MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Laboratory for Computer Science

Personal Information Architecture Note 1 12/1/93

BodyTalk and the BodyNet: A Personal Information Infrastructure

Olin Shivers
shivers@lcs.mit.edu

Abstract

The current evolution of personal information appliances, such as cellular telephones, personal digital assistants, and notebook computers, can be made more effective if re-structured into a personal network architecture. This architecture is based upon two central components: a hardware communications system, the *BodyNet*, and a common interface language, *BodyTalk*.

This architecture will allow personal digital devices to communicate and cooperate in powerful ways, and open a path to portable, wearable network interfaces that are as comfortable and personal as an old and familiar sweater.

1 Introduction

Suppose that seated at the restaurant table next to you is a businessman, armed with all the hi-tech acoutrements that modern tech-oriented professionals cannot seem to live without: a pager, a cellular telephone, and a "digital diary." If this man drives a car, he also will have a remote arming device for its anti-theft system in his pocket; if he rides public transportation, he will be carrying a walkman for the ride. A notebook computer is a likely possibility, and we might even consider his digital watch to be a part of the portable suite of technology he is carrying around.

The first thing to notice about all of these gadgets is how poorly they interact with one another as well as with their master. In order to return a

	display	keyboard	pen	spkr	mike	memory	cpu	xcvr
watch	√	?				?	?	
pager	√					√		√
cell phone	√	√		√	√			√
walkman				√		√		
walkman/fm	√			√		√		√
diary/PDA	√	√	√	?		√	√	?
notebook	√	√	√	?		√	√	
⋮								

Figure 1: Functional duplication in personal appliances

phone call, our businessman must read the phone number off the pager's display and type it into the cell phone's keyboard. He is, in effect, acting as a low-bandwidth communication channel for two digital appliances. He must do likewise for phone numbers stored in his digital diary. If he receives a phone call while listening to his walkman, he has to remove one set of audio hardware in order to use a different one.

A second thing to notice is how much functional duplication there is among these components. Every single digital component has a keyboard and a display of some sort—although their owner can use only one keyboard at a time, and can only attend to a single display at a time. Counting the pager, the cellular telephone, and the walkman's FM receiver, our businessman is also carrying three independent communications systems.

All in all, our businessman is carrying two headsets, three communications systems, five keyboards and six displays—that's one headset, two communications systems, four keyboards and five displays too many (fig. 1). The main reason for all of this duplication is, again, the lack of cooperation among these devices. As a result, each one must implement a complete set of functions and interfaces for the user. This lack of cooperation and resultant functional duplication go well beyond the hardware to the software side of these systems: since the cellular telephone and the digital diary can't cooperate, they must both implement incompatible phone number databases.

If we could somehow get all of these devices to interoperate, we could eliminate the functional duplication, reduce the systems complexity, lighten the burden of equipment, and still increase the overall utility.

2 The BodyNet

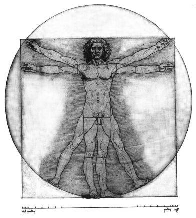

What is needed is a personal communications infrastructure: a low-power wireless communications system with a limited (*e.g.*, six-foot) range. Six feet is just enough range to get a signal from head to toe, and thus link together all of the various gadgets that someone would carry on his person. By restricting the range, we can reduce packet collision from nearby wireless nets belonging to other people, and reduce the power requirements of the transceivers, allowing them to be small and cheap.

A wireless net of this kind allows us to restructure our assembly of personal gadgets along a network architecture. Once devices can help each other out, each bit of hardware on the BodyNet doesn't have to do everything poorly. It needs only concentrate on providing its single basic function well. Other services will be provided by other components on the net.

Under this approach, our example personal system would be reconstructed from a collection of the following cooperating components:

Computer A user only needs a single small computer, presumably carried someplace convenient, such as on the belt or in a purse or briefcase. This computer will provide local database and computing services to any other device on the BodyNet. The computer is a featureless box, perhaps the size of a pack of cigarettes. It would have neither display or keyboard, performing all I/O over the BodyNet.

