

Ericsson REVIEW

THE TELECOMMUNICATIONS TECHNOLOGY JOURNAL

NO 3, 1998

Bluetooth—The universal radio interface for *ad hoc*, wireless connectivity

Internet directory services with click-to-dial

Jambala—Intelligence beyond digital wireless

ERION—Ericsson optical networking using WDM technology

Access 910 system



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
S-126 25 Stockholm, Sweden

Tel: +46 8 719 00 00

Fax: +46 8 681 27 10

Internet address:

<http://www.ericsson.com/Review>

Publisher: Lars A. Ståhlberg

Editorial board: Hans Alberg, Greger Berg, Tom Boothe, Mikael Bäck, Philippe Charas, Jonas Hermansson, Anders Hidmark, Lena Krogstad, Filip Lindell, Bo Malmberg, Mats Nilsson, Eric Peterson, Göran Rasmussen, Michele Schmidt, Ole Segtnan, Sture Sjöström, Bengt Stavenow, Lars-Gunnar Sundin, Peter Svärd

Editor: Eric Peterson

Production manager: Eva Karlstein

Layout: Paues Media, Stockholm

Printer: Ljungföretagen, Örebro

ISSN: 0014-0171

Volume: 75, 1998

Cover: Bluetooth technology substitutes a universal short-range radio link for the many proprietary cables that are presently used for connecting devices. But more than this, Bluetooth radio technology provides a universal bridge to existing data networks, a peripheral interface, and a mechanism for forming small *ad hoc* groupings of connected devices away from fixed network infrastructures.

© Telefonaktiebolaget LM Ericsson

Contents

Bluetooth—The universal radio interface for *ad hoc*, wireless connectivity

Imagine a cheap, power-efficient radio chip that is small enough to fit inside any electronic device or machine and that provides local connectivity. Bluetooth is a universal radio interface in the 2.45 GHz frequency band that enables portable electronic devices to connect and communicate wirelessly via *ad hoc* networks. **Page 110**

Internet directory services with click-to-dial

Internet directory services make available vast resources of the Internet, helping users to save valuable time in finding names, e-mail addresses, and so forth. Click-to-dial services allow users to invoke calls by clicking on the phone number they retrieve using directory services. Ericsson's solution sets up and transmits calls via the PSTN, but it is equally compatible with VoIP. **Page 118**

Jambala—Intelligence beyond digital wireless

Jambala is the next-generation application platform that operators need to provide new, high-value services in an increasingly segmented end-user market. Jambala provides a unique combination of availability, reliability, scalability and Internet readiness—all using commercially available hardware. **Page 126**

ERION—Ericsson optical networking using WDM technology

Ericsson maintains a simple, pragmatic approach to networking as they further explore and exploit dense WDM technology and optical networking. Ericsson's next-generation transport-network technology—ERION—enables operators to derive maximum benefit from investments in client technology, while simplifying networks and improving traffic protection and routing functionality. **Page 132**

Access 910 system

Ericsson's Access 910 is a general-purpose, access-network system that provides PSTN, Internet, VoIP, ATM, and switched video broadcast capabilities to a wide range of service networks. The support it provides for practical, cost-effective migration from narrowband to broadband services makes it ideal for present-day and future telecom environments. **Page 138**

Contributors

In this issue



Jaap Haartsen



Mats Eriksson



Fred Jones



Magnus Grenfeldt



Michael Begley



Lars-Olof Haster



Sture Roos

Jaap Haartsen has worked for Ericsson Mobile Communications since 1991—two years at Research Triangle Park, USA, and four years in Lund, Sweden. He is currently located in Enschede, Holland, where he works with present-day and future Bluetooth design applications. He holds an MSc and PhD in electrical engineering from the Delft University of Technology, Holland.

Mats Eriksson, who has a background of working for telecom and Internet operators, joined Ericsson in 1997. He currently works as product business manager in the Tele Internet area, which emphasizes services that combine the Internet and telecom networks. He holds an MSc degree in industrial organization and innovation from the University of Halmstad.

Fred Jones is product manager of D-AMPS HLR/AC within the Wireless Intelligent Network Applications & Platforms product unit at Ericsson Research Canada. Since joining Ericsson in 1991, he has also worked with design, system engineering and technical management of the D-AMPS product line. He holds a Master of Engineering degree from Dublin City University, Dublin, Ireland.

Magnus Grenfeldt currently works as product manager of ERION (Ericsson Optical Networking). Since joining Ericsson in 1994, he has also worked as a designer of optical hardware and as a product manager of SDH products. He holds an MSc in materials physics from Uppsala University.

Michael Begley has a Bachelor of Engineering degree in computer engineering from

the University of Limerick, Ireland. During his five years with the System Management Department at Ericsson Systems Expertise Ltd., Atlone, he has worked with OA&M and element management of access nodes.

Lars-Olof Haster, who joined Ericsson in 1979, works with system design of access-network products for AXE. An expert in V5, he participates in ETSI and ITU standardization work. He holds an MSc in applied physics and electrical engineering from the Linköping Institute of Technology.

Sture Roos is presently a senior expert in system architecture hardware. Apart from taking leave of absence for an ITU assignment in Brazil, he has worked steadily for Ericsson since 1967, participating in AXE and digital cross-connect development.

From the editor

Eric Peterson

At the office and at home I work with numerous electronic devices: computers, PDAs, modems, mobile phone, CD-ROM and minidisc players, headphones, microphones, and so on. I admit it. I'm a techno-freak—I like gadgets. But, oh, the cables—that's an entirely different story! Can't someone, please, do something about all the cables? Actually, thanks to Bluetooth technology, my *I-detest-all-these-cables* woes may soon be over, for all time!

The Bluetooth consortium is still quite new—it was formed in February—but already it has gained the support of several big names in the mobile-telephony, laptop-computing and datacom-network industries: Ericsson, 3Com, Casio, Compaq, IBM, Intel, Lucent, Motorola, Nokia, Qualcomm, Samsung, Siemens and Toshiba, to name a few. The reason? Bluetooth is a universal radio interface that enables portable electronic devices simultaneously and without wires to connect to and communicate with up to seven other units via short-range, *ad hoc* networks. Its applications are truly without limit. Mark my words: Bluetooth technology will change the way we interact with machines and other human beings. To learn more about the technology behind Bluetooth, read the article in this issue authored by Jaap Haartsen, the father of Bluetooth!

The ability of electronic devices to communicate wirelessly with each other opens the door to innumerable new device designs as well as applications and services. If you are an end-user, the future looks very bright indeed. But operators and service providers also have reason to cheer as they look to the future, because emerging applications and services are certain to give them additional revenue. It is in this context that Ericsson's next-generation application platform, called Jambala, will have a central role. The Jambala platform represents a unique com-

bination of availability, reliability, scalability and Internet readiness. Its use of Java and CORBA technologies gives operators a platform-independent environment with which to support applications and deliver new services.

Picture an electronic version of the Yellow Pages, White Pages and much, much more, available at your fingertips anytime, anyplace. Ericsson's Internet directory query solution enables operators to provide (and charge for) access to names, phone numbers, e-mail addresses, and other resources in Internet/intranet directories. Operators who combine click-to-dial functionality with Internet directory query enable subscribers to initiate their telephone calls simply by clicking on a name or number—which they retrieved using the IDQ service.

New devices, devices that communicate with other devices, plus more and better services, all mean that operators must have a better means of managing services and related traffic levels. Here, too, Ericsson offers innovative access- and transport-network solutions. The Access 910 system, for example, is a general-purpose, access network that provides PSTN, Internet, VoIP, ATM, and switched video broadcast capabilities over a broad array of customer-network interfaces and to a wide range of service networks. Similarly, Ericsson's next-generation transport-network solution, called ERION, introduces optical subnetworks (dense WDM technology) alongside client-layer technologies. Optical subnetworks—which are used for carrying ATM and IP, SDH, and PDH—simplify network management while improving security and increasing throughput.

As always, I hope you find the topics in this issue of Ericsson Review relevant and timely.

Eric Peterson
Editor



Bluetooth—The universal radio interface for *ad hoc*, wireless connectivity

Jaap Haartsen

Bluetooth is a universal radio interface in the 2.45 GHz frequency band that enables portable electronic devices to connect and communicate wirelessly via short-range, *ad hoc* networks. Each unit can simultaneously communicate with up to seven other units per piconet. Moreover, each unit can simultaneously belong to several piconets.

Bluetooth technology—which apart from Ericsson, has gained the support of Nokia, IBM, Toshiba, Intel and many other manufacturers—eliminates the need for wires, cables and connectors for and between cordless or mobile phones, modems, headsets, PDAs, computers, printers, projectors, local area networks, and so on, and paves the way for new and completely different devices and applications.

Before guiding us through frequency-hopping technology and the channel, packet and physical-link definitions that characterize the Bluetooth air interface, the author briefly describes the conditions that led up to the development of Bluetooth. He then acquaints us with the networking aspects of Bluetooth technology, describing piconets and scatternets, connection procedures, and inter-piconet communication.

Imagine a cheap, power-efficient radio chip that is small enough to fit inside any electronic device or machine, that provides local connectivity, and that creates a (worldwide) micro-scale web. What applications might you use it in?

In 1994, Ericsson Mobile Communications AB in Lund, Sweden, initiated a study to investigate the feasibility of a low-power, low-cost radio interface between mobile phones and their accessories. The intention

was to eliminate cables between phones and PC cards, wireless headsets, and so forth. The study was part of a larger project that investigated multi-communicators connected to the cellular network via cellular telephones. The last link in the connection between a communicator and the cellular network was a short-range radio link to the phone—thus, the link was called the multi-communicator link or MC link. As the MC link project progressed, it became clear that there was no limit to the kinds of application that could use a short-range radio link. Cheap, short-range radios would make wireless communication between portable devices economically feasible.

Current portable devices use infrared links (IrDA) to communicate with each other. Although infrared transceivers are inexpensive, they

- have limited range (typically one to two meters);
- are sensitive to direction and require direct line-of-sight;
- can in principle only be used between two devices.

By contrast, radios have much greater range, can propagate around objects and through various materials, and connect to many devices simultaneously. What is more, radio interfaces do not require user interaction.

In the beginning of 1997, when designers had already begun work on an MC link




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

transceiver chip, Ericsson approached other manufacturers of portable devices to raise interest in the technology—for the system to succeed, a critical mass of portable devices must use the short-range radio. In February 1998, five promoters—Ericsson, Nokia, IBM, Toshiba and Intel—formed a special interest group (SIG). The idea was to achieve a proper mix of business areas: two market leaders in mobile telephony, two market leaders in laptop computing, and a market leader in core, digital-signal-processor (DSP) technology. On May 20 and 21, 1998, the Bluetooth consortium announced itself to the general public from London, England; San Jose, California; and Tokyo, Japan. Since then, several companies have joined as adopters of the technology (Box B).

The purpose of the consortium is to establish a *de facto* standard for the air interface and the software that controls it, thereby ensuring interoperability between devices of different manufacturers. The first products to use MC link technology will emerge at the end of 1999 in mobile phones, notebook computers and accessories (Figure 1).

Box A Abbreviations

ACL	Asynchronous connectionless
ARQ	Automatic retransmission query
CVSD	Continuous variable slope delta
DSP	Digital signal processor
FEC	Forward error correction
FH	Frequency hop
FSK	Frequency shift keying
HEC	Header error correction
HPC	Handheld personal computer
IrDA	Infrared Data Association
ISM	Industrial Scientific Medical
MAC	Media access control
MC	Multicomunicator
PC	Personal computer
PDA	Personal digital assistant
RF	Radio frequency
SCO	Synchronous connection-oriented
SIG	Special interest group
TDD	Time division duplex
TDM	Time division multiplex

Box B The Bluetooth consortium—promoters and adopters

The promoters of the Bluetooth* consortium formed a special interest group (SIG) at Ericsson Inc., Research Triangle Park, North Carolina, on February 4, 1998. The consortium was announced to the public on May 20 and 21, 1998. Many companies have since joined the consortium as adopters of the technology (status as of July 11, 1998):

Ericsson	Promoter	Plantronics
Intel	Promoter	Psion
IBM	Promoter	Puma Technologies
Nokia	Promoter	Quadrige
Toshiba	Promoter	Qualcomm, Inc.
3Com		Samsung Electronics Ltd.
Axis		Siemens Forsvarsystem A/S
BreezeCOM		Symbian
Casio		Symbionics Ltd.
Cambridge consultantsLtd.		T-Span System
CETECOM GmbH		Temic Semiconductor
Cirrus Logic		TDK
Compaq Computer Corp.		TTP Communications Ltd.
Convergence Corporation		Universal Empowering Technologies
Dell Computer Corp.		VLSI Technology, Inc.
InnoLabs Corporation		Xircom
Jeeves Telecom Ltd.		
Lucent Technologies UK Ltd.		
Metricom		
Motorola		
NeoParadigm Labs, Inc.		

* The name, Bluetooth, was taken from Harald Blåtand, a Danish Viking king from the early Middle Ages.




Figure 2
Frequency-hop/time-division-duplex channel.

The Bluetooth air interface

The focus of user scenarios envisioned for first-generation products is typically on traveling business people. Portable devices that contain Bluetooth radios would enable them to leave cables and connectors at home (Box C). Before the air interface for Bluetooth could be designed, however, certain requirements had to be settled:

- The system must operate worldwide.
- The connection must support voice and data—for instance, for multimedia applications.
- The radio transceiver must be small and operate at low power. That is, the radio must fit into small, portable devices, such as mobile phones, headsets and personal digital assistants (PDA).

