Ericsson

REVIEW

4/1999

THE TELECOMMUNICATIONS TECHNOLOGY JOURNAL

1924 75 1999 75 YEARS IN PUBLICATION

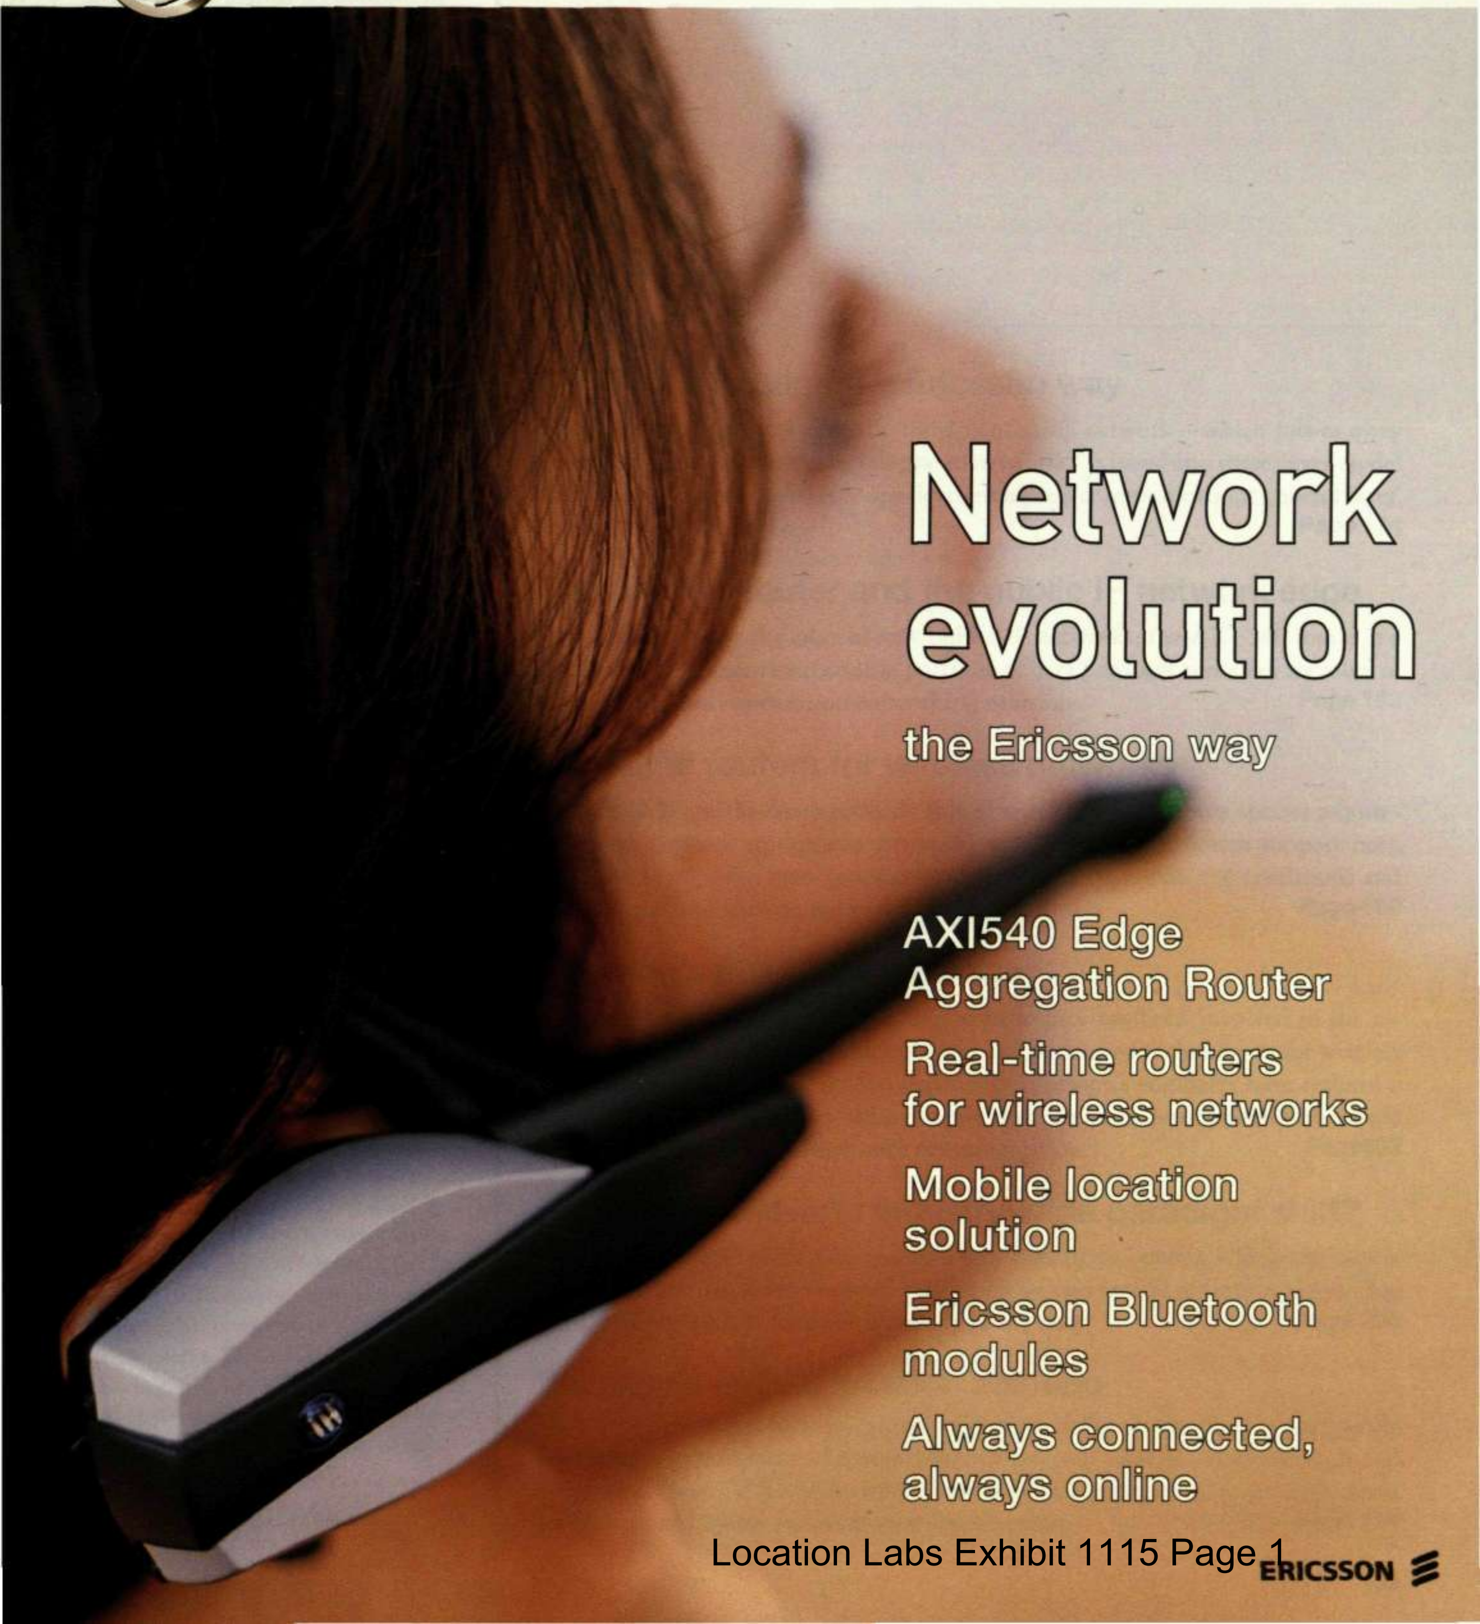

Network evolution

the Ericsson way

AXI540 Edge Aggregation Router

Real-time routers for wireless networks

Mobile location solution

Ericsson Bluetooth modules

Always connected, always online

ERICSSON

Ericsson
REVIEW

THE TELECOMMUNICATIONS TECHNOLOGY JOURNAL

The purpose of Ericsson Review is to report on the research, development and production achievements made in telecommunications technology at Ericsson. The journal is published in English and Spanish and distributed quarterly to readers in more than 130 countries.

Address:
Telefonaktiebolaget LM Ericsson
SE-126 25 Stockholm, Sweden

Phone: +46 8 719 00 00

Fax: +46 8 681 27 10

Internet:
http://www.ericsson.com/review

New subscriptions:

http://www.ericsson.com/review

Address changes and distribution:

Fax: +46 8 681 27 10

Publisher: Jan Uddenfeldt

Editor: Eric Peterson

Production manager: Eva Karlstein
eva.karlstein@lme.ericsson.se

Layout: Paues Media, Stockholm

Illustrations: Claes-Göran Andersson, Hara

Printer: Ljungföretagen, Örebro

ISSN: 0014-0171

Volume: 76, 1999

Cover: Ericsson has begun manufacturing Bluetooth chips and sees a giant market ahead. "There is enormous interest in Bluetooth. This could be our biggest product ever", says Sigrun Hjelmquist, president of Ericsson Components. Ericsson's Bluetooth Headset, first announced on November 15, 1999, is the first hands-free accessory to incorporate Bluetooth technology.

Contents

Ericsson Review is also published on the Web. Now with a new look and improved search capabilities. Visit us at http://www.ericsson.com/review.

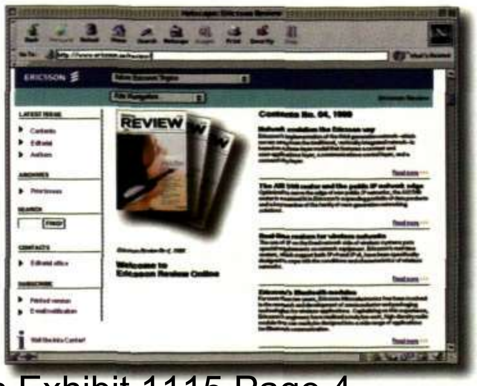

Location Labs Exhibit 1115 Page 4

Editorial

Eric Peterson

Summer 1979. I had recently returned home from my freshman year at college and begun working at a summer job for the city's Public Works department. My first assignment was to clear the community cemetery of flowers in the wake of Memorial day. For the next two weeks, I would systematically march up and down row after row after row of headstones, stopping and stooping in front of each to fill my arms with wilted and decaying flowers and then cart them off to an adjacent refuse pit. Certainly, the setting was serene, but the assignment was dismally monotonous—I desperately sought some form of stimulation that would help carry me through the day. What I needed, I decided, was music. Accordingly, I connected a pair of headphones to a rather cumbersome radio/cassette player (approximate dimensions: 40 cm wide, 25 cm high, and 10 cm deep) that I could just coax into a college backpack/book bag that I wore on my back. "Wouldn't it be great," I thought to myself, "if someone could make a mini-player that clipped onto your belt or fit in your pocket?" Little did I know that Sony had, just two months earlier, announced its first Walkman—a cute little stereo cassette player called the *Soundabout*. I got my first look at one a year later in Tokyo.

Summer 1995. I had more or less mastered the basics of my new PDA—an Apple Newton Message Pad—and began exploring some of its more advanced features and capabilities. For one thing, this device contained two slots for PCMCIA cards—for example, for modems and extra memory, or for connecting to a GPS device or a bar-code reader. As I considered the possibilities, my imagination ran wild. "Wouldn't it be great," I thought, "if I could plug in a wireless modem and connect to a data network? Why, I would be connected and online at all times!" Wholly unbeknownst to me at the time, a solution was already in the making. Best of all, one that is much bolder and promises to deliver much more than the basic messaging services I had envisioned.

Summer 2003. I was on assignment to monitor a two-day conference in Anytown, USA. Ordinarily, I'm fairly careful about what I eat, but whenever I have the opportunity to visit the USA, I make it a point to order an American-style pizza (thin crust, extra cheese, pork sausage and fresh mushrooms). This was my first stay in Anytown, so I didn't know my way around. No problem. I stepped out onto the sidewalk outside my hotel, slipped on my earpiece, pushed a button to open a dialog with my phone/PDA (in my coat pocket), and spoke:

Me: "Computer, find pizza restaurants within a ten-block radius of my location."

PDA: "Three restaurants located. Reading list: 1. *Guido's Italian Pizzas*, range 600 meters. 2. *Pizza Palace*, range 735 meters. 3. *The Stuffed Olive*, range 817 meters." (I consider these choices, and then continue.)

Me: "*The Stuffed Olive*, audio directions." (Note: I could also have requested visual directions via a map interface on my PDA.)

PDA: "Turn right and proceed along Presidential Drive to Franklin Avenue, range 43 meters." (The computer monitors my movements.) "Turn right on Franklin Avenue and proceed 782 meters." (Haven't I walked far enough? Maybe I made a wrong turn?)

Me: "Computer, I think I'm lost. *The Stuffed Olive*, audio directions."

PDA: "Proceed straight ahead on Franklin Avenue, range 38 meters."

Me: "Oh, there it is! Good. Computer off."

The Bluetooth earpiece I wore in the above scenario has already been announced and will be available on the market in mid-2000 (to learn more about the making of Bluetooth devices, read *Ericsson's Bluetooth modules*).

Ericsson's T28, T28 World, and R320 phones already include voice recognition capability and I'm betting we'll see a great deal more of this technology by 2003.

At Telecom 99, I saw a host of "third-generation" mobile terminals, which implies that manufacturers believe in and are gearing up for the mobile Internet. The immense popularity of i-Mode (Japan) and Mobitex (North America) also implies that the mobile Internet—like the Internet before it—has the potential to turn our world on its ear (see *Mobile Internet—An industry-wide paradigm shift?*). Ericsson has staked out its strategy (see *Network evolution the Ericsson way*) and developed an extensive portfolio of technical solutions (for example, see *AXI 540 router and the public IP network edge* and *Real-time routers for wireless networks*), including those for mobile positioning systems (see *Ericsson's mobile location solution*). We have also taken the initiative to create an applications alliance, to help developers to get their ideas to market.

Eric Peterson
Editor

Contributors

In this issue


Christoffer Andersson


Henrik Arfwedson


Johan Börje


Steinar Dahlin


Hans-Åke Lund


Gordon Saussy


Rob Sneddon


Göran Swedberg


Patrik Svensson


Anders Wirkestrand


Erik Örnulf

Christoffer Andersson joined Ericsson Erisoft, Luleå, in 1997. He is currently based in Silicon Valley, as Technical Expert of the GPRS Application Alliance, which he joined in May 1999. Prior to this, he participated in WCDMA system development and top-level algorithm design—for the WCDMA evaluation system as well as for the commercial product. He holds several patents pertaining to third-generation mobile systems.
EUSCRAN@am1.ericsson.se

Henrik Arfwedson is Application Manager of Bluetooth Products in the RF Modules unit at Ericsson Components. He joined Ericsson in 1995 and has worked as a hardware designer and application engineer for digital and analog circuitry. He holds M.S. degrees in electrical engineering from the Uppsala Institute of Technology and in science from Linköping University.
Henrik arfwedsson@eka.ericsson.se

Johan Börje is currently in charge of marketing Ericsson's real-time routers for wireless networks at Ericsson Packet Switching Systems. Before this, he was the Director of Product Marketing, Network Operators and Service Providers at Ericsson Belgium. He has also served at Corporate Technology at Telefonaktiebolaget LM Ericsson, where he helped coordinate Ericsson's participation in R&D programs of the European Union. He holds degrees in industrial engineering from the Linköping Institute of Technology and in management from the Cranfield Institute of Technology.
johan.borje@era.ericsson.se

Steinar Dahlin is a Senior Expert of System Design and the Research Director at Ericsson Research. He worked with the design of Land Mobile Radio Systems from 1967 until 1982, when he was appointed head of AMPS systems design. During the latter part of the 1980s and throughout the 1990s, Steinar has worked with the specification and product architecture of GSM, D-AMPS (TDMA/136), PDC and UMTS.
steinar.dahlin@era-t.ericsson.se

Hans-Åke Lund is currently Manager of Router System Technology at Ericsson Mobile Data Design in Göteborg, Sweden. He joined Ericsson in 1995 as a data communication software designer, and recently worked with the design and technical management of router functionality in the GPRS computer platform. He holds an M.S. in electrical engineering from the Chalmers University of Technology, Göteborg, Sweden.
Hans-Ake.Lund@erv.ericsson.se

Gordon Saussy is Vice President of Marketing at Ericsson IP Infrastructure, formerly Torrent Networking Technologies. He has previously held vice president of marketing positions at 3Com Corporation and Chipcom Corporation, and has worked with the development and marketing of data networking products since the early 1980s. Mr. Saussy has an M.S.E.E. from the Massachusetts Institute of Technology and holds several patents in data networking.
gsaussy@torrentnet.com

Rob Sneddon, who joined Ericsson in 1991, is currently Design Manager of the analog and mixed-signal ASIC design groups at Ericsson Components' Swindon design center—the center has been heavily involved in the design and development of Bluetooth ASICs. Rob has 21 years' experience of ASIC design and management of analog, mixed-signal and RF design. He holds a B.S. in theoretical

physics from Strathclyde University, Scotland, and an M.S. in the physics and technology of materials from Glasgow University, Scotland.
rob.sneddon@eka.ericsson.se

Göran Swedberg is currently Senior Marketing Manager at Ericsson Radio Systems, Mobile Location Solutions. He has also served as Product Manager of Prepaid Systems at Ericsson Radio Systems and as Marketing Manager of Market Intelligence Value-added Services. Mr. Swedberg is a frequent author in the international telecom press and has been named "Best Speaker" at the GSM World Congress. He holds a master of science and an executive M.B.A.
goran.swedberg@era.ericsson.se

Patrik Svensson, who joined Ericsson in 1990, is currently the Director of the GPRS Applications Alliance. Before this, he played an active role in establishing units and initiating projects for the development of data and telecom applications. At Ericsson Messaging Systems, he was Director of Product Management and Marketing. He has also served as Manager of the GSM Business Mobile Advantage program, whose purpose is to supply business users with GSM network applications.
patrik.svensson@era.ericsson.se

Anders Wirkestrand works as Product Manager for real-time routers at Ericsson Packet Switching Systems. He joined Ericsson Research in 1995 to work on IP router and ATM switch architecture, and later served as the software project manager of the AXI 510 Edge Router at Ericsson Datacom Networks and IP Services. Mr. Wirkestrand has an M.S. in applied physics and electrical engineering from the Linköping Institute of Technology.
anders.wirkestrand@era.ericsson.se

Erik Örnulf is presently Vice President of Systems Technology Management at Corporate Technology, Telefonaktiebolaget LM Ericsson, a position he has held since 1998. From 1992 to 1998 he was Vice President of Research at Ericsson Radio Systems. Before this, he was responsible for the systems architecture group for cellular systems. Mr. Örnulf began his career at Ericsson in 1969. He holds an M.S. in engineering from the Lund Institute of Technology.
erik.ornulf@lme.ericsson.se

Network evolution the Ericsson way

Steinar Dahlin and Erik Örnulf

Mobile communication is changing society's behavior. Mobile phones have become an everyday accessory for hundreds of millions of people, and they are also increasingly being used as the one and only means of personal voice telecommunication. But the emerging third generation of mobile communication means more than anywhere-and-anytime voice-communication capability. It makes information services instantly available and introduces a more powerful and flexible way of doing business. Mobile multimedia services and mobile or wireless office solutions will simplify the implementation of virtual enterprises.

In this changing environment, telecom operators and service providers will want to exploit innovative technologies and create revolutionary new services while maintaining their existing customer base. Wholly in keeping with these requirements, Ericsson is now working on its implementation of the third-generation multiservice network.

The authors describe Ericsson's move away from the traditional, vertically integrated network to the horizontally layered network, which features a content and user applications layer, a communications control layer, and a connectivity layer. They also highlight several pioneering network elements: the media gateway and its versatile use in the core network; the high-performance ATM switch; and powerful high-speed and real-time routers—all of them scalable to address every key requirement for long-term growth.

BOX A, ABBREVIATIONS

AAL	ATM adaptation layer	LSP	Label-switched paths
API	Application program interface	MGW	Media gateway
ATM	Asynchronous transfer mode	MPLS	Multipoint label switching
AUC	Authentication center	MSF	Multiservice Switching Forum
BGP	Border gateway protocol	O&M	Operation and maintenance
CATV	Cable TV	OSA	Open service architecture
CDMA	Code-division multiple access	PVC	Point-to-point virtual circuit
DWDM	Dense wavelength-division multiplexing	QoS	Quality of service
		RBS	Radio base station
EDGE	Enhanced data rates for GSM and TDMA/136 evolution	RNC	Radio network controller
		RNS	Radio network server
EIR	Equipment identity register	SCP	Service control point
ERAN	EDGE radio access network	SDH	Synchronous digital hierarchy
ETSI	European Telecommunications Standards Institute	SGSN	Serving GPRS support node
		SMS	Short message service
GCP	Gateway control protocol	SONET	Synchronous optical network
GGSN	Gateway GSN	SSF	Service switching function
GMSC	Gateway MSC	STM	Synchronous transport module
GPRS	General packet radio service	TDM	Time-division multiplexing
GSM	Global system for mobile communication	TDMA	Time-division multiple access
		UMTS	Universal mobile communication system
GSN	GPRS support node		
HLR	Home location register	UTRAN	UMTS terrestrial radio access network
HTTP	Hypertext transfer protocol	UWCC	Universal Wireless Communications Consortium
ILR	Interworking location register		
IMT	International mobile telecommunication	VHE	Virtual home environment
		VLR	Visitor locaton register
IP	Internet protocol	VPN	Virtual private network
ITU	International Telecommunication Union	WCDMA	Wideband code-division multiple access
IWR	Interworking functions		

Introduction

The information society is evolving into a globally networked economy—a development shaped by the convergence of computing, communication and broadcasting technologies. The emerging third generation of mobile communication is ushering in a new paradigm. While mobile communication is now voice-centric, offering the benefits of person-to-person voice contact anywhere and anytime, personal telephony is rapidly being transformed into a mass market of personal mobile multimedia services and terminals. Third-generation mobile communication will do much more than bring voice telecommunication capabilities to our pockets. It will make information services instantly available, including the Internet, intranets, and entertainment services. For instance, third-generation terminals might function as a video camera from which end-users can send electronic postcards and audiovisual clips.

Since terminals can be used as tools for mobile electronic commerce (e-commerce), end-users will in effect have a shopping mall in their pocket, with the ability to buy tickets, make banking transactions, and pay for goods. Third-generation mobile communication will also introduce a more powerful, flexible and efficient way of doing business. Mobile multimedia services and mobile or wireless office solutions will simplify the implementation of virtual enterprises. Similarly, appliance-to-appliance and appliance-to-people telecommunication applications will grow in importance, vastly improving security and efficiency (Figure 1).

The third-generation network architecture

The near-future envisioned

The convergence of telephony, data communication and media technologies will transform the entire telecom business. Users, operators and suppliers will be affected as an ever-growing amount of the IT world's information flow is switched to the Internet. Data traffic over the Internet will have overtaken voice traffic via telecom networks by the beginning of the year 2000, and this growth is likely to result in an enormous upswing for Internet-based systems in general, and wireless data in particular.

In this changing environment, telecom

operators and service providers will be looking for network solutions that facilitate the future introduction of a wide range of new services. They will want to be able to exploit novel technologies while maintaining the existing customer base that is served by present-day wireline and wireless telecom networks.

The development of wireline and wireless access and the inclusion of third-generation wideband radio-access technologies offer new opportunities through the evolving third-generation multiservice network.

Voice and real-time services provided over this network will be circuit-switched, whereas data traffic—using EDGE and wideband code-division multiple access (WCDMA) technology—will be packet-switched with best-effort capabilities. In a subsequent phase, IP technology will be taken all the way to the terminals—even for voice traffic (Figure 2).

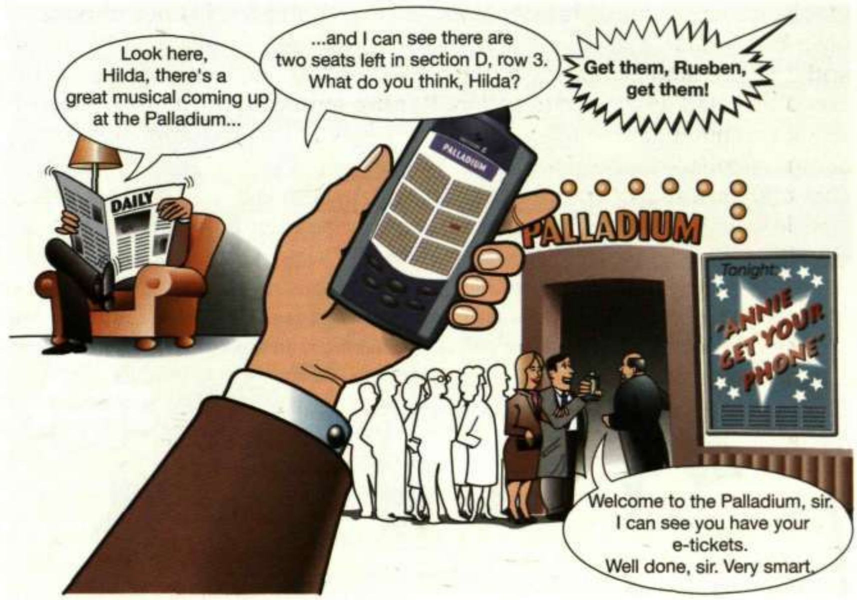

Figure 1
Buying tickets.

