

ECSCW 2007

Liam J. Bannon • Ina Wagner • Carl Gutwin •
Richard H.R. Harper • Kjeld Schmidt
Editors

ECSCW 2007

Proceedings of the 10th European Conference on
Computer-Supported Cooperative Work, Limerick,
Ireland, 24-28 September 2007

Liam J. Bannon
University of Limerick, Ireland

Ina Wagner
Technical University of Vienna, Austria

Carl Gutwin
University of Saskatchewan, Canada

Richard H.R. Harper
Microsoft Research Cambridge, UK

Kjeld Schmidt
IT University of Copenhagen, Denmark

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN-13: 978-1-84800-030-8 e-ISBN-13: 978-1-84800-031-5

Printed on acid-free paper

© Springer-Verlag London Limited 2007

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licences issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

9 8 7 6 5 4 3 2 1

Springer Science+Business Media
springer.com

Table of Contents

From the editors	ix
ECSCW 2007 Conference Committee	xi
ECSCW 2007 Program Committee	xii
What did I miss? Visualizing the past through video traces <i>Michael Nunes, Saul Greenberg, Sheelagh Carpendale, and Carl Gutwin</i>	1
Social bookmarking and exploratory search <i>David Millen, Meng Yang, Steven Whittaker, and Jonathan Feinberg</i>	21
Instrumental action: the timely exchange of implements during surgical operations <i>Marcus Sanchez Svensson, Christian Heath, and Paul Luff</i>	41
Prior-to-request and request behaviors within elderly day care: Implications for developing service robots for use in multiparty settings <i>Keiichi Yamazaki, Michie Kawashima, Yoshinori Kuno, Naonori Akiya, Matthew Burdelski, Akiko Yamazaki, Hideaki Kuzuoka</i>	61
Designing family photo displays <i>Alex S. Taylor, Laurel Swan, and Abigail Durrant</i>	79
The Awareness Network: <i>To Whom</i> Should I Display My Actions? And, <i>Whose</i> Actions Should I Monitor? <i>Cleidson R. B. de Souza and David Redmiles</i>	99
“...and do it the usual way”: fostering awareness of work conventions in document-mediated collaboration <i>Federico Cabitza and Carla Simone</i>	119
A safe space to vent: Conciliation and conflict in distributed teams <i>Matt Billings and Leon A. Watts</i>	139

Table of Contents

Semi-synchronous conflict detection and resolution in asynchronous software development <i>Prasun Dewan and Rajesh Hegde</i>	159
Tag-based metonymic search in an activity-centric aggregation service <i>Michael J. Muller, Werner Geyer, Beth Brownholtz, Casey Dugan, David R. Millen, and Eric Wilcox</i>	179
The distributed work of local action: Interaction amongst virtually collocated research teams <i>Dylan Tutt, Jon Hindmarsh, Muneeb Shaukat, and Mike Fraser</i>	199
Bringing round-robin signature to computer-mediated communication <i>Takeshi Nishida and Takeo Igarashi</i>	219
Asymmetrical collaboration in print shop-customer relationships <i>Jacki O'Neill, David Martin, Tommaso Colombino, Jennifer Watts-Perotti, Mary Ann Sprague, and Geoffrey Woolfe</i>	231
Dressing up for school work: Supporting a collaborative environment with heterogeneous technologies <i>Christina Brodersen and Ole Sejer Iversen</i>	251
Exploring cooperation through a binder: A context for IT tools in elderly care at home <i>Alexandra Petrakou</i>	271
Common Information Spaces along the illness trajectories of chronic patients <i>Glenn Munkvold and Gunnar Ellingsen</i>	291
Gifts from friends and strangers: A study of mobile music sharing <i>Maria Håkansson, Mattias Rost, and Lars Erik Holmquist</i>	311
Making the home network at home: Digital housekeeping <i>Peter Tolmie, Andy Crabtree, Tom Rodden, Chris Greenhalgh, and Steve Benford</i>	331

Table of Contents

Behaviours and preferences when coordinating mediated interruptions: Social and system influence <i>Natalia Romero, Agnieszka Matysiak Szóstek, Maurits Kaptein, and Panos Markopoulos</i>	351
Health care categories have politics too: Unpacking the managerial agendas of electronic triage systems <i>Pernille Bjørn and Ellen Balka</i>	371
How-To pages: Informal systems of expertise sharing <i>Cristen Torrey, David W. McDonald, Bill N. Schilit, and Sara Bly</i>	391
Seeing Ethnographically: Teaching ethnography as part of CSCW <i>Barry Brown, Johan Lundin, Mattias Rost, Gustav Lymer, Lars Erik Holmquist</i>	411
Cues to common knowledge <i>N. Bryan-Kinns, P. G. T. Healey, D. Papworth, and A. Vadiuva</i>	431
Index of authors	449

From the editors

This volume represents the proceedings of ECSCW'07, the 10th European Conference on Computer Supported Cooperative Work, held in Limerick, Ireland. This is a significant milestone for our field, an multidisciplinary research community focusing on the understanding and technical augmentation of cooperative work practices embracing the social, technical, and design sciences. With the emergence of ubiquitous and mobile computing, new challenges have been presented to the research community in terms of how we conceptualize and investigate augmented work practices, but the issues of communication, coordination and collaboration within our working lives remain as core concerns for our field, as can be seen in the papers presented here.

The full papers presented in this volume have been selected from a pool of 123 full length paper submissions of excellent quality, making for a difficult, and lengthy, reviewing process. The relatively small number of paper acceptances demonstrates a commitment to high quality and to maintaining a single-track conference programme format, as a way of ensuring a shared awareness among the community of the latest research activity in the field. We are sure that you will find the papers of interest, and a significant contribution to the general CSCW field.

Many people have worked hard to ensure the success of this event, and we wish to briefly acknowledge them here:

All the authors who submitted high-quality papers to the Conference, thus ensuring the overall quality of the event.

All of those who contributed to the conference through taking part in workshops, master-classes, panels, demonstrations, posters, etc.

The Programme Committee and the external reviewers, who gave of their time and energy to perform a selfless task to a very high standard.

The Organizing Committee, who did trojan work behind the scenes to make everything work smoothly, on time and within budget.

The student volunteers who provide support throughout the event.

The many sponsors and supporters of ECSCW'07 for their contributions to the conference and the community more generally.

