Ericsson
1/2001

REVIEW

THE TELECOMMUNICATIONS TECHNOLOGY JOURNAL

Wireless-Trench

More power,
smaller sources

AXE 810—The evolution continues

Open architecture in the core of AXE

Cable modems—Broadband highway to the home

The R380s smartphone

ERICSSON

Ericsson
REVIEW

THE TELECOMMUNICATIONS TECHNOLOGY JOURNAL

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

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

Phone: +46 8 719 00 00

Fax: +46 8 681 27 10

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

New subscriptions:
http://www.ericsson.com/review

Address changes and distribution:
Fax: +46 8 681 27 10
birgitta.blom@lme.ericsson.se

Publisher: Jan Uddenfeldt

Editorial board: Hans Alberg, Magnus Frodigh, Anders Hidmark, Lena Krogstad, Filip Lindell, Lennart Norell, Jonas Noréus, Hans Oskar, Helena Schytt, Sture Sjöström, Bengt Stavenow, Peter Svärd, Ulf Westin, Mikael Westmark

Editor: Eric Peterson

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

Layout: Paues Media, Stockholm

Illustrations: Claes-Göran Andersson, Hara

Printer: Ljungföretagen, Örebro

ISSN: 0014-0171

Volume: 78, 2001

Cover: The new Wireless Trench™ technology makes it possible to operate a power amplifier without any bonding wires to ground (bringing a new meaning to the term *wireless*!), which translates into considerable savings (less pad area, simplified bonding, smaller and cheaper package) without degradation of performance. What is more, it opens the way for the use of silicon in wireless integrated power amplifiers for mobile terminals.

Contents

Contributors

In this issue


Steve Bridges


David LeCorney


Ted Johansson


Mats Jonback

Steve Bridges joined Ericsson in 1997, after having worked many years in the aerospace, environmental controls and personal computer industries. He was the object leader on the R380 development project and served as the Symbian liaison officer before becoming applications support manager for SBU Communicator, in 1999. He is currently working on the new generation of open smartphone and communicator devices and with the Ericsson Developer Zone. Steve holds a B.S. in electronic engineering from Bolton Institute of Technology.
Steve.bridges@eml.ericsson.se

David LeCorney joined Ericsson in 1977 and has worked as a designer, tester, installer, subsystem manager, methodology coordinator, operation and maintenance coordinator, and project manager. He is currently a product manager for Core Switching Products in the AXE Switching Platform. He has worked with AXE for a total of 13 years, in both the control system (APZ) and the traffic-handling system (APT). He has also worked for six years with the AKE system, a predecessor to AXE, and five years with the AXE-N system, which was the forerunner of the current TSP Server Platform system. He holds an M.A. in mathematics from the University of Cambridge.
David.lecorney@uab.ericsson.se

Magnus Enderin joined the Methods, Tools and Training department at Ericsson Utvecklings AB in 1997, and in 1998, he became a product manager for the adjunct processor platform. He is a certified upper secondary school engineer in telecommunication, and holds an M.S. in industrial work environment, production systems from Luleå University and a Licentiate of Engineering in industrial ergonomics from the Royal Institute of Technology in Stockholm.
Magnus.enderin@uab.ericsson.se

Juan E. Figueroa is currently director of the project office for the Cable Access Product Unit within the Home Communications operation of Ericsson's Multiservice Network division. Since joining Ericsson in 1999, he has served as product manager for the PipeRider and for new technologies efforts in the Home Communications group in Lynchburg, Virginia. Juan also previously worked for Bell Labs and Harris Semiconductor (now Intersil). He has a Ph.D. in physics from the State University of New York, Binghamton.
juan.figueroa@ericsson.com

Bill Guzek is currently manager of systems design for the Cable Access Product Unit within the Home Communications operation of Ericsson's Multi-


Magnus Enderin


Juan E. Figueroa


Bill Guzek


Mikael Lindberg


Tomas Lundqvist


Stefan Schultz

service Network division. Since joining Ericsson in 1991, he has held several management positions within Ericsson's R&D units in Lynchburg, Virginia. Most recently, he has been involved in the development of cable modems and in evaluating wired and wireless home networking technologies. He holds an M.S. in computer engineering and an MBA in marketing/innovation management from Syracuse University.
bill.guzek@ericsson.com

□□□□□ **Ted Johansson** is an expert on ■□□□□ RF transistor and process design at Ericsson Microelectronics in Kista, Sweden. He has spent several years developing high-power transistors for base station transmitters. He now divides his time between advanced semiconductor process development for wireless products and the development of power amplifiers in silicon for handhelds. He joined Ericsson in 1989 and holds a Ph.D. in electronic devices from Linköping Institute of Technology, Sweden.
Ted johansson@mic.ericsson.se

□□□□□ **Mats Jonback** joined Ericsson in □■□□□ 1982 to help develop the hardware architecture of the APZ 212 central processor. After several years of mainly working with AXE hardware design (which included a one-and-a-half year period at AMD in Texas), he became a project manager. He currently works as a product manager of AXE Central Processors at Ericsson Utvecklings AB.
Mats.jonback@uab.ericsson.se

□□□□□ **Mikael Lindberg** joined Ericsson □□■□□ as an ASIC designer in 1991. Over the years he has worked with ASIC and hardware system design, always with a focus on hardware development for the core network. He currently holds a line manager position at Ericsson Utvecklings AB, where he works with hardware development of the traffic-carrying parts of AXE. Mikael holds an M.S. in applied physics from the Royal Institute of Technology, Stockholm.
Mikael.lindberg@uab.ericsson.se

□□□□□ **Tomas Lundqvist,** who joined □□□■□ Ericsson as a system engineer in 1994, currently works with hardware system design in the AXE Systems Management department at Ericsson Utvecklings AB. He holds an M.S. in electrical engineering from the Chalmers University of Technology, Göteborg.
Tomas.lundqvist@uab.ericsson.se

□□□□□ **Stefan Schultz,** who joined □□□□■ Ericsson in 1987, currently works as a system manager at Ericsson Utvecklings AB. Prior to this, he has worked with system testing (APZ IO), project management, IO support (two years in Australia), in-service performance, and line management.
stefan.schultz@uab.ericsson.se

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

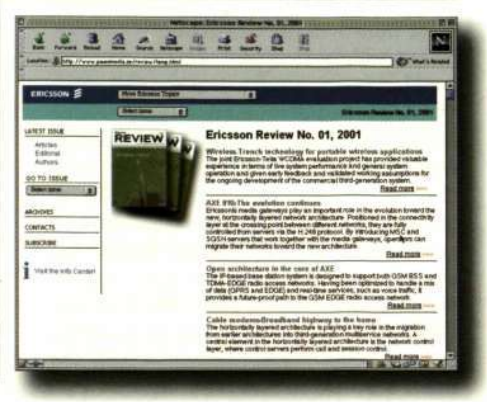

Editorial

Eric Peterson

As you probably know, Japan will be one of the first markets to offer WCDMA on a very large scale, and massive efforts are underway there to set the stage for the roll-out this year. Last summer, I had the opportunity to visit Japan to get a first-hand look at the work being done "behind the scenes." I expected to be impressed (otherwise, I would not have scheduled a visit) but I was thoroughly "blown away" by the sheer magnitude of the undertaking as well as the top-drawer caliber of the people I met and the wonderfully contagious enthusiasm and dedication they exhibited! I was also deeply impressed by the maturity of the market—indeed, if any market is ready for third-generation technology and services, it is Japan!

I lived in Japan from January 1980 to November 1981. In those days, public phone booths were a common fixture outside of train stations and department stores. I remember watching with fascination as hoards of commuters would pour out of the train and stream from the station. Invariably, a queue would form by the phones. As a foreigner in the country, I enjoyed watching as callers concluded their conversations, bowing and bobbing as they mouthed their good-byes over the phone. Today, 20 years later, most of the public phone booths have vanished. Instead, the population relies on mobile phones. Virtually everyone I saw this past summer had one: businessmen, "office ladies," school children, even elderly white-haired ladies wearing traditional Japanese kimono (fortunately, some things haven't changed).

Although truly remarkable, I was not solely struck by the extent to which mobile phones have penetrated and taken root in Japanese society. After all, I'm used to seeing people carry and use mobile phones, since I live in Sweden, which is ranked very high in terms of mobile phones per capita. What really excited me was how, and the extent to which, the Japanese use their phones. As I've already stated, I saw people with mobile phones virtually everywhere I looked. But what I found particularly interesting is that I did not actually hear very many phones ringing or even see so many people conversing on their phones—imagine the cacophony that would arise if all the commuters' phones began ringing during the mad morning rush through Tokyo or Shinjuku station! Instead, I saw people everywhere "reading" or "surfing" (that is, they were using multimedia/data services). I observed people from all walks of life actively consulting their phones while they waited for or commuted on a train or bus, roamed the streets, shopped, and hurried to or from lunch. I even saw a motorcycle deliveryman pull out his phone and scan it quickly while he waited for the light to change.

Thanks to iMode/J-phone and the broad range of available services, the Japanese are already fully accustomed to the concept of the mobile Internet. In terms of technology, services and savvy, the market is very mature and ready to make the transition to full-fledged mobile multimedia. Elsewhere, most markets will be implementing packet data (GPRS, cdma2000) this year. In some measure, this will propel us to the point where Japan now stands, laying a secure foundation for our migration toward WCDMA/UMTS/EDGE. The end-user in me is very encouraged—this year, instead of just reading about it, a lot of ordinary folks outside of Japan will actually begin experiencing the mobile Internet first-hand. It is already a hit in Japan, and I predict that people elsewhere will also be thrilled by it!

Eric Peterson
Editor

Wireless-Trench technology for portable wireless applications

Ted Johansson

Until recently, the use of silicon for the production of wireless integrated power amplifiers has not been viable. Here, however, the author describes how a revolutionary new grounding method makes this technique possible.

Wireless-Trench technology was first applied to an integrated power output amplifier for the DECT system—a digital wireless standard for home and office handsets. Using this same technical approach, Ericsson is currently developing a family of related power amplifier products for GSM 900 and 1800/1900MHz as well as GSM dual/triple-band, PDC, other cordless systems, and Bluetooth.

BOX A, TERMS AND ABBREVIATIONS

DECT	Digital enhanced cordless telecommunications
f_T	Maximum transition frequency
GaAs	Gallium-arsenide
GSM	Global system for mobile communication
HBT	Heterojunction bipolar transistor
LDMOS	Laterally-diffused metal oxide semiconductor transistor
MESFET	Metal semiconductor field-effect transistor
NPN	N-type/P-type/N-type
PA	Power amplifier
PAE	Power-added efficiency
PDC	Personal digital communication
P_{in}	Input power
PNP	P-type/N-type/P-type
P_{out}	Output power
QSOP	Quad small outline package
RF-IC	Radio frequency integrated circuit
SEM	Scanning electron microscopy
Si	Silicon
SiGe	Silicon-germanium
V_{cc}	Supply voltage
VSWR	Voltage standing-wave ratio

An important rule of thumb often used by radio circuit designers says that a maximum transition frequency (f_T)—a general figure of merit for any high-frequency transistor—of at least ten times the operating frequency for the active device is needed to design a smoothly running radio frequency (RF) circuit for modern wireless applications. While this may be true for elementary radio building blocks, the design of the output power amplifier (PA) puts additional requirements on the semiconductor process, thus precluding the use of silicon processes for the power amplifier. For this reason, gallium-arsenide (GaAs) metal semiconductor field-effect transistors (MESFET) and other more advanced technologies have been widely in use in this area up to the present time.

As the wireless communication business continues to grow, there is a great demand for reducing the price of all system parts, as well as for using semiconductor manufacturing processes that can handle very high volumes during the short product cycles of many of the new "gadgets." Silicon bipolar integrated circuit technology (RF-IC) is one suitable candidate for non-digital parts. With today's production processes offering an f_T in the 25 GHz range and higher, silicon can even fulfill the technical requirements needed for power amplifiers. The use of RF-IC technology also facilitates the integration of more functionality in power amplifiers. In particular, power control and active linearization are two important features needed for evolving second-generation and soon-to-arrive third-generation systems.

Ericsson Microelectronics has long enjoyed a good reputation for its development of high-power bipolar and discrete laterally-diffused metal oxide semiconductor transistors (LDMOS) for cellular base stations. As a provider of low-voltage RF-IC and GaAs power amplifiers within Ericsson, and lately for use on the Bluetooth market, the company has wide access to the know-how and technology required to manufacture silicon power amplifiers for high-volume applications. Ericsson Microelectronics has gained this know-how through its own development and in-house production, as well as through its partnerships with other vendors. Today, this wealth of experience is being used to explore silicon RF-IC technology for power amplifiers.

Grounding: the essential element for power amplifier performance

One of the most important factors for achieving high performance in power amplifiers is being able to provide a good conducting path from the active devices to ground. The straightforward solution is to use many bonding wires from the chip's surface to the package. Preferably this includes a grounded lead frame, maybe even with a number of the pins permanently connected (fused) to outside ground connections on the circuit board, or a package in which the back side of the lead frame is exposed and soldered directly to the circuit board.

However, at these high frequencies, the bonding wire from the chip to the package severely limits the circuit's performance. To achieve low resistance and low impedance to ground, the number of ground bonding wires is increased; but since a power amplifier is usually quite a small circuit, the additional pads needed for bonding increases the size of the chip considerably, which can prove to be very costly. Packages with exposed lead frames represent no immediate solution to the bonding wire problem, since they are also considerably expensive.

Device technology

The semiconductor process used for Ericsson's silicon power amplifiers is a half-micron, 25 GHz f_T double-polysilicon bipolar process with additional features for improved wireless performance. This process allows operation up to 5 V and includes NPN and PNP transistors for use in analog and digital designs. For the integrated internal matching network, on-chip capacitors and inductors are used. Four layers of metalization (with a thick top layer

for improved performance of the integrated inductors) are deployed. Advanced, deep-trench isolation is used to obtain small-sized, low-parasitic, high-performance transistors. Figure 1 shows a schematic view of a trench-isolated bipolar NPN transistor.

Wireless-Trench technology

As was pointed out above, a good ground connection is essential for power amplifiers. In addition to the bonding wires, the substrate (the back side of the die) is usually connected to ground (the package), which enables the use of contacts at the chip surface for ground connection. The contacts consist of metal on highly doped semiconductor material. To achieve low resistance, a high processing temperature and long processing time are required. A contact of this type might have to occupy a considerable area to achieve low contact resistance. Furthermore, the semiconductor structure with which it maintains contact consists of several medium-doped layers on the original substrate with more resistance adding up in the current path as a result. In many RF-IC processes, the substrate itself is low-doped, which precludes its use as a low-resistance path to ground.

In Ericsson's RF-IC process for power amplifiers, the substrate selected is as highly doped as possible, close to the solubility limit. Used together with deep trench isolation, the circuit achieves excellent protection against latch-up effects, which might occur in devices that switch large currents, such as power amplifiers.

To assure excellent grounding from the active device to the highly doped semiconductor substrate and the package, a new type of tungsten-metalized substrate contacts have been developed for the front side. These contacts are formed by etching additional deep trenches through the medium-doped upper silicon layers down to the highly doped substrate, filling the trenches with a barrier material, which is annealed for a short time to form a good metal-to-semiconductor contact, finally filling them with tungsten. The new process module is fully compatible with conventional RF-IC processing, and adds only a few more steps to the process flow.

Figure 2 shows a scanning electron microscopy photograph of a double NPN-transistor structure with its deep trench isolation to the left, and the new trench sub-

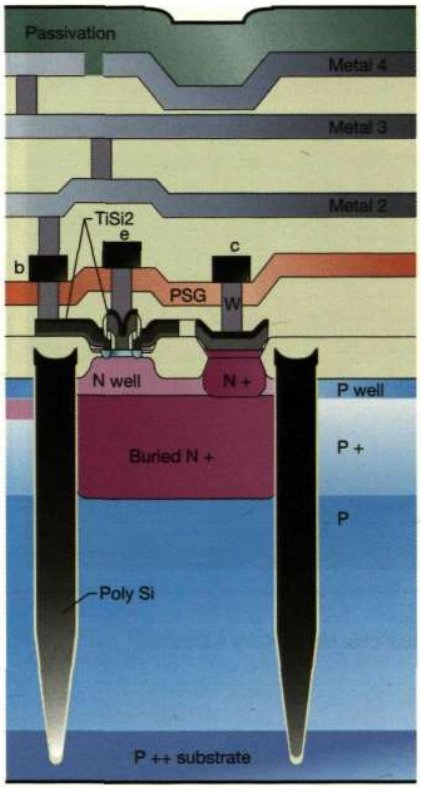

Figure 1
Cross-section of trench-isolated NPN transistor.