Driving forces

Increasingly fierce competition is one of the drivers behind the forthcoming change in the telecommunications industry. The success of network operators hinges on their ability to develop revolutionary solutions, and to build networks that have been optimized to meet the demand for services that have not yet been defined.

The telecom industry is also shifting away from technology-driven developments toward a much stronger element of consumer-driven trends. Consumers want wireless access to the Internet using small, energy-efficient terminals, and communication over broadband networks without annoying delays. To be more specific, the main driver for third-generation wireless communication—universal mobile telecommunication system (UMTS) and IMT-2000—is the ability to supplement the standardized services currently available in GSM and TDMA/136 with wideband services for which third-generation systems will ensure wide-area coverage.

Structure

Ericsson is committed to supporting this development, notably the principle of providing traditional telecommunication services and new Internet-based services over the same network. That is why—wholly in keeping with the ideas adopted by the Multiservice Switching Forum (MSF)—

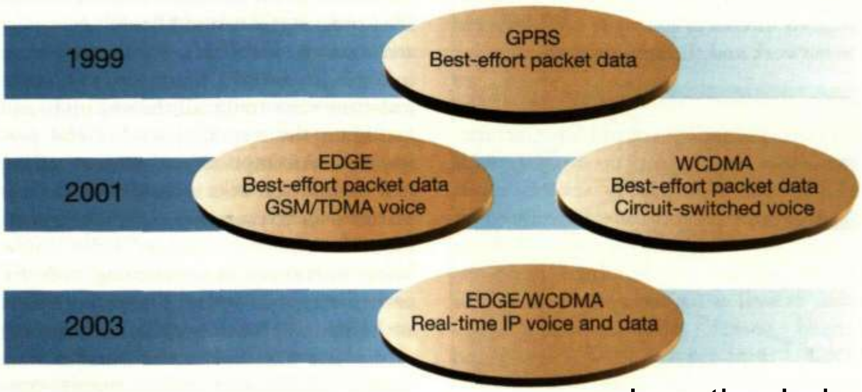

Figure 2
Evolution of wireless IP—time schedule.

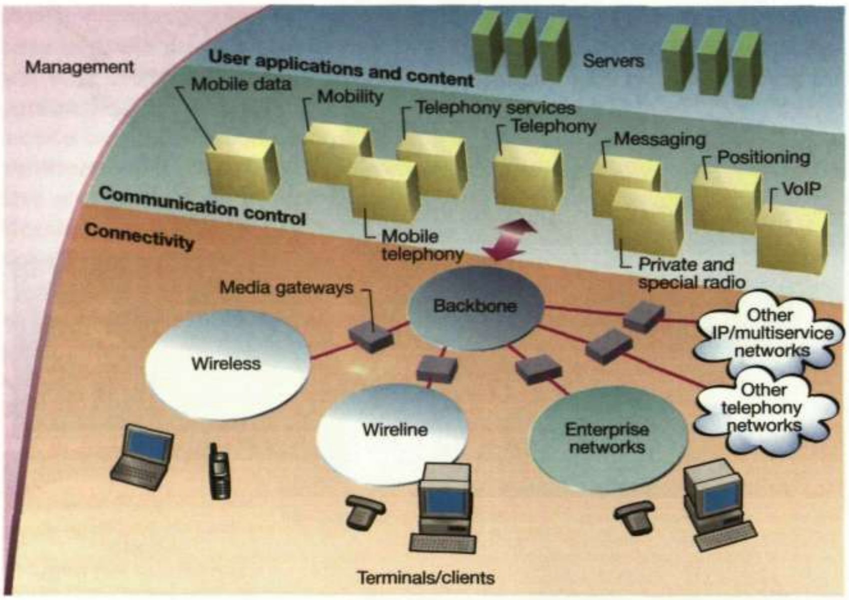

Figure 3
The new horizontally layered network.

Ericsson's implementation of the third-generation network is based on a three-layer model (content and user applications, communications control, and connectivity) that breaks with the traditional vertical integration (Figure 3).

Content and user applications

The content and user applications layer includes applications for information-centric services for which users pay. These services, which may involve e-commerce or any online activity, typically reside on off-network servers. Middleware and open application program interfaces (API) are used between the network and the applications.

Communications control

The communications control layer incorporates all the functionality needed to provide seamless and high-quality services across different public and private networks. This functionality comprises the system logic for cellular and fixed telephony and IP applications, as well as for future systems such as general packet radio service (GPRS/EDGE), third-generation WCDMA, and other multimedia networks.

Connectivity

The connectivity layer is first and foremost a transport mechanism that is capable of transporting voice, data and multimedia information. The layer is built on the Internet protocol (IP) and on asynchronous transfer mode (ATM). The connectivity-layer architecture includes backbone transport and access networks—fixed telephony, cable TV (CATV), mobile communication, and enterprise access.

The backbone architecture incorporates core and edge equipment. Core equipment transports aggregated traffic streams between service nodes on the edge of the backbone network. Edge equipment provides the bit pipe with the intelligence to guarantee appropriate quality of service, and collects all customer and traffic-specific data for accounting and billing purposes.

Ericsson's pioneering development of media gateways (MGW) has been a major contribution toward the fully integrated communications architecture formed by these three layers.

Seamless migration paths

Any viable strategy for merging voice- and data-traffic streams over a multiservice network must be based on seamless migration that retains existing customers. Ericsson's migration strategy for operators ensures business continuity while at the same time delivering new and more cost-effective IP-networking solutions. Through a successive and business-driven transition, services and applications can be shifted to inherently future-proof and scalable solutions.

Node architecture

Radio access network

When evolving into the third generation, the mobile access network will use ATM and IP to provide efficient and flexible transport and routing capabilities. Packet-switching technologies will also be adapted to support real-time voice traffic all the way up to and including the terminals (end-to-end perspective). An important component of the new mobile networks is the Cello transport platform for access products.[1] Initially optimized for mobile technology, Cello is now being introduced as a switching node for packet transport. Several Ericsson products are being built on this platform, including media gateways, IP routers, radio base stations (RBS), and radio network controllers (RNC) for various mobile systems (Figure 4).

The EDGE radio access network

With the introduction of the EDGE radio access network (ERAN, represented by the yellow section of Figure 5), Ericsson provides an IP- and server-based network architecture that has been optimized to handle a combination of packet-switched data (GPRS/EDGE) and today's circuit-switched GSM services. This architecture allows the separation of control and payload, which leads to increased flexibility and permits the platform to be tailored to different network elements; for example, media gateways and radio network servers (RNS).

The ERAN connectivity layer incorporates IP transport technology, enabling independent data traffic to run with lower priority "on top of" real-time voice traffic. Compared to circuit-switched implementations, this solution increases dimensioning flexibility and significantly saves on transmission. It also fully complies with the ETSI standard interfaces to the mobile switching center and the serving GPRS support node (SGSN) in the core network.

The UMTS radio access network

The UMTS terrestrial radio access network (UTRAN, represented by the green section of Figure 5) consists of radio network controllers, radio base stations, and radio access subnetwork operations support.

The radio base station provides radio resources and maintains the radio links to end-user equipment. The main tasks of the radio network controller are to manage radio access bearers for user data transport, to manage and optimize radio network resources, and to control mobility. The radio access subnetwork manager is a set of software for handling operation and maintenance (O&M) tasks.

Two ATM adaptation layers (AAL) are employed in the UTRAN connectivity layer. The new specially designed AAL2 is used for low-delay, real-time connections, and AAL5 is used for non-delay-sensitive packet-switched connections and for control and network-management signaling.

Core network

The third-generation cellular core network, which supports circuit-switched and packet-switched services, contains the hardware and software needed to provide end-users with multimedia UMTS applications. To ensure a future-proof core network, Ericsson is introducing a layered architecture that is based on a modular, scalable design (see the blue section in Figure 5).

Most end-user applications reside in the content and user applications layer—implemented in terminals and application servers. Concepts such as the virtual home environment (VHE) and the open service architecture (OSA) allow operators to differentiate themselves from one another and to provide unique services that enable them to maintain a higher position in the value chain. Aside from being instrumental in the standardization of VHE and OSA, Ericsson has further enhanced the OSA concept, making it an efficient tool for implementing a broad range of multimedia applications. The operator or a third-party provider can develop new features in a standard-computer environment, using the service components specified in the OSA and additional application support servers.

The communications control layer houses a number of "network servers"; for example, the MSC server, the GSN server, the home location register (HLR), service con-

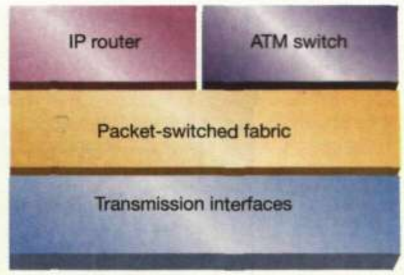

Figure 4
Cello—a common platform for IP routers and ATM switches.

Figure 5
The layered architecture of the third-generation network:
Yellow section: the EDGE radio access network (ERAN).
Green section: the UMTS terrestrial radio access network (UTRAN).
Blue section: the core network.

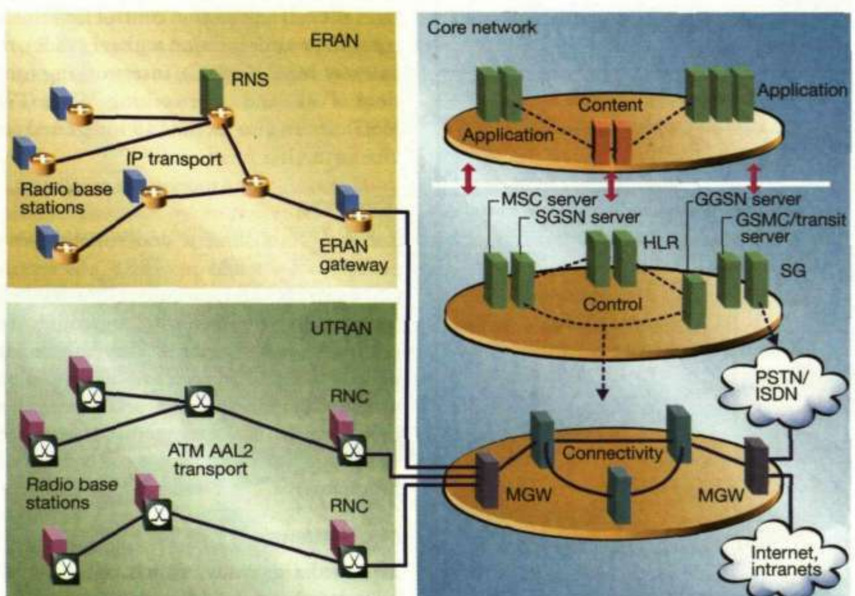

trol points (SCP), authentication centers (AUC), and the equipment identity register (EIR). These servers are responsible for controlling security, mobility management, the setup and release of calls and sessions requested by end-users, circuit-mode supplementary services, and similar functions.

The connectivity layer uses media gateways

- to process end-user data—notably voice, which requires coding and decoding, echo cancellation and multi-party bridging;
- to map quality of service (QoS); and
- to convert protocols.

This layer also serves as an access switch to backbone switches and routers, and is responsible for setting up the bearer connections that carry the media streams in the layer.

Some of the most important core-network elements are described in more detail below.

MSC server

Using the gateway control protocol (GCP), the MSC server controls calls and sessions and their associated resources for circuit-switched traffic in the media gateways. The MSC server, which provides circuit-switched services (including teleservices, and bearer and supplementary services), is responsible for charging, security, mobility management, and connection management. The MSC server is based on Ericsson's well-proven GSM-MSC technology. In addition to the basic functionality of the GSM system, it also features UMTS-specific traits, which means that the MSC server incorporates the call and session control functionality of the visitor location register (VLR), the gateway MSC (GMSC), interworking functions (IWF) and interworking MSCs (IW-MSC). It can also include an integrated service switching function (SSF).

GSN server

For the establishment and supervision of packet-data sessions in GPRS, the serving GSN (SGSN) and gateway GSN (GGSN) servers handle all control signaling. For packet-switched services, they handle signaling to the radio access network and terminal. They also handle session management, mobility management, security, charging, and short message service (SMS) functionality.

Media gateway

The media gateway, which resides at the boundary between different networks, is a self-contained network element, the main purpose of which is to provide all the functionality needed for transforming the media stream in the connectivity layer. MGWs may contain resources for voice transfer (such as voice coders and echo cancellers) and for data transformation: encryption, tunneling and address translation. They also provide signaling gateway functionality for converting lower-layer control protocols. Media gateways are controlled remotely by the MSC and GSN servers by means of the gateway control protocol. Ericsson is working to ensure that this protocol is based on an open standard, namely the ITU Study Group 16 and IETF Megaco H.248 standard, which will soon be finalized.

In Ericsson's core network architecture, a media gateway can lend its resources to any MSC or GSN server, and an MSC or GSN server can use the resources of any media gateway. Figure 6 shows a simplified view of the integrated media gateway system architecture with its external and internal interfaces and main functional modules. The media gateway supports software and hardware upgrades. Its Web-based element-management functionality includes online documentation—using the hypertext transfer protocol (HTTP). The media gateway also employs other standard protocols for transport and signaling and can interact with other nodes in synchronous/asynchronous transfer mode and IP networks. The MGWs are built on the Cello platform.

Network server platform

Network servers are built on the Tango platform (Figure 7). Tango, in turn, is based on an open operating system that executes on commercially available CPUs. The operating system is supplemented with support software to provide the availability and robustness that are needed in real-time telecommunications systems; for example:

- state replication for redundant applications;
- software load, start, upgrade and configuration control;
- O&M equipment handling, including an HTTP-based manager; and
- hardware and software supervision.

The Tango platform also includes a highly scalable, real-time database cluster for the handling of high-performance, transaction-oriented applications, such as home location registers and service control points.

The processing hardware is a cluster of loosely coupled CPUs, which provides ex-

Figure 6
Cello-based media gateway—system structure.

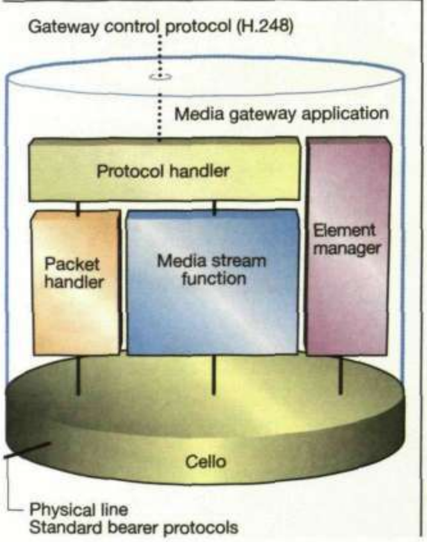

cellent scalability from very small systems up to large clusters of several dozen CPUs.

The main Tango (formerly called TelORB[2]) building blocks are:

- equipment shelves with a switched Ethernet backplane;
- processor boards with or without disk and load medium;
- IP router that interfaces the server cluster with the core network;
- operating system and support software;
- high-performance, real-time transaction software and database; and
- O&M functionality.

Migration from AXE to the new server platform

Ericsson's core network solution allows the operator to make full use of his deployed infrastructure, while at the same time optimizing the combined second- and third-generation network. This cost-effective solution also minimizes the impact of migration on the existing network.

Ericsson is developing a migration package for installed second-generation MSC and GSN nodes. This package includes the functionality needed to operate existing MSCs and GSNs as full-fledged third-generation nodes designed to handle inter-system handover and to support solutions that involve a shared transport network. The server part of the migration package is identical to that of the new UMTS product line. This creates synergies and allows an updated MSC or GSN to function as a server and media gateway (or as a fully integrated network entity) to the rest of the network.

The migration package also includes hardware and algorithms for the user data processing that is needed for upgrading a second-generation node to third-generation status. However, the server and media-gateway structure offers an alternative solution that eliminates the need for hardware additions to existing nodes. In this case—considered as the main alternative—the user data processing functions are provided in a new, stand-alone media gateway that is controlled from the server part of an upgraded second-generation node or, if necessary for reasons of capacity, from another stand-alone server. Figure 8 shows an overview of the two migration paths and the new UMTS product line. A combined GSM/UMTS network can contain a mix of migrated and non-migrated products together with new UMTS products. A co-located media gateway or radio network controller provides the transswitching functions in the optional migration product line.

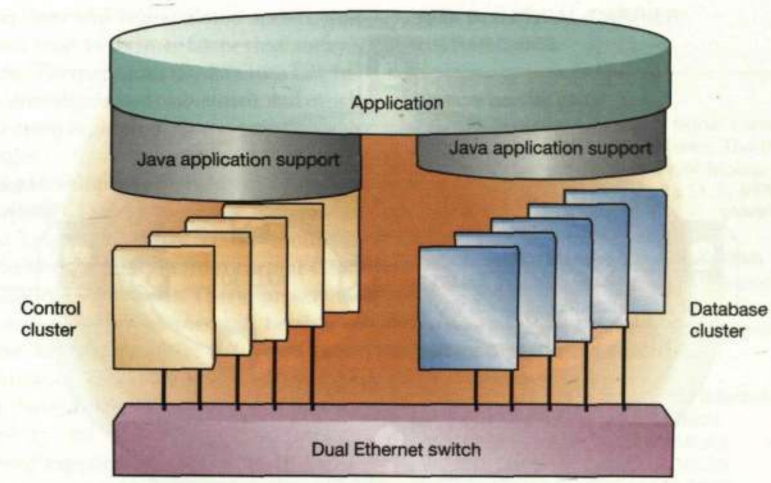

Figure 7
The Tango platform.

Backbone network

The backbone network, which interconnects the media gateways and the entities of the radio access network, is an integral part of the entire network solution. It can use any available transport technology, but the most

Figure 8
Product migration toward third-generation capabilities.

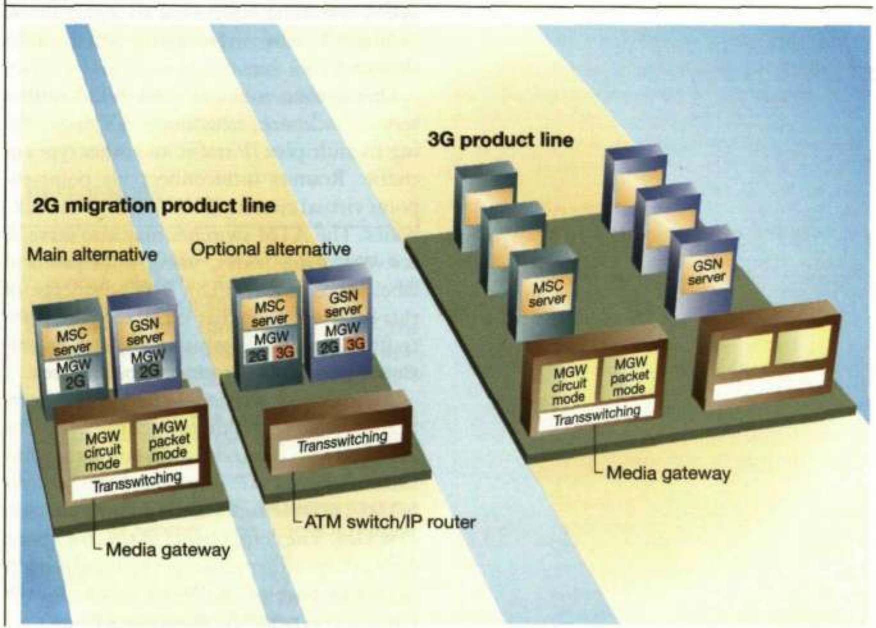

Location Labs Exhibit 1115 Page 13

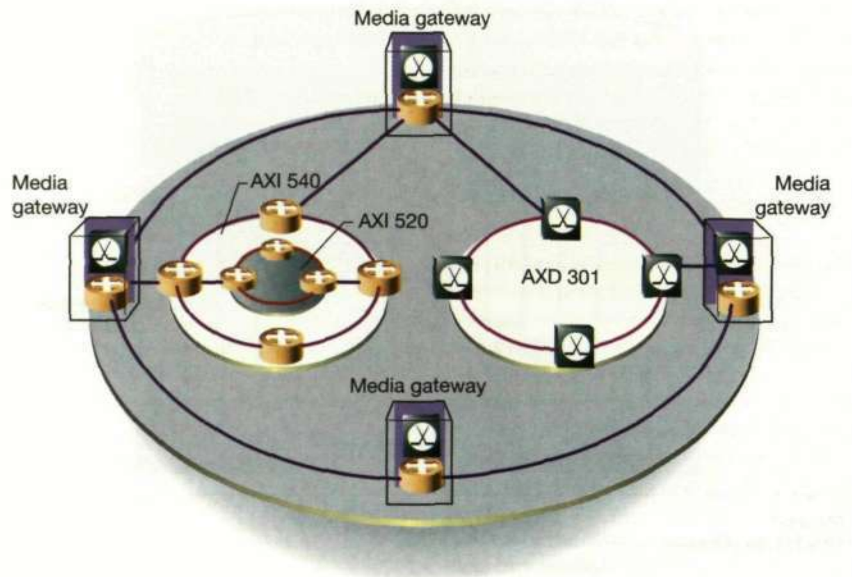

Figure 9
Backbone network. The high-speed IP backbone is based on AXE 500 products; the high-speed ATM backbone is based on AXE 300 products. The media gateways incorporate Cello products.

important future technology will be fiber, typically with a dense wavelength-division multiplexing (DWDM) layer that enables Terabits per second (Tbit/s) of aggregate bandwidth in the network. The backbone network is built as a mesh of IP routing or ATM switching nodes that are richly interconnected by point-to-point links (Figure 9).

One version is the IP-over-ATM multiservice backbone, which uses ATM switching to multiplex IP traffic and other types of traffic. Routers interconnect via point-to-point virtual circuits (PVC) over ATM backbones. The ATM switches may also serve as active routing nodes, using multiprotocol label switching (MPLS). The advantage of this architecture is that it can support voice traffic alongside IP, which facilitates migration toward the next-generation Internet.

Another version—the IP-over-SONET/SDH backbone—eliminates the ATM layer; it establishes point-to-point links between IP routers directly over SONET/SDH rings which, in turn, run over DWDM. The components of this backbone network are either real-time-enabled routers or MPLS-enabled ATM switches. In the radio-access network, these components are mostly embedded in the respective network entities, but autonomous routers and switches of identical design are needed for traffic aggregation. The same advanced quality-of-service mechanisms for voice traffic are introduced into access routers and high-speed backbone routers.

Access routers and switches

Ericsson's real-time, radio-access IP routers and switches can adapt to the specific requirements imposed by radio access technologies. Examples include optimizations to compensate for low-speed links, advanced quality-of-service handling (which involves allowing delay-sensitive voice traffic to override non-real-time traffic), and the distribution of network synchronization. These functions enable operators to make the most of total network capacity, by configuring IP networks for demanding real-time services and using spare capacity for traditional best-effort IP services.

The real-time radio-access IP routers[3] and ATM switches are constituent parts of the Cello platform. This means that they are used as components in various Cello-based products, such as media gateways for packet-data and voice services in several systems, radio network controllers for UMTS and cdma2000, and in radio base stations.

The high-speed backbone

Ericsson offers a broad set of solutions for the high-speed backbone network. The new high-capacity AXI 540 Edge Aggregation Router is designed to operate in fixed-access and point-of-presence (PoP) aggregation networks.[4] The AXI 520 provides even higher routing capacity, in addition to supporting 2.5 Gbit/s link capacity (OC-48/STM 16). For ATM-based networks, the AXD 301 is the natural choice.[5]

The throughput of the AXI 540 router can be scaled to over 20 Gbit/s, or more than 20 million packets per second, aggregated from 40,000 virtual IP interfaces per system and up to 30 OC-12 interfaces. It also supports up to 400 interior or external border gateway protocol (BGP) peers and can accept over six million routes from these peers. Its inherent scalability addresses all key requirements for long-term growth.This demands comprehensive support for routing protocols and the ability to map traffic to supported backbone class-of-service and traffic-engineered topologies. Multiprotocol label switching is a new standard that will play a key role here.

The basic concept of MPLS is to attach an extra label to the IP packet. Instead of performing a complex lookup to determine a packet's destination, intermediate nodes in an IP network can use this label to determine the packet's outgoing path. Since forwarding is based on the label, MPLS readily supports IP-based virtual private networks (VPN). MPLS can also be used to supplement the basic topology with new paths (a technique referred to as traffic engineering) because the process that manages the labels in an MPLS network is decoupled from the topology processes of the network. By providing extra label-switched paths (LSP) for certain classes of traffic, and by employing MPLS LSPs to optimize the use of a complex mesh, network operators can improve the throughput of payload packets and reduce jitter.

