Towards seamless mobility with personal servers

Markus Bylund and
Zary Segall

The authors

Markus Bylund is a Researcher at the Swedish Institute of Computer Science, Kista, Sweden.
Zary Segall is Distinguished Professor at University of Maryland Baltimore County, Baltimore, Maryland, USA.

Keywords

Computer hardware, Control technology, Telecommunications, Mobile communication systems

Abstract

Observes that the future of mobile communication networks lies not only in how successful people are in deploying technologies (2.5G or 3G with a high degree of coverage and roaming between operators), but also in how well people can create a functioning environment and usage situation for end-users in which they can get a homogeneous and continuous usage experience, despite the very heterogeneous world in which they, after all, will live. Concepts are advanced that support this observation (network independence, UI/device flexibility, and user experience continuity), and a possible solution is proposed that would take people in that direction (the personal server).

Electronic access

The Emerald Research Register for this journal is available at
www.emeraldinsight.com/researchregister

The current issue and full text archive of this journal is available at
www.emeraldinsight.com/1463-6697.htm

info
Volume 6 · Number 3 · 2004 · pp. 172-179

DOI 10.1108/14636690410549499

1. Introduction

Seamless mobility (Thorngren *et al.*, 2003), or the I-centric view of mobile networking, is often referred to as the marriage between 2.5/3G and WiFi technologies[1]. With the wide coverage of 2.5G/3G technologies, combined with the high but local performance of WiFi, it is argued that electronic services can be used seamlessly (Sawhney, 2003).

However, we believe that there is more to it than performance and network connectivity coverage. In this paper, we explore the added factors of user interface (UI) and device flexibility and the notion of user experience continuity. In combination, these factors would allow a continuous usage experience of all sorts of electronic services, accessed from many different types of usage contexts, and mediated via a multitude of devices. We further propose the concept of a personal server as a solution to the challenges that these factors add. The concept illustrates that there exist technical solutions to the challenges of achieving seamless mobility. Still, we conclude this paper by discussing complementary regulatory and business initiatives that also need to be realized.

The paper is organized as follows. Section 2 presents network capabilities and coverage, UI and device flexibility, and the importance of true user experience continuity as key issues in achieving seamless mobility. Section 3 describes the personal server approach and its impact on seamless mobility. We conclude with a discussion of the personal server concept's implications and outline a few future research challenges.

2. Expanding the concept of seamless mobility

In this section we highlight three factors that we view as particularly important to achieve seamless mobility: network capabilities and coverage, UI and device capabilities, and the importance of true user experience continuity.

2.1 Network independence

Many, if not most, discussions of mobility have so far been about network issues, ranging from how to provide support for mobility in network protocols (for example Mobile IP (Perkins, 1997)) to discussions of different solutions for wireless networking. The main reason for this is that different technical constraints largely influence what can be done in mobile settings[2]. For quite a few classes of applications, the performance of 2.5G and 3G networks is too low – for example,

info
Volume 6 · Number 3 · 2004 · 172-179

highly interactive applications such as games, where the latency is the most limiting factor. This is also the case for applications as simple as Web browsers[3]. To quantify this problem we performed a number of network traffic measurements, using a Sony Ericsson P800[4]. The experiments revealed that 2.5G networks (GPRS) offer exceptionally poor throughput and response time, compared to solutions that route the network traffic over, e.g. a Bluetooth connection. The latter, in combination with wired network connections, can be used as an alternative to 2.5G connections. It was also interesting to notice that the variation in response time for the slow 2.5G connections was surprisingly high[5].

The performance of network connections, however, is only one of several factors that influence mobility. Roaming between different types of networks and service providers is another factor. One could argue that, on the one hand, this is a technical matter with several existing solutions (Hansen *et al.*, 1998; Perkins, 1997). On the other hand, it is a matter of more practical nature. For users to get seamless access to network connectivity from every possible place, all service providers need to agree on automatic roaming or users need to subscribe to all service providers. Given the diversity and large number of service providers, neither factor is likely to work well enough in order to take mobile network access for granted (Almgren, 2003).