Figure 2
Scanning electron microscopy (SEM) photograph of double NPN transistor with new trench substrate contact to the right.

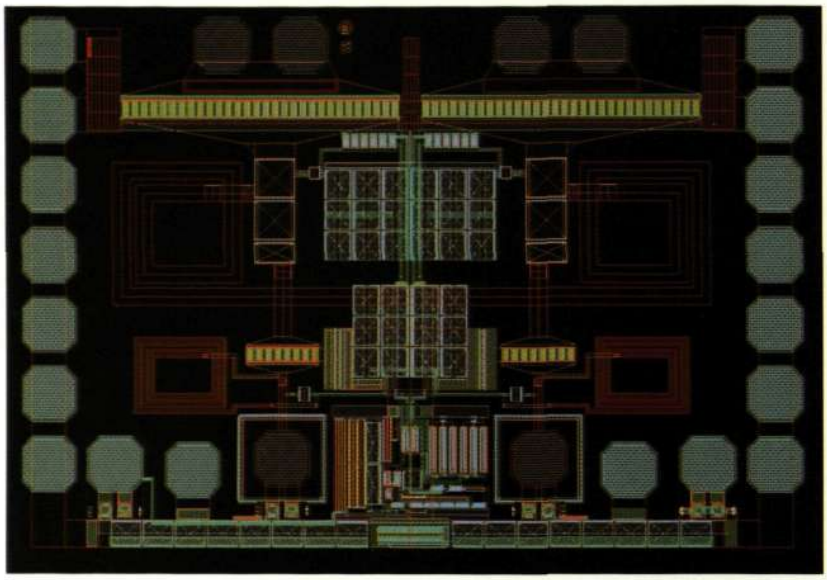

Figure 3
The PBL 403 09 DECT power amplifier chip. Actual chip size is 1.3 x 1.0 mm.

strate contact to the right. Since almost all of the total resistance from the front to the back of the wafer is contained in the top layers—now penetrated by the metal-filled trench—a good ground reference plane and low resistance path to the package are obtained within the chip. Each contact occupies very little area (around one square micron), and can be used in large arrays anywhere on unoccupied circuit areas, say, between blocks or under metalization, to ensure excellent grounding for each part of the circuit. Furthermore, the impedance to ground is now mainly resistive, instead of inductive as is the case when many bonding wires are used. The substrate contacts can also be used to improve shielding and reduce crosstalk between different circuit blocks. In the future, this may prove to be crucial to the integration of radio or control circuitry into the power amplifier chip.

But the greatest potential of the new substrate contact was yet to be demonstrated. If the new metal-filled substrate contact, which provides ground connections from the substrate up to the circuit, is as low-resistive as expected, it should even be possible to use it as a current path from the circuit to ground, instead of the many ground bonding wires usually required in a good power amplifier design. To explore this hypothesis, an experiment was undertaken. A chip used for the DECT power amplifier PBL 403 09 (Figure 3), designed with the new substrate contacts from the beginning, was mounted in an open QSOP16 package. A total of 22 bonding wires were used in this design: fourteen for ground, and eight for supply voltage, input/output signals, and so on. The back-side contact from the chip to the package was of good quality and low-ohmic. The circuit was mounted on its evaluation board, and input/output matching was tuned for optimum performance. The ground bonding wires were then carefully removed, one by one, using a stereomicroscope. When the first wires were removed, the RF performance did not change. With all the ground wires removed, RF performance still remained virtually unchanged. This experiment shows that it is possible to operate a power amplifier without any bonding wires to ground (bringing new meaning to the term *wireless*!) using this new trench substrate contact.

By permanently eliminating the bonding wires to ground, considerable pad area is saved (20% of the chip area for DECT power amplifier), and bonding the chip is simplified (eight wires instead of 22 in the same example). It also becomes possible to select a smaller and cheaper package, thus saving costs and space without any degradation of performance.

This new, unique feature of Ericsson Microelectronics' family of power amplifiers for wireless applications has been dubbed Wireless-Trench technology. Several patents already cover the fabrication and use of the metalized trench contacts.

First application: silicon DECT power amplifier

Wireless-Trench technology was first applied to an integrated power output amplifier for the DECT system. DECT is a digital wireless standard for home and office handsets. Using maximum power of around 0.5 W (27 dBm), the standard communicates at 1.9 GHz, transmitting only 1/24 of the time. The power level and general requirements for the power amplifier are simpler to meet as compared with, say, GSM. A DECT power amplifier was therefore a good starting point for testing the potential of the new silicon RF-IC technology; thus it was chosen as the first product in the new family of silicon-integrated power ampli-

Figure 4
P_{out}, gain and PAE vs. P_{in} at V_{cc} = 3.6 V for DECT power amplifier.

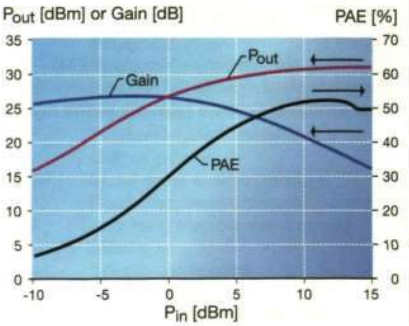

fiers for low-voltage, wireless applications that will be introduced by Ericsson Microelectronics in the near future.

The PBL 403 09 integrated DECT power amplifier, which is housed in a fused QSOP16 package, delivers up to 30 dBm (1 W) of output power at 1900 MHz, with more than 50% power-added efficiency (PAE) The device can be operated up to a 100% duty cycle with minimum performance degradation. Input and output are of the differential type.

Figure 4 shows the input/output characteristics. The two-stage integrated amplifier has on-chip input and inter-stage matching (50 ohms at the input). It also contains biasing circuitry and on/off control signal, and operates using a single 3.6 V supply. It can withstand more than the 5 V supply voltage that occasionally occurs in systems during battery charging, besides withstanding open/short voltage standing-wave ratio (VSWR) conditions to over 5 V with no failures or degradation.

The PBL 403 09 power amplifier is specially designed to interface easily with the PBL 402 15 transceiver chip for DECT. The transceiver has a differential output that interfaces with the power amplifier (using only capacitors for DC-blocking purposes); it also delivers the on/off control signal with appropriate timing. Figure 5 shows how easily the power amplifier is interfaced with the transceiver to form the radio part of a DECT handset or base unit.

The PBL 403 09 silicon power amplifier matches the performance of comparable GaAs products—typically, GaAs products show efficiency in the 40% range, with up to 55% for a few high-performance designs using heterojunction-bipolar-transistor (HBT) technology—Si (35%) and SiGe (30% to 40%). The use of conventional silicon technology makes the use of negative bias or battery switches unnecessary, in contrast to most of the GaAs products that are currently available.

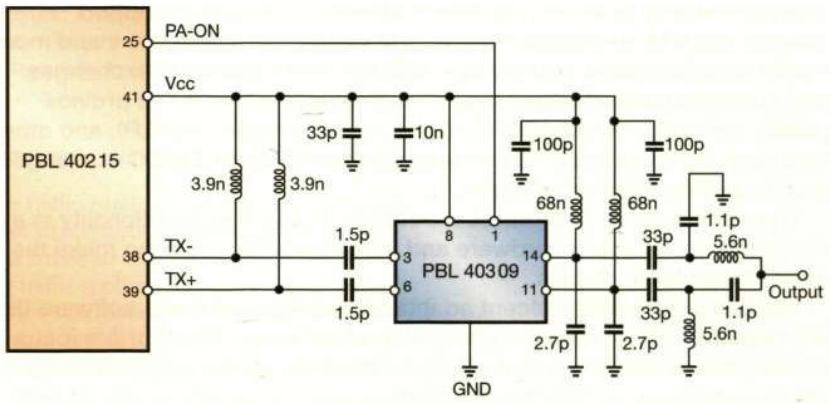

Figure 5
DECT radio with PBL 402 15 transceiver and PBL 403 09 power amplifier.

amplifier for 900 MHz GSM. At 3.4 V supply voltage, the amplifier delivers 35.5 dBm output power with more than 55% efficiency and 31 dB of small-signal gain. The power amplifier, which is a two-stage integrated amplifier designed for 2.7 to 5 V single-supply operation, has integrated input and inter-stage matching and analog input for the power ramp. In power-down mode, it consumes less than 10 µA of standby current. In addition, this single-ended amplifier uses Wireless-Trench technology, which means that it can be operated without ground wires, and without limits to its performance. Similar impressive results have been demonstrated with power amplifiers for other wireless systems.

What's next?

Using the same technical approach as was used for the PBL 403 09, Ericsson is currently developing a family of related power amplifier products for GSM 900 and 1800/1900 MHz as well as GSM dual/triple-band, PDC, other cordless systems, and Bluetooth. Figure 6 shows the electrical data for a pre-production sample of the PBL 403 10, a single-ended power

Conclusion

The results obtained for power amplifiers designed with the new Wireless-Trench technology, as presented in this article, prove the capability of the new technology for low-cost, high-performance wireless power applications. Moreover, this marks the entrance of silicon into power amplifiers—an area that has been dominated by GaAs technologies.

Figure 6
P_{out}, gain and PAE vs. P_{in} at V_{cc} = 3.4 V for GSM 900 power amplifier.

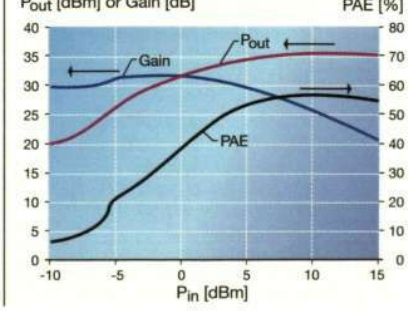

AXE 810—The evolution continues

Magnus Enderin, David LeCorney, Mikael Lindberg and Tomas Lundqvist

Thanks primarily to an architecture that was developed to support change, the AXE exchange continues to evolve. Its architecture and modularity have benefited customers—AXE has served as local exchanges and international exchanges, and even in mobile networks to provide mobile switching centers (MSC), home location registers (HLR), and other functions. This has resulted in a total number of about 13,000 exchanges and an all-time-high growth rate.

The modularity of AXE makes it possible to add new functionality in a cost-effective way, but hardware and software R&D must also make the most of new technologies.

This article describes recent adaptations of hardware and software that will prepare AXE for the next generation of networks. The authors focus on a system architecture that will serve as the basis for migration toward a server-gateway architecture for third-generation mobile networks and the next generation of multiservice networks. Adaptations will also enable improvements in existing networks where traffic is growing quickly.

Introduction

In the next few years, networks will evolve toward today's vision of an "all-IP" network. We can see two possible scenarios: one for traditional networks, and one for "next-generation" networks. For traditional networks, the scenario describes what will happen in multiservice networks and in second-generation mobile networks, such as GSM. The next-generation network scenario describes the development of mobile Internet and fast Internet access in fixed networks.

In traditional networks, evolution is driven by never-ending growth in the need for processing capacity; in mobile networks, by growth in the number of subscribers and usage levels per subscriber. The wireline network is also experiencing a sharp increase in traffic because of Internet "surfing."

In next-generation networks, traditional telephone and data networks will converge to become general networks designed for many different services and modes of access. The convergence of data and telecommunications makes it possible to combine the best of two worlds. Some requirements, such as the need for heightened performance, are fulfilled more easily when development is based on data communications products. Also, the variety of access modes—via second- and third-generation mobile networks, the multiservice networks, and broadband—will necessitate the coexistence of different transmission formats. Thus, gateways will be required at network interconnection points.[1]

BOX A, TERMS AND ABBREVIATIONS

AAL	ATM adaptation layer
ACS	Adjunct computer subsystem
ALI	ATM link interface
AM	Application module
AP	Adjunct processor
APC	AM protocol carrier
APIO	AXE central processor IO
APSI	Application program service interface
ASIC	Application-specific integrated circuit
ATM	Asynchronous transfer mode
BICC	Bearer-independent call control
BIST	Built-in self-test
BSC	Base station controller
C7	CCITT (now ITU-T) no. 7, a common-channel signaling system
CAS	Channel-associated signaling
CP	Central processor
cPCI	Compact peripheral component interconnect
CPP	Cello packet platform
DL	Digital link
DLEB	Digital link multiplexer board in the GEM
DSA	Dynamic size alteration
DTI	Data tranmission interworking
ECP	Echo canceller in pool
ET	Exchange terminal
ETSI	European Telecommunications Standards Institute
FTP	File transfer protocol
GCP	Gateway control protocol
GDDM	Generic datacom device magazine (subrack)
GDM	Generic device magazine (subrack)
GEM	Generic Ericsson magazine (subrack)
GS	Group switch
GSM	Global system for mobile communication
GSS	Group switch subsystem
HDLC	High-level data-link control
HSB	Hot standby
IO	Input-output
IOG	Input-output group
IP	Internet protocol
IPN	Interplatform network
ISDN	Integrated services digital network
ISP	Internet service provider
IWU	Interworking unit
MGW	Media gateway
MIPS	Million instructions per second
MSC	Mobile switching center
MSCS	Microsoft cluster server
MSP	Multiplex section protection
MTP	Message transfer part
MUP	Multiple position (timeslot)
NGS	Next-generation switch
OSS	Operations support system
PCM	Pulse code modulation
PDH	Plesiochronous digital hierarchy
PLMN	Public land mobile network
PSTN	Public switched telephone network
PVC	Permanent virtual circuit
RAID	Redundant array of independent disks
RAM	Random access memory
RISC	Reduced instruction set computer
RLSES	Resource layer service specification
RM	Resource module
RMP	Resource module platform
RP	Regional processor
RPC	Remote procedure call
RPP	Regional processor with PCI interface
SCB-RP	Support and connection board - RP
SDH	Synchronous digital hierarchy
SES	Service specification
SONET	Synchronous optical network
SS7	Signaling system no. 7
STM	Synchronous transfer mode
STOC	Signaling terminal open communication
TCP	Transmission control protocol
TDMA	Time-division multiple access
TRA	Transcoder
TRC	Transceiver controller
TSP	The server platform
TU 11, 12	Typical urban 11 (12) km/hr
UMTS	Universal mobile telecommunications system
VCI	Virtual channel identifier
VPI	Virtual path identifier
XDB	Switch distributed board
XSS	Existing source system
WCDMA	Wideband code-division multiple access

Typical of next-generation network architecture is the separation of connectivity and communication or control services. This new new architecture will appear first in GSM and WCDMA/UMTS core networks, where AXE exchanges will continue to serve as MSCs. For multiservice networks, telephony servers will become hybrid nodes which consist of an AXE exchange that uses an AXD 301 to handle packet-switched data.

Based on the traditional and next-generation scenarios, network evolution will demand increased processing capacity, greater switching capacity, and conversion to packet-switched technology. In addition, new ways of doing business will emerge as virtual operators pay for the right to use telecom equipment owned by other operators. Similarly, more operators are expected to enter the market, putting additional demands on charging and accounting functions.

To succeed in today's telecommunications market, operators must be able to provide their users with new functionality in networks and complete coverage by mobile systems at the same or a lower cost as time progresses. Operators put demands on the return on investment, cost of ownership, footprint (multiple nodes per site), plug-and-play functionality (short lead times in the field), and quality of software packages (quality of service).

This article describes the latest developments made in the AXE platform to meet these requirements, and what new products will be launched. For example, to shorten the software development cycle, a layered architecture was introduced in the early 1990s, resulting in application modules (AM). Current work on application modules focuses on the server-gateway split. AXE hardware is also undergoing a major architectural transformation to further reduce the footprint, power consumption, and number of board types.

Ericsson's goal is to cut the time to customer for AXE by targeting improvements in production, transportation, and installation. Far-reaching rationalization has been achieved through the introduction of the generic Ericsson magazine (GEM), an open, flexible, scalable, high-capacity magazine (subrack). A new distributed, non-blocking group switch (GS890) has also been introduced, as have new devices, such as the ATM link interface (ALI) and ET155-1, which enable AXE to serve as a gateway to an ATM network. Also, a new input-output (IO) system, called the APG40, has been developed using products from mainstream data communications vendors.

The application platform

The traffic-handling part of AXE has a two-layer architecture: the application platform, which consists of hardware and software; and the traffic applications, which solely consist of software. The application platform can be seen in terms of hardware and software, which present different but complementary views of the architecture.

The software view

The AXE application software has continued to evolve since it was first layered and modularized in the early 1990s. Market demand for shorter lead-times was the main force driving the change to layers and modularization.

Application modules

Traffic applications are encapsulated in application modules that use the application-platform software—called the resource module platform (RMP). There is no direct communication between application modules. All communication between them takes place via protocols, where the protocol carriers are implemented in the resource module platform (Figure 1).

Figure 1
AXE 810 application software architecture.