The AXI 520 IP router and its Internet software offer a comprehensive set of routing protocols developed by a team of professionals with vast experience of Internet scaling issues. The router is modular in design with separate functions running in protected memory storage areas. Reliability of the overall system is maximized by ensuring that modifications of one module have no secondary effects on other software modules. Another advantage of the modular-software approach is that it enables enhancements and maintenance functionality to be developed and installed rapidly in response to customer requests. The AXI 520 fully implements the new MPLS standard.

The AXD 301 is a general-purpose, high-performance ATM switch from Ericsson. It can be used in several positions in a network, scaling from 10 to 160 Gbit/s switch capacity connected through OC-12 links. It supports every service category defined for ATM including the available bit rate category, which is a best-effort service. An advanced buffering mechanism allows services to be mixed without compromising quality. Designed for non-stop operation, the AXD 301 features duplicate hardware and software modularity, which enables individual modules to be upgraded without disturbing traffic. The switching system is easily managed using an embedded Web-based management system.

Conclusion

Wireless communication using the Internet protocol will form the basis of third-generation systems and offer completely new services and transmission speeds that are more than 100 times faster than today's systems. These speeds translate into fast Internet downloads and convenient and quick e-mail services, including the transmission of images.

In its development work on wideband CDMA (WCDMA), Ericsson has placed special emphasis on successive evolution into the next generation from current GSM and TDMA/136 systems. This approach ensures that today's operators will be able to migrate into the future at their own pace, depending on market demands. Ericsson's efforts have made WCDMA the leading technology for third-generation mobile telephony supported by the majority of mobile operators worldwide.

Ericsson's implementation of the third-generation network is based on a three-layer model: the content and user applications layer, the communications control layer, and the connectivity layer. A key technology solution in this implementation is the introduction of a server and media-gateway architecture, which includes two new system platforms in addition to the evolved AXE. Media gateway products are based on Cello, and the new multimedia control servers are built on the Tango transport platform. Another key solution is the introduction of real-time IP capabilities into all router products, through coordinated quality-of-service mechanisms.

Ericsson has taken a leading position with interests in all areas of network and service management and is also leading the way in the centralized management of the entire network. An integrated network of this kind meets operator requirements for cost effectiveness and end-user demands for anywhere, anytime communication.

REFERENCES

1 Reinius, J.: Cello—An ATM transport and control platform. Ericsson Review Vol. 76(1999):2, pp.48-55.
2 Hennert, L. and Larruy, A.: TelORB—The distributed communications operating system. Ericsson Review Vol. 76(1999):3, pp.156-167.
3 Börje, J., Lund, H-Å. and Wirkestrand, A.: Real-time routers for wireless networks. Ericsson Review Vol. 76 (1999):4, pp. 190-197.
4 Saussy, G. C.: The AXI 540 router and the public IP network edge. Ericsson Review Vol. 76(1999):4, pp. 182-189.
5 Blau, S. and Rooth, J.: AXD 301—A new generation ATM switching. Ericsson Review Vol. 75 (1999): 1, pp. 10-17.

BOX B, OPTICAL-CARRIER HIERARCHIES

Optical carrier (OC)
An optical carrier for optical signal transmission has been defined as follows: The transmission base rate, OC1, is 51.84 Mbit/s. OC3 is three times that of OC1, or 3 x 51.84 = 155.52 Mbit/s.

Synchronous transport module (STM)
STM is a format in the synchronous digital hierarchy (SDH) that specifies the frame structure for lines that carry ATM cells. The transmission rate for STM-1 is 155.52 Mbit/s.

Transmission rates for optical interfaces

OC1	51.84 Mbit/s
OC 3/STM-1	155.52 Mbit/s
OC 12/STM-4	622.08 Mbit/s
OC 48/STM-16	2.48832 Gbit/s
OC 192/STM-64	9.95328 Gbit/s

The AXI 540 router and the public IP network edge

Gordon Saussy

In April 1999, Ericsson acquired Torrent Networking Technologies, a USA-based start-up manufacturer of high-end Gigabit routers for public IP networks. The team at Torrent—now the Ericsson IP infrastructure group within the Datacom Business Unit—began rolling out its new AXI 540 Edge Aggregation Router during the fourth quarter of 1999. The AXI 540 router, which has been optimized to serve the edge of new public IP networks, is a natural fit in Ericsson's expanding portfolio of data products.

In this article, the author discusses the requirements for routers at the edge of the new public IP network, and presents the key technologies and architectural capabilities of the AXI 540 router, which was built to fulfill these requirements.

From the Internet to the public IP network

Fueled by the popularity of the World Wide Web (WWW) and the increased use of intra-company as well as inter-company e-mail, the growth of the Internet during the past five years has been phenomenal—to the point of shaking the foundations of the data communication and telecommunications industries. These industries are aggressively moving toward a common vision of IP convergence that is based on the Internet protocol (IP) and a packet-switched infrastructure.

Today, the Internet is an *ad hoc* overlay built on top of existing telecommunications and wide-area data communication networks. For the most part, entrepreneurial Internet service providers (ISP) have built the Internet around existing carrier networks, with little participation from carriers (network operators). The Internet scores points for its ubiquity—it is accessible from virtually anywhere on the planet via modem and over high-speed telecommunications services in all major markets—but it falls far short of meeting the requirements for reliability and performance associated with traditional wide-area data communication services. Consequently, the vision of IP convergence, and the many bold predictions regarding IP telephony, electronic commerce and the transformation of the workplace will not be fulfilled by the Internet in its present state.

Nonetheless, new public IP networks are already being built to supplant the Internet along with major elements of the telecommunications network. The public IP network is not being built simply as an overlay by independent ISPs; instead, it is an integral part of the carrier network that is built and delivered by a hybrid of innovative, entrepreneurial carrier-service providers. The public IP network will maintain the ubiquity of today's Internet, augmented with the reliability of telecommunications service and performance enhancements that enable the most aggressive applications to operate seamlessly.

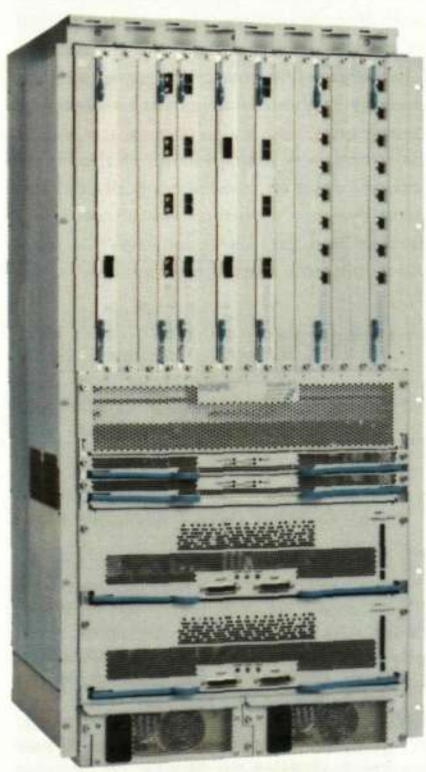

Figure 1
Rear panel view of the AXI 540 router.

Converged IP services

To some, the term IP convergence describes the transition from present-day telecommunication infrastructure technologies (SONET/SDH TDM and ATM/Frame Relay switching) to new, IP-centric solutions. This concept, which could be termed the IP infrastructure convergence, is technologically fascinating and will certainly have a major impact on carrier backbones during the next five to ten years. To others, IP convergence describes a transformation of the current multiprotocol service model to one of bundled services based on IP; that is, IP service convergence. This field is evolving very quickly and will have a highly visible impact in the near term.

The notion of IP service convergence is a simple one. Today, the typical business contracts with a variety of providers of diverse communications services. For instance, a medium-sized corporation might have different providers for basic telephony (POTS, fax), videoconferencing service, private wide-area networking (WAN), remote local area network (LAN) access, and Internet access. Each of these services potentially requires its own WAN connection and customer premises equipment (CPE). Similarly, each service probably requires subscribers to master some level of user expertise. By contrast, IP service convergence promises a simplified model that delivers all these services (and many more yet to be imagined) via a single (or optionally redundant) connection to an IP network. Consequently, subscribers will be able to access a full menu of communication services via a wireline or wireless connection from one service provider.

The Internet as we know it today can be used to prototype new services, but it is not equipped to deliver converged IP services on a large scale. The routers used to forward IP traffic within the Internet and the IP protocol suite that runs as a distributed operating system within these routers are weak in two key areas: performance and topology. The new public IP network will rely on innovations in each of these areas to enable a bundled service model.

Poor performance issues in the Internet today are easy to observe: anyone who has used a browser to surf the Web has firsthand experience of them. Statistical multiplexing of bandwidth is an inherent characteristic of packet-switched networks, and contention for use of the network creates variations in latency ("jitter") and data loss. Present-day IP routers make no effort to improve on these issues for any type of traffic; that is, all service is unreserved, best-effort and connectionless. The converged service model will require more intelligent handling of different classes of traffic within the public IP network, thereby ensuring that certain applications (such as voice) are given priority over other classes of traffic.

Topology is currently a more subtle issue than performance in the Internet. Present-day routers maintain a single topology database (routing table) for the global Internet. From any specific router to any point in the Internet, there is one unique path, dynamically maintained through the exchange of topology information between routers. The protocols currently used to exchange this information

- do not accommodate private traffic (such as a corporate WAN backbone);
- are unable to reroute traffic around congestion points; and
- do not easily allow topologies to be constrained by commercial factors.

The public IP network will require enhancements to be made to topology management, by introducing IP-based virtual private networks (VPN), topologies based on class of service (CoS), and simpler implementation of commercial constraints (for example, the preference of one backbone alternative over another due to lower transport costs).

With enhanced quality and class of service (QoS/CoS) capability and more flexible topologies that support IP-based VPNs, the public IP network will be able to deliver the full range of telecommunications and data communication services, with public Internet access as a bundled feature. New architectures, products and standards are being developed that will be used to build the public IP network and make the convergence of IP services a near-term reality.

BOX A, ABBREVIATIONS

ALU	Arithmetic logic unit
ARP	Address resolution protocol
ASIC	Application-specific integrated circuit
ATM	Asynchronous transfer mode
BGP	Border gateway protocol
CoS	Class of service
CPE	Customer premises equipment
DA	Destination address
DS	Differentiated Services ("DiffServe")
DSCP	DiffServe code point
DVMRP	Distance-vector multicast routing protocol
DWDM	Dense wavelength-division multiplexing
EBGP	Exterior BGP
FR	Frame Relay
IBGP	Interior BGP
IETF	Internet Engineering Task Force
IP	Internet protocol
IS-IS	Intermediate system-to-intermediate system
ISP	Internet service provider
LAN	Local area network
LSP	Label-switched path
MPLS	Multiprotocol label switching
NTP	Network timing protocol
OC-48	Optical carrier 48 (2.5 Gbit/s)
OSPF	Open shortest path first
PIM-D	Protocol-independent multicast, dense mode
PIM-S	Protocol-independent multicast, sparse mode
PoP	Point of presence
POTS	Plain old telephone service
PPP	Point-to-point protocol
PVC	Point-to-point virtual circuit
QoS	Quality of service
RIP	Routing information protocol
RISC	Reduced instruction set comprocessor
SA	Source address
SDH	Synchronous digital hierarchy
S-G	Source address-to-multicast group
SONET	Synchronous optical network
STM-16	Synchronous transport module 16
TDM	Time-division multiplexing
ToS	Type of service
VPN	Virtual private network
WAN	Wide area network
WWW	World Wide Web
xDSL	Digital subscriber line (various types)

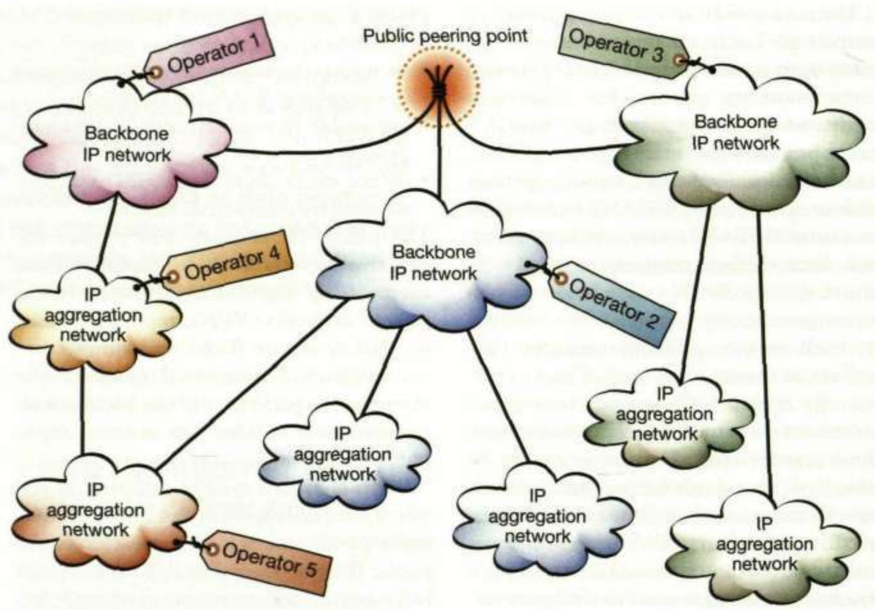

Figure 2
The public IP network: multiple tiers, multiple operators.

Building the public IP network

As with the Internet (as well as telecommunications networks), the public IP network will not be a single network; instead, it will be composed of a mesh of parallel networks that are interconnected at major peering points. Each of these parallel networks, which will be an independent IP network, may in turn be broken down into different elements that are owned and operated by different carriers. The structure of each parallel IP network can be broken down into two major elements:

- backbone networks; and
- aggregation networks.

Backbone networks

Backbone networks will be operated nationally and internationally by large carriers and service providers. They will always run over large fiber plants, typically with a dense wavelength-division multiplexing (DWDM) layer that enables Terabits per second (Tbit/s) of aggregate bandwidth in the network. These networks are always built as a mesh of interconnected IP routing nodes that are richly interconnected by point-to-point links. There are, however, *numerous ways of layering the IP network* over the fiberoptic backbone. These solutions can be loosely grouped as follows:

- IP-over-ATM multiservice backbones—these networks use asynchronous transfer mode (ATM) switching to multiplex IP traffic with other traffic across the backbone. Routers interconnect via point-to-point virtual circuits (PVC) over ATM backbones. In some cases, the ATM switches also function as active routing nodes using multiprotocol label switching (MPLS) technologies. The advantage of this architecture is its ability to support existing backbone traffic (non-IP data, non-data) alongside IP, which greatly facilitates migration toward the next-generation Internet. However, it also carries a price in terms of loss of throughput, due to ATM overhead, and in terms of ATM network management. Not all network architects believe this price is offset by the benefits of ATM.
- IP-over-SONET/SDH backbones—these networks eliminate the ATM layer, implementing point-to-point links between IP routers directly over SONET/SDH rings (which, in turn, run over DWDM). If present, non-IP traffic is carried over separate point-to-point connections in the same SONET/SDH structure. As with the IP-over-ATM approach, this solution can be implemented today, based on proven standards.
- IP-over-DWDM backbones—networks of this kind are still purely theoretical. The idea is to replace the SONET/SDH layer with a new, lightweight physical

layer that maps IP traffic directly to DWDM fibers. Proponents of this solution argue that because much of the structure of SONET/SDH is optimized for circuit switching, a simpler approach that is optimized for IP packets will result in better price-to-performance.

The vision of IP infrastructure convergence calls for rapid migration to IP-over-DWDM backbones. *In reality, however, all three* backbone solutions will evolve, and the migration toward a common solution in the backbone will take many years.

IP aggregation networks

Carriers (large and small) and ISPs will operate IP aggregation networks within service areas, which might encompass an industrial park, a few states, or even an entire country. In terms of logic, an IP aggregation network will function like a funnel: thousands or tens of thousands of subscriber connections will be transported via carrier switching and multiplexing networks into an IP aggregation point, where powerful aggregation routers map streams of subscriber traffic to backbone connections. The functionality provided by these new IP aggregation networks can be viewed in three domains: subscriber access termination, subscriber traffic processing, and backbone integration.

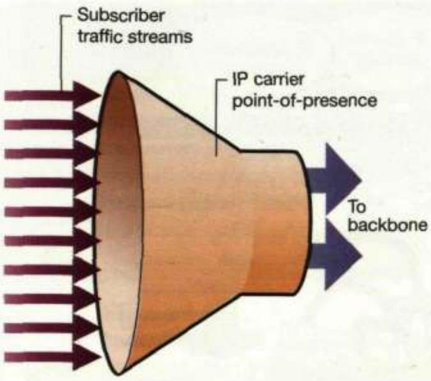

Figure 3
IP aggregation in public IP networks.

Subscriber access termination

The logical connection between the subscriber and the IP aggregation point will be a point-to-point IP connection. Carrier networks will use diverse Layer 1 and Layer 2 multiplexing and switching technologies to deliver thousands of these connections within a service area. Alternatives will range from high-speed leased lines and ATM or Frame Relay PVCs to IP/PPP/ATM connections over various digital subscriber line (xDSL) networks to wireless and cable-modem networks. The IP aggregation routers will thus need to offer tens of thou-

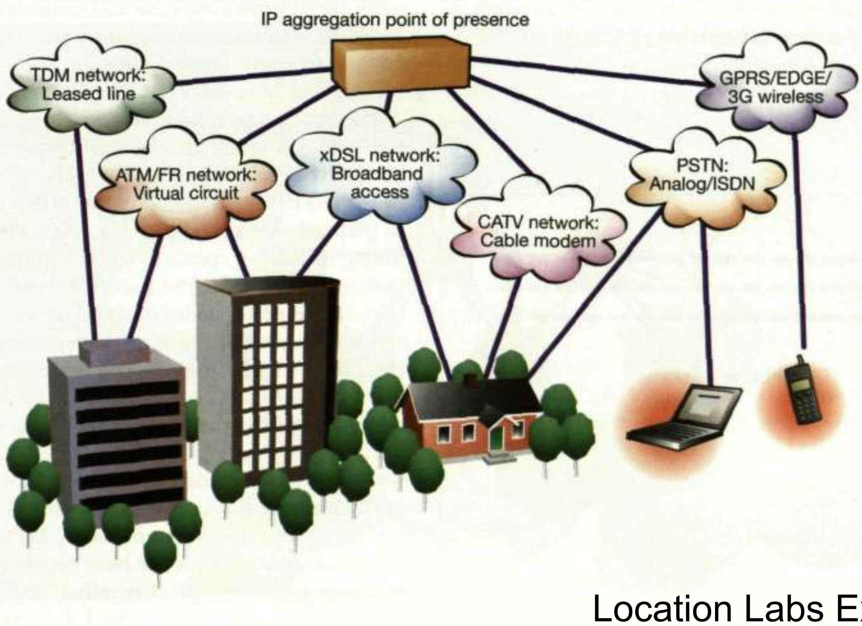

Figure 4
IP subscriber access: diverse alternatives.

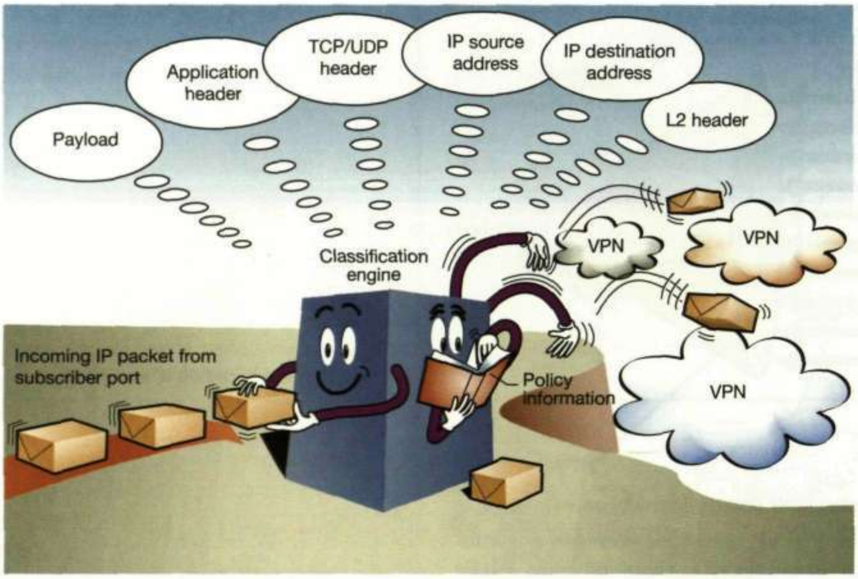

Figure 5
Classifying and handling IP subscriber traffic.

sands of virtual IP interfaces on a variety of physical ports, to ensure easy integration with diverse networks. Below the IP aggregation router, the carrier's access network can operate much as it does today—as a Layer 1 or Layer 2 switched network with minimal IP awareness.

Subscriber traffic processing

Many thousands of subscriber streams are terminated and aggregated at the IP aggregation point. Thus, the IP aggregation router must quickly be able to classify received packets according to predefined policies, per subscriber and application. This extends beyond simple classification into a finite pool of priorities—each traffic class may require mapping to a different VPN, and may require unique shaping and prioritization.

Backbone integration

A primary function of the IP aggregation routers is to route aggregated traffic onto IP backbone networks. This demands comprehensive support for routing protocols (OSPF, IS-IS and BGP-4) as well as the ability to map traffic to supported backbone class-of-service levels and traffic-engineered topologies. Two new standards for IP networks will play a key role: DiffServ and MPLS.

The new Differentiated Services, or DiffServ, standard allows IP traffic to be marked for preferential handling by the network. DiffServ redefines a byte in the existing IP header (the type of service, or ToS, byte) to include a six-bit "DSCP" (DiffServ code point) that indicates the service requirements for the packet. DiffServ-capable nodes examine the value of this field and perform forwarding operations accordingly. Of the 64 possible DSCP values, the Internet Engineering Task Force (IETF) intends to define up to 32 as "global" service classes. The remaining 32 service classes will be left open to network-specific definition. Nodes can also rewrite DSCP values in transit. The DSCP value allows certain packets to be prioritized ahead of others at each network node, thereby reducing jitter and increasing "goodput" (throughput of payload packets) for specific traffic streams (albeit at the expense of less important traffic).

The multiprotocol label switching initiative, which is much broader than DiffServ, has evolved into a family of standards within the IETF. The basic concept of MPLS is to prepend an extra label to the IP packet. Thus, instead of performing a complex lookup to determine the packet's destination, intermediate nodes in an IP network can simply process the label to determine the packet's egress path. Since forwarding within an MPLS "cloud" is solely based on the label, MPLS readily supports IP-based virtual private networks. Moreover, because the process that manages the labels in an MPLS network is decoupled from the basic topology processes of the network, MPLS can be used to augment the basic topology with new paths—a capability called "traffic

Figure 6
Integrating with IP backbones.

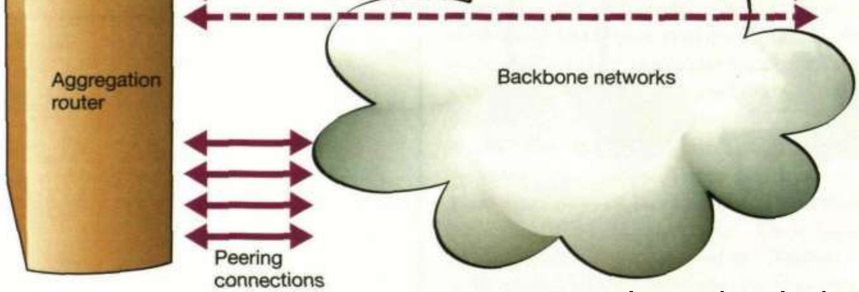

engineering." By engineering extra label-switched paths (LSP) for certain classes of traffic, and by using MPLS LSPs to optimize the use of a complex mesh, network operators can generally improve goodput and reduce jitter.

While MPLS and DiffServ represent enhancements to basic IP networks, neither makes any presumptions about the underlying protocol sandwich. Standards have been defined to allow MPLS to run directly over ATM or Frame Relay backbones as well as over any network that carries IP traffic in point-to-point protocol (PPP) frames. Because DiffServ operates strictly within the IP packet, it can be used in any IP network. Infrastructure convergence is not a requirement for benefitting from these new standards—the only criterion is that the IP routing nodes within the infrastructure be capable of supporting the new standards.

Ericsson's portfolio of solutions for new public IP networks is rich. Indeed, Ericsson has taken a leadership position with regard to the new MPLS and DiffServ standards. Ericsson's product range spans from the IP "core" all the way through the aggregation layer and into the diverse access networks (xDSL, wireless, TDM, ATM/FR, and so on) that deliver subscriber traffic to the network. Ericsson's infrastructure solutions for new IP backbone and aggregation networks are particularly powerful.

