Supporting ubiquitous information on very small devices is harder than you think

David M. Hilbert, Jonathan Trevor

FX Palo Alto Laboratory
3400 Hillview Avenue
Palo Alto, CA 94304 USA
{hilbert, trevor}@fxpal.com

Bill Schilit

Intel Research
2200 Mission College Blvd.
Santa Clara, CA 95052 USA
bill.schilit@intel.com

Abstract

A basic objective of ubiquitous computing research is ubiquitous information: the ability to utilize any content or service, using devices that are always at hand, over networks that don't tie us down. Although much progress has been made, the ideal remains elusive. This paper reflects on the interrelations among three dimensions of ubiquitous information: *content, devices,* and *networks*. We use our understanding of these dimensions to motivate our own attempt to create a ubiquitous information system by combining unlimited World Wide Web content with mobile phones and mobile phone networks. We briefly describe a middleware proxy system we developed to increase the usefulness of very small devices as Internet terminals. We conclude with a post-mortem analysis highlighting lessons learned for others interested in information systems for very small devices.

1 Dimensions of ubiquitous information

While people have made inroads toward ubiquitous information, the ideal remains elusive. One reason is that the very notion of "any information, anytime, anywhere" places conflicting requirements on the *content*, *devices*, and *networks* that make up information systems.

Consider, for instance, laptop computers and local area networks. The combination of powerful processors, flexible user interfaces, high-resolution displays, speakers, and high-bandwidth networking make it easy to interact with rich and interactive content. However, these same characteristics place practical limits on device portability and network mobility. Mobile phones and networks, on the other hand, make the opposite trade-off. They provide extreme device portability and network mobility at the cost of greatly reduced content capabilities, due to limited user interfaces and lower bandwidth networking. Palm-sized computers with wireless WAN cards provide a middle ground: they are more portable than laptops, and more usable and content-capable than mobile phones, but far less pervasive than phones. It seems you can't have it all.

It gets worse. Current trends suggest these tensions will grow in the future. Desktop computers and wired networks keep getting faster and more powerful, meaning ever-richer content will flourish on the Web. At the same time mobile phones keep getting smaller and lighter, meaning tiny user interfaces, processors, and wireless networking will continue to be hard-pressed to handle the content. While wireless devices and networks are increasing in power too, the smallest and lightest devices will be the most ubiquitous, ensuring that usability and interactivity will continue to be problematic. In essence, we face *natural tensions* whenever we attempt to maximize "ubiquity" along all dimensions at once.

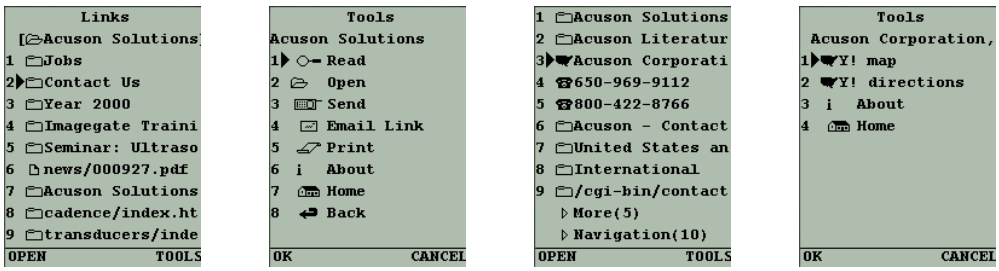

a b c d

Figure 1: M-Links on a Neopoint 1000 web phone. M-Links retrieves requested Web pages and returns native format (WML, HDML, CHTML) screens. In this example, M-Links presents the user with a list of links to navigate the homepage of Acuson, a medical equipment manufacturer (Figure 1a). The open folder indicates the page the user is currently navigating. The closed folders indicate links to other pages. File icons indicate links to non-HTML content, such as PDF or multimedia. When the user presses the "TOOLS" soft key, the interface switches to a list of actions the user can perform on the current link (Figure 1b). For instance, pressing "Read" would allow the user to read all the text on the page. Pressing "CANCEL" returns the user to navigation mode (Figure 1a) again. Selecting "Contact Us" navigates to Acuson's contact information page (Figure 1c). M-Links displays not only the links but other useful information such as phone numbers and addresses. The "Navigation" item (bottom of screen) collects links that repeat across many Web pages. The "More" item (second from bottom) shows the remaining links on the Web page that would not fit on the current screen. In (Figure 1c), the user selects Acuson's street address and presses the "TOOLS" soft key to show a list of actions appropriate for addresses (Figure 1d). For instance, if the user selects "directions", m-links passes the street address to Yahoo! Maps and returns the directions to the user in text form. The "About" and "Home" options allow the user to get more information about the current link and to return to the m-Links home page respectively.