Pricing of network connectivity is yet another factor that influences how users make connections in mobile settings. For example, in the case of GPRS, the cost of service is often a combination of a flat and a traffic-based rate, while the cost of most wired connections (for example cable TV and ADSL) is only based on a flat rate. This means that users who access both GPRS and ADSL connections have reasons to plan their usage of bandwidth-intensive services for times and places where ADSL is available. A possible hypothesis is that some users reject GPRS only because of the variable cost associated with its usage.

All these factors result in the uncertainty of adequate network connection being universally available. Considering this fact, we argue that to achieve seamless mobility we must not only resolve the issue of how to provide for mobile network connectivity, but also decide how to enable users to operate services independently of (Internet) network connectivity. This would partially unburden users from issues such as low bandwidth, high latency, roaming, and high traffic costs. We recognize that some services require a real-time connection, but most services include parts that could work without such network connection.

2.2 UI and device flexibility

It is not productive to talk about seamless mobility if one does not consider that the user context, while mobile, is bound to vary over time. However, different contexts and situations require different types of devices and user interfaces. This has been acknowledged by the manufacturers of access devices. There exists today a wide range of devices for mobile computing – ranging from smart phones to tablet PCs and laptops. These devices differ in a great number of ways, but most often their differences are well motivated. The screen of a smart phone, and its adapted UI, is much smaller than that of the desktop computer, simply because it would be burdensome to carry a large display while walking around. Nevertheless, while sitting in an office, most users prefer a device with a large display. These differences can be seen on all levels of design of the artifacts that we use to access electronic services. Some artifacts have hard buttons that immediately trigger specific applications, while others have full-size keyboards. Some devices have scroll wheels for navigation while others have mice, and so on.

Judging by the diversity of mobile devices available today, it would seem that the range of available artifacts provides good support for mobility as it is. This is not the case, however, because the device flexibility is not there yet. If one is to take full advantage of the differences between various artifacts, one must also be able to switch seamlessly between them. This is seldom possible with the infrastructure for electronic services and range of artifacts available today. Usually, it is not even possible to use the same service and data on different artifacts at all, let alone in a seamless manner. In some cases it is possible to synchronize data between similar, but different, services on different devices. This is only true for the data, though; the state of the service is seldom, if ever, included in the process.

Therefore, we argue that in order for seamless mobility to become a reality, we must find the means to support flexibility in the choice of device and UI by allowing access to the same services and data from many different types of devices.

2.3 User experience continuity

As stated above, seamless mobility requires some kind of network independence and flexibility in choice of user access device. These two factors would make it possible to introduce continuous user experience for mobile services. This would allow users to start working with a task on one device with a particular network connection, to continue the work on another device completely without network connection (on an airplane or at a hospital for example), and finally, to finish the task on a third device and network connection.

From the users' point of view, this kind of continuity can take two different shapes: continuity as in the remote access case, and continuity with adapted UIs. In the first case, the different access devices can be seen as remote controls to the electronic services of the user. The UI is identical in all cases, but the hardware running the UI changes. The electronic services of the user execute on one and the same computer (usually a desktop computer or a server) and, as the user switches access devices, only the stream of screen, keyboard, and mouse events needs to change origin and destination. The execution state and data of the service are kept the same over time, since they are never moved between the devices[6].

The second alternative, to adapt the UI to different user access devices, is more complex. In this case the UI of the electronic service is adapted to the capabilities and constraints of the device currently in use. A laptop may provide a wide overview of all functionality of the service, including full access to input features, while a much smaller smart phone might offer a simplified UI. Output may be filtered (certain kinds of media may, for example, not be possible to present) and the means for input may be reduced because of the limited keyboard. On some devices, the adaptation may go so far as to abandon the prevailing desktop and Windows metaphor in favor of something more suitable. This could, for example, be the case for a purely voice-based interaction device. However, this alternative for providing continuity is expensive and complicated to realize, mainly because it requires implementation and maintenance efforts that grow rapidly with the number of devices that should be supported. Another problem is that when a new device is introduced, a new adaptation (or version of the service) needs to be implemented[7].