Communicator A cellular telephone or some sort of network link is an important part of the system. A wireless network link would also be a small, featureless box, tucked away someplace convenient. Such a system doesn't have its own database of phone numbers, or complex operating modes. Nor does it have a keyboard, display, microphone, or speaker. It simply provides wide-area communication services to the rest of the BodyNet. These communication services would be used by applications such as voice telephony or voice mail, email, stock market trading, banking, and remote database access. The privacy and authentication services needed by some of these communications applications would be provided by a computer on the BodyNet.

Visual display Depending upon circumstances, a user might choose one of several alternative displays. The surface area of a wristwatch could be used for a small display. Such a display system would be simply a BodyNet receiver and a display. Time could be displayed on it by the processor node as one of the more trivial BodyNet applications. A page-sized LCD tablet could be carried for more detailed display. A pair of "magic glasses"—a small, high-resolution display built into a pair of glasses, rear-projecting onto the lenses—could well provide the ultimate general-purpose display. Magic glasses are in some sense the "general purpose" display because one never has to look in a particular direction to attend to the display. No matter where the user is looking, the information comes to him.[1] Note that these displays do not require local computation or memory. They simply provide a single service—visual output—to the user's BodyNet.

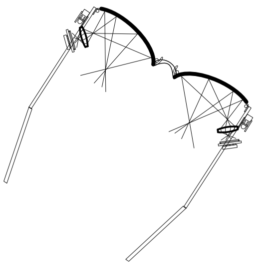

Audio I/O A pair of headphones, either on a headband or mounted on a pair of glasses, would provide an output device for digital audio streams on the BodyNet. Bone-conduction speakers could also serve as unobtrusive personal audio devices. Air microphones can be mounted on earpieces or mike booms, or clipped to collars. Direct-contact throat mikes can also be clipped around the neck, touching the throat just to the side of the voice box. Direct-contact mikes allow for externally sub-audible private voice input.

Gestural input Keyboards, chording devices, pen tablets, and virtual keyboards can all provide for gestural input, including text entry.

More exotic devices A key point of a network architecture is that it is extensible. The following items are all possible elements of a BodyNet

[1]Such a display is under active development at MIT[Alvelda]. The Personal Information Architecture project will seek to integrate these devices into our experimental prototypes as they become available.

ensemble:

GPS sensors A host of interesting applications becomes possible if the software can be aware of user's location. They include navigation, planning, restaurant recommendation, taxi dispatch, real-time displays of subway car locations at the closest subway station, and so forth.

Eye trackers Some existing consumer cameras already use built-in eye-trackers to determine the focus region for auto-focus. An eye tracker, mounted in a pair of glasses, could serve as a hands-free mouse or function in augmented-reality applications.

Head trackers Head trackers are vital for augmented-reality and immersive environment applications.

Medical monitors People with health problems might carry sensors for vital measurements such as pulse rate, blood pressure, blood sugar, or insulin levels.

Ring button A small ring, with one or two buttons mounted on the palm side, could be used by the user to give replies to simple queries. A ring button is the keystroke analog of magic glasses in that it is an input device that is always available: no matter where your hands are, the ring's button is immediately available.

Note that although the above list emphasizes the single-device-single-function approach to constructing a complete BodyNet system, many devices will quite likely bundle various capabilities together. An LCD pen tablet, for example, would provide both input and output capability. Many devices would probably have embedded processors capable of handling simple computational tasks independently. The important point is that these devices be able to cooperate and share resources.

2.1 BodyNet or monolithic device?

Instead of pushing for a network architecture like the BodyNet, we might instead push for a monolithic solution. Perhaps our hypothetical technologically-overwhelmed businessman might find salvation in a single digital appliance that does everything for him. After all, digital hardware gets smaller all the time, and the day may come when everything can be rolled together into a single package—a personal digital assistant with a cellular phone link and perhaps a few other sundry capabilities.

We do not believe in this eventuality over a broad class of users, for reasons of demand individualisation and supply dynamics. Individual users will persist in remaining individual—the system requirements of one will not satisfy the needs of another. The last two decades have taught us that the whims and preferences for user interfaces are quite individual. Proficient keyboard users may want a keyboard for text entry; people less familiar with computers may feel more comfortable with voice-recognition. Only diabetics will want to buy blood-sugar monitors, and so on, along the established patterns of mass individualised consumption that are expected to become even more dominant in the next century.