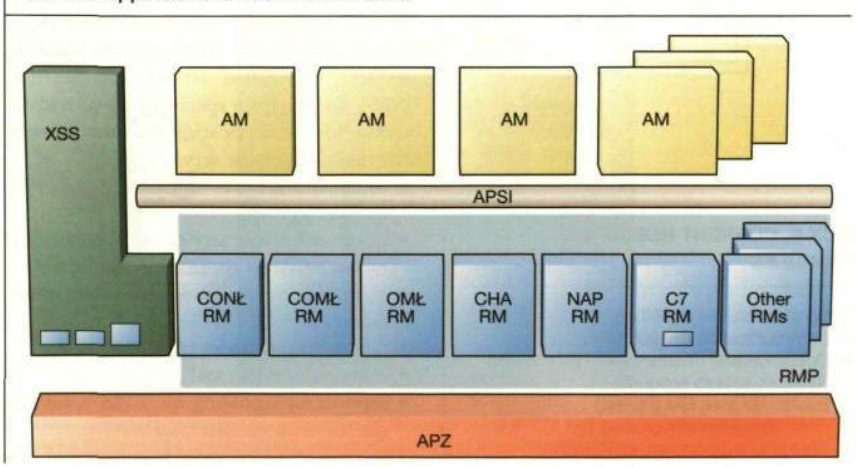

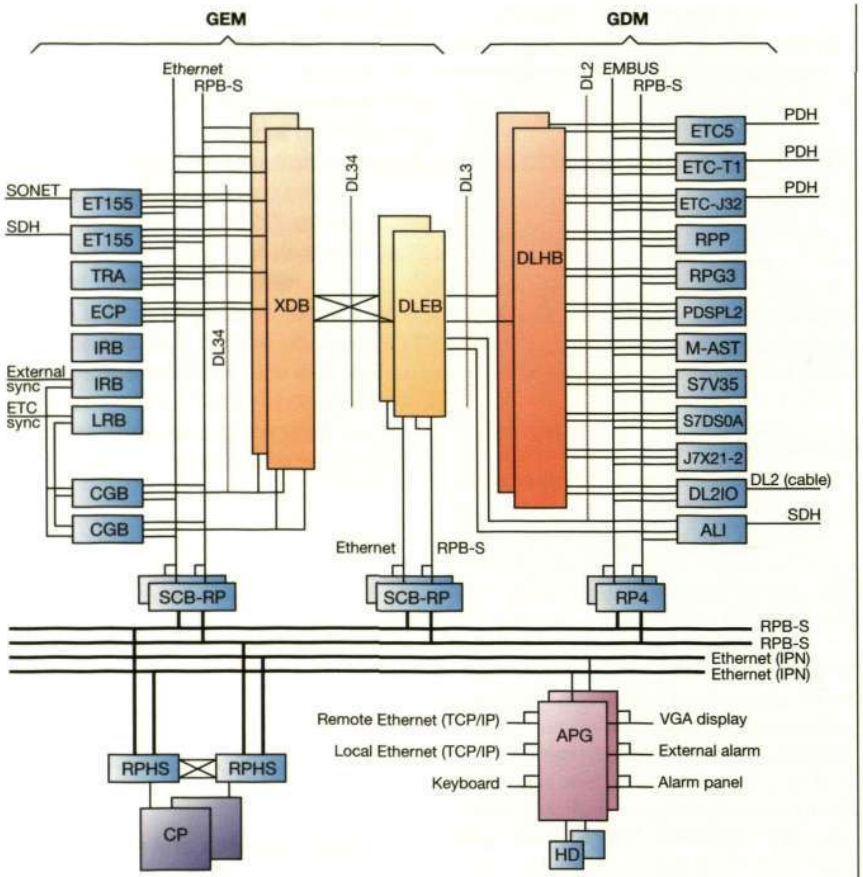

Figure 2
Hardware architectural overview.

The interface of the resource module platform with the application modules, called the application platform service interface (APSI), is a formalized client-server interface. Subsets of the APSI, or individual interfaces, are called service specifications (SES). Each service specification provides a service for the application modules, some of the most important services being:

- connection service—for setting up a physical connection;
- communication service—for communication between application modules;
- object manager—for operation and maintenance;
- AM protocol carrier (APC);
- charging service; and
- services for signaling no. 7 (SS7), such as the message transfer part service.

The total concept, which consists of the RMP, APSI, the AMs, and the inter-AM protocols, is known as the application modularity concept. It is largely thanks to this concept that the system lifetime of AXE consistently exceeds expectations—the AXE system is 30 years old, yet we see no end to its potential evolution.

Resource modules

When the RMP was first introduced, it was small in comparison to the large volume of existing software, called existing source system (XSS). Following several RMP development projects, the migration from the earlier architecture to the AM architecture is now virtually complete. The resource module platform, in turn, has become large and complex, leading to a refinement of the AM concept and the division of the RMP into several resource modules (RM). The interfaces provided by the resource modules are formalized as resource-layer service specifications (RLSES).

A priority in the development of RMs and AMs (known collectively as system modules) is to specify the interfaces (SES and RLSES) early in the development process. When the interfaces are "frozen", the separate system modules can be designed independently of one another, often at different geographical locations, as is common in the Ericsson organization.

The second major advantage for development of applications is that application-platform hardware is now associated with specific resource modules and controlled by them. Application modules simply request services via the APSI. Hardware is "owned" by the software in the respective resource modules.

Hardware

The hardware (Figure 2) revolves around the GS890 group switch, which has 512 K multiple positions, each with a bit rate of 64 kbit/s.

The biggest change in the hardware architecture is the addition of a new exchange interface (the DL34) to existing DL3 and DL2 interfaces. The DL34 interface supports from 128 to 2,688 timeslots to each device board, in steps of 128, via a backplane interface in the generic Ericsson magazine running at 222 Mbit/s. This interface made it possible to improve the devices connected to the group switch, further reducing input power, cabling, the number of board types, footprints, installation time, and other parameters. The high-speed DL34 in-

BOX B, CURRENT RESOURCE MODULES

- Connection RM (CONRM)
- Communication RM (COMRM)
- Operation and maintenance RM (OMRM)
- Network access products RM (NAPRM)
- Pooled devices RM (PDRM)
- Common channel signaling RM (C7RM)
- Charging support RM (CHARM)

terface has also facilitated a more efficient version of the group switch itself, also located in the GEM.

Another architectural change (Figure 2) is the interplatform network (IPN). In the first phase, the IPN will be used to speed up communication between the central processor (CP) and adjunct processors (AP). The interface is based on fast Ethernet. In a second phase, it will be upgraded to Gigabit speed.

All devices that support the DL34 interface can be mixed in the GEM more or less without limit. The maximum capacity of each GEM is 16 K multiple-position timeslots (MUP), which corresponds to eight STM-1 ET boards, for example. Physically, there is space for 22 device boards in one GEM.

If a switch is needed that is larger than 16 K MUPs, or when the number of devices exceeds 22, additional GEMs must be added. These can be configured without interruption while the system is processing traffic. An exchange based on the GEM is linearly expandable. The maximum switch size is 512 K MUPs at normal rate (64 kbit/s) or 128 K MUPs at a subordinate rate (n x 8 kbit/s).

By inserting a pair of digital link multiplexer boards (DLEB) in the GEM, we can convert the DL34 backplane interface into four DL3 cable interfaces. This allows all of today's devices in the generic device magazine (GDM), generic datacom device magazine (GDDM), and other formats to be used with the new switch core.

Software evolution

The original AXE concept has enabled ongoing evolution and modernization. Legacy software can be reused or modified, subsystems and system modules can be modernized or completely replaced, interfaces can be left unchanged or redesigned to meet new needs, and new architectures can be introduced. Yet AXE remains fundamentally AXE.

Several software development programs are under way in parallel. The APZ processor is being adapted to serve as a multiprocessor CP to multiply call-handling capacity. Cost of ownership is being reduced. Ericsson is improving its internal management of software to reduce time-to-market and time-to-customer. These activities are in addition to the four programs described below.

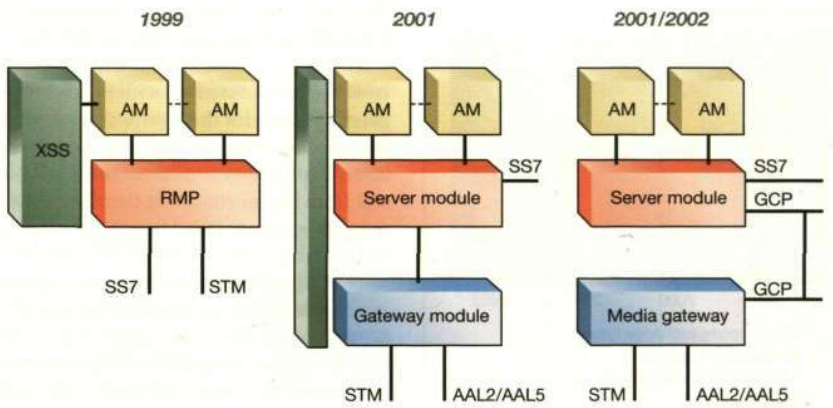

Figure 3
Software evolution toward the server-gateway split.

Software migration

The network architecture is changing. In particular, we see two new types of node: the server (or media-gateway controller) and the media gateway itself, resulting in what is known as the "server-gateway split." In order to meet the demands of the new architecture, the software in the application platform is being migrated from a traditional telecom environment—which is STM-based, mainly for voice traffic—to-

Figure 4
System architecture, ATM interworking function.

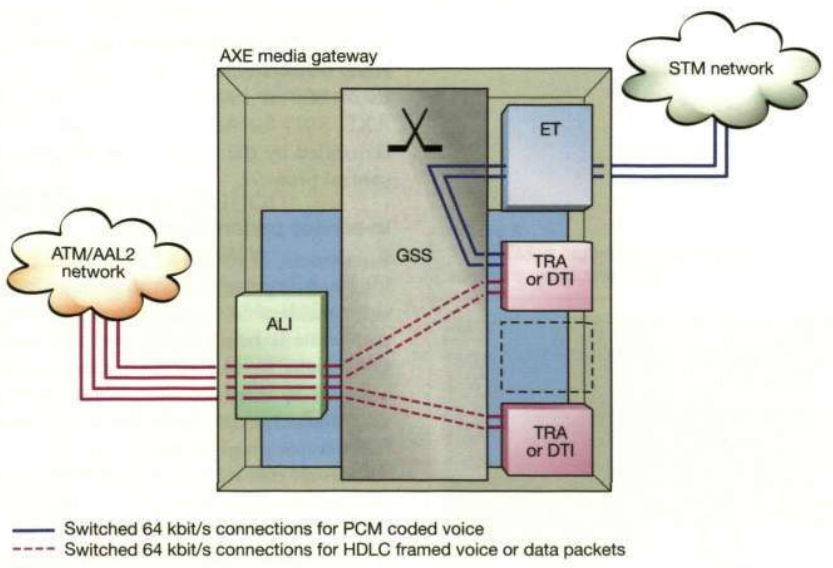

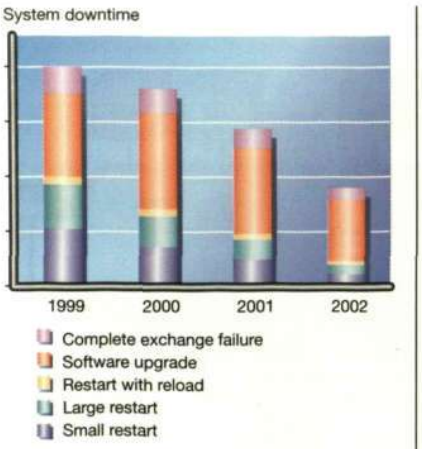

Figure 5
Reduction in system downtime.

ward a packet-based environment, initially for ATM, and subsequently for IP.

In this new architecture, the server controls the call or session, whereas the media gateway provides the bearer service, transmitting information over the connectivity layer. An advantage of separating call control from bearer control is that bearer technology can be upgraded from STM via ATM to IP without affecting the call control.

The original AM architecture has proved itself to be future-proof and is well suited to accommodate the server-gateway split. The split has been implemented in system modules (new and modified AMs and RMs) without any fundamental changes to the software architecture. For example, new RMs for bearer-access service and for TDM access are currently being developed for the WCDMA/UMTS program.

Two separate migration tracks are currently being followed. For the WCDMA/UMTS program, AXE is the natural choice of technology for the server, thanks to its high capacity and high level of functionality. AXE is also available as a combined server/media-gateway node, where the server part controls its own gateway part. By contrast, the Cello packet platform (CPP) is the preferred choice of technology for ATM-based media gateways. In such a network, the server nodes communicate using bearer-independent call control (BICC). The server controls the gateway using the gateway control protocol (GCP). ATM-based gateways communicate using Q.AAL2.

The second migration track is for the next-generation switch (NGS) program for the multiservice network. The server node is a hybrid node made up of co-located AXE and AXD 301 systems. The gateway node consists of an ATM switch (also an AXD 301) for AAL1 connections, and is controlled by the server using the gateway control protocol.

In-service performance

Robustness, or in-service performance, has long been a priority in AXE development, and considerable improvements have been and are still being made. No software is 100% reliable, but robust systems can be achieved by means of recovery mechanisms that minimize the effects of software malfunctions or program exceptions.

One of the most powerful robustness mechanisms is known as *forlopp*, which is a recovery mechanism at the transaction level. (Forlopp is an anglicization of the Swedish word *förlopp*, roughly the "sequence of events.") The forlopp mechanism is a low-level recovery mechanism by which only the transaction affected by the software malfunction is released. This minimizes the disturbance to the overall system. The forlopp mechanism, which has been refined in several stages and is now applied to all traffic handling in AXE and to many operation and maintenance functions, has significantly reduced the number of recoveries at the system level.

System restart is used as a recovery mechanism for certain faults and for system upgrades. The restart mechanism itself has been optimized in several ways, so that in the few cases that still require restart, its duration is as short as possible.

Activities for restarting software in the regional processor (RP) have been especially improved. For example, regional processors are restarted through minimal restart by default—that is, with the suppression of unnecessary actions. Complete restarts are performed on regional processors only when necessary. Regional processors are restarted asynchronously, which means that no regional processor has to wait for any other regional processor to become ready for a restart phase. These improvements have significantly reduced the time consumed when restarting application hardware.

If necessary, AXE can, via the IPN, be reloaded from the backup copy of the system on the APG40 file system. This approach is ten times faster than the design that applied before the APG40 and IPN were introduced. Also, the time it takes to make a backup copy of the system on APG40 is one-fifth of what it was before.

The restart duration for a system that is upgrading to new software has also been reduced. The most recent improvements involve RP software, which is now preloaded prior to an upgrade, instead of during the upgrade, thus saving restart time.

Major improvements have been made to the function-change mechanism used for system upgrades. Inside the central processor, the time needed for data transfer prior to a changeover to new software has been cut from hours to typically ten minutes. (The data transfer does not disturb traffic handling, but exchange data cannot be changed during data transfer.)

Another improvement applies to the retrofit of new CPs that replace old ones. The bandwidth of the connection between the processors has been expanded by using re-

gional processor with industry-standard PCI interface (RPP) and Ethernet connections, thereby reducing data-transfer times ten fold. As a consequence of these many improvements, system downtime has been reduced significantly in recent years, a trend that is sure to continue (Figure 5).

Software licensing

The main force driving software licensing is customer preferences to pay for access to specific functions, features, and capacity. Thus, "software licensing" is really the "licensing of features and capacity." The second driving force is Ericsson itself, because it is in the company's interest to deliver standard nodes (those that contain standardized software and hardware).

With software licensing, Ericsson delivers a standard node for a particular market or market segment, such as an MSC for a cellular system using time-division multiple access (TDMA). This standard node consists of a complete software configuration with a standard hardware configuration that is deliberately over-dimensioned. Ericsson then limits the call capacity and functionality of the node by means of software keys. When a customer requests more call capacity, Ericsson personnel execute a password-protected command to increase the capacity in the node to the new level.

This method of increasing capacity, referred to as "traffic-based pricing" or "pay-as-you-grow," is much simpler than the traditional method, in which Ericsson personnel would deliver, install, and test new hardware on a live node. The commands can also be executed remotely from a maintenance center, making a visit to the site unnecessary. This arrangement benefits customers and Ericsson, especially in the mobile market where growth can be rapid and unpredictable.

Similarly, in the future, software will contain functions and features for which customers can purchase licenses on an *ad hoc* basis. The software for these functions and features will be unlocked using password-protected commands. This method of delivering software is much simpler than the traditional method, by which new software is delivered in the form of software upgrades.

The licensing of call capacity is already available on AXE nodes, and the introduction of a general mechanism to handle all software licensing of functionality and capacity is planned. Some of the characteristics of this general mechanism are as follows:

- Software licensing will be common to AXE 810, TSP, CPP, and other Ericsson platforms.
- Ericsson will maintain a central license register, from which new licenses can easily be issued.
- License keys will be distributed electronically, in encrypted form.