From the user perspective, these two alternatives have both advantages and disadvantages. The remote access alternative is preferable because the UI always stays the same – users immediately recognize the service when using it on a new device and there is no learning time that needs to be accounted for. However, since no adaptation of the UI is done, the usability of some services may suffer. An example of the latter is the use of a complex application such as Microsoft Word on a smart phone. The small screen provides an unacceptably poor view of all actions and capabilities of the rich UI, and input is extremely awkward. The adapted UI alternative is powerful since it allows services to take full advantage of the unique features of each device, thereby building on the knowledge and design expertise that was put into the development of each device. This means that special features such as hard buttons and scroll wheels can be assigned functionality that ties closely to the intentions of the designer of the device. However, since the same service will appear differently on different devices, some learning time will be required for each new device being used.

3. Personal server

The term ubiquitous computing (Weiser, 1991; Kortuem and Segall, 2003) refers to a vision of invisible computers being embedded in our environment and participating in our lives. The envisioned usage, however, is far from the way that we use computers today. Instead of having multi-purpose computers such as desktop or laptop PCs, computers will blend into our environment and turn into invisible and special-purpose devices that will help us to accomplish tasks everywhere, not only when sitting at a desk. Using computers for various purposes, according to the vision, will be as much of an unconscious activity as using the nowadays-ancient technology of writing for long-term storage of information.

Recognizing the relevance of the three factors of seamless mobility outlined above, we anticipate that next-generation personal computers will support a variety of more interactive and proactive modes of computation than we can see today. As embedded systems and networks become mobile, personal mobile computers will act as both agents and intermediaries between their users and the embedded system infrastructure. Miniature, portable server systems will allow their users to seamlessly access networks, retrieve data from embedded sensors, control embedded actuators, perform functions proactively on the user's behalf, and interact with other people's personal systems as people come into proximity. For such systems, the locus of control moves through the environment, interacting proactively and often autonomously in response to real-time, environmental inputs. For the purposes of this paper, we will call such a personal control system a personal server (Bylund and Segall, 2003).

Personal servers provide a variety of new, dynamic modes of operation and interaction depending on the environment and the user's needs – in other words, they are potentially highly suitable for realizing seamless mobility. In the simplest case, we can visualize the use of a personal server in terms of the user equipped with such a device, walking up to a monitor, keyboard and other peripherals with wireless interfaces and beginning to work. His or her personal server

info
Volume 6 · Number 3 · 2004 · 172-179

connects to the surrounding peripherals, devices, and networks, allowing access to local data or control systems. Since the locus of control moves with the user, it arrives fully customized and able to personalize the user's interaction with environmental systems (e.g. to the user's job, capabilities, goals, age, etc.) and employ the user's own set of application software. The type of hardware is not important on this level; it could be a desktop computer just as well as a virtual server on a multi-user machine, or a PDA. The important aspects of the concept are instead that it is uniquely tied to a single user, and that the user only needs to view the collection of functionality and data that the server constitutes as one single entity[8].

To use the services of the personal server, user access devices of different kinds are needed. Note that the user access device can be, but is not necessarily, the same device that hosts the personal server. In contrast to the uniqueness of the personal server (from the user perspective), user access devices are many, one for each context, situation, or perhaps even application[9].

3.1 Personal server characteristics

The personal server concept encompasses a large number of solutions for electronic service mediation. In order to evaluate the concept from a seamless mobility perspective, we need to explore some of its characteristics somewhat more deeply. The most obvious characteristic is whether the personal server is remote (i.e. something that is placed at a fixed location and accessed remotely) or mobile (i.e. something that the user can bring along and access locally).