License-free band

To operate worldwide, the required frequency band must be available globally. Further, it must be license-free and open to any radio system. The only frequency band that satisfies these requirements is at 2.45 GHz—the Industrial-Scientific-Medical (ISM) band, which ranges from 2,400 to 2,483.5 MHz in the US and Europe (only parts of this band are available in France and Spain), and from 2,471 to 2,497 MHz in Japan. Consequently, the system can be used worldwide, given that the radio transceivers cover the frequency band between 2,400 and 2,500 MHz and that they can select the proper segment in this band.

Frequency hopping

Since the ISM band is open to anyone, radio systems operating in this band must cope with several unpredictable sources of interference, such as baby monitors, garage door openers, cordless phones and microwave ovens (the strongest source of interference). Interference can be avoided using an adaptive scheme that finds an unused part of the spectrum, or it can be suppressed by means of spectrum spreading. In the US, radios operating in the 2.45 GHz ISM band are required to apply spectrum-spreading techniques if their transmitted power levels exceed 0 dBm.

Bluetooth radios use frequency-hop (FH) spread spectrum, since this technology better supports low-cost, low-power radio implementations. Frequency-hop systems divide the frequency band into several hop channels. During a connection, radio transceivers hop from one channel to another in a pseudo-random fashion. The instantaneous (hop) bandwidth is small in frequency-hop radios, but spreading is usually obtained over the entire frequency band. This results in low-cost, narrowband transceivers with maximum immunity to interference. Occasionally, interference jams a hop channel, causing faulty reception. When this occurs, error-correction schemes in the link restore bit errors.

Channel definition

Bluetooth channels use a frequency-hop/time-division-duplex (FH/TDD) scheme (Figure 2). The channel is divided into 625 μs intervals—called slots—where a different hop frequency is used for each slot. This gives a nominal hop rate of 1,600 hops per second. One packet can be transmitted

Box C User scenarios

Three-in-one phone—use the same phone everywhere

When you are at the office, your phone functions as an intercom (no telephony charge). At home, it functions as a cordless phone (fixed-line charge). When you are on the move, it functions as a mobile phone (cellular charge).

Internet bridge—surf the Internet regardless of the connection

Use your portable PC to surf the Internet anywhere, regardless of whether you are connected wirelessly through a mobile phone (cellular) or through a wired connection (PSTN, ISDN, LAN, xDSL).

Interactive conference—connect every participant for instant data exchange

In meetings and at conferences, you can share information instantly with other participants. You can also operate a projector remotely without wire connectors.

The ultimate headset—a cordless headset keeps your hands free

Connect a headset to your mobile PC or to any wired connection and free your hands for more important tasks at the office or in your car.

Portable PC speakerphone

Connect cordless headsets to your portable PC and use it as a speakerphone regardless of whether you are in the office, your car, or at home.

Briefcase trick (hidden computing 1)

Access e-mail while your portable PC is still in

the briefcase. When your portable PC receives an e-mail message, you will be notified by your mobile phone. You can also use the phone to browse incoming e-mail and read messages.

Forbidden message (hidden computing 2)

Compose e-mail on your PC while you are on an airplane. When you land and are allowed to switch on your mobile phone, the messages are sent immediately.

Automatic synchronization (hidden computing 3)

Automatically synchronize your desktop computer, portable PC, notebook (PDA or HPC) and mobile phone. As soon as you enter the office, the address list and calendar in your notebook automatically updates the files on your desktop computer or vice versa.

Instant postcard—send photos and video clips instantly from any location

Connect a camera cordlessly to your mobile phone or to any wire-bound connection. Add comments from your mobile phone, a notebook, or portable PC and send them instantly to a recipient anywhere in the world. Suitable for professional and personal use.

Cordless desktop—connect all peripheral tools to your PC or the LAN

Connect your desktop/laptop computer cordlessly to printers, scanners and the LAN. Increase your sense of freedom through cordless mouse and keyboard connections to your PC.

per interval/slot. Subsequent slots are alternately used for transmitting and receiving, which results in a TDD scheme.

Two or more units sharing the same channel form a piconet, where one unit acts as a master, controlling traffic on the piconet, and the other units act as slaves. The frequency-hop channel is determined by the frequency-hop sequence (the order in which hops are visited) and by the phase in this sequence. In Bluetooth, the sequence is determined by the identity of the piconet master and phase is determined by the master unit's system clock (Figure 3). In order to create the master clock in the slave unit, the slave may add an offset to its own native clock. The repetition interval of the frequency-hop sequence, which is very long (more than 23 hours), is determined by the clock. If every participant on a given channel uses the same identity and clock as input to the hop-selection box, then each unit will consistently select the same hop carrier and remain synchronized. Every piconet has a unique set of master parameters which create a unique channel.

The channel makes use of several, equally spaced, 1 MHz hops. With Gaussian-shaped frequency shift keying (FSK) modulation, a symbol rate of 1 Mbit/s can be achieved. In countries where the open band is 80 MHz or broader, 79 hop carriers have been defined. In countries where the band is narrower (Japan, France, and Spain), only 23 hop carriers have been defined (Table 1). On average, the frequency-hop sequence visits each carrier with equal probability.

Packet definition

In each slot, a packet can be exchanged between the master unit and one of the slaves. Packets have a fixed format (Figure 4). Each packet begins with a 72-bit access code that is derived from the master identity and is unique for the channel. Every packet exchanged on the channel is preceded by this access code. Recipients on the piconet compare incoming signals with the access code. If the two do not match, the received packet is not considered valid on the channel and the rest of its contents are ignored. Besides packet identification, the access code is also used for synchronization and compensating for offset. The access code is very robust and resistant to interference. Correlation of the access code by recipients provides similar processing gains as direct-sequence spreading.

A header trails the access code. It contains important control information, such as a




Figure 3
Hop selection scheme: In the selection box, the master identity selects the sequence, and the clock selects the phase. Combined, they give the hop carrier to be used.

Parameters	Values
Modulation	G-FSK, $h \leq 0.35$
Peak data rate	1 Mbit/s
RF bandwidth	220 kHz (-3dB), 1 MHz (-20 dB)
RF band	2.4 GHz, ISM band
RF carriers	23/79
Carrier spacing	1 MHz
Peak TX power	≤ 20 dBm

Table 1
Radio parameters.

three-bit media-access-control (MAC) address, packet type, flow control bits, bits for the automatic-retransmission-query (ARQ) scheme and a header-error-check (HEC) field (Figure 5). The header, whose length is fixed at 54 bits, is protected by a one-third rate forward-error-correction (FEC) code.

Payload may or may not trail the header. The length of the payload may vary from 0 to 2,745 bits.

To support high data rates, multi-slot packets have been defined. A packet can cover one slot, three slots, or five slots. Packets are always sent on a single-hop carrier. For multi-slot packets, the hop carrier is used as applied in the first slot. After the multi-slot packet, the channel continues on the hop as dictated by the master clock. For example, let us consider four slots: k , $k+1$, $k+2$ and $k+3$. Ordinarily, these would be associated with hop frequencies f_k , f_{k+1} , f_{k+2} and f_{k+3} . However, a three-slot packet that starts in slot k uses f_k for the entire packet. The next packet begins in slot $k+3$ and uses f_{k+3} .

Physical link definition

Two types of link have been defined for supporting multimedia applications that mix voice and data:

- synchronous connection-oriented (SCO) link;





Figure 4
Fixed packet format.

Figure 5
Header fields.



polling scheme. A slave is only permitted to send in the slave-to-master slot when it has been addressed by its MAC address in the preceding master-to-slave slot. A master-to-slave packet implicitly polls the slave; that is, an ordinary traffic packet addressed to a slave polls the slave automatically. If no information to the slave is available, the master can use a POLL packet to poll the slave explicitly. POLL packets consist of an access code and header only. This central polling scheme eliminates collisions between slave transmissions.

Establishing connection

When units are not participating in a piconet, they enter standby mode, from which state they periodically listen for page messages. From the total set of 79 (23) hop carriers, a subset of 32 (16) wake-up carriers has been defined. The subset, which is chosen pseudo-randomly, is determined by the unit identity. Over the wake-up carriers, a wake-up sequence visits each hop carrier once: the sequence length is 32 (16) hops. Every 2,048 slots (1.28 s), standby units move their wake-up hop carrier forward one hop in the wake-up sequence. The native clock of the unit determines the phase of the wake-up sequence. During the listening interval, which lasts for 18 slots or 11.25 ms, the unit listens on a single wake-up hop carrier and correlates incoming signals with the access code derived from its own identity. If the correlator triggers—that is, if most of the received bits match the access code—the unit activates itself and invokes a connection-setup procedure. Otherwise, the unit returns to sleep until the next wake-up event.

Units connecting to a unit in standby mode must know the standby unit's identity and preferably its native clock

- to generate the required access code (which constitutes the paging message);
- to derive the wake-up sequence;
- to predict the phase of this sequence.

Since paging units cannot accurately know the native clock of a recipient, they must resolve the time-frequency uncertainty. They do so by transmitting the access code continuously—not only in the hop in which they expect the recipient to wake up, but also in hops before and after. For a period of 10 ms, paging units transmit the access code sequentially on several hop frequencies around the expected hop carrier. Note: the access code is only 72 bits long (72 ms). Therefore, many codes can be sent in the




Figure 7
Continuous variable slope delta (CVSD) waveform coding.

space of 10 ms. The 10 ms train of access codes on different hop carriers is transmitted repeatedly until the recipient responds or a time-out is exceeded.


When a paging unit and recipient select the same wake-up carrier, the recipient receives the access code and returns an acknowledgement. The paging unit then sends a packet containing its identity and its current clock. After the recipient acknowledges this packet, each unit uses the paging unit's parameters for hop selection—thereby establishing a piconet in which the paging unit acts as the master.

To establish a connection, the paging unit must obtain the identity of units within transmission range. Therefore, it executes an inquiry procedure: the paging unit transmits an inquiry access code (which is common to all Bluetooth devices) on the inquiry wake-up carriers. When a recipient receives the inquiry, it returns a packet containing its identity and clock—the very opposite of the paging procedure. After having gathered each response, the paging unit can then select a specific unit to page (Figure 8).

Scatternet

Users of a channel must share capacity. Although channels are 1 MHz wide, as more and more users are added, throughput per user quickly drops to less than a few tens of kbit/s. Furthermore, although the medium available bandwidth is 80 MHz in the US and Europe (slightly less in Japan, France and Spain), it cannot be used efficiently when every unit must share the same 1 MHz hop channel. Therefore, another solution was adopted.

Figure 8
Connection-establishment procedure and maximum time associated with establishing a connection.






Figure 9
A scatternet of four piconets applied to the scenario described in Figure 1.

Units that share the same area and that are within range of one another can potentially establish *ad hoc* connections between themselves. However, only those units that truly want to exchange information share the same channel (piconet). This solution permits several piconets to be created with overlapping areas of coverage. Each piconet adheres to its own hopping sequence through the 80 MHz medium. The channel in each piconet hops pseudo-randomly over the carriers in the 80 MHz band. The users in each piconet have only a 1 MHz hop channel at their disposal.

A group of piconets is called a scatternet. Aggregate and individual throughput of users in a scatternet is much greater than when each user participates on the same piconet with a 1 Mbit/s channel. Additional gains are obtained by statistically multiplexing hop channels and by reusing channels. The 1 MHz hop channel in any given

piconet need only be shared by users of that piconet. Because individual piconets hop differently, different piconets can simultaneously use different hop channels. Consequently, units in one piconet do not share their 1 MHz channel with units in another piconet. The aggregate throughput (the total throughput accumulated over all piconets) increases as more piconets are added. Collisions do occur, however, when two piconets use the same hop channel simultaneously. As the number of piconets increases, performance in the frequency-hop system degrades gracefully. Simulations of a scatternet consisting of 10 piconets indicate that reduction in throughput per piconet is less than 10%. In the scatternet, the radio medium is shared; in a piconet, the channel and information are shared.

Since every piconet uses the same bandwidth, each shares the 80 MHz in an average sense. Provided they select different hop channels, however, no two piconets must simultaneously share the same 1 MHz channel.

Let us assume there are 100 users. If each belonged to the same network, all 100 users would have to share the same 1 MHz channel. Thus, average throughput per user would be 10 kbit/s and aggregate throughput would be 1 Mbit/s. However, if not everyone wanted to talk to each other, we could split the piconet into independent piconets. For example, if the users separated themselves into groups of five, then we could create 20 independent piconets. With only five users sharing the 1 MHz hop channel, throughput per user increases to 200 kbit/s and aggregate throughput increases to 20 Mbit/s. Obviously, this assumes ideal conditions, where no two piconets select the same hop channel at the same time. In reality—because the piconets hop independently—collisions will occur, reducing effective throughput. Nonetheless, the final throughput obtained from multiple piconets exceeds that of a single piconet.

The maximum number of units that can actively participate on a single piconet is eight: one master and seven slaves. The MAC address in the packet header, which is used to distinguish each unit, is limited to three bits. Figure 9 illustrates the scatternet approach applied to the scenario shown in Figure 1.

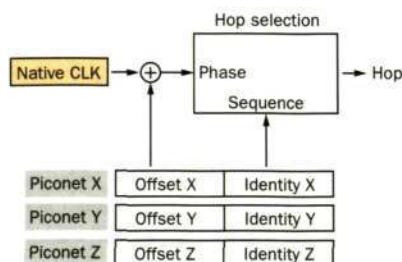


Figure 10
Hop selection in inter-piconet communication.

