

◆ Using Data on Digital Cellular and PCS Voice Networks

Shankar Narayanaswamy, Jianying Hu, and Ramanujan S. Kashi

Many new digital cellular and personal communications services (PCS) networks that provide both voice and data communication capabilities are being installed in the United States and around the world. Current cellular and PCS handsets primarily target voice telephony, but the next generation will also support data services. In this paper we review the current offerings in wireless handsets that support both voice and data services and also present our design for a next-generation wireless handset. We discuss the tradeoffs in selecting an operating system and input device, and in designing a user interface. The physical design of our user interface and the wired network infrastructure support applications such as telephone directory, electronic mail, and Web browsing. The result is a small, user-friendly handset featuring handwriting recognition, foldable video graphics adapter (VGA)-sized liquid crystal display (LCD) screens, and the Inferno™ operating system.

Introduction

Since cellular telephones were introduced in the 1970s, their usage has grown by leaps and bounds. In the last decade, growth rates in the United States have ranged from 35% to 60% per year.¹ Concurrently, prices for cellular service have been dropping. The Personal Communications Industries Association projects a total of 88.3 million subscribers to personal communications services (PCS), paging, cellular, and other mobile radio services by 1998, climbing rapidly to 167.4 million subscribers by 2003.² The number of cellular digital packet data (CDPD) users is also growing, mainly in mobile sales and service areas of large corporations, despite the relatively high cost of CDPD terminal equipment. AT&T currently provides CDPD service in 70 major metropolitan areas across the United States.

Recently, new standards for digital cellular telephones have appeared, including:

- IS-136,³ based on time division multiple access (TDMA) technology;
- IS-95,⁴ based on code division multiple access (CDMA) technology; and

- GSM⁵ (Global System for Mobile Communications), based on frequency hopped multiple access (FHMA) and TDMA technology.

Systems using these standards are in the process of being deployed. All three standards feature data channels that support data communication at the wireless channel's data rate (19.6 kb/s for CDPD and up to 14.4 kb/s for IS-136, IS-95, and GSM). Without a data channel separate from the voice channel, data must be acoustically encoded and transmitted over the voice channel, as in wireline computer modems. The voice channel in a digital cellular telephone network uses speech compression algorithms that are not efficient for data, resulting in lower achievable data bandwidth and higher error rates. The new PCS networks are meant primarily for data services, and all the major service providers are deploying digital technologies that support both voice and data.

The Internet has experienced tremendous growth in the past few years. Its applications are classified into three categories:

- Communications, including electronic mail

(e-mail), electronic bulletin boards, USENET, and Internet voice telephony;

- Data transfer, including transfer of programs, data files, and documents using the file transfer protocol (FTP), wide area information server (WAIS), and Gopher; and
- Document/data retrieval, done on the World Wide Web, using a browser to display local and remote documents that are encoded in HyperText Markup Language (HTML).⁶ The hypertext transport protocol (HTTP) transfers files between computers over the (wired) Internet.

The information sources accessible via the Web are large and diverse. Many corporations maintain publicly accessible Web sites to disseminate information on their particular corporation and its products and operations. Information sources such as newspapers and magazines disseminate information through on-line services that enable free or subscription-based public access to their articles. Other content providers make documents or services publicly accessible on the Web. These providers include city and state governments, which place statutes and minutes from public meetings on the Web; universities, which make research results and publications available; and individuals, who maintain personal Web pages.

Internet search engines such as Yahoo!* support retrieval of Web documents via a keyword search. This is a powerful way to find information in cyberspace. Search engines are used by people in both technical and nontechnical fields to find documents containing information on topics of interest. The Web browser is also becoming a popular interface to corporate databases on corporate intranets and as a generalized document access mechanism.

In the past few years, the market for personal digital assistants (PDAs) has been growing. Several vendors now offer PDAs and related products. The International Data Corporation says that 5.1 million hand-held computers were sold in 1997, and sales are expected to rise to 8.2 million in 1998. PDAs currently on the market include the Newton* MessagePad,* 3Com PalmPilot,* Sharp Zaurus,* and the new Windows* CE devices. Their strengths lie in novel user

Panel 1. Abbreviations, Acronyms, and Terms

AMPS—Advanced Mobile Phone Service
ARDIS—Advanced Radio Data Information Systems
CDMA—code division multiple access
CDPD—cellular digital packet data
CT2—Cordless Telephone 2, a second-generation digital cordless system
DECT—Digital European Cordless Telephone
e-mail—electronic mail
FHMA—frequency hopped multiple access
FTP—file transfer protocol
Gopher—a program that searches for file names and resources on the Internet and presents hierarchical menus to the user.
GSM—Global System for Mobile Communications
GUI—graphical user interface
HCW—handwriting capture widget
HCWC—handwriting capture widget commander
HDML—Hand-held Device Markup Language
HDTP—hand-held device transport protocol
HMM—hidden Markov model
HPC—hand-held PC
HTML—HyperText Markup Language
HTTP—HyperText transport protocol
ICAP—interactive communicating applications protocol
ISDN—integrated services digital network
ISP—Internet service provider
JVM—Java virtual machine
LCD—liquid crystal display
PACS—Personal Access Communications System
PCS—personal communications services
PDA—personal digital assistant
PHS—Personal HandyPhone System
PIM—personal information management
PSTN—public switched telephone network
RAM—random access memory
ROM—read-only memory
SDC—stroke-direction code
SNA—simple network architecture
TCP/IP—transmission control protocol/Internet protocol
TDMA—time division multiple access
TEW—text entry widget
URL—universal resource locator
VGA—video graphics adapter
WAIS—wide area information server

interfaces; small, efficient operating systems; long battery life; a small size and aspect ratio, or *form factor*; and well-thought-out, integrated applications. Recently, several companies have introduced PDA accessories to support wireless communications such as the Minstrel* CDPD modem for the PalmPilot.

The developments we have described affirm that wireless providers are making mobile data access available and that consumers are buying mobile data services. It is therefore feasible to successfully build and market a wireless device that supports data applications over digital cellular and PCS networks.

This paper examines the issues and choices that arise when designing such a mobile terminal and its supporting infrastructure. We believe the key to success is useful, easy-to-use applications on a user-friendly hardware platform. Without useful applications, the user has little incentive to buy the terminal. Without a good user interface to the applications and a user-friendly platform, user acceptance will be low.

In this paper, we describe the current state of the art, provide some detailed background, and present our own solution, which is being prototyped at Bell Labs. We review some of the smart phones currently on the market, as well as the more successful PDAs, and take a look at wireless data service and the bandwidth available to applications under various digital wireless systems now and in the near future. We also describe some options for the software infrastructure of the terminal, such as the operating system and the communications protocol. In our discussion of user interface issues for portable wireless terminals, we cover both the terminal hardware and the interactive software. The presentation of our solution begins with a description of the high-level system design and architecture, including the infrastructure required to provide data services on the wired network. Next we present the terminal hardware and the user interface, followed by the software infrastructure that supports the user interface. We conclude with projections for future work.