On the supply side, manufacturers will continue to innovate, producing ever more novel and powerful devices, increasing the plethora of hardware that drove our harried businessman to distraction in the first place. If the devices don't cooperate, "more" is not necessarily "better"—it's just more intractable complexity.

A network architecture, in contrast, is flexible. It allows users to adapt their systems to their needs and resources, which may vary over time. It also allows users to upgrade their interface hardware as new devices become available.

With the BodyNet architecture, the complete system that is the user's wearable computer is actually something that is emergent from the collection of the interface devices, local computational resources, communications links and managing software that each user chooses to carry on his person. In this sense, each system could be as individual as a fingerprint.

Factoring the capabilities and resources of the user's personal interface into their primitive functional components allows them to be connected back together in multiple and often initially unforseen ways. For example, in the BodyNet, a "cellular telephone" is now something that arises from three basic pieces: a wireless network link, an audio interface, and an application that uses these to provide a "cell phone" service for the user.

This emergent structure allows the power of the devices to be multiplicative, rather than additive. For example, adding a communications link to a BodyNet system provides wide-area network capability to *all* the other devices in the system, and to all the applications running on these devices. Our communication link can be used in other ways than as a cellular telephone—email, stock market, plane reservation, entertainment, and other applications can also use the link.

3 Beyond Personal Computing

A personal digital interface is not just an interface to a single computer, with a single box's worth of information. In a networked environment, it is an interface to the world. While a BodyNet system is a useful entity as a standalone system, once a communications link to the global Net is included it becomes tremendously powerful—an ant on a vastly more complex beach.

The rise of the global digital network as the world's primary communications medium in the next decade or so is a development that rivals the introduction of the telephone as a technological agent of social change. In this picture, the BodyNet serves to span a gap of enormous scale: on one side will be the global Net, with its terabytes of data, billions of users, teraflops of computation, and untold aggregate bandwidth; on the other side will be the individual user navigating this information sea. The BodyNet interface and the applications that hook it into the global network are the technology that will match the planetary-scale resources of the Net to the human-scale needs of the individual.

4 Intimate Interfaces

A computer interface should be like a wristwatch—always there, always on. You don't have to take it out, turn it on, and boot it up. You are simply always on-line, always connected. A wristwatch is an example of a technology that is sufficiently useful that most people in modern societies are willing to attach one permanently to their body. The device is comfortable and useful enough to warrant this.

A computer/communications system, on the other hand, is much, much more useful than the simple time service provided by a watch. Suppose that such a system were also to become comfortable enough to warrant being treated like a watch? This kind of "intimate interface" to the world-wide information infrastructure would make a qualitative shift in the way we interact with our communications systems and their services. The BodyNet project is intended to explore this new territory.

The wristwatch analogy is interesting for another reason, as well. For hundreds of years, wristwatch manufacturers were engaged in competition to produce more and more accurate wristwatches. However, the advent of the modern quartz-crystal movement has effectively shut down the push for new accuracy in portable time-keeping devices. To make an accurate watch is not a technological challenge in the late twentieth century.

As a result, today's watches are chosen not so much for their time-keeping quality, but for their decorative value. This move from high-tech to jewelry is the sign of a mature technology. It took watches several hundred years to make this transition completely. Computers will quite likely do it within the next decade.

5 BodyTalk

Integrating disparate information appliances into a cohesive personal system requires more than communications hardware. A common interchange language is also necessary. The Internet, for example, is more than wires and routers—it is the definition of the TCP/IP communications protocol that allows systems to communicate over the network hardware. The MIDI language serves a similar function in the music industry, as the *lingua franca* that allows musical systems to interoperate. Just as TCP/IP enabled the current global explosion that is the Internet, MIDI enabled the musical appliance industry to blossom and grow into a multi-billion dollar market.

In order for displays, microphones, pen tablets, wristwatches, communicators and the rest of the personal appliance bestiary to cooperate and be unified into a single, flexible system, a similar standard needs to be defined. Let us call such a language "BodyTalk."