Liam Bannon, Ina Wagner,
Carl Gutwin, Richard Harper, and Kjeld Schmidt

ECSCW 2007 Conference Committee

Conference Co-chairs:

Liam Bannon, University of Limerick,
Ireland
Ina Wagner, Technical University of
Vienna, Austria

Program Co-chairs:

Richard Harper, Microsoft Research
Cambridge, UK
Carl Gutwin, University of
Saskatchewan, Canada

Conference Committee:

Workshops:

Mark Rouncefield, Lancaster
University, UK

Demos & Videos:

Annette Aboulaflia, University of
Limerick, Ireland

Posters:

David Pinelle, University of Nevada
Las Vegas, USA

Doctoral Colloquium:

Alex Taylor, Microsoft Research
Cambridge, UK
Susanne Bødker, University of Aarhus,
Denmark

Organizing Committee:

Local Arrangements:

Gabriela Avram, University of
Limerick, Ireland
Luigina Ciolfi, University of Limerick,
Ireland
Anne Murphy, University of Limerick,
Ireland

Student Volunteers:

Michael Cooke, University of
Limerick, Ireland
Emanuela Mazzone, University of
Central Lancashire, UK
Anders Sigfridsson, University of
Limerick, Ireland

ECSCW 2007 Program Committee

Liam Bannon, University of Limerick,
Ireland
Susanne Bødker, University of Aarhus,
Denmark
Barry Brown, Glasgow University, UK
Peter Carstensen, IT University of
Copenhagen, Denmark
Monica Divitini, NTNU Trondheim,
Norway
Geraldine Fitzpatrick, University of Sussex,
UK
Connor Graham, University of Melbourne,
Australia
Nicholas Graham, Queen's University,
Canada
Christine Halverson, IBM, USA
Jon Hindmarsh, King's College London, UK
Shahram Izadi, Microsoft Research
Cambridge, UK
Oskar Juhlin, Interactive Institute, Sweden
Dave Kirk, University of Nottingham, UK
Kari Kuutti, University of Oulu, Finland
Hideaki Kuzuoka, University of Tsukuba,
Japan
Du Li, Texas A&M, USA
Paul Luff, King's College London, UK
David Macdonald, University of
Washington, USA
Gary Marsden, University of Cape Town,
South Africa
Andrew Monk, University of York, UK
Michael Muller, IBM, USA
Kenton O'Hara, HP Labs, UK
Mark Perry, Brunel University, UK
Volkmar Pipek, University of Siegen,
Germany
Wolfgang Prinz, Fraunhofer, FIT, Germany
Dave Randall, Manchester Metropolitan
University, UK
Mark Rouncefield, Lancaster University,
UK

Nicolas Roussel, Université Paris-Sud (LRI)
& INRIA Futurs, France
Colston Sanger, London South Bank
University, UK
Albrecht Schmidt, Ludwig-Maximilians-
University Munich, Germany
Kjeld Schmidt, IT University of
Copenhagen, Denmark
Carla Simone University of Milano-Bicocca,
Italy
Danae Stanton Fraser, University of Bath,
UK
John Tang, IBM, USA
Alex Taylor, Microsoft Research
Cambridge, UK
Ina Wagner Technical University of Vienna,
Austria
Leon Watts, University of Bath, UK
Volker Wulf, University of Siegen,
Germany

External reviewers

Sandrine Balbo
Richard Banks
Peter Benda
Giorgio de Michelis
Mattias Esbjörnsson
Beki Grinter
Pamela Hinds
Claudia Ignat
Anab Jain
Anders Mørch
Leysia Palen
Wally Smith
David Sun
Bettina Törpel
Alexandra Weilenmann

What Did I Miss? Visualizing the Past through Video Traces

Michael Nunes¹, Saul Greenberg¹, Sheelagh Carpendale¹ and Carl Gutwin²

¹University of Calgary, Calgary, AB, T2N 1N4 Canada

²University of Saskatchewan, Saskatoon, SK, S7N 5C9 Canada

{saul, nunes, sheelagh}@cpsc.ucalgary.ca, gutwin@cs.usask.ca

Abstract. Always-on media spaces broadcast video between collaborators to provide mutual awareness and to encourage casual interaction. This video can be easily recorded on the fly as a *video trace*. Ostensibly, people can review this video history to gain a better idea of the activities and availability of their collaborators. Such systems are obviously highly contentious, as they raise significant privacy concerns. However, the ease of capturing video means that video trace systems will appear in the near future.

To push the boundaries and encourage debate about video trace technologies within the CSCW community, we created TIMELINE, a highly effective visualization system that combines ideas in *slit scanning* as used in interactive art to allow people to easily and rapidly explore a video history in detail. We describe its design and implementation, and begin the debate by offering preliminary reflections on how it can be used and misused. To encourage this debate, TIMELINE is freely available for others to try.

Introduction

Video media spaces (VMS) are always-on video channels that connect people and places (e.g., Bly 1993; McEwan 2005). Their primary purpose is to provide collaborators with awareness that leads to casual interaction. Once the subject of esoteric research requiring specialized equipment and networks (Bly 1993), the wide availability of inexpensive web cameras combined with the Internet and powerful home computers now let people easily create their own media spaces.

Researchers have argued that VMS are valuable for distributed groups with a real need or desire to recreate the kinds of interactions that normally happen when

they physically work close to one another (e.g., Kraut, Egido and Galegher 1990; Whittaker, Frolich, and Daly-Jones 1994). The video acts as a surrogate for interpersonal proximity by bringing distant people closer together: one can see other people's presence and activity over the video channel. This interpersonal awareness creates opportunities for people to engage in light-weight casual interactions through the video channel at (hopefully) appropriate times and in an appropriate manner. While such always-on video raises an Orwellian specter of Big Brother, VMS are increasingly accepted by everyday computer users, e.g., when friends stay connected with one another for long periods of the day through the free digital video capabilities of some Instant Messengers, or when small communities share a collective n -way media space through experimental systems such as the COMMUNITY BAR (McEwan and Greenberg 2005).

Video is quite good at providing rich 'at-a-glance' awareness of activities of others, which in turn lets people estimate availability for conversation. The problem is that video still misses much when compared to everyday situations where people co-habit a space (Hudson and Smith 1996). Most importantly, video demands *foreground attention*, while in everyday life people notice the activities of others in the *background periphery of attention*: for example, people see others as they walk by and their comings and goings, and they hear the sounds they make. That is, VMS only works if people are looking at their computer screen. The result is that people who use VMS do not get as good a sense of others' activities over time, which in turn affects how they can interpret what others are doing and how available they are for conversation.

To partially solve the problems of current VMS, Hudson and Smith (1996) suggest that media spaces could be augmented by giving people a visualization that displays the recent patterns of activities – the *activity history* – of others, which also reveals rhythms in people's behaviors over time (Begole et. al. 2002).

One way to reveal this activity history is through a *video trace* – a video history visualization that reveals important media space events over time, and that also allows easy inspection of those events. Using video is powerful, as the raw images potentially provide a more expressive record of past activity when compared to abstracted activity information (see §Related Work). Video is captured as it arrives (a simple matter with current computer technologies), and is reconstructed as an interactive visualization that reveals an overview of the video history. The idea is that people can infer relevant patterns in the visualization, and even explore it in detail to acquire a better idea of the past activities of their collaborators and what it suggests about their current availability.

Video trace systems are obviously highly contentious, as they raise significant privacy concerns. Brief embarrassing or private actions are now captured, easily found and replayed at full fidelity. Previously hidden time-based work patterns are revealed, such as how many hours someone is actually working over a day. While some would argue that we should not encourage research in this area, we

strongly believe that the simplicity of the idea along with the ease of capturing video means that – useful or not – video trace systems will appear in the near future, e.g., created and disseminated by grass roots developers, or included in commercial products such as Instant Messaging systems, or marketed for domestic surveillance. The real problem is that there has been little debate about such systems. Part of this problem is that prior video trace prototypes do not provide the richness necessary to really explore their expected uses (see §Related Work). Consequently, they do not incite detailed debate about their expected uses, and the tradeoffs between awareness-gathering needs and privacy.