Improved handling

Several improvements to the operational handling of AXE have recently been made, such as parameter handling and hardware inventory. Other handling improvements are being planned, such as plug-and-play functionality. These improvements reduce operational costs and also often have a positive impact on the ISP, since they reduce the risk of human error.

One of the most important improvements to handling is called dynamic size alteration (DSA). The size of many data records in AXE is traffic-related, since the number of individuals, or instances, in the record varies from node to node, as well as over time, in accordance with the capacity demands put on the node. Thanks to dynamic size alteration, the number of individuals in the record is increased automatically, up to a preset limit, without any intervention from the maintenance technician. However, as a warning, AXE issues an alarm when, say, 75% of the reserved data capacity has been used up, so that the technician can raise the limit as necessary.

Dynamic size alteration also simplifies the handling of this kind of alteration, since the alarm (referred to above) specifies an action list of the data records whose size needs to be altered (increased). The technician merely enters a single command to increase the sizes of all these data records.

Hardware evolution

Generic Ericsson magazine

Another far-reaching improvement is the GEM, a high-capacity, flexible, and scalable magazine (subrack) that anticipates future developments (Figure 6). GEM-based nodes will be smaller, dissipate less power, and have greater maximum capacity. Their implementation will dramatically improve the cost of ownership and cut time-to-customer for AXE.

In previous versions of AXE, each function was located in a separate magazine, and

Figure 6
The GEM, a high-capacity, flexible and scalable magazine (subrack).

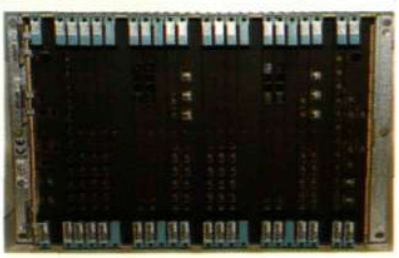

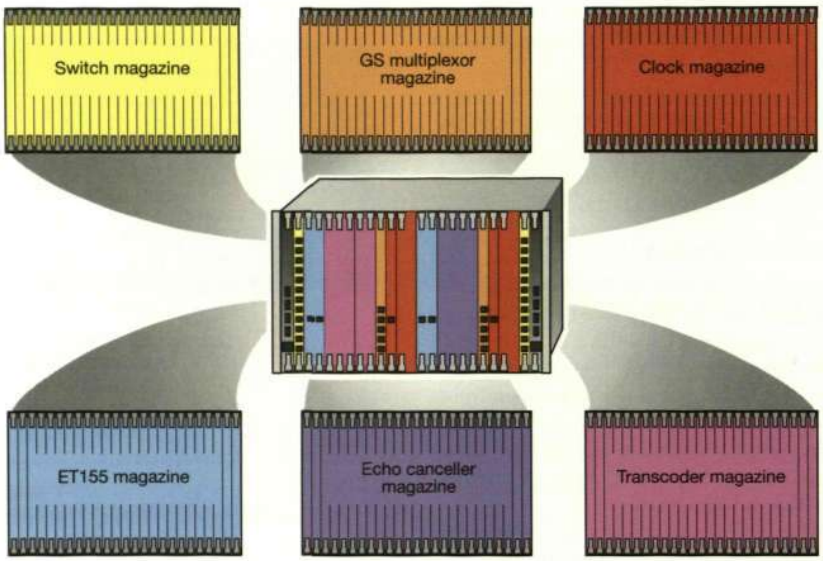

Figure 7
Many different functions can be mixed in the GEM.

the associated regional processors were on separate boards in the magazine. The GEM and the boards developed for the GEM represent fundamental change.

Several functions are mixed in a single magazine, and the RP function is integrated on each board. This makes the most of advances in technology to modify the architecture rather than merely shrink the hardware. One advantage of this change is that considerably fewer magazine and board types are needed in each node. A second advantage is that many internal AXE interfaces have been moved to the GEM backplane, reducing the need for cables. As shown in Figure 7, many different functions can be mixed in the GEM.

Boards developed for the GEM have high capacity, and because more functions are concentrated on each board, the total number of boards can be reduced. Mixing these high-capacity boards in a single magazine and removing many cables by moving the interfaces to the GEM backplane produces a more compact system.

The GEM concept can be used to build anything from an extremely small node to an extremely large one, using the same boards and magazine (Figure 8). The smallest node will consist of one GEM, and the largest, 32 GEMs. As capacity requirements grow, the node can be expanded smoothly and without interruption of traffic by adding one or more GEMs with a suitable combination of devices. A GEM can house

- two SCB-RPs, providing an Ethernet switch, control function, maintenance support, and power distribution;
- two switch boards (XDB), providing a 16 K plane duplicated group switch;
- up to 22 device boards with 15 mm spacing, such as ET155, ECP, or TRA;
- pairs of DLEBs, providing multiplexing functionality for DL3 cable interfaces, placed in device slots; and
- CL890 clock modules placed in device slots.

The GEM, which has been developed for use in the BYB 501 equipment practice, provides several infrastructure functions via the backplane for the boards housed in the magazine:

- duplicated power distribution;
- duplicated group switch connection;
- duplicated serial RP bus;
- maintenance bus for hardware inventory and fault indication;
- duplicated Ethernet, 10 or 100 Mbit/s; and
- extremely robust backplane connectors.

The option to use Ethernet as the control protocol makes the GEM an open magazine prepared for future developments. Another advantage of the GEM is that it is generic and can be used in non-AXE products from Ericsson.

Figure 8
The GEM concept can be used to build anything from an extremely small node to an extremely large one, using the same boards and magazine.

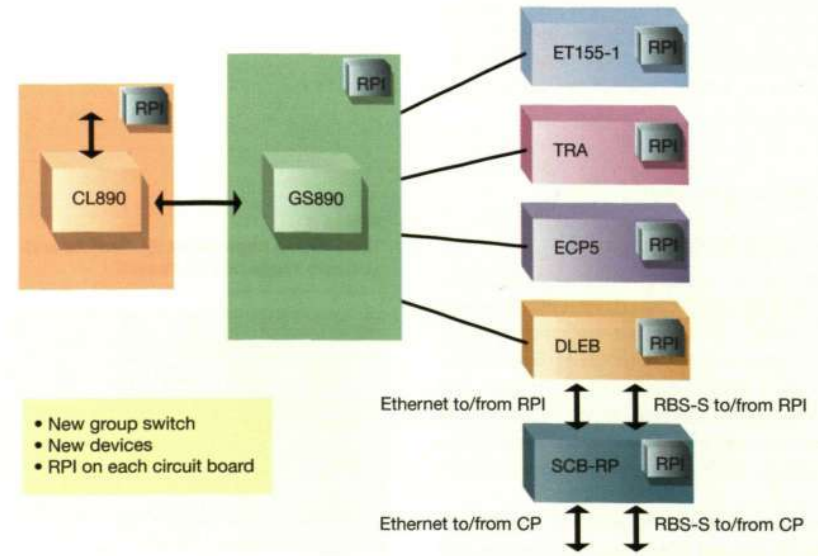

The GS890 group switch

The GS890 fully supports the GEM concept and its goal of dramatically improving cost of ownership and time-to-customer. The figures in Table 1 show an extraordinary reduction in power dissipation and number of boards and cables. This was achieved through a combination of recent advances in ASIC technology and high-speed interfaces with architectural modifications.

The maximum size of the GS890 has been increased considerably, making it possible to build larger nodes. For many network configurations, it would be better to build and maintain a small number of large nodes. The switch is also strictly non-blocking, so many limitations on the configuration of the node and the surrounding network have been removed.

The subrate switch function in AXE, which enables switching at 8 kbit/s, has also been modified. Primarily used in GSM networks, this function makes efficient use of pooled transcoders and transmission resources surrounding the base station controller (BSC). The subrate function was previously implemented as a pooled function, but in GS890 (Box C) it has been integrated into the switch core and increased to 128 K. The advantages are that much larger BSCs are feasible, less equipment is needed, and the traffic selection algorithms are easier and faster because the pooled switch has been removed.

GS890 group switch connection

All existing group switch interfaces are supported, making it possible to connect all current devices to the GS890. The DL2, which is a cable and backplane interface, carries 32 timeslots. Similarly, the DL3, which is a cable interface used by the ATM interworking unit (IWU) and GDM, carries 512 timeslots. The new DL34 backplane interface supports new high-capacity devices developed for the GEM.

TABLE 1, IMPROVEMENTS IN THE GROUP SWITCH

Characteristics	GS12 (Max 128 K)	GS890 (Max 512 K)	Reduction
Power dissipation in 128 K MUP group switch core	1200 W	250 W	81%
Internal cables in a 128 K MUP group switch core, and including clock distribution	1,024 (12-pair cables) and 48 (4-pair cables)	88 (4-pair cables)	92%
Number of boards in a 128 K MUP group switch core, excluding clock boards	320	16	95%
Equivalent figures for a 512 K group switch	--	1,000 W, 448 (4-pair cables) and 64 boards	--

GS890 hardware

The main functionality of the GS890 is implemented in two new ASICs: one core ASIC, which handles the switching; and one multiplexer ASIC, which concentrates traffic into the core. Both ASICs have been designed using 0.18 µm CMOS technology.

The core ASIC contains 4 Mbit of RAM, runs at up to 333 MHz, has 500 K gate logic and 20 high-speed interfaces running at 666 Mbit/s, allowing the ASIC to handle up to 13 Gigabits of data every second. In a fully expanded switch matrix, 64 of these ASICs will be needed per plane. Standard four-pair cable is used to interconnect the switch matrix elements located in different magazines. The cable can be up to 20 meters long and carries 1.3 Gigabits of bidirectional data per second, two 666 Mbit/s interfaces in each

BOX C, MAIN CHARACTERISTICS OF THE GS890

- Maximum size 512 K MUP, equivalent to 524,288 (64 kbit/s) ports
- Strictly non-blocking architecture, regardless of traffic type
- Fully distributed group switch, with switch matrix distributed among up to 32 GEMs
- Integrated subrate switching capability up to 128 K MUP, equivalent to 1,048,576 (8 kbit/s) ports
- Improved maintenance support by logic BIST and RAM BIST in all ASICs, reducing maintenance time and the time that the system runs on one plane
- Hardware support for fast dynamic fault isolation
- New and flexible high-speed-device interface
- Device protection support with no wasted capacity (used for ET155-1)
- New high-speed, internal group switch interface
- Rear cabling removed

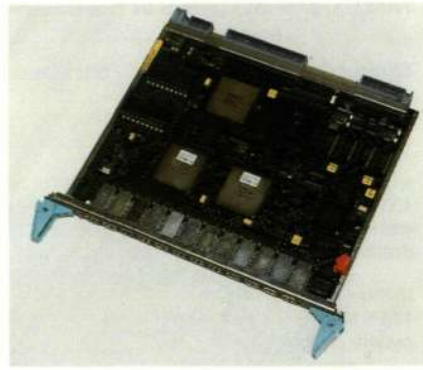

Figure 9
Photograph of the switch distributed board (XDB).

direction. The new hardware requires two new board types:

- the switch board; and
- the multiplexing board, which makes it possible to connect existing devices.

In addition, two more boards have been developed to support a smooth migration from the GS12 to the GS890. The migration consists of replacing one board in each GS magazine, thereby converting it into a multiplexing magazine. The concentrated traffic is then connected to the new GS890 core. These boards help safeguard investments already made in AXE. The GS890 (Figure 9) supports the following types of traffic:

- normal rate, 64 kbit/s;
- non-contiguous wideband, *n* x 64 kbit/s up to 2 Mbit/s;
- subrate, 8 kbit/s or 16 kbit/s;
- wideband on subrate, *n* x 8 kbit/s up to 256 kbit/s; and
- broadband connections up to 8 Mbit/s.

The CL890 synchronization equipment, also developed to support the GEM concept, consists of

- duplicated clock modules;
- up to two highly accurate reference clock modules; and
- up to two incoming clock reference boards for connecting additional clock references.

Devices

RPG3

A new version of the RPG2 with slightly improved CPU performance has been designed that only occupies one slot instead of two. The RPG3, which is still located in the GDM, is mainly used for different kinds of signaling application.

ET155-1

The ET155-1 is an SDH exchange terminal that will occupy one slot in the GEM and will be available for ETSI and ANSI/SONET standards. The ETSI variant will carry up to 63 TU12s and support both optical and electrical line interfaces. In ANSI/SONET mode, the board will be able to handle up to 84 TU11s.

The board fully supports channel-associated signaling (CAS) in both standards by extracting the signaling information from the incoming PDH frames and sending it through the switch to pooled signaling terminals. The ET155-1 can work either as a non-redundant single-board exchange terminal or in tandem, supporting MSP 1+1 and equipment protection.

One motherboard, which contains the main functionality, and two generic line interface modules, containing the optical respectively electrical line interface, have been developed. An ET155-1 board is produced by mounting the required module on the motherboard. The line interface modules are generic and will also be used in non-AXE products from Ericsson.

Transcoders and echo cancellers

A big improvement has been made in the transcoder area. A new board designed for the GEM supports 192 channels for all the codecs used in GSM and TDMA systems. The same board is also used for echo cancellation, in which case it supports 128 channels.

Figure 10
Block diagram of the ATM link interface (ALI).

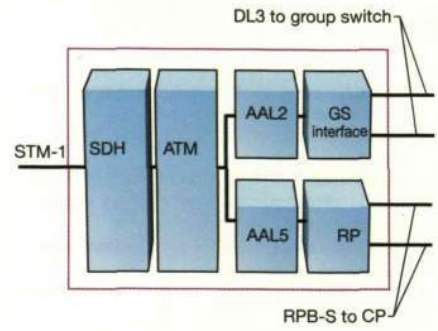

A fully equipped GEM with transcoders supports 4,224 transcoder channels; when fully equipped with echo cancellers, the GEM handles 2,816 channels.

Future GEM devices

The devices described here were designed for the first release of the GEM. Future releases will include more devices ported from today's GDM and GDDM. Technology allows more traffic to be handled by each board in AXE, so the boards can efficiently take advantage of the configurable bandwidth of the DL34.

ATM interworking function

With the introduction of WCDMA/UMTS, AXE will be used as a combined server and media gateway and purely as a server. In the downlink direction (to the radio access network), AXE will be attached to the *Iu* interface, a new interface based on ATM. ATM will also be used in the core networks.

An ATM interworking unit has been implemented to add an ATM interface to new and legacy AXE nodes. The interworking unit enables existing AXE-based MSCs to be upgraded easily to handle signaling and the user plane in ATM networks. The interworking unit is connected directly to the group switch interface for the user plane, and the signaling information is transferred to the central processor via the RP bus (Figure 10).

The interworking unit provides a 155 Mbit/s ATM/SDH interface to AXE. New hardware and software for AXE enable it to perform as a gateway to an ATM network capable of handling voice, data, and signaling. The first release of the ATM interworking unit supports three types of payload in the ATM network:

- coded voice on AAL2;
- circuit-switched data on AAL2; and
- signaling on AAL5.

The device that implements the interworking unit is called the ATM link interface (ALI). This hardware unit is divided into three boards that are mounted together on a plug-in unit which can be located in a full-height GDM that provides RP bus, –48 V and maintenance bus via the backplane. The unit is then connected via DL3 cables to the group switch. The fiber for the STM-1 interface is connected to the external equipment.

Adjunct processor platform

Statistics from the telephony or radio parts of a network must be collected to optimize network efficiency. At the same time, virtual operators are willing to pay for the right to use telecom equipment owned by other operators, and network operators are demanding fast, almost instant billing (such as with prepaid cards). These requirements demand processes with close to real-time execution.

In terms of operation and maintenance, it is better to have one big node in a network instead of several small ones, although this demands greater switching capacity. To reduce operators' cost per subscriber, big switches are being planned. This has spurred the development of new high-end AXE switches. There are numerous methods of

BOX D, ATM VIA AXE

To transmit coded voice over ATM, a permanent virtual circuit (PVC) is set up between the ALI and the remote end. The PVC is specified by its VPI/VCI address. The first ALI release allows a total of eight PVCs to be set up in the user plane.

Within each PVC, AAL2 connections (based on the I.363.2 standard) can be set up and released using Q.2630.1 (Q.AAL2) signaling. Up to 248 AAL2 connections can be set up in each PVC.

The C7 signaling links for supporting Q.AAL2 and other user parts are configured as another type of PVC. For each ALI, up to 64 signaling PVCs can be configured to carry adaptation layer AAL5, based on the I.363.5 standard. Each of these AAL5 channels carries a C7 signaling link using SSCOP (Q.2110) as the link protocol.

