

PICASSO: Pervasive Information Chronicling, Access, Search, and Sharing for Organizations

Sinem Güven¹

Mark Podlaseck²

Gopal Pingali²

¹Columbia University,
Department of Computer Science,
New York, NY 10027.
sinem@cs.columbia.edu

²IBM T.J. Watson Research Center,
19 Skyline Drive,
Hawthorne, NY 10532.
{podlasec,gpingali}@us.ibm.com

Abstract

Several researchers have pointed out and begun to demonstrate that it is now possible to digitize an entire lifetime of experiences into a pocket-sized storage device, and thus to create a rich and portable electronic chronicle of an individual's life and activities. This possibility has brought with it a host of significant challenges: How can the creation of these electronic chronicles be natural and effortless? How should these chronicles be organized? What kinds of navigation and search tools unleash the potential of the chronicled data allowing the user to receive and retrieve the information most relevant to their context at any time? How can such chronicles impact business organizations? How can people in organizations share chronicles and effectively combine chronicled data from different individuals in collaborative settings?

Our research is beginning to address these issues with a focus on the business organizational setting. In this paper we introduce PICASSO, our work on pervasive chronicle creation and exploitation in enterprise applications. PICASSO enables the capture of rich user context (including PC/PDA interactions, audio, video, images, and location) and provides tools that enable users to search, navigate, share, and merge personal events, both from desktop computers and from mobile devices.

1. Introduction

With ever-diminishing form factors and ever-increasing storage volumes of digital devices, it is becoming possible for individuals to capture more and more aspects of their lives in wearable accessories. This has led several researchers to start exploring what it means to comprehensively capture life experiences, (including everything seen and heard, every document worked on, every email exchanged, etc. [16]), thus resulting in a digital record of an individual's life. This vision, in fact, stretches as far back as 1945 [6]. While

capture and storage of information through a variety of pervasive sources is fast becoming a reality, it is becoming a challenge to effectively manage and navigate these vast repositories of information both at a personal and at an enterprise level.

At the personal level, inadequacies of folder/directory based organization of information have long been recognized, resulting in a significant body of work on personal information management [3, 10, 11, 12, 14, 15, 20, 25, 27, 31, 33, 36, 37, 38, 40], as well as information organization techniques and their assessment [2, 7, 8, 17, 18, 19, 24, 28, 29, 32, 34, 39, 40]. Many of these approaches have emphasized the importance of the temporal dimension in organizing and browsing personal information [15, 35, 36, 37] while others have applied information retrieval techniques to personal repositories [12, 20, 38]. Based on these trends, the vision of a multimedia *electronic chronicle* [21, 22, 23] has recently emerged, which includes notions of logging all user activity through various mechanisms and presenting this multimedia information in virtual time-machine fashion, to enable users to review important episodes and gain crucial insights.

With the advent of PDAs, PDA-cell phones, and similar mobile personal devices, many researchers and commercial systems have begun to address information management and navigation interfaces on such small factor devices [4, 5, 30]. However, much of this work has focused on capture and organization of information at a personal level with little emphasis on an enterprise context. We find that in enterprises today, much valuable information about employee activities, interactions, relationships, experiences and insights is not i) captured, ii) effectively organized for browsing and retrieval, iii) shared or published, iv) available for searching and mining, nor v) integrated into business processes. This is the motivation for the work presented in this paper on PICASSO.

Figure 1 shows our vision of PICASSO functioning in an enterprise context. As seen on the left of Figure 1, at an individual employee level, PICASSO enables the