To push the boundaries and encourage debate about video trace technologies within CSCW and other communities, we created TIMELINE: a highly revealing visualization system that allows people to easily and rapidly explore a video history in detail. As we will see, TIMELINE applies the overview and detail approach from information visualization to *slit scanning* as used in interactive arts. In subsequent sections, we describe TIMELINE’s design and implementation so others can replicate it, and how it relates to prior work. We then begin the debate by offering preliminary reflections, people’s reactions to it, as well as potential uses and abuses. To encourage further debate, TIMELINE is available for others to try: <http://groupiab.cpsc.ucalgary.ca/cookbook/> (select TimeLine in the sidebar).

Finally, while our own motivation was to see how TIMELINE provides additional awareness cues in a media space (something not yet proven), it is not constrained to this domain. Indeed, the debate has ramifications to other domains that could find video trace technology valuable.

TimeLine

The TIMELINE visualization of a video history trace lets people do the following:

- a) immediately see patterns of activity within a video history via a technique called slit-scanning;
- b) use minute, hour, day and week visualizations to present longitudinal overviews views of the history at different time granularities;
- c) explore patterns across different parts of the scene by moving the slit;
- d) rapidly explore event details within a large video stream by scrubbing;
- e) retrieve further details of the far past by selecting times of interest.

This section describes these features. We stress that the static images in this paper are a poor substitute for actually using the system. We highly recommend that viewers either try our download or watch our on-line companion video (Nunes, et al., 2006, also at <http://groupiab.cpsc.ucalgary.ca/papers/>). As well, we recommend viewing the paper’s images in color (electronic PDF or color printing) vs. as a grey-scale printout.

Slit scanning

TIMELINE uses an existing technique based on *slit scanning* (Levin 2006, Davidhazy 2007) to create a composite image of video activities over time. Slit scanning, originally developed in photography (see §Related

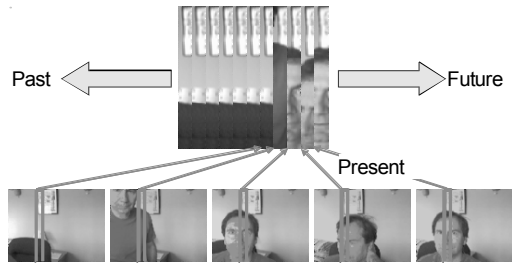


Figure 1. How video slices implement slit scanning.

Work), exposes film to only a narrow slit from a scene; while panning the camera smoothly captures a normal scene, interesting images are created by irregular panning (spatially distorted scenes), or when objects moving in the scene are seen as motion over space. The same approach is realized in video by *video slicing*.

Video slicing first extracts a scan line from a video frame, and then adds that line to a composite image over time. Figure 1 illustrates this process over 5 frame sequences. To exaggerate the effect, we illustrate a slice that is several pixels wide and ~ 1 second intervals between sequences. The same area in each frame is extracted from the captured video frames (red boxes, bottom), and added to the right side of the composite image (top). Thus the top image portrays a pattern showing an empty room (first 7 slices), and then the person arriving in the room and sitting down behind the camera. Slices need not be vertical.

TIMELINE, illustrated in Figure 2, implements video slicing using 1 pixel wide vertical slices captured at 17 frames per second (fps) to give a smooth scanning effect. The full-sized real-time video stream, shown at the bottom left of the Figure 2, is displayed within a floating window. The vertical red line within that window is the *slit focus bar*, and specifies which pixel column slice is being extracted from the video frame. The slice is added to the right side of the top row. In this case the face is ‘blurred’ as the person is moving back and forth across the scan line during capture.

Minute, Hour, Day and Week Visualizations

Somewhat similar to the Last Clock (Angesleva and Cooper 2005), a week-long timepiece is created by selectively adding slices at different time intervals to other rows. In this way, TIMELINE shows the last minute (1st row), hour (2nd row), day (3rd row), and week (4th row) of captured video. This process is done continuously. Thus the visualization itself, in combination with the dynamics of how new frames are added to it over time, gives viewers an overview of the short- and long-terms rhythms of activity across the composite image.

Each row captures different granularities of video frames, which means a single slice (especially in the week row) can actually represent many frames collected in a large interval of time. In these cases, and as described in




Figure 2. TIMELINE in action. The video history shows a personal workspace over several days.

§Implementation, TIMELINE uses an image differencing method that chooses the most information-rich frame to display from the many available frames.

Adjusting the Slit

One way in which TIMELINE extends on previous systems such as the Last Clock (Angeleeva and Cooper 2005) is by allowing viewers to adjust the focus area of the slit. As mentioned above, the red line is the slit focus bar, which indicates the column from which the video slices are taken. Viewers can interactively focus the visualization on a different area of the scene by moving the slit focus bar over a new column in the video frame. As the bar is being moved, the entire visualization is updated immediately and smoothly. Thus one can ‘scan’ the scene, where interesting patterns emerge within the minute, hour, day and week rows as the bar is moved. For example, Figure 3 illustrates the same overview as Figure 2. Here, the viewer has moved the slit focus bar over the doorway in order to see people as they come and go. The visualization clearly shows this activity, where slow scans of moving people are seen as they enter and leave in the minute row. The hour row shows the same activity around the doorway, except that the movement is compacted into single columns of activity. That is, they are seen as ‘disturbances’ in the hour timeline.

Scrubbing

Another significant difference between TIMELINE and other systems is that TIMELINE allows for very rapid and detailed exploration of the video history. When a person drags the mouse to move the cursor across any of the visualization rows, the corresponding video taken at that particular point in time is rapidly displayed in the floating window. For example, quickly scrubbing over the entire day row will replay the events of that day at the same speed, i.e., a 1 second scrub replays the whole day in a second.

Scrubbing is illustrated in Figure 3, where the viewer is scrubbing back and forth over an image of a person in the minute row (under the cursor and as indicated by the translucent band) to see replay details of that person entering the room in the floating window.

This scrubbing capability of TIMELINE is important, and distinguishes it as an awareness system. The constructed visualization lets people notice activities, variations and disturbances in the scene over different time scales, as represented in the minute, hour, day and week rows. Quickly scrubbing over an area of interest reveals the actual activity that produced that visual. That is, the visualization provides an overview of activity during the week, while scrubbing lets people quickly investigate details of possible interest.

Retrieving Details of the Far Past

Scrubbing the minute view always shows all stored frames, which were captured at 17 fps. To see beyond the last minute, one must move to the hour/day/week rows. However, the video sampling rate and thus the granularity of the playback detail in these rows is much coarser than in the minute row, as more time elapses between selectable frames. In TIMELINE, for example, each slice in the hour row actually represents about 3.3 seconds of activity. In other words, while 56 frames were seen by the system in these 3.3 seconds, only 1 is stored in the hour view. Similarly, each day slice represents 1 1/3 minutes (1344 frames), and each week slice about 9 1/3 minutes (9408 frames). Scrubbing will show the frame represented by that slice, but not the others within that interval. Yet those other frames could be important for understanding what is going on.



