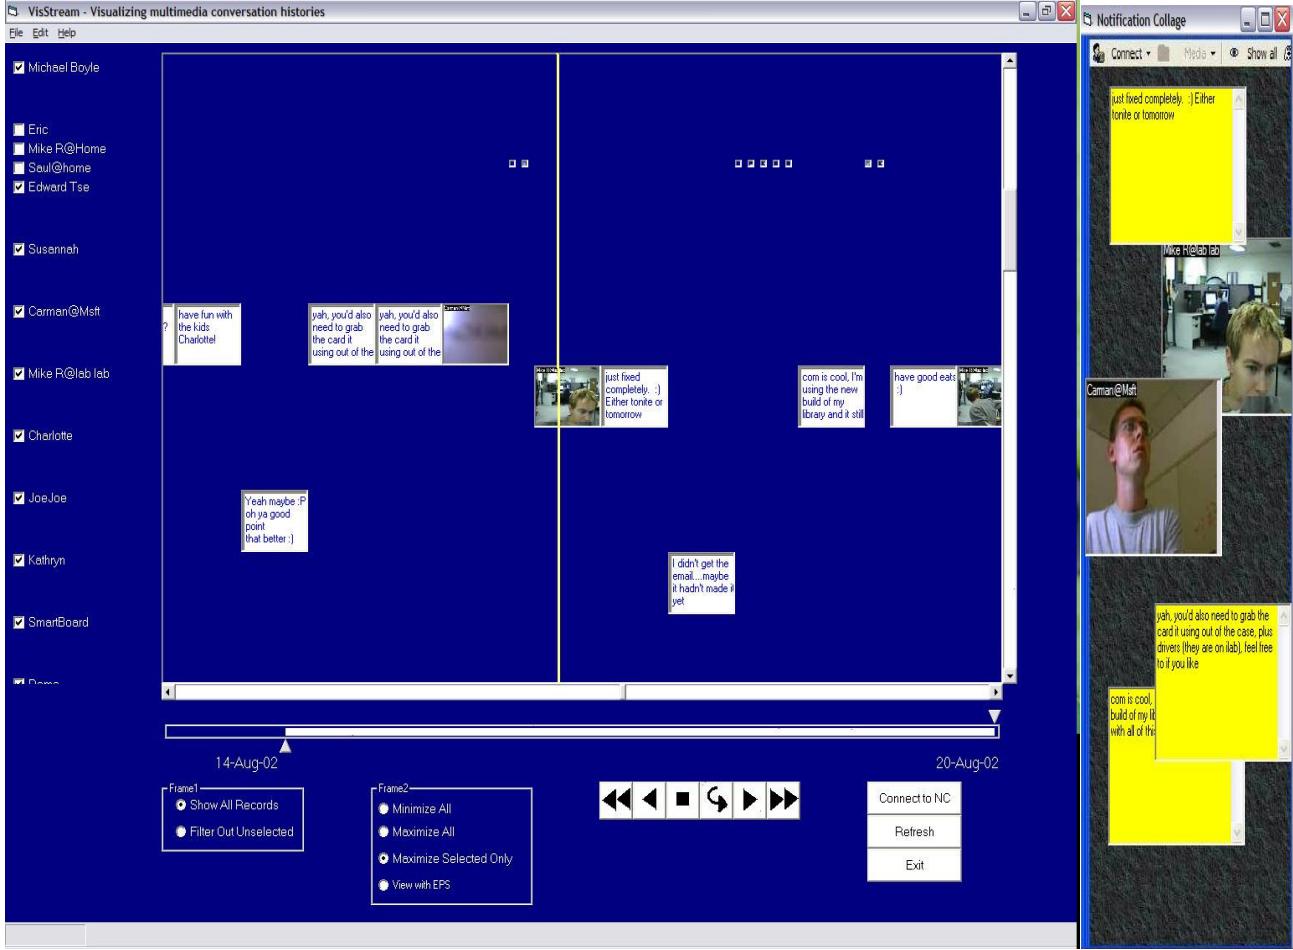


Figure 5: A snapshot of VisStream, which also allows playback on the Notification Collage (shown here in compressed form).

control visible in the lower part of the figure. Media elements outside this range are removed from view. As well, selecting rows (i.e. people and particular media streams) also affect what appears in the Notification Collage playback window; unselected rows are not displayed in it.

Collectively, these techniques allow one to detect broad patterns in the data, zoom in on particular regions to gain a better understanding of what is going on, remove uninteresting threads by filtering, get details on demand, and play back selected and filtered conversation segments at various speeds.

While the VisStreams prototype shows potential, our actual use of it to review histories left us feeling that it could be

improved. We recognized that we needed a better understanding of what people would actually do with such a tool. This is what motivated the study described in this paper. The study results now reveal the problems with VisStreams. On the surface, it is best at giving detailed views of short time segments. Yet closer examination reveals that the timeline does not reveal quite enough contextual detail. VisStreams also omits most other fundamental tasks. While we could dynamically filter the view, it was very difficult to view conversational units. It had no way of displaying rhythms and trends, nor was there any way to visualize social networks or easily discover community norms.