

**“Forget-me-not”
Intimate Computing in Support of
Human Memory**

Mik Lamming and Mike Flynn

Technical Report EPC-1994-103

*Published in:
Proceedings of FRIEND21,
'94 International Symposium on Next
Generation Human Interface,
2-4 February 1994,
Meguro Gajoen, Japan.*

Copyright © Institute for Personalized Information Environment 1994.

Rank Xerox Research Centre
Cambridge Laboratory
61 Regent Street
Cambridge CB2 1AB

Tel:+44 1223 341500
Fax:+44 1223 341510

“Forget-me-not” Intimate Computing in Support of Human Memory

Mik Lamming and Mike Flynn

RANK XEROX RESEARCH CENTRE
CAMBRIDGE, ENGLAND.

INTRODUCTION

At RXRC we have been trying to understand how anticipated developments in mobile computing will impact our customers in the 21st century. One opportunity we can see is to improve computer-based support for human memory — ironically a problem in office systems research that has almost been forgotten. Considering how often computers are presented as devices capable of “memorising” vast quantities of information, and performing difficult-to-memorise sequences of operations on our behalf, we might be surprised at how often they appear to have *increased* the load on our own memory.

The Forget-me-not project is an attempt to explore new ways in which mobile and ubiquitous technologies might help alleviate the increasing load. Forget-me-not is a memory aid designed to help with everyday memory problems: finding a lost document, remembering somebody’s name, recalling how to operate a piece of machinery. It exploits some well understood features of human episodic memory to provide alternative ways of retrieving information that was once known but has now been forgotten.

We start by introducing a model of computing in the 21st century which we call the “Intimate Computing” model and talk about some of the opportunities and problems we anticipate it will provoke. After a cursory introduction to the basics of human episodic memory, we describe the architecture and user interface of Forget-me-not. We end with a few preliminary conclusions drawn from our early experiences with the prototype.

INTIMATE COMPUTING

Our project team has been looking forward to technologies we anticipate will be readily available in the early 21st century. Mark Weiser, in his landmark article in *Scientific American*, described a world filled with computers, some small and portable, others embedded in the electronic devices that populate our homes, businesses and the world at large [Weiser, 1991]. He called this the “Ubiquitous Computing” model. Whilst we believe this vision is still several years away from being fully realised, we resonate with it, and can identify elements emerging already.

At RXRC we have confined our attention primarily to understanding the role for personal mobile computers: Personal Digital Assistants (PDAs). Our vision of the PDA is *not* confined to a device the size of a notebook, but includes something much, much smaller, perhaps the size of a watch or piece of jewellery — a device that can be *worn* and taken everywhere. Indeed, we expect to see elements of the PDA embedded in most current portable devices — cell phones, for example. These tiny PDAs will include wireless communication facilities allowing them to collaborate with other similar devices and nearby services. Our interest is in understanding the opportunities presented by a world in which we can rely on a large proportion of our users having a powerful computer with them *at all times*.

One possible consequence of wearing a computer is that it can be much more useful to you personally. Since it always accompanies you, and nobody else, it makes special sense to tailor its behaviour to your own special needs. Moreover, because it will be involved in many of your activities, it can become intimately familiar with them, and adapt to them like a personal assistant. We call this the *Intimate Computing Model*. The more the intimate computer knows about you, the greater its potential value to you. While personal computing provides you with access to its own working context — often a virtual desktop — intimate computing provides your computer with access to your *real* context. As we shall soon see, this is a crucial feature of Forget-me-not.

The Data Avalanche

Another important consequence of having a computer with us all the time is the ease with which information can be gathered. Data can be ‘beamed’ to us from nearby electronic devices or from the PDAs of other people we encounter, *where and when* we encounter them. And since the PDA is itself active, it can be commanded to pick up information automatically, as data become available, without the need to issue explicit instructions each time. The range of data we might profitably gather automatically is enormous: receipts from sales transactions, telephone numbers dialled, even overheads from presentations we attend. These are all things we might ask our PDA to capture, just in case the information might be useful someday. But there is an obvious snag.