Adaptation layers for both ETSI (SSCF-NNI, Q.2140) and Japanese C7 standards (JT-Q.2140) are provided.

The architecture for supporting circuit-switched data is similar to that for coded voice except that the payload packets are routed to a data transmission interworking (DTI) unit instead of a transcoder. The first release of ALI supports a total of 32 data channels.

In the future, the ATM interworking concept might be used to implement new functions such as 64 kbit/s voice on AAL1 or on AAL2 based on I.366.2. Another function in the works is "soft PVCs," which can be operated using B-ISUP signaling.

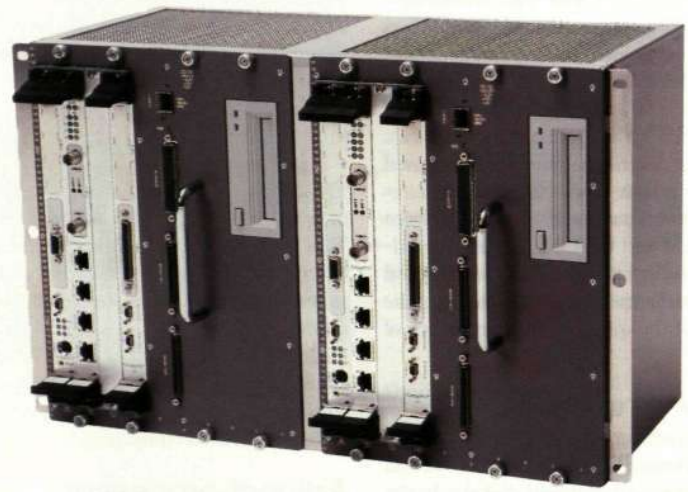

Figure 11
Photograph of a large APG40 cluster.

enhancing performance, since AXE processing power is distributed among several processors (CP, RP, and AP), and research and development can be pursued along different lines in parallel.

APG40 platform for near real-time applications

The most widely deployed AXE IO system, IOG20, was introduced in 1997. The IOG20 is a proprietary product for the CP file system, man-machine communication, boot server, and statistics functionality. In 1998, the first adjunct processor—the APG30—was released, primarily as a billing system. The APG30 is widely deployed in TDMA systems as a platform for billing but also for operation and maintenance of AXE as an IOG replacement. The APG30, which is based on a commercial computer from an external vendor, uses MIPS RISC microprocessors and NonStop-UX O/S.

The APG40, successor to both the IOG20 and APG30, is a platform for near real-time applications. Three needs have driven the development of the APG40 (Figure 11). One need is to offer a platform for new data-handling tasks and applications that are best performed at the network element level by a standard computer system. Another need is for improved performance, as a result of increased traffic per switch and new types of data extracted from each switch. The third need is to reduce time-to-market and time-to-customer.

The APG40 is a platform for the AXE central processor IO functions (APIO) that were inherited from the IOGs. It is also a platform for billing (for example, FOS) and statistic data collecting (for instance, STS), storage, processing and output from the AXE switch. For example, the APG40 can be a platform for collecting data related to in-service performance from the central processor. It can also format that data for distribution to an operations management

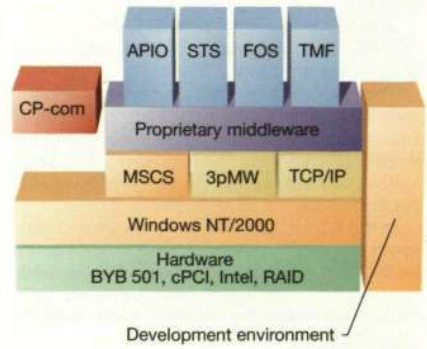

Figure 12
APG40 building blocks and development environment.

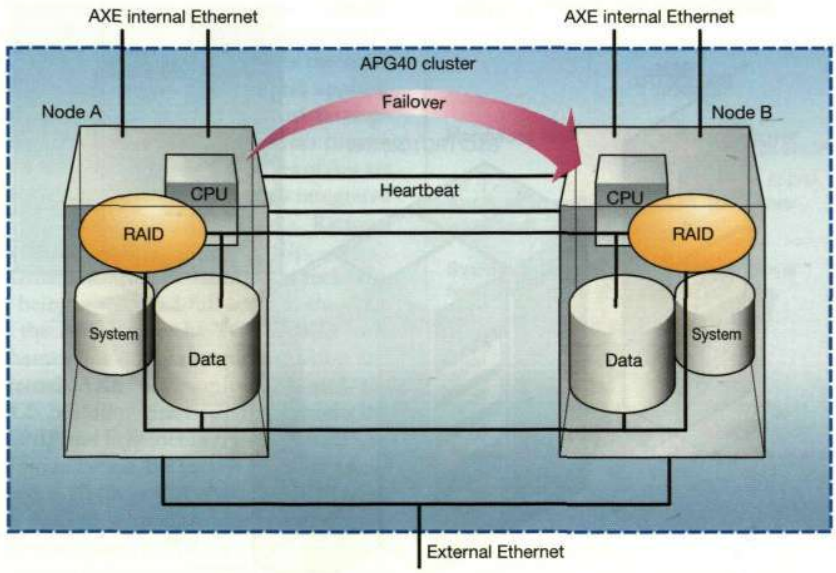

Figure 13
The RAID hardware system.

center, such as the operations support system (OSS), for preventive measures or for sending an alarm when a critical level is reached.

The applications that run on the APG40 do not communicate with one another and are not interdependent. The platform provides basic functionality, such as highly available processing capacity, safe storage, and data output mechanisms.

The adjunct computer subsystem (ACS) (proprietary middleware in Figure 12) provides the software platform for applications in the APG40. The ACS offers Windows NT, Microsoft cluster server (MSCS), and communication protocols, such as FTP, RPC, TCP/IP, and Telnet.

In-service performance

Operators require reliable input-output systems that store AXE system backup copy safely and perform reloads fast. A switch must always be available for operation and management from an operations support system. To ensure the in-service performance of the APG40 itself, and AXE as a whole, concepts introduced in the IOGs and APG30 have been brought to this new platform. In the APG40, mainstream technologies, such as MSCS and RAID hardware, form the basis for telecommunications-grade in-service performance.

The APG40 is a cluster server (Figure 13). In its first release, a cluster will consist of two nodes. Each node has its own memory and runs the operating system. All applications running on an APG40 cluster are active on one node only—the executive node; the other node is a standby node. Thus, if a major fault occurs, the standby node starts and takes over the application. A watchdog in the MSCS ensures that the adjunct processor software does not simply stop working. When an error is detected, either in hardware or software, the cluster server can restart the failed resources in the other node. This process is called *failover*.

The physical size of the APG40 is a half magazine for a small cluster and one full magazine for a large cluster. The small APG40 has one 18 GB disk per node, while the large APG40 can house up to three 18 GB disks per node. Later on, it will be possible to introduce disks with a capacity of 36, 72, or 144 GB.

Thanks to the RAID hardware system (Figure 13), the APG40 has a higher disk-integration speed than the IOG20. The RAID hardware system is used for safe storage. Only the executive node can write data to the data disk. To avoid loss of data due to a disk error, the executive node writes data to a Windows NT file system, and the RAID system writes to the data disk in both the executive node and the standby node. The executive node thus owns all data disks in both nodes. In the IOG20, on the other hand, each node owns its own set of data

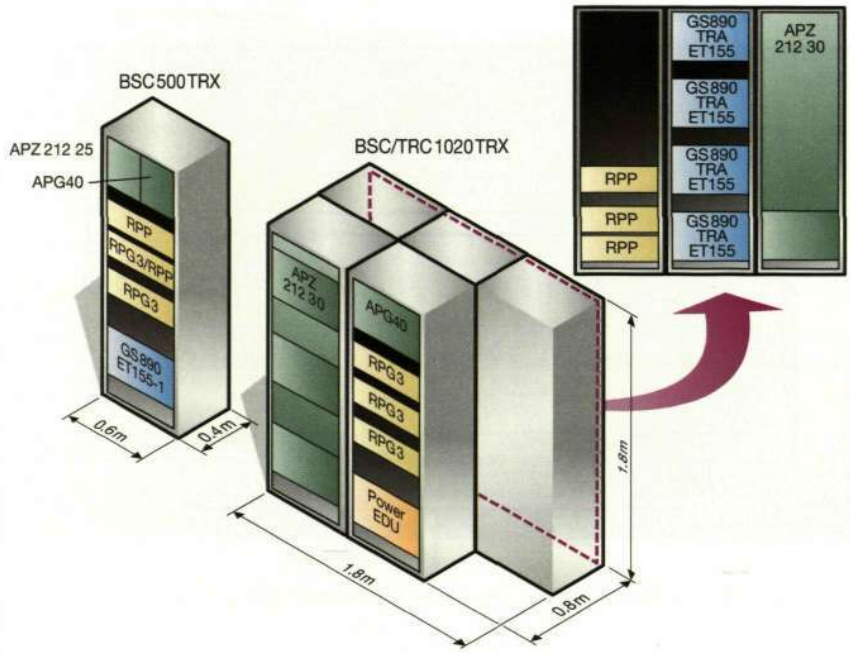

Figure 14
Node layout for BSC 500 TRX and BSC/TRC 1020 TRX.

disks, which must be re-integrated each time a failover occurs. Re-integration in the APG40 occurs when a node is replaced or the RAID system fails.

The AXE IO platforms listed in Table 2 are based on three different high-availability concepts. In the IOG20, the passive node is *hot-standby*, whereas in the APG40, the passive node is *warm-standby*. In this context, the main difference between hot- and warm-standby is that a larger part of the software is up and running in hot standby. In the APG30—with its fault-tolerant concept—identical processes are executed in parallel on separate processor units. If a hardware failure occurs, an immediate failover is initiated from the failing processor unit.

Behind the different availability concepts in Table 2 lie different assumptions about the most probable source of failure. If a hardware failure happens more often than a software failure, the fault-tolerant technology gives better in-service performance. A failover in the APG30 executes in an instant, because the same processes, including the applications, are running in parallel on both sides.

Nevertheless, because the risk of hardware failure is considered extremely low, the case for dimensioning has been based on software failure, and the high-availability concept used in the APG40 has been introduced as the preferred solution. Given the risk of hardware faults, application software faults and system software faults, the failover period has been held to a mininum. On the APG40, these three kinds of fault are well below the minimum CP buffer capacity of two minutes, which means that no data is lost.

Performance and functionality enhancements

The demands put on the adjunct processors are not the same as those put on the central processor. The central processor provides non-stop computing, so the adjunct processor concept is based on and provides high-availability computing. Consequently, new processor systems can be obtained from mainstream data communications vendors. The first generation APG40 is based on Intel Pentium III 500 MHz processors and Microsoft NT4 clustering technology. By adapting these mainstream products to spe-

cific needs, Ericsson benefits from ongoing research and development conducted by others. This approach also simplifies the rapid introduction of new features and applications using Microsoft and Rational development tools integrated in an adjunct processor design environment. Examples of this are debugger and Rational test tools integrated in Microsoft Visual C++, Rational ClearCase, and HTML online help.

Other examples of mainstream technology being introduced full-scale in the AXE by the APG40 are the TCP/IP, FTP, and Ethernet (for external communication and internal AXE communication), and the cPCI building practice. By introducing TCP/IP and Ethernet in the IPN for CP-AP communication, the central processor reload speed is 10 times that obtained when using the APG40 with signaling terminal open communication (STOC).

The introduction of TCP/IP as the external interface to the OSS is a major step toward replacing the propriatary interfaces based on MTP and X.25. The external Ethernet board will enable a TCP/IP connection of up to 100 Mbit/s.

Telnet has been introduced for handling commands. Similarly, FTP and RPC are used for handling files. Support for a component object model-based (COM) OSS might be introduced in the future.

For local element management of the AXE 810, WinFIOL and Tools 6.0 (and later versions) have been adapted to the new APG40.

Results

Best-in-class hardware embodies compact layouts and easy extensions, exemplified in Figure 14 by one BSC and one combined BSC/TRC for the GSM system.

The BSC, which has a capacity of 500 transceivers, occupies only one cabinet. At the top, the CP (APZ 212 25) and APG40 share one shelf. Below this are three magazines with a mix of RPPs and RPG3s. The C7, transceiver handler and PCU applications run on this hardware. At the bottom of the cabinet, one GEM contains ET155, synchronization, DL3 interface boards (DLEB) and the group switch.

The combined BSC/TRC has a capacity of 1,020 transceivers and approximately 6,000 Erlang. The APZ 212 30 is needed for this larger node, occupying two cabinets. Another cabinet houses four GEMs, each of which contains a mix of transcoders, ET155, synchronization, DL3 interface boards and the group switch. The next cabinet contains one APG40, the RPG3s (for transceiver handling) and C7. The power equipment is located at the bottom. PCUs can be included in the last cabinet if the BSC/TRC is intended to handle GPRS traffic.

TABLE 2, THREE AXE IO SYSTEMS

	HSB (IOG20)	FT (APG3x)	HA (APG4x)
Hardware fault	Software failover 90 sec.	Hardware failover < 1 sec.	Software failover < 60 sec.
Application software fault	Software failover 90 sec.	Process restart 22 sec.	Process restart < 1 sec.
System software fault (major)	Software failover 90 sec.	Reboot 80 sec.	Software failover < 60 sec.

Conclusion

The AXE architecture allows us to continue to improve the different parts of the system as required by network evolution. This ensures that operators always have a network with the best possible cost-of-ownership. A limited number of high-capacity nodes with a limited footprint keep site costs down. State-of-the-art hardware with reduced cabling also ensures quick, high-quality installation procedures. Operation and maintenance is kept economical through low power consumption, an interface consistent with the OSS, and a high-capacity, user-friendly extension-module platform. The software architecture allows migration toward next-generation networks, whether fixed or wireless.

REFERENCES

1 Ekeroth, L. and Hedström, P-M.: GPRS support nodes. Ericsson Review Vol. 77 (2000):3, pp. 156-169.

2 Hägg, U. and Lundqvist, T.: AXE hardware evolution. Ericsson Review, Vol. 74 (1997): 2, pp. 52-61.

3 Hallenstål, M., Thune, U. and Öster, G. ENGINE server network. Ericsson Review Vol. 77 (2000):3, pp. 126-135.

4 Andersson, J-O. and Linder, P. ENGINE Access Ramp—The next-generation access architecture. Ericsson Review Vol. 77 (2000):3, pp. 146-147.

Open architecture in the core of AXE

Mats Jonback and Stefan Schultz

The APZ 212 40 is the first central processor of a new generation based on industry-standard microprocessors. It introduces a new *warm-standby/hot on-demand* system principle: the CP has two independent sides, each of which contains two processors—one that serves as an instruction processing unit and one that serves as a signal processing unit—which run as a two-way symmetrical multiprocessing computer. The APZ virtual machine handles ASA execution and is the middleware that guarantees telecommunications-grade availability.

The authors describe the APZ 212 40 hardware and software, the *warm-standby* concept for high availability, and differences and similarities between this and previous APZ processors.

Introduction

Being based on industry-standard microprocessors, the APZ 212 40 central processor (CP) is a milestone in Ericsson hardware implementation. The latest in a line of central processors, the APZ 212 40 is the first central processor of a new generation, and the platform for future multiprocessor solutions (Figure 1).

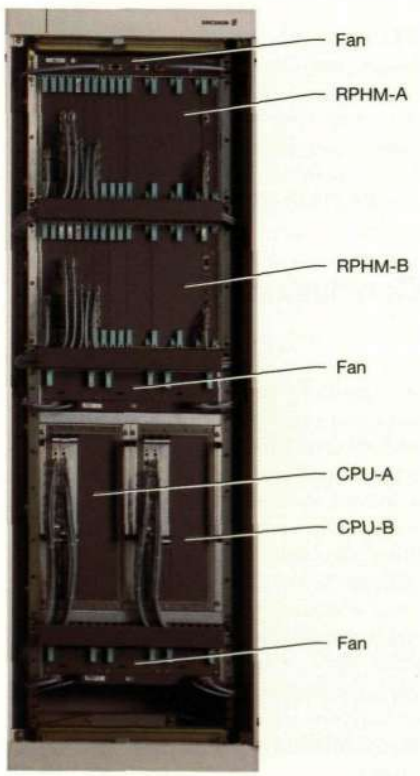

Figure 1
APZ 212 40 central processor cabinet.
Note: CDU panel not shown.

The APZ 212 40 introduces a new system principle. In place of the parallel synchronous (lock-step) machine, we are introducing a *warm-standby/hot on-demand* principle. The CP has two independent sides, A and B, that are loosely coupled via a high-speed interconnect and a maintenance channel. Each side contains two processors: one that serves as an instruction processor unit (IPU), and one that serves as a signal processor unit (SPU), running as a two-way symmetrical multiprocessing (SMP) computer. The hardware and the operating system define a two-node cluster of two-way SMP computers.