Figure 3. Moving the slit focus bar immediately updates the visualization. The cursor is also scrubbing a time.

To remedy this, TIMELINE provides detailed exploration of the distant past by allowing a person to select an area of interest by right clicking in one of the coarser-grained rows, illustrated in Figure 4. TIMELINE then retrieves the detailed video around that time period, populates the more detailed rows with that video, and freezes the visualization to show how it would have



Figure 4. Regions in coarser grained rows can be selected for detailed exploration in the finer grained rows.

appeared during that moment in time. Figure 4 shows a person selecting a time of interest in the day row (cursor, far left). The detailed frames around that time period are retrieved and used to populate the hour and minute rows. Feedback to show the relationships between these time periods are indicated by the green braces between these rows. Right clicking returns to the ‘live’ view.

Reading the TimeLine Visualization

Previous researchers have suggested that knowing others’ activities over time could help the viewer identify opportune moments to make contact with them, i.e., not only when others are reachable, but when they are likely amenable to being contacted (e.g., Begole et. al. 2002; Fogarty et. al. 2005). Begole et. al. (2002) observe that many activities often recur over days and weeks as long-term patterns, or ‘work rhythms’. Patterns typically indicate: when remote colleagues arrive and depart for the day or when they take breaks; whether they are working; how visitors enter and leave the area and how long they stay; the ebb and flow of meetings and phone calls over the day; and differences in activities between work days. Fogarty et. al. (2005) and Johnson and Greenberg (1999) further argue that details of activities are highly correlated with non-interruptibility. These include knowing things such as: whether they are talking on the phone and/or to guests; the number of guests, if a person is just leaving or entering, and so on.

Many of these activities are hinted at in the overview visualization, and are easily seen during scrubbing (as long as they are in the camera’s field of view). At first, images produced by the visualization may appear difficult to understand. Yet viewers quickly learn to read them. The full-sized frame in the floating window also serves as a good point of reference to help people understand what different patterns mean as they are being generated (indeed, our experiences are that people enjoy moving in front of the camera to create these patterns).

For example, consider some of the prior figures. In Figure 2, we clearly see that the person is sitting in front of their computer. The partial scans of that person's



Figure 5. The visualization reveals a brief visit.

face suggest they are moving their head only slightly, i.e., they are concentrating intently on the screen. Glancing at the hour row, we see that this person has been there for about 20 minutes. Earlier in that hour, the constant pattern suggests that the person was away. Yet there are several visual disturbances (the vertical white and grey lines), which are likely people walking by in the background. The day and week rows give a broader overview. The day row shows (from left to right), the previous afternoon, lights going off for about an hour but back on again (someone has likely left and reentered the room later), and then off again over the nighttime until the person returns in the morning. The week shows almost 5 days of day and night activity: we clearly see (going left to right) that this person has worked two partial days (the weekend) and then two full days.

As another example Figure 3 shows the TIMELINE with the focus column set on the distant doorway, revealing people as they enter and exit the room. The minute line visualization shows a recognizable person as they pass through the doorway, while the hour shows these changes as single column perturbations.

Figure 4 is somewhat similar the Figure 2, except in this case we see more ‘coming and going’ activity by the person. We know it is the same person because the colors and patterns of his shirt are the same (indeed, looking at the week view we see that he has worn the same shirt for several days).

Figure 5 extracts a portion of an Hour row from a home telecommuter. Here we see a more or less regular pattern of the person working at the computer over this hour. However, a tall red line at the left of the row differs from other colors and patterns in the scene (i.e., the telecommuter is wearing an orange shirt, not a red one). This disturbance suggests that a 2nd person has briefly entered the scene. Scrubbing over that area reveals that it is the tele-commuter’s wife coming to give him a quick kiss on the cheek (as seen in the focus window).

Figure 6a+b shows how camera angles can reveal quite different information. As evident in the focus window of Figure 6a, the worker has positioned his camera to capture a side view of his desk: the telephone, keyboard and partial view of a seated person are all visible. The annotated minute view in Figure 6b reveals activity around the desk as interesting patterns. From left to right, we see an arriving person (blue jeans and shirt are visible), then fingers as he types on the key-

board. The person then picks up his phone, as indicated by the black bar (the phone pattern) changing to tan desk color and the fingers disappearing from the keyboard. We then see the phone returning and the person resume typing for a few moments. The following solid pattern suggests that the person is still there but no longer typing (because there is no ‘leaving body’ pattern).

Implementation

TIMELINE is developed in C# .NET and uses two home-grown toolkits. EASYIMAGES provides a camera class that makes it very easy to retrieve frames from a webcam video stream attached to a particular computer. .NETWORKING is a notification server that lets people publish multimedia data to a shared dictionary data structure; the system automatically takes data posted by one client, and propagates it to other clients that have subscribed to that data (Boyle and Greenberg 2005a). The TIMELINE CAPTURE client captures the webcam video stream from a person’s computer. It then publishes each video frame to the TIMELINE VIEWER, which processes and displays it as described below.

TIMELINE is extremely responsive. To achieve this, the TIMELINE VIEWER stores in memory all relevant video frames that have a 1:1 correspondence with the currently displayed slices. This makes it possible to instantly regenerate the visualization as people moved the slit focus bar (by quickly extracting the relevant pixel slices from all frames), or to immediately display the appropriate frame in the floating window while scrubbing.

Internally, each row is represented by an object that independently samples the video input stream at a sampling rate appropriate for its time frame, and only stores the ones that are represented as a slice on the display in an array as an uncompressed frame. For pragmatic reasons, we chose a constant frame array size of 1020 for each row, and a frame size of 320 * 240 pixels in dimension – a common format for webcam streams. This provides a ‘reasonable’ balance between frame rate for the minute row (17 fps * 60 seconds = 1020 frames / minute) and memory use (1020 frames * 320 * 240 pixels/frame * 24 color bits/pixel \approx 225 MB / row or 900



Figure 6a. The scene




Figure 6b. The minute row captured by the camera angle of Figure 6a.

MB across all rows). In addition to the currently visible frames, each line also stores a single 1020 * 240 pixel image that represents the currently visible video slice visualization itself. *TIMELINE* then resizes this image to fit a window, regardless of that window's resolution.

Next, recall that each row has a different sampling rate: only the slices in the minute row shows a real time sampling of the video, whereas each slice in the minute, hour and month rows represent increasingly longer interval containing many frames. As we go from the hour to the day and the week rows, the interval between sampling successive frames lengthens, and we increase the chance of not capturing significant though brief events occurring within the omitted frames. Thus *TIMELINE* uses a *change detection algorithm* to select the 'best' frame from a series of frames in an interval as the most likely to contain a significant event: this is the frame whose pixels differ the most from the previously displayed frame. When the sampling interval elapses, the row adds that frame to its frame array, and that row's video slice image is updated by shifting it left one pixel and drawing the appropriate column from the new frame in the rightmost column.