Inter-piconet communication

Different piconets adhere to different frequency-hop sequences and are controlled

by different masters. If a hop channel is temporarily shared by independent piconets, packets can be distinguished by the access codes that precede them—access codes are unique for each piconet. Piconets are uncoordinated and hop independently; synchronization of different piconets is not permitted in the unlicensed ISM band. Nonetheless, units may participate in different piconets on a time-division-multiplexing (TDM) basis. That is, a unit can sequentially participate in different piconets, provided it is active in only one piconet at a time.

Inter-piconet communication is achieved by selecting the proper master identity and clock offset in order to synchronize with the channel of the desired piconet (Figure 10). A corresponding set of identity and clock offsets is available for each piconet. Whenever a unit enters a piconet, it adjusts the clock offset to account for minor drifts between the master clock and the unit's native clock. A unit can thus act as a slave in several piconets. When leaving one piconet for another, a slave informs the current master that it will be unavailable for a predetermined period. During its absence, traffic on the piconet between the master and other slaves continues as usual.

A master unit can also periodically jump to another piconet and act as a slave (were it to act as a master in the new piconet, that piconet would have the same channel parameters as the "old" piconet—therefore, by definition the two would be indiscernible). When a master unit leaves a piconet, all traffic on the piconet is suspended until it returns.

Figure 11 shows a slave participating in two piconets. Piconet X consists of master X and slaves A_x and B_x . Piconet Y consists of master Y and slaves A_y , B_y , and D_y . Slave C_{xy} participates in piconets X and Y. The clock of each unit is also shown. Positive offset (indicated by blue) or negative offset (indicated by red) has been added for synchronization with the master clock. Slave C_{xy} contains a native clock and two offsets for the master units X and Y respectively.

Authentication and encryption

To ensure user protection and information secrecy, the system must provide security measures that are appropriate for a peer environment—that is, each unit in Bluetooth must implement authentication and en-




Figure 11
Participation of slave C_{xy} in two piconets.

crypton algorithms in the same way.

A base level encryption has been specified that is well suited to silicon implementation, and an authentication algorithm has been specified that provides a level of security for devices lacking in processing capabilities. Future support for ciphering algorithms will be backward-compatible.

The main security features are:

- a challenge-response routine—for authentication;
- stream cipher—for encryption;
- session key generation—session keys can be changed at any time during a connection.

Three entities are used in the security algorithms: the Bluetooth unit address, which is a public entity; a private user key, which is a secret entity; and a random number, which is different for each new transaction. As described above, the Bluetooth address can be obtained through an inquiry procedure. The private key is derived during initialization and is never disclosed. The random number is derived from a pseudo-random process in the Bluetooth unit.

Conclusion

Bluetooth is a system for providing local wireless connectivity between portable devices. It is particularly suitable for *ad hoc* networking. The air interface has been optimized to provide maximum immunity against interference in the 2.45 GHz band. The system is supported by several leading manufacturers of personal computers and telecommunications equipment. The first consumer products to support Bluetooth are expected to appear on the market around year-end 1999.

Internet directory services with click-to-dial

Mats Eriksson

Internet directory services make available vast resources of the Internet, helping users to save valuable time in finding names, e-mail and postal addresses, telephone numbers, and so forth. Click-to-dial services allow users to invoke calls by clicking on the telephone number they retrieve using an Internet directory service.

Ericsson's solution sets up and transmits calls through the PSTN, but it is equally compatible with VoIP. Operators can provide IDQ and click-to-dial together or separately, charging for each service on a per-user basis, by subscription, or through sponsorship (advertising).

The author describes the IDQ and click-to-dial services, and the components that make up Ericsson's small- and large-scale solutions.

number and invoke a call with a simple click of the mouse. Calls are transmitted through the switched telephone network, which will continue to be a major voice carrier.

Today and tomorrow

To set up an ordinary phone call today, we must dial the number digit by digit. To find the number, we rely on memory, a phone book, an electronic directory on an in-house intranet, or on a public directory inquiry service. Although the idea of controlling phones through a computer interface is not new, when applied in public environments, such as the public switched telephone network (PSTN) and the Internet, it will add a new dimension in terms of accessing information and speeding up communication.

Telecom operators and other content providers earn substantial revenue from directory information services by selling directory inquiry services and advertising in various book directories, such as the Yellow Pages. Internet directory inquiry (IDQ) and click-to-dial functionality constitute an alternative way of providing these services. What is more, they directly benefit Internet users by allowing them to look up phone numbers in operator directory databases.

Effective directory lookup is a necessity in today's society, in which people and organizations tend to move more frequently and over greater distances than ever before. Hav-

As a global source of information, the Internet has been around for a short time, compared with the circuit-switched telephone network for voice transmission, which long ago earned its place in modern history. But both media require users to follow certain manual procedures for reaching another party. To send e-mail, for example, you must type the address of the person you want to contact; likewise, to phone someone, you must dial that person's number.


However, by using the Internet's vast information resources and combining it with the telephone network, we can simplify manual processes and achieve new functionality—offering new services and bringing in higher revenues.

Internet directory services with click-to-dial functionality enable users to choose a

Box A Abbreviations

API	Application program interface
AUC	Authentication server
CDR	Call data record
CS1+	Modified capability set no. 1
DAG	Directory access gateway
IDQ	Internet directory inquiry
IN	Intelligent network
LDAP	Lightweight directory access protocol
PSTN	Public switched telephone network
SCP	Service control point
SIG	SCP interconnection gateway
SSP	Service switching point
TCP/IP	Transmission control protocol/Internet protocol
VoIP	Voice-over-Internet protocol

Figure 1
Imagine being able to combine the best that the Internet and telephone networks have to offer; for example, to find information on people (addresses, phone numbers, etc.) and to call them automatically.



ing access to a vast directory gives users a better chance of finding the right number at once.

Internet-based directory services also convey added economic benefits, since on-line self-service costs less to maintain. Thus, this service is a winning proposition for operators and their customers. And when the service is combined with click-to-dial, users can place a call without having to transfer information (for example, by jotting down the telephone number on a piece of paper) or dial (entering the number one digit at a time).

The business user

On a typical business user's desk, we find a phone connected to the PSTN and a computer connected to data networks (Figure 1). Depending on what the local operator offers, the business user may use Internet directory inquiry by itself, or IDQ and click-to-dial. When used together, IDQ and click-to-dial provide optimum efficiency and convenience, but because the applications are open, operators may provide the services separately.

If a business user prefers to use the operator's click-to-dial service by itself, he must have access to a local directory of phone numbers. Operators can charge calls either on a per-search basis or against a subscription. IDQ and click-to-dial services are also fully compatible with voice over IP (VoIP). Similarly, the directory service can remain in place and be used with Internet-specific numbering plans.

The residential user

Today, residential users are increasingly likely to have two telephone lines installed at home. In addition, a growing number of households have personal computers connected to the Internet. Nevertheless, because the directory services currently available are costly to maintain, the price of such services remains high for frequent use by private consumers. By providing a Web-based service, such as Internet directory inquiry, operators can offer directory services at prices that will increase usage by residential users. Furthermore, operators can keep or even increase their profit margins, since the services are available on an on-line, self-service basis.

The residential-user service can be charged for each search made:

- Charges can be added to a subscriber's regular phone bill.

- Charges can be debited to a sponsoring advertiser.
- Charges can be debited to the subscriber and a sponsoring advertiser.

In the future, when residential users begin using VoIP, the demand for Internet directory inquiry will increase further. However, if residential users solely rely on the IP connection to provide them with telephony and Internet services, the click-to-dial service will require a different solution.

Low impact on operator networks

Operators who introduce Internet directory inquiry and click-to-dial will find that these services have a low impact on their networks. As the services mature, however, millions of people will use them every day, which will surely have some effect on such factors as network load.

The initial introduction requires only a minor upgrade of one service control point (SCP). Therefore, whether they do business on a small or large scale, operators can implement these services without having to change their networks significantly and without having to invest heavily in new equipment and software. Further, the size of any future investments they might make can

Figure 2
Directory look-ups are made from a Web interface. Click-to-dial uses this same interface to control call setup. The call proceeds over a regular phone.

Business Area	Address	Phone Number
Infocom Systems	LM Ericssons väg 8 126 25 Stockholm	08/719 00 00
Mobile Systems	Torshamnsgatan 23 164 80 STOCKHOLM	08/757 00 00
Ericsson Components AB	Isafjordsgatan 164 81 Stockholm	08/757 50 00
Telefonaktiebolaget LM Ericsson	Telefonvägen 30 126 25 Stockholm	08/719 00 00

Name Anders Andersson	Role
Organization 	Location Stockholm

Additional search:

Number of Hits: Pages:

A-number 08-757 50 00

Minutes: Numbers:

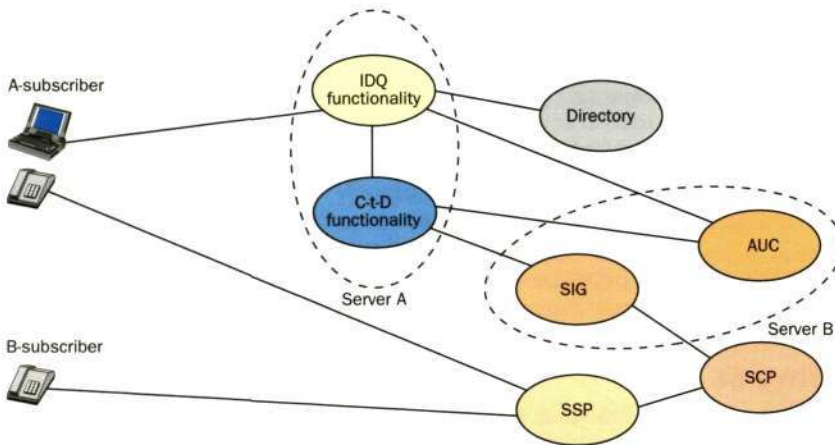


Figure 3
In the small-scale solution, the IDQ and click-to-dial (C-t-D) servers are co-located on the same server hardware. Other applications are located on a separate server. This division is made for security reasons.

be proportionate to market demand and inflow of revenue.

Small-scale solution

The small-scale solution (Figure 3) to providing IDQ and click-to-dial requires two separate UNIX machines, each of which is connected to a directory and the intelligent network (IN). The first machine, which acts as the front server, runs the Web server (for the IDQ service) and the call setup server (for the click-to-dial service). The IDQ Web server connects to the directory by means of the lightweight directory application protocol (LDAP). The second machine handles user and charging functions and runs the

signaling gateway to the intelligent network. It is also connected to an SCP that receives call setup orders and prompts appropriate service switching points (SSP) to set up calls.

Large-scale solution

An operator's decision to implement IDQ and click-to-dial on a large-scale basis is always dependent on the scope of the network and its environment as well as on the maturity of the customer base. In addition to this, other factors differentiate a small-scale solution from a large-scale solution.

Each application requires its own dedicated machine to satisfy the need for computing resources (Figure 4). The front Web servers can be distributed to local areas where use is heavy, or they can be based on a high-availability platform.

A directory access gateway (DAG) must also be introduced for handling multiple directories—information is seldom stored in a single directory; instead it is divided into geographical areas, market segments, or into other categories. Several operators, each with directories of their own, can share information services through a directory access gateway.

In large-scale solutions, the server that handles user and charging functions is located on a high-availability platform. SCP interconnection gateways (SIG) between the TCP/IP environment and the No. 7 signaling network are in parity with the SCPs. This configuration enables charging data to be accurately kept and ensures the availability of services.




Figure 4
The large-scale solution consists of several physical servers. IDQ and click-to-dial servers are distributed throughout the network. The DAG, AUC and SIG are centrally located together with the SCP.

Components

IDQ server

As the end-user interface to the service, the IDQ server allows users to make directory inquiries. The IDQ server (Figure 5) consists of

- a Web server at the front end;
- an LDAP server at the back end;
- a RADIUS client that checks whether a user is authorized to access the Web server.

Requests for directory searches are received at the Web server end and sent on as LDAP requests via the back end to a directory server. When a user has successfully performed a directory search, the RADIUS client computes charging information. The RADIUS client processes information through an authentication server (AUC).




Figure 5

The IDQ server is a Web server that interfaces client look-ups. It authorizes and charges clients using the RADIUS client and the AUC, and passes look-ups through the LDAP server to the DAG or to the actual directory.

Click-to-dial server

The click-to-dial server (Figure 6) constitutes the interface that allows users to invoke the call setup service. Users do not see any difference between the IDQ and the click-to-dial interface, since both are activated when the click-to-dial option is selected.

The click-to-dial server contains a Web server and a RADIUS client, which checks authorization and processes charging information, much as the IDQ server does. This makes it possible to charge click-to-dial calls on a per-click basis. Through its back end, the server communicates with the SCP interconnection gateway using an application program interface (API) that sends the call setup orders over a TCP/IP-based network.




Figure 6

The click-to-dial server receives call setup orders through its Web server. The RADIUS client, which interfaces with the AUC, performs authorization and charging functions. Call setup orders are forwarded through an API to the SIG.

Directory access gateway

As the name implies, the directory access gateway provides access to multiple directories. This is necessary because operators generally store information in more than one directory. Notwithstanding, these various sources are compiled and displayed to the user as a single, efficient directory.

Just as operators may at times want to share directories with one another, at other times—especially in competitive, deregulated markets—they want to keep certain information to themselves. In that case, the DAG can serve as a neutral inter-connector between operators (Figure 7).