Information overload is already a familiar problem we all have to contend with. When these extra data start flowing in, the task of organising them will become even more daunting. It is unlikely we will suddenly find the time or develop the inclination to index all of it. In any case, we often don’t know what information we have been given. For example, we often receive information that some person, or some system, thinks will be useful to us — we then find we can’t file it away because we have not read it, but nevertheless we are reluctant to throw it away, just in case it might be needed.

Much of this information snowstorm will be difficult for the computer to index automatically on our behalf based simply on its content. For example, sketches, photographs and video are notoriously difficult for computers to analyse. It looks like intimate computers, when used as automatic data gatherers, could make our lives much harder. If this is the prospect, then we still won’t be able to make good use of the information we encounter — and just won’t bother to collect it.

Context as a Retrieval Key

Fortunately, the very features of PDAs that contribute to information overload can also come to the rescue. The user’s context can itself provide a valuable key for indexing information automatically. A detail from a past event in which the user was involved might be difficult to recall, the name of a document, for example. But the context of the event can be easier to remember. For example, we may be able to recall: the place where the document was received, the people present when it was handed over, or the task being carried out at the time.

The idea that physical context can be a powerful cue for recall is not a new idea. Researchers in psychology have even developed theories about this sort of memory, called episodic or *autobiographical memory*. They observe that we naturally organise our memories for past events into episodes [Barsalou, 1988], and that the location of the episode, who was there, what was going on, and what happened before or after, are all strong cues for recall [Saywitz, Bornstein & Geiselman, 1992; Tulving, 1983]. Studies by Eldridge et al. [Eldridge, Barnard & Bekerian, 1994; Eldridge, Lamming & Flynn, 1992] have confirmed these findings, and moreover, have led us to believe that it might be possible to construct a prosthetic episodic memory device — a so-called *memory prosthesis*.



Using a prosthetic episodic memory device.

The basic idea is as follows: assume we could construct a device that accompanied the user everywhere, and which captured important data and context from his or her life. Furthermore, assume it would organise these data into a form that mimicked the episodic memory structures created naturally by the user. Needing to recall a detail from a past event, and armed with our device, the user could then draw upon his or her own, possibly fading, episodic memory, to locate similar episodes and data stored in the permanent memory of the device. In this way, the user could use the small things he or she could remember about the context of the event to retrieve the details that had been forgotten.

FORGET-ME-NOT

Some Requirements

Forget-me-not is our first attempt to create a working prototype of a portable episodic memory aid. Prior to starting design work, a number of small scale prototypes were constructed to investigate various aspects of the basic concept. [Eldridge, Lamming & Flynn, 1992; Lamming, 1991; Lamming & Newman, 1992; Newman, Eldridge & Lamming, 1991]. This led us to define a fairly extensive set of broad requirements, a few of which we will now revisit [Lamming, et al., 1994]. In designing Forget-me-not, we had the same aspirations as most other user interface designers. Any user interface needs to be easy to learn, easy to use, reliable, trustworthy and so forth. However these requirements take on special significance for a device designed to support human memory.

For example, we generally design user interfaces to be easy to learn and rarely quantify what this means in practice. But in the case of Forget-me-not, we have an obvious yardstick for measuring ease of use: it must be easier to remember how to use Forget-me-not than trying to remember the forgotten fact.

Availability is clearly a special issue, too. Memory failures do not *only* occur when the user is conveniently seated in front of a powerful workstation. We needed to design a user inter-

face to which the user can gain access anywhere and anytime. In consequence, we aimed to create a user interface that might plausibly fit on a wristwatch or other wearable object.

A tiny display, capable of fitting on a wristwatch, for example, mitigates in favour of a graphical user interface where every pixel can be used to convey useful information. But the meaning of any symbols used must be easy for the user to recognise or they will unduly contribute to the memory load.

As we shall see, these three requirements had the greatest influence over the design of Forget-me-not.

The Wearable Device



From the user's point of view, Forget-me-not has a very simple architecture. The Forget-me-not program and all of the data resides in a small portable device called a ParcTab.