BodyTalk would insulate applications from the variety of interfaces permitted by the BodyNet architecture. For example, suppose an application needs a yes/no input from the user. There are many, many ways that a user might return that one bit of information: typing "y" on a keyboard, saying "yes" into a microphone, rolling his eyes towards a "yes" box projected onto a pair of glasses, or touching the circumference of a ring button.

Which of these various forms of reply a user would choose would depend upon what particular devices that he had purchased, and perhaps upon other conditions as well, such as privacy. Voice input, for example, is very convenient, but it is not possible in a movie theatre or a crowded office. It also requires more computational resources than some users might be willing to buy. The point is that there is no way for an application to have built into it all the different kinds of interface permitted by a flexible network architecture such as the BodyNet. The BodyTalk standard must insulate applications from this variation of interface modalities.

BodyTalk must be broad-spectrum, allowing new devices to be easily integrated into a user's personal system. For example, if the user picks up a camera, it should automatically insert itself into his BodyNet, announce its type, describe what sorts of operations it supports, and perhaps make its user manual or help system available for the user to call up as needed. The user should be able to switch the camera's information displays to his glasses. If the user has the software and enough processing power on his BodyNet to support speech recognition, he should be able to simply say "Aperture priority, fill flash," and have the camera respond appropriately.

BodyTalk should "open up" the internal architecture of the devices it connects. For example, the interface on a modern-day cellular telephone is limited by the imagination and design skills of the telephone's designer, who can never anticipate all the useage patterns of every user. However, if the telephone exposes its basic functional elements to external control, then individual users can customize their interfaces, program common patterns, and choose control software written by third-party developers for particular niche markets. In this sense, BodyTalk provides an "open architecture" within devices as well as between components.

Opening up a device in this way raises issues of security—for example, given complete control of a cellular phone, a user can listen in to other channels, or make phone calls that are billed to another user. This does not mean that devices shouldn't expose their internal functionality, only that the exported interface should be carefully chosen to prevent misuse.

BodyTalk is not a universal applications protocol. BodyTalk specifically focuses on the management of *human interfaces*. It is used by applications to interact with a human user. Thus the domain of BodyTalk includes text I/O, visual displays, and image capture. It does not include the semantics of the applications. Phone calls, nutritional advice, stock market transactions, mortar fire control: these things are outside the language, and are defined by applications that run on top of BodyTalk.

Although a businessman has been serving as our motivating example, the utility of a wearable computer and a language that can coordinate the various devices worn on the body extends beyond this arena. Whether it is a doctor in an operating room, a soldier dropping rounds into a mortar, a vacationer enjoying audio/video entertainment, or a shopper selecting food from the supermarket for later delivery, all of these individuals need both to efficiently manage their personal hardware and to interface to a larger communications system.

The commercial development of BodyNet systems will also provide a particular efficiency to the more specialised applications areas. Military and medical systems are typically very expensive; however, proliferation of the technology in the consumer arena will drive the price of the technology down to very low levels across the board.

6 Research

The design of BodyNet and BodyTalk raise a number of research questions and issues.

- **Language design**
  The basic linguistic structure of BodyTalk language is a challenging problem. The flexibility required by the various interfaces argues for some kind of an object-oriented approach. The message-passing architecture of object-oriented systems allows for different devices to handle messages in different ways—so our "yes/no" example above can be handled appropriately by interfaces of various modalities.

  The distributed nature of the BodyNet also argues for a message-passing paradigm. It is a very local distributed system, with nodes that come on line and drop off the net dynamically—quite possibly during the execution of a single task. An actors models seems well adapted for this kind of computational environment. For example, just the simple act of picking up a piece of hardware such as a keyboard and having it automatically insert itself into the user's BodyNet would cause the runtime environment to reconfigure itself.

- **Interface taxonomy**
  What is the right taxonomy to describe I/O done with various interfaces? This will be the single most challenging aspect of the design effort.

- **Encryption**
  A BodyNet user doesn't want his email showing up on his neighbor's display. Nor does he want his neighbor's keystrokes interfering with his typing. Preventing eavesdropping, and keeping adjacent BodyNets from interfering will require some sort of encryption.