Each row also keeps an archive of video frames on disk; this archive allows previously seen frames to be selectively displayed when a person wants to retrieve details from the distant past (it is far too expensive to keep these in memory). All frames are stored as a series of MPEG4 v2 compressed video files that are in 1020 frame numbered segments, i.e., that match a unit of frames that can fit in a particular row. As new incoming frames replace old ones in a row, the row object removes them both from the display and from memory and writes them to disk. We keep separate archives for each row rather than a single monolithic archive; this speeds up the process of retrieving and displaying region details in the distant past. When a person selects a past point in time to review, the archives comprising the rows at that particular moment of time are retrieved. Still, the operation is expensive: there is a noticeable delay (typically a few seconds) to rebuild all retrieved row images from disc. As a side note, we are careful to retain a copy of incoming video even when people are reviewing sequences from the distant past, which means that a person can switch back to the live visualization with no loss of information.

TIMELINE is memory intensive. It requires about 900 MB for the basic fully populated display. When it reads in from the archive, it is adding to what is already stored in memory: up to a maximum of around 2.2 GB if people request archives from the week view (as this reads in all other rows). Memory requirements could be brought down significantly by reducing the frames per line, the frame resolution, or frame bit depth.



Figure 7. Hammer thrower by George Silk. Reproduced from National Gallery of Australia gallery of Silk's work: www.nga.gov.au/Silk/Gallery.htm




Figure 8. LAST CLOCK. Reproduced from Angesleva and Cooper (2005).

Related Work

Photography and Interactive Art

The idea of using slit-scans to capture people's activity evolved as a method in photography, film and interactive art installations. There are far too many examples of its use to cover here: Levin (2006) has an excellent compilation.

Slit-scan methods were historically used for creating photographic distortions: a fine slit was moved past the film as the picture was being exposed (Davidhazy, 2007). An example photo is shown in Figure 7: Silk's Hammer Thrower from the U.S. track team Olympic tryouts, published in *Life* magazine on July 18, 1960.

Various art projects extend this idea to live digital video. For example, Romy Achituf's *Pixel Present* (1998) used slit-scan to capture and display live digital video of audience members walking by a large screen. The closest visualization to our work is Angesleva and Coopers' (2005) *LAST CLOCK* (Figure 8). As with other artists, they use video to capture slit-scan images of people moving around an area. However, they stretch the idea of time by fashioning the visualization as a series of concentric rings – a clock – that shows 12 hours of footage. There are three rings: the outer one is the composite image of the last 60 seconds, the middle the last 60 minutes, and the inner ring is the last 12 hours. The regions where new slit scans replace the old become the second, minute and hour hand of a clock, e.g., the white lines in Figure 8 gives the time 5:45:55.

Unlike *TIMELINE*, most approaches to slit-scan art (including *LAST CLOCK*) create a static image that cannot be explored further. To quote: "Slit-Scan imaging techniques are used to create *static images* of time-based phenomena" (Levin 2006, emphasis added). Still, a few artists allow rudimentary navigation. *ARTIFACTS OF THE PRESENCE ERA* (Viegas et. al. 2004) used a rock formation

metaphor as the visualization. Rows of slices taken over time are layered atop one another, where older layers were visually flattened and compressed to mimic strata. Visitors could crudely navigate between layers: turning a knob would display a single video frame representing an entire layer. This is equivalent to a person only being able to retrieve a single frame for an entire TIMELINE row.

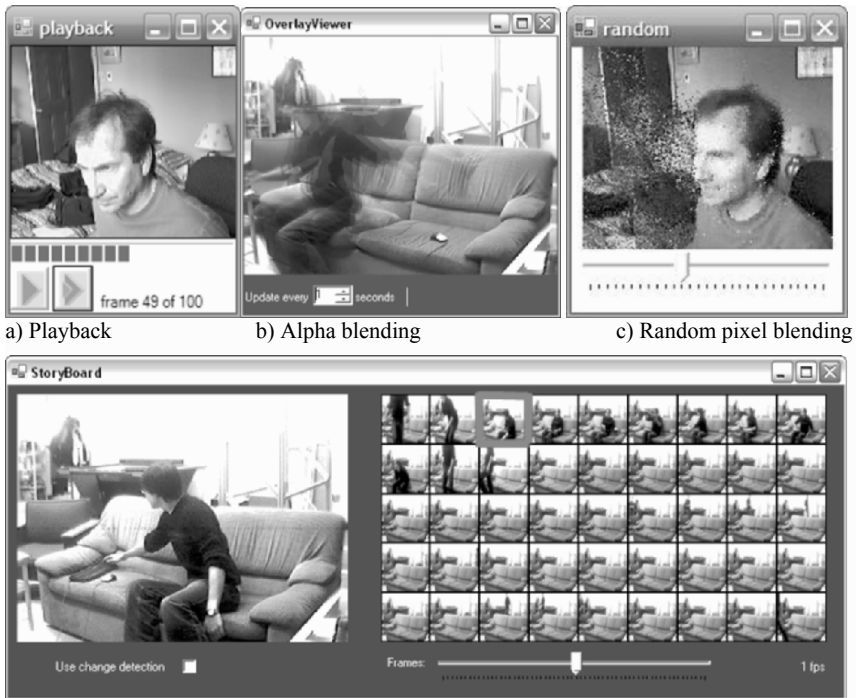
Other researchers used the idea of representing captured video as a volume, and then exploring this volume by passing a plane through it. That is, instead of capturing images by a slit-scan, it is the slicing of the volume with the plane that creates a slit scan visualization showing portions of successive frames over time. Elliot's (1993) VIDEOSTREAMER transformed video into a variety of unusual viewing streams and shapes (e.g., a 3d cube). One could see the edges of the cube (which creates a slit-scan of the sides and tops of each frame), and then navigate the shape by mousing over it. Fels, Lee and Mass (2000) let people slice through the volume at any angle and position. Again, their purpose was art: None of these systems have the level of interactivity provided by TIMELINE.

Video Media Spaces

A handful of researchers in Computer Supported Cooperative Work have experimented with the notion of somehow capturing people's activity over time, and using this information to augment a media space. Hudson and Smith motivated this in 1996 "to provide a more general idea of recent patterns of activity without requiring the constant attention of the receiving user" (p. 255).

One approach uses activity graphs: video frames are analyzed for differences, and a graph visually portrays the amount of change in the video over time. For example, an empty office will be seen mostly as a flat line, a person entering or leaving will be seen as a spike, while a seated person will be seen as a wavy line reflecting that person's small motions. Hudson and Smith (1996) introduced such activity graphs as a way to augment a media space, while Lee, Girgensohn and Schlueter (1997) recommend using such graphs *instead* of video transmission as a way to include people who refused to use video due to privacy concerns.