The APZ virtual machine (VM), a new part of the system, handles the ASA execution and is the middleware required to achieve telecommunications-grade availability. It provides recovery and repair for hardware faults, software upgrades and the like while minimizing traffic disturbance. At the same time, Ericsson is introducing an industry-standard hardware equipment practice, cPCI, and a standard operating system. Parts of the CP core have been rewritten in standard C++ code using commercial development tools.

The APZ 212 40 will be deployed as a high-capacity central processor for all kinds of application. It offers processing capacity three times as great as the fastest APZ 212 30, and is fully backward-compatible—that is, it can run all existing AXE application software.

APZ 212 40 hardware

Hardware compatibility and structure

Processor technology is advancing rapidly, and Ericsson has decided to make the most of research and development in the mainstream of the computer industry by using industry-standard microprocessors. This reduces Ericsson's time-to-market, especially for upgrades, which can be made available as improvements are achieved in third-party development. At the same time, this makes it possible to concentrate in-house design efforts on developing services, improving robustness, and integrating solutions.

Each APZ 212 40 central processor magazine (CPUM, see Figure 3) contains five boards and a power module (Box A). The CPU board contains the industry-standard microprocessors. Connections are provided by

- the regional processor handler magazine

interface (RPHMI)—through cable connections;
- the base input-output (IO) board—through access ports and some general support functions; and
- the update bus board (UPBB)—through interfaces such as the new Ethernet bus and cable connections.

The new cPCI backplane is connected to the CPU backplane via a PCI interface module (PIM).

Ethernet and the interplatform network

To speed up central processor-to-APG40 communication, and thus reduce reload time, boot time, backup time, and the like, the APZ 212 40 uses direct communication

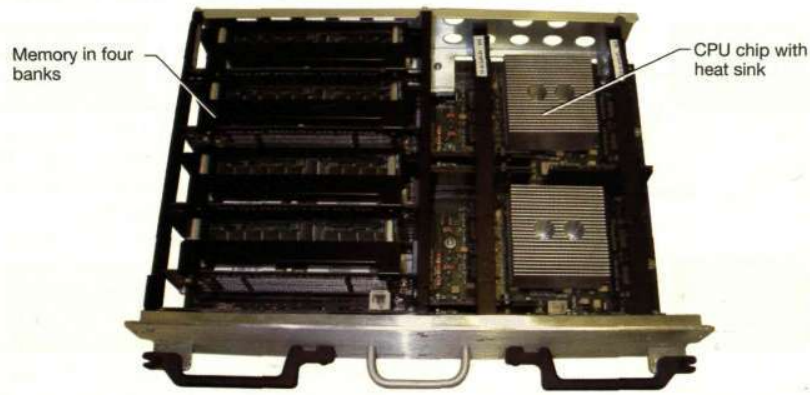

Figure 2
CPU magazine.

BOX A, APZ 212 40 BOARDS

CPU board
- 2 high-performance (GHz) microprocessors, each with 8 MB level 2 cache
- 8 GW (16 GB) SDRAM memory

RPHMI board
- Cable connections to the magazines (RPHB) for the A- and B-sides RPH-A and RPH-B
- Cable connection to the RPHMI on the other CP side for error information and for control of system states (WSB).

Base IO board
- Access port (connected to UPBB) for obtaining detailed low-level system-error information from the microprocessor system (CPU board)
- Some general support functions for the CPU board and power module

UPB board
- One 1 Gbit/s Ethernet optical fiber connection to the other CP side for updating
- Two 100 Mbit/s Ethernet cable connections to the adjunct processor over the interplatform network (IPN), connected via the IPNX Ethernet switch in RPHM
- One 10 Mbit/s Ethernet cable connection to the adjunct processor, a processor test bus (PTB) for access with the central processor test (CPT) command interface, connected via the IPNX Ethernet switch in RPHM
- Cable connection to the central display unit (CDU) panel at the top of the cabinet
- Cable connections to the CPU magazine fan units for monitoring and controlling the fans

The PCI interface module (PIM) card (not visible from the front) includes functionality for connecting the CPU backplane with the cPCI backplane.

Power module (in bottom left part of the CP subrack)
- Duplicated –48V external connection
- DC/DC converters for internal system voltages

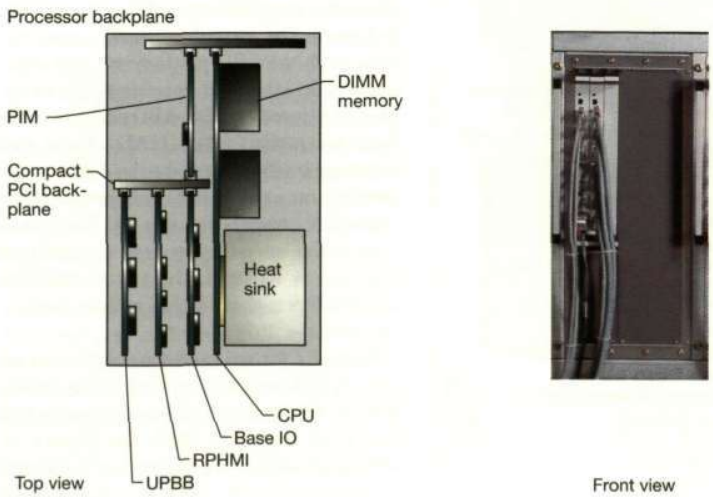

Figure 3
CPU subrack.

via Ethernet. Figure 4 shows the Ethernet switch used to bridge communication between the APG40, CP, AXD 301, and other equipment connected to the interplatform network (IPN). Communication between the central processor and APG40 uses cross-connected 100 Mbit/s Ethernet links. The two CP sides are connected by end-to-end 1 Gbit/s Ethernet link optical fiber.

The central processor test (CPT), which is used for maintenance and repair activities, is also connected to the Ethernet switch. All Ethernet connections to and from the central processor are connected to the UPBB board via the Ethernet switch. The use of Ethernet for CP-to-IO (APG40) communication also decreases the load on the regional processor handler bus (RPHB).

Cross-connected regional processor handler bus

As in previous models of the APZ, an RPHB runs between the signal processor unit and the RPHM. However, unlike previous APZs, the RPHB is cross-connected in the APZ 212 40 for redundancy. This connection is between the RPHMI board in the central processor unit magazine and the regional processor input-output (RPIO) board in the RPHM. With the exception of the RPIO board, the two RPHMs are identical to those developed for the APZ 212 30. The RPIO board has been updated to support the cross-connection.

Of the two CP sides, one is always the executive side; the other is on standby. The executive CP side has control over one or both of the RPHMs. In a normal traffic situation, the executive side controls both RPHMs. Because one of the two RPHMs is always active (used in traffic), if an error occurs, the standby RPHM can take over immediately. The RPHMs can also be separated, and each can be assigned to its own CP side.

To avoid accidents, such as the inadvertent disconnection of the active RPHB, the state of the RPHMs is shown on the CP control and display unit (CDU) panel at the top of the cabinet.

APZ 212 40 software

Layered software structure

The structure of the APZ 212 40 can best be described as layered (Figure 5). From the bottom, we have the CPU hardware platform, with microprocessor, memory, firmware console, privileged architecture library (PAL) code, and other functions. On top of that, we find a suitable operating system that interfaces with the hardware and the application—in this case, software developed by Ericsson.

The APZ virtual machine is a PlexEngine that executes traffic. It is the interface between the existing application software and the hardware platform, similar to a microprogram. The virtual machine communicates with the CPU platform via two additional thin layers: an operating system application interface (OS API) and the hardware abstraction layer (HAL). These two layers are intended to make the APZ less dependent on a particular microprocessor architecture or operating system. The idea is that if the operating system or hardware platform is replaced, the virtual machine and the layers above it can remain more or less intact.

The APZ virtual machine accesses the resources and services provided by the underlying CPU hardware and operating system via HAL and OS API. In Figure 5, PLEX/ASA indicates where in this layered structure the APZ software units and the

BOX B, TERMS AND ABBREVIATIONS

AMU	Automatic maintenance unit
AOT	Ahead of time
AP	Adjunct processor (external IO system used for man-machine communication, external boot media and charging output)
APG40	The latest generation of adjunct processor most suitable for the APZ 212 40
ASA	Programming language assembler, also machine language in earlier APZ CPs
BOOTP	BOOT protocol
BUMS	Backup in the main store
CDU	CP control and display unit
CP	Central processor (includes RPHM, CPUM and cabling)
cPCI	Compact peripheral component interconnect
CPIO	Central processor input-output
CPT	Central processor test
CPU	Central processor unit. In this machine, same as CPUM
CPUM	CPU magazine (subrack), one CPU side
DIMM	Dual in-line memory module
DS	Data store
FTP	File transfer protocol
GW 16	Gigaword (1,073,741,824 words) 16-bit word length
HAL	Hardware abstraction layer
IO	Input-output
IP	Instruction processor
IPN	Interplatform network
IPNX	IPN switch
JAM	Jump address memory
JIT	Just in time
MAU	Maintenance unit
MW 16	Megaword (1,048,576 words) 16-bit word length
OS	Operating system (commercial)
OS API	OS application program interface
PAL	Privileged architecture library
PCI	Peripheral component interconnect
PIM	PCI interface module
PLEX	Proprietary programming language for exchanges
PS	Program store
PTB	Processor test bus
RP	Regional processor
RPHB	RP handler bus
RPHM	RP handler magazine (subrack)
RPHMI	RPHM interface board
RPIO	Regional processor input-output
RS	Reference store
SDRAM	Synchronous dynamic random access memory
SMP	Symmetrical multiprocessing
SP	Signal processor
UPBB	Update bus board
VM	Virtual machine
WSB	Working state bus

rest of the application can be found. In the actual target machine, only ASA code is executed. PLEX-to-ASA compilation is done off target.

The APZ virtual machine employs the operating system application interface to communicate with the operating system and to use its functions when required. The OS API layer is intended to make the system less dependent on particular operating system platforms. If, for some reason, it becomes beneficial to change the platform, the main changes will be made in the OS API. The central processor design is based on a commercial CPU, making it feasible to use a commercial operating system that

- is available for the chosen microprocessor architecture;
- supports multiple processors;
- uses 64-bit processor architecture;
- supports dynamic-link libraries; and
- has memory protection between processes and threads.

Two ASA compilers

The APZ 212 40 CP compiles ASA binary code using two methods: just-in-time (JIT) and ahead-of-time (AOT). Thus, the APZ includes a JIT compiler and an AOT compiler to compile the ASA binary code into native CPU code. This is transparent to the user.

Just-in-time compiler

The APZ virtual machine calls the JIT compiler when there is no compiled code—that is, when neither AOT-compiled nor JIT-compiled code is available. The JIT program compiles and stores a "compilation unit," which is a small part of the ASA code—a sequence of instructions out of which there are no jumps (short compilation time is the main focus of the JIT compiler). The next time that piece of code is called, the compiled code will be executed directly. If anything happens that affects the ASA binary code, such as a correction being inserted or a trace on an instruction address using the test system, the generated code is invalidated and must be recompiled the next time it is called.

Compared to earlier versions of APZ, the JIT-compiled code has full register coherency. Consequently, the JIT compiler is the one to use during troubleshooting. It has full support for the test system and the current correction handling—the code is JIT-compiled until the AOT compilation is finished.

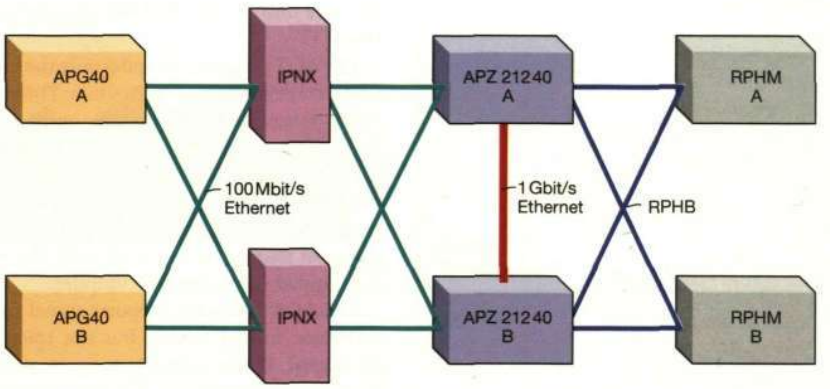

Figure 4
Ethernet links and cross-connected RPHM.

Ahead-of-time compiler

The AOT program compiles entire software units of ASA binary code in background execution, including all currently loaded ASA corrections. The AOT compilation is performed in a separate process on the SPU processor. Its purpose is to generate code that is as efficient as possible. The AOT compiler is allowed longer execution time to produce code. Thus, unlike the JIT compiler, it can optimize over an entire ASA

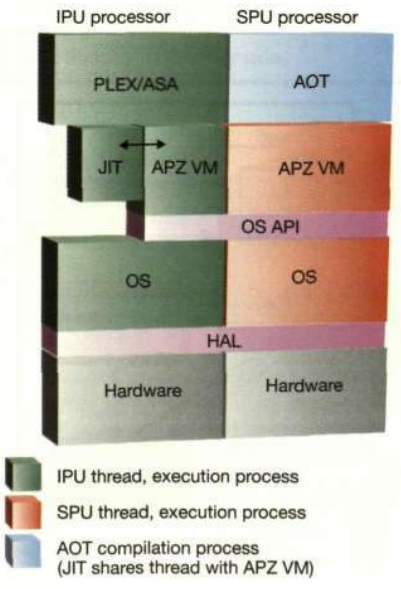

Figure 5
Central processor structure of the APZ 212 40.

software unit. Capacity is the main focus of the AOT.

The AOT program compiles only the most frequently executed software units. These are either selected by the system or decided manually—that is, they are preset. The number of blocks is estimated to be 10-15% of the total number, or around 150 blocks. This should be sufficient to cover at least 90% of the most frequently executed code, the rest is compiled using the JIT compiler.

The AOT compiler supports signal trace and some forlopp traces, but not trace on out-signal. If any other traces are requested by the test system, the JIT compiler will take over program execution. The *trace on every jump* option using the test system is no longer available.

The switch between AOT-compiled and JIT-compiled code is transparent to the user. For instance, when a trace is activated, the system switches from AOT-compiled to JIT-compiled code. When the trace is deactivated, the system automatically switches back to the AOT-compiled code.

Record-oriented data-store architecture

The APZ 212 40 introduces a record-oriented data store (DS) architecture. Record-oriented data allocation is based on the observation that after one piece of data has been read into memory, another variable in the same record (with the same pointer value) will likely be read soon after.

In contrast to existing central processors, modern microprocessors rely heavily on caching, so they spend much less time reading data if they have just read the preceding data. Thus, when the APZ 212 40 reads data, some data that follows immediately after is also read and put in the cache memory. The introduction of the record-oriented data-store architecture only affects a few PLEX or ASA software units in the APZ operating system and not the application.

APZ virtual machine

The APZ virtual machine is the compiled and linked module of C++ code that executes the application. The APZ virtual machine, which is compiled into native code off-line, is the first module loaded after the operating system when the system is powered on. This microprogram substitute takes care of the incoming signals from the regional processors (RP) and from the IPN network—it schedules, dispatches, and distributes the signals that result from the jobs. The APZ virtual machine also offers special services or support functions to the software layers above it (Figure 5).

The APZ virtual machine executes in a process with two threads, referred to as the instruction processor (IP) thread and the signal processor (SP) thread. The instruction processor thread handles all job scheduling (Figure 6). The main loop of the instruction processor thread starts by polling the job buffers for incoming jobs. The jobs are then dispatched and executed. The main parts of the instruction processor thread are

- Exec—which is the virtual machine kernel, dispatcher, scheduler, and signal distributor;
- error-handling and recovery;
- the run-time log;
- the JIT compiler; and
- an interface to the AOT-compiled code.

The signal processor thread handles external communication and controls the job table. Basically, it runs in an infinite loop checking the external and IPU interfaces for new signals or messages to store in job buffers—for subsequent execution in the instruction processor thread. The main parts of the signal processor thread are

- job-table scanning;
- error-handling and recovery;
- communication support (transports and external signaling);
- RP signaling; and
- IPN signaling.

Figure 6
APZ virtual machine.