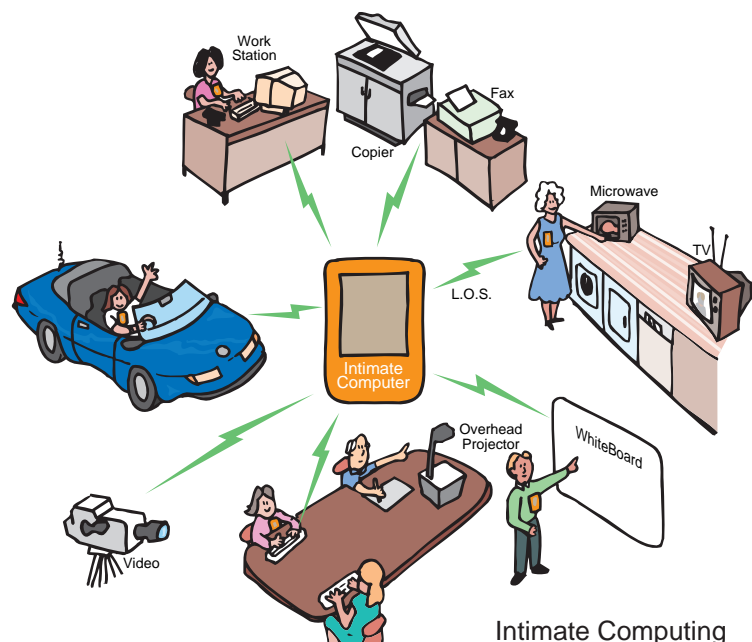
The ParcTab was built by Roy Want and his colleagues in the Computer Science Laboratory at Xerox PARC [Schilit, et al., 1993]. It is just light enough that it can be clipped to the belt or carried in a pocket. It has an LCD bitmap screen measuring 128 by 64 pixels which is touch-sensitive. There are three buttons, a tone generator,

and a bi-directional infra-red wireless transceiver with which the ParcTab can communicate with other ParcTabs and stationary equipment.

User's Conceptual Model

From the user's point of view, the ParcTab collects information about selected aspects of the user's activities, and organises these data into a personal biography. In our simple prototype, the user provides Forget-me-not with a list of the devices, or classes of device, from which data are to be collected. As the user encounters and interacts with each of these devices, Forget-me-not automatically gathers up information describing the device's name and location. It records salient details of each operation the user performs using the device, and, finally, it appends a time-stamp and stores the records away in its own memory. For example, if the user makes a phone call, the calling and called numbers would be recorded along with the start time and call duration.

When the user encounters another person, presumably also wearing a ParcTab, their Forget-me-not systems may exchange information and save that away, too. For example, it might record



that the user “encountered Bill Smith in Bill’s office”. Note, though, that Forget-me-not refers to *rules of revelation* previously set up by the user before disclosing the identity or any other personal details relating to its owner to any other person or device.

Information about each subject is stored in a personal biography which is private. A biography in Forget-me-not is represented as a sequence of encounters with people, places and devices. An *encounter* means different things in different situations. Each meaning is reasonably intuitive. Encountering a person means being in the same area, normally a room, for a short while. On the other hand, encountering the phone means *using* it and encountering a document means *processing* it in some way: receiving it, passing it on, reading it, or writing it.

Data Capture: the Current Status.

At the time of writing, our repertoire of captured activities is small. But we expect it to expand and want to make sure that our database schema and data capture mechanisms can be generalised to handle the wide variety of activities users might want to record about themselves. We have considered activities as diverse as paperwork, travel, watching TV and shopping. We have also looked at ways to capture data sources that might simply facilitate recall of past events, such as the news headlines and the weather. At present, though, we capture only the following information:

Personal Location and Encounters with Others

Forget-me-not can infer the current location of the user by asking the nearest room-based wireless transceiver to report its location. Another mechanism we use relies on data collected from Active Badges worn by each of our users [Want, Hopper, Falcao & Gibbons, 1992].

In principle, as we mentioned above, two or more Forget-me-nots in the same room can hail each other, exchange identity details and thus keep track of whom the user encounters. In practise, and once again for simplicity and ease of implementation, we centralise location logs and deduce when encounters occur. We take special care only to reveal information the user would know anyway — as a result of being in the same room, for example.