Hudson and Smith's (1996) WHEN DID KEITH LEAVE? uses multiple video frames to show activities over time as well as an activity graph. The system collects a small series of still images over time that characterize the flow of activity in the space. Their algorithm selects and displays a small number of frames (e.g., five) from a video stream, where the chosen frames are those that show significant visual differences in activity. Whenever a new frame is captured, it adds it to the series (while removing the oldest one) when at least 20% of the new image had changed. Thus the series of images need not be linear with time. Gutwin (2002) provides an alternative to this visualization, where he suggested (but did not implement) a short-term video trace where several video snapshots taken in the recent past are alpha-blended onto the current video frame. Hudson and Smith (1996) advocated a somewhat similar technique except that changes were shown



d) Displaying past frames as storyboard miniatures.

Figure 9. Other earlier approaches tried by the authors on visualization a video trace

as shadows, thus showing some activity (darkness = movement) while masking its details. In all these works, the concept of video traces was a side issue that was not explored in depth: Gutwin (2002) was investigating traces as a way to mitigate network issues in telepointer tracking, while Hudson and Smith 1996 were primarily focused on privacy in VMS. In other related work, Terry (2004) explains a method for simultaneously showing multiple points of time from the same scene in a single image. More generally, Roussel introduces the notion of *multiscale communication* as a communication system that supports variable degrees of engagement.

In parallel with video traces, a variety of other researchers try to extract and visually portray patterns collected from computer usage logs and from physical sensors. For example, Begole et. al. (2002) describe how they generate rhythms of personal activities as actograms by analyzing and modeling a person's computer activity, their mail, phone and instant messaging use, and their online calendar appointments. Fogarty, Hudson, et. al. (2004) argue that sensor data can be used to generate models that differentiate between interruptible and non-interruptible situations. While their purpose is to use this data mostly for automating whether a person is interruptible at a particular instance in time, the same data could easily be used to generate temporal patterns of activity.

Our Early Investigations

Our investigations prior to building TIMELINE led to other visualizations of activity history. First is *rapid playback*, inspired by Dietz and Yerazunis's (1991). When a person moves and then returns the cell phone back to the ear, the phone replays the missed conversational passage as high-speed pitch-preserved audio. Similarly, our video-based approach, illustrated in Figure 9a, displays the current image while automatically capturing the last n video frames (we used 100 frames captured at 1 frame every two seconds, or 2 1/3 minutes of past activity). The viewer can play back this captured video stream at 20x or 40x normal speed (i.e., as a 10 second or 5 second movie) by pressing a particular speed replay button (Figure 9a, buttons on bottom left – this screen capture is in the middle of rapidly replaying the sequence). This effect works quite well for quickly reviewing the very recent past, and can be extended to include longer time periods and different frame rates. It is limited because no overview is given of past activities, and the person has to actively decide to review the video.

We then realized Gutwin's (2002) unimplemented idea of *frame alpha-blending*, that composites several past frames onto a single video frame, as illustrated in Figure 9b (this varies Hudson and Smith's 1996 Shadow method). The further into the past, the more faded a change appeared. As before, a person could set the frame rate. While this did give a sense of activity, it did not satisfy. It did not scale well beyond a modest number of frames (as the changes in the alpha blended images proved too translucent). If the frame rate was high, so was the glimpse into the past (i.e., activity was usually realized as motion blur). If the frame rate was low, then the composite images were disjoint and hard to interpret.

With *random pixel blending*, we overlaid a percentage of randomly chosen pixels from the just-taken video frame onto the displayed image (Figure 9c). The actual percentage is specified via the slider. If the percentage is small (e.g., 15%), then motions are realized as a pixel scatter blur effect. If a person remains somewhat still, details are slowly filled in. Because backgrounds change rarely, they come in at full view. For example, Figure 9c shows the same image as in Figure 5a. The telecommuter is sitting fairly still behind his computer – thus his image (as well as the room background) is more or less complete. His wife just walked in and gave him a kiss on the cheek; because she is moving, this is seen as scattered pixels as she came through the doorway (top left). A hint of her clothes color is just visible (red shirt and blue jeans) at the bottom left, as is the top of her head as she leans over to kiss him (middle left). Thus privacy is somewhat protected: people are aware of stable events, but details and rapidly done actions are obscured. As with alpha blending this works only with the very recent past.

Figure 9d illustrates our *storyboard display*, a variation of Hudson and Smith's "When did Keith leave" approach (1996). An end user sets a sample rate for extracting frames from the video stream, and the last n samples are displayed as miniatures in a visual storyboard. For example, Figure 9d has a sample rate of 1

per second; thus the 45 miniatures in the storyboard reveal the last 45 seconds of activity. If it was set at (say) 1 sample per minute, the storyboard would reveal the last 45 minutes of activity. As with *TIMELINE*, it displays the frame that differs the most from the last storyboard sample. New samples also over-write old ones in the wrap-around sequence (the latest sample is outlined in red, i.e., the 3rd image in the top row in Figure 9d). Compared to “When did Keith leave”, this approach conserves time as a true linear stream, and the smaller low resolution images adds some privacy protection. Yet our own impressions of this storyboard technique were not favorable. Because of the many images, it was hard to tell ‘at a glance’ what was going on. As well, setting a high sample rate (e.g., 1 sample per second) proved distracting and did not provide a long enough history window to justify the screen space. Yet setting a low frame rate (e.g., 1 sample per minute) omits multiple key events that could happen within the sampling period, e.g., 2 visitors arriving and leaving within a few moments of each other.

Reflection

This paper described *TIMELINE* as a very efficient method for seeing and reviewing past events. We placed *TIMELINE* within the context of prior work, provided enough details of its implementation so that others can replicate it, and also make it freely available for others to try.

Yet our motivation is not to provide *TIMELINE* as a solution. Rather, we want to use it as a case study that pushes the extremes of what is possible, and to provoke debate about video trace technologies. In this spirit we re-ignite the debate, started by Hudson and Smith (1996) and Begole et. al. (2002), by offering a preliminary reflection on *TIMELINE*’s use. We base our reflections on the reactions of many people (including ourselves) to the *Timeline* system during numerous live demonstrations and through self-trials. First, we ran *TIMELINE* as a video mirror within a public interactive installation. The installation was located in a highly visible public part of our research laboratory on a 60” touch-sensitive plasma display (using a webcam attached to this display), and hundreds of visitors saw and tried it as part of several open-house days. Second, we used *TIMELINE* ourselves, on personal workstations within our laboratory as well as at home.

People’s Positive Reactions and Suggested Uses.

Readability. The live installation verified that the visualization is easily readable. With only a brief introduction to how it works, visitors could comprehend and read the visualization display, especially because they could see how their actions were immediately represented within it. They were able to spot themselves in the visualization, pick out events such as the arrival of a crowd for a demonstration session, and see the rhythms of activity between night and day. People found it

easy to reason about what the display was showing. They could find odd events of interest, such as a brief period of light in the middle of the night. Scrubbing and moving the slice focus bar let people determine how patterns in the visualization were created, enhancing their understanding as well as their ability to read other parts of it. While static images produced using slit-scanning techniques can appear to be distorted and strange, augmenting them with this level of interactivity brings them into comprehension.

Self-Reflection and Playfulness. People were intrigued to see themselves within Timeline. This replicates prior experience in slit-scan art installations. People made patterns within the visualization by waving arms or colourful items, or even slowly turning around in front of the camera to produce a flat scan of their head all the way around. Visitors would often pose in front of the installation to leave their mark in the visualization (also noted by Viegas et. al. 2004).

