

Using Symbiotic Displays to View Sensitive Information in Public

Stefan Berger, Rick Kjeldsen, Chandra Narayanaswami,
Claudio Pinhanez, Mark Podlaseck, Mandayam Raghunath
IBM Research, T.J. Watson
berger, fcmk, chandras, pinhanez, podlaseck, mtr@us.ibm.com

Abstract

Symbiotic displays are intelligent network devices that can be enlisted on demand by mobile devices, offering higher resolutions and larger viewing areas. Users of small mobile devices can consequently map various kinds of information to the displays that are most appropriate for comfortably viewing them. The availability of these displays in public areas, where passersby may glimpse another's sensitive information, introduces a privacy issue. To address this, we suggest a technique that combines the blurring of sensitive words on the public symbiotic display with a user interface that allows these words to be read on the display of a personal mobile device. We describe three ways in which such interactions can be enabled and discuss the architectural trade-offs associated with these approaches. We also describe a prototype implementation of an email reading application utilizing an IBM WatchPad [10] and an Everywhere Display projector [12]. This prototype was demonstrated at PerCom'04.

1. Introduction

Personal mobile computing devices offer several useful features such as the ability to communicate with others while on the move, access information relevant to the task at hand, and help manage one's time in a better fashion. This trend can be expected to continue due to increases in storage capacity, CPU speed and other technological advances. However, in another dimension that relates to people, we see a flat or even a falling trend with age in factors such as visual acuity, willingness to carry heavy and large devices, and ability to operate small devices. As a result, improving the user experience on small devices is central in making further progress in personal mobile computing. In this paper we consider one specific aspect of the user experience, namely accessing visual information from a mobile device. The sizes of mobile computing devices often restrict the physical dimensions of displays that can be supported, and limitations of human visual acuity essentially limit the amount of useful information that can be presented on the display.

In a recent paper [13], three of the authors described a

futuristic scenario where mobile computers establish symbiotic relationships with stationary devices in the environment in order to offer users a combination of best attributes of both systems, namely personalization of content due to mobile computers as well as large, easy to read, high quality displays. Environmental displays become intelligent network objects that offer their services just like today's network printers. Mobile computers discover such displays, communicate with the displays to ascertain display characteristics, and securely transmit information.

The primary focus of our paper is on one of the technical challenges associated with the symbiotic mode of operation, namely the issues that can arise when information is presented on an environmental display that is potentially visible to others in the vicinity. For instance, many of us have been in presentations where a personal instant message intended for the speaker flashes on the screen during a presentation and is visible to the audience. In order for mobile users to leverage large environmental displays for a better user experience, this and similar problems of information privacy need to be solved.

In particular, we address here the problem of viewing personal email on large, public environmental displays. We chose email as the application to help focus and motivate our research efforts for a variety of reasons. Email is generally considered to be extremely personal in nature. Users want to be able to access their email while they are mobile. There is a wide variation in the types of email messages and attachments that one exchanges and the display characteristics of mobile devices are often inadequate to read and fully respond to email. In addition to the display problem, there are several systems issues such as security, network bandwidth, and power management on the mobile device, that need also to be addressed as part of a comprehensive solution.

We built a demonstration of such a scenario showing a symbiotic relationship between one specific mobile device, the IBM WatchPad prototype [10] and an Everywhere Display projector [12] (ED-projector). The demonstration scenario involved a user reading email received on his watch by displaying it on displays created by the projector. We showed this demonstration at PerCom'04 [13] to gauge the receptivity to the concept. The response to the demo served to confirm our

Table 1. Typical resolution of displays and maximum possible resolution according to limitations on human visual acuity, using 20/20 human visual acuity, i.e., 1 minute of arc (adapted from [14]).

display type	user distance (inches)	typical width (inches)	typical width (pixels)	typical resolution (dpi)	maximum resolution (dpi)	maximum width (pixels)
cell phone panel	10	1	100	100	350	350
PDA display	12	2	300	150	291	582
laptop display	16	10	1200	120	218	2180
desktop monitor	20	15	2000	133	175	2625
meeting room screen	230	80	1200	15	15	1200

confidence that the information display problem is an important one to solve, as well as to affirm the symbiotic display approach we were taking to solve this problem. Based on this feedback we have continued to enhance the different aspects of this interaction paradigm, both from a user experience perspective as well as the underlying software technologies used to create the demonstration. We describe below how we created the initial demonstration that we showed at PerCom'04, and the progress since then. Finally, we outline research issues we are currently addressing and other possible applications for the proposed technique.

2. Symbiotic Displays

Many researchers have explored the idea of using environmental displays to supplement the display of handheld devices [10, 16, 18]. The key observation leading to this approach is that every display has intrinsic limitations that create usage barriers. For example, human visual acuity imposes an upper bound on display resolution. Given that even people with perfect vision cannot resolve details smaller than one minute of visual arc angle, increasing display resolution beyond that point does not contribute significantly to improvements in the readability of shown information. Table 1 (adapted from [14]) shows the typical resolution of current displays and the maximum “meaningful” resolution as a function of visual acuity and their size.