2 A real world ubiquitous information system

We wanted to create a real world ubiquitous information system using content, devices, and networks available today. We knew that in the future, wearable displays and novel input devices would remove *some* of the usability limitations of portable devices, and publicly accessible embedded computers with usable displays and interfaces will someday be more ubiquitous. However, these future possibilities aren't helpful in creating a real system now. So we decided to build upon Wireless Web phone technologies because of their trend towards ubiquitous devices and networks.

2.1 The M-Links middleware proxy

While the current Web browsing model works well for desktop and laptop computers, it is less well suited to phone-tops. Cell phones typically accommodate only three to twelve lines of text, and their design emphasizes portability and features such as battery life, audio clarity, and ease of selecting names from a phonebook. Web interaction has, so far, been a secondary concern. Thus, we developed M-Links, a middleware transducing proxy, to help mobile phone users access and actually *do* things with a wider range of Web content than before. M-Links differs from other Web transducers (Brooks et al., 1995) in that it factors Web browsing into two separate interfaces: one for navigating links and the other for performing actions on links. See Figure 1.

When users access a Web page using m-Links, they see a list of links from that page and can dig through the list in the same way they dig through folders on their desktop to locate files. When they find a link, they may invoke services, analogous to right clicking on a document and using the context menu on a desktop interface. Although users can't do much directly on the phone with

content such as large PDF documents or MPEGs, m-Links users can always do *something* with any content they find, even when their device is not equipped to handle the content itself. For instance, once users have located content of interest, e.g., a digital movie or novel, they can always send the content (or URL) via email for later use on the desktop or other device by simply selecting the link and applying the e-mail service. Alternatively, users can use M-Links to invoke third-party content translation services, for instance, to convert PDF documents into text formated for the phone, once these services become available.

M-Links is both simple to use and powerful. Its simplicity comes from the *navigation interface* with its server-side data-detectors for bubbling-up useful bits of information, such as phone numbers and addresses: separating links from page content makes navigation a matter of selecting a link from a list. Its power comes from the *action interface* with its open systems architecture for incorporating new network-based services, similar to browser plug-ins*:* because users can apply various Web-based services to any link, they can do more with content than simply read it on their phones. For more detailed discussion of the M-Links infrastructure and user interface, see (Schilit et al., 2001) and (Trevor et al., 2001) respectively.

3 Lessons Learned

Our research approach differs from traditional approaches in that we were teamed with business people to develop a working prototype in conjunction with a business plan that we could show to potential users, investors, and business partners. Feedback from all these people helped identify areas in which M-Links could be improved.

3.1 Technical Issues

We quickly realized **the need for speed** could not be overestimated. A basic assumption all along was that we needed to optimize the interface for devices with limited displays and input mechanisms. This led to our separating links from content to speed navigation, and our use of server-side data detectors to bubble-up interesting bits of information such as phone numbers and email addresses. However, people still found it painful to navigate to find content. They said things like “I wish I could find all the phone numbers on this site at once” or “I know what I’m looking for is a PDF file, but I can’t remember where it is!” This led us to develop yet more data detectors (e.g., for street addresses) and suggested we might also increase speed by providing a *filtering* interface, that would instantly bubble-up all the links on a site of a given type specified by the user.

Another lesson was that **no device is an island**. People felt frustrated having to register using their email address via the phone keypad, and wanted to know why they couldn’t set up their account from their desktop PC. This suggested a substantial support infrastructure at the desktop might greatly increase the chances of user acceptance. We began developing a Web site for account management and have begun exploring other mechanisms for leveraging desktop interactions to facilitate phone-top interactions. For instance, instead of specifying URLs using key entry at the phone, users might select from a list of “favorites” or “recently visited” sites gleaned from their desktop Web usage. Another issue is that people often had a hard time re-locating content on the phone that they had already located at the desktop before. This led to the idea of adding another filtering feature: namely, “show me all the links I’ve accessed from this site before”.

As time passed, we also began to appreciate the importance of **hands-free use**. We repeatedly heard that U.S. mobile phone users, in contrast to Japanese users, use their phones while otherwise engaged, such as while driving. In Japan, Wireless Web use is pervasive on public transportation, as users are free to direct their attention to their phones. This led us to investigate a voice interface

for M-Links. However, poor integration between the voice and data capabilities of mobile phones at the time made developing such an interface problematic. We expect this limitation to decrease in the future with 3G mobile phone networks.

Finally, we learned that **infrastructure limitations** were more significant than we had expected. The idea that users could instantly access the Wireless Web simply wasn't true. It often took several seconds to establish a data connection. Furthermore, the reliability and performance of the networks left users frustrated by dropped connections and slow page retrievals. Finally, while users always had their phones with them, phone batteries didn't live up to the promise of "anytime" use. While all of these are limitations of the devices and networks upon which M-Links was built, and not M-Links itself, we underestimated the power of frustrated expectations to reduce users' willingness to consider new capabilities. Again, many of these problems will likely be improved with future advances in mobile phone service.

3.2 Non-Technical Issues

In addition to technical issues, numerous non-technical issues also inhibited our success.