Workstation Activities

At EuroPARC many users spend several hours each day using their workstations. We have had some success instrumenting the workstation to record which programs are running, and which files are read or updated. Software to summarise this huge amount of data is the subject of ongoing work [Flynn, 1991]. At present, we have focused on e-mail transactions, this being a typical and pervasive workstation activity. Both incoming and outgoing mail items are captured in the biography.

File Exchange and Printing

Forget-me-not also has facilities for exchanging electronic files with other Forget-me-not systems. Thus it is relatively simple to keep track of when, where and from whom documents were received.

Telephone Calls

Like most modern PBX installations, our telephone exchange has a call logging facility. We collect the data emitted on the logging channel and, subject to the privacy rules currently in force, add them to the database. At the present time, the phone numbers of incoming external calls are not recorded.

The User Interface

In parallel with automatic data collection, Forget-me-not can search and display all of the data that have been collected. The user interface for invoking these functions is in an early stage of development, and we expect it to evolve considerably as we learn more about usage patterns.

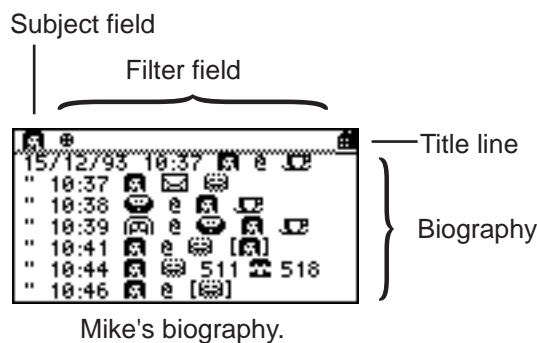
Screen Layout

The screen is divided into two main areas. The top line is called the title line, and the area below is called the biography. The *biography* is divided into seven lines, each of which represents a single event in which the subject was involved. Events are presented in time order. Each line contains one or more icons representing people, places, things or actions together with a time stamp to say when the event occurred.

The *title* line contains a *subject* field and a *filter* field. As its name suggests, the subject field contains the name of the subject of the biography currently being viewed. Icons are placed in the filter field by the user to specify which events should be included in the biography. Below we show a typical screen display from Forget-me-not.

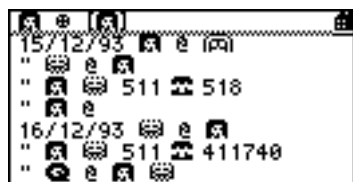
Icons are designed (or chosen) by the user to represent familiar people, places or objects. The subject in our example is Mike (👤), and his icon can be seen in the title line at the extreme top-left. The filter field is blank.

Forget-me-not shows Mike's biography for the 15th of December 1993 starting at 10:37 when Mike went to the Kitchen (👤 📍 🚪). Shortly after that he received some mail from Grouch (👤 📧 📧). At 10:38 Peter (👤) came to the Kitchen, followed at 10:39 by Marge (👤). At 10:41 Mike went back to his office where he encountered Grouch (👤 📍 📍 📍). At 10:44 one of them used telephone extension 511 to call extension 518 (👤 📞 511 📞 518), and at 10:46 Mike went to Grouch's office.



Command Language

Forget-me-not has only five commands: (1) scroll up, (2) scroll down, (3) filter, (4) inspect, and (5) change subject. The first two commands scroll the biography forwards or backwards and are activated by the top and bottom Tab buttons.



Mike's biography, filtered to show only those events that occurred in his office.

A biography often contains vast amounts of information. This would be very difficult to browse using just the scroll buttons, so Forget-me-not provides a third command to *filter* a biography to reveal just those events with particular properties of interest. To filter a biography, the user indicates which icons must appear in an event in order for it to be of interest and thus be displayed. This is done by dragging the icon in question to the filter field. In response, Forget-me-not re-formats the display to show just those events involving the chosen icon. Several icons can be placed in the filter field indicating that all must be present in each event. Here we show Mike's biography filtered to show only those events that occurred in his office.