It is easy to see from Table 1 that technological advances will not improve cell phone displays to a state where reading large amount textual content becomes an easy task. For instance, if we consider the maximum possible useful resolution of 350 pixels (reading 10 inches away from the eye), it is hard to fit more than 3-4 words per line. Also, on the cell display, in the best of scenarios it is possible to fit only 10 lines in a 500 pixel long display, resulting in a reading process involving more than 10 scroll-down operations for the average 400-word email message. Moreover, a lot of the layout of the content is lost when the text is fit onto the small screen, especially when it contains tables, bullet lists, etc. Besides

that, content such as spread-sheets and large images add an extra layer of difficulty for the small screens of handhelds.

If a mobile device can opportunistically leverage large displays in the environment to display content to the user, the inherent limitations of the mobile device do not stand in the way of a richer user experience. Of course, as noted by Raghunath et al. [14], establishing the infrastructure required for display symbiosis is not an easy task itself. As it was the case for network-accessible printers, it is likely that it will take many years until a suitable infrastructure is in place in many areas and symbiosis becomes a seamless and natural interaction mechanism. Our research efforts are currently directed towards creating the underlying technologies to make display symbiosis happen in a controlled laboratory environment. We hope to extend these technologies to larger and more general environments over time.

An extremely important property that is required for such symbiotic interactions is *display amnesia*. When the user notifies the display that she has finished viewing the content, the display must destroy any state information or copies of displayed content that it obtained on behalf of the user. In a limited environment users may trust that all displays obey this property. However for more widespread symbiotic interactions between handhelds and public displays we may need a scheme where displays are certified by a certification authority as being amnesic. Mobile devices are expected to check these certificates as part of the initialization process.

3. Sensitive Content on Public Displays

In addition to the underlying technologies for display symbiosis we also need to address the issue of hiding sensitive content from others who may be in the vicinity of the large display. While larger displays such as computer monitors and projection displays offer a better user experience for the mobile user, they also come with the risk of revealing sensitive information. This is an area of concern especially when viewing private content such as electronic mail messages.



Figure 1. Different sensitive levels for the same example text: a) Maximum level of sensitivity; b) Medium level of sensitivity; and c) Full text.

The approach we take to solving this problem is to obscure the portions of the content that need to be hidden from others on the large display. We simply replace such sensitive portions of the document with blurred blobs. When the user wants to know what a blurred blob on the environmental display represents, the user selects the blob, and the words corresponding to the blob are shown on the display of the mobile device. We expect that the small display of the mobile device is unlikely to be seen by others in the vicinity.

Blackening out portions of documents as a security technique has been used by governments, businesses, and privately for a long time. The method was extensively used during World War II in the letters sent home by soldiers in the field, and it is still routinely used by the US government to partially declassify documents. It is considered a simple and effective way to provide widespread access to information without compromising the security of specific portions of documents, in spite of its occasional vulnerability to simple code-breaking methods [3]. Our approach follows some of the ideas suggested in [16] in the context of handling sensitive information in websites.

Figure 1 shows a portion of an email message used in our demonstration with three different blurring levels. In this case, the selection of the words being blurred was performed manually. Notice that the majority of the blurred words correspond to names, dates, locals, and numbers. Essentially any content that is sent to the user's mobile device needs to be marked with sensitivity levels. In the case of textual content, we expect that sensitivity level is an attribute associated with portions of the text

just like font style, font size, etc. If the mobile user is viewing content on a large display in a public setting the mobile device should apply a filter that blurs all words that are above a sensitivity threshold. If the display is not as large and is being viewed in a relatively secluded setting the threshold may be set at a higher level so that more of the content is directly visible on the environmental display.

In general practice, marking content with sensitivity levels may be difficult. Ross et al. [16] suggest several ways in which one may automatically process text to identify portions that one may want to keep private. One may also scan email messages for dates, email addresses, names, price values, etc., using mechanisms such as those described in [11, 17]. One further level of complexity could be added by examining the text grammar structure. Grover et al. [5] have demonstrated LT TTT, a tokenization tool that parses text into sentences, sentence structures, and tokens defined by rules corresponding to dates, bibliographical entries, and medical terms.

Another possible technique for determining candidates for suppression is to select words that have very low usage in everyday language. The assumption here is that less frequent words are likely to convey more information, and therefore their suppression will increase security. Similarly, corpuses of words that should be suppressed in specific knowledge areas could be created by corporations, governments, or through open source collaborative methods. In this scenario, the mail server would access a specific corpus as part of mail processing and create the privacy level annotations. It could also be the case that suppressing guidelines would be developed

by corporations and governmental agencies, establishing rules not only on words but also on what parts of an email are more sensitive. Moreover, corporate guidelines could include rules forcing potentially offensive words to be hidden, contributing to create a harassment-free environment. In some cases even the sender could be asked to mark up sensitive material if this type of viewing becomes widespread.