  Cryptographic authentication will also be very important in the role of BodyNets as intimate interfaces to the global Net and its services. In order to place long-distance phone calls, access bank accounts, bid on contracts, play the stock market, bet on horses, or send love letters over open radio links, some form of secure communications protocol will be required.

7 Related Work

The area of advanced mobile interfaces to computing services is a fertile one, strongly pushed by the explosive growth of the global Net as a motivating content provider. There are several related research efforts currently underway, discussed below.

7.1 Ubiquitous computing

The notion of having computers scattered into the nooks and crannies of everyday life in a cooperative fashion was stated as early as 1976[Newell]. The well-known Ubiquitous Computing project at Xerox PARC[UbiComp] has undertaken this effort. Aimed at weaving computation into the fabric of our physical world, the Ubiquitous Computing project has focused on three primary interfaces, the so-called "Tabs," "Pads," and "Screens."

However, PARC's Ubiquitous Computing project does not have the same focus on modular extensibility that is central to the BodyNet project. The Ubiquitous Computing project is primarily concentrated on interfaces to a large network of computing and computing services; its central research goal is coordinating these wide-area resources to provide a consistent view to the user. Note also that a BodyNet system does not necessarily need to act as an interface to the global Net. BodyNet systems without external links can still be valuable, useful tools.

7.2 Wearable computers

Sieworek's Vu-man and Navigator projects have constructed prototype wearable computers at Carnegie Mellon University. The Navigator system is a belt-mounted PC-class computer coupled to a head-mounted display used for navigation on the Carnegie Mellon campus.

Again, the Navigator project is focussed on the design of a particular wearable computer, while the BodyNet project aims to allow for an infinite variety of wearable computers to arise by the simple assembly of hardware resources, with cooperation mediated by the BodyTalk standard. The modularity and consequent extensibility of the BodyNet architecture are some of its distinctions.

7.3 Virtual reality

Virtual reality is a research area that has captured the popular imagination in the last few years. Mention of head-mounted displays usually brings to mind notions of immersive environments or augmented-reality applications. However, these applications, like network interfaces, are not essential to the idea of the BodyNet, which can function quite usefully with much simpler interfaces.

Virtual reality is hard; hard in many ways that are irrelevant to the basic premise of BodyNet interfaces. Several difficult basic problems, such as precise, low-latency head and eye tracking, need to be solved in order to deliver good virtual-reality and augmented-reality environments. VR applications also require major computational resources to handle their rendering, sensing, and simulation requirements—perhaps more computational resources than will portably and affordably be available in the near-term. In this respect, VR capability can be viewed as a cherry on the BodyNet sundae: users who want VR capability can invest in the necessary hardware to support it, whereas users who wish to economise or travel light can still get a lot of utility from simpler systems.

In many cases, the resources necessary to support VR applications can be provided by local, sessile systems. We can reasonably assume that most people will have much more serious computer systems in their offices and livingrooms than they carry on their belt. These systems can take care of the compute-intensive tasks that a VR application requires. Moreover, if a set of reflective spots is placed on eyeglass-mounted displays, then ceiling-mounted trackers might be able to do the necessary head-tracking. Reflective spots are cheap and don't weigh anything. Thus when the user is walking down the street, he can use his glasses display as a simple interface, but as soon as he walks into his office, he can switch over to a sophisticated augmented-reality interface, by having his glasses function as simple front-end for his office VR system.

This kind of architecture is exactly backwards from that of current notebook-based systems. With current notebook architectures, a user carries around a single computer, but accesses it through various interfaces: on the road, he uses a lightweight LCD display and a cramped keyboard; in his office, he plugs the same computer into a docking station and uses a better display and keyboard. A BodyNet system allows the user always to have the same interface hardware, and to use it to access computational resources of varying degrees of power. Instead of hooking the same computer up to different interfaces, we hook the same interface up to different supporting computers.

7.4 Wide-area network systems

There are many research projects underway to explore the architecture of the coming global digital network and its services. The I95 project[I95], underway at MIT, is one of the more large-scale efforts to provide a substrate for general-purpose, wide-spectrum use of the Net: business transactions, entertainment, communication and so forth. Singapore's IT2000 project, based on some of the early I95 work, and the SmartValley project are other efforts in this vein.