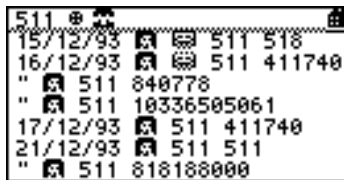
Since icons in the filter field must be present in every line of the biography displayed beneath, they are factored out to save space.

Icons are a concise way of representing the object or activity they stand for. To find out more details about the object behind an icon, Forget-me-not provides a fourth command, *inspect*. The user clicks on an object's icon, wherever it may appear. An inspector, tailored for the selected object, takes over the display.



An inspector for a person.

You can switch to look at the biography of another subject by using the fifth command: *change subject*. To do this, you drag an icon to the subject field. The result is the new subject's biography *viewed through the rules of revelation he or she has set up about you*. So for example, although you may be viewing Bob's biography, it will usually only display Bob's encounters with you.



Biography for extension 511.

Now it must be clear that *every* object in the Forget-me-not system has its own biography — not just the user and other people. This is one of the most powerful features of Forget-me-not. Users can 'teleport' long distances through their own biography by following the biography of another subject known to be encountered relatively infrequently. Here we show the biography for Mike's telephone filtered to show the start time of each phone call.

Example 1: Searching for a document.

Let's join Mike (icon) who is trying to locate a document he passed to Marge (icon) during a meeting involving Grouch (icon), Peter (icon) and Professor (icon). Mike finds occurrences of the icons for Grouch, Peter and Professor and moves them to the filter line. Forget-me-not displays all the meetings involving at least those people. Mike also remembers that the meeting was held in his office (icon) and moves its icon to the filter line too. Forget-me-not further restricts the set of displayed events to match the pattern field. Scrolling through these events, Mike is able to identify the event where he gave a document to Marge (icon → icon). He might then select the document (icon) and print it.



Meetings with Grouch, Peter, and Professor in Mike's office.

Example 2: Searching the document's biography

When Marge returned the document, Mike gave it to someone else but can't remember who. He decides to look at the document's own biography by moving its icon to the subject line.



Biography of a document.

The biography for the document is now presented and shows when it was edited (icon) or read (icon). Although the document clearly existed before Mike initially received it, the biography shows no details prior to that time. Mike can scroll through the biography and see what happened to it. He sees that on December 1st, he passed the document to Marge. Mike sees no further events listed until Marge passed it back to him the next day. He sees that, after reading, it he passed it on to Smiley (icon).

CONCLUSION

A great deal of the technology and infrastructure upon which Forget-me-not depends is not yet commercially available. Necessarily then, Forget-me-not has to simulate some aspects of the environment we anticipate will exist in the near future. We are fortunate in that one feature of RXRC that makes this sort of exercise tractable is the environment itself. RXRC comprises a single, relatively small, four-storey building. As a consequence, we have been able to provide cellular wireless communications facilities almost everywhere in the building, in addition to the LAN and switched audio-video networks that already cover the building [Buxton & Moran, 1990].

Nevertheless, we have had to make a number of compromises. For example, the portable device we actually used, called a ParcTab is somewhat larger than the wristwatch sized device we would like to use. Also, much of the functionality we would expect to reside inside the Tab, lives in a non-mobile server — in reality the user's personal workstation on the desk in his or her office. In addition, for simplicity and ease of implementation, user biographies are stored in a central shared database, rather than on board the ParcTab or in the private workstation. But in designing the user interface for Forget-me-not, we have tried to keep in mind the need to move towards a *personal, private* but *distributed* solution of the kind outlined in the introduction.

Forget-me-not has only been working for a few months. Consequently we have not yet started a formal evaluation. Nevertheless we are already forming some useful conclusions.

It is clear that privacy and security will be a major concern for users of Forget-me-not. At this stage, we have not devoted very much energy to finding technical solutions to even the most obvious problems. Instead we have tried to tease out the real issues, which we recognise are complex. We expect that the solution will involve a mixture of technology and legislation.