Finally, it is possible to simplify the reading process by blurring words and references that are commonly used by the sender and the recipient, but have a low usage in the general population. This can be performed automatically by analyzing emails previously exchanged by the sender and receiver and look for unusual words frequently used by both. For example, in Figure 1 the sender and receiver may often exchange emails concerning treasure hunting, so the word “treasure” on the top center is frequently used between them, and therefore blurred. We conjecture here that, in such conditions, guessing the meaning of those words is easy for the email receiver, while it is still hard for other people.

4. Architectures for Symbiotic Displays

In a typical symbiosis scenario, mobile users establish a network connection between their mobile device and the display, sending it content that they wish to see. The display shows the content and the user may interact with what they see either directly using the input mechanisms supported by the display or using buttons or other input mechanisms on the mobile device.

For display symbiosis to work, displays need to become intelligent network objects that have some or all of the following attributes:

- Displays are directly connected to the wired network infrastructure and are addressable on the network.
- Displays support several data formats such as ASCII, HTML, PDF, Postscript, PowerPoint, JPEG, GIF, MPEG, Flash, etc., and are able to negotiate with other network devices, in a manner similar to how browsers express “accept” tags to web servers.
- Displays support direct user interaction via keyboards, mice, touch sensitive screens, or other forms of gesture recognition to enable users to perform simple operations such as scrolling, pausing video, etc.
- Displays support downloadable code such as JavaScript, Java applets etc., for richer user interactivity.
- Displays preferably support a short-range wireless network interface to communicate with mobile devices, and to make it easy for mobile devices to discover those displays that are in the immediate vicinity.




Figure 2. Configuration of our environment.

Another architectural consideration is that many mobile devices are supported by network-connected computers that either provide content to the mobile device or at least maintain a copy of the content on the mobile device. Instead of directly transferring the content from the mobile device to the display over the wireless network, it may be better if the display can fetch the content directly from the network copy. Wireless communication usually has a high energy cost, and redirecting the display to go to the wired network instead of the wireless link may save precious battery resources on the mobile device. Further, fetching content from a wired network may be faster, than transferring large amounts of data over a slow wireless link.

In order to allow the display to fetch content over the wired network on behalf of a user, we need to address a few issues. The mobile device must first present to the display the address of the content server and a reference to the copy of the content the user wishes to view. The display must be able to reach the server, i.e., the server should either be on the publicly accessible Internet or the display should be within the same administrative domain as the server. The server should be configured to serve up the copy. And we also need the mobile device to authorize the display to fetch the content from the server on behalf of the user, perhaps by sending a credential token to the display to be presented to the server.

4.1 System Architecture Trade-offs

The preferred high level system architecture of the symbiotic email reading system is shown in Figure 2. The user carries a mobile device that has both a long-range wireless interface as well as a short-range wireless interface. The long-range wireless connection is used by the mobile device to receive messages from the mail server and interact with the email server. The short-range wireless interface is used for communication with displays in the immediate vicinity of the mobile device.

Depending on whether the mobile device directly

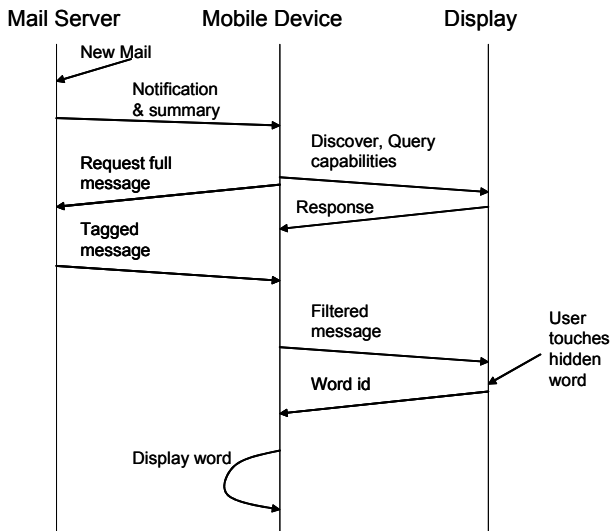


Figure 3. Mobile receives full email message.

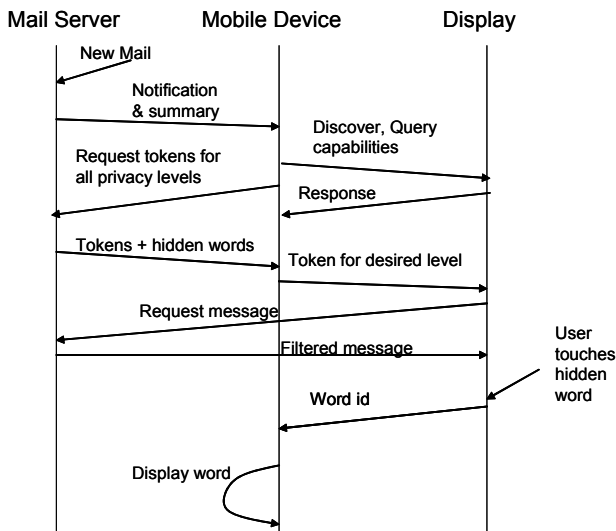


Figure 4. Mobile receives tokens and hidden words.