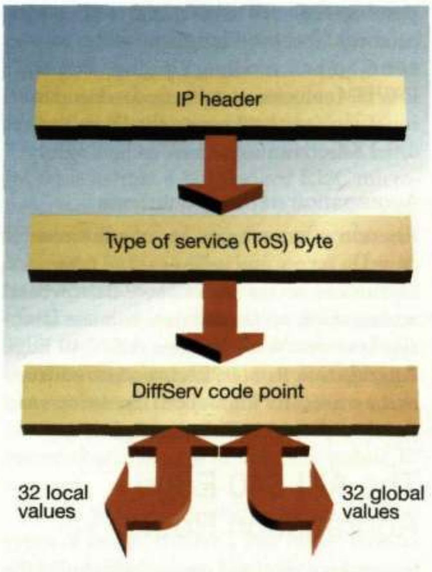

Figure 7
The DiffServ code point (DSCP) in the IP header.

Backbone networking solutions

The AXI 520 IP core router, which can be used to build IP core networks at speeds up to OC-48/STM-16 (2.5 Gbit/s) over SONET/SDH connections, fully implements the new MPLS and DiffServ standards. The AXD 301 IP/ATM switch sup-

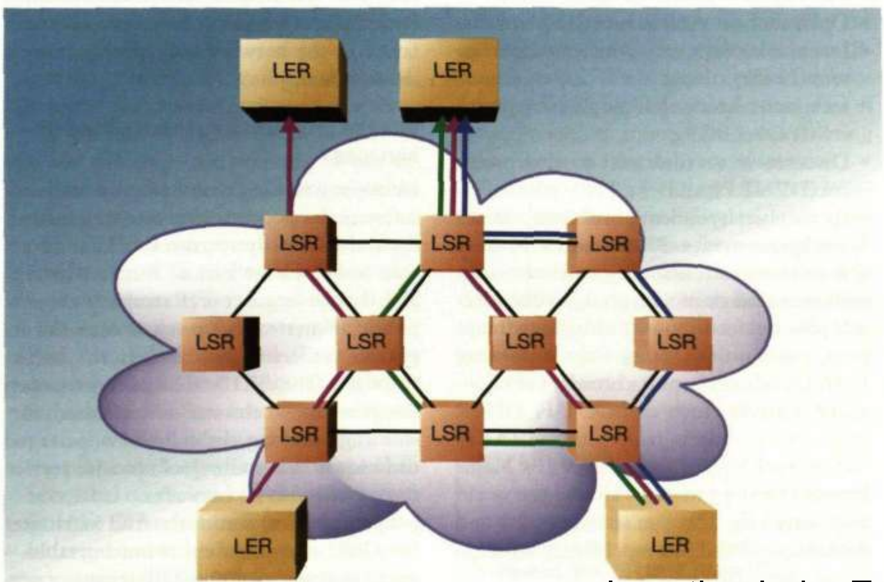

Figure 8
Multiprotocol label switching (MPLS) in the backbone.

BOX B, IPACTION ROUTING SOFTWARE

The AXI 540 router runs all routing software on a centralized and redundant "route processor," which is essentially a single-board computer that has been optimized to handle all the background operations for the router. For optimum resiliency and smooth integration of new features, the routing software runs as a series of independent processes over a streamlined UNIX operating system. Processes with the IPaction software suite include unicast routing, multicast routing, QoS handling, SNMP management, NTP service, and others.

Any router's most important characteristic is the breadth and depth of its routing software. The IPaction suite (currently at revision 2.1, the fifth production release) is rich with features and supports all popular routing protocols. It has been tested extensively in major ISP networks around the world, and demonstrated stability and interoperability in production networks. As the engineering team continues its work on future IPaction releases, this software will evolve into one of the industry's most comprehensive IP routing code bases.

ports ATM cell switching and multiprotocol label switching, and scales to over 100 Gbit/s of switching capacity.[1] Ericsson's DWDM solutions can be used to build metropolitan-area backbones for IP as well as other telecommunications technologies.

Aggregation network solutions

Ericsson offers a broad set of solutions for new IP access and aggregation networks. Optimized access nodes serve narrowband access, xDSL access and new wireless IP access systems. With the new AXI 540 Edge Aggregation Router, Ericsson also addresses the market for high-speed fixed access and point-of-presence (PoP) aggregation.

The AXI 540 Edge Aggregation Router

Service providers and carriers who build the public IP network have very specific criteria for their IP aggregation routers:

- Can the product route?
- Will the product scale?
- Is the product reliable?
- Will the product deliver advanced IP services?

Ericsson's AXI 540 Edge Aggregation Router and its IPaction routing software have been optimized to address these criteria.

Can the product route?

Ericsson's IPaction routing software supports a complete range of routing protocols (Box B):

- Routing information protocol (RIP);
- Open shortest path first (OSPF);
- Intermediate system-to-intermediate system (IS-IS);
- Interior or exterior border gateway protocol (IBGP/EBGP);
- Distance-vector multicast routing protocol (DVMRP); and
- Protocol-independent multicast, dense and sparse modes (PIM-D and PIM-S).

The software suite also implements a comprehensive range of advanced border gateway protocol features, including route mapping, communities, route flap dampening and route reflectors, plus a broad set of value-added features, such as BOOT-P, DHCP relay, proxy address relay protocol (ARP) and network timing protocol (NTP). Major Internet service providers around the world have tested the IPaction software suite and demonstrated full interoperability with existing router backbones.

Will the product scale?

The architecture of the AXI 540 router supports scaling in every important dimension. It scales throughput to over 20 Gbit/s or to more than 20 million packets per second, aggregated from as many as 40,000 virtual IP interfaces per system. It also supports as many as 400 interior or external BGP (IBGP/EBGP) peers and can accept over six million routes from these peers. The active routing table maintained at each port can hold as many as 400,000 network prefixes or 64,000 source address-multicast group (S-G) pairs for multicast routing. The overall system supports up to 50,000 classification filters at each port, mapping traffic to as many as 120,000 independent queues across the system's switch fabric. The inherent scalability of the AXI 540 router architecture addresses every key requirement for long-term growth.

Is the product reliable?

Ericsson's AXI 540 routers deliver fully resilient routing at the hardware, software, system and network levels. The AXI 540 platform uses redundant power supplies, redundant fan trays, redundant 20 Gbit/s switching fabrics and redundant route processors to ensure maximum availability at the hardware level. Ericsson's IPaction software partitions all major capabilities into separate software tasks that run in protected memory over a multitasking operating system, thereby ensuring that a fault in one task cannot affect the operation of others. What is more, the overall design of the AXI 540 system uses SONET/SDH alarms and multipath routing features to identify faults in the network and quickly reroute around them.

Will the product deliver advanced IP services?

Delivering routing that works is a matter of software development and testing; scalability and reliability require a solid architecture and implementation. But to truly enable the convergence of IP service in the new public IP network, a product must incorporate visionary technological breakthroughs. The AXI 540 Edge Aggregation Routers use a custom silicon-based forwarding path that embodies new, patented techniques for route lookup and service management in IP networks.

Every physical port on the AXI 540 router has a local copy of the entire routing table—up to as many as 400,000 destination network entries—and a dedicated ASIC that searches the table and performs packet-processing operations. Within this ASIC, specialized circuits implement a new algorithm (developed by Torrent's founding engineering team) for ultra-fast, ultra-efficient "longest-prefix match" searches in a large routing table. Since every port has its own dedicated ASIC, the system easily scales to higher and higher speeds—every new card added to the AXI 540 router carries all the "horsepower" it needs to forward packets at wire speed to any destination in the global public IP network. Ericsson was recently issued a patent on the techniques developed for this route lookup technique, called the ASIK algorithm (ASIK is an anagram of the inventor's initials).

Advanced IP services require more than just high-speed routing. The custom per-port silicon used in the AXI 540 also performs various simple table lookups, to compare each packet's entire IP header to a list of up to 50,000 filters. The packet's IP destination address (DA), IP source address (SA), Layer 4 protocol and port numbers, and type-of-service label (that is, DiffServ CodePoint) can be matched against installed filters to classify a packet for special handling. Once classified, the AXI 540 silicon-based forward path can

- discard packets immediately—this is useful if you have traffic that you do not want to forward (in implementing IP security, however, router-based packet filtering is not a substitute for a full-fledged firewall);
- police packets to a specified byte-per-second rate—the AXI 540 router allows you to implement either
  - strict policing (the immediate discarding of packets that match a filter when a specified rate has been exceeded); or
  - soft policing (the marking of out-of-profile traffic for early discard when the output port is congested);
- prioritize and queue traffic in the switch fabric—the AXI 540 switch fabric supports four priority levels and up to 8,192 independent queues at every slot; traffic that matches a specified filter can be queued independently of other traffic;
- remark the DiffServ label—at the egress port, the AXI 540 silicon-based forwarding path can change the value of the DiffServ label on the packet; this allows traffic from different filters to be aggregated for core bandwidth handling; and
- originate MPLS tunnels—at the egress port, the AXI 540's silicon-based forwarding function will prepend specified packets with MPLS labels and originate label-switched paths into an MPLS-enabled backbone network.

The ASIC technology used in the AXI 540 router is neither a full-fledged RISC microprocessor nor a simple set of state machines. It is best described as a custom arithmetic logic unit (ALU) with a unique set of built-in IP-specific operations that runs a short set of microcoded instructions on every packet. This approach, which allows some functionality to be enhanced over time through new software, maintains the wire-speed routing and classification capability that is essential at the edge of the new public IP network.

With wire-speed silicon routing for all types of traffic, the AXI 540 router ensures that every application from every subscriber sees the performance it needs. This powerful set of capabilities, together with the scalability and reliability of the system and the power of IPaction routing software, makes the AXI 540 router the clear technology leader for new IP aggregation applications.

Figure 9
AXI 540 router (front view).

Conclusion

The rapid evolution of the Internet during the past five years has shaken the foundations of industries, created and eradicated fortunes, and catalyzed social and economic change. But this is only the beginning. Ahead lies a new era of IP-based communications—the era of the public IP network.

The public IP network will transform the communications industry and enable the convergence of a variety of services to an IP-centric model. It will require powerful "core routers" for IP backbones and a new class of IP aggregation routers for consolidating subscribers and the delivery of services.

Ericsson, the leader in mobile communications worldwide, is delivering a comprehensive set of solutions for new public IP networks, from the core to the edge. The AXI 540 Edge Aggregation Router is a key member of this powerful family of next-generation networking solutions.

REFERENCES

1 Blau, S. and Rooth, J.: AXD 301—A new generation ATM switching system. Ericsson Review, Vol. 75 (1998):1, pp.10-17.

Real-time routers for wireless networks

Johan Börje, Hans-Åke Lund and Anders Wirkestrand

Through the introduction of a novel breed of routers that have been optimized for use in wireless networks, Ericsson has opened a new era of wireless IP. Being based on existing and emerging industry standards, the real-time routers will enable operators to reap the full benefits of all-IP networks—for example, the routers feature a dual IPv4/IPv6 stack for optimum scalability and security, an advanced quality-of-service mechanism for delay-sensitive traffic, hardware accelerators that yield outstanding performance, automatic configuration for reduced operational expenditure, and operation on a true telecommunications-grade platform with no single point of failure.

The Internet protocol is making its way into wireless networks. Having been enhanced in terms of efficiency and quality-of-service support, the Internet protocol has matured and can now be used to transport services in wireline as well as wireless systems. Nonetheless, the use of IP on the fixed network side of wireless systems puts special requirements on networking equipment, such as IP routers, base stations, radio network controllers and gateways.

The authors describe Ericsson's real-time routers, which have been specifically designed to cope with the special conditions of wireless networks, including network synchronization, narrowband links, and applications that are sensitive to delay.

What is a real-time router?

A router is a packet-switching device used to interconnect several different networks to form a common network that is based on IP networking technology. Based on its understanding of the networks to which it is connected, the router decides how each packet is to be forwarded. In addition, real-time routers must

- be able to differentiate between high-priority and low-priority packets;
- forward high-priority packets with low latency as well as low delay variation;
- always be free of internal congestion; and
- contain mechanisms that prevent large packets with low priority from blocking small packets with high priority on slow links.

Real-time routers enable operators to build IP networks for demanding, real-time services, while exploiting the capacity of the entire network and providing traditional, best-effort IP services using spare capacity.

Real-time routers for wireless networks

To the operator, all-IP networking has several advantages:

- simplicity and reduced costs—a unified network is easier to manage, requires less capital, and results in lower operating expenditures;
- transmission gains—packet-based technology makes more efficient use of links;
- profit from the development of sublayer technologies—new and better transmission and link solutions can easily be introduced into the network;
- new sources of revenue—a unified network layer spurs the development of applications and facilitates accessibility; and
- IP for remote access—secure virtual private networks (VPN) can generate substantial savings for corporations and residential users.

To realize the vision of a system of all-IP networks, advanced, real-time routers must be introduced throughout the wireless access network, to provide the necessary transport of real-time services. Embedded routers in base stations or stand-alone aggregation routers serve as concentrators that optimize the use of transmission links by means of statistical multiplexing.

If, however, the access network is not based on the Internet protocol, but the system to which it belongs employs an IP-based connectivity backbone, then real-time router functionality must be introduced at the gateway nodes that interface the IP connectivity backbone. Inside the IP connectivity backbone, real-time routers handle the flow of aggregate traffic from the wireless access networks. Due to the higher capacity of the links, the requirements for efficient handling are lower in this network;

BOX A, ABBREVIATIONS

ATM	Asynchronous transfer mode	MIB	Management information base
BSC	Base station controller	MPB	Main processor board
CLI	Command line interface	MSC	Mobile services switching center
DEN	Directory enabled networking	NCMS	Network configuration management system
DHCP	Dynamic host configuration protocol	NPMS	Network performance management system
DiffServ	Differentiated Services standard	O&M	Operation and maintenance
ET-FE	Exchange terminal with forwarding engine	OSPF	Open shortest path first
FEHW	Forwarding engine hardware	PPP	Point-to-point protocol
FPGA	Field programmable gate array	QoS	Quality of service
FTP	File transfer protocol	RMON	Remote network monitoring
GSM	Global system for mobile communication	RNC	Radio network controller
GSN	GPRS service node	RNS	Radio network server
GUI	Graphical user interface	SCB	Switch core board
HLR	Home location register	SNMP	Simple network management protocol
HTML	Hypertext markup language	SXB	Switch extension board
ICMP	Internet control message protocol	TCP	Transmission control protocol
IP	Internet protocol	UDP	User datagram protocol
IPsec	Internet protocol security protocol	UMTS	Universal mobile communication system
LDAP	Lightweight directory access protocol	VPN	Virtual private network

that is, the granularity of quality of service (QoS) separation is less critical, due to lower transmission delays for packets on high-speed links (Figure 1). Real-time routers for wireless networks must efficiently be able to handle small packets, low-speed links, delay-sensitive traffic, synchronization, a large number of nodes, and continuous, always-online connections.

Small packets

Present-day cellular networks transport voice in short speech frames that typically carry 20 ms of compressed voice. Each speech frame contains between 10 and 40 bytes of data. Given such small packets, the overhead of IP and transport layers becomes an issue. Thus, an effective header compression scheme must be implemented in the router. In particular, this applies to access networks. Nonetheless, header compression might also be employed to save bandwidth as the small packets enter the core network. A more important consideration, however—to minimize router load—is to ensure that the routers of the core network are not required to route each individual voice packet. In short, IP-level voice trunking must be implemented to benefit from the high speeds of the links in the core network.

Low-speed links

Access links in present-day radio networks are typically low-speed—often 2 Mbit/s or less. On these links, small, delay-sensitive packets compete for bandwidth with large, low-priority data packets. For example, if a high-priority voice packet arrives at the router on a 384 kbit/s link just as the router begins transmitting a large packet, then the voice packet runs the risk of being delayed by up to 30 ms. Obviously, this is unacceptable, since the delay budget in wireless networks is very tight for voice transmission—in particular, because other equipment, such as transceivers and transcoders, has already consumed most of the delay budget. Thus, an effective mechanism must be implemented in the router to ensure that the transmission of low-priority packets does not unduly block or delay voice packets.

Delay-sensitive traffic

A large portion of the traffic carried over wireless IP networks consists of voice and other delay-sensitive services. Accordingly, this traffic and its related control data must be given priority over other traffic. Given the low-speed links of wireless access networks, more effective QoS separation must be introduced with more classes than those required in the backbone network. That is, compared to those of ordinary data communication networks, routers for wireless networks must include more extensive and intelligent quality-of-service implementations.

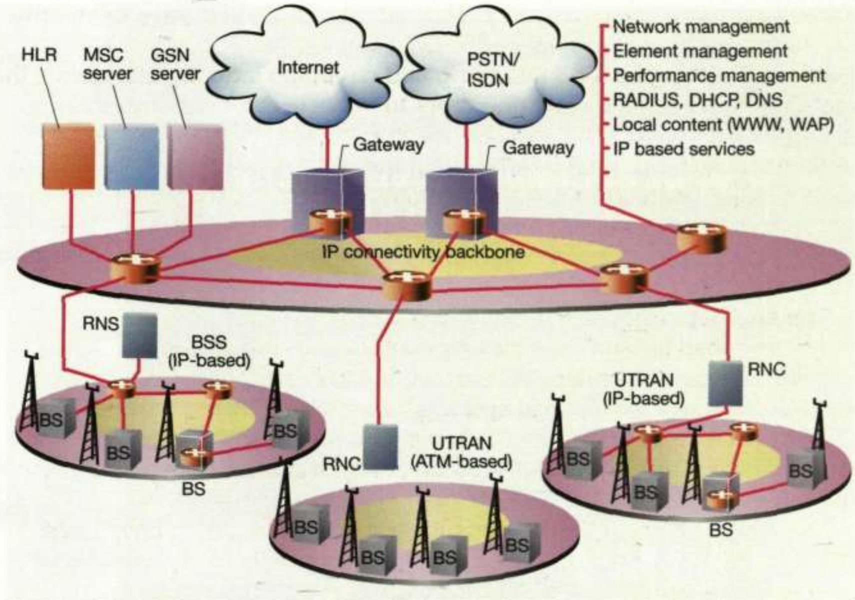

Figure 1
Future wireless IP networks will play a key role in transporting services. In this example, the Internet protocol carries voice and data in the operator network. Media gateways are used to interface external networks. The real-time router—pictured as a stand-alone router in the radio network and as an integrated router in base stations and gateways—ensures real-time characteristics throughout the network.

Synchronization

As elsewhere, radio base stations in wireless IP networks must be synchronized. The routers for wireless networks must be able to provide radio base stations with a high-quality synchronization signal that is distributed via transmission links between routers and base stations. This requirement has not been considered or cannot be met by most present-day datacom router products.

Large number of nodes

Implementations of wireless IP networks will consist of thousands of routers, many of which will be embedded in radio base stations. Consequently, new approaches must be applied to the automatic address and node configuration, which integrates advanced traffic engineering and network planning tools. Work in this area stands to benefit

Figure 2
With dual-switch cores (in right-most and left-most positions) and 26 generic slots for processors and interfaces, Ericsson's real-time router embodies high density and redundancy.

from the use of directory-enabled networking (DEN) by means of lightweight directory-access-protocol (LDAP) configuration servers. Likewise, advanced features for automatic configuration will be made available if IPv6 is chosen as the IP networking layer (Box B).

Always connected, always online

Looking ahead, we foresee that the number of Internet-enabled mobile terminals (mobile phones, smart phones, communicators, PDAs, and so on) will greatly multiply. Moreover, we will see many embedded devices and the widespread use of the Internet protocol in wireless consumer devices. Ideally, every device worldwide should have the opportunity to be connected and online at all times. Simple, convenient and immediate wireless connectivity is a key enabler for rapid growth in the market. Nonetheless, the connected society will require a huge number of IP addresses. Here again, the capabilities of IPv6 are expected to play an important role.

Ericsson's real-time routers

The Ericsson real-time router, which is based on the Cello platform[1], can be deployed as an embedded router or as a stand-alone real-time router. As an embedded router, it can serve as an integral part of other products based on Cello, such as

- radio network controllers (RNC) for UMTS;
- radio base station controllers (BSC) for cdma2000; and
- media gateways for packet data services and voice services.

When used as a stand-alone device, Ericsson's real-time router targets the aggregation of traffic in the wireless access network close to base stations. It contains all necessary functions for interfacing base stations in various IP-based systems; for example, optimizations of low-speed links, advanced QoS handling, and the distribution of network synchronization. It can also serve as an IP connectivity network router, providing up to a few Gbit/s total capacity.

Regardless of whether the router is deployed in an embedded or stand-alone configuration, it features the same routing software, interfaces, and operation and maintenance (O&M) support. To the network layer, it is merely another router, regardless of other functionality in the node.

BOX B, IPV6—NEXT-GENERATION INTERNET

The Internet protocol version 4 (IPv4) was developed in the early 1980s. With the ever-increasing success of the Internet, it became apparent to the Internet Engineering Task Force (IETF) that a new version of the protocol would be needed to handle growth and new requirements stemming from new applications. To this end, the IETF drafted a family of specifications—request for comments (RFC)—that make up the IPv6.

The main differences between IPv4 and IPv6 relate to scalability, simplicity and security. One of the shortcomings of the IPv4 is the limited number of addresses. In theory, the total number of addresses is four billion (32-bit address field), but the way in which addresses are allocated greatly reduces the actual usable amount. With the promise of communication between PDAs, embedded devices, and always-connected, always-online wireless devices, the global demand for addresses is expected to increase dramatically. Consequently, IPv6 with its 128-bit address field is a most welcome enabler.

For the sake of simplicity, IPv6 provides new automatic configuration mechanisms that make the building and reconfiguring of networks much less cumbersome. This is especially significant to

- wireless networks—where the potential number of IP nodes is large;
- companies who might want to change Internet service providers; and
- companies who might want to merge their operations.

IPv6 provides built-in support for IPsec. Thus, together with a unique and publicly routable address for each individual user, IPv6 can provide a higher level of end-to-end security than IPv4 implementations, which often have to rely on private, non-publicly routable addresses.

While IPv6 was under development, several amendments were made to IPv4. While these workarounds will extend its life, they do not resolve the fundamental issues of scalability and end-to-end security for every user.

Ericsson's recent acquisition of Telebit S/A, gives the company a commanding lead relative to routers that are capable of handling IPv4 and IPv6 simultaneously.

For more information on IPv6, see http://www.ipv6.org.

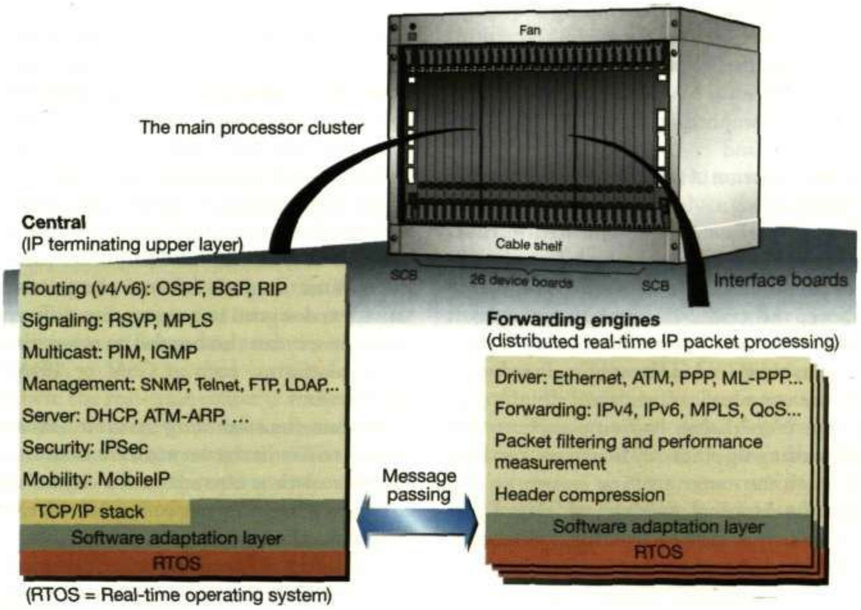

Figure 3
The real-time router comes with a distributed architecture with a feature-rich suite of software protocols that has been optimized for wireless conditions. The use of distributed forwarding engines (implemented in hardware) enables the router to be scaled linearly. This guarantees real-time characterstics at all interfaces simultaneously.

The real-time router, which is designed for dual IPv4/IPv6 operation, features a truly distributed architecture with a scalable main processor cluster and interfaces with hardware support for IP by means of forwarding engines (Figure 2) that have been implemented in hardware (FEHW).

A benefit of the distributed router architecture is that the system scales well, since packet-processing power is augmented with each added interface (Figure 3). The central processor cluster can thus use one or more main processor boards (MPB). Each MPB makes up a modern software-execution environment used primarily for central tasks in the router, including

- unicast routing protocols for IPv4 and IPv6;
- multicast routing protocols for IPv4 and IPv6;
- network management support (simple network management protocol, SNMPv3, and command line interface, CLI, through Telnet or console);
- performance management support (remote network monitoring, RMON);
- IP security (IPsec) and key distribution; and
- the handling of Internet control messages (Internet control message protocol, ICMP).