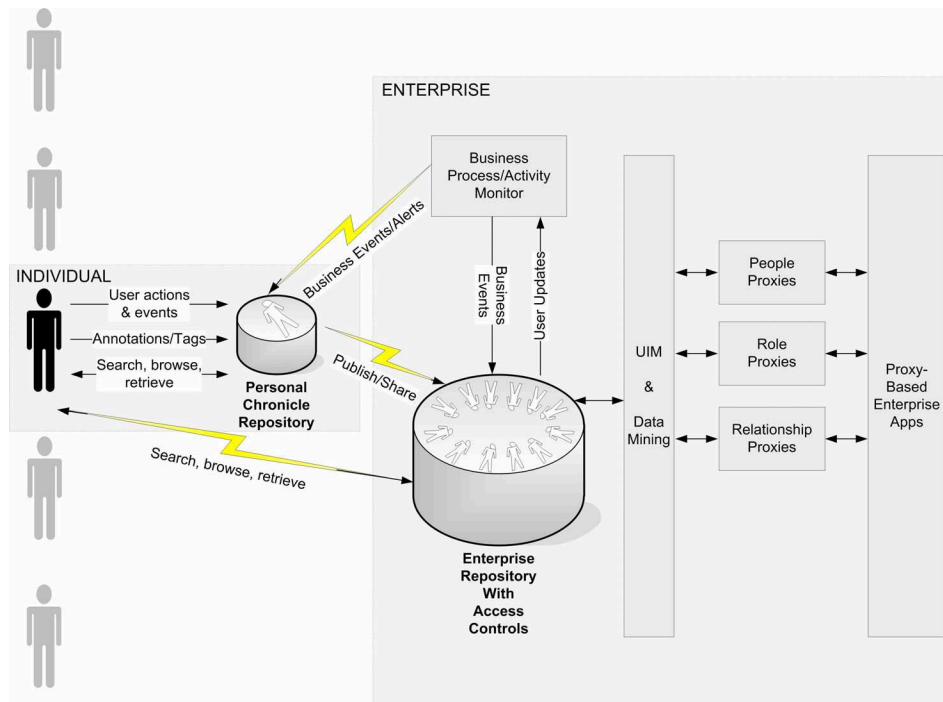


Figure 1 PICASSO in the Enterprise

capture, annotation, searching, and browsing of information, interactions, actions and events. This information is stored in a local repository accessible only by the user. As seen on the right of Figure 1, information from personal chronicles is transferred to enterprise repositories when the user chooses to publish/share this information and with the access specified by the user. As rich data about individuals becomes available in the enterprise repository, powerful data mining and unstructured information management tools can be applied to it. We envision three important representations resulting from this: i) people proxies – rich dynamic representations of people in the enterprise, ii) relationship proxies that capture associations between individuals, and iii) role proxies that capture information related to different roles in the enterprise and relate these to specific individuals. A number of proxy-based enterprise applications can then be built on these representations. In addition, we note the growing importance of business process modeling and business activity monitoring tools at the enterprise level. Data from individual chronicles can provide such tools with ongoing updates of user activities, while business events and alerts can be communicated back to relevant individuals and stored in their chronicles, cross-indexed with the rest of their personal information and events.

PICASSO builds upon our initial work on Personal Chronicling Tools [26] where we developed desktop PC based tools for enterprise workers with emphasis on monitoring all user interaction with the PC, and interfaces for tagging/annotation. In this paper, we extend our work

to mobile devices such as PDAs, include capture through a variety of multimedia devices, and lay special emphasis on unified interfaces for both navigating chronicles and sharing them across the enterprise using either PC's or PDA's.

2. Motivating Scenario

As a motivating scenario for PICASSO, consider a business consultant, John Doe. John works for XYZ Corp. and is traveling to close a business opportunity with a client, the JKL Corp. This familiar situation offers the following opportunities for improvement:

John's calendar entry reads: "Travel to Meeting with JKL from 1 to 2 pm" and is accompanied by his flight information. After the trip, the entry contains the same (now useless) information. Why not augment the entry with details on how John actually spent his day? PICASSO aims at utilizing calendar entries as indices into the past, as well as the future.

On his trip, John met several people for the first time. John is bad with names and is certain to forget which name belongs with which face by the time of his next trip. PICASSO's support for digital voice recordings, images, and videos enables John to capture name/face/role relationships as well as impressions that occur as business events unfold.

As John travels, he wants to prepare for the meeting by reviewing relevant emails, notes, documents, instant messages, calendar entries, address book entries, and internal web sites. Currently, these diverse materials are

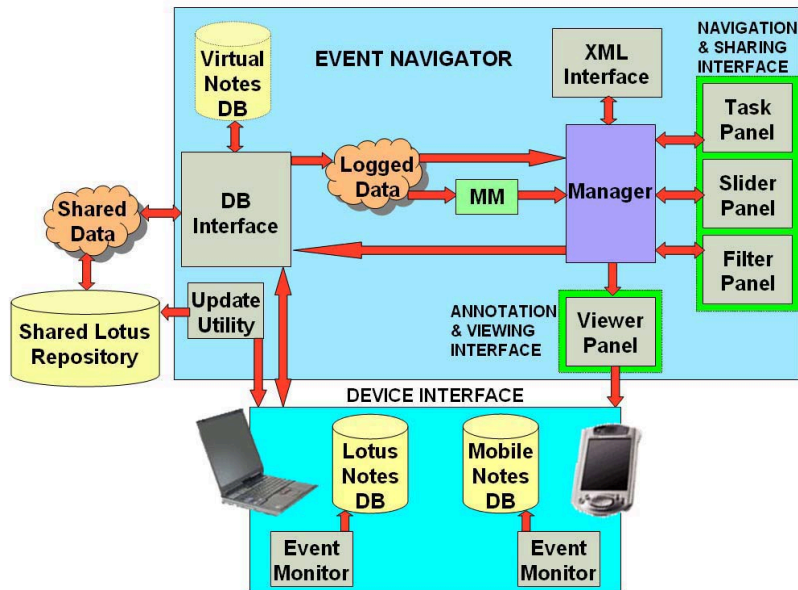


Figure 2 PICASSO System Architecture

not easily consolidated. PICASSO aims to provide the search tools to gather relevant information, regardless of its format, and the navigation tools to peruse it seamlessly.

No matter how rich the capture of activities at an individual level, the chronicle of a trip or a meeting is far more complete when it includes information from the other people involved. John's record for the meeting with JKL would benefit immensely if it also included the impressions and background material of colleagues who attended the meeting. By the same token, the others benefit from the information shared by John with them. PICASSO aims to simplify sharing of information between people in the organization, and to provide effective ways of combining one's own chronicle with the chronicles of others.

During his trip, John obtained crucial information and data about an important new initiative at JKL. The significance of this opportunity may extend well beyond his project. In a large organization such as John's, many people, who could benefit from this information, do not even know of John's existence. PICASSO aims to provide tools for selectively broadcasting information across an enterprise, for searching and mining for such information at an organizational level, and mechanisms to alert appropriate people in the organization based on their roles and profiles.