transfers the content to the display or requests the display to fetch the content from a server we follow the message flow shown in Figure 3 or Figure 4. In Figure 3, when email arrives at the user's email server, the email server sends the mobile device a notification along with a brief summary of the message. This communication occurs over the long-range wireless network via the wireless gateway (which is omitted from the figure for brevity). The user notices the incoming message and decides that she would like to read it and requests that the entire message be sent to the mobile. She also looks for an environmental display to read the message. As described

above, the handshake between the display and the mobile includes an exchange of information that conveys the capabilities of the display, acceptable file formats and other necessary parameters.

The full message arrives at the mobile device after some time and, as described above, portions of the message are suitably tagged with the levels of sensitivity. Depending on the type of display the user has selected, she indicates to the mobile the sensitivity level at which she would like to view the message. The user's mobile device uses the tags in the message and the desired level of sensitivity to filter the message. All words that are at a higher level of sensitivity compared to the sensitivity level associated with the particular display interaction are removed from the message and replaced with blobs. Each blob is associated with a word id.

Once the message is filtered in this manner, the mobile device converts the message into one of the formats acceptable to the display. This filtered message is sent to the display. When the user reads the message on the large display, the user can select a blurred blob. The display recognizes the selection and sends the associated word id to the mobile device. The mobile device then looks up the word by its id and displays it on the screen of the mobile device for the user to read on the small screen. In order to support this mode of interaction, the content the mobile device may need to include a piece of code that sends the word-id to the mobile, in a manner similar to current web pages that include JavaScript code that which reacts to mouse events on a page by displaying balloon help.

Figure 4 illustrates the case where the mobile does not receive the entire message over its long-range wireless network. This approach is suitable for mobile devices which have a low bandwidth long-range connection such as cell phones or two-way pagers. In this case, the mobile device requests a set of tokens from the email server, one token for each level of sensitivity. The server creates ephemeral tokens and returns them to the mobile device along with the set of sensitive words in the email message. Subsequently, the user of the mobile device decides to view the message on a particular display at a chosen privacy level. The mobile device then sends the corresponding token and the URL of the email server to the display. The display connects to the email server and requests the message providing the token as a credential authorizing the request. The email server filters the message, replacing the sensitive words with blobs and tagging the words with appropriate word ids, and sends the filtered message in a format that is acceptable to the display. When the user interacts with the displayed content, the word ids are sent to the mobile device and the appropriate words are displayed on the mobile device as in Figure 4. As noted above, this approach may result in a much faster user response time since the entire message is not sent over the relatively slow long-range wireless

network, but this approach is feasible only when the display's network is more capable and the display has permissions to interact directly with the email server.

A third architectural model can be used to support extremely resource-constrained mobile devices or mobile devices when the long-range wireless coverage is spotty. If the mobile device is out of coverage range when the email message arrives, the email server may be unable to notify the mobile device. The notification should be held by the network and delivered when the mobile is back in coverage. This function is preferably implemented in a user specific proxy machine in the network that monitors the connection state of the user's mobile device. The proxy machine is also a good place to implement the code that automatically process the user's email and adds the sensitivity level tags. The proxy machine can also orchestrate the symbiotic display interaction when the mobile device is highly constrained. For instance, the proxy machine can receive user blob selections from the display and send the words one at a time to the mobile device. The proxy, the mobile device, and the email server have a mutual trust relationship; the display is the only one that needs to be brought into this group, which is done using the ephemeral tokens in a manner similar to Figure 4 once the mobile device identifies the display to the proxy.

4.2 Eliminating the Short-Range Interface

The primary function of the short-range interface on the mobile device is to enable the user to easily locate and connect to displays in the immediate vicinity. As display symbiosis becomes more widespread, we expect that there will be several intelligent displays that are deployed in a building or campus, and selecting the one that is nearby may be not be that easy. Using a short range wireless interface is a simple way to narrow down the display candidates to those that are nearby, without having to solve hard problems such as indoor location tracking. In addition to solving the display identification problem, short-range wireless interfaces have several other advantages in the area of security as discussed by [2].

Nevertheless, mobile devices may find it challenging to support two different network interfaces. If we eliminate the short-range wireless network from the mobile device, the display discovery and identification problem becomes a bit more complex. In spirit, this is essentially very similar to locating and identifying network attached printers in a campus. Users may pre-configure a list of favorite displays that are found in areas they frequently visit. If they are in a new area of the campus the display selection process is likely to be a bit more cumbersome. Intelligent displays may sport name tags that display their network identifiers, just like

printers today. Users may need to type in the network identifier on the mobile using the small input devices on the mobile device.