Smart Phones

Currently, only two cellular telephones on the market provide data services—the Nokia* 9000

Communicator⁷ and the AT&T PocketNet telephone.⁸ The Nokia 9000 Communicator is a sophisticated palmtop computer built into a GSM cellular telephone. The AT&T PocketNet* telephone is a regular analog cellular telephone with a built-in CDPD modem. This section describes both phones in detail.

The Nokia 9000 Communicator, available in Europe, is a full-featured GSM cellular telephone capable of receiving faxes, e-mail, and short messages. Its GEOS* 3.0 operating system⁹ runs on a 24 MHz Intel 80386 processor whose 8 MB of memory are divided into 4 MB for its operating system and application software, 2 MB for program execution, and 2 MB for user data storage.

The 9000 has a clamshell design. When closed, the 9000 operates as a normal cellular telephone with a 3×10 character LCD screen and a 2×6 character soft key area. When opened, it reveals an LCD screen on one side and a keyboard on the other. The 4.5×1.4 in. eight-level gray scale display has 640×200 pixels of resolution. The system includes a loudspeaker, a microphone, and a 55-character QWERTY keyboard, not supported by a mouse.

The 9000 supports full PDA functionality. It can run applications ranging from address book, calendar, and to-do list to calculator and world clock. It even supports Internet access via a Web browser and telnet. The 9000 weighs 14 ounces with external dimensions of $6.8 \times 2.5 \times 1.5$ in.

The design of the Nokia 9000 has several limitations. First, the keyboard is very small, making it difficult for users to type with all ten fingers. Second, users must navigate the screen by means of buttons, which are inferior to a pointing device. Third, the screen is small, limiting the amount of information it can display at one time. Last, it is rather bulky, making it uncomfortable to hold when used as a cellular telephone.

The AT&T PocketNet telephone is built for AT&T by PCSI, Inc. A regular Advanced Mobile Phone Service (AMPS) cellular telephone, it incorporates a CDPD modem. Embedded in the telephone is a transmission control protocol/Internet protocol (TCP/IP) stack that enables third-party applications to access the Internet. It has a 5-line \times 20-character display, weighs 11.6 oz., and measures $6 \times 2.3 \times 1.3$ in. The

Table I. PDAs currently on the market.

Manufacturer	Model	CPU	Operating system	Screen size (in pixels)	Other comments
3Com	PalmPilot*	Motorola 68328 DragonBall*	Palm operating system	160 x 160	Graffiti* handwriting recognition
Apple	MessagePad* 130	ARM 610	Newton* Intelligence 2.0	320 x 240	Paragraph and Apple handwriting recognizers
Casio	Z-7000	Custom (x86 compatible)	GEOS*/PenRight!*/HOPE!*	320 x 256	Handwriting recognition supported
Hewlett-Packard	OmniGo* 100	80186 compatible	GEOS 2.1	240 x 240	Keyboard and Graffiti handwriting recognition
Motorola	Envoy 150	Motorola Dragon 68349	Magic Cap* 1.5	480 x 320	Two-way wireless communicator
Philips	Velo* 1	MIPS* R3900	Windows* CE	480 x 240	Pen/touch for navigation, QWERTY keyboard
Hewlett-Packard	Palmtop PC 300/320LX	Hitachi SH-3	Windows CE	640 x 240	Pen/touch for navigation, QWERTY keyboard
Sharp	Zaurus* ZR-5800	16-bit proprietary	Synergy	320 x 240	Pen/touch for navigation, QWERTY keyboard

PocketNet telephone comes bundled with the UP.Link* Web browser from Unwired Planet.¹⁰ Used on the telephone to view documents formatted in the Handheld Device Markup Language (HDML), this browser supports some interactivity. A message received on the telephone can prompt the user to respond to the sender with one of four replies.

The limitations of this device include its screen and the lack of an efficient input mechanism. A larger bit-mapped screen would be useful for displaying data, and a keyboard or other input device would support richer interactivity.

Personal Digital Assistants

There are many PDAs on the market, some with communications facilities. **Table I** lists the major PDAs.¹¹ By far the most popular PDA is the 3Com PalmPilot, a quarter million units of which were shipped in the first half of 1997.¹² Inexpensive, small, and light, it costs less than US\$300 and weighs 0.36 lb. in a 4.7 × 3.2 × 0.7 in. package. This device, which uses handwriting recognition for user input, enables users to write in unistrokes (single strokes), which are then processed by the Graffiti* recognition engine. The shapes of the unistrokes were chosen to be as differentiable as possible, while still retaining

some similarity to the character represented.

The PalmPilot requires users to write one unistroke character at a time in a predefined area of the screen. The handwriting recognizer inserts recognized characters into the current entry box, making the recognizer's task easier at the expense of a less intuitive interface. However, the sales figures indicate that this is acceptable to users. PalmPilot applications include calendar, appointment book, notebook, calculator, and all the other standard PDA applications. Users can back up their files onto a PC via a connector on one edge of the device.

The Newton MessagePad also uses handwriting recognition, but it does not constrain the user in the way the PalmPilot does. Instead, all entry boxes are handwriting enabled. The user writes directly into the entry box using normal cursive or printed letters. The recognizer can handle mixed-mode handwriting (that is, mixed cursive and printed letters), although the error rate is higher than in the PalmPilot. Expensive and bulky compared to the PalmPilot, the MessagePad costs US\$950 and weighs 1.3 lbs. in a 8.25 × 4.5 × 1 in. package.

Several PDAs use standard QWERTY keyboards instead of handwriting recognition. Sharp has been in this market for several years with the Zaurus series.

Table II. Channel bandwidths in digital wireless networks.

Standard name	Technology	Maximum data channel bandwidth	Comments
ARDIS	Slotted CSMA	19.2 kb/s	Packet data. User data rate is 8 kb/s. Service in U.S. provided by IBM and Motorola.
Mobitex	Slotted CSMA	8 kb/s	Packet data. Service in U.S. provided by RAM Mobile Data.
CDPD	Slotted DSMA/CD	19.2 kb/s	Packet data. AT&T has a large network in the U.S.
AMPS	Analog	14.4 kb/s	Analog cellular modem required.
IS-136	TDMA/FDMA	9.6 kb/s	Cellular telephone standard popular in the U.S. and Asia.
IS-95	CDMA	14.4 kb/s	Cellular telephone standard popular in the U.S. and Asia.
GSM	TDMA/FHMA	9.6 kb/s	Cellular telephone standard popular in Europe and Asia.
PACS	TDMA	32 kb/s	For indoor and microcell use.
DECT	FDMA/TDMA	32 kb/s	Cordless telephone standard.
CT2	FDMA	32 kb/s	Cordless telephone standard superseded by DECT.
PHS	TDMA	128 kb/s	Cordless telephone system for microcell/indoor use. Popular in Japan.
Metricom	FHSS	28.8 kb/s	Wireless data service provider serving the San Francisco Bay area and other metropolitan areas.

CD – Collision detection

CSMA – Carrier sense multiple access

DSMA – Digital sense multiple access

FDMA – Frequency division multiple access

FHMA – Frequency-hopped multiple access

FHSS – Frequency-hopped spread spectrum

TDMA – Time division multiple access

Zaurus machines, which support all the standard PDA applications, weigh about 0.75 lb. in packages that measure $6.5 \times 3.7 \times 1$ in.