3. PICASSO System Design & Architecture

As shown in Figure 2, the PICASSO system architecture comprises core PICASSO software components as well as enterprise software utilities such as

Lotus Notes, Mobile Notes and XML. The core software consists of two main modules, namely the *Event Navigator* and the *Event Monitor*. Both can run on a PC or a PDA. The *Event Monitor* runs in the background to automatically monitor a user's activity and logs events in a user's local database. The *Event Navigator* provides two key functions: a *Navigation & Sharing Interface* and an *Annotation & Viewing Interface*. From the architecture perspective, the Event Navigator comprises the following components: Manager, Task Panel, Slider Panel, Filter Panel, Viewer Panel, Virtual Notes DB, Update Utility, DB Interface and XML Interface. The rest of this section provides more details about each component.

The Manager acts as a central control mechanism that communicates with each component, coordinates their interactions, and passes the resulting graphics to target displays. When PICASSO Event Navigator is first started, the Manager tries to connect to the relevant databases (both the local Lotus DB and the shared Lotus repository) through the DB Interface. If it succeeds, the DB Interface then reads the data and passes it to the Manager along with the relevant multimedia files. The Manager then populates the internal data structures, writes the multimedia files to the hard disk, and passes the metadata to the XML Interface, which organizes it into an XML file, and writes it to the hard disk. Finally, the Manager invokes the creation of all four panels. The Task Panel and the Filter Panel set up the task bar and the filters panel respectively. The Slider Panel sorts the data and creates the sliders. The Viewer Panel responds to user's interaction with the Slider Panel and displays the relevant information.

When the user shares data, the Manager writes the data

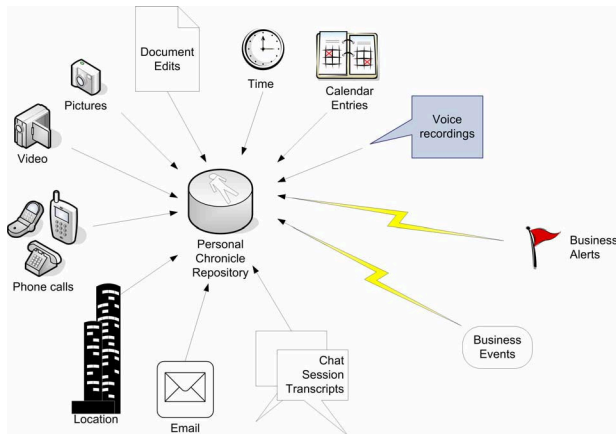


Figure 3 Data Sources for Chronicling

to the Shared Lotus Repository via the DB Interface. Similarly, when the user updates some information, such as changing an annotation, the Manager interacts with the DB Interface to write the data back to the Lotus DB. In this case, the created XML file is also updated via the XML Interface.

Encoding the data in an XML file enables use of PICASSO on a PC or handheld device without the need to have a wireless connection or Mobile Notes software running. For example, when PICASSO running on a PC creates an XML file, the user can synch that XML file along with associated multimedia files to a handheld device. When PICASSO runs on the handheld device, the Virtual Notes DB component can simulate the Mobile Notes DB, and simply pass on the synched XML file to the DB Interface, which treats it as a special case and forwards it to the XML Interface via the Manager. Similarly, if PICASSO running on the PC fails to connect to the Lotus DBs, the Virtual Notes DB can again be used to store and share information offline. To store/share information in an offline manner, the Manager interacts with the XML Interface and Virtual Notes DB (via the DB Interface), and creates a new XML file to store the updates and the shared data. When the device later has internet access or is synched to the PC running Lotus Notes, the XML file and the relevant multimedia files are uploaded. The Update Utility then can write the data to the appropriate Lotus DB.

Users can interact with PICASSO via Lotus Notes or Mobile Notes using a PC or a handheld device. As a standalone application, PICASSO interacts with the Lotus DB via the Lotus API. For platform independence and ease of porting, the PICASSO software module is written in J2ME Personal Profile using the Web Sphere Device Developer [41] environment. As a result, the same application runs on a handheld device as well as a PC without any modifications. Like any J2ME application, PICASSO needs an additional run time environment, such as J9 [41], to run on a handheld device.

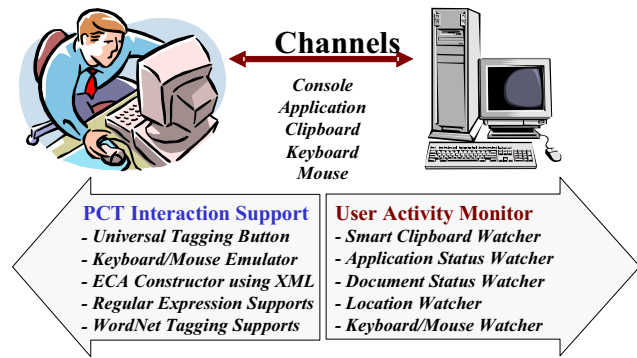


Figure 4 Event monitoring tools on the PC

4. Pervasive Chronicling

PICASSO supports the capture of various aspects of individual activity across many applications and with a variety of devices, as seen in Figure 3. We have focused on the PC and the PDA as the two primary user devices. On the PC, we have developed tools [26] to monitor every channel of communication between the computer and its user, and to thereby log all user interactions with the computer (see Figure 4). This includes every keyboard/mouse event, every email received, every instant message, all clipboard activity, any change in status of any user document, and the change in status of any user application.

Besides such passive logging of user actions, active annotation by the user can result in valuable metadata that is not available through any other source. Although PICASSO comprehensively captures events on a user's PC (and soon on the PDA), it also provides the user with the ability to mark or "tag" events that are particularly interesting.