What differentiates the DAG from other types of directory is that it uses indexes of information instead of actual data records. When a request reaches it, the DAG shows




Figure 7

The DAG is used when multiple directories are connected to the system. It knows in which directory to search.

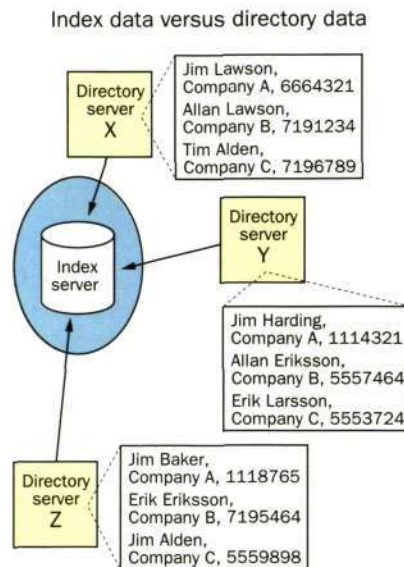


Figure 8
Directories contain structured information, such as names, addresses, phone numbers, e-mail addresses, and so on.

which directory or directories should be searched. But because indexes only present an overview of what the databases contain, it is impossible to access the databases directly. In itself, the DAG contains several interfaces for client requests. However, when the DAG is used in an IDQ or click-to-dial solution, it answers queries from the IDQ server through *LDAPv3*. Thus, users of the DAG can access *LDAPv2*, *LDAPv3* and *Whois++* directories.

Directories

One of the first things that comes to mind when we think of directories is the "white pages," which list subscriber names, organizations, addresses, and phone numbers. But other types of information can also be included in the IDQ service (Figure 8). For example, an e-mail directory or the kind of information we usually associate with the "yellow pages."

Directories connected directly to the IDQ server support *LDAPv3*. However, if the directory is offered from a directory access gateway, *LDAPv2/v3* and *Whois++* may be used without adaptation. Most types of X.500 directory provide LDAP access.




Figure 9
The SIG receives call setup orders from the click-to-dial server, converts the orders from TCP/IP to No. 7-based signaling, and forwards them to the SCP.

SCP interconnection gateway

The SCP interconnection gateway is the interface between the TCP/IP-based network and the No. 7 signaling network. The SIG receives call setup orders from the click-to-dial server via the API and converts the orders from TCP/IP to CS1+ commands. When a call has been set up, the SIG is released from the call process. However, if the call cannot be set up—for instance, if the B-subscriber is busy—the SIG signals back for further action.

Authentication server

The authentication server communicates with the IDQ and click-to-dial servers through their RADIUS clients (Figure 10). One of two major components in the AUC is the RADIUS server, which serves RADIUS clients. The RADIUS server verifies and authorizes user identities and collects accounting data from records of calls.

The other major component in the AUC is the formatting server, which reads the RADIUS clients' accounting statistics and formats them into charging information for the billing system. Charging information can also be formatted into standard call data records (CDR) and sent to the telephony billing system. That way, the charges for

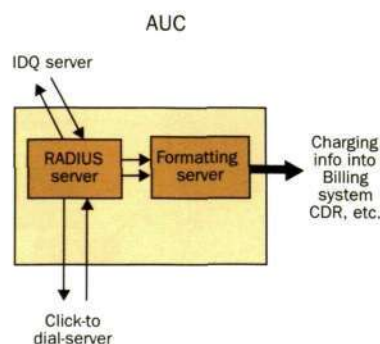


Figure 10
The AUC authorizes and collects charging data via its RADIUS server. Charging data is then formatted; for example, into a call data record (CDR).

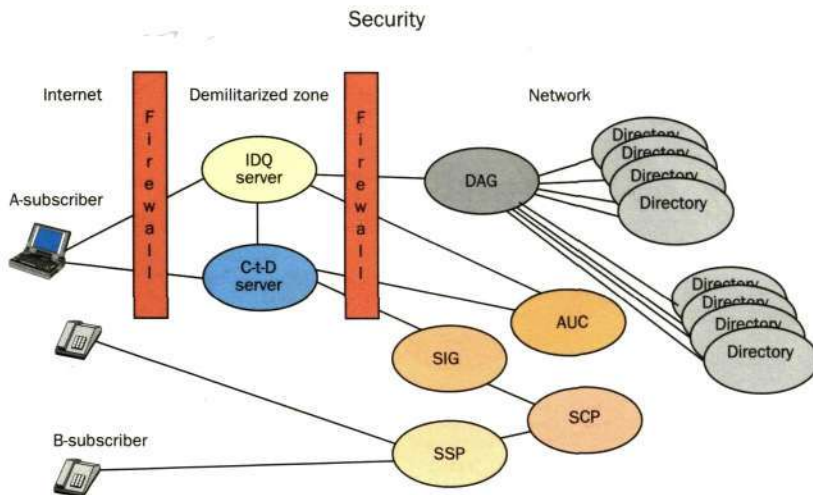


Figure 11
Access to the IDQ and click-to-dial servers is restricted by a firewall, which creates a "demilitarized zone" for the servers. To communicate with internal servers, signals must pass a second firewall. Because the servers can only communicate with their natural neighbors, this configuration gives the internal systems a high level of security.

IDQ and click-to-dial can be added to an ordinary phone bill.

Charges for IDQ and click-to-dial services can be computed in three different ways. The first two methods—which are handled by the AUC, and which involve directory searches and call setup—call for charging per use, by subscription (or a combination thereof), or in combination with advertising. The third method involves computing charges on the basis of how regular traffic is charged and when a call is established. This variant is based on the regular charging system for traffic billing.

Service control point

The service control point (SCP) can run a wide range of applications, but in the scope of IDQ and click-to-dial, it has a particular role. To deploy the new services, operators who already use SCPs need only install a software module.

In operation, SCPs receive call setup orders from the relevant server through the SCP interconnection gateway. The SCPs assign accurate service switching points (SSP) and manage call setup until a call has been established.

Service switching point

The service switching point is the final touchstone in processing the service. SSPs receive call setup orders from the SCP and establish the call. SSPs also take care of the third part of charging, which involves calculating charges for traffic.

By utilizing click-to-dial via the intelligent network and the SSPs, tromboning (in-

efficient routing) is minimized or eliminated, depending on how the SSPs are implemented in the network.

Security

Security plays an important role in solutions where the telephone network connects to the open Internet. The level of security used in the implementation is dependent on operator policy and network architecture.

In a generic security solution (Figure 11), the front Web servers are located in what might be described as a *demilitarized zone* between the public Internet and the telephone-signaling network. At the front, a first firewall is positioned between the user (Internet) and the IDQ and click-to-dial servers. A second firewall forms a barrier between (a) the IDQ and click-to-dial servers and local directories and (b) the internal network, the AUC and SIG.

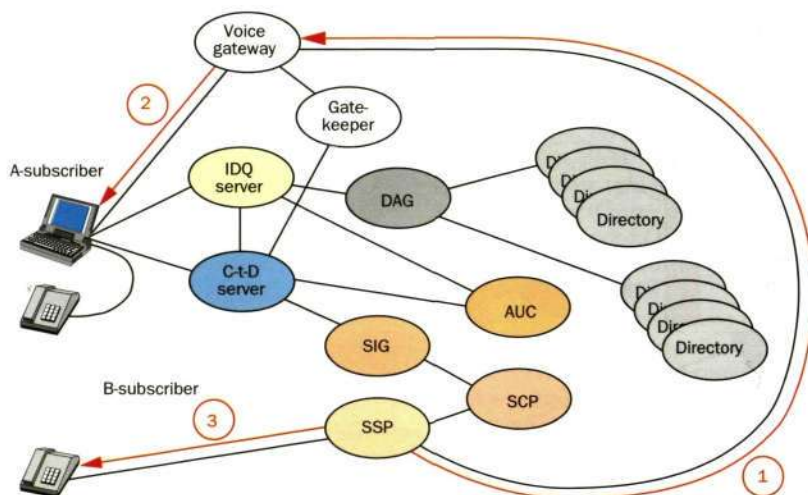
Each of the physical servers solely communicates with natural neighbors, has a stripped operating system, and can only perform the functions that the services require.

The system is managed by means of encrypted communication.

The future—voice over IP

In this context, the IDQ and click-to-dial products focus on current circuit-switched networks and installed bases of hardware and software. Nevertheless, we should not neglect the growing base of voice-over-IP networks. With that in mind, IDQ and

Figure 12
Click-to-dial is provided via voice-over-IP. Call setup is sent from the SSP (1) to the voice gateway and to the user (2) and, finally, to the B-number (3).



click-to-dial were also designed for use with VoIP (Figures 12-14).

In environments where users have migrated to VoIP, the directory service (which is based on IP networks) operates as described above. The tendency today is to use E.164 numbering on the Internet, which is compatible with current directory data schemes.

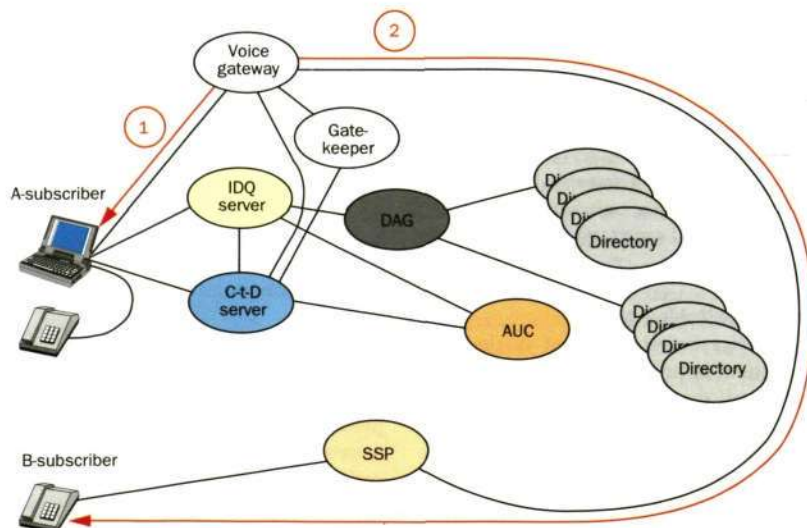
Setting up a call from an IP-based terminal is different from setting up calls from terminals in a circuit-switched network, since IP-based terminals handle both signaling and voice.

- Having clients drive the call setup

process keeps control of the call in the network. This option can be managed from the IN architecture, where calls are set up through the SSP; the IP-based part of calls is treated as a regular incoming VoIP call.

- Call setup can be ordered from an application node in the IP network. In reality, this means that the call must first be set up in the IP-based terminal, and then from the IP-based terminal to a server in the circuit-switched network.
- In the future, the terminal might be used to drive call setup—whenever a number

Figure 13
The click-to-dial server functions as an application node. Call setup is sent as an ordinary call via the voice gateway to the user (1) and to the B-number (2).






Figure 14
The terminal drives the call across the network. The call is sent to the voice gateway (1) and on to the B-number (2).

is encountered in client software, the click-to-dial function could proceed as a regular VoIP call.

Conclusion

The telephony and Internet markets are growing rapidly. At the same time, strong forces are at work to bring the two into convergence.

Competition is on the rise in deregulated markets and users are demanding more and better services; in particular, they want transparent services.

Internet directory inquiry and click-to-dial services satisfy end-user demands for convenience and market-driven pricing. Thanks to improved cost-effectiveness, these services also satisfy operator demands for better cost margins and greater profitability. Appropriately priced services attract new customers, thereby generating new revenues.

Operators with a large installed base need not modify their network. With the addition of a few servers (hardware and software), the new services can easily be introduced on a small or large scale.

Jambala—Intelligence beyond digital wireless

Fred Jones

New, high-value services that target an increasingly segmented end-user market are instrumental in retaining and expanding an operator's subscriber base. Scalable and flexible application platforms are needed to provide the basis for network nodes that cater to such needs.

Jambala is a next-generation application platform that truly fulfills these requirements. In its initial installation, it accommodates HLR/AC and SCP applications and services for D-AMPS/AMPS networks. Future releases will address the area of mobile-to-mobile and fixed-to-mobile convergence.

In describing Jambala's architecture and implementation, the author shows how the platform provides a unique combination of availability, reliability, scalability and Internet readiness, all based on commercially available hardware. He further explains how platform support for Java and CORBA technologies yields a flexible, high-value, platform-independent environment in which services can be created as rapidly as the technology will allow.




Figure 1
Block diagram of Jambala components, the applications it can run, and the external applications and operating systems with which it interacts.

Application platforms of different sizes and configurations are needed for building efficient and manageable intelligent network (IN) nodes. Operators must be able to expand these platforms smoothly and flexibly in response to market needs. Sustaining network growth means ensuring that node growth can be scaled within broad limits, while at the same time maintaining the redundancy, reliability, and availability of performance that is required in network nodes. This must be done in a way that is cost-effective and compatible with future modifications, in order to provide support for operators' rapidly changing needs.

The principal features and benefits of Jambala are:

- a cost-effective solution;
- efficient, cost-effective operation, administration and maintenance (OA&M);
- zero system downtime;
- flexible and future-compliant architecture;
- multi-application support;
- open service-creation environment (SCE);
- support for convergence.