Microsoft's Windows CE operating system for hand-held computers has inspired several new hand-held PCs (HPCs). Philips was first on the market with the Velo* 1, whose clamshell design includes a built-in fax/data modem and a QWERTY keyboard. Although it comes with a 480×240 pixel touch screen, it has no native support for handwriting recognition, the absence of which is common across all Windows CE machines. The 0.91 lb. Velo 1 measures $6.7 \times 3.65 \times 1$ in.

Although most PDAs support fax/data modems and direct PC connections, none offer built-in wireless communications capabilities. Conventional cellular modems may be used in place of external wireline modems. Motorola's Envoy and Marco were PDAs that included wireless communications, but they did not succeed in the marketplace. The Simon, produced by IBM and BellSouth, suffered a similar fate.

Wireless Data Service Providers

This section describes some of the services, standards, and protocols used by providers to transmit data over wireless networks. A wireless connection has limited bandwidth, and each digital cellular telephone standard defines a data channel and a voice channel that share the radio bandwidth.^{13,14} **Table II** shows data channel bandwidths for each digital cellular standard and for some other wireless networks.

Advanced Radio Data Information Systems (ARDIS) is a two-way packet data service operated by Motorola. A customer's mainframe computer is linked to an ARDIS host using either simple network architecture (SNA), X.25 dedicated circuits, or some other means. Remote users access the system from laptop terminals through wireless modems. The service is suitable for two-way transfer of data files containing less than 10 KB.

Mobitex is a nationwide, interconnected trunked radio network for voice and data developed by Ericsson and Swedish Telecom. RAM Mobile Data has operated data-only Mobitex service in the U.S. since

1991. It operates in much the same way as a cellular telephone system, except that handoffs are handled by the mobile terminal rather than the network.

CDPD uses a full 30 kHz AMPS channel on a shared basis to provide mobile packet data connectivity to existing data networks and other cellular systems without any additional bandwidth requirements. It directly overlays the existing analog cellular infrastructure and uses existing base station equipment, making it simple and inexpensive to install. AT&T operates a large CDPD system in the U. S., supporting broadcast, dispatch, e-mail, and field monitoring applications. Mobile users can connect to the Internet or the public switched telephone network (PSTN) via intermediate systems, which act as servers and routers.

Subscribers to the analog cellular telephone network, called AMPS, may transmit and receive data using cellular voiceband data modems, just as PCs use wireline voiceband data modems. Although these modems work well in practice, they are expensive.

IS-136, IS-95, and GSM are digital cellular telephone standards that provide data and voice on the same physical channel. GSM is widely deployed in Europe and Asia, whereas IS-54, an earlier version of IS-136, is deployed in the U. S. and Asia. IS-95, a standard widely used in Korea and now being deployed in the U.S. and Asia, uses circuit-switched data channels. A companion standard for packet-switched data channels—IS-657¹⁵—has been developed for IS-95.

The Personal Access Communications System (PACS) was originally developed and proposed by Bellcore in 1992. It supports voice, data, and video for indoor and microcell use. The Cordless Telephone 2 (CT2) system and the Digital European Cordless Telephone (DECT) system are European standards for cordless telephones. The main function of DECT is to provide local mobility inside a building.

Users of the Personal HandyPhone System (PHS)—a cellular system widely deployed in Japan that supports voice and data—must be close to a base station to access the network. PHS was designed for microcell and indoor use, so handoffs are supported only at walking speeds.

Metricom, an Internet service provider (ISP), has deployed a proprietary wireless data network covering

the entire San Francisco Bay area. Now expanding into other cities and campuses, this network does not carry voice traffic and has high latency.

Dedicated data services are accessed through a single service provider, whereas digital cellular data services are offered by more than one service provider. Most PCS service providers are deploying systems based on digital cellular standards, so it makes sense for designers of new mobile data terminals to target data services using public digital cellular and PCS networks rather than proprietary data networks.

Table II shows that digital cellular telephones can attain up to 14.4 kb/s on the data channel. For low bandwidths it is important to design applications and interfaces that will not overload the radio channel. Future cellular networks will possess greater data transmission capacity, but they will not have as much bandwidth as the wired network.

A mobile terminal with a small screen can display small bitmaps and a few characters at a time, requiring only a low radio bandwidth to transmit data to be displayed. Gauging how low a bandwidth is acceptable is a function of the screen size, the currently active applications, the operating system, and the system's windows. We believe that with proper system design, low wireless bandwidth will not be a serious impediment to the effectiveness of most applications.

On the other hand, the quality of service on the data channel is of more serious concern. Wireless channels in urban areas exhibit deep fades and are therefore prone to burst errors and loss of connectivity. We must compensate for these impairments with good protocol design on the wireless data link.

Software Choices

To support data services, a handset needs an operating system, applications, and wireless data communications capability, and today's marketplace offers several choices. Lucent Technologies' InfernoTM system¹⁶ provides a complete solution, including an operating system, a communications protocol, and a programming language. Unwired Planet¹⁰ offers HDML, the hand-held device transport protocol (HDTP), and the UP.Link platform and browser. The Java* programming language,¹⁷ supported by Java

chips, may be used in the mobile terminal. GEOS,⁹ by GeoWorks, is a popular operating system for small devices. The interactive communicating applications protocol (ICAP*) of Wink* Communications¹⁸ supports interactive graphical applications.

Inferno Operating System

The Inferno operating system runs on many hardware platforms, including the ARM 7500, Intel 386 and higher, MIPS R3000 and R4000 series, and SPARC processors. Applications are written in Limbo™, a C-like programming language whose compiled byte code is interpreted by Inferno's DIS® virtual machine, making the applications highly portable. DIS byte code is similar to machine code, so the interpreter is fast.

Inferno applications may implement a graphical user interface (GUI) using the Tk widget set,¹⁹ which is built into the kernel. Applications may be multithreaded, with threads that communicate with each other locally or over the network using the Styx™ communications protocol. Communications channels are authenticated and encrypted, enabling users to access remote services securely. Inferno uses file semantics for access to local and remote resources. All resources are mounted in a single hierarchical name space structured like a directory tree. Leaf nodes in the tree can be local or remote; the access semantics are the same. This is a powerful way to abstract the network away from the application programmer and the user. Inferno software is also compact, requiring a total of less than 1 MB for the operating system and applications.

UP.Link Platform

The UP.Link platform includes HDML for developing applications and services, HDTP for serving HDML documents, the UP.Browser* Web browser for viewing HDML documents, and UP.Mail* for e-mail. The UP.Link server connecting mobile devices to the Web is located on the wired network.

HDML is not a true markup language, but rather a set of commands or statements that specify how a hand-held device interacts with a user. An HDML document is a deck of cards with four types:

- *Entry cards* display a message and prompt the user to enter a string of text,

- *Choice cards* display a list from which the user can choose one item,
- *Display cards* simply display information, and
- *No-display cards* execute an action but do not appear on the telephone display.

HDML documents are stored in compressed form for transmission over the wireless channel and are decompressed on the terminal before they are executed.