To flag PC-based events at the time they are encountered, tagging is triggered by the use of the "T-button," a software button attached to the title bar of every window on the desktop (Figure 5a). For events encountered on the PDA, the device's "iTask" button (Figure 5b) triggers the tagging. In both cases, the Event Navigator (see Section 5 for details) is launched and brought to the foreground with the tagged event's screen capture and associated metadata. At this point, it is possible to confirm the tag and dismiss the Navigator, having simply flagged the event as significant. Alternatively, one may choose to annotate the event with some comments. This is accomplished by clicking on the screen capture itself and then typing. A yellow background unfolds behind the comment being entered. Clicking off the yellow background ends the editing and labels the entry with the user's initials. Clicking back into the yellow background reenables editing of the entry.



Figure 5

- a) Tagging (“T”) button on a PC (which appears on every Window opened on the PC)
- b) Tagging button on PDA (the physical iTask button on the iPAQ)

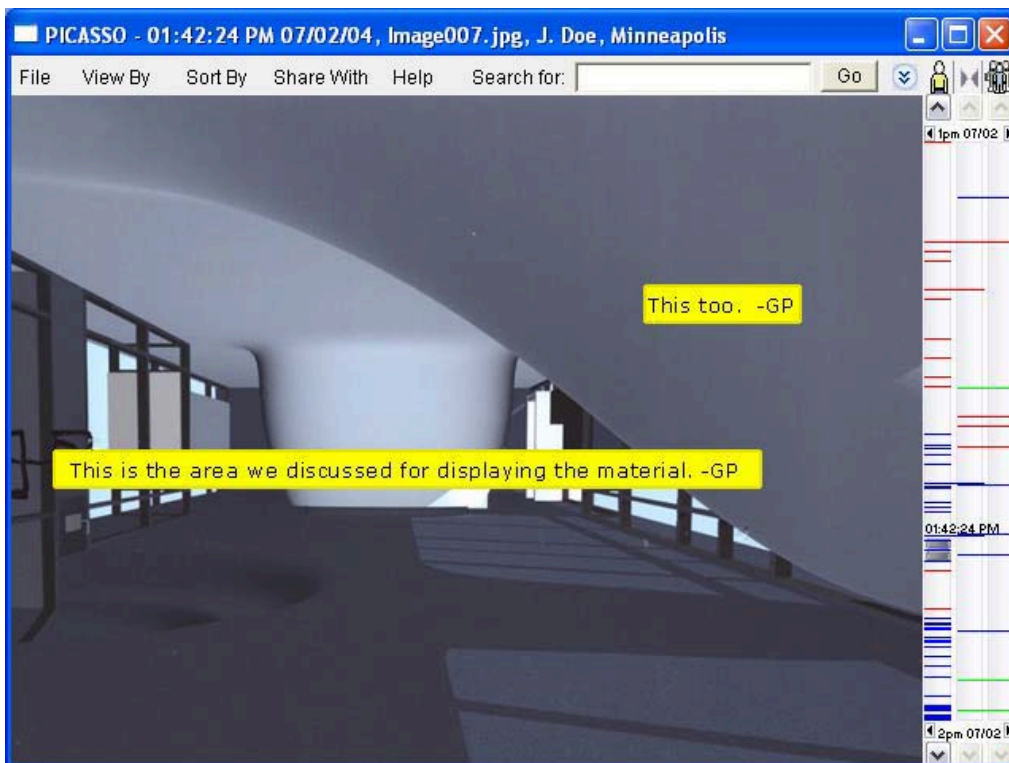


Figure 6 Annotation Interface in the Event Navigator. Slider Panel is seen on the right showing events for 1 hour. The image shown represents the selected event at 01:42:24 PM. User annotations have a yellow background.

Clicking on another area of the capture begins a separate annotation. Figure 6 depicts a photograph that has been tagged and annotated in the Event Navigator.

It is also possible to tag and annotate events which have *previously* occurred using the Event Navigator in precisely the same manner: When a past event of interest is displayed, one simply begins typing on it. Events that have been shared by other users can be annotated similarly; the appended initials of each annotator help to distinguish who wrote which comment. Words from all comments are extracted by PICASSO, which maps these to a common shared dictionary, WordNet, an electronic lexical database [13].

Any user activity, active or passive, results in an event being recorded in a local database – a Lotus Notes database in our current implementation. At a database level, each event has a type associated with it (email, instant messaging, clipboard, etc.) and has associated data including time, location, and other metadata associated with each specific application. User-tagged events have a special type, and all events have a notion of data access, which is personal by default, but can be extended through Lotus Notes mechanisms to include people in a group or public access.