The forwarding engines (FEHW), which are used on the interface boards for IP (ET-FE) to accelerate routing and link handling (Figure 4), use configurable hardware technology (FPGA) to implement their functions. This guarantees flexibility now and in the future, since installed boards can be updated remotely. The forwarding engines feature advanced functions for individual IP packet handling:

- IP forwarding with longest prefix match;
- support for fine granular QoS separation in accordance with the Differentiated Services, or DiffServ, standard;

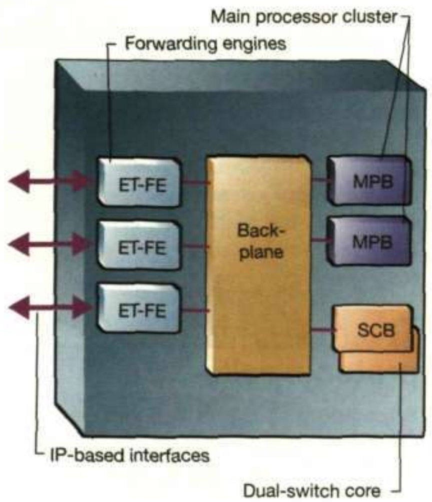

Figure 4
Single subrack configuration.

Location Labs Exhibit 1115 Page 27

- prioritized packet scheduling with support for expediated forwarding and multiple assured forwarding classes and levels;
- header compression for E1, T1 and J1 interfaces; and
- the collection of statistics for performance monitoring.

These functions are implemented in hardware in order to support a high rate of packets per second in the router. As mentioned above, the traffic in wireless networks is often dominated by small voice packets—which, compared to the traffic of traditional data communications, yield a high packet rate. Nonetheless, hardware-supported IP forwarding significantly lowers packet delay through the router, making it easier to predict the duration of delays, which is a prerequisite for delay-sensitive applications, such as telephony.

The real-time router is unusually configuration-friendly; that is, it accommodates numerous configurations in a flexible manner. The physical structure consists of

- a 19-inch rack with 26 generic slots for processor boards or interfaces;
- single or duplicate switch cores with up to 16 Gbit/s throughput capacity;
- FEHW-supported interface boards, currently for the point-to-point protocol (PPP) over a range of interfaces from 1.5 to 155 Mbit/s;
- support for 1.5 to 155 Mbit/s ATM interfaces used together with IP routing in pooled devices; and
- FEHW-supported Fast-Ethernet interfaces.

Additional interface types will be added to the real-time router as required. Interfaces with speeds above 155 Mbit/s will be added when the networks targeted by the real-time router require them.

Multiple subrack configurations (whose extra slots allow interfaces or processors to be connected) can be constructed by linking the switch core boards (SCB) of several racks or by adding switch extension boards (SXB). The system has high board density (small footprint) and no single point of failure (telecommunications-grade reliability).

The real-time router protocol suite

Ericsson's real-time routers are founded on a common software base for internetworking protocols. All protocols are available in modular distributed software. Certain time-critical functions that relate to per-packet handling, are also available as a protocol subset implemented in hardware for use directly in the interfaces with the FEHW. The software architecture features an "any-process-at-any-processor" design, which permits distributed processing of protocols in the Cello processor cluster and yields a very scalable solution to which new processors can be added to handle specific protocols. What is more, because the protocol suite was designed to be ported to different platforms, it can also be used on other, non-Cello platforms, such as GSM or TDMA base stations.

The benefits of running the same software on all routers in the network are that

- true real-time characteristics can be obtained, thanks to network-wide wireless optimizations;
- new software functionality can be incorporated into the entire network as new protocols are introduced, ensuring rapid time to market;
- uniform O&M support gives consistent handling of routers;
- uniform and comprehensive information is available to management systems, which facilitates the use of powerful tools; and
- interoperability and stability are ensured.

From the start, the protocol suite was developed for both IPv4 and IPv6 networking. That is, the structure has been optimized for handling either standard, and features transition mechanisms between the two. To work successfully with IPv4 and IPv6, some protocols require special handling. The aim has been to support either standard equally well, and to allow the network operator to decide which one (or both) will run in the network.

Another strong feature of the software is that it meets stringent security requirements; for example, it supports IPsec in combination with advanced and updateable encryption and authentication algorithms.

Real-time router management

General approach

In IP-based networks, embedded and standalone routers provide a horizontal network layer service to all involved nodes. In terms of management, this network layer constitutes a logical whole that is managed independently of other parts and services (for example, the radio resources) of the network. This approach benefits the operator in several ways:

- the same systems and procedures for managing IP resources and nodes can be applied to radio-access and core networks, thereby creating a single IP network;
- simplified end-to-end monitoring of QoS parameters (packet loss and so forth) over the access and core networks; and
- the separation of radio and transmission resource management creates an efficient division of labor and reduces training costs.

Where the real-time router is embedded in a node, it shares some resources and information with other subsystems. The implementation allows the operator to configure how this is handled; for instance, how alarms relating to the shared hardware platform should be filtered and correlated.

Element and IP layer management

The management solution for the Ericsson real-time router consists of an element manager and an IP layer manager. These products were designed with the following key principles in mind:

- minimization of O&M costs;
- optimum control of network performance;
- comprehensive support for network planning and optimization;
- highest level of security; and
- open, industry-standard interfaces.

IP layer management

A great number of routers must be used in wireless networks. To reduce the costs associated with setting up and configuring these networks, Ericsson's engineers have designed an advanced network configuration management system (NCMS) that uses an intuitive graphical user interface (GUI) together with the latest LDAP/DEN technology for the automatic configuration of routers. Configuration data is entered at a central location, after which affected nodes are instructed to fetch new configuration data from a central server. In this way, a comprehensive job, such as that of changing a network's open-shortest-path-first (OSPF) area setting, can be reduced to a minimum of operator commands entered at a central location (Figure 5).

A criterion for deploying packet technology in wireless networks is that end-to-end, real-time sessions must be achievable for delay-sensitive applications and signaling. At the same time, businesses and other premium users will conclude service-level agreements with operators, to obtain prior-

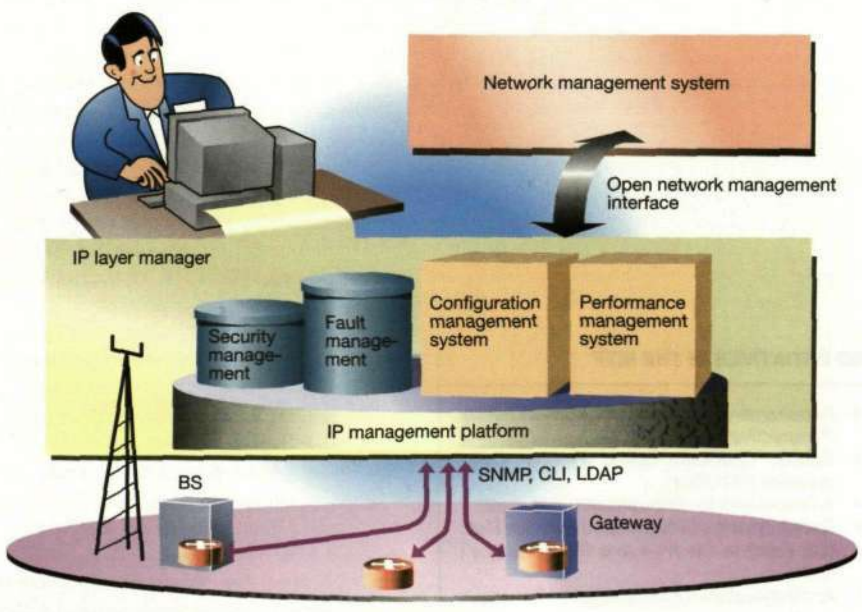

Figure 5
Operators enter configuration data at a central location. Affected nodes are then instructed to fetch new configuration data from a central server. In this way, a comprehensive job can be reduced to a minimum of operator commands.

Location Labs Exhibit 1115 Page 29

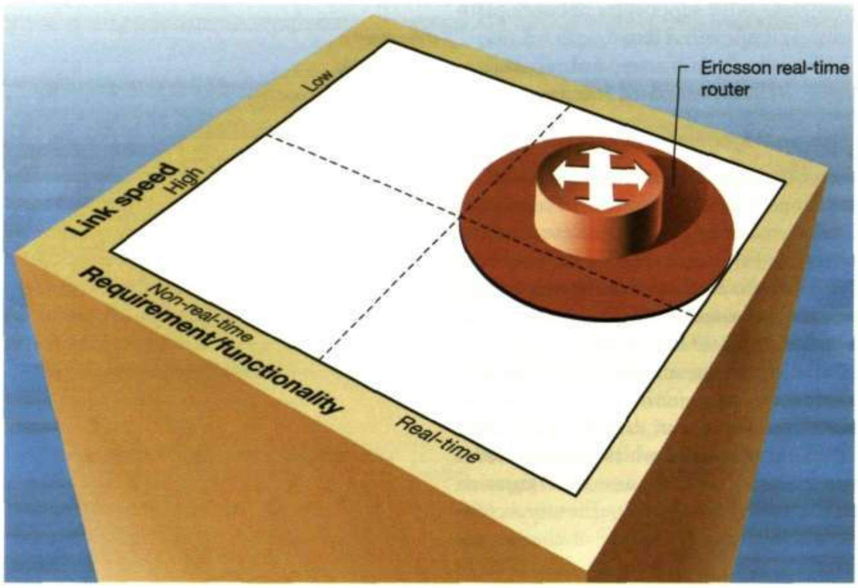

Figure 6
Positioning of Ericsson's real-time router.

ity treatment in the network. In each example, a sophisticated means of monitoring network performance is needed.

Ericsson's software developers designed the network performance management system (NPMS) to collect statistics from the network using standard SNMP and remote-monitoring management information bases (RMON MIB). Once collected, this information is employed to generate performance-related alarms and to serve as the basis for a wide range of powerful offline tools with which the operator can

- evaluate congestion levels on individual links;
- evaluate network-wide performance for DiffServ classes;
- differentiate payload trends from periodic or *ad hoc* congestion;
- obtain an overview of route choices and potential connectivity problems per traffic flow;
- evaluate "what if" conditions by simulating the introduction of extra routers and links, configuration changes, and link or node failure; and
- apply different network-optimization algorithms to determine the best configuration.

If offline analyses and modeling identify a new optimum configuration, the network can automatically be updated thanks to the integration between the network performance management system and the network configuration management system. Each of these systems will also be able to manage other vendors' nodes, provided the nodes comply with industry standards.

Element management

The real-time router implements SNMPv3. This newly standardized version of the simple network management protocol contains authentication and access control enhancements that make it possible to configure the

BOX C, WIRELESS COMMUNICATION-RELATED INITIATIVES IN THE IETF

Ericsson (sometimes in cooperation with other companies) has actively contributed to the IETF, to ensure that the standards issued by that body will evolve in a direction that facilitates all-IP solutions in wireless networks. The most important proposals include:

- Real-time Traffic over Cellular Access Networks[2];
- Load Control of Real-time Traffic[3];
- Requirements for Mobile IP from a Cellular Perspective[4];
- Robust Checksum-based Header Compression (ROCCO)[5];
- A Framework for Differentiated Services[6];
- Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers[7]; and
- An Architecture for Differentiated Services[8].

router in a secure way. The application of IPSec to management traffic may further enhance security. Moreover, different users can be assigned different privileges, and the signing and logging of events makes it possible to trace measures.

The element manager, which is a thin-client Java implementation that resides on the router itself or on a separate server, can be downloaded to any Java-enabled HTML browser. Using the router's SNMP interface, the element manager provides task-oriented GUIs for most configuration tasks. The router may also be configured using the command line interface, either over Telnet or at the console. The CLI, which provides the same level of security as SNMP, uses industry-standard commands and supports extensive scripting.

As stated above, the real-time router supports LDAP, which when assisted by services such as the dynamic host configuration protocol (DHCP) and file transfer protocol (FTP), enables the router to fetch image and configuration files from designated servers. Thus, all software-related tasks (upgrades, configuration changes, and so on) can be handled remotely and with a minimum of human intervention. This also applies to the software settings of the forwarding engines.

Relevant, standard management information bases are implemented for performance monitoring. In addition, a subset of the RMON MIB enables the supervision of local thresholds without the polling overhead associated with conventional SNMP monitoring. The element manager also supports advanced fault management using standard SNMP traps.

Conclusion

All-IP networking offers several operator advantages. However, to realize these advantages in wireless networks, advanced, real-time routers must be installed to satisfy the special needs of transmission equipment in the networks.

Ericsson's real-time routers are unique in two respects: they have been designed specifically to meet the requirements of the wireless environment, and that they will be available as embedded parts of other nodes as well as stand-alone products. To the customer, this means that the same components (interfaces, processors, and switches) can be reused for embedded and stand-alone routers, which lowers maintenance and training costs. Similarly, operators who use the same O&M system for all routers—regardless of their location (as stand-alone or embedded)—can greatly simplify network administration and employ advanced tools for network planning and configuration. Advanced features for automatically configuring the real-time routers further reduces the time and cost of deploying real-time IP networks.

Ericsson's real-time routers support both IPv4 and IPv6, which means they will support the ongoing expansion of the Internet into the wireless "always-connected, always-online" era.

TRADEMARKS

Java™ is a trademark owned by Sun Microsystems Inc. in the United States and other countries.

REFERENCES

1 Reinius, J.: Cello—An ATM transport and control platform. Ericsson Review Vol. 76 (1999):2, pp. 48-55.
2 draft-westberg-real-time-cellular-00.txt
3 draft-westberg-loadcntr-01.txt
4 draft-ietf-mobileip-cellular-requirements-02.txt
5 draft-jonsson-robust-hc-01.txt
6 draft-ietf-diffserv-framework-02.txt
7 RFC 2474
8 RFC 2475

Ericsson's Bluetooth modules

Henrik Arfwedson and Rob Sneddon

In just over a year, more than 1,600 companies have adopted Bluetooth, making it one of the fastest growing technologies ever. By 2002, it is estimated that more than 100 million mobile phones, computers and other types of electronic equipment will incorporate this technology. Bluetooth, which is a short-range radio communication link developed and introduced by Ericsson, is currently in the process of becoming an accepted world standard.

The authors describe various technical constraints and considerations associated with the development of Ericsson's Bluetooth modules. When these modules are used together with the accompanying OS-independent software stack, customers can focus their R&D resources on developing specific hardware and software for their Bluetooth applications. The authors also discuss the main issues that developers need to consider when designing Bluetooth-enabled equipment.

Background

In 1994, Ericsson Mobile Communications launched an initiative to study low-power, low-cost radio interfaces between mobile phones and their accessories. A radio-access solution was sought because it would eliminate the need for cables and overcome line-of-sight restrictions. It was also recognized that a low-cost solution would usher in wireless connectivity for a multitude of new applications and give rise to a host of associated components and devices.

For more than ten years, Ericsson Microelectronic has been involved in the research and development of semiconductor and packaging technologies for wireless applications. Capitalizing on this experience, Ericsson's engineers have realized a truly low-cost, high-density solution—Bluetooth.

Industry-wide commitment

Bluetooth was envisioned to be an open, global standard. To achieve a critical mass and to promote a joint worldwide standard, Ericsson approached IBM, Intel, Nokia and Toshiba with its concept. In May 1998, the Bluetooth Special Interest Group (SIG) was announced, and within 18 months more than 1,600 companies had signed the Bluetooth Adopters Agreement. Today, based on massive global interest and support for the technology, analysts are predicting that by 2002 over 100 million mobile phones, computers and other electronic devices will incorporate Bluetooth technology. Bluetooth specifications are issued under the authority of the Bluetooth SIG.

Universally recognized technology

Each Bluetooth unit has a unique identity. Equipment enabled with Bluetooth technology automatically searches the vicinity for other Bluetooth-compliant equipment. On contact, information is exchanged allowing the systems to determine whether or not to establish a connection. At this first encounter, the Bluetooth devices transmit a personal identification number (PIN). After that, no further identification process is necessary. Up to eight devices can operate at the same time in a Bluetooth cell. Moreover, each Bluetooth device can be active in several cells at the same time.

Basic operation

To minimize the risk of disturbance from other radio devices, such as microwave ovens, garage door openers, and so on, a frequency-hop scheme causes the system to switch 1,600 times per second through 80 channels. This means that if one frequency channel is blocked, the disturbance to Bluetooth communications will be limited. It also means that several Bluetooth networks can run concurrently without disturbing one another. First-generation Bluetooth-enabled equipment can exchange information at a rate of 1 Mbit/s at a range of up to

Figure 1
User model with local wireless connectivity. Applications envisioned for the near future.

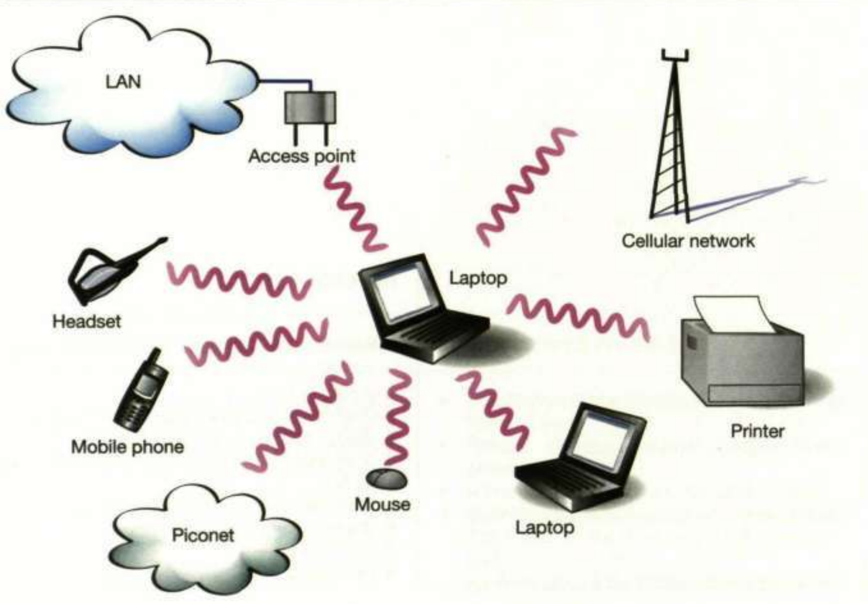

10 meters. Also, because Bluetooth operates in the global industrial-scientific-medical (ISM) band at frequencies between 2.4 and 2.5 GHz, it has the potential to become a universally recognized standard.[1]

Market confidence

Having convinced both the telecommunications industry and the wider electronics industry of the viability of the Bluetooth concept, it has been recognized that products based on Bluetooth technology must be developed, demonstrated and delivered quickly in order to maintain the interest and confidence of these market sectors.

The first goal—to produce actual product demonstrators—has already been achieved. Through enormous efforts of many diverse groups of design and development engineers, Ericsson has been first to successfully demonstrate Bluetooth compliance in real telecommunication products at recent trade shows. With market confidence in the technology established, efforts are now being focused on Bluetooth products and product support for the two major markets—that is, for

- applications within telecommunications; and
- applications within PCs and other electronic equipment.

Bluetooth-enabled products from Ericsson

Within Ericsson, more than twenty separate projects for "Bluetooth-enabled" products are currently under way. These projects, which cover a wide spectrum of applications, are being coordinated through the Ericsson Bluetooth technology council—to ensure that common Bluetooth features and key support activities are properly controlled, and that effective use of resources is maintained. Within these Ericsson products, the Bluetooth function can be integrated into the application to achieve maximum system optimization.

At Mobile Focus and COMDEX/Fall'99, in Las Vegas, USA, Ericsson unveiled the Bluetooth Headset, a practical headset that connects to a mobile phone by a radio link instead of a cable. This is the first hands-free accessory to incorporate Bluetooth technology. The Bluetooth Headset will be available on the market in mid-2000.

The Bluetooth Headset is a lightweight, wireless mobile phone headset with a built-in Bluetooth radio chip that functions as a connector between the headset and a Bluetooth plug on the Ericsson phone. To answer or initiate a call (using a voice-recognition-capable phone, such as the T28, T28 World or R320), the user presses a key on the headset. Because the Bluetooth radio link covers distances of up to 10 meters, the phone can be left in the user's briefcase, a coat pocket, or even in another room. Weighing a mere 20 grams, the Bluetooth Headset sits comfortably on either ear.

BOX A, ABBREVIATIONS

ANSI	American National Standards Institute	LTCC	Low-temperature, co-fired, ceramic
API	Application program interface	OS	Operating system
ASIC	Application specific integrated circuit	PCM	Pulse code modulation
BGA	Ball grid array	PIN	Personal identification number
CMOS	Complementary metal-oxide semiconductor	PSTN	Public switched telephone network
EMC	Electromagnetic compatibility	RF	Radio frequency
ETSI	European Telecommunications Standards Institute	RFCOMM	Serial Port Emulation based on ETSI TS07.10
FCC	Federal Communications Commission	RX	Receiver
HCI	Host Controller Interface	SDP	Service discovery protocol
IC	Integrated circuit	SIG	Special interest group
IrDA	Infrared Data Association	TX	Transmitter
		UART	Universal asynchronous receiver and transmitter
		USB	Universal serial bus
		VCO	Voltage-controlled oscillator

Figure 2
Ericsson's Bluetooth Headset is the first hands-free accessory to incorporate Bluetooth technology.

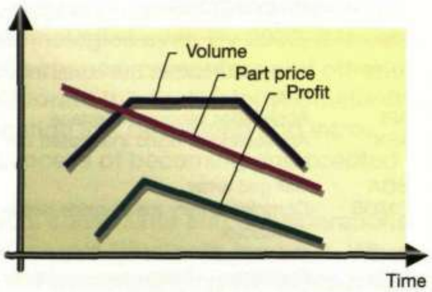

Figure 3
A typical product cycle.

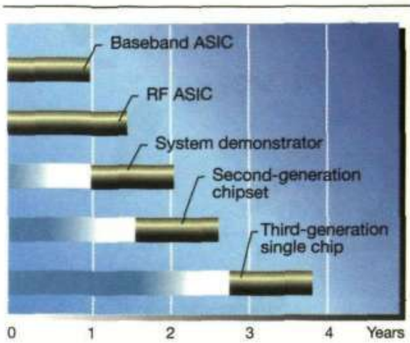

Figure 4
Typical time scale for developing a family of ASIC chipsets from scratch. The first-generation end product takes approximately two years to develop. Each subsequent generation takes about one year to develop.

Bluetooth products and product support for developers

As already stated, in order for Bluetooth to become accepted, many telecommunications, PC, and other electronic system developers must adopt the technology within their products. In terms of Bluetooth integration, these developers range from the very sophisticated to the complete novice. Accordingly, Ericsson has developed a multi-track strategy that addresses the needs of diverse customers. Ericsson offers a Bluetooth Starter Kit and a complete Bluetooth Developer Kit for Bluetooth evaluation. Finally, to support the software stack, Ericsson has partnered with Symbionics.

Bluetooth partitioning

The simplest solution for most developers is to buy a complete Bluetooth module from Ericsson. The module includes radio frequency (RF) and baseband circuits, FLASH memory, and the associated software stack. However, if developers want to integrate the Bluetooth baseband (EBC) themselves, they need only buy a Bluetooth RF module, which is designed to be compliant to the Bluetooth RF interface specification. Before undertaking their own integration, developers should carefully consider time to market, product cycles, and profitability.

Microelectronics for Bluetooth

The Bluetooth module development roadmap is closely associated with the ASIC development roadmap. On first-generation Bluetooth modules, the RF ASIC has been designed on biCMOS technology; the baseband ASIC has been designed on complementary metal-oxide semiconductor (CMOS) technology. The general trend followed by ASIC silicon technologies is toward processes with smaller geometries and lower supply voltages. Apart from reducing cost per function, this approach also permits greater levels of integration. However, although many future-generation products will benefit from the implementation of single-chip modules, there will continue to be applications in which the optimum level of integration employs a two-chip solution.

Optimum levels of ASIC and module integration

Several factors that could influence decisions concerning the optimum level of integration are:

- circuit compatibility—digital and RF circuits do not mix well;
- functional stability—future products may require baseband ASIC migration to accommodate additional functionality or to take advantage of a less expensive process while the RF interface remains stable; and
- cost—depending on the choice of technology, a multi-chip solution might cost less than a single-chip solution.

Product cycle and profitability

Sales volumes grow for a short period after a new product has been introduced. This introductory period is followed by a long period during which sales volumes remain stable. Finally, sales volumes decline as market demand decays, or the product is superseded or becomes obsolete (Figure 3).

Initial production volumes do not begin to generate profits until development costs have been recovered. Once this milestone has been passed, profit per unit increases, due to improvements in production and yield. In time, however, profit per unit declines, due to price erosion and increased competition.

The key point to be made is that the end of the product cycle should be thought of as being fixed in time. Consequently, any delay in delivering volumes to the market will result in loss of profit during the most profitable phase of the product cycle. Figure 4 shows typical time scales for the development from scratch of a family of ASIC chipsets. The first-generation product takes approximately two years to develop. Each subsequent generation takes about one year to develop. The ASIC development cycles are on the critical path for product development.

Figure 5
Typical technology roadmap.

ASIC technology and chipset roadmaps

Figure 5 shows the ASIC technology roadmap for Bluetooth ASICs. Where RF chips are concerned, the commercial availability of the manufacturing process determines which ASIC technology is used for each generation. Because digital CMOS technology is nearly always more advanced, each generation of baseband ASICs can generally be designed on smaller geometries. However, for second-generation products, the baseband ASIC design remains on the 0.25 micrometer technology node, to accommodate the option of merging the baseband ASIC with the RF ASIC on a biCMOS single-chip solution. All third-generation ASICs are based on 0.15 micrometer CMOS technology. In this case, the main reason to maintain a separate RF ASIC is to accommodate developers who only require the RF module.