3.2 Remote server

The advantage of having a remote server is that services get access to a wired, and possibly high-quality, network connection at all times. Services can execute continuously assisting the user, even when the user is not in contact with the remote server (e.g. a broker service buying and selling stocks on behalf of its user). However, being remote also implies the need of a network connection between server and user access device in order for the user to interact with the services. The latter violates the desire to gain independence of network connection as described above, but depending on user needs this may be balanced by the possibility to have personal services executing continuously, regardless of the user's whereabouts. The remote server can be further divided into two groups, depending on whether it is shared or personal. A shared server would typically be owned and maintained by a service provider or a corporation that would host services for individual users, while the personal server could be an individual user's desktop computer at home or at work. This distinction, however, is not significant when evaluating the concept from a seamless mobility perspective.

3.3 Mobile server

The advantage of the mobile server is that its user can bring it along. This allows the user to access services, even when there is no network connection available at all. Of course, parts of the functionality of some services require a network connection in order to work, but most services can at least offer some functionality in a completely disconnected mode. A chat service, for example, inherently requires a network connection to other chat peers in order to function. But the service could at least offer history-browsing and the editing of preferences when no network connection is available. As discussed above, the user might even prefer a local off-line mode even if a network connection is available, in order to minimize the cost of network traffic. In such a case, services can be programmed to minimize network traffic by only sending and receiving necessary information. An e-mail client, for example, could be set to download only the headers of new e-mails, and the e-mail bodies only on explicit user request (just as most e-mail clients on PDAs and smart phones already work). Another network-related advantage of the mobile server is reduced latency. Highly interactive services such as games benefit from local execution, especially if the only available network connection is a GPRS connection with latency in the range of seconds (Figure 1).

The mobile servers can be further divided into portable or wearable depending on how mobile they are. A portable server is typically a laptop computer that can be brought along by the user, but it is large enough to include full-size user I/O peripherals such as screen and keyboard. It is difficult to use when on the move or, for example, when driving a car; but if it is close to the user (e.g. in a back-pack or in the trunk of a car), more

Figure 1 A taxonomy of different types of personal servers

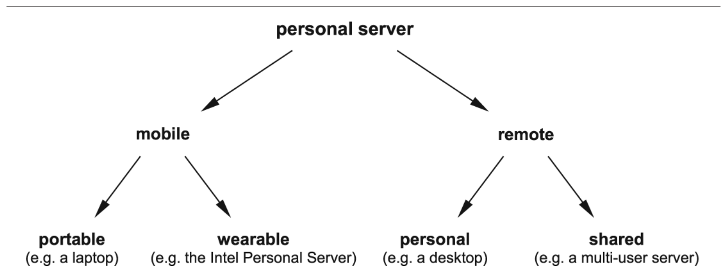

info
Volume 6 · Number 3 · 2004 · 172-179

suitable user access devices can be used. In this case, all that is needed is a short-range wireless network connection such as Bluetooth or peer-to-peer WiFi. In contrast, the wearable server (Segall, 2002) is small enough to actually be worn by the user at all times. This is likely to imply that there is no room for I/O peripherals, in which case all user interaction needs to be maintained via user access devices (just as in the case of the portable server on the move). In some cases, such as the aWare Messenger (Bylund and Segall, 2003), special-purpose devices can provide dedicated UIs to certain services running on the wearable server[10].

3.4 User access devices

In some cases, for example the portable server, the same hardware that executes the services also provides user I/O capabilities such as screen, keyboard, and mouse. When speaking of personal servers in general, though, this is not the case. Instead, users need to rely on user access devices in order to interact with their services. As suggested above, these should be chosen on the basis of the context and situation that the user is in at the time of usage. When sitting in an office at work, a desktop computer with a large screen, a full keyboard, and a mouse is probably one of the better alternatives. However, while on the move a smart phone or perhaps a headset and a head-up display is more suitable.