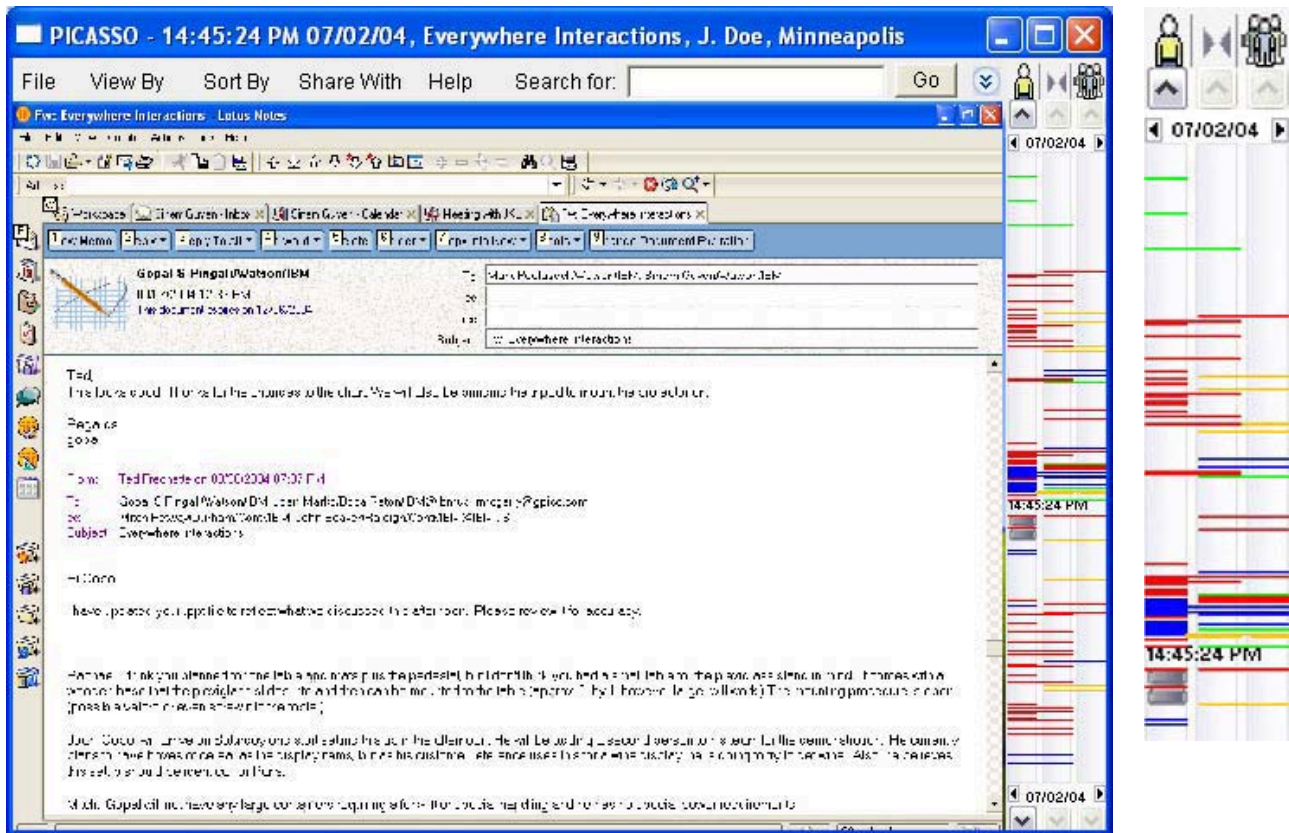


Figure 7 (a) Event Navigator running on a Desktop. The main window (View Panel) shows the selected event, a received e-mail in this case. Notice that the title bar displays the meta data about the selected event including the name and the location of the user. (b) Slider Panel on the right shows all events for a selected timeframe. The events include (left to right): 1) Personal Events, 2) All Shared Events (both incoming and outgoing) 3) Incoming Events shared by other people. Each line represents an event and its color represents the user's location at that time.

5. Chronicle Navigation

The PICASSO Event Navigator depicts each event captured by the Event Monitor as a mark on augmented timeline sliders. Different colored marks represent different locations where the events took place. The span of a timeline slider is adjustable by zooming, which makes it possible to view chronicles for intervals of a minute up to a year. Navigating through the marks on a timeline slider will display screen shots of the corresponding events, as well as associated metadata such as date and location.

The PICASSO Event Navigator can be used in one of the following ways:

- As a Desktop Navigator
- As a PDA Navigator
- As part of a Lotus Notes / Mobile Notes Calendar Entry

5.1. The Desktop Navigator

PICASSO can run as a stand alone application on a laptop or a desktop. Figure 7 shows the PICASSO Desktop Navigator in action.

5.1.1. Task Panel

The Task Panel consists of *File*, *SortBy*, *ViewBy*, *ShareWith* and *Help* pull-down menus in the window's task bar. The *SortBy* menu enables users to sort the chronicles by *time*, *date*, *author*, *document name*, and *location*. The *ViewBy* menu provides means of viewing chronicles at different zoom levels such as *by year*, *by month*, *by week*, *by day*, *by hour*, and so on. Finally, the *ShareWith* menu makes it possible to share personal chronicles with other users.

5.1.2. Filter Panel

If a simple keyword search is not adequate, clicking on the double arrow button (next to the search text box) on the task bar adds a Filter Panel directly below the taskbar.

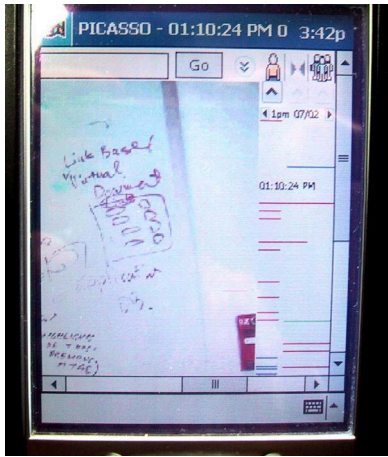


Figure 8 Event Navigator running on a PDA. (Photo from low-resolution iPAQ, soon to be replaced with HP 47xx with 640x480 resolution)

To accommodate more complex queries, multiple Filter Panels can be added. The timeline sliders are dynamically updated by user queries and filtering.