Commercial hardware

The application platform on which Jambala is implemented consists of Ericsson's TelORB operating system/middleware (Box B) combined with off-the-shelf hardware components:

- Two OA&M UltraSPARC 300 MHz processors;
- 200 MHz Pentium boards with 512 Mbytes RAM;
- Two 100 Mbit/s Ethernet switches;
- One CD-ROM drive;
- One 40 GB hard drive;
- One SS7/C7 stack;

Figure 1 shows a block diagram of the different components on the Jambala platform, some of the applications it can run, and the external applications and operating systems with which Jambala can communicate. Network nodes—such as a service control point (SCP) or a home location register/authentication center (HLR/AC), which require a high degree of scalability and must meet stringent, real-time, fault-tolerance requirements—are implemented directly in Jambala.

Other IN nodes with less strict performance criteria, such as service-creation environments (SCE) and service-management systems (SMS), can be implemented in off-the-shelf operating systems and hardware.

Box A Abbreviations

API	Application program interface
BHCA	Busy hour call attempt
C7	CCITT signaling system no. 7
CMISE	Common management information services element
CORBA	Common object request broker architecture
D-AMPS	Digital advanced mobile phone service
DBMS	Database management system
GIOP	General inter-ORB protocol
HLR/AC	Home location register/authentication center

IIOIP	Internet inter-ORB protocol
IN	Intelligent network
MIB	Management information base
OA&M	Operation, administration and maintenance
ODMG	Object database management group
SCE	Service creation environment
SCP	Service control point
SMS	Service management system
SS7	Signaling system no. 7
TCP/IP	Transmission control protocol/Internet protocol
TMN	Telecommunications management network
WIN	Wireless IN

These nodes can make use of the open CORBA 2.0 interface (Figure 3) to access resources and communicate with applications in Jambala.

Zero-downtime operation

To implement traffic- and memory-intensive network nodes, such as HLRs or SCPs, high levels of reliability, redundancy and availability must be built into the nodes. Jambala offers an extensive range of features that satisfy these requirements:

- Automatic recovery from software errors
- Data fault tolerance and redundancy
- On-line backup
- Adaptive hardware configuration
- Smooth software upgrade
- $n+1$ hardware redundancy
- Hot-swap replacement
- Geographical node redundancy

The execution environment provides mechanisms for handling faulty code and erroneous data that can lead to failures in execution. These mechanisms ensure automatic recovery from execution failure, or that execution is terminated without affecting network node operation.

The integrated database supports distributed transactions, thus providing fault tolerance by continuously replicating data. For example, HLR and SCP data is spread out over several processors, where it is partitioned, distributed and replicated as distributed transactions (Figure 4). In the event of processor failure, traffic processes or OA&M processes in the nodes are automatically diverted, allowing the nodes to make use of replicated data in other processors.



Figure 2
Photograph of Jambala hardware components.




Figure 3
The CORBA 2.0 interface stack used for accessing resources in Jambala.

Box B TelORB

TelORB is a distributed operating system for large-scale, embedded (soft), real-time applications that require very high availability. This requirement can be met thanks to

- the TelORB operating system (OS) kernel with appropriate mechanisms;
 - a real-time database;
 - software-configuration control;
 - an associated development environment.
- In addition, a CORBA-compliant object request broker and a Java virtual machine run on top of TelORB.

Features and benefits of TelORB

- Execution environment—programs are executed in processes that
 - are distributed over several processors;

– interact primarily by invoking remote operations.

- Real-time, object-oriented database in primary memory—data is replicated at several physical processors. This configuration makes the system tolerant to hardware faults and serious software faults.
- Software-configuration control—this ensures that the processes in the system execute on available processor resources, which is especially useful when hardware fails.
- A separated management layer for TMN-like operation and maintenance through an object model stored in the target system. The object model is displayed via a user interface on a workstation.
- A driver concept that allows application-specific hardware to be used without having to modify the operating system kernel.

- A built-in CORBA 2.0-compliant object-request-broker core that supports the IIOP for inter-operability with other systems.
- Run-time environments for C/C++ and Java code.
- A software-structuring model in which interfaces can be maintained as separate product items (which aids the development of large software systems). The model is supported by tools that
 - compile the source code in the structure;
 - build load modules;
 - collect information that is needed for configuring the target system.
- In-service upgrades of software and hardware with little or no disturbance to operation. This includes the option of adding processors to increase capacity.

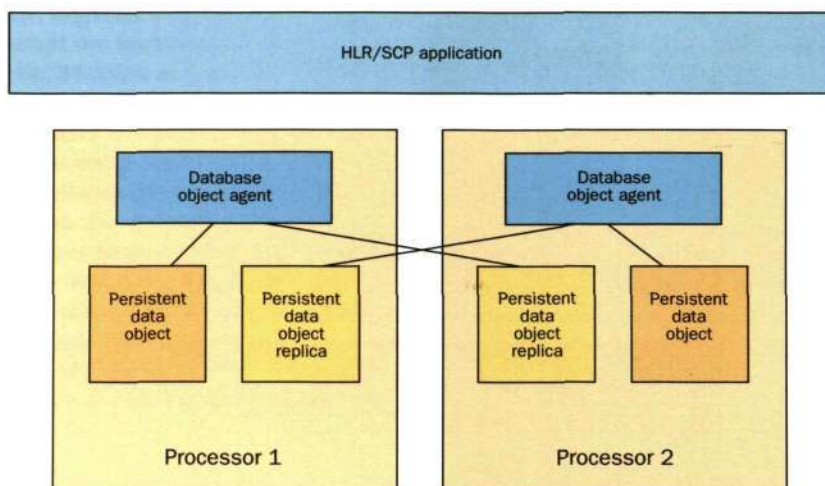


Figure 4
Replication of application data on separate processors.

HLR and SCP data stored in the database is log-marked, which means it can be backed up on line at regular intervals. Backups can be made automatically by the operating system's backup scheduler, or initiated manually by a system administrator. The application data is copied to the hard disk, where it is sorted into files in the partitioned backup archive of the file system. Because it is scaled to use additional hardware resources, the hardware platform configuration can eliminate any adverse effects that making backups might have on HLR or SCP performance.

When a processor board is added to the platform, the platform configuration manager automatically reallocates data and processes to the new processor according to configuration data. The configuration manager monitors platform processors and triggers appropriate alarms should a failure occur. Moreover, if a failure is detected, the configuration manager will attempt to recover the faulty processor by automatically reloading processor data and processes with replicas kept in other processors. If these attempts fail, the configuration manager searches its information to find an alternate configuration. The faulty processor's data and processes are then automatically reallocated to the processors defined in the alternate configuration, and the faulty processor is blocked.

Upgrading application software in Jambala does not require any system downtime. During an upgrade, the new version of application software coexists with the old ver-

sion—the platform allows for concurrent execution of old and new software in the same processor. All software processes running while the upgrade takes place are allowed to continue using the old software, whereas all new processes are channeled into the new software.

Jambala offers $n+1$ redundancy for all processing elements, memory units, and signaling links contained in a network node, which means that a Jambala network node can have no single point of failure. The OA&M processor operates in hot-standby mode to ensure continuous OA&M system availability. The network's signaling links and processor boards operate in load-sharing mode—a highly suitable mode for distributing traffic and providing continuous high-capacity operation. The application platform can also be readily deployed in a configuration based on geographical redundancy, which ensures that network node failure resulting from, say, a natural catastrophe at the site of an HLR or SCP will not have an impact on overall network performance.

Linear scalability


Increasing numbers of wireless subscribers and predicted increases in cellular voice, data and services are pressing HLR and SCP vendors to offer solutions with scalable memory, processing, and signaling capabilities. The Jambala platform offers a linearly scalable solution for all of these areas (Figure 5).

Jambala's scalability allows operators to expand network nodes to accommodate traffic growth patterns. The HLR, for example, is used in areas where subscriber levels are high but where the number of active subscribers is relatively low, typically 25% to 40% of the total number. Where this ratio applies, operators can configure the initial system with tailored memory/processing capacity and a signaling capacity that matches the activity level. As the number of subscribers and subscriber activity increases, operators can expand the system in a cost-effective way.

Internet-ready OA&M

Internet-based OA&M solutions are becoming more prevalent in the mobile-communication industry. Flexibility, platform-independence, low cost, and rapid application development are among the

Figure 5
The Jambala platform scales in a linear fashion to match subscriber growth and activity.



strongest drivers toward solutions of this kind. Jambala offers an OA&M solution using a telecommunications management network (TMN) model and a management information base (MIB). The features of this OA&M system include:

- a managed object view of node resources;
- a manager-agent system that performs node-resource operations;
- CCITT-based notification formatting;
- Common management information services element-based (CMISE) manager-agent interaction;
- CORBA-based alarm/notification service;
- alarm management functions;
- security management functions;
- performance management functions;
- system management functions;
- Web-based graphical user interface;
- open SCE for the differentiation and customization of services.

Communication with the OA&M agent takes place over an interface that is compliant with CORBA 2.0 (Figures 7 and 9). CORBA compliance is achieved by using inter-ORB (the general inter-ORB protocol, GIOP) together with the Internet inter-ORB protocol (IIOP). The GIOP specifies a set of message formats and common data representations for communicating with other ORB-based systems or nodes. GIOP messages are exchanged over TCP/IP networks by means of the IIOP, which allows the Internet/intranet to serve as a bridge that customer operating systems can straddle to perform on-line, on-site and remote OA&M operations. This gives operators a Web-based OA&M solution and the ability to develop and use new OA&M applications.

Java support

The Jambala platform contains an integrated, real-time Java virtual machine (Figure 10) that allows Java applications to be developed on or ported to the application platform. The SCP and associated wireless IN (WIN) services are examples of applications that have been developed in Java.

In addition, the built-in database provides an application program interface (API) that is compliant with the object database management group (ODMG), thus enabling application portability across different database management systems (DBMS). The API adds to the Java programming language the binding that Java applications need for interfacing directly with the



Figure 6
Jambala OA&M user interface.

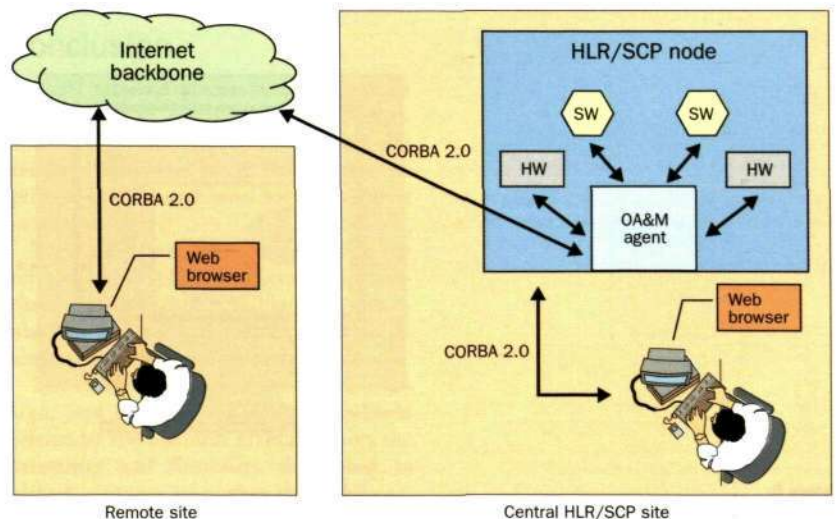



Figure 7
Remote and on-site communication with the OA&M agent over a CORBA-compliant interface.

Figure 8
With Jambala, convergence is possible at signaling-protocol and transport-protocol levels across multiple network standards.



database and for exploiting its strong reliability, availability and redundancy characteristics.

Support for convergence

Support for convergence is becoming increasingly important in a world of deregulation, new end-user demands, multiple network operators, and merging technologies. The key features of a convergence node are adaptability and flexibility; that is, the node's ability to reuse existing infrastructure and technology as well as to incorporate new developments and technologies. Jambala's Internet readiness and scalable support of real-time Java applications and multiple protocols make it an ideal mobile-to-mobile and fixed-to-mobile convergence node (Figure 8).

Applications

HLR/AC

One of the first applications to be implemented on the Jambala platform will be an HLR/AC for digital, advanced mobile phone service (D-AMPS). Operators will observe that the scalability of memory and signaling links offered by the new platform allow them to dimension the HLR/AC to match subscriber and traffic levels in their mobile networks.

SCP

A WIN Phase 1-capable SCP (PN-3660, Figure 9) has also been implemented on the Jambala platform. This application uses the platform's distribution mechanisms to trigger WIN services in the SCP, a network node that is implemented using the Java virtual machine and the platform's built-in DBMS API. By using the Java virtual machine to implement the SCP, operators can also use JavaBean technology and commercially available software development tools (such as Java Studio) to design and implement end-user services. An important feature of JavaBean is its modular nature, which lends itself to reusability when operators create multiple services. For example, JavaBean provides the basic blocks that operators need to build WIN services.