Java Programming Language

Very popular, C-like, and multithreaded, the Java object-oriented programming language is designed for programming on the Internet. Java programs (including a document viewer) can be embedded into documents to enable anyone to view them. Software can be taken from any remote site and run locally on any platform. Like Inferno programs, Java programs are compiled into byte code interpreted by the Java virtual machine (JVM) to support portability. Because Java is on the average about 20 times slower than C,²⁰ specialized processor chips—called Java chips—may be used for executing Java applications.

GEOS Operating System

The GEOS operating system, which runs only on Intel 80x86 processors, was designed for devices with a small memory and slow general-purpose processors. The GEOS system is real time, object oriented, and multithreaded, with an operating system kernel that directly supports messaging. Application development is isolated from specific user-interface implementation, so systems incorporating GEOS can have different specific user interfaces and still maintain compatibility with any GEOS application. Application developers do not place user interface elements on the screen. Instead, developers provide information to the user interface library about how they would like to see the elements arranged. GEOS code is similar to C-language code, but it has object-oriented extensions.

Very popular, GEOS is currently used on the Nokia 9000, Hewlett-Packard OmniGo,* and the Casio Z-7000. It requires 3 MB of read-only memory (ROM) and 1 MB of random access memory (RAM) for the operating system and 17 to 20 applications. GEOS supports several handwriting recognizers, and TCP/IP sockets enable interprocess communication. The operating system supports speech recognition in the user

interface, although no products currently on the market use this facility.

Wink Engine

ICAP, a compact protocol for defining and delivering interactive graphical applications, requires low bandwidth and a small amount of memory. The Wink engine that runs in the handset requires 128 KB of ROM and 32 KB of RAM and runs on top of the handheld device's operating system and transport.

The Wink engine interprets ICAP applications, maps objects and user interface elements to the target system, and executes scripts or data stream commands. ICAP records may contain application object definitions, resource definitions, scripts, updates to objects and resources already received, and control commands.

Windows CE

Windows CE is a compact operating system based on Windows 95. Several palmtop computers, or HPCs, that run this operating system are available from Philips, Hewlett Packard, and others. All of them feature at least 8 MB of RAM and 8 MB of ROM, which is much more than what PDAs running other operating systems have. Standard applications include small versions of Windows applications, such as Microsoft's Internet Explorer, Word, and Excel.

All these products feature a keyboard and an LCD screen in a clamshell casing. Third-party software supports pen input and handwriting recognition, but none of the current products use a pen as more than a mouse replacement.

User Interface Issues

An effective, user-friendly interface is crucial for user acceptance of any consumer product. A product with a non-user-friendly interface will not succeed, regardless of its underlying technical excellence. For a portable device, the first tier of usability factors is application suite, size, weight, and battery life. These factors determine whether the user will even look at the product. The next tier is the interfaces to the applications themselves, including the look and feel of the device, the clarity of the screens, the ease with which data is entered into the applications, and their responsiveness.

The user interface has two aspects: user input and

application output. User input drives applications and the operating environment. Application output returns data to the user based on the user's commands. This section addresses issues relating to the quality of the user interface.

Input Devices

There are several possible modes for user input. The keyboard is good for mass text entry (such as programming, word processing, and data entry), but it is heavy and bulky, requiring training for fast text entry. A small, light keyboard is feasible (as seen on some PDAs and the Nokia 9000 GSM smart phone), but its small size makes it difficult to use. If a graphical user interface is employed, a pointing device is needed to supplement the keyboard.

A touch screen or pen digitizer is inefficient for mass text entry, but it does not require any user training, as some handwriting recognizers do (for example, training for a modified alphabet). A pen digitizer also integrates the pointing and text entry functions into a single input device that adds little to the weight or bulk of the device (a resistive touch screen adds 1.4 mm to the thickness of the LCD panel). The pen works well with applications that have a point-and-tap interface. The handwriting recognizer for text entry is the weakest part of a pen interface because it is error prone. A good interface design, however, can minimize the effect of recognition errors on the user.

Speech recognition has advanced enough to make it possible to perform many tasks using spoken input. Commercial isolated-word recognizers exist (for example, DragonDictate,* Kurzweil Voice,* and IBM VoiceType*), and research labs such as those in Lucent, Carnegie Mellon University, IBM, and Bolt, Beranek and Newman are conducting research in continuous speech recognition. Speech recognition offers the advantage of hands-free operation, but it is sensitive to interfering noise, room acoustics, and microphone placement. It usually requires a head-mounted boom microphone, vocabulary specified in advance, and a recognizer trained for that vocabulary. The user interface must be able to deal with recognition errors, and the applications must be designed to reduce recognition errors. Speech recognition also requires a powerful processor, making it impractical for portable devices.

Visual Display

Advances in LCD screen technology have made it inexpensive to put a small screen on a portable terminal. It costs about US\$20 to US\$30 to add a 640×200 pixel LCD screen to a device. A color screen, while very pleasing to the eye, requires a backlight for visual clarity, which adds weight and at least 2W of power consumption. If no backlight is used, power consumption remains low (about 15 mW). Keeping the screen small makes it lighter and less bulky, which is ideal for portability, but a larger screen can display more information, thereby improving the user interface. A bit-mapped screen supports a GUI, which is greatly superior to a character-based interface.

Although most computer screens operate in landscape mode, where the width is greater than the height, some PDA screens operate in portrait mode. Similar to the orientation of a printed page, portrait mode is considered superior for PDAs. In contrast to laptops, most PDAs use monochrome, non-backlit screens to save on power consumption and bulk. Screen resolution (dot-pitch), contrast, and brightness are key parameters that affect the perceived quality of the screen. A high refresh rate makes it possible to display animation and video on the screen.

Audio Feedback

Modern desktop computers use beeps and tones for audio feedback. They can also play audio messages such as voice mail and can operate as speakerphones. A loudspeaker on the handset is a mixed benefit, because it generates noise and needs volume control to prevent it from becoming a nuisance.

Alternatively, a headset (possibly with a boom microphone) can deliver audio feedback. The disadvantage of a headset is that the user may feel inconvenienced by having to carry it around, put it on, and plug it in.

System Architecture

This section describes the design objectives and high-level architecture of our handset. The next three sections examine its implementation in greater detail.

Our objective is to augment the digital cellular telephone network to provide access to wireless data any time, anywhere. We take advantage of the popu-

larity of the cellular telephone by providing applications that support communication. Our solution requires enhancing a current base station to connect directly to a data network containing an Internet connection and to recognize wireless data packets not destined for the wired telephone network. The base station needs additional processing power for intelligent processing of data packets moving to and from the handset. For example, Web pages may be transcoded to transmit only relevant parts of a page, thereby conserving radio bandwidth.

Applications

Compelling applications are the key to the success of the system. We consider two tiers of applications: primary and secondary. The primary applications take advantage of the popularity of cellular telephones by supplementing the telephone function. Secondary applications support information retrieval and personal information management (PIM) functions. In keeping with our philosophy that it is better to build a few excellent applications than many mediocre ones, we designed the system around the primary applications.

Primary applications. We selected three primary applications for our system: directory, e-mail, and Web browsing. These applications directly complement the telephone part of the handset to create a general communications device.

The directory enables a user to obtain contact information such as an e-mail address or telephone number from the white pages, yellow pages, a corporate internal telephone directory, or a personal directory. A user may query the directory via a name search or other mechanism and then automatically dial the number obtained. The directory returns the address of the search subject, which can be shown on a map with directions from the user's current location. Because the directory is a document, it can also be accessed through the Web browser.