5.1.3. Slider Panel

The Slider Panel provides chronicle visualization using three dynamic timeline sliders [1]. The left slider is used to navigate a user's personal chronicles. The right slider navigates chronicles that were sent or shared by others. Finally, the middle slider navigates the merged set of a user's shared personal chronicles *and* those shared by others, creating a comprehensive group narrative. We believe that this middle slider might provide meaningful information that could otherwise be difficult to ascertain. For example, if two users shared events in the same time frame and the event marks have the same color, the users may have been in the same location. The up/down arrows at the top and bottom of each slider allow the users to scroll from one event to the next without having to drag the slider cursor. The small right/left arrows advance the visualization to the previous or subsequent time interval.

5.1.4. View Panel

The View Panel is where the user can visualize relevant information regarding the currently selected event. The information displayed consists of a screen shot of the event as well as metadata and annotations. The metadata is displayed on the Title Bar, and gets updated every time the user moves the cursor to another event. Annotations are layered on top of the screen shot. (The user can use the View Panel to add annotations to the currently selected event as described in Section 4.) Finally, double clicking on the View Panel launches the

relevant application and opens the file pointed to by the currently selected event.

5.2. The PDA Navigator

PICASSO is designed to run as an application on a handheld device. The User Interface and the interaction mechanisms are identical to that of the Desktop Navigator. Figure 8 shows the PICASSO PDA Navigator in action.

5.3. Lotus Notes / Mobile Notes

PICASSO also can run as an embedded applet as part of a Lotus calendar entry [9]. When a user opens up a calendar entry in Notes, PICASSO can populate it with all the events that occurred in the time frame defined by the entry. This can help the user track their activities as well as interactions that took place during that time frame without having to perform any search or filtering. Figure 9 depicts an *enhanced calendar entry*.

6. Sharing Chronicles within the Enterprise

Events are shared with others in the enterprise by means of the *ShareWith* pull-down menu found in the Event Navigator's task bar. Our current implementation is simple. A user's *ShareWith* pull-down is populated with the names of individual colleagues found in that user's Lotus Sametime instant messenger buddy list. In addition, groups of colleagues are pulled from the user's address book and added to the same list.

When a user encounters an event to be shared with one or more colleagues, she tags it using the windows-based "T-button" or the PDA's iTask button. This brings up the Event Navigator, as previously described in Section 4. The user here specifies the individual or group with which she wishes to share the event via the *ShareWith* pull-down. The event is then replicated to a shared Notes repository through the DB Interface, from which her colleagues can access it by means of their Event Navigators. The relevant XML file is also updated through the XML Interface.

7. Enterprise Usage Examples

In this section, we present an example of how the various navigation tools provided by PICASSO can be used effectively in an enterprise scenario. For this purpose, we refer to the motivating scenario presented in Section 2.

John is interested in retrieving some information from PICASSO on the way back from his visit to JKL. In particular, he is interested in retrieving all e-mails sent