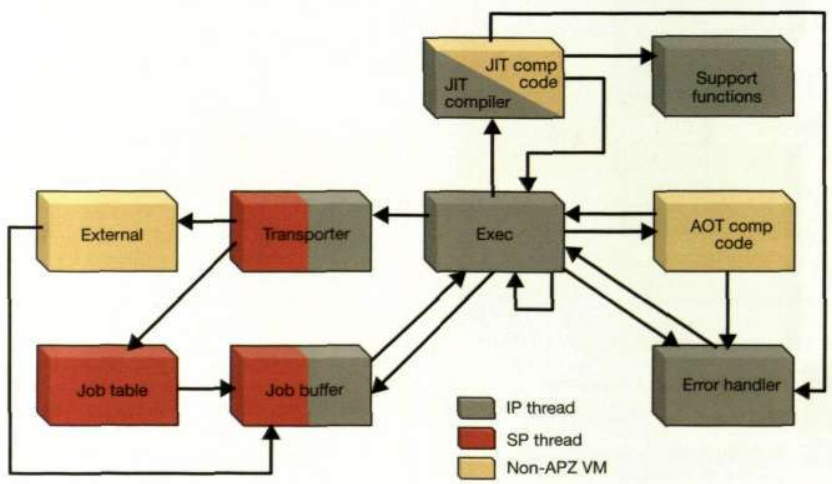

Warm-standby concept for high availability

The CPU sides do not run in parallel synchronous (lock-step) mode. Instead, a *warm-standby/hot on-demand* principle is applied. In normal traffic situations, one CP side executes and the other side is on standby (normal). The same system that is executing on the executive side is preloaded into memory on the standby side. The standby side is also regularly updated with transient data.

Likewise, the standby side is updated through the automatic backup function. When a side is scheduled to be switched or a function is scheduled to be replaced—for example, when testing, performing maintenance, or adding hardware—the standby side is "heated up"; that is, the entire memory is copied from the executive side to the standby side, so that the standby side becomes a mirror image of the executive side. A 1 Gbit/s Ethernet link (the updating bus) between the two CP sides provides the capability to heat up the standby side. The copy procedure can be divided into two phases (Figure 7): background copy (raw copy) and traffic-oriented copy (high-level copy). Memory is copied page by page using a fault-on-write mechanism. This means that a copied page is set to write protect. Then, when a write order to that page is received, the page is logged before it is written. The page is copied again later (Figure 8).

Measurements and estimates have been made for copying data between CP sides. With one 1 Gbit/s link, the capacity will drop marginally for a period less than 15 seconds. At the end of the copy interval, the system is frozen for around 0.5 seconds, allowing the CP states and the APZ virtual machine to be copied. During this interval, traffic from the RPs is buffered, so no data is lost.

If the standby side is put into operation because of a debilitating hardware fault in the executive side, a restart with reload is ordered. In the normal state, a system will already have been loaded, so what actually occurs is a nearly instantaneous reload followed by a large restart. The command log is later read from the adjunct processor without affecting downtime.

The advantage of the *warm-standby/hot on-demand* concept is in-service performance during software recovery, which fully compensates for any negative in-service performance from CPU hardware recovery. As with previous APZs, planned events do not result in any traffic disturbance.

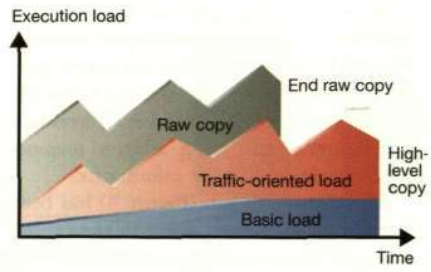

Figure 7
"Heating up."

Previous APZ processors—similarities and differences

New concepts in APZ 212 40

The *warm-standby* concept, which is new to APZ design, obviates the need for maintenance hardware to detect matching errors between CP sides. Consequently, the APZ 212 40 has no separate maintenance unit (MAU).

As explained above, the APZ central processor is based on a commercial microprocessor that requires a substitute for the old

Figure 8
Data transfer to standby side.

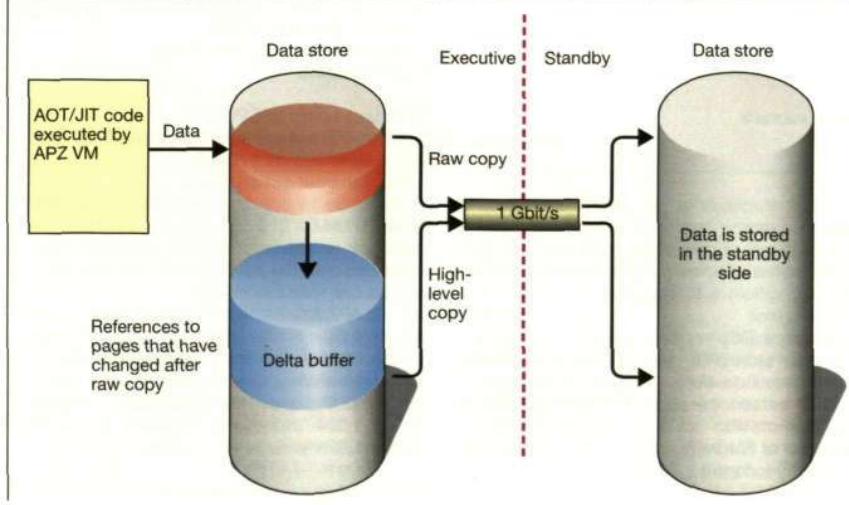

microprogram. The substitute is found in the APZ virtual machine and the new compilers. Hence, the information that the system provides in fault situations and other specific states differs from that of earlier CP versions. For example, the jump address memory (JAM) contains different information.

An Ethernet switch (Figure 4) has been introduced to connect ports in the central processor to the adjunct processor and other systems, such as the AXD 301.

Hardware dissimilarities

The smallest replaceable hardware unit in the new central processor is the entire CPUM—that is, a complete CP side. For the RPHM, the smallest replaceable hardware unit is the same as in previous versions: individual boards. The CP working state logic is implemented on the RPHMI board; a special working state bus (WSB) cable connects the RPHMI boards in the two CPU sides.

The memory can be configured in steps of 1 GB. The normal way of upgrading or increasing physical memory is to update a configuration file on the adjunct processor, replace the CPUM, add physical memory, and then use regular size-alteration commands to inform the applications that more memory is available. As usual, this process is ended with a backup.

A new CP control-and-display unit panel indicates which side is executing and which side is on standby. It also indicates which RPHM is handling traffic at any moment. The CDU panel indicates normal system state when the executing side is switching traffic and the standby side is ready to take over control. The standby side is ready to take over when

- it has the same system generation preloaded in system memory as the executive side; and
- when it is receiving transient data.

A new state (normal, NRM on the CDU panel) has been defined for the standby side to indicate when these criteria are satisfied.

Software dissimilarities

The C and C++ parts of the system are treated separately from the rest of the system. These separate files on the adjunct processor are loaded using BOOTP and FTP. They are not part of the system backup.

To ensure short reload time, backup in main store (BUMS) must be active. When

BOX C, SYSTEM CHARACTERISTICS AND TECHNICAL SPECIFICATIONS

Equipment practice	BYB 501		
Footprint	600 x 400 mm (23.62 x 15.75 inches)		
Height	1800 mm (70.87 inches)		
Cabinet subracks	2 CPU subracks (one each for CPU-A and CPU-B)		
	2 RPH subracks (RPHM) (one each for sides A and B)		
Electromagnetic compatibility (EMC) class:	Fulfills EMC Class B		
Cooling type	Fans		
CPU subrack	2 x 3 packages with 2 fans each		
RPH subrack	1 package with 3 fans each		
CPU subrack (one of the CP sides)	1 processor board including memory 1 interface board to RPHM 1 interface board to other CP side 1 base IO board		
RPH subrack	1 interface board to CPUM Max. 16 RPH boards (parallel type) or Max. 8 RPH boards (serial type) 1 IPNX Ethernet switch board		
Memory size	8 GW SDRAM		
Power supply	–48 V		
Power consumption	~1400 W		
System limit	APZ 212 20	APZ 212 30	APZ 212 40
Data store (DS)	1.5 GW	4 GW	7.7 GW
Program store (PS)	64 MW	96 MW	256 MW
Reference store (RS)	2 MW 32 bits	16 MW 32 bits	16 MW 32 bits
Number of RPs	1,024	1,024	1,024
Number of EMGs	1,024	1,024	1,024
Capacity increase from the APZ 212 20	1.0	3.5	10

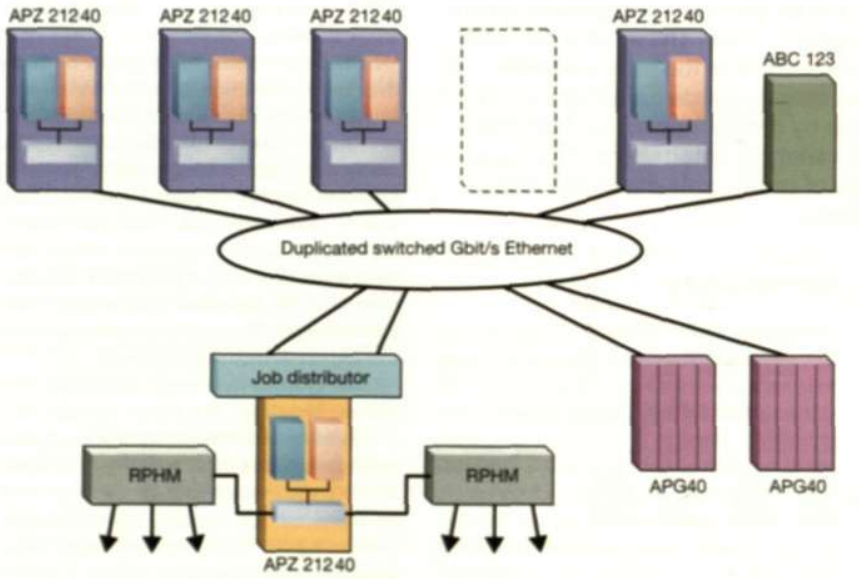

Figure 9
Example of a node in the future.

backup of the system data is first made, it is stored in the backup area of primary memory. As part of the backup function, the system is then transferred to external media—the APG40. The automatic backup function is routed via the standby side to the adjunct processor. In the standby side, an automatic reload is initiated to keep the standby side "warm."

When the APZ VM, JIT, HAL, OS API, AOT, or OS needs to be updated, the new version is loaded as a normal file onto the adjunct processor. A boot is then initiated in the CP standby side, and the new versions are loaded. The CP sides (roles) are then switched and the sequence is repeated.

Conclusion

The new APZ hardware architecture makes the most of increased processing speed through the use of industry-standard microprocessors and innovations such as internal Ethernet buses. The mainstream sourcing of certain processors allows AXE users to benefit from general advances in microprocessor technology.

The latest version of Ericsson's APZ central processor introduces several innovative concepts, but backward compatibility with previous versions of AXE hardware and software has been ensured, thereby safeguarding operator investments.

BOX D, KEEPING PACE WITH FUTURE DEVELOPMENTS

With its architecture and use of commercial microprocessors, the APZ 212 40 allows us to take advantage of rapid evolution in the computer industry. Development will keep pace with the mainstream computer industry and enable improvements in characteristics, such as capacity and robustness, with limited changes in software. The APZ central processor design is also open enough to simplify moves between different microprocessor suppliers and the operating systems they support.

The introduction of the interplatform network (IPN) via directly connected Ethernet makes the APZ 212 40 the obvious platform for future multiprocessor systems. In a distributed multiprocessor system, replicated processors can improve in-service performance and simplify handling when adding capacity to the node. Figure 9 shows an example of a multiprocessor system.

A powerful duplicated processor handles job distribution between a number of replicated processors (call handlers). The APG40, for IO communication, is connected to the Gigabit Ethernet used for communication. Other equipment can be attached to the network based on the needs of applications.

Cable modems—Broadband highway to the home

Juan Figueroa and Bill Guzek

Cable operators are eyeing a huge business opportunity in providing residential customers with high-speed Internet access. The market for cable modems is expected to increase dramatically over the next few years. This increase will be the result of rising user demands for access to the services and capabilities made possible by greater speed, such as high-quality voice, video on demand, and a variety of entertainment services.

The authors describe the technology of using cable plants to bring high-speed Internet access into the home.

Introduction

Surfing the World Wide Web typically offers residential Internet users a click-and-wait experience rather than an interactive extravaganza. Because connection speeds are typically limited to 53 kbit/s or less, frustrated residential online users are demanding higher-speed connections. However, in spite of the slow narrowband speeds currently available through existing dial-up telephone modem connections, residential Internet and online usage continues to grow rapidly.

Local service providers currently offer residential ISDN services that provide connection speeds up to 128 kbit/s. To offer downstream speeds in excess of 1.5 Mbit/s, service providers will have to look to alternative solutions, such as digital subscriber line (DSL) technologies, faster downstream data connections from direct broadcast satellite (DBS), fixed wireless access, and high-speed cable modems.

Broadband coaxial cable passes by more than 105 million homes in North America, and more than 75 million of these subscribe to cable TV. Coaxial cable connections offer nearly universal coverage and a powerful platform for providing high-speed data access to residences and small businesses. However, to support advanced communications services, one-way cable television systems must be upgraded into modern two-way networks. This is a technically complex and capital-intensive proposition.

Cable systems were originally designed to deliver broadcast television signals efficiently to subscribers' homes. To ensure that consumers obtain cable service via the same TV sets with which they receive over-the-air broadcast TV signals, cable operators recreate a portion of the over-the-air radio frequency (RF) spectrum within a sealed coaxial cable line or CATV network designed and used for cable TV distribution. The system must be upgraded with bidirectional amplifiers in the cable-distribution or CATV network before signals can flow in two directions. Most CATV networks are a hybrid of fiber and coaxial cables. Signals are passed through fiberoptic cables from the head-end center to locations near the subscriber. The signals are then transmitted in coaxial cables that run to the subscriber premises. Higher-frequency signals flow toward the subscriber and lower-frequency signals flow toward the broadcasting head-end.

Cable plant architecture

Several elements are required to bring high-speed data over cable (Figure 1). A cable plant consists of a hub or ring of hubs. In a typical cable plant, one or more of these hub facilities serve as the collection points. The head-end hub gathers television signals from various sources—primarily satellite TV transponders. The TV signals are picked up, decoded, and down-converted to selected channels. These channels are then combined into a local fiberoptic network for local distribution or placed on a higher-capacity optical network between the regional hubs. One or more of these hubs also serve as the main interface to the Internet and are tied to it via high-speed optical links.

Each hub has a head-end that uses smaller fiber bundles to distribute the television

BOX A, TERMS AND ABBREVIATIONS

ADC	Analog-to-digital converter
ARP	Address resolution protocol
ATM	Asynchronous transfer mode
BER	Bit error rate
BPI	Baseline privacy
Cable network	Refers to the cable television plant that would typically be used for data-over-cable services.
CATV	Cable TV
CM	Cable modem
CMCI	Cable modem-to-CPE interface
CMTS	Cable modem termination system
CPE	Customer premises equipment
DBS	Direct broadcast satellite
DHCP	Dynamic host configuration protocol
DOCSIS	Data-over-cable service interface specification
DSL	Digital subscriber line
DSP	Digital signal processor
FEC	Forward error correction
Head-end	Central distribution point for a CATV system
HFC	Hybrid fiber-coaxial
ICMP	Internet control message protocol
IEEE	Institute of Electrical and Electronics Engineers
IF	Intermediate frequency
IP	Internet protocol
ISDN	Integrated service digital network
ISP	Internet service provider
LAN	Local area network
LLC	Logical link control
MAC	Media access control
MCNS	Multimedia Cable Network System Partners Ltd.
MSO	Multiservice operator
NAT	Network address translation
NSI	Network-side interface
OSI	Open systems interconnection
PC	Personal computer
PDU	Packet data unit
QAM	Quadrature amplitude modulation
QoS	Quality of service
QPSK	Quadrature phase-shift keying
RF	Radio frequency
RFC	Request for comments
RFI	RF interface
RISC	Reduced instruction set computing
SID	Service ID
SNAP	Subnetwork access protocol
SNMP	Simple network management protocol
SU	Subscriber unit
TFTP	Trivial file transfer protocol
UDP	User datagram protocol
USB	Universal serial bus
VoIP	Voice over IP

signal to smaller local distribution amplifiers. The fiber is then terminated and the signal is converted back into electrical signals that are sent over coaxial cable into the neighborhood (Figure 2). This scheme of using a mixed fiber and coaxial cable distribution scheme is called a hybrid fiber-coax (HFC) system. The amplifiers in this part of the network must be bidirectional.

In the past, the cable plant typically distributed signals from the head-end system to the customer premises using coaxial cable. Today, however, nearly all operators, including multiservice operators (MSO), have upgraded or are in the process of upgrading their plants to hybrid fiber-coax systems. These systems provide clean, high-quality signals to the neighborhood without replacing all of the cable to the customers' home.