The Web browser, which also supports e-mail, enables users to specify an e-mail address on a form or select it from an electronic address book. In addition to the primary applications, a serial port and the wireless channel can back up and synchronize local data on a desktop or network computer. A separate application can monitor battery life and radio channel status.

Secondary applications. The hardware and software also support secondary applications, but we do not make a special attempt to optimize for them. They are less important than the primary applications, and their user interface requirements often conflict with the requirements of the primary applications. Many of these applications are interactive, requiring text input from the user.

A billing support application allows the user to determine the cost of making a telephone or fax call to any domestic or international number at a specified time. It advises the user of the cheapest time to make the call and schedules a fax transmission for that time. Of course, users may receive faxes at any time. They can also browse previous billing statements, obtain customer service, and make payments electronically.

Most PDAs implement a number of PIM applications. These include appointment diary, expense diary, general calendar with scheduling, automatic appointment reminders, world clock, notepad, calculator, car mileage tracker, and others. Other practical uses include simple games and a voice recorder for notes and messages.

In addition, it is possible to build software agents that take advantage of the wireless connection to perform remote tasks on the user's behalf. Such services include banking, shopping, and appointment scheduling.

Access to the Wired Data Network

A wired PSTN call can take several seconds to set up, and a cellular call takes even longer. As a result, an application that downloads programs or data from the network has a long initial delay. After that, it has a continuous circuit-switched connection, even though the radio channel is idle most of the time. Because the channel is priced by connect time rather than by bandwidth used, obtaining services over the wireless channel is expensive. Ideally, the data channel should be a packet-switched multiple access channel rather than a circuit-switched channel. Even if a connection-oriented topology is used, it should support quick call setup similar to that of integrated services digital network (ISDN) connections.

Fortunately, new digital cellular packet data stan-

dards are evolving. IS-657¹⁵ is a packet data standard for CDMA. These standards will help speed the deployment of packet data capabilities in new digital cellular and PCS networks.

In our system, a wired data network directly connects to the wireless base station (see **Figure 1**), with several base stations sharing each data network. Data traffic is routed to and from the wireless channel by the data network rather than the PSTN. On the uplink, the wireless base station detects data packets and passes them directly to the packet-based wired data network. The call setup time is therefore very fast.

Local Versus Remote Data

Data may be synchronized and backed up on the network. Because data storage on the handset is limited, user data may also be stored on the network. This adds greater security and reliability, because data will not be compromised if the handset is lost or stolen.

Handset Hardware

This section describes the handset hardware, including its telephone operation, PDA operation, LCD screens, audio, and other hardware. It also presents an alternative handset design.

The handset features a clamshell casing which, when closed, looks and operates like a regular cellular telephone (see **Figure 2**). When the clamshell casing is open, it reveals two touch screens and works like a PDA. The two screens behave as one video graphics adapter (VGA)-sized screen divided into two contiguous halves. The seam between the two screens is oriented horizontally and is made as narrow as possible to decrease disruption of the visual flow of the logical screen from one physical screen to the other.

Telephone Operation

When the handset is closed, it looks and acts like a normal cellular telephone, with its numeric keypad on the front, the usual extra buttons (send, store, recall, answer), and a small LCD screen. The earpiece and microphone are also in the usual places, allowing users to make and receive calls as they would on a regular cellular telephone. The form factor is about the same as that of other cellular telephones.

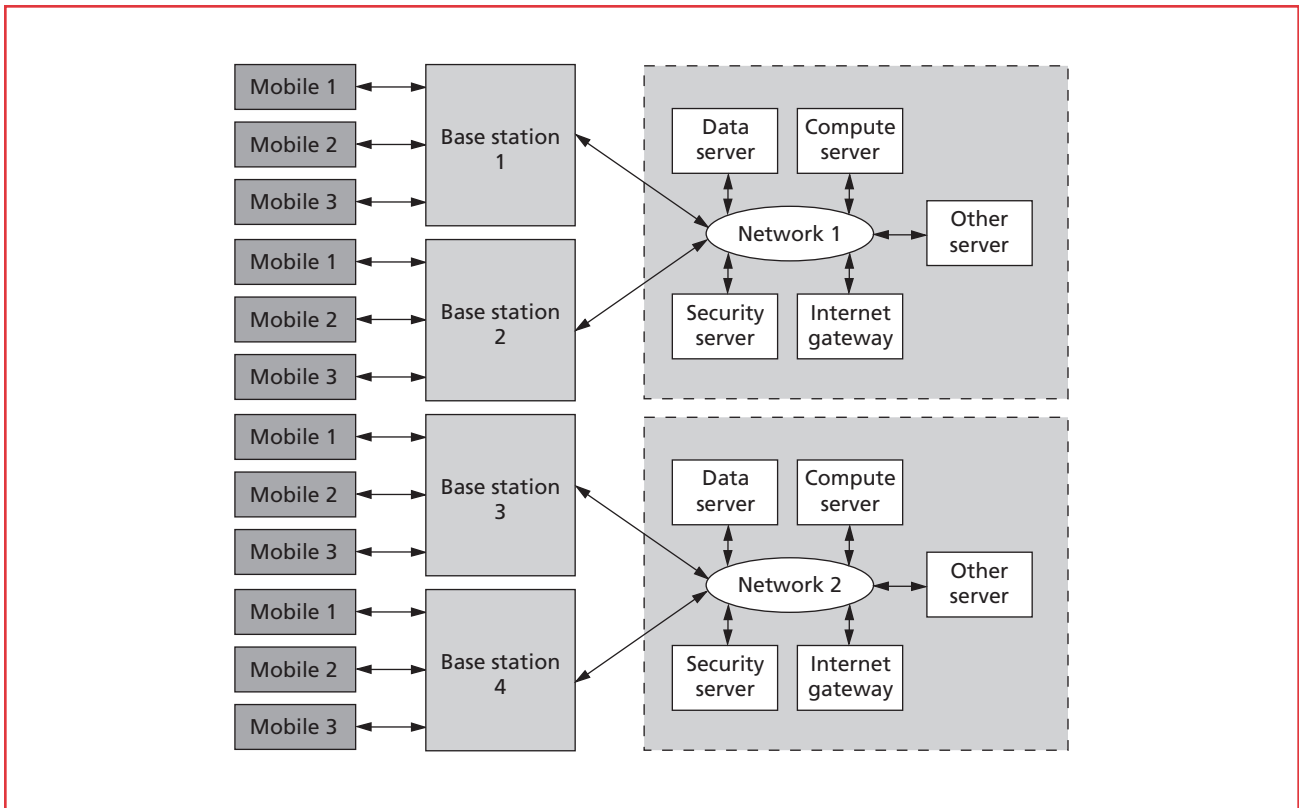


Figure 1.
Network architecture, excluding PSTN connections.

PDA Operation

When the handset is open, it operates as a PDA. The PDA incorporates a digitizer—a resistive touch-sensitive membrane laid over the screen—and a passive pen, used to press soft keys, manipulate menus, and write on the touch screen. When not in use, the passive pen resides in a slot on the handset’s casing. If the pen is lost, any blunt-pointed instrument—even a finger—works as a substitute, although it is not precise enough for handwriting recognition. In this case, a software keyboard is available for text entry.