Examples of currently available user access devices are 2.5/3G smart phones, PDAs, Tablet PCs, and desktop computers in the form of Web kiosks. The strengths of the smart phones are their small form factor that makes them highly mobile, and the nearly constant network access over 2.5/3G connections. These two features are also the weaknesses of the device – being so small also makes them unsuitable for many tasks (for example word processing, reading text, and viewing images), and the network connection is unreliable, slow, and with high latency. Many PDAs are slightly bigger than smart phones. This makes them somewhat less mobile but at the same time more useful for certain tasks. PDAs are usually better equipped in terms of network connectivity – PDAs with both WLAN and Bluetooth, and in addition GPRS as an optional feature, have been available for some time. This further makes them more capable in a mobile setting since more network solutions are available. Tablet PCs resemble PDAs in how they are used, but they are much larger, usually with full-size displays. They are also more powerful and thus often capable of hosting services locally (i.e. acting as a portable server) in addition to acting as access devices. Finally, Web kiosks can be found in many public places ranging from cafés to airports and public libraries. For users without access devices of their own, these can provide access to remote servers with HTML UIs. They can also be used in cases where other access devices are too limited in terms of I/O capabilities.

3.5 Software support for personal servers

There have been few attempts to provide serious software support for personal servers as described herein. Alan Dearle describes ubiquitous environments (Dearle, 1998) which could work as a software platform for personal servers, and some technical problems (from a software point of view) that need to be solved in order to realize them. Dearle also lists a number of platforms that address at least parts of what is needed in order to implement them (for example Grasshopper (Dearle *et al.*, 1994), Telescript/Odyssey, and Aglets)).

For the Intel Personal Server (Want *et al.*, 2002a, 2002b), Want *et al.* have chosen to rely mostly on already existing support such as Web servers and file-sharing mechanisms as a means for connecting the personal server with user access devices. This makes the software platform for the personal server trivial – an ordinary OS (in their case Linux), a Web server, and services with HTML UIs are enough. This choice also makes the personal server quite limited. The number of access devices, for example, is limited in this case since they must have a Web browser. This also limits how well the UI can be adapted to the access device at hand, and hard buttons and other special I/O peripherals cannot be assigned service-specific functionalities.

In the sView project, however, we have developed a system (Bylund and Espinoza, 2000; Espinoza, 2003) that supports all aspects of personal servers as described above. The system builds on the notion of personal service environments (Bylund, 2001) that store electronic services and data of individual users. When residing on a computer, the services can be accessed locally via the I/O peripherals of the computer, or remotely via network-connected access devices (such as PDAs, cell phones, and Web kiosks). This allows the system to implement both mobile and remote servers. The service environments are mobile, which makes it possible to move all services to a different computer if needed. During the migration, the execution state of each service is saved in order to achieve continuous user-service interaction. This makes it possible to combine the qualities of mobile and remote servers by running the service environment on a server (acting as a remote server) when

info
Volume 6 · Number 3 · 2004 · 172-179

needed, and moving it to a mobile device (acting as a mobile server) when that is more suitable.

As part of our work with sView we have also developed solutions to support device-independent software components (Nylander *et al.*, 2004) and peer-to-peer communication between personal service environments (Espinoza and Hinz, 2003). The former is a key component when realizing continuity as in the case of adapted UIs (described above), while the latter is vital when implementing services that rely on direct communication between different personal servers.

4. Discussion and summary

Mobile computing devices were first introduced as a complement to traditional mainframes and personal computers. We can now sense a convergence between different kinds of computing devices – mobile and wearable computing devices are becoming powerful enough to assist users in ways similar to stationary devices. Stationary computing devices, however, are (and will continue to be) more powerful than mobile devices in most respects. But the convergence of stationary and mobile devices is not necessary for realizing a convergence from a user perspective – this could just as well be done on a virtual level. The important thing is that stationary and mobile computing devices are equipped in a way that allows for seamless switching between different modes of operation.

With networking, on the other hand, we foresee a complementary mix of wired solutions and a multitude of wireless solutions, in contrast to a complete paradigm shift from wired to wireless technologies (Sawhney, 2003). However, if this kind of complementary mix of technologies is to be successful, the means to perform all kinds of adaptation to variations in the quality of network connectivity must be available.