Once the network address of the intelligent display has been conveyed to the mobile device, all messages between the mobile device and the display shown in Figure 3 and Figure 4 can be sent over the long-range wireless interface, so the message exchanges are essentially the same as before.

4.3 Eliminating the Long-Range Interface

One of the primary uses for the long-range interface is to send notifications to the user, for instance, when new email arrives. If this long-range interface is replaced with a short range interface, it is far more likely that the user is in regions where network connectivity is unavailable, so asynchronous notifications may be undeliverable.

Nevertheless, long-range wireless is expensive from an energy perspective and several mobile devices may only be capable of supporting short-range wireless communication. In this case, intelligent displays should preferably support the short-range wireless network, and also support a routing function where mobile devices are able to treat intelligent displays like network access points. Alerts can be delivered to a mobile device when it is in the vicinity of an intelligent display and the mobile device has established a connection with the server via the display. In this mode of operation, all communication between the mobile and the server are routed through the intelligent display. The mobile device and the server may wish to encrypt their communication for additional security.

With only a short-range network interface, the reach of the mobile will certainly be spotty. In this case, the third architectural model is the preferred choice.

5. The WatchPad-ED Email Reader

The first goal in our research on symbiotic displays has been to understand what the key aspects of the user experience of multiple displays are. We focused the initial steps of our work on determining the user experience requirements of symbiotic display applications in order to guide the development of the underlying technology and systems. As a first application we chose to address a typical situation that reflects a current limitation of handhelds receiving and reading long emails that include pictures and other attachments.

Our email reader scenario follows the infrastructure model outlined in [14]. We assume that the user carries a personal device able to communicate with the Internet and to run the local portion of the email reading application. We also assume that there are display devices in the environment that support the offer of *intelligent*