LCD Screens

Although a high-contrast, backlit, high-resolution color LCD screen renders the best display, considerations of power, size, and cost dictate that we use a monochrome reflective screen. Hardware screen drivers divide the logical screen into two parts and display each half in the appropriate physical screen. Each physical screen is 640×240 pixels, with a dot pitch of 0.2 mm. This results in a screen size of 5×1.9 in., giv-

ing the entire handset a form factor of 6×2.5 in. when closed and 6×5 in. when open.

Each screen operates in landscape mode. Portrait mode is favored by many PDAs because it is more natural for user input. But our primary applications are for communications and viewing documents, so the user interface is optimized for displaying information. Screen width is a critical parameter in document display, and landscape mode eliminates the need for horizontal scrolling. In addition, orienting the seam horizontally between the screens minimizes visual disruption.

Audio

The outside of the handset contains a telephone microphone and an earpiece speaker. The telephone microphone is used for audio input in PDA mode by mounting it with apertures both outside the casing (for telephony) and inside (for PDA and speakerphone use). The audio quality on this microphone is sufficiently good for voice telephony, but it is not adequate




Figure 2.
Handset design.

for speech recognition. Mounted on the inside of the casing is a speaker that supports speakerphone applications and plays audio clips, including voice mail.

Other Hardware

Besides the usual cellular telephone components, the handset contains RAM, ROM, a serial port, and an Intel 80386 microprocessor for running the operating system and applications. The choice of microprocessor is heavily influenced by the choice of the operating system, because a part of the operating system must be available for this processor. Intel x86 processors have the greatest software support at present. The serial port may be used for data synchronization, backups, and optional peripherals.

An Alternative Design

The greatest disadvantage of the physical design of our PDA is that all its functionality is built into a single expensive device, forcing the user to carry around the extra hardware, even when it is not needed. An alternative is to build a plug-in module that the user can carry in a briefcase or in the car rather than in a

pocket. The PDA module could be attached to a regular cellular telephone when data services are required, and the handset could be sold separately from the PDA module. The user would be able to choose from several plug-in alternatives with various screen sizes.

This solution, however, would require the user to carry around not only the PDA module, but also the telephone, making it similar to other PDAs, which use a cellular modem for wireless communications. In addition, it is inconvenient to switch between the telephone function and the PDA function, because the two pieces must be connected and disconnected during the switch. **Figure 3** shows a sketch of this design.

Handset User Interface

The handset supports a GUI driven by pen/touch input. Applications are designed to use a point-and-click interface as much as possible, making most functions accessible by means of buttons and pull-down menus. For text input, a writing area called the handwriting capture widget (HCW) is maintained at the lower side of the screen. Having all the handwriting




Figure 3.
An alternative handset design, with a PDA module that attaches to a phone.

captured in one place on the screen provides maximum screen visibility, because the user's writing hand does not obstruct the rest of the screen. Another input option is a software keyboard, drawn on the lower screen. The user taps keys with his or her fingers or the pen and can move the cursor either with the arrow keys on the soft keyboard or by pointing with the pen.

The applications were written to recognize the small screen size and to adapt the interface to the screen size, precluding the use of fonts that are too large or excessive window ornamentation. They followed normal GUI design principles and ensured that only one application could be visible at any one time. This prevents user confusion and eliminates the need for placing and resizing applications on the screen. For applications that interact, text and graphics buffers enable convenient transfer of data.

Handwriting Capture Widget (HCW)

The bottom of **Figure 4** shows the HCW. It has a baseline on which the user writes and tick marks to

separate the handwritten (printed, not cursive) characters. The arrow buttons enable the user to scroll for long input words. The HCW captures characters and initially displays them as electronic ink. As soon as a character is completed, it is sent to the handwriting recognizer, which replaces the electronic ink with recognized text.

All editing is performed in the HCW. If an error occurs in recognition, the user taps on the erroneous character to bring up a corrections menu, which contains the five most probable alternatives returned by the recognizer. The user either chooses the correct character, rewrites the character, or brings up a soft keyboard.

When the user taps the "Done" button, recognized text is inserted at the cursor position in the current text entry widget (TEW), described in the next section. The button labeled "Lower" indicates that only lowercase characters are expected. When the user taps this button, the input mode changes to uppercase or digits. Punctuation symbols are always active in the vocabu-

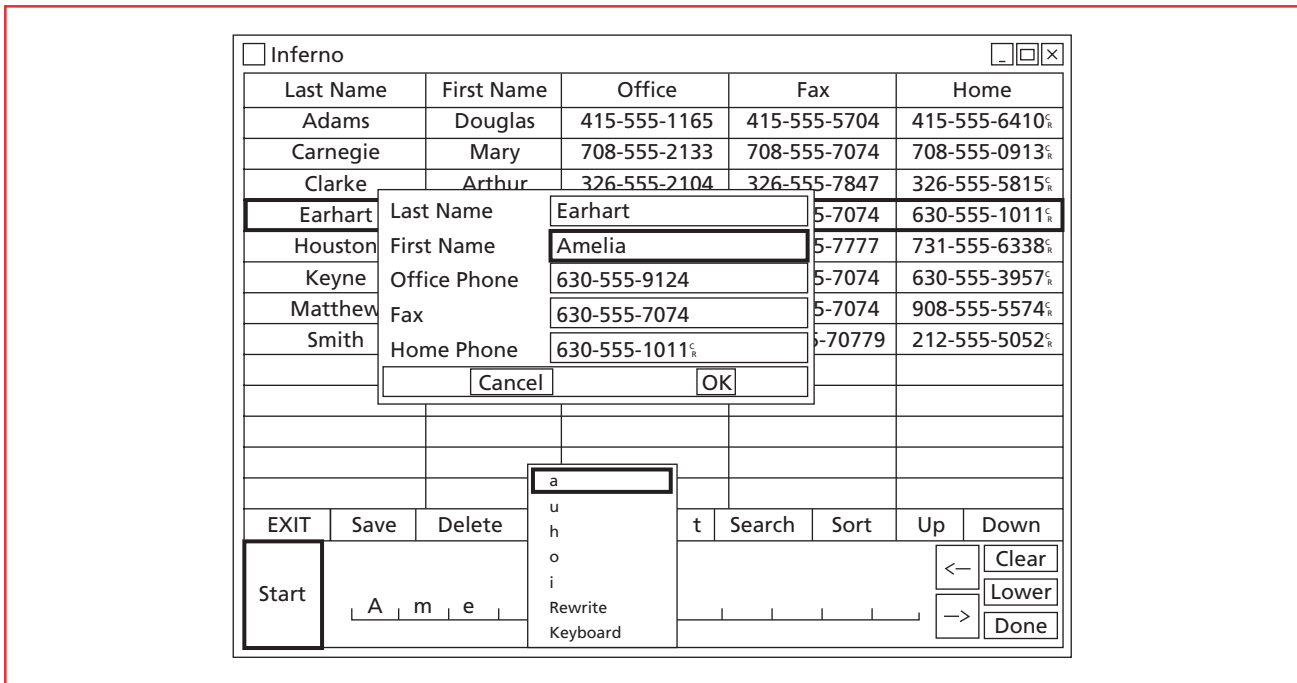


Figure 4.
A user interface with the directory application active and the correction menu visible in the HCW.

lary. The “Clear” button clears the selected characters or all input if no character is selected.