All in all, we argue that the concept of personal servers in general, and wearable servers in particular, is highly attractive as a general solution for providing seamless mobility. First, the concept allows services to be (partially) independent of network connection and quality. Second, the concept is built on the notion of UI and device flexibility as a means of treating the variation in user needs that a mobile setting offers. And third, the concept allows coexistence of different solutions for providing user experience continuity. Further, the implementation of personal servers may range from wearable computing to future generations of smart phones (Table I).

Nonetheless, since personal servers are such a wide concept, not all variants of them are equally well suited to handling all challenges of providing seamless mobility. In Table I we list a few such challenges in order to compare how well the two main categories of personal servers can handle them. The first two challenges (service access and autonomous execution) are mostly related to network independence. They address the solution's degree of dependence on network connectivity in order for the user to access services and for the services to operate autonomously. The last two categories (remote control and adaptive UI) address how well the solutions support continuity, as in the remote control case, or adaptive UI respectively.

When it comes to shifting focus from network technologies to a more holistic view of seamless mobility, though, technical solutions in isolation will not suffice. This is illustrated by the involvement of the European Commission (EC) in deploying GSM, which was driven by an urge to harmonize cellular networks within Europe. This technical goal was motivated by political goals of reinforcing European integration, which eventually would lead to a boost in competition since the deployment of GSM would break up old (national) PTT monopolies. As a side-effect, region-wide instead of national players would be able to enter the stage, which would lead to the possibility of challenging industry leaders in the US and Asia (Steinbock, 2002). While the EC, to a great extent, was successful in this quest, the strength of the region-wide players has turned out to be counterproductive for the development of seamless mobility. The strong telecom operators have the power to shut off local competition in mobile service deployment and provisioning, by exercising strict control over the service portfolio offered to their customers. Therefore, technical solutions to seamless mobility, such as the one offered in this paper, are doomed to be crippled as long as this situation prevails. For this reason, we can see an opening for regulators to support the same kind of initiatives that was once exercised to create the strength and dominance of a few players, in order to allow for solutions that provide true convergence not only of network technologies, but also of the functionality and user experience offered to end-users.

However, the most important catalyst to technical innovation cannot be provided by regulators – but rather by the market leaders of the telecom sector. In order to really boost the evolution of seamless mobility, these players must dare to face competition to a greater extent than we see today. By competing in excellence on an open

Table I A comparison between the two major categories of personal servers with regard to some seamless mobility challenges

	Mobile server	Remote server
Service access	Great opportunities since services can execute locally	Limited opportunities since a network connection is required in order to access the services. Off-line execution is not possible
Autonomous execution	Good opportunities, but services requiring a continuous network connection may suffer	Great opportunities since services can rely on a continuous high-speed network connection
Remote access	Great opportunities since remote control UI typically require low latency and at least medium high bandwidth	Good opportunities, but with connections between server and user access device with high latency and/or low bandwidth, the performance may suffer
Adaptive UI	Equal opportunities. However, services with modest or no need of networked data that output large amounts of data to the user (e.g. single-user games) benefit from mobile server execution. On the other hand, services that require large amounts of networked data but only output a limited amount of data to the user (e.g. a personal search engine) benefit from remote server execution	

(technical) arena, instead of locking up customers in closed and proprietary solutions, the room for technical innovations that create seamless mobility grows. This can be exemplified by the rise of the WWW. Before the Web, transfer of information between different systems was both impractical and costly, due to great differences in both hardware and software (Berners-Lee, 1996). Hypertext as a solution to some of these problems had been a hot topic in both academia and industry for several years, and several proprietary products had been proposed (for example Apple's HyperCard). However, it was not until after the WWW, proposed by Tim-Berners Lee at Cern (Berners-Lee *et al.*, 1992), had started to grow in popularity simply by the force of its own simplicity and openness, that market leaders really started to believe in business models building on the concept. Thus, what is essentially needed in order to rise above the present situation is for the industry to believe that this can happen again – that there is a market in open solutions to achieving seamless mobility.