SCE/SMS

By using JavaBean technology, operators can use commercially available design envi-

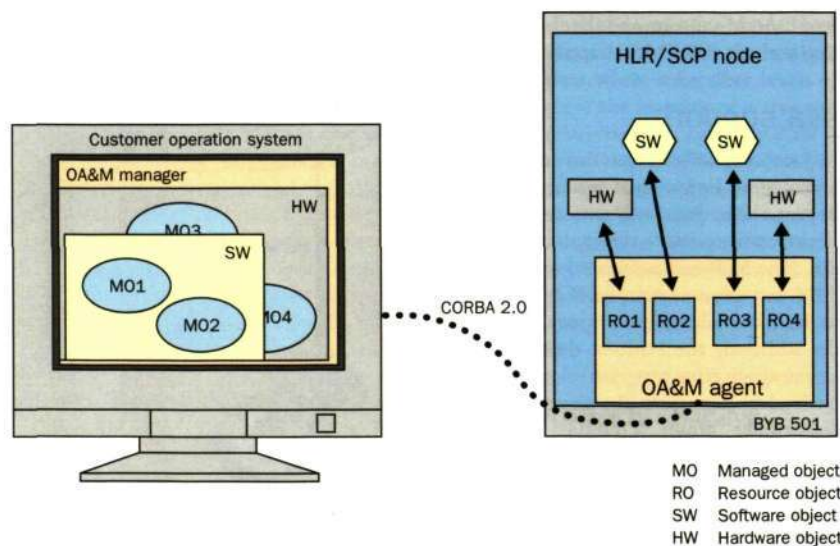


Figure 9
Manager-agent OA&M support.

ronments as SCEs, creating, implementing and customizing their own WIN services. JavaBean technologies provide a library concept of sorts, whose components can be used for customizing and enhancing end-user services. Furthermore, operators can use other technologies, such as push technologies, to remotely update and add to their JavaBean libraries. Because SCE/SMSs have less stringent real-time and fault-tolerance requirements than SCPs or HLRs, they can exploit the application platform's ability to allow applications to be run on other operating systems, as well as its capacity for interfacing with the platform's SCP via the CORBA 2.0 interface.

Future architecture

Since the architecture of the application platform is modular, each component in the system can be independently evolved and upgraded. Therefore, the evolution of the system is best discussed in terms of its individual components.

Processing capability

The processing system consists of several commercially available processor boards that can be replaced without affecting system availability. Thus, operators can capitalize on innovations in processor technology that are gaining ground in the computer industry, thereby maximizing the price-to-performance ratio for node growth in their own networks.

Signaling capability

The signaling system provides support for SS7 and C7 signaling. Because the system consists of commercially available network signaling products, operators can integrate improvements and upgrade the system in a cost-effective manner.

TelORB operating system/middleware

TelORB is a real-time, fault-tolerant and scalable operating system/middleware solution. The ongoing in-house development of this product, which is supported by Ericsson's partnerships in the industry, ensures that the product and its external interfaces will evolve in response to market demands.

OA&M support

The system's OA&M component will be the first commercial implementation of a TMN-based interface that is compliant with CORBA 2.0 IIOP. In addition, it can easi-

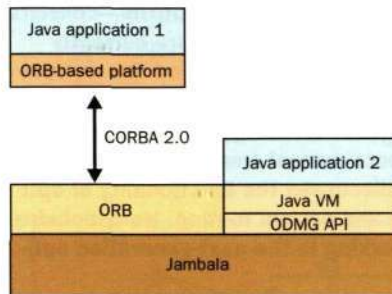


Figure 10
Local and remote support for Java applications.




Figure 11
The wireless intelligent network (WIN) reference model.

ly be integrated into other operating systems; for example, those of new or established operators.

Conclusion

Jambala provides operators with a unique combination of cost-effectiveness, performance and scalability. It uses commercially available hardware, which ensures that its price-to-performance ratio is second to none in the industry.

The TelORB operating system gives the robust Jambala platform outstanding, real-time node performance, and Jambala's modular hardware design provides flexible scalability within broad limits.

The combination of object-oriented design, Java support, and Internet readiness offered by the platform gives operators the versatility and flexibility they need to remain competitive as they enter the next century.

ERION—Ericsson optical networking using WDM technology

Magnus Grenfeldt

Ericsson's next-generation transport-network technology—ERION—consists of products based on WDM technology. Several customers are currently evaluating ERION. Two others—BT and Sonera—have qualified the product and are deploying it in volume.

The author describes the benefits of ERION technology, how it can be introduced into networks, and how it affects current transport technologies, such as SDH, PDH, ATM and IP. He also discusses the functionality of optical networks, primarily optical protection and optical routing. He concludes by describing the role of optical subnetworking in the next-generation optical networks.

- each time a signal is analyzed, split, or filtered;
- to compensate for signal degradation—the result of signals traveling in an optical fiber.

EDFAs boost signal strength without terminating signals into the electrical domain. Without EDFAs, we could not implement dense WDM technology in telecom networks.

The optical network

There are several schools of thought regarding optical networking. In the early 1990s, the prominent view emphasized networks based on optical cross-connects, which enable the provisioning of dynamic bandwidth and facilitate restoration when failures occur.

Ericsson has been involved in the research of optical cross-connects for several years. To date, most of their activities have been carried out in collaboration with European operators and equipment vendors participating in the RACE and ACTS programs, whose research has been well received. One Ericsson project won the RACE award for outstanding contribution to technological progress. This and other research activities within the company have guided product development in the ERION program, leading to a networking strategy that influences and supports the products.

As a network component, an optical cross-connect differs from an electrical cross-connect. For instance, optical techniques have limited flexibility and grooming granularity.

Wavelength division multiplexing (WDM) has been in use for some time in telecommunications networks. Early applications, called non-dense WDM, used 1,300 nm and 1,550 nm wavelengths (second and third transmission windows), which could potentially double the transmission capacity of a single optical fiber.

Breakthroughs in WDM technology—such as the development of erbium-doped fiber amplifiers (EDFA), advanced filtering technology, and improved wavelength stability of telecom lasers—have blazed the way for new applications (for example, applications with different and closely spaced wavelengths—dense WDM, DWDM) and increased transmission capacity 16 to 32 times.

Channel separation in DWDM—or the difference in wavelength between adjacent channels—is typically less than one nanometer. This means that some 30 to 40 channels can be squeezed into the active region of an EDFA.

To maintain the transmission quality of optical systems, signal strength must be boosted

Box A Abbreviations

ACTS	Advanced Communications Technologies & Services
ATM	Asynchronous transfer mode
DWDM	Dense WDM
DXC	Digital cross-connect
EDFA	Erbium-doped fiber amplifier
ERION	Ericsson optical networking
IP	Internet protocol
O&M	Operation and maintenance
OADM	Optical add/drop multiplexer
OSNR	Optical signal-to-noise ratio
PDH	Plesiochronous digital hierarchy
RACE	Research and technology development in advanced communication technologies in Europe (1987-1995)
SDH	Synchronous digital hierarchy
WDM	Wavelength division multiplexing

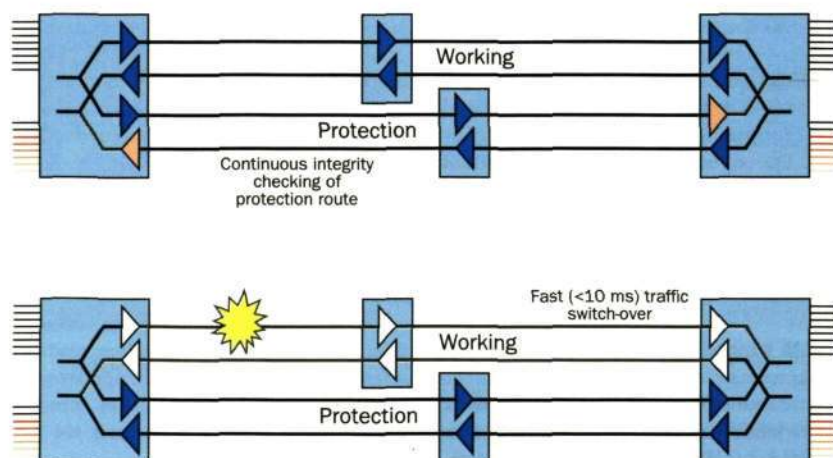


Figure 1
Optical multiplex section protection. Traffic is sent on working and standby sections simultaneously, and the receiving side determines which section is to be used. This allows the system to help monitor the standby section continuously.

Optical subnetworking

Although researchers envision the day of all-optical networks, analysis of requirements for next-generation transport networks points to solutions based on electronic switching. Due to the shortcomings of optical techniques, some functions, such as grooming and flexible path-provisioning must be performed in the electrical domain. Today, these functions are executed in the synchronous digital hierarchy (SDH) layer in multiplexers and digital cross-connects. Electrical techniques, such as asynchronous transfer mode (ATM), also enable network bandwidth to be allocated dynamically or to be taken from a resource pool to satisfy temporary demands.

In terms of network throughput, however, operators can significantly lower costs by moving certain functions, such as traffic protection, from the electrical domain to the optical domain. ERION products are the first of their kind to implement optical multiplex section protection. BT (a valuable reference customer of ERION products) uses this technique to increase transport network availability.

Traffic routing is another function that can be moved from the electrical to the optical domain. To understand the benefit of doing so, let us examine the flow of traffic in present-day networks. Of the traffic that passes through an isolated network node, only a small portion needs to be terminated at the node. This is because

- traffic is demultiplexed and sent to a different layer in the network (for example, to an access layer or to a high-level trunk);
- traffic is groomed into a signal at a high bit rate;
- an end user is connected to the site.

Residual traffic is sent straight through the node and terminated at a different site in the network.

By moving functions to the optical domain, operators can greatly simplify their networks, which in turn reduces operation and maintenance (O&M) costs.

The introduction of optical routing permits through traffic and terminated traffic to travel at different wavelengths. Each connection in the network is allocated its own wavelength, and only terminating connections are dropped into the underlying switch or cross-connect. Some operators claim that 30% to 80% of the traffic in each network node is transient. Thus, the potential savings from adding optical routing are enormous.




Figure 2
Optical routing eliminates the need for terminating transient traffic in intermediate nodes. Each separate connection in the network is allocated its own wavelength.

What is more, optical routing frees up valuable switch and cross-connect capacity, which operators can use to generate additional revenue.

Readers may question whether optical-routing functionality is currently necessary. The answer to this question depends in part on an operator's needs to extend the capability of cross-connects (typically 20 to 80 Gbit/s) and ATM switches—doing so is generally a complex undertaking. The addition of optical-routing functionality compensates for low capacity in cross-connects and ATM switches by increasing bandwidth.

Multi-wavelength transmission, in which optical signals are terminated only once every 600 kilometers (to provide an electrical gasket of sorts between each span), is more difficult to manage than traditional

Figure 3
The introduction of optical routing can save cross-connect capacity. In this example, the introduction of optical routing would reduce load on the cross-connects by one half.

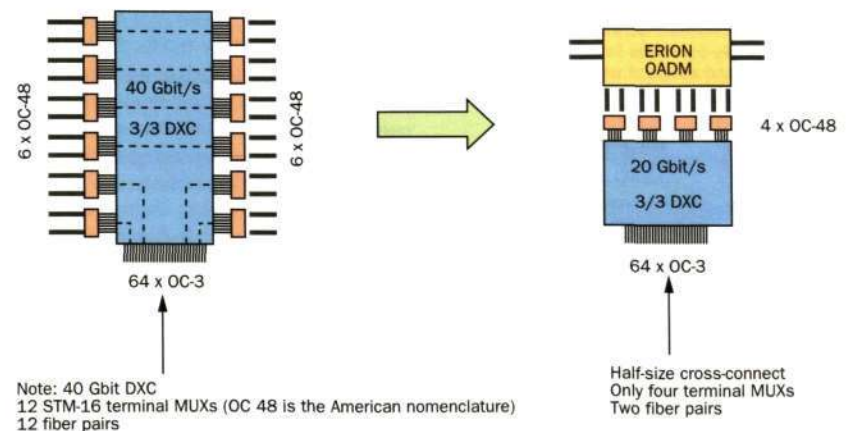


Figure 4
The diagrams show the spectrum at the beginning of an optical section (top) and after it has passed through several consecutive optical line amplifiers (bottom). Noise has been added to the system, and gain tilt from the amplifiers has affected the balance of channel power.





Figure 5
Optical subnetworks interconnected with electrical cross-connects form a reliable, manageable and efficient network structure.



synchronous digital hierarchy-based or pleiochronous digital hierarchy-based (PDH) networks. This is due to the EDFAs: as a network component, the amplifiers introduce and amplify noise, which limits most network signal-to-noise performance. Furthermore, they do not amplify all wavelengths alike. This behavior, called gain tilt or gain flatness, drowns one or more wavelengths in noise, and over-amplifies the others, causing overload in far-end receivers.

These characteristics put stringent demands on the engineering of optical networks. For each installation, engineers must measure losses between amplifiers and calculate signal-to-noise performance. As with many phenomena affected by analogous components, such as the EDFA, engineers apply the weakest-link-in-the-chain principle. For each wavelength of a system, engineers must simulate and track broadband noise and signals and calculate the bit error ratio that results from passing through numerous amplifiers in the fiber.

Standardization bodies are still debating which quality parameters should be used for determining the quality of signals that travel through an optical network. Some groups claim that the optical signal-to-noise ratio (OSNR) should be monitored and measured. Unfortunately, this parameter neither describes the noise generated in network receiver modules nor reports the

penalty for distortion, such as dispersion and non-linear effects. In the end, measurement of the bit error ratio may be the only feasible means of measuring signal quality. In terms of optical networking, this method eliminates the need for spectrum analyzers and similar components in network elements. It also strengthens the relationship between the optical network and its clients. Certain other features traditionally incorporated into SDH, such as bit-error-ratio monitoring and path-trace analysis (which determines signal origin), might also be required in optical networks.

To facilitate the engineering of optical networks, Ericsson has developed several tools that simulate customer applications. The tools primarily help in creating standard offerings (which simplify volume roll-out in networks) and in designing network elements.

You might wonder how optical techniques are supposed to simplify a network when their transmission properties are so complex. They do so through the application of optical subnetworks, and through simplified synchronization, topology, and support-network (for example, data-communication) management.