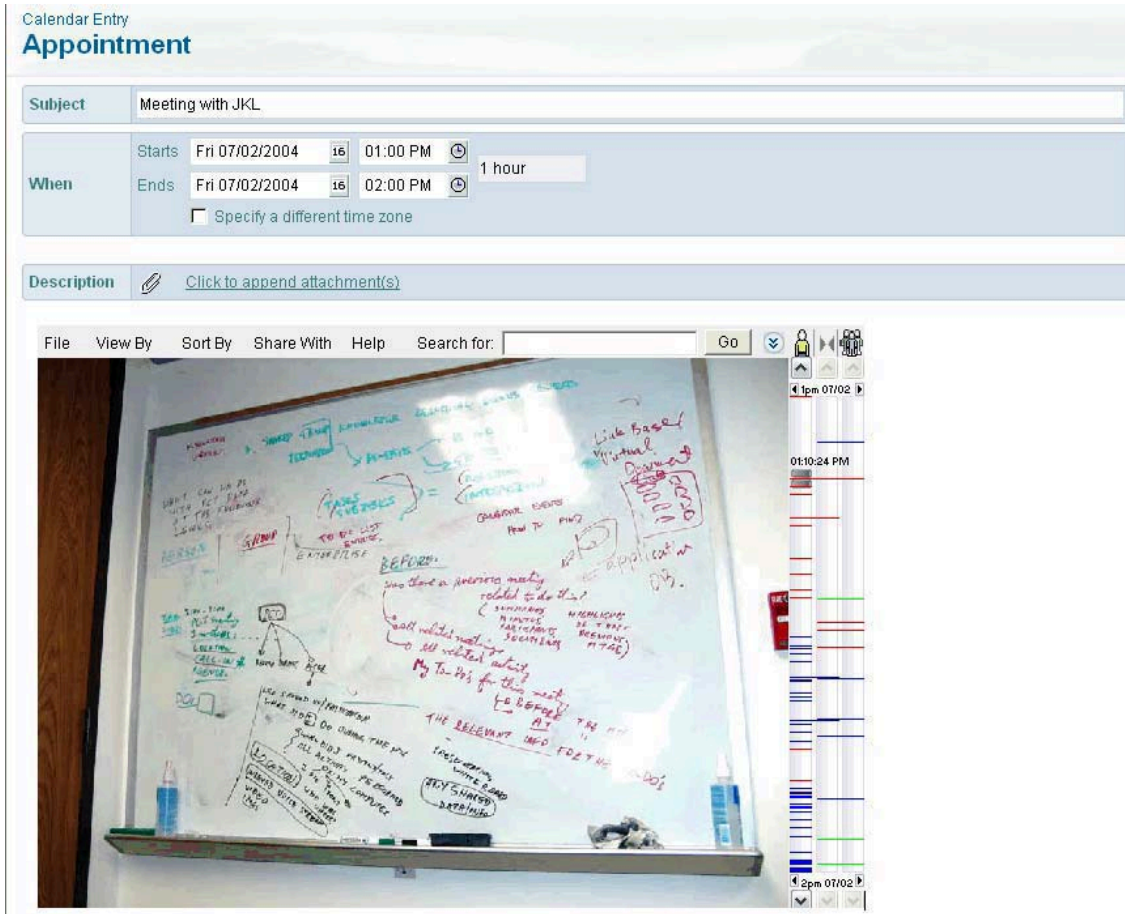


Figure 9 Event Navigator embedded in Calendar.



Figure 10 Searching for E-mails from Joe OR URLs shared by Joe. Notice the query being specified with the expanded Filter Panel.

by Joe, and all the URLs that were shared by Joe during his visit to JKL.

In order to get to the desired set of events, John first searches for all events that have the keyword "Joe". Since this yields too many results, he then decides to expand the Filter Panel and add further constraints. This time, he specifies that he is looking for *e-mails from Joe* and, having remembered that he actually tagged that e-mail, John also turns on the *tagged* filter.

Finally, John also decides to retrieve all the *URLs shared by Joe*. He then expands the Filter Panel one more time and specifies this additional constraint. Figure 10 depicts the expanded Filter Panel.

8. Conclusions & Future Work

In this paper we have presented the PICASSO system, giving a broad overview of its various aspects including i)

pervasive information chronicling and capture through various sources, ii) access of the information from different devices, iii) consistent and intuitive search and event navigation interfaces on both PCs and PDAs, and iv) the sharing and merging of chronicled information between different people in the organization. We believe a system like PICASSO can provide significant business execution and productivity benefits, and can influence the next generation of dynamic capture, archival, and exploitation of information in the enterprise.

We view the work presented here as still in the initial stages of realizing the vision presented in Section 1 on PICASSO in the enterprise. We have developed a system architecture for PICASSO and implemented the first versions of chronicling, navigation, and sharing tools. We have placed special emphasis on designing interfaces that are easy to use and have taken care to ensure that published information is searchable via powerful enterprise mining and search tools. Our next steps are to deploy these tools with groups of users in field trials, to conduct studies on their ease of use, effectiveness, and to measure actual and perceived benefits.

Finally, we need to further develop and study the impact of these tools on the enterprise including the development of people, role, and relationship proxies and their integration with business activity monitoring tools and enterprise applications.

9. Acknowledgements

The authors would like to acknowledge Claudio Pinhanez and Stefan Hild of IBM for their valuable suggestions at the early stages of this work, and for significantly influencing the direction of this work. The authors also acknowledge the work by Pilho Kim on the PCT system, which provided the launch pad for PICASSO.

10. References

[1] C. Ahlberg, C. Williamson, and B. Schneiderman, "Dynamic queries: an implementation and evaluation," *Proceedings of ACM Conference on Human Computer Interaction (CHI'92)*, Monterey, CA, May 1992.

[2] Y. Arens, C. Y. Chee, C. Hsu, and C. A. Knoblock, "Retrieving and Integrating Data From Multiple Information Sources," *International Journal of Cooperative Information Systems*, 1993.

[3] D. Barreau and B. Nardi, "Finding and reminding: File organization from the desktop," *SIGCHI Bulletin*, 27(3), July 1995.

[4] B. B. Bederson, A. Clamage, M. P. Czerwinski, and G. G. Robertson, "A fisheye calendar interface for PDAs: Providing

Overviews for Small Displays," *Proceedings of ACM Conference on Human Computer Interaction (CHI'03)*, Ft. Lauderdale, Florida, April 2003.

[5] S. Björk, J. Redström, P. Ljungstrand, and L. E. Holmquist, "POWERVIEW: Using information links and information views to navigate and visualize information on small displays," *In Proc. HUC2k (Handheld and Ubiquitous Computing 2000)*, Bristol, U.K.

[6] V. Bush, "As we may think," *Atlantic Monthly*, 176 (1), January 1945, pp. 641-649.

[7] I. Cole, "Human Aspects of Office Filing: Implications for the electronic office," *Proc. Human Factors Society -- 26th Annual Meeting*, 1982.

[8] S. Cousins, A. Paepcke, T. Winograd, E. Bier, and K. Pier, "The Digital Library Integrated Task Environment," *Technical Report SIDL-WP-1996-0049*, Stanford Digital Libraries Project, Palo Alto, 1996.

[9] "Domino Designer 6: A Developer's Handbook," *IBM Redbooks*, ISBN- 073842658X, 2002.

[10] P. Dourish, W.K. Edwards, A. LaMarca, and M. Salisbury, "Presto: An experimental architecture for fluid interactive document spaces," *ACM Transactions on Computer-Human Interaction*, 6(2), 1999, pp. 133-161.

[11] S. Dumais, E. Cutrell, JJ Cadiz, G. Hancke, R. Sarin, and D.C. Robbins, "Stuff I've Seen: A System for Personal Information Retrieval and Re-Use," *Proc. SIGIR '03 (Special Interest Group on Information Retrieval)*, Toronto, August 2003.