Figure 6 shows the equivalent chipset roadmap. The third-generation ASICs have

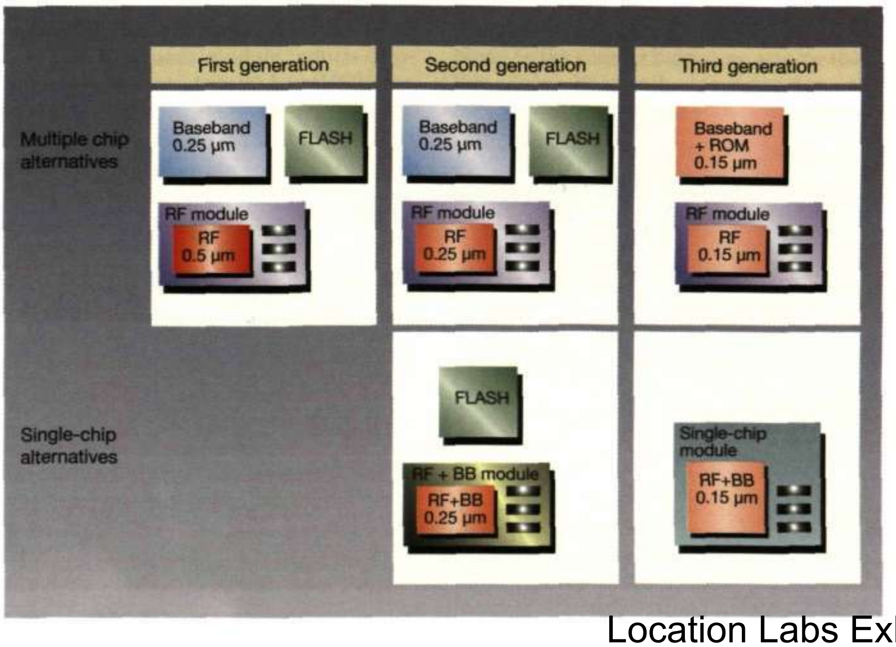

Figure 6
Typical chipset roadmap.

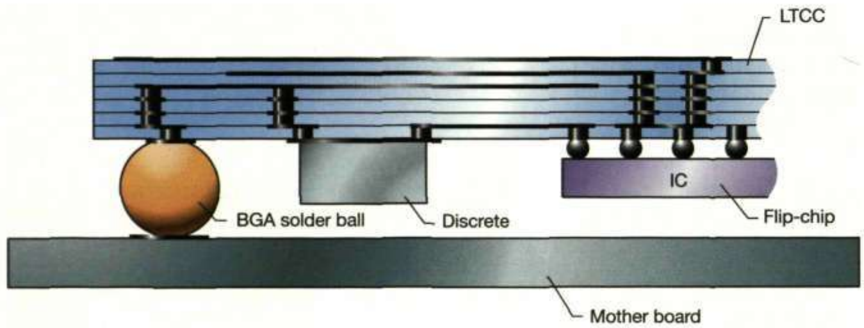

Figure 7
Illustrated profile of the Bluetooth radio module from Ericsson Microelectronics (height: 1.66 mm; width: 10.2 mm; depth: 14.0 mm).

absorbed the external FLASH memory of previous generations. This cannot be done until the product functionality is known to be stable.

Packaging

Recent developments in flip-chip and module technologies can greatly benefit the control and minimization of circuit parasitics, reduce overall costs, and introduce less stringent requirements regarding the placement of ASIC bonding pads.

RF module design

The radio transceiver function consists of an RF integrated circuit (IC) and some 50 discrete components that have been integrated into a single module (Figures 7 and 8).

Important factors that affect the selection of technology and design are:

- small dimensions, including a maximum height of 1.6 mm;
- low cost and ability to manufacture large volumes; and
- a reusable module that can be included in numerous applications.

The decision to package the function as a single-surface mountable device was made very early. The size requirements of the RF IC led developers to choose a flip-chip assembly process. This technology, which has been in use for 25 years, is considered the

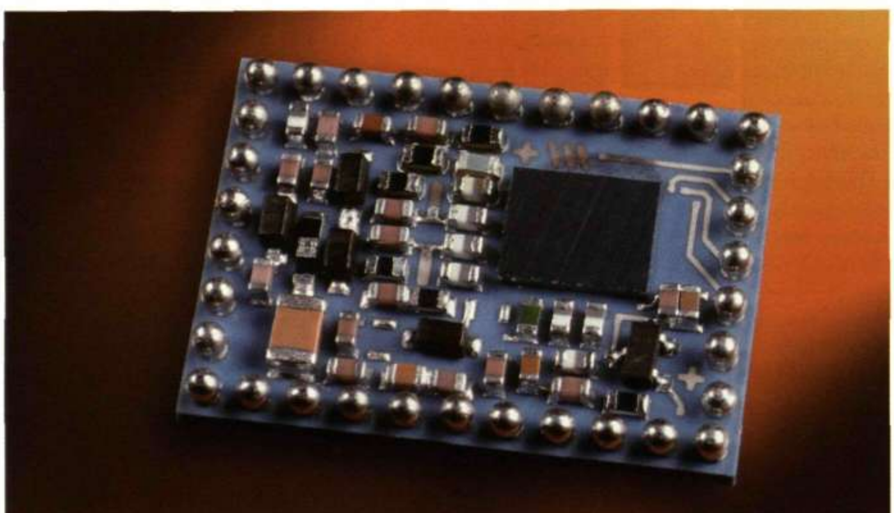

Figure 8
Photograph of the Bluetooth radio module.

ultimate assembly and interconnect solution. Until recently, however, it was an exclusive and costly technology that could be handled by only a few companies.

The small size, and requirements for flip-chip compatibility and high-frequency performance and mechanical properties led developers to select low-temperature, co-fired, ceramic (LTCC) substrates. This decision was based on the availability and maturity of the LTCC technology, vendor compatibility, and the promise of denser integration in the future.

The LTCC technique allows engineers to integrate microwave structures for the antenna filter and the receiver (RX) and transmitter (TX) baluns in the substrate. Laser trimming of the critical voltage-controlled oscillator (VCO) tank gives high performance. Dense component packing and the integration of functions in the substrate result in well-balanced utilization of the substrate's top and inner layers.

Simplicity and mainstream technology characterize the design. Because the module is assembled on one side (single-side assembly), all components can be soldered in a single operation (reflow soldering). The ground plane in the substrate and ball grid array (BGA) balls in the periphery of the substrate form electrical shielding.

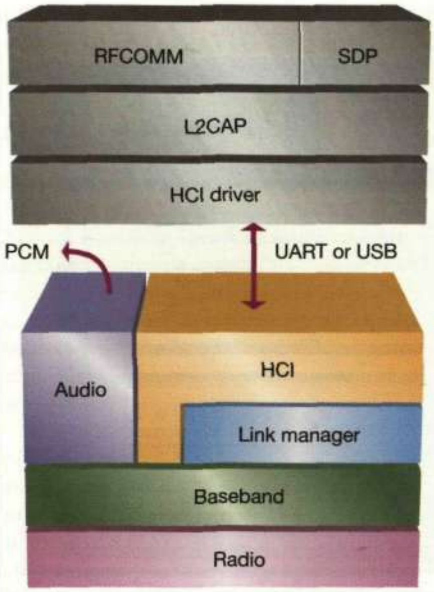

Figure 9
Ericsson's Bluetooth module and OS-independent software stack.

The Bluetooth module

Ericsson has developed a short-range Bluetooth module that consists of three major functional parts mounted on a printed circuit board. These are a baseband controller, FLASH memory and the RF module. The Bluetooth module, which complies with universal serial bus (USB) and universal asynchronous receiver/transmitter (UART) or pulse code modulation (PCM) interfaces, supports voice and data transmission. Suggested application areas that can be addressed by this module include:

- portable computers;
- handheld devices;
- cameras;
- mobile phone accessories;
- computer peripherals; and
- interfaces to the fixed-line access network.

Designing a Bluetooth-enabled product

Three important areas need to be considered before designing a complete Bluetooth-enabled product:

- Bluetooth hardware, firmware and software;
- antenna design and production; and
- application-specific software programming.

Note: in the context below, we assume that Ericsson's Bluetooth module (or an RF module-plus-discrete solution) and the operating system-independent (OS-independent) stack have been used (Figure 9).

The application

Many major applications have already been identified for Bluetooth, most of which presently communicate and transfer data via cable or Infrared Data Association (IRDA) connections (or in some instances via floppy disks). However, Bluetooth has the potential to create completely new areas of applications. When deciding how to implement an application using Bluetooth technology, engineers must first address the Bluetooth profile: all Bluetooth devices must conform to one or more profiles as stipulated by the Bluetooth profile specification

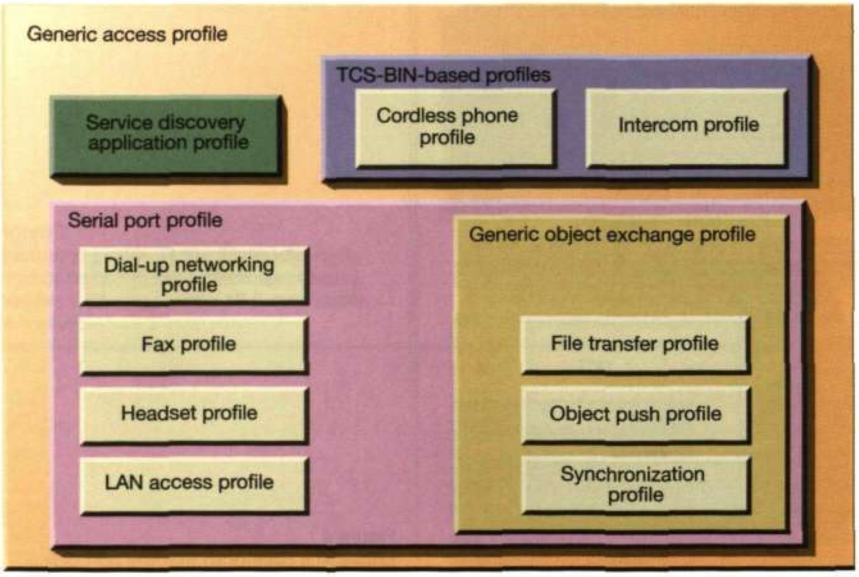

Figure 10
Profiles in Specification 1.0A.

(Figure 10), which defines mandatory and optional messages, procedures, and features. Interoperability is the main reason for defining different profiles for different usage models. New services that do not comply with a standard profile can be implemented, but they must be referred to as non-standard Bluetooth profiles.

Mechanics

The definition or selection of device packages should be started as early as possible during the product development phase. Prime considerations are

- package materials;
- parasitic effects; and
- size.

Ordinarily, between six and twelve months are needed from design to volume production of a new package. Obviously, the mechanics of the product put constraints on the choice of components and antenna implementation.

Bluetooth components

The height and size constraints of the application will influence the decision to use either a discrete or modular solution when implementing the physical hardware with Ericsson Bluetooth components, including

- the Ericsson Bluetooth module; or
- the Bluetooth RF module, plus baseband FLASH, and associated components.

Thanks to Ericsson's Bluetooth module and the OS-independent software stack, most of the developer's work on the software and hardware core can be devoted to the application level.

Antennas

Antenna design is affected by the availability or lack of space (physical volume). Developers must thus obtain input from the antenna designers early on in the project. Important parameters are space (volume), package shielding of the product, and production tolerances. When designing the mechanical structure of the product, developers need to have decided whether they will use an omnidirectional or a high-gain antenna—Ericsson can provide information and support.

Ericsson Bluetooth development tools

The development kits provided by Ericsson greatly simplify the design of application-specific hardware and software. Two separate hardware platforms are available:

- the Ericsson Bluetooth Developer Kit, which contains discrete components and provides access to all hardware interfaces. This kit could be regarded as the standard platform for Bluetooth product development; and
- the Ericsson Bluetooth Starter Kit, which uses the Bluetooth module for implementing Bluetooth functionality.

The following application program interfaces (API) will be available in each kit:

- API on HCI;
- API on RFCOMM; and
- API on SDP.

Most profiles depend on the RFCOMM protocol and the application-specific software that works on top of RFCOMM.

Ericsson Bluetooth hardware and firmware

Three hardware interfaces are used with the Bluetooth module or a discrete equivalent solution:

- UART and PCM voice over the PCM interface;
- USB voice over the USB interface; and
- UART or USB is used for control and data transmission.

A host-controller-interface (HCI) driver for the chosen hardware interface is needed to drive the Bluetooth module, which provides the functionality for receiving HCI commands from, and sending HCI events to, the host application; for setting up links; and for handling timers. The module also handles the actual transmission of data and voice traffic. The formats of HCI commands, events, data packets, and other Bluetooth protocols have been formalized *in Specification of the Bluetooth System vl.OA.*

OS-independent software

High-layer protocols are implemented in source code. Ericsson will license the portable OS-independent software stack (written in ANSI C code) for mass production of devices using Ericsson Bluetooth products. Thanks to the OS-independent stack, software developers need only work with high application-specific layers. Nonetheless, all implementations of the application software must comply with the profile specification: *Specification of the Bluetooth System Profiles vl.OA.*

Qualification

The Bluetooth SIG has outlined a formal qualification program. To pass this program, all products must adhere to the following Bluetooth-specific criteria (use of the Ericsson Bluetooth module will fulfill the first two criteria):

- radio link;
- protocols (lower layers);
- profiles; and
- end-user information.

The qualification program protects the value of the Bluetooth brand by ensuring product interoperability and radio link quality. Manufacturers whose products pass the qualification program may take advantage of IP licensing and use the Bluetooth brand name.

Type approval

Before a product may be introduced into the market of any given country, it must be "type approved" according to national or regional regulations. Regulations are grouped according to

- spectrum constraints;
- radio and telecommunication requirements;
- electromagnetic compatibility (EMC) and safety;
- connection to the public-switched telephone network (PSTN) and encryption; and
- labeling, test reports and procedures.

Ericsson's Bluetooth module will be pre-certified according to regulations by the Federal Communications Commission (FCC) and the European Telecommunications Standards Institute (ETSI). Moreover, it contains the procedures for putting the device into all necessary test modes.

Conclusion

Bluetooth—which is a wireless technology standard introduced by Ericsson, IBM, Nokia, Intel and Toshiba in May 1998—uses short-range radio signals to connect two or more (up to a maximum of eight) devices, such as mobile phones, computers, handheld devices, household appliances, automobiles, consumer electronics and office equipment.

The Bluetooth module is a standard component from Ericsson Microelectronics that can easily be designed into a wide range of applications for Bluetooth communication. Thanks to the advanced flip-chip assembly of the integrated circuit, a low-temperature, co-fired ceramic substrate, and ball-grid array assembly to the application PBA, the radio module is relatively small. Ericsson will begin large-volume manufacturing of the module during the first quarter of 2000.

Companies who use Ericsson's Bluetooth module, the OS-independent stack, and other components for designing Bluetooth-enabled devices can focus their R&D resources on developing application-specific hardware and software.

In May, 1999, Ericsson launched the Ericsson Bluetooth Developer Kit, a "toolbox" of equipment that provides developers with a flexible environment for designing Bluetooth applications in less time and at lower cost, and for integrating Bluetooth technology into a wide range of electronic devices.

REFERENCES

1 Haartsen, J.: Bluetooth—The universal radio interface for ad hoc, wireless connectivity. Ericsson Review Vol. 75 (1998): 3, pp. 110-117.

Mobile Internet—An industry-wide paradigm shift?

Christoffer Andersson and Patrik Svensson

The Internet has had an overwhelming effect on the way we do business. It has simplified many of our daily tasks and engendered new lifestyles. As a marketing channel, it has displaced entire organizations and forced corporations to overhaul their operations. The Internet has also changed our way of viewing software, inspiring a new approach to application development. The question is, will the introduction of the mobile Internet give rise to a new paradigm?

Does the mobile Internet represent a true business opportunity? Today, data usage on mobile networks accounts for less than 1% of all traffic. Why, then, should we expect the majority of mobile users to subscribe to the mobile Internet? If the market for the mobile Internet is to take off in a big way, when will this occur, and how?

As with any new trend or innovation, the mobile Internet has a stable of stalwart evangelists. Notwithstanding, the initial reaction of most onlookers is "wait and see." But after having witnessed the introduction and growth of the Internet, we have also learned not to pass hasty judgement on the unknown.

In this article, the authors share a vision, and point out numerous indicators that signal the imminent arrival, of the mobile Internet.

The mobile internet

The infrastructure of the mobile Internet will be based on the Internet protocol (IP), which means we can reuse the hardware, software and competence we have devoted to building the Internet. The infrastructure will also be built around general packet radio services (GPRS) technology, which introduces packet data into GSM and TDMA (IS-136) systems. GPRS enables operators to use the same coverage for data and voice services, and establishes the IP network backbone that will be used in third-generation mobile systems. The key feature of GPRS technology is that it enables users to be connected and online at all times: "always connected, always online."

As with the Internet, the mobile Internet will consist of a wide variety of services and applications that run on top of the infrastructure. However, the reader should note that the mobile Internet will not be a slow copy of the Internet. Instead, applications for the mobile Internet will mainly capitalize on the fact that its users are mobile and continuously connected to the network (if they so desire). In time, today's Internet will probably become a subset of these new services. The power of mobility is characterized by communication that is carried out as stipulated by the user.

I-mode and Mobitex

Even now, we see examples of how the mobile Internet is being successfully implemented. In February 1999, the Japanese mobile operator, DoCoMo, introduced I-mode, which is built on IP-over-packet PDC (P-PDC; that is, packet data for the Japanese standard). By September 1999, I-mode boasted 1.4 million subscribers, and the service continues to grow at an astonishing rate of 90,000 new subscribers per week—by March 2000, I-mode is expected to have over four million subscribers.

Today, I-mode supports some 230 applications and services. Examples include banking, ticket-reservation (concerts, movies and other events), travel-reservation services (flights, buses, hotels), and a variety of entertainment services. Several factors have contributed to the success of I-mode:

- it is simple;
- it supports only one size of phone screen and a subset of the hypertext transfer protocol (HTTP);
- traffic and subscription fees for the service are inexpensive; and
- it targets the Internet generation, which has the necessary technical savvy and purchasing power. Note: this marketing strategy represents a fundamental departure from traditional approaches, which have targeted traveling businessmen—in many circles, this group of consumers is considered to represent the early adopters of mobile communication technologies.

The Mobitex network in North America is another example of successful implementation of the mobile Internet. Initially, this network was solely used for vertical industry solutions, such as dispatch, utilities and field service. However, a little over a year ago, when the focus of the network shifted to messaging and two-way paging, the market began to surge. Palm Computing, Inc.,

BOX A, ABBREVIATIONS

EDGE	Enhanced data rates for GSM and TDMA/136 evolution	LAN	Local area network
GAA	GPRS Applications Alliance	MOA	Mobitex Operators Association
GATE	GPRS applications test environment	PDA	Personal digial assistant
GPRS	General packet radio services	P-PDC	Packet data over personal digital communication
GPS	Global positioning system	SMS	Short message service
GSM	Global system for mobile communication	UMTS	Universal mobile telecommunications system
HTTP	Hypertext transfer protocol	WAP	Wireless application protocol
IP	Internet protocol	WCDMA	Wideband code-division multiple access

a subsidiary of the 3Com Corporation, played a key role when it introduced the Palm VII—the Palm VII is a Palm computer with a Mobitex radio that supports Web clipping, a simple but effective way of utilizing the narrowband mobile Internet. The new success of Mobitex can be attributed to more mature products but also to a maturing market that understands and is ready to use the technology.

From these two examples we see that it is possible to create a substantial market using simple technology and low bandwidth. We also see that success comes from being user-oriented and taking a new marketing approach.

World market

However interesting these two examples might be, neither the I-mode nor the Mobitex market is substantial enough to make us believe in a large global market. To motivate investments for market sizes that exceed $100 billion per annum, we must see more compelling evidence that points to a major worldwide market.

In January 1999, the Internet consisted of an estimated 43 million servers. This count did not include servers behind firewalls, in intranets, or behind dial-up connections. A server constitutes at least one content provider and one service, but more often than not each server hosts at least ten services. In other words, the Internet constitutes millions upon millions of services. In January 1999, the Internet had at least 150 million users, and experts feel certain this number will reach one billion by 2004. The Internet is an indisputable market of the masses, but as with nearly all markets, the early growth was slow.

Mobile telephony is another major market. Today, there are an estimated 250 million GSM subscribers worldwide, and the number of mobile subscribers (all standards) is expected to exceed one billion by 2004. Just think—two major markets with similar growth projections! What is more, recent market studies conducted in Europe indicate that more than 50% of all GSM subscribers also subscribe to the Internet (and visa versa). Thus, we see a huge potential market for mobile Internet. However, to create a market that attracts and retains these users, several conditions must be met:

- good coverage;
- mass-market pricing;
- endorsement by major players (who represent end-users, financiers, and suppliers);
- pioneering users who lead the way in buying products and services, and who are willing to put up with the pitfalls of a nascent market.
- easy-to-use devices; and
- a plenitude of applications and services that meet the needs of individual mobile end-users.

Coverage

In GSM, the field strength of the radio signal for a standard call with acceptable voice quality must be at least 8 or 9 dBm. However, simulations of GPRS indicate that the signal strength only needs be 5 or 6 dBm to send packet data. The GPRS standard actually extends GSM coverage for data services, since data is less sensitive to signal strength than voice. Persons who use the short message service (SMS) might already have experienced this effect.

Mass-market pricing

The addition of GPRS functionality to GSM is essentially an upgrade. Depending on the supplier, the extent of the upgrade (in terms of hardware, software and labor) will vary. Assuming the upgrade is to a network supplied in full by Ericsson, we can draw a number of positive conclusions:

- A comparison of investments in an existing GSM network—either to accommodate additional voice subscribers or to add packet data capabilities—shows that operators stand to gain six times as many data subscribers as voice subscribers per

BOX B, MOBITEX

Work on the Mobitex specification began at Sweden's Televerket (now Telia), in the early 1980s. In 1987, the first Mobitex network was taken into commercial operation. In 1988, Ericsson, in cooperation with the Mobitex Operators Association (MOA), became responsible for further development of the Mobitex infrastructure. MOA is responsible for the Mobitex Interface Specification, which specifies the air interface and the protocol used by permanently connected terminals.

Mobitex is a network technology designed exclusively for two-way, wireless data communication. Mobitex has, since its commercial start in 1987, evolved into a system of very high quality and with numerous functions and applications. The bit rate, however, is only 8 kbit/s, which puts severe restraints on usage. Nonetheless, thanks to the huge success of the first basic information services, such as stock quotes and news, companies are now scrambling to introduce a host of new services. Two examples include the United Parcel Service, which gives subscribers up-to-date information on the location of parcels, and Merriam-Webster, which offers complete online versions of its dictionaries.

In the USA, the Palm VII personal digital assistant (PDA) uses the Mobitex radio interface to access packet data services.

invested dollar. Furthermore, the typical investment scenario will probably be carried out in two or more phases. The first phase is to offer full data coverage throughout the entire network. This done, the operator will increase network capacity as needed. In GSM networks, this can be accomplished while maintaining a positive cash flow.

- Data traffic in a GSM network can be offered on a best-effort basis; that is, data is transmitted in time slots that are not occupied by voice traffic.

Students of price-elasticity theory know that in an optimum market only 10% of the market is willing to buy a product or service at one price level and as much as 90% of the market will buy at another price level. Ordinarily, the initial price of a product or service is set somewhere between these two levels. For the mobile Internet market to take off in a big way, the subscription price must be set at a level that attracts approximately 50% of the market during the first few years after its introduction. During this period, however, operators must also be able to sustain a reasonable profit margin.

Figure 1
The user has easy access to online services via a touch-screen that is connected to a GPRS-enabled terminal. The touch of one icon brings up a map of the area; the touch of another icon calls up a report on current traffic conditions; the touch of still another icon fetches current events in the area (concerts, parades, and so on).

Investments and endorsements from major players

Are content providers and major user communities willing to invest in the mobile Internet? The automotive industry—in particular that of Europe—has begun to view drivers and passengers in a new light. Today, many millions of people spend over an hour a day commuting in their automobiles. Taken as a whole, this body of commuters constitutes a huge, but still relatively unharvested market. For instance, in Los Angeles, with a population of eleven million, approximately nine million people drive their cars everyday. Some products developed for this market by early entrepreneurs include navigational aids and safety systems. Many other products and services will soon follow. Examples include the automatic transfer of engine-related statistics prior to service, and up-to-the-minute information on traffic conditions—accidents, road construction, and alternate routes.

Many automobiles currently employ a computer to regulate engine power, fuel consumption and other operating attributes. Some automobile manufacturers are eyeing the possibility of downloading auto-performance upgrades to these on-car computers. Today, only luxury automobiles provide a phone as a standard accessory, but many of the auto manufacturers with whom we have spoken plan to make GPRS phones standard in all new cars.