We conclude from the above arguments that all-optical networks are impractical to engineer and manage. For this reason, Ericsson advocates the concept of optical



Figure 6
The ERION Networker. Today, system capacity is 40 Gbit/s per fiber pair.


subnetworking, whereby small optical domains are interconnected by digital cross-connects, ATM switches or similar methods, to ensure an efficient and manageable network. The interconnection of various subnetworks enables operators to make maximum use of client-layer technologies. Optical protection and routing guarantee efficiency and the availability of a total network solution.

Some might say this approach excludes optical cross-connects from networks. The truth, however, is that optical subnetworks can be built from, and interconnected with

each other by means of, optical cross-connects. The immediate objective is to derive maximum benefit from investments made in client technology and the features it has to offer (such as SDH).

Today, several alliances in the telecommunications industry are beginning to acknowledge the notion of optical subnetworks, and are joining Ericsson in promoting this concept in relevant standardization bodies. As operators look to the future, one of their top priorities is simpler network management. Optical subnetworks will help them to realize this objective.

Figure 7
The future of optical networks? Here we see a network serving several client technologies in accordance with the type and origin of the underlying services.



Relationship to client technologies

Will SDH technology disappear from networks as a result of WDM and optical networking? Speculation over which client technology will dominate is a popular subject at conferences and seminars. One consideration favored by many parties is Internet protocol (IP) over WDM.

Looking ahead, we find no easy or straightforward answers. In looking back, however, we recall that SDH technology was implemented for a specific reason: its predecessor, PDH, could not offer required levels of flexibility or availability. With that in mind, we should observe that the same is true of all-optical networks. In all likelihood, SDH will continue to play a role in future networks—even if the function of

network elements changes. When analysts talk of dynamic bandwidth allocation, which is a driver of the all-optical network, they envision complementing SDH, or replacing it, with ATM in backhaul networks. But optical technology is not likely to replace SDH—at least not in the near future.

Many consider IP to be the next unifying network protocol. Recently, gigabit Ethernet and other interfaces have begun to appear on the market, and some groups claim that these products, by allowing raw IP to be carried over optical networks, will eventually spell the demise of SDH and ATM. Others, however, claim that IP will not be able to support real-time services without ATM or SDH.

Seen from an optical networking perspective, WDM technology will probably be used as a transport technology serving several clients. The idea of multi-client networks supports the transfer of functionality from the electrical domain to the optical domain. Optical protection, for example, is an attractive and low-cost way of increasing the availability of data communication networks.

Ericsson's solution

Although optical subnetworks may assume any shape and topology, developers of ERION products emphasize ring topology. Rings are the predominant choice of net-

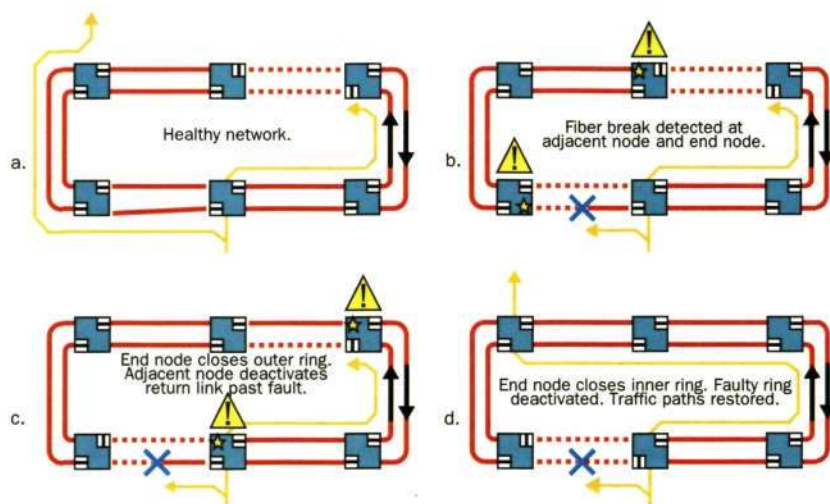


Figure 8
Reaction of Ericsson's FlexBus to a break in the fiber network. The inactive section of the ring is shifted to encompass the failure.

work topology, since they offer a superb means of restoring traffic and always provide two paths between each node pair. Optical rings differ from SDH rings in the following ways:

- They cannot use electrical-to-optical termination in each node (which is an inherent part of SDH networks) to prevent signals and noise from circulating around the ring.
- Transmission impairments severely limit the length of the path for any signal in the ring.
- The switching technology used in self-healing rings must fulfill the reliability requirements of telecom networks.

Ericsson's ring solution, called the flexing bus optical self-healing ring, was derived from applied research activities within the ERION program. The flexing bus can deactivate one section of the ring, which solves the problem of circulating signals and channels and introduces a unique ring-protection mechanism. The only requirement put on the fiber infrastructure is that the physical ring must be closed.

Deactivation of a segment of the physical ring creates a logical bus. Optical add-drop multiplexers can then be used to control or shift the location of the inactive segment, thereby creating multiple logical-bus configurations.

The bus reconfigures to encompass faults within the inactive segment. Wavelengths,

or channels, are routed simultaneously in each direction around the ring. Network elements can receive from either direction, but the inactive ring segment ensures that only one complete path exists between the end points of any given connection. This ring-protection method is fast and does not require a network-management system to intervene. However, restoration events are reported to, and may be inhibited by, the management system.

Conclusion

Ericsson maintains a simple, pragmatic approach to networking as they further explore and exploit DWDM technology and optical networking:

- Large, complex problems are broken down into smaller components with applicable network solutions.
- Optical subnetworks are introduced to provide the benefits of DWDM alongside client-layer technologies.
- Traffic-protection and traffic-routing functions can be transferred to the optical domain—offering cost benefits compared with the business-as-usual approach to networking.

Client systems will continue to play an important role in the network. Optical networks will become a medium for carrying client technologies, such as ATM and IP, SDH, and PDH.




Figure 9
The ERION Networker is commercially available and has already been shipped to several customers.

Access 910 system

Michael Begley, Lars-Olof Haster and Sture Roos

The Ericsson Access 910 is a general-purpose, access-network system that provides PSTN, ISDN, Internet, voice-over-IP, ATM, and switched video broadcast capabilities over a broad array of customer-network interfaces and to a wide range of service networks.

The Access 910 system provides support for practical, cost-effective migration from narrowband to broadband services, making it ideal for present-day as well as future telecommunications environments. Using the remotely controlled MDF—which ensures low operation costs—operators can use the Access 910 system to upgrade individual customer services. Standardized V5.2 and ATM Forum interfaces ensure system flexibility.

The authors describe the Access 910 system, its main functions and supported services; introduce the Access 910 equipment practice, which was designed especially for the Access 910 system; and conclude with a description of the Access 910 management system.

The Access 910 can be configured to operate as a stand-alone system connected, via a V5.2 interface, to the local exchange, or it may be configured to operate and be maintained as part of the AXE local exchange. In either case, the functionality provided by the Access 910 system is the same.

In designing the Access 910 system—Ericsson's new, general-purpose, access-network system—engineers have implemented all applicable standards and standard interfaces to the subscriber and net-

work sides. The system includes integrated plug-in modules for all types of access. Its remotely controlled main distribution frame (MDF) enables service providers to upgrade user services without having to visit Access 910 equipment (Figure 1).

The Access 910 system facilitates migration from narrowband services to future broadband services. There are several reasons why operators choose to update or replace access equipment:

- By upgrading existing copper connections, operators can provide advanced services and reduce maintenance costs.
- Operators can perform planned maintenance and provide individual customer services.
- Operators can upgrade customer subscriptions from the public switched telephone network (PSTN) to an asymmetrical digital subscriber line (ADSL) or add PSTN or integrated services digital network basic access (ISDN-BA) to ADSL.

System functions

The Access 910 system provides on-demand narrowband and broadband services. All supplementary narrowband services are transparent and controlled over the V5.2 interface from the local exchange (in Figure 2, an AXE 10 for narrowband services).

Remotely controlled MDF

The remotely controlled MDF allows operators to respond quickly to subscriber de-

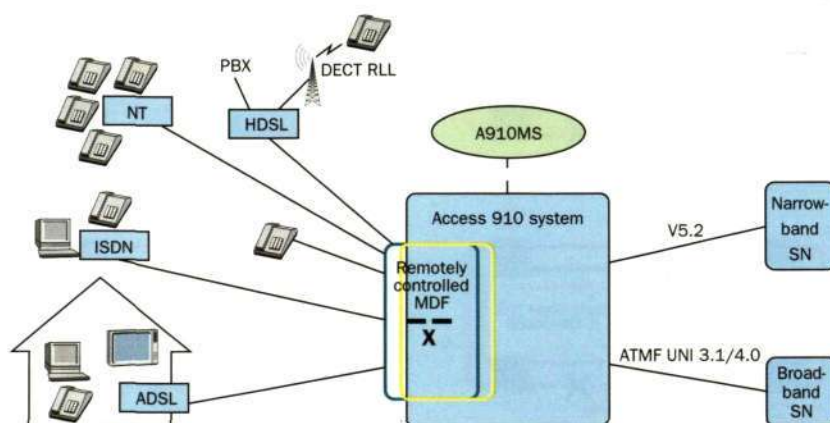


Figure 1
The Access 910 system in the access network.

mands: operators can use it to change subscriber services (for example, from the PSTN to ISDN) remotely (that is, from the service center) without having to visit the Access 910. The operator can then schedule a visit to manually reconfigure the subscriber service in the MDF at a later, more convenient time.

Core network connection

Broadband and narrowband services can share the aggregate synchronous-transport-module (STM) interface to the network, or depending on the narrowband network or the size of the Access 910 system, it can provide separate broadband and narrowband STM-1 streams.

Redundancy

As an option, the Access 910 system may be equipped with standby, narrowband plug-in units. Through a procedure known as protection switching, these automatically take over for access units (AU) and access unit switches (AUS) that have failed. This procedure also generates alarms, identifying the faulty unit or units.

Supported services

ATM traffic concentration

Ericsson's Access 910 systems collect narrowband and broadband traffic from the network for delivery to the exchange. V5.2 interfaces connect the exchange to PSTN/ISDN subscribers. ATM traffic is de-

livered to an ATM network, usually via an ATM access switch concentrator (AAS-C), which connects several remote lines to the ATM access switch (AAS) and concentrates ATM traffic. In general terms, the ATM access switch

- optimizes the use of available access network bandwidth;

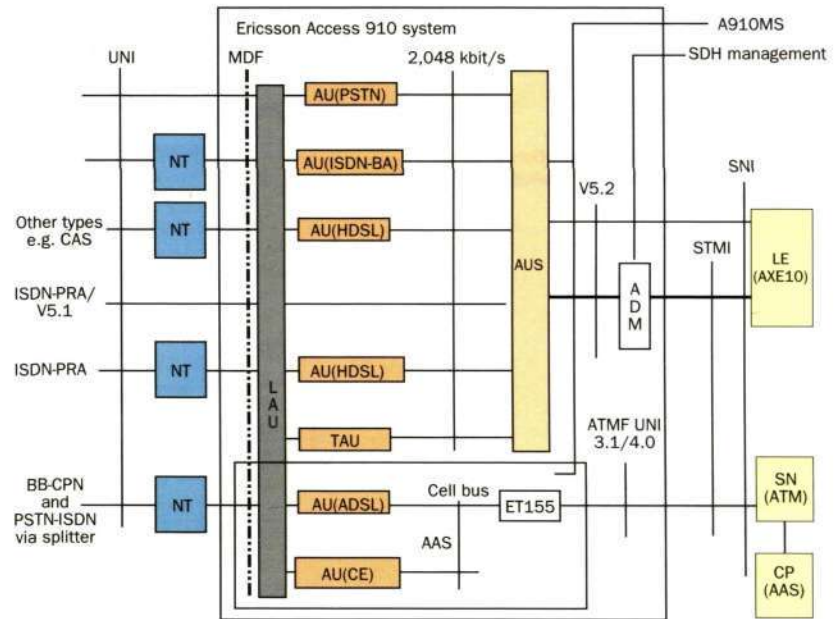


Figure 2
Functional architecture of the Access 910 system.

Box A Abbreviations

A910IMD	Access 910 interface mediation device	DXC	Digital cross-connect	NE	Network element (also called network unit, NU)
A910LCT	Access 910 local craft tool	EM	Element manager	NMDS	Narrowband multiservice delivery system
A910MS	Access 910 management system	EPS	Equipment protection switching	NT	Network termination
AAS	ATM access switch	ET155	155 Mbit/s exchange termination	OAM	Operation, administration and maintenance
AAS-C	AAS concentrator	ETSI	European Telecommunications Standards Institute	OAM-CP	OAM control processor
ADM	Add/drop multiplexer	FTP	File transfer protocol	PBX	Private branch exchange
ADSL	Asymmetrical digital subscriber line	GUI	Graphical user interface	PSTN	Public switched telephone network
ATM	Asynchronous transfer mode	HDSL	High-speed DSL	SDH	Synchronous digital hierarchy
ATMF	ATM Forum	IAS	Internet access server	SN	Service node
AU	Access unit	ICS	Internal communication support	SNI	Service node interface
AU-ADSL	Access unit ADSL	I/O	Input-output	SNMP	Simple network management protocol
AUS	Access unit switch	ISDN	Integrated services digital network	STM	Synchronous transport mode
BB-CPN	Broadband customer premises network	ISDN-BA	ISDN basic access	TAU	Test and administration unit
CAS	Channel associated signaling	ISDN-PRA	ISDN primary rate access	TCP/IP	Transmission control protocol/Internet protocol
CE	Circuit emulation	ISP	Internet service provider	UNI	User network interface
CP	Control processor	LAU	Line access unit	VPC	Virtual path connection
DDF	Digital distribution frame	LCT	Local craft tool	VP-CC	Virtual path cross-connect
DP	Digital processor	LE	Local exchange		
DSL	Digital subscriber line	MDF	Main distribution frame		
		MXC	Mechanical cross-connect		

Figure 3
Example of user A and B connection to the Internet service provider (ISP) using V5 as normal call set-up.