[12] T. Erickson, "The design and long-term use of a personal electronic notebook: A reflective analysis," *Proceedings of ACM Conference on Human Computer Interaction (CHI'96)*, 1996, pp. 11-18.

[13] C. Fellbaum, "WordNet: An Electronic Lexical Database," *MIT Press*, 1998.

[14] S. Fertig, E. Freeman, and D. Gelernter, "Finding and reminding reconsidered," *SIGCHI Bulletin*, 28(1), January 1996.

[15] E. T. Freeman and S. J. Fertig, "Lifestreams: Organizing your electronic life," *AAAI Fall Symposium: AI Applications in Knowledge Navigation and Retrieval*, 1995.

[16] J. Gemmell, G. Bell, R. Lueder, S. Drucker, and C. Wong, "MyLifeBits: Fulfilling the Memex Vision," *ACM Multimedia '02*, Juan Les Pins, France, December 1-6, 2002, pp. 235-238.

[17] D. K. Gifford, P. Jouvelot, M. Sheldon, and J. O'Toole, "Semantic file systems," *13th ACM Symposium on Operating Systems Principles*, October 1991.

[18] D. Goldberg, D. Nichols, B. Oki and D. Terry, "Using collaborative filtering to weave an information tapestry,"

Communications of the ACM, 35 (12), December 1992.

[19] W. Hill and J. Hollan, "History-enriched digital objects: Prototypes and Policy Issues," *The Information Society*, 10(2), 1994, pp. 139-145.

[20] D. Huynh, D. Karger, and D. Quan, "Haystack: A platform for creating, organizing and visualizing information using RDF," *Proc. WWW02*, 2002, Hawaii.

[21] R. Jain, "Media Vision: eChronicles," *IEEE Multimedia*, July 2003.

[22] R. Jain, "Structuralizing Multimedia Data," *IEEE Multimedia*, January 2001.

[23] R. Jain, P. Kim, and Z. Li, "Experiential Meeting System," *ACM ETP'03*, 2003, pp. 1-12.

[24] S. R. Jones and P.J. Thomas, "Empirical assessment of individuals' personal information management systems," *Behavior and Information Technology*, 16(3), 1997, pp. 158-160.

[25] S. Kaasten, S. Greenberg, and C. Edwards, "How people recognize previously seen WWW pages from titles, URLs and thumbnails," *Proceedings of ACM Conference on Human Computer Interaction (CHI'02)*, 2002, pp. 247-265.

[26] P. Kim, M. Podlaseck, and G. Pingali, "Personal Chronicling Tools for Enhancing Information Archival and Collaboration in Enterprises," To appear in *ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE 2004)*, New York, October 2004.

[27] M. Lansdale, "The psychology of personal information management," *Applied Ergonomics*, 19.1, 1988, pp. 55-66.

[28] H. Lieberman, "Letizia: An agent that assists web browsing," *Proceedings of the International Joint Conference on Artificial Intelligence*, Montreal, August 1995.

[29] T. Malone, "How do people organize their desks?" *ACM Transactions on Office Systems*, 1 (1), January 1993, pp. 99-112.

[30] M. Masoodian and D. Budd, "Visualization of Travel Itinerary Information on PDAs," *5th Australian User Interface Conference (AUIC2004)*, Dunedin, New Zealand, 2004.

[31] R. Mander, G. Salomon, and Y.Y. Wong, "The 'Pile' metaphor for supporting casual organization of information," *Proceedings of ACM Conference on Human Computer Interaction (CHI'92)*, 1992, pp. 627-634.

[32] J. McHugh, S. Abiteboul, R. Goldman, D. Quass, and J. Widom, "Lore: A database management system for semistructured data," *SIGMOD Record*, 26(3), September 1997, pp. 54-66.

[33] T. Nelson, "The Zigzag system," <http://xanadu.com/zigzag/tutorial/ZZwelcome.html>.

[34] T. Phelps and R. Wilensky, "Towards active, extensible, networked documents: multivalent architecture and applications," *Proceedings of ACM Conference on Digital Libraries (DL'96)*, Bethesda, MD, March 1996.

[35] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman, "Lifeline: Visualizing Personal Histories," *Proceedings of ACM Conference on Human Computer Interaction (CHI'96)*, Vancouver, 1996, pp. 221-227.

[36] J. Rekimoto, "Timescape: A time machine for the desktop environment," *Proceedings of ACM Conference on Human Computer Interaction (CHI'99)*, Pittsburgh, 1999.

[37] M. Ringel, E. Cutrell, S. Dumais, and E. Horvitz, "Milestones in time: The value of landmarks in retrieving information from personal stores," *Proceedings of Interact 2003*, Zurich, 2003.

[38] B. Rhodes and T. Starner, "Remembrance Agent: A continuously running automated information retrieval system," *Proceedings of the First International Conference on the Practical Application of Intelligent Systems and Multi Agent Technology*, London, April 1996.

[39] C. Silverstein, M. Henzinger, H. Marais, and M. Moricz, "Analysis of a very large Alta Vista query log," *SRC Technical Note 1998-014*, Palo Alto, October, 1998.

[40] S. Whitaker and C. Sidner, "Email overload: Exploring personal information management of email," *Proceedings of ACM Conference on Human Computer Interaction (CHI'96)*, New York, 1996, pp. 276-283.

[41] Websphere Studio Device Developer, <http://www-306.ibm.com/software/wireless/wsdd/>.