The trading industry also perceives the mobile Internet as a significant new business opportunity. For example, Charles Schwab, one of the largest brokerage firms in the USA, plans to offer mobile trading. Customers who subscribe to mobile, online trading services will not be bound by the constraints of time (office hours) or location to check on or update their portfolios. Instead, they will be able to access all necessary information and trade whenever they like.

The banking industry has also identified business opportunities in the mobile Internet, including

- automated banking services;
- clearing houses for electric commerce (e-commerce); and
- competing as service providers.

As Internet banking becomes more prevalent, banks are looking beyond the immediate horizon. They see mobile devices as being a tool from which users can manage their finances. Indeed, because it is easier to integrate secure transactions into a mobile device, many banks actually prefer the mobile Internet to the fixed Internet.

To enable the commerce of goods and services over the Internet, a neutral, third party must clear the financial transactions and control the integrity of buyer and seller. No one is better set up do this than banks and financial institutions. Likewise, e-commerce can further automate banking services and reduce costs, which makes it an attractive way of doing business. Some banks even envision the day when mobile e-commerce will relieve us of bills and other costly payment methods.

Recent mergers in the banking and insurance industries have created efficient and competitive service brokers. Banks are very competent at managing, marketing, billing and charging for any kind of professional service. Consequently, apart from selling banking and insurance services, many banks have also begun purveying telecommunications, travel, and utility-oriented services.

Other sectors that stand to benefit from the mobile Internet are utilities and safety and social services. For instance, via the mobile Internet, service technicians can obtain descriptions and instructions while they are in the field. In a like manner, fire brigades, police, and ambulance personnel can retrieve information on buildings (for example, blueprints), dangerous substances, criminal records, and medical conditions. Social services personnel can receive assignments and send and retrieve data on patient conditions.

The entertainment industry is espying new opportunities to sell music, books, films, and games. For instance, the industry has plans to sell entertainment to travelers who are waiting at terminals and while they sit "captivated" en route to their destinations.

In Southeast Asia, mobile gambling is being seen as a major business opportunity. When combined with secure payment methods, the mobile Internet will enable users to bet on horses, lotto, cricket, or whatever they like wherever they are.

Yet another interesting use of the mobile Internet is machine-to-machine communication—also referred to as telemetry—of which there are two obvious applications:

- More efficient service and maintenance of field equipment—for example, Coca-Cola has begun to deploy machines that can communicate their current stock of goods, which makes maintaining the total field inventory much more cost-effective.

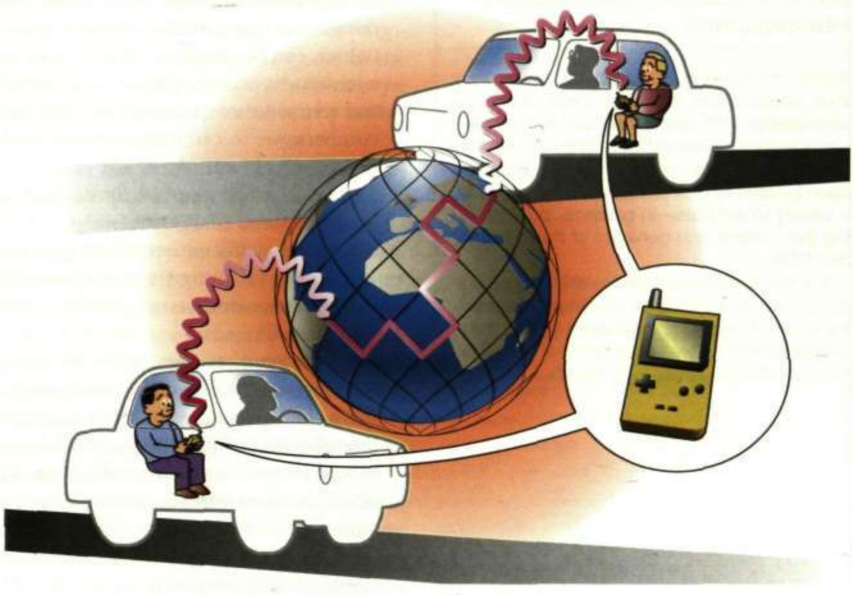

Figure 2
Two users in different parts of the world share a game. Could this mean that parents will finally see their children happily occupied and content during extended road trips or while on vacation?

- The ability to communicate with any device, anywhere and at any time:
  - The largest elevator service firm in Europe has plans to install GPRS terminals in all its elevators.
  - Weather stations and security systems are being outfitted for mobile service.
  - The combination of positioning systems (such as GPS) and GPRS makes it easier to track and maintain equipment, devices, and other valuable goods.[1]

Availability of technology

Unless the devices and applications associated with the mobile Internet are simple, reliable, easy to use, and offer reasonable power consumption, then few people will be inclined to use them:

- Devices and applications must be intuitive and simple to use.
- Services must be easy to personalize; that is, users should only have to select the applications they want to have on their terminals—they should not have to download, unzip and install files if they do not want to.

BOX C, HANDSETS TO MATCH USERS' PERSONALITIES

Today, many users buy products that match their personalities. Many of them also buy accessories with which they can further personalize the exterior of their phones, PDAs, and so on. Taking this one step further, users will soon be able to pick and choose from among a variety of services—in essence, personalizing the content and behavior of their phones and PDAs.

For some, the mobile device will function like a phone or as a device for sending and retrieving e-mail. For others, it will be an instant messaging device as well as a terminal on which to play network games with friends.

- Mobile terminals must be as rugged and reliable as present-day mobile phones (which can be dropped on the floor and withstand a certain amount of moisture, and temperature extremes) but not more cumbersome to carry than keys, a wallet or a watch.
- Terminals must operate for at least 48 hours between charges.

Many device manufacturers claim to be well on the way to fulfilling these requirements, and numerous indicators confirm their claims (Box C):

- Soon the number of portable computers will exceed that of desktop computers.
- Palm Computing, the most successful manufacturer of portable handheld organizers, recently introduced the Palm VII, which features wireless access. The company also successfully breached the "leading-edge" segment and won the hearts of early adopters.
- Although it is heavy and cumbersome compared to most mobile phones, the Nokia Communicator is used by many workforces who need to communicate with the home office from a mobile work environment.
- At Telecom 99, Microsoft demonstrated that, using its Windows CE operating system, it can connect to the Internet via a GSM phone.
- Bluetooth simplifies connections between terminals (communicating devices) and palmtop devices (computing devices).
- Many mobile phone manufacturers exhibiting at Telecom 99 demonstrated the potential of the wireless application protocol (WAP).

Pioneering end-users

Today, virtually every new GSM phone comes ready to handle packet data—SMS works right out of the box. Unfortunately, because the keypad on most terminals is not user-friendly for inputting text, the majority of users do not use SMS (Figure 3). Notwithstanding, Telia (Sweden) recently reported that it derives some 5% to 10% of its revenue from SMS.

Given the deep penetration of GSM in some countries, it is now more the rule than the exception for teenagers to own a mobile phone. Teenagers use SMS to communicate in the classroom, to communicate at nightclubs (where the music is too loud to permit normal conversation), for games, and to communicate with their partners.

The younger generation makes up a market

Figure 3
Ericsson's Chatboard.

segment that has the competence and savvy to use the mobile Internet. Moreover, this group is willing to put up with the initial pitfalls and shortcomings of mobile data, and it has buying power. Finally, this group of consumers is not bound to older paradigms that blind it from seeing new uses and solutions for making life simpler.

Applications

Today, most debates on wireless data center around the wireless application protocol, which is a suite of protocols that will aid applications in wireless environments, including new wireless technologies (GPRS, WCDMA, EDGE, cdma2000). Thanks to a superior, aided user interface and transmission over the air, the wireless application protocol successfully bridges the difficulties that arise when data services are used over wireless systems. However, WAP alone will not usher in the mobile Internet. Some applications simply cannot run over WAP, but require more of the client—more than a Web-based interface can provide. Similarly, some devices, such as laptop computers that run on Windows95/98, do not support WAP.

In all likelihood, many applications will be developed in a thin-client version (in WAP for small devices) and in a more advanced version that targets a specific environment (or perhaps in Java, to enable the application to run in a variety of systems).

Despite the wireless data success of I-mode and Mobitex, and a growing momentum for WAP applications, to date no GPRS-specific applications have been developed (although most IP applications run "okay" over GPRS). Several industry segments are monitoring this market closely, however, and some have already begun investing in it. In all likelihood, GPRS will be a significant enabler of WAP, by making it available to the masses.

The GPRS Applications Alliance

Ericsson has taken the initiative of driving the development of GPRS applications and services by setting up the GPRS Applications Alliance (GAA). As the leader in GPRS, with over 45 trials and deployment contracts, Ericsson is ideally positioned to drive the applications-development process forward.

The new applications and services needed to propel the GPRS market will be developed by new entrants on the mobile communication scene, such as software developers and content providers. The role of the GAA is to serve as a catalyst, by raising the awareness of opportunities offered by the introduction of GPRS.

Ericsson has established two GAA centers: one in Stockholm and one in Silicon Valley. Additional GPRS development centers—up to thirty in all—will be established at various sites around the world during the next few months. Some of these will be set up by Ericsson; others by partner companies who will provide local or regional support. Each center will offer the same training, development, and testing competencies as the centers in Stockholm and in Silicon Valley.

The GAA centers will support application developers at every step. Technical training will enable developers to assess the impact of GPRS on their current applications, the majority of which will have been developed for the Internet in the fixed environment. The centers will also help developers to create business strategies that address the world of GPRS. Technical support will provide the answers to questions such as "How will my desktop PC application work in a vehicle that is traveling at 100 kilometers per hour?"

One of the most important services provided by the GAA centers is application testing. After a developer has created an application, it must be thoroughly tested to ensure that it will operate effectively in GPRS environments. To this end, Ericsson has created the GPRS applications test environment (GATE). Running on a laptop computer, GATE simulates the impact that varying radio environments have on the IP stream, and thus on the application. Different scenarios can be simulated: for example, city centers with high levels of traffic, good coverage and competition for bandwidth as well as rural areas with poor coverage and little traffic but high bandwidth availability. The tests can be repeated as many times as necessary under identical conditions.

To further assist application developers, the GAA also offers a process for verifying products before their commercial introduction. If the GAA determines that an application meets all stipulated criteria, then the Alliance will work with the developer to market his or her product. The GAA is also setting up (at its first two centers) a radio test bed, which is essentially a live GPRS network that provides the final test of a new

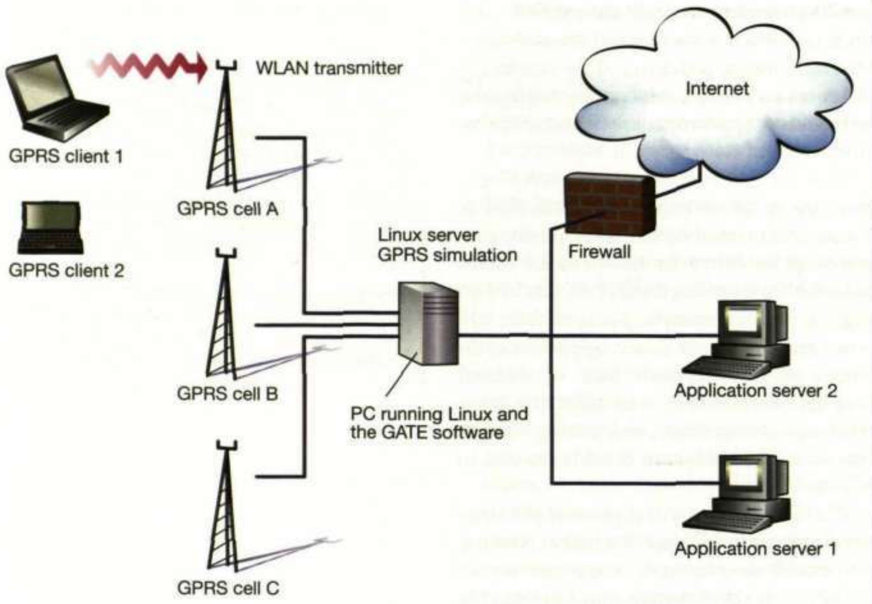

Figure 4
The GATE setup used for demonstrations.

TRADEMARKS

Chatboard™ is a trademark owned by Telefonaktiebolaget LM Ericsson, Sweden.
Coca-Cola® is a registered trademark of the Coca-Cola Company.
Palm VII™ is a trademark of Palm Computing.
Windows CE® is a registered trademark of Microsoft Corporation.
Bluetooth™ is a trademark owned by Telefonaktiebolaget LM Ericsson, Sweden.
Java™ is a trademark owned by Sun Microsystems Inc. in the United States and other countries.
The trademark MOBITEX™ is owned by Telia AB and licensed by Telefonaktiebolaget L M Ericsson, Sweden.

application. After new applications have been tested and verified they are listed on the GAA website (www.gprsworld.com).

Since it was announced last November, the Alliance has attracted over 80 members (the charter members of the GPRS Applications Alliance are Ericsson, IBM, Lotus, Oracle, Palm Computing, and Symbian). The GAA has already begun working with several major software and application-development partners.

The creation of a successful applications environment for GPRS is an essential precursor to the evolution of UMTS. That is, the success of UMTS will be ensured through the introduction of GPRS applications and services. Ericsson's GPRS Applications Alliance initiative serves as a catalyst for creating a world of packet-switched applications and services.

The impact of technology

Experiences from GATE

The GPRS applications testing environment is a Linux-based simulator. Its name, however, is somewhat deceptive, since GATE also aims to support EDGE and UMTS. GATE filters packets at the IP level and drops or delays packets according to the mobile network models that have been assigned to it. The objective is not to simulate every aspect of a full-fledged network, but to give a valid view of the mobile link, as perceived by the application. The GATE computer is connected to a local area network (and, optionally, to the Internet) that hosts various application servers. GATE's other network interface is connected to an emulated GPRS mobile terminal—a laptop computer or another IP-enabled device, such as the Ericsson MC218. This connection (Figure 4) can either be

- a wireless LAN—to demonstrate mobility and handover; or
- a wired Ethernet—if the intent is to test a specific application.

The testing procedure

When a company opts to join the GAA, the GAA technical team contacts the member technical team to discuss the applications to be tested and the technical details of the test. Early on, each party must understand what can and cannot be accomplished. Once all the details have been sorted out, the member team visits one of the GAA test labs. During phase-one testing (usually one day per application), the application is installed in the GAA lab and each party gets acquainted with the setup. When the application works as it was intended to do (most applications run out-of-the-box but with varying quality), some manual stress tests are applied through the GATE interface—for instance, simulated changes in network load, and out-of-coverage simulations. The technical teams then draft a suitable roadmap for future tests. The member usually submits a test specification for the application, which shows how the developer intended the application to work. The rigorous tests run the application through a series of repeatable scenarios. The resulting test protocol serves as the basis for improving the application. The intent is to make testing a cooperative effort between the member and the GAA: the GAA provides the testing facilities and wireless expertise, and the member provides expertise in software development (and specific knowledge of the application).

An often-heard comment at the end of the testing phase is: "This knowledge can be applied to all our products, not just GPRS or wireless products!" Indeed, all network or

Internet products should have efficient protocols that are fast and do not send excessive data. Accordingly, the process of optimizing a product and of making it more robust also prepares it for the conditions of bearers of greater bandwidth. This lowers the cost for users in systems that charge for bandwidth (which might be the case for WCDMA/EDGE/cdma2000) and gives better user experience.

Interruptions

Which of us has not had his or her e-mail software freeze up unexpectedly, for instance while downloading an attachment? Why does this happen? Are these incidents attributable to network failure, conflicts with the operating system, or a lack of system resources? Interruptions represent a complicated challenge, since—apart from the many ways radio conditions can affect an application—numerous other factors can prevent an application from reaching the server. Some technical solutions are available, however. And apart from these, much can be done to improve perceived performance by keeping the user informed and putting him in charge (Box D).

Interruptions become even more apparent in wireless environments, where varying delays and the temporary loss of connection occur more frequently than in fixed-wire environments. Smart, user-friendly software ensures that the end-user stays informed about what is happening and feels confident that he or she will not be forced to wait indefinitely or abort. A meter that indicates the progress of a download, or a dialog box that suggests options for troubleshooting a problem could make the difference between a satisfied end-user and one who ceases to use the software altogether.

BOX D, INTERRUPTION SCENARIO

Graham P. Richardson is a typical wireless data user. One sunny day in November, 2000, he and his wife depart on a well-deserved vacation from their home in Silicon Valley. Their destination is Las Vegas. Judy, Graham's wife, drives as Graham finishes up some work on his laptop computer. Graham's laptop is connected to the Pacific Bell GPRS network—consequently, he can connect wirelessly to the corporate LAN.

This morning, Graham needs to report the latest sales figures and synchronize his files with those on the server. But as the car approaches the Nevada desert, the coverage becomes very weak and is finally non-existent. The GPRS modem loses the connection and the application freezes in wait for a response from the server. Frustrated, Graham cannot stop the task. When he comes back into range, the application does not understand that it can reconnect to the server. In the end, Graham has to reboot his computer and much of his input to the application is lost.

The events we have described in this scenario can and will happen unless developers have the properties of the wireless environment in mind when they design their applications.

Conclusion

Will the introduction of the mobile Internet give rise to a new paradigm—one in which the user is always connected, always online? We think so. Several indicators signal that the day of the mobile Internet is at the door:

- We are seeing simple but successful implementations of the mobile Internet—for instance, I-mode in Japan, and the Mobitex network in North America.
- We see a huge potential market where the subscriber to mobile telephony communications and the Internet subscriber are one and the same person.
- Mass-market pricing can be offered at cost levels optimized for GSM operators.
- Many user communities and content providers are already investing—"gearing up" for the mobile Internet.
- The younger generation, which already uses SMS and the Internet, constitutes a leading-edge segment that has the buying power to kick-start the mobile Internet.
- Mobile terminal and communicator manufacturers are fully ready to provide handsets for the mobile Internet.

The pieces of the puzzle are falling into place, opening the way for the mobile Internet. Being "always connected, always online," users will, simply by clicking or tapping, be able to manage their business and private affairs at any time and from any place.

REFERENCES

1 Swedberg, G.: Ericsson's mobile location solution. Ericsson Review Vol. 76 (1999):4, pp. 214-221.
2 Haartsen J.: Bluetooth—The universal radio interface for ad hoc, wireless connectivity. Ericsson Review Vol. 75 (1998):3, pp. 110-117.
3 Arfwedson, H. and Sneddon, R.: Ericsson's Bluetooth modules. Ericsson Review Vol. 76 (1999):4, pp. 198-205.

Ericsson's mobile location solution

Göran Swedberg

We live in a changing world. Landscapes and skylines shift as cities expand and new buildings are erected. People are also changing: over time we tend to dress differently, modify our habits, find new interests, and refine our views on music, politics, and so on. On the face of it, everything is in a state of change. Moreover, the pace at which change occurs seems to be picking up. Under the surface, however, many fundamental constants remain, oblivious to external change—the changing world of fashion preys on human vanity (a constant); and the recent shift from wired to wireless or mobile communication is indicative of mankind's as yet unfulfilled need to communicate (another constant).

In other words, although our physical environment or our outward expression and behavior fluctuate, our innermost needs—which are the drivers of change—do not. This is why Homer's *The Iliad* and *The Odyssey*, Shakespeare's *Hamlet*, and Strindberg's *Miss Julie* remain timeless and poignant: these works speak to our innermost (and oftentimes hidden) needs. They relate to values, such as love, respect, comfort, and so forth. No matter how much we seem to have changed, it can be argued that our basic needs do not. How fortunate we are to live in an age when technology is maturing to the point that it can offer services that fulfill these basic needs. The mobile location solution is just the sort of breakthrough that addresses needs through personal, customized services.

In this article, the author describes some of the driving forces behind mobile location solutions and introduces the main service categories and their target groups. He then gives an overview of Ericsson's technical solution, which allows operators to achieve immediate and 100% market penetration.

Figure 1
The mobile location solution has been designed to handle a variety of positioning methods and application interfaces.

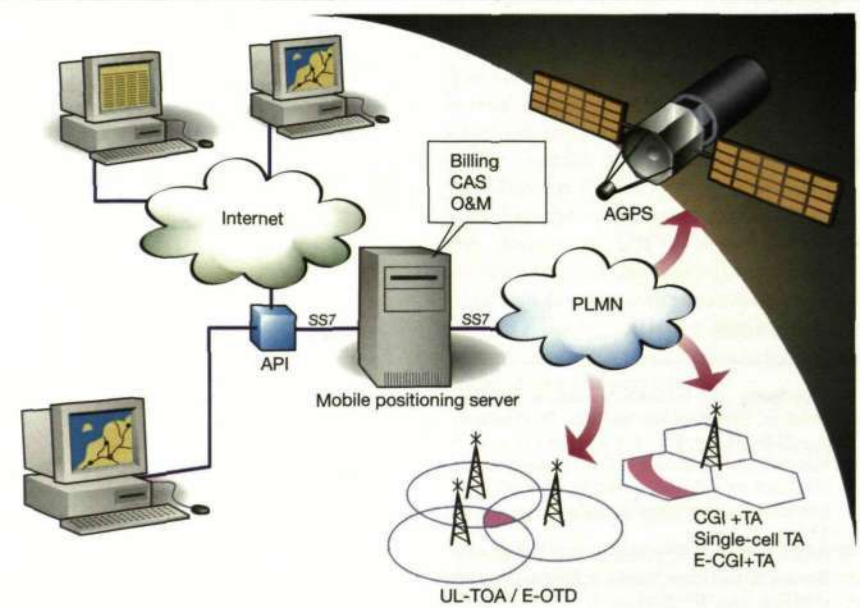

Background

Mobile location solution (MLS) is Ericsson's name for a location system, including applications, that determines the geographic position of mobile subscribers and provides them with relational information and services (Figure 1).

In the early days of telecommunications, we called to a specific location in the hope that the party with whom we wanted to speak would come to the phone. With mobile telephony, however, we know to whom we are calling, but not necessarily his or her whereabouts. Other recent developments, such as the growth and development of information technology and databases (that is, the Internet and associated portals), have introduced completely new prerequisites for the use of information technology. Looking ahead, by combining positional mechanisms with location-specific information, we can offer truly customized personal communication services through the mobile phone or other mobile devices.

Driving forces

Legal aspects

In the USA, legal aspects have acted as the driving force behind GSM positioning standardization. The main player in this market has been the Federal Communications Commission (FCC). Organizations that represent fire brigades, hospitals and other emergency centers have also participated in formulating the FCC requirements. Authorities and industry have agreed on the pace and standards for positioning systems that will serve the market. The requirements have not been completely finalized, but an outline and characteristics have been specified: the FCC regulations differentiate between terminal-based (handset-based) and network-based solutions. A terminal-based positioning solution (Box B) relates to positioning intelligence that is stored in the terminal or its SIM card. These kinds of positional mechanism require a new terminal, a new SIM card, or both. In practice, this means that once the system has been installed, subscribers will have to replace their handsets or SIM cards to benefit from it. Market penetration will increase gradually as handsets and SIM cards are replaced over a period of, say, four to five years. Examples of terminal-based solutions are the network-assisted global positioning system (A-GPS), SIM toolkit, and enhanced observed time difference (E-OTD).

By contrast, network-based positioning solutions (Box C) do not require positioning intelligence to be built into the handset (mobile terminal), which means that market penetration is 100% from the day the system or service is launched. Examples of network-based solutions include the cell global identity and timing advance (CGI+TA) and uplink time of arrival (UL-TOA) methods. Because terminal- and network-based positioning systems have different characteristics, the FCC has stipulated separate requirements for each. Today, the only methods that satisfy these requirements are the UL-TOA solution, which is representative of a network-based system, and the A-GPS solution, which is representative of a terminal-based system.

Commercial aspects

Outside the USA, the development of positioning systems is mainly driven by commercial considerations. These considerations, however, are every bit as strong as the legal aspects that drive the development in the USA. There are three main commercial reasons why operators would invest in positioning services:

- Differentiation—by adding positioning capabilities, operators can offer their subscribers new and attractive services. Operators who do so can compete from a more favorable strategic position.
- Reduced costs—operators who introduce positioning systems can optimize their networks to trace unsuccessful calls. With this information, they can adapt their networks (without waste or overdimensioning) to match calling patterns.
- Increased revenues—the potential of commercial services that use positioning information is truly infinite. What is more, professional and private subscribers are willing to pay for these services.

Location applications

Service categories

Location-based services are categorized by type of application.