The BodyNet project intends to use these sorts of projects, particularly the I95 project, to find driving applications for the interface it will provide. Your primary interface to the global Net is the one that is always with you, so BodyNet-based front ends to network service and content providers are certainly worth exploring.

In a similar vein, the Livingroom of Tomorrow project being started by the Personal Information Architecture group will also be concerned with providing an interface to the avalanche of bits soon to be pouring into modern homes. The BodyNet project will be in close cooperation with the LRoT project to develop these two different interfaces to the global Net in harmony.

8 Linkages

The BodyTalk project will benefit from important linkages, both inside and outside MIT. We intend to use these linkages both to influence the project's course and to transfer its results to the outside world.

8.1 Inter-laboratory links at MIT

The BodyTalk project is being undertaken by the Personal Information Architecture group, a research group at MIT that bridges between the Laboratory for Computer Science and the Media Laboratory. As a result, the project can combine the traditional systems and architectural strengths of LCS with the Media Lab's understanding of digital media and interface design. We expect to recruit graduate students from both labs, and to draw upon the research expertise of both centers.

Inside LCS, we expect to form linkages with many of the other research groups. Zue's Spoken Language Systems project can provide speech understanding support for BodyNet interfaces. The network infrastructure being developed by Clark's Advanced Network Architecture group and the nascent I95 project will provide driving applications for the "intimate interfaces" the BodyNet will develop. The Telemedia and Library 2000 efforts will also deliver content for the BodyNet applications to manipulate. An overview of these research projects can be found in the LCS annual progress report[LCS].

8.2 The BodyTalk Consortium

In order for BodyTalk to be successful, it must be adopted by many diverse manufacturers of digital devices and software—the power of BodyTalk comes from its ability to link varied systems. So it is of vital importance that BodyTalk be broad-spectrum and vendor-neutral. By undertaking this project in a university setting, we hope to act as impartial designers, providing a continuously upgraded specification that will be vendor-neutral. Ensuring that the language is broad-spectrum is a research and design problem, one that will benefit from consideration of a wide array of potential interface devices. The Laboratory for Computer Science intends the specification of the BodyTalk standard to be open by making it available to all interested parties, but will retain control of the definition to keep it well-defined and coherent.

To develop and promote the BodyTalk standard, we intend to form a BodyTalk consortium of interested industrial partners. We are interested in having as consortium members manufacturers of consumer electronics, computers, interface hardware, applications software, and communications systems, as well as other relevant industrial concerns working with us to develop the standard. Members of the consortium will be able to represent particular interests in the various interface functionalities and modalities, and influence the design of the protocol.

The BodyTalk Consortium will also act as a two-way conduit for transferring the technology out from LCS into the commercial arena, and for bringing specific feedback from users and manufacturers into the Laboratory. Workshops and conferences will be held on a regular basis to disseminate the BodyTalk specification out into industry. LCS will capitalise on the experience gained by its successful X Consortium to manage the BodyTalk Consortium.

9 Proposed Project

We propose to develop a BodyNet system, with a view towards establishing it as a commercial standard. The principal result will be the definition of a communications specification, covering both a hardware transmission layer, BodyNet, and a language layer, BodyTalk. Demonstration systems and applications will be constructed as test beds for the ideas in BodyTalk and as proofs of concept.

9.1 The InfoSpigot

The initial, driving application will be *InfoSpigot*, a mobile interface to the World Wide Web[WWW, WWW-Intro]. This is an interesting application for several reasons:

- It involves using the BodyNet as a mobile interface to the general resources of the Net. This forces the design to handle issues of external communications links.
- It can take advantage of recent linkages established between the WWW project and LCS. As WWW development progresses, both at MIT and CERN, the BodyTalk effort will be able to profit by and influence new capabilities in the Web.
- It leverages off of the sizeable and growing amount of data available through WWW.
- WWW is a multimedia system, handling sound, images, and text. BodyNet is about interface, so the richer the set of modalities the application has to handle, the more completely the BodyNet design will be exercised.
- WWW has hooks to allow general server-side processes to provide arbitrary services and database accesses. There are plans at MIT to further extend this support, which will allow a successful implementation of the InfoSpigot to generalise into an even more powerful tool.