First, there were significant **business-related challenges**. Substantial investment (in time and capital) is required to move research prototypes to commercial products. In some cases, a large marketing campaign may be required to draw attention to, or promote, a lifestyle in order to sell a product, as in the case of AT&T's recent mLife campaign. Furthermore, the lack of a micro-payment system for Wireless Web services (like the DoCoMo iMode model in Japan) means there is no simple mechanism for third-party service providers (such as ourselves) to charge small amounts of money per use of their service without arranging profit-sharing agreements with the big players. Thus, reaching a large number of customers requires relationships with telephony service providers (such as Sprint PCS and AT&T), applications providers (such as Yahoo), or infrastructure providers (such as Inktomi). We were beginning to establish relationships with key partners in all of these areas when the Internet bubble burst, leaving us in an inhospitable business climate.

Second, we now realize that limited **marketplace maturity** and **customer readiness** also impeded our progress. In a nutshell, our vision depended on people actually using their Wireless Web phones to access the Wireless Web. While this condition was met in Japan and other countries, the U.S. has continued to lag: people use their phones to place calls, not access information. This makes it hard for "add-on" information services such as M-Links to take off, and continues to fuel the desire for ever-smaller phones, a trend at odds with phones being used as information devices.

3.3 The dimensions reconsidered

In choosing a ubiquitous information platform, we now realize we should have asked ourselves: **ubiquitous for what?** In the U.S., mobile phones are used ubiquitously as *communications* devices, but not as *information* devices. We were seduced into thinking our platform would include all those devices people were already using as communications devices, when in fact our platform was really only those devices already being used as information devices—a far smaller pool. Given this perspective, another approach would have been to pick devices used less pervasively in general, but more pervasively *as* information devices, such as palmtop computers.

We also should have asked ourselves: **ubiquitous for whom?** By deciding up-front that we wanted to support ubiquitous information for everyone, instead of selected groups, we were forced into providing generic services for a horizontal market, instead of targeted services for vertical markets. A problem with generic services is they often imply lower value, which in turn implies

the need for very low adoption costs. This intensified the pressures on infrastructure issues including network speed, reliability, and performance—attributes beyond our control. Another strategy would have been to focus on specific applications with higher value for particular types of customers, thus increasing their willingness to tolerate current infrastructure limitations. After achieving success in vertical markets, we might then have spread out into adjacent markets to increase ubiquity.

In short, we believe there was **a cultural catch-22** at work that we didn't fully appreciate at the time. Since our parent company is Japanese, we were acutely aware of (and motivated by) the huge success of mobile phone information services in Japan. Additionally, analysts in the U.S. were saying a key reason American's weren't satisfied with the Wireless Web was that they, to a far greater extent than Japanese, were accustomed to accessing WWW content on their desktops, and thus felt frustrated they couldn't access the same content on their mobile phones. So there appeared to be a great opportunity: combine the runaway success of Wireless Web usage (as demonstrated in Japan) with the ability to access a much wider range of content on the WWW (as is the custom in the United States). Ironically, while Japanese are enthusiastic about using the Wireless Web, they are far less interested than Americans in accessing WWW content. At the same time, while Americans are enthusiastic about accessing WWW content, they are far less willing than Japanese to use mobile phones to do it. In other words, we envisioned a community of users as enthusiastic about Wireless Web usage as Japanese, and as enthusiastic about accessing WWW content as Americans: a community that doesn't currently exist—at least not in the U.S.

4 Conclusions

Despite its limitations, M-Links achieved many of the goals we set out for it. Namely, it substantially increases the content capabilities of highly portable (ubiquitous) devices operating over very wide area wireless (ubiquitous) networks. However, if we were to do it over again, we would consider retargeting M-Links to take advantage of devices that are more pervasively used as information devices today (such as palmtop computers), and perhaps for specific markets (such as mobile sales and repair professionals) in which users may be more willing to tolerate limitations in current infrastructure technologies. Nonetheless, by leveraging extensible network-side services, our approach still offers substantial value to small Internet device owners by allowing them to exploit the computing resources and network connectivity of larger, more powerful devices to increase the ways in which they can display, share, and otherwise manipulate Web content using very small devices.

5 Acknowledgements

We would like to thank T.K. Koh for her development support and Debra Go and Patrick Ahern for their business expertise.

References

Charles Brooks, Murray S. Mazer, Scott Meeks, and Jim Miller (1995). Application-specific proxy servers as HTTP stream transducers. *Proceedings of the 4th International World Wide Web Conference (WWW '95)*.

Bill N. Schilit, Jonathan Trevor, David M. Hilbert, and Tzu Khiau Koh (2001). M-Links: An infrastructure for very small Internet devices. *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom 2001).*

Jonathan Trevor, David M. Hilbert, Bill N. Schilit, and Tzu Khiau Koh (2001). M-Links: From desktop to phonetop: A UI for Web interaction on very small devices". *Proceedings of the 14th Annual Symposium on User Interface Software and Technology (UIST 2001).*