In summary, we have expanded the concept of seamless mobility to include UI/device flexibility and user experience continuity. The novel concept of a personal server has, for the first time, been suggested as an approach to achieve fully-fledged seamless mobility. We have presented a taxonomy of personal servers and performed experiments to characterize their mode of operation and performance requirements. Nevertheless, there is substantial research work yet to be done. To mention a few outstanding research issues: development of measures of user experience continuity, new UI and devices that are designed for UI flexibility, and new operating systems' capabilities that are supporting the personal server concept and its modes of operation.

Notes

1 See also proceedings from the Tokyo Mobile Roundtable 2002 and the Stockholm Mobility Roundtable 2003.
2 Bandwidth, for example, can easily vary by a factor of 1,000 depending on network connection (from GPRS with less than 100kb/s to wired networks with 100mb/s or more). Variations in latency are almost as dramatic – a factor of 100 can easily be found (from several seconds with GPRS down to milliseconds with wired networks).
3 A Web page is typically made out of several dozens of data entities which all need to be fetched with HTTP requests. For example, it typically requires several minutes to load the first page of the Swedish newspaper Dagens Nyheter (a total of about 300kB and several dozens of entities to download) to a Sony Ericsson P800 smart phone over a GPRS connection. Loading the same page on a desktop computer with a high-speed wired connection requires less than a second. By introducing a proxy close to the wired-wireless border, the performance of GPRS can be improved (Chakravorty and Pratt, 2002), but the gap to WiFi technologies and wired networks is still huge.
4 A simple Java application running on the phone made a series of HTTP requests with the phone connected to the Internet via GPRS as well as Bluetooth and USB cable via a laptop computer (which was connected to a high-speed wired LAN). The response time between the network connections turned out to be about equal for Bluetooth and USB but many times higher for GPRS. A single HTTP HEAD request, which included less than one kilobyte of content, required several seconds to complete when operating via GPRS. When the request was changed to HTTP GET, which in addition downloaded a 55-kilobyte entity, response times were raised to 15 seconds on average. The response times for Bluetooth and USB were in the range of a few hundred milliseconds (about two seconds for the GET request). A complete description of these tests can be found in a separate technical report (Bylund, 2003).

5 These variations could be explained by variations in GPRS channel availability, but also TCP retransmissions due to the high latency of the connection (Chakravorty and Pratt, 2002).
6 Applications such as the Virtual Network Computing (VNC) client (Richardson *et al*., 1998) and Microsoft's Remote Desktop Connection can provide exactly this. Thin clients who are small enough to execute on PDAs and smart phones (in addition to desktop/laptop computers) mirror the screen, keyboard, and mouse of another computer.
7 In Nylander *et al*. (2004), we describe a methodology, including a fully functioning system, which reduces these difficulties when adapting UIs to different types of devices.
8 The second aspect should be put in contrast to something that is (possibly poorly and incompletely) duplicated on several different devices, such as the calendar on the desktop computer, the one in the PDA, and the one on the screen fridge at home.
9 Other definitions of the personal server concept exist, such as that of Want *et al*. (2002a, b) for example. Our concept is somewhat more comprehensive than that of Want *et al*., although theirs is included in the personal server and identified as a wearable server.
10 So far, there are few examples of hardware that can realize wearable servers of this kind, at least if high requirements on performance are placed. One exception is the Intel Personal Server (Want *et al*., 2002a, b), which is a computer of about the size of a deck of cards. It comes completely without user I/O peripherals, but includes a Bluetooth module that makes it possible to connect the server to user access devices.

References

Almgren, G. (2003), "Roaming between wireless ISPs: closed club or commons?", Stockholm Mobility Roundtable, Stockholm, May 22-23.