Voyeurism occurs when people get pleasure observing other people, and this certainly occurs within Timeline use. People were often fascinated by TIMELINE'S ability to see what others were doing in the past, where they found scrubbing and cyclic replay of a scene compelling.

Rhythms. People commented on the aesthetics of the TIMELINE visualization as a history mirror: it clearly showed the cyclic rhythms of activity and how changes occur over time within a space.

Surveillance. While the public demonstrations were set up as an interactive installation, people also suggested practical uses for it. Many said that TIMELINE had great potential as a surveillance system used for security purposes. Ideas ranged from its use by trained security personnel, to home monitoring (nanny cams), to property protection when away, and to equivalents of baby monitors.

Analysis Tool. People also suggested that TIMELINE could be useful as a research analysis tool for detecting patterns and counting key events in captured video streams collected for research TIMELINE.

Video editing. Could TIMELINE be used within a video editing context? We don't have the answer to this, for the constant motion of a camera would give quite a different visual effect than what has been shown in our paper. Still, as an alternate visualization TIMELINE may be useful to find scene changes, or moments when the camera's point of view have drifted off its central target to a new target in the scene.

A low-bandwidth ambient distributed awareness. TIMELINE may also work as a pure peripheral awareness display that doubles as an art installation. Instead of sending video frames, we could just send columns (the vertical slice) which is quite low bandwidth. This visualization could work as a 'long-distance relationship' awareness display that is less intrusive than full frame video, where its lower fidelity protects privacy better than full video.

Negative Reaction: The Privacy Issue.

Now we turn to the dark side. Privacy is an extremely serious issue, since *TIMELINE* sometimes does its job too well. Hudson and Smith (1996) believe systems such as these demonstrate how privacy issues might occur when providing awareness information through a video trace: transient activities of a person are no longer lost if the system records and displays significant events. Our own uses of the system reflect these concerns.

Reluctance and perception of risk. Even as creators we were reluctant to use it in broadcast mode for long periods of time. *TIMELINE*'s power makes it invasive. Potentially embarrassing, private or inappropriate behaviors are not only captured, but easily found and replayed. A trade-off arises: while we and our collaborators could benefit by using the system to see each other's recent activities and events, it also serves as a (perhaps unintended) surveillance system (see Boyle and Greenberg 2005b discussion of privacy issues in VMS).

It captures more than you. We often used *TIMELINE* as a local mirror within our research laboratory, an open space that included other workers. Yet this immediately incited concerns from other inhabitants of the laboratory, for they were very uncomfortable with the idea of being permanently captured on video for others to see (even though all actions were in a public space). There were meetings and discussions about this, and eventually people were willing to have it run in very limited situations because they valued its use for research. Even so, we found that we inadvertently captured other people in embarrassing situations, e.g., a cleaner who went to sleep on a lab couch during a night work shift.

Similarly, one author telecommutes from his home, and regularly uses a traditional VMS. He was willing to use *TIMELINE* in work situations with his distant colleagues, but was concerned that other family members using the home office would be captured unintentionally. This would not only be embarrassing for a family member if caught in a compromising situation, but also for the distant viewer who could unintentionally see that situation in *TIMELINE*.

Tacit information becomes explicit. Another downfall is that *TIMELINE* also allows people to very easily compare their activity with others (if multiple instances are running), and in turn they may become overly self-critical of their own work habits, or of others. Tacit information becomes explicit. The system allows observations such as "did I work as many hours as my colleagues?" This could lead to undue stress on the workers' part, which would be magnified if the worker knew their video stream was being broadcast.

Distraction. There is also the issue of distraction with *TIMELINE*'s full-screen view. This can be solved by embedding it into more traditional media spaces. For example, the inset figure shows it redesigned to fit within Community Bar (McEwan and Greenberg 2005). People see a traditional media space view in a side bar, but can raise a video trace by mousing over it for an overview of hourly activities. Clicking this trace then raises the full *TIMELINE* system.

Yet we can control what is captured. Adjusting the camera angle can make a large difference in what is captured, and consequently can reduce privacy invasion. For example, Figures 6a+b show how pointing the camera at a desk highlights when a person is at their computer and/or on the phone, but does not transmit head shots or surrounding areas. Similarly, positioning the camera outside an office doorway (as done by Buxton 1997) means that the remote viewer can only see what people walking by a hallway can see, and that closing the door protects oneself in both the physical and virtual space.



Next Steps

There is much that we have not done in this paper. We did not formally evaluate TIMELINE'S usability as a visualization, nor did we formally compare it to other time-based approaches. We speculated, but did not prove, that a video trace is actually useful as an availability tool. We argued that privacy is a problem but have not really deployed it to see how people would really adjust and/or negotiate its use in their real world context. Clearly, there are many future avenues of research (but this should not preclude debate). A few are described below.

Usefulness as an awareness tool. Our original motivation was that a video trace would be a good way to augment a media space, where it would provide people with addition cues to determine availability and interruptability. Yet the jury is still out on this. It could be that a glimpse into the past adds little, or that very simple techniques showing the last few moments and/ or key events would suffice. Indeed, there is active debate within CSCW concerning exactly what information suffices as a good predictor of interruptibility (Fogarty et. al. 2005).

Comparitive usefulness as a video history tool. We argued that TIMELINE is a good visualization tool that provides both an overview and query-in-depth. While our informal observations suggest that this is in fact the case, we have not compared it to other 'conventional' techniques. Perhaps TIMELINE offers insufficient fidelity of historical information (narrow field of view, image quality, lack of audio). Or perhaps other approaches for visualizing video-based activities over time may prove more effective, e.g., Terry (2004), Hudson and Smith (1996), Begole et. al. (2002). This comparison clearly needs to be done.

Art was the original motivation for slit scanning *vs.* purposeful applications. We believe that bringing high interaction capabilities to bear on slit-scanning could enhance overall artistic effect. However, we are not artists, so next steps could include involving artists in the project to see how they modify and re-purpose TimeLine.

Novelty *vs.* Real Use. Is the attraction of TIMELINE due to a novelty effect, or does it have long-standing value? Once people learn the visualization, how do

they use it to interpret a scene and to look back at history? What canonical tasks emerge over time, and how can we modify *TIMELINE* to suit these tasks? Obviously, these questions can only be answered through a long-term longitudinal study of *TIMELINE*'s use in real world contexts by real people.

Privacy protection is obviously important, but the question remains of how one can balance privacy and awareness. Perhaps a good starting point is to embed and evaluate other privacy preservation image processing techniques within *TIMELINE* e.g., context-sensing for controlling when video is taken (Neustaedter and Greenberg 2003), image blurring in low-risk situations (Boyle and Greenberg 2005b), or novel techniques such as those proposed by Hudson and Smith (1996).