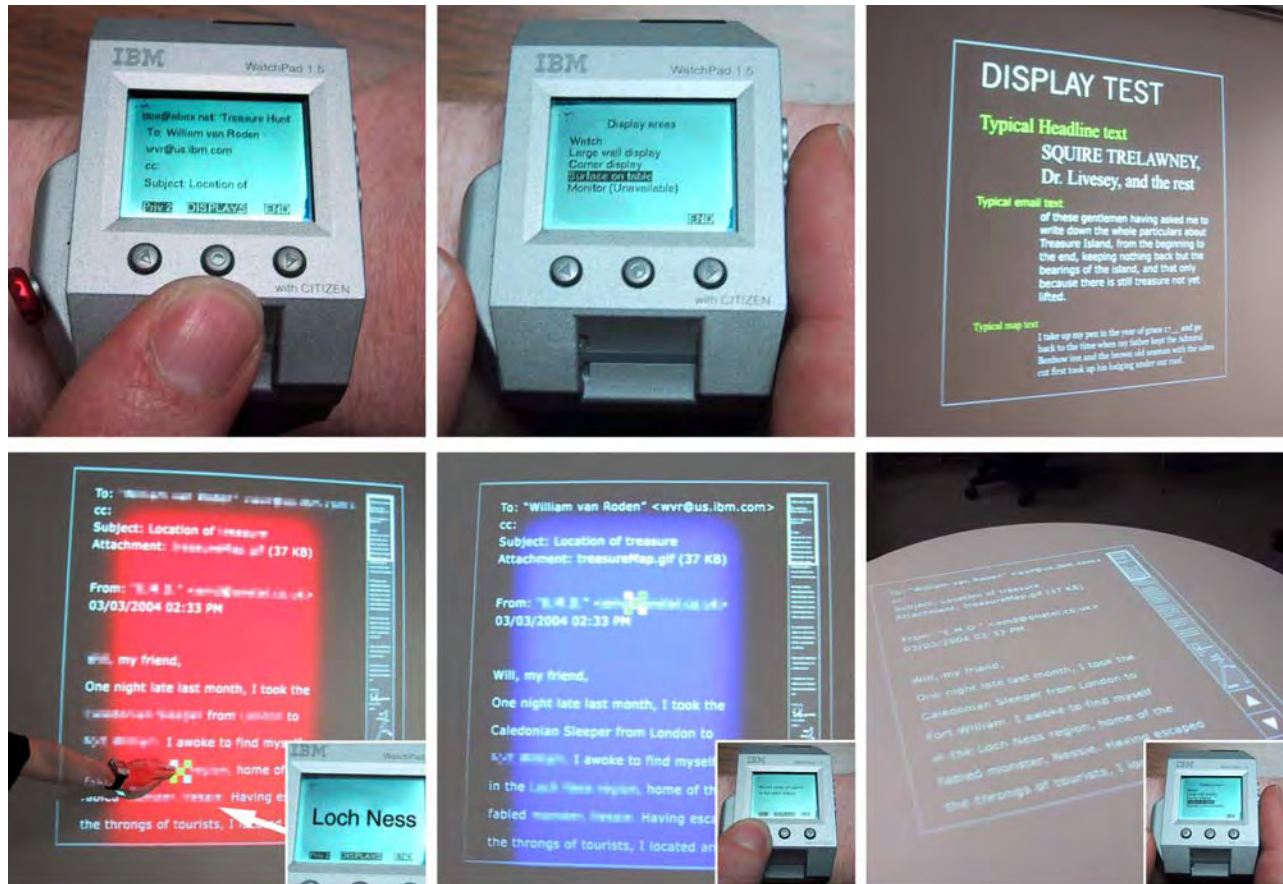


Figure 5. Using the email reader: a) an email is received by the handheld, and to more easily read the-email, a user asks the watch to find public displays in the environment; b) a list of all available displays in the environment is displayed; c) the user display the test text onto the different public display available to select the most appropriate display; d) the email is rendered onto the display with some words blurred to increase privacy; the content of blurred words can be privately checked by selecting the blurred area and reading its corresponding content on the watch display; e) the user can easily change the level of sensitivity – the lower the level, the smaller the number of blurred words; f) at any moment, the email reader can be moved to any other display in the environment or hidden.

display services. These services can be requested and used by the user as she needs more or specialized display power than afforded by her own personal device.

In our implementation, the personal device is the IBM WatchPad [10], a wearable computer with the form factor of a wristwatch with a 320x240 pixel display. The WatchPad prototype does not have a long-range network interface, but supports Bluetooth. It has 8MB of memory, which is used to run the Linux Kernel, X11, and other applications.

To simplify the task of creating an array of intelligent display services, we chose the Everywhere Display steerable projector [12] (ED-projector) as the first display device to add intelligence to. The ED-projector is a projector augmented with a pan-tilt mirror that can be used to move the projected image to one of many surfaces. Software on the display pre-distorts the image to counteract the effects of oblique projections. A single ED

projector is capable of simulating several displays in the environment with a wide range of different display sizes and resolutions depending on how far the projection surface is from the projector and the setting of the zoom level. In effect, by using a single physical ED projector we can study the effects of symbiotic relationships with a number of different “virtual displays” with varying display characteristics.

The ED projector system also has the ability to recognize user interactions with the display using a pan-tilt camera that uses a vision system looking for touch-like gestures on the displayed image [7]. More importantly the task of adding intelligence to the ED-projector system is aided by the fact that it already has a good deal of intelligence and is configurable to different surfaces and display layouts using commands specified in an XML format [8].

5.1 User Experience

The user experience with the email reader basically comprises four different steps:

Notification of the Arrival of a New Email Message: When a new email message arrives, the user's WatchPad beeps and also vibrates. The wristwatch display provides a simple option of whether to read it immediately or later.

Reading on the Watch: If the user decides to read the incoming email immediately, the text of the email is formatted to the display characteristics of the WatchPad. Given the limitations of the screen area (320x240 pixels), allowing only 3-4 words to be displayed in a reasonable typeface (see Figure 5.a), the reading of long emails becomes extremely difficult. In this situation, the user presses the middle button to check the publicly available displays in her immediate surroundings (see Figure 5.a).

Public Display Discovery and Selection: A list with the names of the available displays is shown on the watch screen (see Figure 5.b). In our setup, the list of surfaces on which the ED is able to project. By scrolling through this using the watch's stem switch (the button on the bottom-left corner of the WatchPad), the user activates a test screen on each available display. To simplify the selection of a display with appropriate characteristics, an image is produced on the inspected display showing text segments rendered with different font sizes for different elements of the text such as subject or body (see Figure 5.c). The goal is to provide the user with a simple and intuitive way to determine whether the display is appropriate. Moving the watch's stem switch and selecting a different display immediately causes the image to disappear from the previous display and appear in the new one, facilitating the localization of the newly selected display. No technical details about resolution or size are given to the user; the selection process is done directly by this process of "moving" a text pattern across displays. The user can go back and forth this list as much as necessary, comparing display capabilities until selecting a particular one by pressing the stem switch.

Reading on the Public Display: Figure 5.d shows the rendering of an email message on a public display. As discussed earlier, the basic mechanism to handle sensitive information is the suppression of selected words and phrases in the text. To access the text of a specific suppressed word, the user touches the projected image as shown in Figure 5.d. The vision system interface of the ED-projector recognizes the touch, and the suppressed word appears on the watch display, where it can be safely read. Alternatively, the user can scroll through the list of suppressed words using the watch's stem switch: for a given word on the watch's display, a cursor on the projected image marks the position of the corresponding

blurred word on the public display.

The user also has the option to reduce the number of blurred words in case he believes that the current situation of the environment is safe to do so. By clicking the left button of the watch a lower sensitivity level is selected. In our prototype, three sensitivity levels are available, corresponding to situations where all sensitive content is suppressed, only very sensitive words are blurred, and full text. The current sensitivity level is signaled by the color of the background, respectively red, blue, and black. Figure 5.e shows the middle level of sensitivity corresponding to the text on Figure 5.d (see Figure 1).