Berners-Lee, T. (1996), "WWW: past, present, and future", *IEEE Computer*, Vol. 29 No. 10, pp. 69-77.

Berners-Lee, T., Caillau, R., Groff, J.-F. and Pollermann, B. (1992), "World Wide Web: the information universe, electronic networking: research", *Applications and Policy*, Vol. 2 No. 1, pp. 52-8.

Bylund, M. (2001), "Personal service environments – openness and user control in user-service interaction", PhLic thesis, Uppsala University, Uppsala.

Bylund, M. (2003), "An empirical evaluation of the performance of mobile network connections", *SICS Technical Report T2003:06*, Swedish Institute of Computer Science, Kista.

Bylund, M. and Espinoza, F. (2000), "sView – personal service interaction", 5th International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology (PAAM'2000), Manchester, April 10-12.

Bylund, M. and Segall, Z. (2003), "Seamless mobility with personal servers", Stockholm Mobility Roundtable, Stockholm, May 22-23.

Chakravorty, R. and Pratt, I. (2002), "Performance issues with general packet radio service", *Journal of Communications and Networks*, Vol. 4 No. 2, pp. 266-81.

Dearle, A. (1998), "Toward ubiquitous environments for mobile users", *IEEE Internet Computing*, Vol. 2 No. 1, pp. 22-32.

Dearle, A., di Bona, R., Farrow, J., Henskens, F., Lindström, A., Rosenberg, J. and Vaughan, F. (1994), "Grasshopper: an orthogonally persistent operating system", *Computing Systems*, Summer, pp. 289-312.

Espinoza, F. (2003), "Individual service provisioning", PhD thesis, Stockholm University and Royal Institute of Technology, Kista.

Espinoza, F. and Hinz, L. (2003), "Generic peer-to-peer support for a personal service platform", Saint, IEEE Computer Society, Orlando, FL.

Hansen, J.S., Reich, T., Andersen, B. and Jul, E. (1998), "Dynamic adaptation of network connections in mobile environments", *IEEE Internet Computing*, Vol. 2 No. 1, pp. 39-48.

Kortuem, G. and Segall, Z. (2003), "Wearable communities: augmenting social networks with wearable computers", *IEEE Pervasive Computing*, Vol. 2 No. 1, pp. 71-8.

Nylander, S., Bylund, M. and Waern, A. (2004), "The ubiquitous interactor – device-independent access to mobile services", paper presented at the Computer-Aided Design of User Interfaces (CADUI'2004) Conference, Madeira, January 13-16.

Perkins, C. (1997), "Mobile IP", *IEEE Communications Magazine*, Vol. 35 No. 5, pp. 84-99.

Richardson, T., Stafford-Fraser, Q., Wood, K.R. and Hopper, A. (1998), "Virtual network computing", *IEEE Internet Computing*, Vol. 2 No. 1, pp. 33-8.

Sawhney, H. (2003), "WiFi – networks and the rerun of the cycle", *Info*, Vol. 5 No. 6, pp. 25-33.

Segall, Z. (2002), *Wearable Server*, University of Oregon, Eugene, OR.

Steinbock, D. (2002), "Wireless R&D: from domestication to globalization", *Info*, Vol. 4 No. 6, pp. 27-49.

Thorngren, B., Andersson, P. and Boman, M. (2003), "Seamless mobility", position paper, Stockholm Mobility Roundtable, Stockholm, May 22-23.

Want, R., Borriello, G., Pering, T. and Farkas, K.I. (2002a), "Disappearing hardware", *IEEE Pervasive Computing*, Vol. 1 No. 1, pp. 36-47.

Want, R., Pering, T., Danneels, G., Kumar, M., Sundar, M. and Light, J. (2002b), "The personal server: changing the way we think about ubiquitous computing", 4th International Conference on Ubiquitous Computing (UbiComp'2002), Göteborg, September 29-October 1.

Weiser, M. (1991), "The computer for the 21st century", *Scientific American*, Vol. 265 No. 3, pp. 94-104.