Forget-me-not is a special kind of hypertext system providing many implicit links for traversing what would otherwise be an impossibly large database. One well-known feature of Hypertext is the ease with which users can get disoriented and lost [Nielsen, 1990]. However a feature of Forget-me-not that is quite encouraging is that getting lost seems to be less of a problem, since the user is almost guaranteed to be familiar with the topology of his or her own biography.

On the engineering front, two obvious issues confront us. It is imperative that systems like Forget-me-not are extremely reliable. Nothing would be more frustrating than to find that the data you had expected to be captured, were lost or inaccessible because of a system failure. That would be memorable indeed. However, it isn't clear to us that current distributed systems technology can always meet these requirements at a reasonable cost.

The other systems issue is speed. Biographies may become extremely large, and thus searching them may become very slow. At present, our prototype uses a commercial database to store the raw activity data and performance is less than adequate for events more than a few months in the past. Indeed, it may be the case that systems like Forget-me-not need to be based upon databases with a better understanding of temporal logic. For the present, we seek solace in our intuition that most queries concern target events that occurred in the near past.

Lastly, and perhaps most importantly, we are now starting to see how challenging is the task of reducing the raw data down to human-recognisable episodes. The problem is made more complex by the apparent, but as yet unproven, need to provide descriptions at many different levels. The description for a day's worth of activities is quite different from the description for a year's worth. These problems will be the focus for the next stage of our research.

REFERENCES

- Barsalou, L. W. (1988). *The content and organization of autobiographical memories*. In U. Neisser & E. Winograd (Eds.), *Remembering reconsidered: Ecological and traditional approaches to the study of memory*. pp. 193-243. Cambridge: Cambridge University Press.
- Buxton, W., & Moran, T. (1990). *EuroPARC's Integrated interactive intermedia facility (iiif): Early experiences*. In IFIP WG8.4 Conference on Multi-user interfaces and applications. pp. 11-34. Herakleion, Crete: North-Holland.
- Eldridge, M., Barnard, P. & Bekerian, D. (1994). *Autobiographical memory and daily schemas at work*. Memory (to appear March '94).
- Eldridge, M., Lamming, M. & Flynn, M. (1992). *Does a video diary help recall?* In A. Monk, D. Diaper, & M. D. Harrison (Ed.), *People and Computers VII*, VII pp. 257-269. York: Cambridge University Press.
- Flynn, M. (1991). *Workstation Activity Monitoring Status Report*. Rank Xerox EuroPARC.
- Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., Robinson, P. & Sellen, A. (1994). *The Design of a Human Memory Prosthesis*. Computer Journal (to appear).
- Lamming, M. G. (1991). *NoTime - A Tool for Notetakers*. EuroPARC Technical Report. 91-135.
- Lamming, M. G. & Newman, W. M. (1992). *Activity-based Information Retrieval: Technology in Support of Personal Memory*. In F. H. Vogt. (Ed.), *Information Processing 92, Personal Computers and Intelligent Systems*, Vol. 3 pp. 68-81. Madrid: Elsevier.
- Newman, W., Eldridge, M. & Lamming, M. (1991). *Pepys: Generating Autobiographies by Automatic Tracking*. In Proceedings of the second European conference on computer supported cooperative work. Amsterdam: Kluwer Academic Publishers.
- Nielsen, J. (1990). *Hypertext and Hypermedia*. San Diego, CA: Academic Press.
- Saywitz, K., Bornstein, G. & Geiselman, E. (1992). *Effects of cognitive interviewing and practice on children's recall performance*. Journal of Applied Psychology, 77(5), 3-15.
- Schilit, B. N., Adams, N., Gold, R., Tso, M. & Want, R. (1993). *The ParcTab Mobile Computing System*. In Proceedings of the Fourth Workshop on Workstation Operating Systems, Napa, CA: IEEE.
- Tulving, E. (1983). *Elements of Episodic Memory*. Oxford University Press.
- Want, R., Hopper, A., Falcao, V. & Gibbons, J. (1992). *The active badge location system*. ACM Transactions on Office Information Systems, 10(1), 91-102.
- Weiser, M. (1991). *The Computer for the 21st Century*. Scientific American(September), 66-75.