At any moment, the user can click the middle button of the watch and return to the display selection process. If the user is already reading email, the image rendered on each display is the most recently seen email page, rendered with the highest sensitivity level for safety as shown in Figure 5.f.

5.2 Implementation Issues

To enable symbiosis between the ED and the watch, we added a Bluetooth interface module to the PC that drives the ED projector. We implemented the proxy functionality described in section 4.1 as a custom written application that runs on the same PC. During the demo, the watch and the ED first connect to each other and establish a TCP/IP connection, with the display conceptually assigning the watch an IP address. For the demo since the proxy and the display are co-located, we simply established the connection directly with the proxy. This simplified the implementation of the demo eliminating the two step process of first connecting to the display and notifying the proxy that the watch had a network connection.

The new email notification is an XML formatted message that the proxy sends to the watch over the TCP connection, which the watch conveys to the user by beeping. As part of the initial connection set up phase the display sends a list of surfaces on which the ED projector can display. In response to user input indicating intent to read the incoming message, the watch displays the list of surfaces to the user, adding the local watch screen as one of the options. The user can scroll through the list using the stem switch and the watch sends each wheel click to the proxy which communicates with the ED-controller program using an XML command language called EDML [8]. When the user presses the stem switch to select a display, the proxy automatically sets the privacy threshold to the highest level, and sends the appropriately filtered message to the display. User interaction with the displayed image is recognized by the vision system and converted by the ED-controller program into messages to the proxy that identifies the words that were touched and sends the words to the watch. The watch then displays the

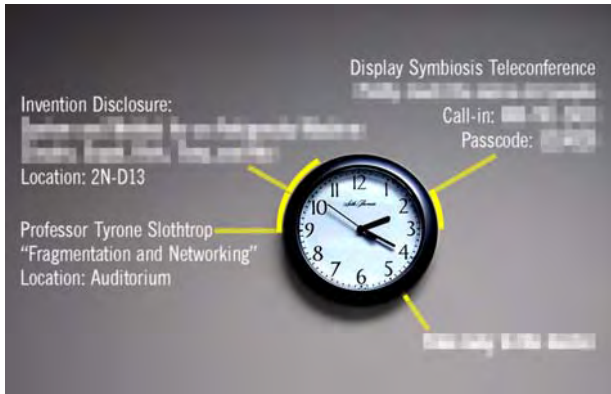


Figure 6. Symbiotic calendar viewing application superimposed onto a real object in the environment; blurred areas correspond to sensitive information in the calendar (design proof).

word or group of words on the watch face.

To extend this implementation to multiple displays we would simply move the proxy program to another server machine that is reachable on the wired network from the different displays. Each display will have its own Bluetooth interface and the display will implement a routing function that enables the watch to reach the proxy via the display's wired network connection. The trust relationship may be enabled through the ephemeral tokens as described earlier. If the environment is benign, the proxy may simply observe the route that the packets take to get to the watch and implicitly trust the display that at the last hop before the watch. The proxy server can initiate a network connection the display controller program and orchestrate the interaction in the same manner as above.

6. Other Applications and Future Work

While we have considered viewing email in detail, our approach applies to a broader set of applications. For example, viewing personal photographs or calendars in public spaces are good candidates to consider.

With respect to viewing one's calendar in public we could leverage some of the mechanisms and fields provided already in iCal specifications [4] or with features supported in Lotus Notes and Microsoft Outlook. The simplest mechanism would be to block out private appointments when viewing in public. In addition, specific fields in the appointment such as location, other participants, passwords for teleconferences and video conferences can be annotated at different levels of sensitivity. These fields could be unblocked on request just as in the email application. We believe that calendar entry tools have to provide simple options to mark the sensitivity level of fields in each appointment and not just whether the entire appointment is sensitive or not.

Figure 6 shows a design proof of a calendar viewing system we are developing. Here, we take advantage of projection to seamlessly integrate the calendar with a clock in the environment (such as in [15]). This prototype expands an idea proposed in [6] by incorporating mechanisms to view sensitive information in the calendar entries. As we see in Figure 6, some calendar entries are totally blurred and some have only the meeting participants hidden. Important information such as call-in numbers and passwords are also normally blurred.

Another interesting application to consider is viewing photographs on public displays. Though the resolution of pictures captured on digital cameras is increasing, the size of the built-in display continues to remain small. An emerging practice is to show the pictures to the subjects in the pictures immediately after shooting them. If symbiotic displays are easily available, it may be possible to view one's pictures soon after capture on a large display. This scenario is examined in detail in [9]. However, users may not wish to have some of their pictures put on large public displays. One way to realize this application would be to first show the thumbnail on the camera. The user can then decide whether this image can be viewed on the larger display. The camera also serves as an intelligent remote control for the viewing experience. Possible ways to determine sensitivity level automatically could use context data such as location where picture was taken, date and time, and who is in the picture.