9.2 The Modular Phone

Another application will be a cellular telephone system implemented as an open collection of cooperating hardware and software agents on a BodyNet. The user should be able to interface his audio hardware to his office phone, or to use a cellular communicator when mobile—all with the same basic interface. The hardware design will emphasize portable, wearable components, such as throat-contact mikes and belt-mounted communicators. The emphasis will be on components that could comfortably be worn a large portion of a working day.

The system should allow for voice control, keyboard input, or control by software agents. As an open system, it should be easy to expand or customise the interface in ways that current cellular telephones do not allow:

- Plug in phone databases from other systems.
- Program the phone to try a sequence of phone numbers when calling someone. (*"First try the office phone. If the voice mail answers, hang up and try the cell phone. Then query me to see if it is OK to try the home phone."*)
- Intelligently strip local area codes and prefix required 0's or 1's from phone numbers, depending upon location and roaming system requirements.
- Quickly shift between various pre-set modes, based upon the user's common venues. (*" 'Movie' mode is for movies and my lectures: do not ring, but record a message. If a discreet information channel, such as a vibrating buzzer, a pair of headphones, or a glasses-mounted display, is available, then give me an indicator."*)

Any of these features could be offered by a manufacturer of cellular telephones, but by factoring the architecture into its basic functional modules, we can allow for arbitrary extension and customisation in software.

Note that pieces of the Modular Phone system should be applicable to different applications, such as remote email, document faxing, voice notebooks, or "walkman" music applications.

9.3 Applications payoff

The significance of the InfoSpigot and Modular Phone implementations is not the applications themselves, but the presence of a realistic environment that exercises the design decisions made for BodyTalk and BodyNet. The BodyTalk design should allow the applications to run successfully with a variety of interface modalities: speech recognition, keyboard input, headphones, visual displays, and so forth. The user's choice of interface should be insulated from the application by the BodyNet. The InfoSpigot implementation will be useful because it will place stress on this design goal, so that cracks in the solutions will show early. The Modular Phone implementation will be useful because telephones admit multiple different interfaces, and because useful system building blocks should emerge from its development. The components for managing voice I/O that are developed for the Modular Phone should be generally applicable to other tasks.

10 Benefits

BodyNet technology is pure infrastructure. Like an airport or a regional power grid, its benefit lies in what it enables.

- BodyNet will enable a new industry, much as the development of the MIDI communications protocols enabled multi-billion dollar growth in the musical device industry.

- It will be a force for systems integration at the device level, transforming the power of individual devices by allowing cooperation.

- It will change the nature of human/computer and human/network interaction.

If a BodyNet-like architecture is not developed, then personal information appliances will continue to evolve into balkanised, insular devices with inconsistent interfaces, limited functionality, and lack of cooperation. If BodyNet is developed, then the path is opened to a new generation of personal user interfaces.

References

[Alvelda] "VLSI μDisplay Technology," Phillip Alvelda and Thomas Knight. MIT AI Laboratory, ARPA Grant No. N00014-94-1-0249.

[I95] *The Viewstation Collected Papers*, ed. David Tennenhouse. LCS TR-590 MIT LCS, November 1993.

[LCS] *Laboratory for Computer Science Progress Report 30*, MIT LCS, June 1993.

[Newell] "Fairytales," Inaugural speech, Allen Newell, 1976. Copies are available from the Carnegie Mellon School of Computer Science.

[UbiComp] "The Computer for the 21st Century," Mark Weiser. *Scientific American*, September 1991, pp. 94–104.

[WWW] "The World Wide Web Initiative: The Project," URL `http://info.cern.ch/hypertext/WWW/TheProject.html`. Tim Berners-Lee, CERN.

[WWW-Intro] T. J. Berners-Lee, R. Cailliau, J-F Groff, B. Pollermann, CERN, "World-Wide Web: The Information Universe", published in *Electronic Networking: Research, Applications and Policy*, Vol. 2 No 1, pp. 52-58 Spring 1992, Meckler Publishing, Westport, CT, USA.

Also available in Postscript form from `info.cern.ch` as file `/pub/www/doc/ENRAP_9202.ps`.