Figure 4
Users A and B connected to an access server with leased-line connection to the Internet service provider.



- enables various kinds of tributary interface.

It is an ATM multiplexer/demultiplexer whose ATM cell-multiplexing function concentrates and multiplexes cell streams to the ATM switch.

Internet narrowband

An Internet access server (IAS) is used in the Access 910 to separate data communication and Internet traffic from traditional speech or telephony calls. Narrowband Internet traffic based on dial-up connections is transparent and connected through the Access 910 system to access server functions in AXE (Figures 3 and 4).

Traffic handling

In accordance with the European Telecommunications Standards Institute's (ETSI) V5 interface specifications, the Access 910 system is transparent in the control plane. Consequently, the Access 910 does not perform any kind of traffic control for narrowband traffic.

Modern releases of AXE introduce the narrowband multiservice delivery system (NMDS). Ericsson has marketed this system under the name Eristream. The Eristream solution is connected transparently to AXE through Access 910 systems.

The Access 910 equipment practice

The integration of electronics technology—in particular that of processors, memory and switching fabrics—yields smaller and smaller core components in the hardware of




Figure 5
System evolution from row to chip.




Figure 6
Line interconnect board (L-board) view.

telephone exchanges. One example of this is the mini-APZ in AXE.

In the Access 910, the control-and-switch board AUS is essentially a digital cross-connect (DXC) or input-output (I/O) system with several standardized 2,048 kbit/s system interfaces. By the same token, each access unit in the Access 910 is an independent system: the magazine or subrack level of architecture has lost all significance. This trend in system evolution is illustrated by Figure 5.

The growing number of 2,048 kbit/s links per unit implies that some form of redundancy must be provided. For backup in the Access 910, $n+1$ redundancy has been implemented.

Since access equipment accounts for 90% of the hardware that connects subscribers, maintenance costs are a major consideration in the system design.

- One way of reducing maintenance costs is to build robust equipment that is easy to maintain—even by unskilled personnel.
- Low power consumption helps keep equipment cool and lessens wear and tear. To reduce overall power consumption, designers have implemented a power-down function, which means that the system regulates its own power needs relative to traffic volume. Moreover, several access units share the same cooling system,

thereby benefiting from statistical traffic load and power consumption. The cooling of subracks, for example, must necessarily disperse heat during maximum power on each board—this is because individual boards are too small to enable statistical calculation.

- Designers have eliminated virtually all internal cabling. System connections have been made through the main distribution frame or the digital distribution frame (DDF). Most remaining cabling is between the MDF and the DDF. Because cabling accounts for much of the installation and material costs, many parts have been integrated into the equipment, particularly in new installations. If this trend continues, the size of MDFs and DDFs will soon exceed that of the equipment they serve.
- External cabling from the copper network has been installed in a way that prevents it from taking up space or from making equipment excessively bulky.

Physical description

The Access 910 equipment practice was developed as an integral part of the Access 910 system. Mechanical protection and protection against electromagnetic discharge has been provided at the board level. Each unit is packed into a sealed box that is placed be-

tween two cooling fins (Figure 7). These fins, which are included in every unit, make it possible to average power consumption for a very large group of access units. Averaging techniques are used for variable traffic loads and different access types. The integrated MDF/DDF design easily enables operators to rearrange equipment to suit access types and traffic load.

The access unit consists of two parts: the access side and the line side. The line side contains optimum functionality, such as MDF, protection switches for connecting spare units, and a connection for a remotely controlled MDF/DDF. The MDF blocks are equipped with pre-connections to minimize jumpering during installation. The pre-connections are opened when a jumper wire is inserted either on the line or the circuit side. The division lines have been placed at 5 mm intervals, which means that units can be incremented in steps of 5 mm. Today, each box is 20 mm wide. All connections are made via jumpering through the MDF/DDF. As an option, backplanes may be added for equipment protection and a remotely controlled MDF/DDF.




Figure 7
Full-size Access 910 system.

For very high power consumption, temperature-controlled fans may also be added. Because cooling is indirect, air filters are not needed.

The Access 910 is available in the following dimensions:

- 1,800 mm (full size);
- 900 mm (half size).

Access 910 management system

The Access 910 management system consists of an element manager (A910MS) and a local craft tool (A910LCT). Interfaces provided by an interface mediation device (A910IMD) are used to integrate the Access 910 management system into the service and network management layer along with other management systems. The management solution supports

- low-level, entry-point deployment suited to small networks;
- large-scale deployment.

The element manager, which includes functionality for V5 provisioning, configuration, fault and security, consists of a management server and several client stations located in the maintenance center. The management server communicates with operation, administration and maintenance control processors (OAM-CP) co-located with the Access 910 system. The OAM-CPs coordinate and store provisioning information but do not handle traffic.

A TCP/IP link is used for communication between the element manager (EM) and the Access 910. Ethernet and an RS485 bus are used for internal communication between the functional units of the Access 910 system.

The management server can be connected to several OAM-CPs, each of which can handle up to 30 Access 910 systems.

The A910MS provides functionality for managing the Access 910 system. Its major functions are

- V5 provisioning;
- configuration management;
- alarm handling;
- test management (fault management);
- security management.

A graphical user interface (GUI) enables operators to configure the Access 910 system for V5.2 connections. No restrictions apply to the number of V5.2 interface identities that can be allocated to a single system. For each interface, however, operators must define the following information:

- V5.2 interface identity
- V5.2 link identities
- Primary- and secondary-link allocation
- Physical C-channels
- Logical C-channel identities
- User port identities

Operators must also specify the logical C-channels on which the various types of data are to be transmitted.

Operators may define two distinct variants of the above configuration, where one variant, defined as standby, can be activated by command from the element manager or on request from the local exchange. When a variant has been activated, newly configured user ports may be added to it. The standby variant remains unchanged, however.

Equipment is provisioned through an automatic plug-and-play configuration or a configuration defined by the operator. From the GUI, operators can specify the type of plug-in unit to be inserted into each slot. If the hardware detected by the system does not match this configuration, an alarm is issued. The equipment-configuration function allows operators to set the administrative state of all equipment entities in the Access 910. Operators may also set the administrative state of V5.2 links and user ports.

Operators enjoy considerable flexibility in configuring the Access 910. Only the number of pre-defined slots into which an access unit switch may be inserted restricts the positioning of plug-in units. Test and administration units and access units may be inserted into any available slot (including slots available for access unit switches). Neither do restrictions apply to the interconnection of plug-in units.

Another important feature of the element manager is that it displays accurate, real-time information from all parts of the Access 910. For ease of use, the current physical structure of managed network elements is displayed together with a graphical representation of the system and board fronts. Operators can use the facilities of the management system to compile an inventory of the Access 910 system. The system also provides support for remotely downloading software to the OAM-CPs using the file transfer protocol (FTP). Operators may then distribute the software to one or more of the system's plug-in units.

Alarm handling

The Access 910 system contains extensive alarm and monitoring capabilities. Each

plug-in unit in the Access 910 is self-monitoring and reports abnormal situations to the management system. In addition, every unit is polled regularly and communication links are monitored. The management system detects communication-link failures and loss of contact with plug-in units.

Alarm summary and event/alarm logging

By default, the main alarm-and-event application window displays a list of current alarms in the Access 910 system. Operators may toggle this view to display every event stored in the log.

The alarm-summary function produces an alarm summary (on-demand) and allows operators to filter alarms according to

- type (communications, quality-of-service, processing, equipment, environment);
- severity (critical, major, minor, warning);
- time interval.

The alarm system provides a range of specific information on detected problems, including identification number, alarm type, alarm location, equipment or process affected, severity of the problem, and date and time of each alarm. An isolated hardware fault in one part of the system does not disturb unrelated parts of the system.

Equipment protection switching

Equipment protection switching (EPS) is the name of a mechanism that offers $n+1$ redundancy within the system. Besides improving availability, EPS also facilitates planned maintenance. The mechanism applies to access units and access unit switches.

Protection switching uses standby plug-in units that automatically take over when an active plug-in unit fails. The protection-switching function can also be invoked manually; for example, when carrying out repairs.

Test management (fault management)

Subscriber-line-maintenance functionality provides a mechanism for detecting and locating faults that have an adverse effect on traffic in the system. This is accomplished in several different ways:

- Operator-requested testing permits operators to select one or more defined user ports for testing. Operators may select different tests for each group of selected user ports. Faults detected during a test are added to the fault list.

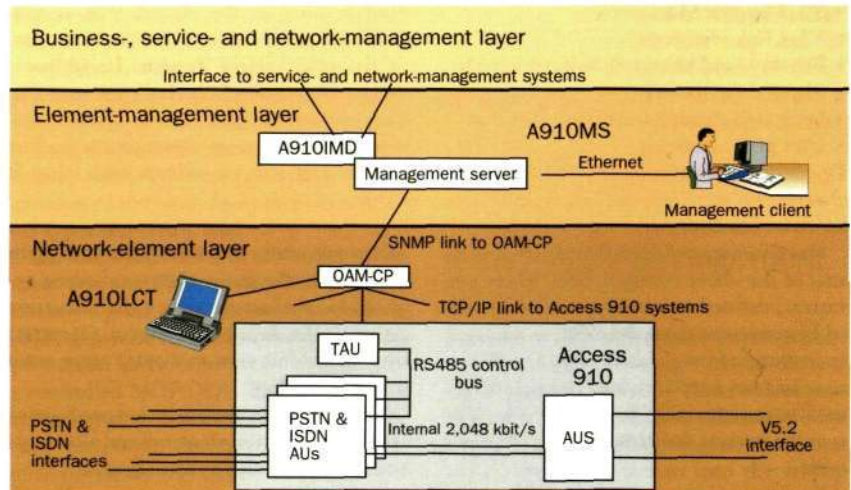


Figure 8
The Access 910 management system.

- Operators may start a schedule of tests for the user ports of a particular access type (for example, PSTN, ISDN-BA and high-speed DSL). For each schedule started, operators may request different tests (line or line circuit) and intervals between tests. Faults detected during a test are added to the fault list.

The fault-management function maintains a list of faulty user ports. If a user port fails a test, the list is updated with the identity of the port, the type of fault detected, and the date and time it was detected. When a fault has been corrected, or if the faulty user port is locked, it is removed from the list. Operators can view the fault list and apply filters to have it display faults of a particular type or date.

Security management

External access to the management system is controlled by standard security mechanisms. Users must log in and supply a password to gain access. Three different security levels exist, each with separate access.

Local craft tool

The local craft tool (LCT) invokes Telnet sessions with the OAM-CP via Ethernet, providing functionality for on-site repair activities—for example, for equipment and fault management (Figure 8).

Interface mediation device

The interface mediation device is an optional product that provides interfaces to

any management system at system and network-management levels.

Conclusion

Ericsson's Access 910 system provides narrowband and broadband capabilities over a broad array of customer-network interfaces and to a wide range of service networks. It supports ATM traffic concentration, the PSTN, ISDN and the Internet. Internet services are provided through the ATM network, the ISDN network, and dial-up connections in the PSTN. Using the remotely controlled main distribution frame, operators can use the Access 910 system to upgrade individual customer services.

The Access 910 management system comprises an element manager and a local craft tool. The element manager enables V5 provisioning, configuration management, alarm handling, test management, and security management for managing the Access 910 system. The local craft tool invokes Telnet sessions with the control processor, facilitating on-site repair activities.

Interfaces provided by the A910IMD integrate the Access 910 management system into the service and network-management layer.

The Access 910 equipment practice and the remotely controlled MDF combine to yield a small physical footprint and reduce the need for manual operations.

Previous issues

No. 4, 1997

Phone Doubler—A step towards integrated Internet and telephone communities

Public Intranet

Geographic information systems

The NDB Cluster—A parallel data server for telecommunications applications

Environment, for better or worse (part 2)

No. 1, 1998

ATM for enhanced network capabilities

AXD 301—A new generation ATM switching system

ATM traffic management and resource optimization

A scalable ATM switching solution

Multiprotocol label switching in ATM networks

Voice and telephony networking over ATM

An element manager based on Web technology

Environment, for better or worse (part 3)

Internet access services, special issue 1998 (English only)

Network solutions for Internet access services

Dial-up service via the integrated access system

Home internet Solution—Always-on Internet access

ANx—High-speed Internet access

Ericsson's scalable frame relay/ATM solution

WISE™ connects GSM to the Internet and intranets

No. 2, 1998

Toward third-generation wireless communication

Wireless intelligent networks—The flexible future

Mobile crosstalk control—Enhancing speech quality in digital cellular networks

GDM-based generation of AXE core switching devices

Telecom management as a competitive tool

Telefonaktiebolaget L M Ericsson

S-126 25 Stockholm, Sweden

Tele: +46 8 7190000

Fax: +46 8 6812710

ISSN 0014-0171

Ljungföretagen, Örebro 1998