Conclusion

TIMELINE is best viewed as an investigation into the extreme, where it asks the question “what would it be like if we could easily see and explore a video trace of a distant person or scene? Currently, it is difficult to predict where it might be best used and the cultural practices that would evolve around it. Perhaps *TIMELINE* could be suited to areas that are largely accepted as public rather than private offices and workspaces. Perhaps it would work well between tight teams or social intimates with a strong desire to stay connected. Perhaps it can be used as a base reference to other approaches, e.g., to see if low fidelity actogram graphs (Begole et. al. 2002) provide comparatively sufficient awareness information while still safeguarding privacy. Or perhaps *CSCW* is the wrong venue; it may better serve as a security system, as art, or as video analysis. This is clearly an area for future – and perhaps controversial – research. Let the debate continue.

Acknowledgements

This research was partially funded by NSERC through its scholarship, Discovery Grant and NECTAR Networks grant program, and by Alberta Ingenuity and iCORE. We also thank our industrial sponsors: Microsoft Research and Smart Technologies, Inc.

Software. *TIMELINE*, *EasyImages*, and *Networking* available at <http://grouplab.cpsc.ucalgary.ca/cookbook/>.

References

- Achituf, R. (1998): ‘*Pixel present: A gestural scanner.*’ Accessed March, 2007. <http://www.gavaligai.com/main/sub/installation/PixelPresent/PixelPresent.html>
- Angesleva, J. and Cooper, R. (2005): ‘Last Clock’. *IEEE Comput. Graph. Appl.* 25, 1, 20-23.
- Begole, J., Tang, J. C., Smith, R. B., and Yankelovich, N. (2002): ‘Work rhythms: analyzing visualizations of awareness histories of distributed groups’. *Proc ACM CSCW '02.* 334-343.
- Bly, S. Harrison, S. and Irwin S. (1993): ‘Media Spaces: Bringing people together in a video, audio, and computing environment’, *Comm. ACM*, 3, 1, 28-47.

- Boyle, M. and Greenberg, S. (2005a) 'Rapidly Prototyping Multimedia Groupware'. *Proc Distributed Multimedia Systems (DMS'05)*, Knowledge Systems Institute, IL, USA.
- Boyle, M. and Greenberg, S. (2005b): 'The language of privacy: Learning from video media space analysis and design'. *ACM Trans. Computer-Human Interaction (TOCHI)*. 12 (2).
- Buxton, W. (1997): 'Living in augmented reality: Ubiquitous media and reactive environments'. In K. Finn, A. Sellen & S. Wilber (Eds.). *Video Mediated Communication*. LEA Press.
- Davidhazy, A. (accessed March, 2007) *Slit-scan photography*. School of Photographic Arts and Sciences, Rochester Institute of Technology. <http://www.rit.edu/~andpph/text-slit-scan.html>.
- Dietz, P. H. and Yerazunis, W. (2001): 'Real-time audio buffering for telephone applications'. *Proc ACM UIST*. 193-194.
- Elliot, E. (1993): 'Watch * Grab * Arrange * See: Thinking with motion images via streams and collages'. MS Thesis, Master of Science in Visual Studies, January, MIT.
- Fels S., Lee, E. and Mase, K. (2000): 'Techniques for Interactive Video Cubism'. *Proc ACM Multimedia*. 368-370.
- Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., and Yang, J. (2005): 'Predicting human interruptibility with sensors'. *ACM TOCHI* 12 (1).
- Gutwin, C. (2002): 'Traces: Visualizing the immediate past to support group interaction'. *Proc Graphics Interface*. 43-50.
- Hudson, S. E. and Smith, I. (1996): 'Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems'. *Proc ACM CSCW*. 248-257.
- Johnson, B. and Greenberg, S. (1999): 'Judging People's Availability for Interaction from Video Snapshots'. *Proc. HICSS Distributed Group Support Systems Minitrack*, Jan., IEEE Press.
- Kraut, R., Egidio, C., Galegher, J. (1990): 'Patterns of contact and communication in scientific research collaboration'. *Intellectual Teamwork: Social and Technological Foundations of Co-operative Work*. LEA
- Lee, A., Girgensohn, A., and Schlueter, K. (1997): 'NYNEX portholes: initial user reactions and redesign implications'. *Proc ACM GROUP*. ACM Press, 385-394.
- Levin, G. (2006) 'An informal catalogue of slit-scan video artworks'. Last edited December 29, 2006. http://www.flong.com/writings/lists/list_slit_scan.html.
- McEwan, G. and Greenberg, S. (2005): 'Supporting social worlds with the Community Bar'. *Proc. ACM Group* 21-30.
- Neustaedter, C. and Greenberg, S. (2003): 'The design of a context-aware home media space. *Proc UBICOMP'03*, 297-314, LNCS Vol 2864, Springer-Verlag.
- Nunes, M., Greenberg, S., Carpendale, S. and Gutwin, C. (2006): 'Timeline: Video Traces for Awareness'. *Video Proc. ACM CSCW'06*, ACM Press. Viewable at <http://grouplab.cpsc.ualgary.ca/papers/>
- Roussel, N. (2006) 'Towards multiscale communications systems'. Rapport de Recherche 1439, LRI, Université Paris-Sud, France.
- Terry, M. (2004) 'Making space for time in time-lapse photography'. *Siggraph Tech. Sketches*.
- Viégas, F. B., Perry, E., Howe, E., and Donath, J. (2004): 'Artifacts of the Presence Era: Using Information Visualization to Create an Evocative Souvenir'. *Proc IEEE Infvis'04*. 105-111.
- Whittaker, S., Frolich, D., and Daly-Jones, O. (1994): 'Informal workplace communication: What is it like and how might we support it?' *Proc ACM CSCW*, 131-138

Social bookmarking and exploratory search

David Millen, Meng Yang, Steven Whittaker*, Jonathan Feinberg
IBM Cambridge, MA, USA 02142, *Sheffield University, Sheffield, S1 4DP, UK
{*david_r_millen,yangmeng,jdf@us.ibm.com*}, *s.whittaker@shef.ac.uk*

ABSTRACT. In this paper, we explore various search tasks that are supported by a social bookmarking service. These bookmarking services hold great potential to powerfully combine personal tagging of information sources with interactive browsing, resulting in better social navigation. While there has been considerable interest in social tagging systems in recent years, little is known about their actual usage. In this paper, we present the results of a field study of a social bookmarking service that has been deployed in a large enterprise. We present new qualitative and quantitative data on how a corporate social tagging system was used, through both event logs (click level analysis) and interviews. We observed three types of search activities: community browsing, personal search, and explicit search. Community browsing was the most frequently used, and confirms the value of the social aspects of the system. We conclude that social bookmarking services support various kinds of exploratory search, and provide better personal bookmark management and enhance social navigation.

Introduction

In recent years, there has been tremendous growth in shared bookmarking applications. Introduced in 2003, the *del.icio.us* social bookmark website was one of the first of this kind of application, and has enjoyed an early and large base of committed users. A flurry of similar offerings has since been unveiled [see (Hammond, et al., 2005) for a recent review].

These internet oriented social bookmarking services have been adapted for use in large organizations. Examples include the *dogear* (Millen, et al., 2006) and *onomi* social bookmarking services (Damianos, et al., 2007). Both of these enterprise-ready bookmarking services support bookmarking of both internet and in-