Information services

Information services make use of an information bank where information is filtered according to the relative position of a user and the applications he or she has selected. Examples of information services include

BOX A, ABBREVIATIONS

A-GPS	Assisted GPS
BTS	Base transceiver station
CGI	Cell global identity
E-OTD	Enhanced OTD
FCC	Federal Communications Commission
GPS	Global positioning system
GSM	Global system for mobile communication
GUI	Graphical user interface
LMU	Location measurement unit
LCS-client	Location service client
MLS	Mobile location solution
MPC	Mobile positioning center
O&M	Operation and maintenance
OTD	Observed time difference
PLMN	Public land mobile network
RTD	Real time difference
SIM	Subscriber identity module
TA	Timing advance
TDOA	Time difference of arrival
TOA	Time of arrival
UL-TOA	Uplink TOA
WAP	Wireless application protocol

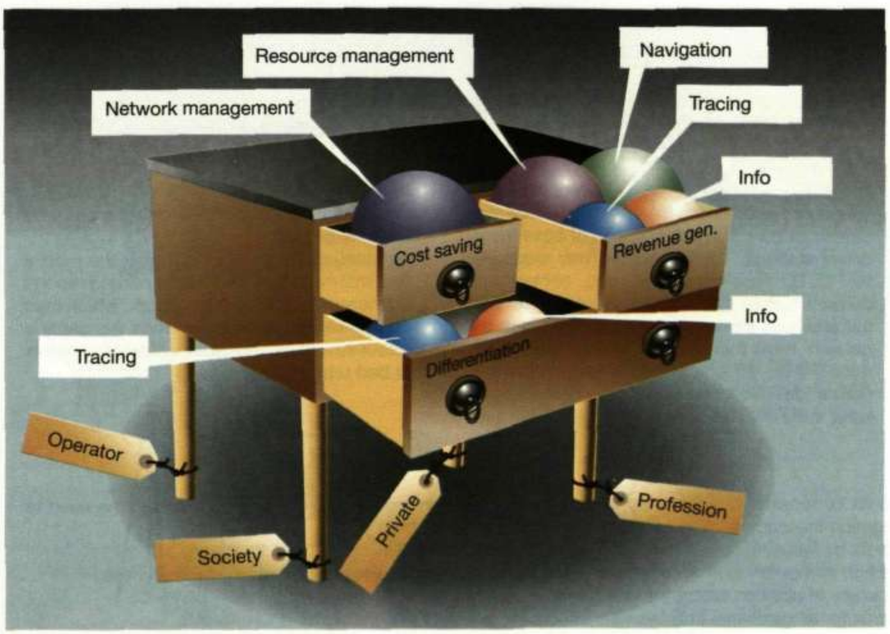

Figure 2
The chest of drawers illustrates how different applications can be grouped strategically for use by their beneficiaries.

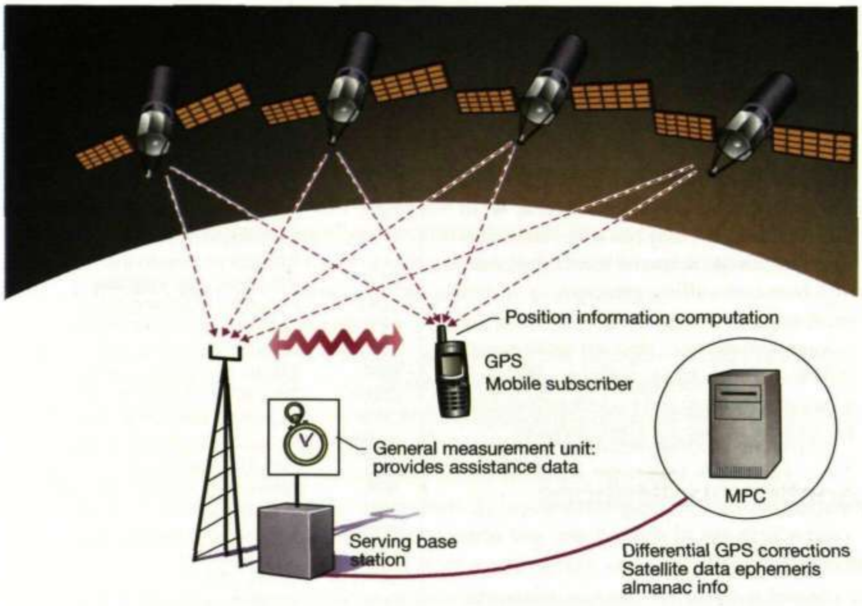

Figure 3
Compared to traditional GPS systems, network-assisted GPS (A-GPS) is a completely different product with very attractive characteristics.

location-based yellow pages, events, and attractions (for example, "What is happening today in town near here?")

Tracing services

Different services can use location-based information to trace mobile terminals (devices), to provide safety, to prevent theft, to facilitate delivery, and so on. Examples of this kind of service include the tracing of a stolen car, helping paramedics to locate persons quickly in an emergency situation, and giving a towing service or automobile repair shop the location of a motorist in need (out of gas, flat tire, dead battery).

Resource management

Resource-management applications are used to manage the logistics of vehicle fleets, freight, and service staff—for instance, repairmen with different skills and qualifications. Examples of resource management include taxi fleet management, the administration of container goods, and the assignment and grouping of railway repairmen.

Navigation

Navigation applications are used to inform subscribers how they can best move from point A to point B. Applications of this kind can be adapted to vehicle or pedestrian navigation.

BOX B, TERMINAL-BASED SOLUTIONS

A-GPS

The global positioning system (GPS) is commonly used for navigation purposes. Stand-alone GPS units are frequently found in cars, boats and airplanes. A GPS unit receives signals from four or more satellites. Each signal contains a time stamp and a description of the position of the satellite. By comparing this information, the GPS unit can calculate its own position. The main drawback of GPS is that satellite signals are relatively weak and may not always provide adequate coverage to *all* environments. However, the GSM network can provide assistance information that gives integrated GPS receivers better coverage than stand-alone GPS receivers (Table 1).

Different kinds of location measurement units (LMU) are used to collect assistance data. In order to provide satellite ephemeris and differential GPS correction, one LMU must be deployed every 300 km in the network. This enhancement provides accuracy within 10 or 20 m. To further increase the coverage of GPS (limited indoor coverage), a highly accurate time reference must be provided. However, this requires the deployment of one LMU in approximately every third BTS.

E-OTD

The enhanced observed time difference (E-OTD) method is based on the measured OTD between arrivals of bursts of nearby pairs of base transceiver stations. The mobile terminal measures the OTD. Synchronization and normal and dummy bursts can be measured. Since BTS transmission frames are not synchronized, the network must measure the relative time difference (RTD). To obtain accurate triangulation, OTD and RTD measurements are needed for at least three distinct pairs of geographically separate base station transceivers. Based on the measured OTD values, the location of the mobile terminal can be calculated in the network or by the mobile terminal itself, provided it has all the necessary information. The E-OTD method can be either:

- network-assisted, in which case the mobile terminal measures the OTD signal and computes its own location (to do so, the network must provide the terminal with additional information, such as BTS coordinates and RTD values); or
- handset-assisted, in which case the mobile terminal measures the OTD signal and reports its measurements to the network, which then computes the terminal's location. Accuracy is about 60 meters in rural areas and 200 meters in bad urban areas.

Table 1

Type of assistance	**Benefit**
Satellite ephemeris	Improves time-to-fix or sensitivity, or both, by focusing acquisition. Improves time-to-fix by eliminating the need to demodulate navigation messages.
Frequency accuracy	Improves time-to-fix by focusing acquisition.
Location estimate	Initializes the position computation procedure. Improves the acquisition of second and subsequent signals.
Differential GPS correction	Improves the accuracy of position estimates (10-20 m).
Time reference	Improves time-to-fix for all receivers. Improves sensitivity for receivers in poor signal environments.

BOX C, NETWORK-BASED METHODS

CGI+TA, enhanced CGI+TA

The single-cell timing-advance positioning method uses the cell global identity (CGI) and the timing advance (TA) parameter to determine the location of mobile terminals. The CGI identifies the cell in which the mobile terminal is located. A cell can be a circular or triangular sector. The TA parameter is an estimate of the distance (in increments of 550 m) from the mobile terminal to the base station. The measurement is based on the access delay between the beginning of a time slot and the arrival of bursts from the mobile terminal. The access delay is proportional to the distance between the base transceiver station and the terminal. The accuracy of this method varies according to the size of the cell. The radius of a cell may vary from 100 meters to 35 km (CGI). The width of an arc is 550 meters.

As seen in Figure 5, the CGI-TA positioning method is merely an enhancement of the CGI. The estimated location is reported in terms of longitude, latitude, and an uncertainty shape within which the mobile terminal is located.

Researchers are currently investigating ways of improving the accuracy of the CGI+TA method; this enhanced technique uses the CGI+TA method as a base in its positioning determination. Test results indicate a granularity of between 100 and 200 meters.

Uplink time of arrival (UL-TOA)

The uplink TOA positioning method is based on measuring the time of arrival of a signal from a mobile terminal to four or more measurement units. Ideally, the signal is a training sequence of a random access burst but it could also be a normal burst. The UL-TOA positioning method works with all existing mobile terminals—that is, no handsets need to be modified. Location measurement units (LMU) located at the base stations receive the bursts and measure the value of the uplink time of arrival (UL-TOA). The mobile position center (MPC) calculates the time difference of arrival (TDOA) by subtracting pairs of UL-TOA values.

Prerequisites for calculating position by means of the UL-TOA positioning method are as follows:

- The geographical coordinates of the measurement units are known.
- The timing offset between the measurement units is known—for instance, by the use of absolute GPS time at the measurement units, or by using reference measurement units (also referred to as reference terminals) to determine the real time difference (RTD).

The MPC delivers a position estimate and an uncertainty estimate to the application. The accuracy of this method varies according to the environment and the number of location measurement units employed. Accuracy typically varies between 50 (rural) and 150 meters (bad urban).

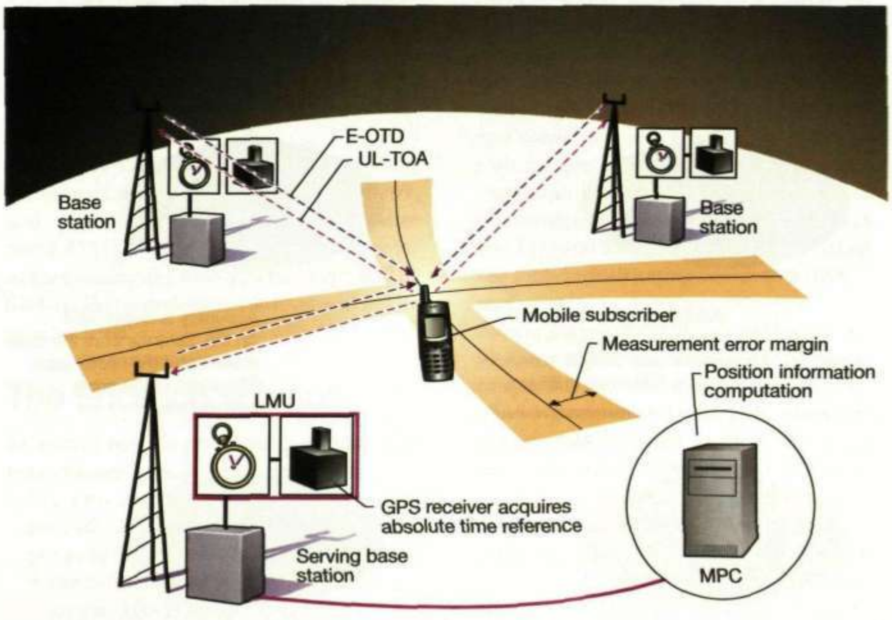

Figure 4
Although they measure in different directions, the UL-TOA and E-OTD methods each use the triangulation of time difference between base stations and the terminal to determine positions.

Other services

Some positioning applications, such as network planning, map services and telematics, do not fall into the above categories. Examples of these services include hot-spot tracing ("How is my network used?"), and location-based charging.

Target groups

It is interesting to note that the characteristics of location-based services apply equally well to commercially driven and legally driven developments—their application categories address some of mankind's most basic needs, including safety, comfort, and the need to communicate. Moreover, because these needs are so basic, they apply to all four target groups.

- Society's (government) interest in location-based services stems from its role as a provider of safe and equitable environments for its citizens.
- Operators want to introduce location-based services into their networks because they create revenue, cut costs, and project a positive business image to the market.
- Professional end-users are attracted by the efficiency with which location-based services enable them to manage their resources. These services also function as marketing amplifiers, since location-related information is, by nature, already very specific and can readily be adapted to specific customer groups. Some professional end-users resell location-based services directly to the end-user.
- Private end-users want location-based information and services because they tend to heighten the user's overall sense of comfort and well-being—truly personalized services match customers' requests according to their profile, location, and time of day.

About positioning systems

System principles

There are two kinds of positioning system:

- overlay systems; and
- integrated systems.

Overlay systems are built on top of an existing network system. Since they use few core network features, these systems can be implemented in mixed networks. However, because overlay systems will not be standardized, they are neither future-proof nor does their design meet legal requirements. In summary, the market window for overlay systems is very short—this window will close when the standardization work has been completed.

Once the standards have been finalized, it

Figure 5
The CGI+TA solution (which offers 100% market penetration) can serve as a fall-back method, provided the preferred positioning method fails to determine the position of the mobile terminal.

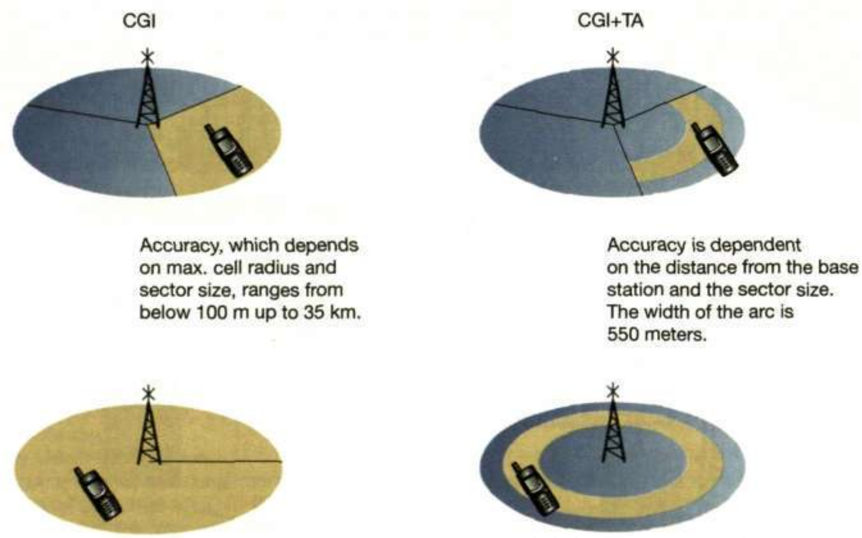

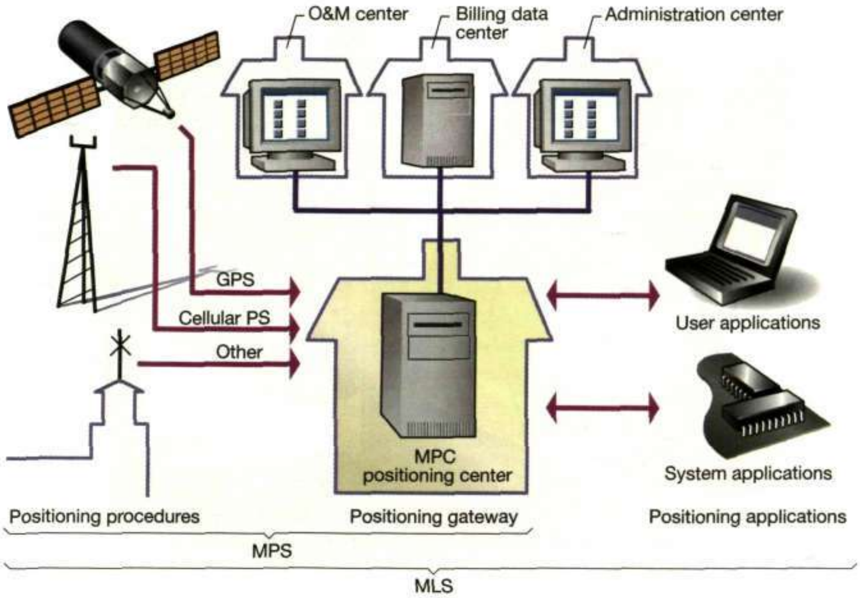

Figure 6
Overview of Ericsson's mobile location solution (MLS). A mobile positioning system has been logically integrated into the network.

will be possible to launch any compliant network-based solution in any network. What is more, the network-based solutions have been designed

- to meet legal and commercial requirements; and
- to follow the GSM evolutionary roadmap toward third-generation systems.

Standardization

The CGI+TA and UL-TOA positioning techniques were standardized in May 1999, and plans have been made to standardize the E-OTD and A-GPS positioning techniques during the first quarter of 2000. The SIM-toolkit positioning solution will not be standardized.

The Ericsson solution

Ericsson's mobile positioning system consists of three logical subsystems.

1. The positioning subsystem can use a variety of techniques to determine and supply geographic coordinates:
   - the cellular positioning system; for example, UL-TOA and E-OTD;
   - the network-assisted global positioning system (A-GPS);
   - other techniques, such as the SIM tool-kit.
2. The positioning gateway subsystem (MPC), which functions as a mediation device between the public land mobile network (PLMN) and the location service client (LCS-client), retrieves data from positioning subsystems (such as the E-OTD) and converts it into positioning information for the LCS-client. It also provides the operator with a GUI (based on a Java platform) for administering the node—in the Ericsson solution, this node is called the mobile positioning center tool (MPC-tool).
   Note: because the MPC is a mediation device, it monitors and can register the usage of specific location applications, thereby allowing operators to charge for them via applications installed in the MPC or via interfaces to the billing system. This is not the case when traditional global positioning systems are used. In those systems, the network solely serves as a link for transporting positioning data between the terminal and the central application-

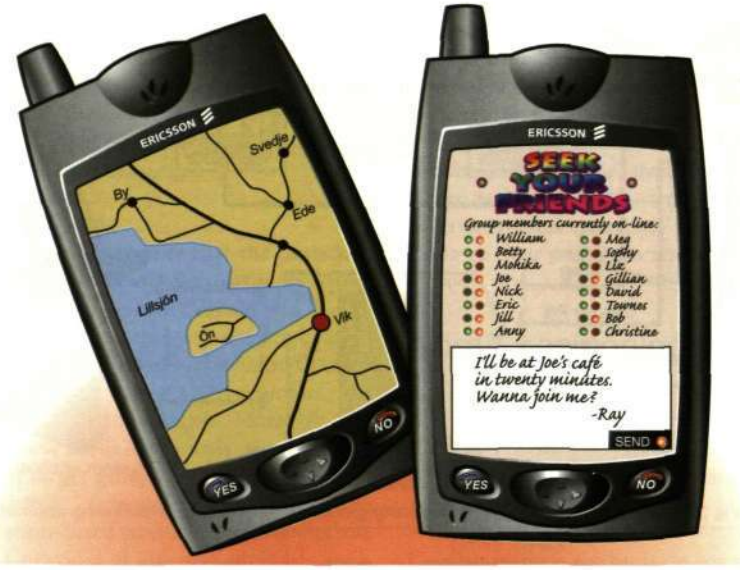

Figure 7
New generations of terminals will provide an unprecedented quality of information in an easy-to-use, "pick-and-choose" fashion.

management site. Therefore, operators or service providers cannot charge specifically for their service.

3. The LCS-client, which is a subsystem of the MLS, contains applications that make use of positioning information. Internal applications (for instance, emergency calls) are coded into the GSM system according to the GSM standard. External applications are supplied to the system by system vendors, the operator, and third-party application developers. Some applications might even be "interactively customized" by the end-user.

Market impact

The demand for services

Can we—by determining the driving forces behind mobile location solutions, the beneficiaries of their services, and the probable number of service categories—judge the extent to which mobile positioning systems will succeed? No. These exercises merely suggest which needs the systems and their services might fulfill. The market impact of MLS services is determined by the end-users' willingness to pay for them.

By means of the Internet, WAP technology, portals, and so on, we are being inundated with information. The challenge at hand is to filter this information to make it practical for use by the man or woman in the street. Modern society demands instantaneous service that has been tailored to fit individual or group needs. The service must take into account the end-user, the group to which he or she belongs, and where he or she is located at any given moment. These are, in fact, the very things that Ericsson's mobile location solution does: it selects information and, with the help of WAP technology, serves it up to the end-user through an easy-to-use, interactive interface (Figure 7).

Market potential

In a report issued by OVUM Ltd., the global market value of MLS between 1999 and 2005 is expected to reach USD 25 billion. Today, there are some 200 million GSM subscribers, and this number (including next-generation GSM) may well reach one billion by 2004. Using this information, let us project a simple forecast of our own: A basic directory service costs USD 0.50 per request. Assuming that 400 million subscribers will use the service 20 times a year, the service will generate an annual revenue

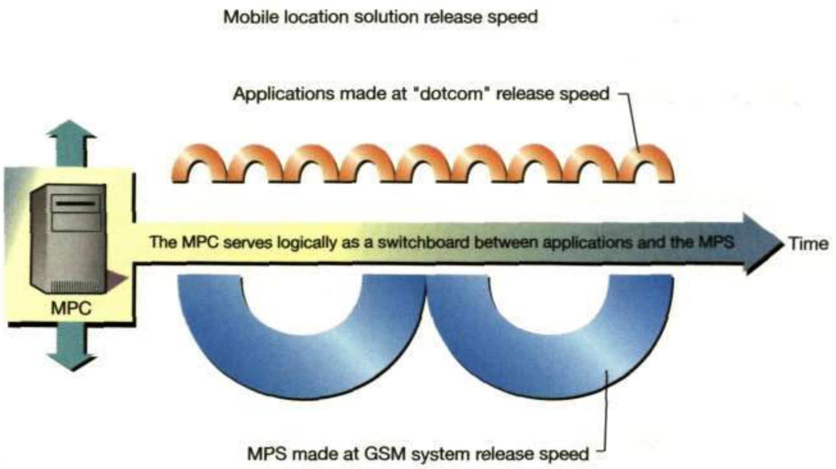

Figure 8
The design of Ericsson's mobile location solution allows developers to create applications independently of the mobile positioning system. This ensures system reliability and guarantees that services can be customized rapidly.

of USD 4 billion. Thus, would not a more advanced yellow-pages service have the potential to generate even more revenue?

Ericsson's strategy

Given the tremendous potential of MLS, both for consumers and the industry, vendors must ensure that their products meet market demand in terms of matching services, timing (time to market), quality, and long-term compatibility. To determine what these demands are, Ericsson has taken an active, leading stance toward MLS, by pushing the standardization work ahead, and by becoming the first vendor of a cellular positioning system.

Maximum revenues are achieved by rapidly changing the product line-up that meets the demands of short-lived market windows. This is achieved through a ready-to-go packaging of standard services and by allowing service providers or third-party service vendors to customize and develop new services. To this end, Ericsson's MLS employs alternative location mechanisms that give subscribers a range of positioning techniques from which to choose, as well as a best-effort technique (provided the requested positioning technique is unavailable). The system is easy to manage and includes "top-notch" billing capabilities. A service portfolio ensures that subscribers receive the right services at the right time (Figure 9).

Conclusion

Mobile location solutions, which combine mobile location systems and location-specific services, represent the next major "killer application" of the mobile industry. Early estimates put the value of the market at USD 25 billion, but given the tremendous demand, these estimates are probably conservative.

The five main service categories for location-based services are information services, tracing services, resource management, navigation, and "other services"—those services that do not fall into one of the first four categories.

The four main target groups of location-based services are government, operators, professional end-users, and private end-users.

Ericsson's mobile location solution has been designed to offer high-end network-assisted (A-GPS) and network-based cell-identity (CGI+TA) positioning techniques, to give operators immediate and 100% market penetration in a future-proof environment.

TRADEMARKS

Previous issues

No. 3, 1999

Third-generation radio access standards
Toward third-generation mobile multimedia communication
Adaptive base-station antenna arrays
Enhancing cellular network capacity with adaptive antennas
Messaging-over-IP—A network for messaging and information services
APZ 212 30—Ericsson's new high-capacity AXE central processor
TelORB—The distributed communications operating system

No. 2, 1999

Cello—An ATM transport and control platform
WCDMA evaluation system—Evaluating the radio access technology of third-generation systems
Tigris—A gateway between circuit-switched and IP networks
GPRS—General packet radio services
Jambala Mobility Gateway—Convergence and inter-system roaming
Professional Services—Meeting the changing needs of network operators

No. 1, 1999

Ericsson's *Pro* products—Adapting mass-market technology to fit specialized needs
Open communication devices using the EPOC operating system
Mobile Advantage Wireless Office—A digital wireless office system for TDMA/136 networks
Edge—Enhanced data rates for GSM and TDMA/136 evolution
Ericsson's *e-box* system—An electronic services enabler

No. 4, 1998

WAP—Wireless application protocol
ADSL Lite—The broadband enabler for the mass market
Competitive broadband access via microwave technology
Integrating spectrum in D-AMPS/IS-136 wireless networks
Wireless LAN systems
GSM on the Net
The SDH interface in AXE

Telefonaktiebolaget L M Ericsson
S-126 25 Stockholm, Sweden
Tele: +46 8 7190000
Fax: +46 8 6812710

ISSN 0014-0171
Ljungföretagen, Örebro 1999