The user can switch between applications by pressing the “Start” button at the lower left corner of the screen. The HCW commander (HCWC) is a special application that enables the user to customize the HCW by modifying its size and the spacing of the tick marks, and by selecting the best one of several handwriting recognizers.

Text Entry Widget (TEW)

TEWs, enhanced Tk entry widgets, are used for text input by all applications. Each data entry box on the form is a TEW that looks like an ordinary text entry widget on the screen. Only one TEW at a time is considered the current one, however, and it is raised over the surrounding window. The application developer can specify a predecessor and a successor for each TEW, thereby organizing the TEWs into an ordered list. This is useful when several entries on a form must be filled out sequentially.

When the user taps the “Done” button on the HCW, recognized text is sent to the current TEW, which gives up currency to its successor. The user may preempt the ordering of the list by tapping on a TEW;

it then becomes the current TEW and the list order continues from there. Inferno file semantics enable communication between the TEWs and the HCW.

Biometric Authentication

When the terminal is turned on, the user is prompted to enter his or her user name and to sign in. A signature verification application authenticates the login locally or over the network, after which the HCW comes up and applications may be run (see **Figure 5**).

Signatures are particularly useful for identification, because each person’s signature is highly distinctive, especially if it is captured on line. The verifier is based on comparing global and local features of the test and template signatures. The global features capture the overall spatial and temporal characteristics of signatures. The feature statistics of a training set of six signatures are used to build a model or template for validating further test signatures.²¹

We extract local information about a signature using stroke-direction code (SDC). More specifically, SDC tries to model hand movements that produce a signature. It treats them as a time-ordered concatenation of a fixed number of strokes and derives informa-

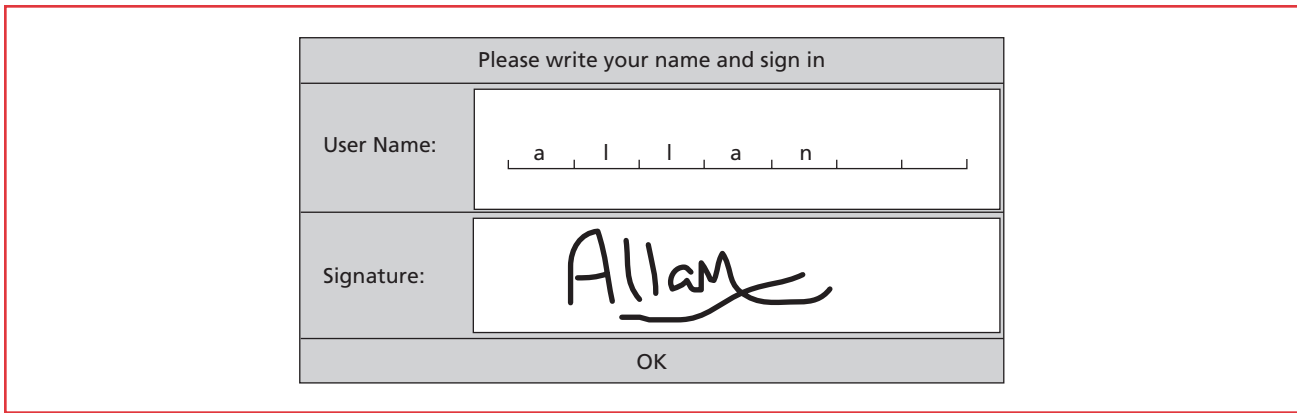


Figure 5.
Login window.

tion about the spatial orientation of these strokes. The SDC obtained from a test signature is compared to the model SDC using elastic matching techniques similar to the ones used in speech techniques.²² This algorithm has a 3% equal error rate and needs about 150 bytes to store a signature model.

Handwriting Recognizer

The user interface requires printed letters rather than cursive or mixed-mode recognition. It does not use speech recognition. Speech recognition introduces a very large overhead in processing complexity, which outweighs its benefits. The current state of the art in speech recognition is not advanced enough for a good general-purpose user interface. Our design, however, does not preclude future use of speech recognition. We chose to support only printed letters because the input text consists mainly of words that are not in the dictionary. The set of input words includes universal resource locators (URLs) for the Web, e-mail addresses, telephone numbers, and personal and place names, making a dictionary less than useful. Because cursive recognizers require a predefined vocabulary for acceptable recognition performance, a cursive recognizer would not be effective. In addition, most of the input consists of single words, eliminating the usefulness of grammar for recognition. Therefore, we are using a printed-letter recognizer, developed at Bell Labs and customized for this system.

The recognizer, based on hidden Markov models (HMMs), supports writer-independent recognition of hand-printed characters of unconstrained styles.²³ An

HMM describes a doubly stochastic process, which generates a sequence of states hidden from observation and an observable process dependent on the underlying state sequence. HMMs have been proven successful in modeling speech and on-line printed lettering. In our recognizer, each character is represented by one or more classes, depending on the number of distinct styles observed for the character. Each class is modeled by a left-to-right HMM with a variable number of states and discrete state-dependent observation probabilities. For any input character, an N -best decoding algorithm (a variation of the Viterbi algorithm) is applied to find the top five closest matches in the alphabet. The HMMs are pretrained on a large number of character samples from many different writers drawn from the UNIPEN database.^{24,25}

Handset Software Infrastructure

The Inferno system, used in our prototype, offers many features that are ideal for our application. The operating system, virtual machine, and applications are small and fast, enabling the handset to use a low-speed processor and a small memory. This is important for size and power reasons. A larger operating system and applications would require more memory chips and would increase the size of the circuit board and the power consumption needed. A slower operating system would require a faster processor, which would also consume more power. Higher power consumption translates to a larger battery or shorter operating time between battery changes, both of which are undesirable.

Inferno has built-in support for authentication and security. Data transfer between machines can be encrypted at one of several levels, and each connection can be authenticated by a signer machine using a password.

Within the handset itself, processor threads communicate with each other using an Inferno channel, which facilitates easy transfer of typed data from one thread to another. Inferno is multithreaded, making it especially useful for the HCW, which must communicate with the TEWs and recognizers. Inferno also supports regular TCP/IP sockets for communications with non-Inferno machines.

The Tk widget set, which is built into the Inferno kernel, supports GUIs. It contains graphical entities such as buttons, menus, entry windows, and other widgets. Programs written in Limbo, the programming language of the Inferno system, can create and manipulate Tk widgets directly.

All the applications and the user interface are highly portable because they are written in Limbo and then compiled into byte code for DIS, the Inferno virtual machine. Inferno provides emulators that run on PCs and UNIX* machines, enabling us to do all our development on desktop machines. Once the handset hardware is ready, it is simple to download compiled Limbo programs into the handset. Emulation also facilitates experimentation with the look and feel of the user interface and the screen size on development platforms.