Another idea that we plan to explore over time is the use of audio headsets for conveying the private information. Mechanisms are used today, for instance at the United Nations, to simultaneously translate conversations into a variety of languages that are delivered to different people using headsets. We propose to extend this concept to deliver sensitive information privately. Other researchers have used stereoscopic shutter glasses to provide different views of the same scene to different viewers [1]. Such ideas may be incorporated if users are willing to wear shutter glasses and headsets.

7. Conclusions

In this paper we describe the approach we are taking towards the realization of symbiotic displays. We have discussed research challenges that we have come across in our quest, how we overcame several of these challenges and also outline the ideas we are currently pursuing to solve many of the remaining challenges.

We first discussed why symbiotic relationships between displays are important in spite of constant improvements in some aspects of mobile devices. We then identified that using display symbiosis gives rise to the important question of how to handle viewing of

sensitive content in public displays. We described a variation of the blackening technique that can be used to mark sensitive content and subsequently allow the user to selectively retrieve the sensitive content on the screen of a personal device. Three possible system configurations were analyzed and the advantages and disadvantages of each method were discussed.

A major contribution of this work is the detailed report of the user experience using display symbiosis for viewing email. We chose email as our representative application since it surfaces many issues associated with sensitive content. An initial version of our implementation for viewing email using the IBM Everywhere Display and the IBM Linux WatchPad was shown at PerCom'04 and received very positive feedback. We are confident from our experience that display symbiosis can be used effectively for a broad set of applications, though several challenges remain.

References

- [1] M. Agrawala, et al., "The Two-User Responsive Workbench: Support for Collaboration Through Individual Views of a Shared Space," in Proc. of SIGGRAPH'97, 1997.
- [2] D. Balfanz, et al., "Talking to Strangers: Authentication in Ad Hoc Wireless Network," in Proc. 9th Network and Distributed System Security Symposium (NDSS'02), San Diego, California, 2002.
- [3] D. Butler, "US Intelligence Exposed as Student Decodes Iraq Memo," in *Nature*, vol. 429, 2004, pp. 116.
- [4] F. Dawson and D. Stenerson, "Internet Calendaring and Scheduling, Core Object Specification (iCalendar)," IETF RFC 2445, Nov. 1998.
- [5] C. Grover, et al., "LT TTT - A Flexible Tokenisation Tool," in Proc. of Second International Conference on Language Resources and Evaluation (LREC 2000), Athenas, Greece, 2000.
- [6] S. Intille, V. Lee, and C. Pinhanez, "Ubiquitous Computing in the Living Room. Concept Sketches and an Implementation of a Persistent User Interface," in Proc. of the Video Session of UbiComp'03, Seattle, Washington, 2003.
- [7] R. Kjeldsen, A. Levas, and C. Pinhanez, "Dynamically Reconfigurable Vision-Based User Interfaces," in Proc. of 3rd International Conference on Vision Systems (ICVS'03), Graz, Austria, 2003.
- [8] A. Levas, et al., "An Architecture and Framework for Steerable Interface Systems," in Proc. of UbiComp'03, Seattle, Washington, 2003.
- [9] C. Narayanaswami and M. T. Raghunath, "Expanding the Digital Camera's Reach," *IEEE Computer*, vol. 37 (12), pp. 65-73, 2004.
- [10] C. Narayanaswami, et al., "IBM's Linux Watch: The Challenge of Miniaturization," *IEEE Computer*, vol. 35 (2), pp. 33-41, 2002.
- [11] M. S. Pandit and S. Kalbag, "The Selection Recognition Agent: Instant Access to Relevant Information and Operations," in Proc. of the 1997 International Conference on Intelligent User Interfaces (IUI'97), Orlando, Florida, 1997.
- [12] C. Pinhanez, "The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces," in Proc. of Ubiquitous Computing 2001 (UbiComp'01), Atlanta, Georgia, 2001.
- [13] C. Pinhanez, et al., "Symbiotic Displays: Combining Multiple Displays to Enhance User Experience," in Demonstration Program of IEEE PerCom'04, Orlando, Florida, 2004.
- [14] M. Raghunath, C. Narayanaswami, and C. Pinhanez, "Fostering a Symbiotic Handheld Environment," *IEEE Computer*, vol. 36 (9), pp. 56-65, 2003.
- [15] J. Rekimoto and M. Saitoh, "Augmented Surfaces: A Spatially Continuous Workplace for Hybrid Computing Environments," in Proc. of CHI'99, Pittsburgh, Pennsylvania, 1999.
- [16] S. J. Ross, et al., "A Composable Framework for Secure Multi-Modal Access to Internet Services from Post-PC Devices," in 3rd IEEE Workshop on Mobile Computing Systems and Applications WMCSA'00, Monterey, California, 2000.
- [17] M. K. Stern, "Identifying and Understanding Dates and Times in Email," IBM Technical Report RC-22875, 2003.
- [18] R. Want, et al., "The Personal Server: Changing the Way We Think About Ubiquitous Computing," in Proc. of UbiComp'02, Goteborg, Sweden, 2002.