Future Work

We are currently prototyping this system, after which we will run several experiments to evaluate it. We are particularly interested in obtaining feedback on the user interface and will conduct experiments with various screen sizes, screen resolutions, color screens, and HCW placement and size. We would like to characterize the data throughput over the wireless channel under different load conditions. Varying the handset form factor will help us determine how to improve it. For example, users may prefer a longer, narrower handset to the one we have currently. Statistics of data bandwidth requirements for each task will help us tune this system and design future ones.

Acknowledgments

Several people at Lucent Technologies contributed to this effort by sharing their knowledge and experience. Sanjay Kasturia was an excellent sounding board for our ideas and provided general guidance. Henry Baird's enthusiastic support helped our collaboration get off the ground quickly. Gordon Wilfong helped clarify the role of the recognizers in this system. Discussions with Thomas Woo, Krishan Sabnani, Razel Haimi-Cohen, Chris Burges, Craig Nohl, Eric Beck, Joel Angiolillo, Nick Katis, Dan McGuire, Charlie Roxlo, Mark Paczkowski, and Mark Cravatts helped us understand the larger issues and the status of related projects within Lucent. Colin Warwick, Ping-Wen Ong, and Tod Sizer gave us invaluable help in understanding the LCD market and technology. We are grateful to all these contributors and to any others we may have inadvertently omitted.

***Trademarks**

DragonBall is a trademark of Motorola.

DragonDictate is a registered trademark of Dragon Systems, Inc.

GEOS is a trademark of GeoWorks Corporation.

Graffiti is a registered trademark and PalmPilot is a trademark of Palm Computing, Inc.

HOPE! is a registered trademark of Mobile Computing Systems Corporation.

ICAP and Wink are registered trademarks of Wink Communications, Inc.

Java is a registered trademark of Sun Microsystems.

Kurzweil Voice is a trademark of Kurzweil Applied Intelligence, Inc.

Magic Cap is a registered trademark of General Magic, Inc.

MessagePad and Newton are registered trademarks of Apple Computer, Inc.

Minstrel is a trademark of Novatel Wireless Technologies Ltd.

MIPS is a registered trademark of MIPS Computer Systems, Inc.

Nokia is a registered trademark of Nokia Corporation.

OmniGo is a registered trademark of Hewlett-Packard Company.

PenRight! is a registered trademark of Grid Systems Corporation.

PocketNet is a trademark of AT&T Wireless Services, Inc.

UNIX is a registered trademark of The Open Group.

UP.Browser, UP.Link, and UP.Mail are trademarks of Unwired Planet, Inc.

Velo is a trademark of Philips Electronics North America Corporation.

VoiceType is a registered trademark of International Business Machines Corporation.

Windows is a registered trademark of Microsoft Corporation.

Yahoo! is a registered trademark of Yahoo! Inc.

Zaurus is a trademark of Sharp Kabushiki Kaisha.

References

1. D. Cox, "Wireless Personal Communications: A Perspective," *The Mobile Communications Handbook*, ed. J. Gibson, CRC Press, Boca Raton, Fla., 1996.
2. R. Schneiderman, *Wireless Personal Communications*, IEEE Press, Piscataway, N.J., 1994.
3. IS-136, *TDMA Cellular/PCS-Radio Interface-Mobile Station-Base Station Compatibility-Digital Control Channel Addendum No. 1*, Electronic Industries Association, Arlington, Va., 1997.
4. IS-95, *Mobile Station-Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular System*, Electronic Industries Association, Arlington, Va., 1995.
5. GSM, *PCS Access Services Interface Specification in Support of GSM-Based (DCS 1900) PCS Routing Service, Database Services, and Message Transport Service*, Bell Communications Research, Red Bank, N.J., 1995.
6. C. Musciano and B. Kennedy, *HTML: The Definitive Guide*, O'Reilly and Associates, Sebastopol, Calif., 1996.
7. Nokia Web site, <http://www.nokia.com/com9000/n9000.html>
8. AT&T PocketNet telephone, <http://www.pcsi.com/pal.html>
9. GEOWORKS' home page, <http://www.geoworks.com>
10. Unwired Planet's home page, <http://www.uplanet.com>
11. "PDA Buyer's Guide," *Pen Computing Magazine*, Vol. 3, No. 11, July/Aug. 1996, pp. 84-85.
12. T. Schmidt, "Tim's Personal Devices Journal," *Pen Computing Magazine*, Vol. 4, No. 16, June 1997, pp. 46-47.
13. J. D. Gibson, *The Mobile Communications Handbook*, CRC Press, Boca Raton, Fla., 1996.
14. T. Rappaport, *Wireless Communications: Principles and Practice*, IEEE Press, Piscataway, N.J., 1996.
15. IS-657, *Packet Data Service Option Standard for Wideband Spread Spectrum Systems*, Electronic Industries Association, Arlington, Va., 1996.
16. Inferno operating system, <http://www.lucent.com/inferno>
17. Java software, <http://www.javasoft.com>
18. Wink's home page, <http://www.wink.com>
19. J. Ousterhout, *Tcl and the Tk Toolkit*, Addison-Wesley, Reading, Mass., 1994.
20. D. Flanagan, *Java in a Nutshell*, O'Reilly and Associates, Sebastopol, Calif., 1996.
21. R. S. Kashi, W. Turin, and W. Nelson, "On-line Signature Verification Using Stroke Direction Coding," *Optical Engineering J.*, Vol. 35, No. 9, 1996, pp. 2526-2533.
22. H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Trans. Acoustics, Speech, Signal Processing*, Vol. ASSP-26, Feb. 1978, pp. 43-49.
23. J. Hu, M. K. Brown, and W. Turin, "HMM Based On-line Handwriting Recognition," *IEEE PAMI*, Vol. 18, No. 10, Oct. 1996, pp. 1039-1045.
24. I. Guyon, L. Schomaker, R. Plamondon, M. Liberman, and S. Janet, "UNIPEN Project of On-line Data Exchange and Recognizer Benchmarks," *Proc. of the 12th Int'l. Conf. on Pattern Recognition*, Jerusalem, Israel, Oct. 1994, pp. 29-33.
25. The UNIPEN project, <http://unipen.nici.kun.nl>

(Manuscript approved May 1998)

SHANKAR NARAYANASWAMY received B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California at Berkeley. As a member of technical staff in the Wireless Technologies Research Department at Bell Labs in



Holmdel, New Jersey, Dr. Narayanaswamy is working on wireless data systems and their applications. In addition, he is interested in the fields of wireless data protocols, mobile communications, and handwriting- and speech-based user interfaces.

JIANMING HU, a member of technical staff in the Language Modeling Research Department at Bell Labs in Murray Hill, New Jersey, received her M.S. and Ph.D. degrees in computer science from the State University of New York at Stony Brook. Dr. Hu is currently working on pattern recognition, document image processing, and information retrieval.





RAMANUJAN S. KASHI holds a B.S. in electrical engineering from the University of Mysore in India and M.S. and Ph.D. degrees in biomedical engineering from Rutgers University in Piscataway, New Jersey. A member of technical staff in the Language Modeling Research Department at Bell Labs in Murray Hill, New Jersey, Dr. Kashi conducts research in shape description and perceptual issues of image analysis and segmentation. ♦