To introduce data into the system, the IP infrastructure must be overlaid on existing systems. Connections to the Internet backbone via concentrators and routers are provided via high-speed optical connections, typically OC-8 and higher-speed connections. These connections are brought into the head-end hub and distributed via routers and optical links to various regional distribution hubs in the MSO network. The IP network is tied into a special adapter called a cable modem termination system (CMTS), which consists of one or more cable modem line cards (CMLC). CMLCs convert the IP data stream into downstream (to the home) and upstream (from the home) RF signals. The downstream signals are sent through an up-converter, which puts them on a specific channel and combines them with the other standard TV signals.

The upstream signals are collected from the subscribers. In a properly designed system, there are several upstream channels for each downstream channel. This is because the upstream data is transmitted at a slower rate than the downstream data. The traffic must be engineered to provide adequate service to users. Voice, streaming video, and gaming services use the greatest amounts of bandwidth. The downstream data rate—approximately 30 to 40 Mbit/s—can be shared by some 500 to 2,000 users. The upstream data rate is approximately 8 Mbit/s per channel. Proxy servers and data cache servers are also employed at the local hub to improve system performance.

Cable service providers in North America and Europe have different transmission requirements for channel width, upstream bandwidth, downstream bandwidth, channel center frequencies, and power limits per channel. After the data-over-cable service interface specification (DOCSIS) had been successfully certified in North America and Europe, cable operators adapted it and most of the specifications set by DOCSIS for Europe. Euro-DOCSIS was formed to address European cable operators, Internet service providers (ISP), and end-user needs (for a

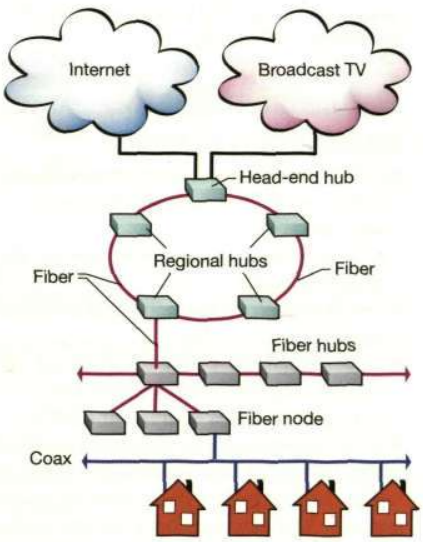

Figure 1
Architecture of the regional hub HFC network.

Figure 2
Architecture of the cable system.

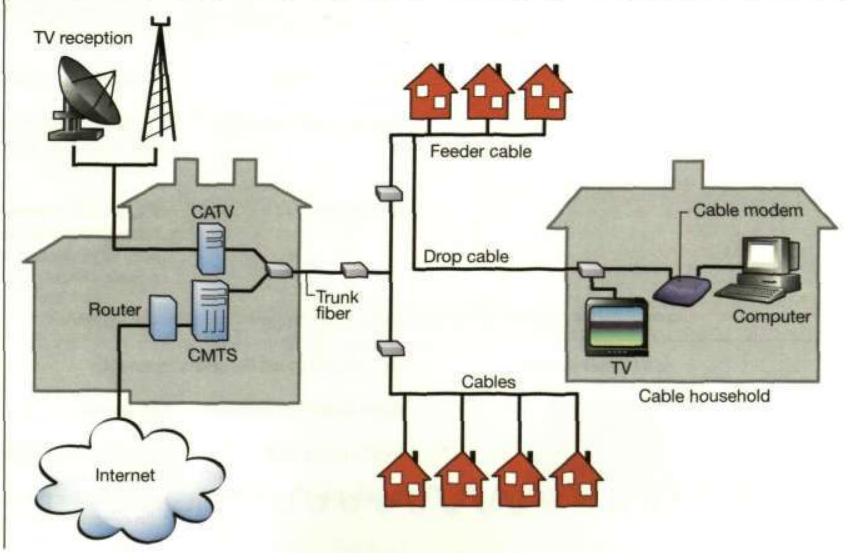

TABLE 1, PROPERTIES OF DOWNSTREAM SIGNALS

Frequency	42 to 850 MHz in North America; 65 to 850 MHz in Europe
Bandwidth	6 MHz in North America; 8 MHz in Europe
Modulation	64 QAM with 6 bits per symbol (normal) 256 QAM with 8 bits per symbol (faster, but more sensitive to noise)

TABLE 2, DATA RATES PER MODULATION SCHEME

	64 QAM	256 QAM
6 MHz	31.2 Mbit/s	41.6 Mbit/s
8 MHz	41.4 Mbit/s	55.2 Mbit/s

comparison of DOCSIS and Euro-DOCSIS, see Tables 1, 2 and 3).

The cable modem termination system

A CATV network consists of six major parts:

- a cable modem termination system;
- a trunk cable;
- a distribution system or feeder in the neighborhood;
- the drop cable to the home and in-house wiring;
- a cable modem, which is connected via coaxial cable and optical fibers; and
- customer premises equipment (CPE).

Figure 3 shows a simple diagram of the data traffic through a data-over-cable system as well as the external interfaces of the key components. The interface between the wide-area network (WAN) and the CMTS is called the network-side interface (NSI); the interface between the CMTS and the cable modem is called the radio frequency interface (RFI); and the interface between the cable modem and customer premises equipment is called the cable modem-to-CPE interface (CMCI). The user connects to the Internet service provider through the WAN interface.

The system shares media for upstream and downstream transmissions. The two most common problems in this architecture are

- the need to control access to shared resources, particularly upstream bandwidth; and
- the injection of ingress noise in the upstream direction.

Traditional coaxial cable systems typically operate with 330 or 450 MHz of capacity, whereas modern hybrid fiber-coax systems have been expanded to 850 MHz. The terms CMTS and head-end are commonly interchanged in contexts that refer to the equipment responsible for communications in the cable network. In reality, the CMTS could be part of the head-end—the equipment from which multiservice operators broadcast television content.

Logically, downstream video programming signals—that is, from the CMTS to the cable modem—begin around 50 MHz, the equivalent of channel 2 for over-the-air television signals. The portion of the spectrum between 5 and 42 MHz is usually reserved for upstream communication—that is, from the cable modem to the CMTS. In North America, each standard television channel occupies 6 MHz of spectrum (in Europe, 7 or 8 MHz). Thus, a traditional cable

TABLE 3, ELECTRICAL INPUT TO CABLE MODEM

Parameter	North American value	European value
Center frequency	91 to 857 MHz ± 30 kHz	112 to 858 MHz ± 30 kHz
Level range (one channel)	-15 dBmV to 15 dBmV	43 to 73 dBµV for 64 QAM 47 to 77 dBµV for 256 QAM
Modulation type	64 QAM and 256 QAM	64 QAM and 256 QAM
Symbol rate (nominal)	5.056941 Msym/sec (64 QAM) and 5.360537 Msym/sec (256 QAM)	6.952 Msym/sec (64 QAM) and 6.952 Msym/sec (256 QAM)
Bandwidth	6 MHz (18% square root raised cosine shaping for 64 QAM and 12% square root raised cosine shaping for 256 QAM)	8 MHz (15% square root raised cosine shaping for 64 QAM and 15% square root raised cosine shaping for 256 QAM)
Total input power (40-900 MHz)/ (80-862 MHz) for Europe	<30 dBmV	< 90 dBµV
Input (load) impedance	75 ohms	75 ohms
Input return loss	> 6 dB (88-880 MHz)	> 6 dB (85-862 MHz)
Connector	F-connector per [IPS-SP-406] (in common with the input)	F-connector per [IPS-SP-406] (in common with the input)

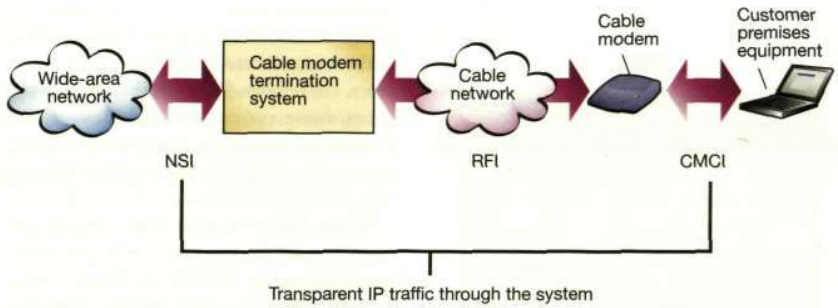

Figure 3
Data traffic through data-over-cable system.

system with 400 MHz of downstream bandwidth can carry the equivalent of 50 to 60 analog TV channels, and a modern hybrid fiber-coax system with 700 MHz of downstream bandwidth has capacity for some 80 to 110 channels.

To deliver data services over a cable network, one television channel (in the 50 to 865 MHz range) is typically allocated for downstream traffic to homes, and one or more channels (in the 5 to 42 MHz band) are used to carry upstream signals. Depending on the availability and the business viability of other channels, the number of cable modem users supported by a head-end can be incremented by commandeering other channels for data and IP transmission. When a channel is used for data, it cannot be used for other conventional, revenue-generating broadcasts, such as commercial TV or pay-per-view services. However, new Inernet-based revenue-generating services can now be offered on that channel.

Figure 4 shows the modulation and demodulation protocols as well as frequency ranges for the CATV system. It also shows the bandwidth and effective bit rates. Figure 5 compares upstream and downstream transmission.

Using 64 quadrature amplitude modulation (QAM) transmission technology, a single downstream 6 MHz television channel can support up to 27 Mbit/s of downstream data throughput from the cable head-end. Speeds can be boosted to 36 Mbit/s using 256 QAM. Depending on the spectrum allocated for service, upstream channels from the home can deliver 0.5 to 10 Mbit/s using 16 QAM or quadrature phase shift key (QPSK) modulation techniques. The upstream and downstream bandwidth is shared by active data subscribers who are connected to a given cable network segment, typically 500 to 2,000 homes on a modern HFC network.

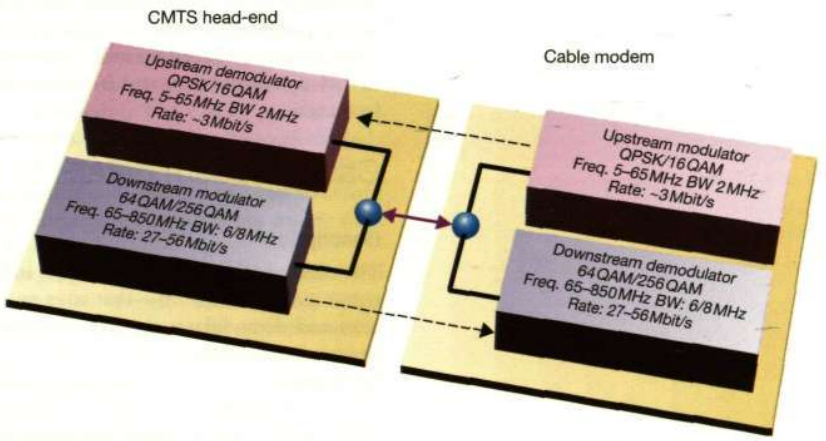

Figure 4
Functional block diagram of the cable modem termination system.

Figure 5
Frequency domains of the upstream and downstream channel positions.

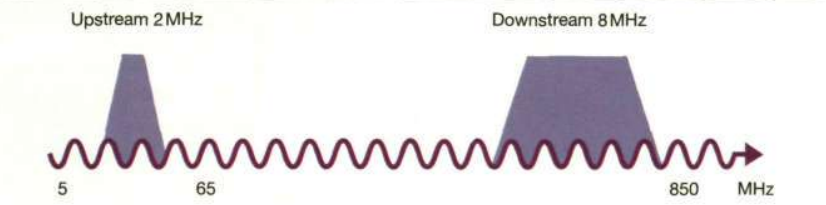

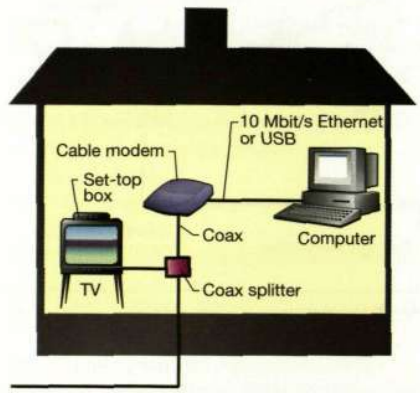

Figure 6
Home environment.

Depending on the network architecture and traffic load, an individual subscriber might experience access speeds from 500 kbit/s to 1.5 Mbit/s or more. Compared to dial-up alternatives, this is blazing performance.

Cable modem details

Description

The cable modem (CM) is a modem in the truest sense of the word—that is, it modulates and demodulates signals. Among its key components are

- a tuner;
- a demodulator;
- an encryption/decryption unit; and
- an upstream modulator.

Cable modems typically send and receive data in two slightly different ways. In the downstream direction, digital data is modulated and then placed on a 6 MHz (North America) or 8 MHz (Europe and PAL system) channel somewhere between 65 and 850 MHz. Upstream transmission is more challenging, since it tends to be very noisy in the 5 to 65 MHz region. Noise or interference is generated by amateur radio operators, citizen band (CB) radios, home appliances, loose connectors and poor cabling. Since cable networks form a tree or branch architecture, noise is aggregated as the signals travel upstream.

When a cable modem is installed, a power splitter and a high-pass filter might be necessary to isolate the TV set from "strong" signals from the cable modem. The filter also blocks upstream ingress noise in the low frequency band. Figure 6 shows how the signal is split from the main cable to TVs and cable modems in the home.

Figure 7 shows the major components of a cable modem implementation. The tuner, including the diplexer, connects directly to the CATV outlet, which provides upstream and downstream traffic to the rest of the cable modem. The tuner solution, which integrates the diplexer into a dual-conversion tuner and digital signal processor, demodulates 64 and 256 QAM signals. The output is a 44 MHz intermediate frequency (IF) signal that is fed into the analog-to-digital converter (ADC) inputs for demodulation and error correction. The media access controller then extracts data and sends it to the customer premises equipment via a universal serial bus (USB) or Ethernet interface.

The tuner also includes a line amplifier for the transmit function, controlled by the media access controller (MAC), which sends upstream signals at the level negotiated by the cable modem and head-end. According to DOCSIS, Reed-Solomon forward error correction (FEC) is recommended. This adds robustness as well as physical layer overhead, which translates into delay.

The heart of the protocol implementation resides in the MAC. The CMTS has control over the allocation of upstream and downstream bandwidth in a dynamic mix of con-

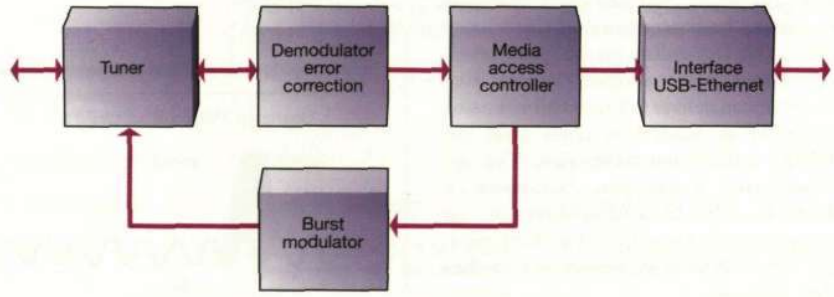

Figure 7
Cable modem architecture.

tention- and reservation-based transmission opportunities. Each cable modem has a unique 48-bit MAC address that is used for registration and authentication. This is entered in a table with the primary service ID (SID) and IP address, once the IP address is assigned. The CMTS then communicates to the modem with the assigned ID.

Significant processing power is needed for

- converting encoded signals into the format needed by the customer premises equipment; and
- conveying messages from the end-user to the destination point through the CATV system.

All available silicon solutions currently implement digital signal processors (DSP) for demodulating downstream transmissions, which are expected to have low bit error rate (BER). The DSPs handle demodulation, whereas regular processing power is devoted to implementing the key components of the cable modem transmission protocol. Powerful RISC processors enforce protocol functions.

Downstream transmission takes place in one of the 6 or 8 MHz channels between 65 and 850 MHz at 25 to 56 Mbit/s, as determined by the modulation scheme (Table 1). The raw data rate depends on the modulation and bandwidth (Table 2).

A symbol data rate of 6.9 Msym/s is used for 8 MHz bandwidth and 5.2 Msym/s is used for 6 MHz bandwidth. Due to error correction, framing and other overhead, the effective data rate is somewhat slower than the raw data rate. Since downstream data is received by all cable modems, the total bandwidth is shared by all active cable modems on the system. This is similar to Ethernet, except that the wasted bandwidth is much greater in Ethernet.

Operation

Initialization

Cable modem communications are set up through a series of initialization steps. After power-up, the modem scans for a downstream channel with which it can synchronize. The CMTS sends synchronization packets to generate a timing reference. Cable modems are synchronized and ranged so that they know when to begin transmission—in order to hit a specific minislot provided by the head-end. The CMTS controls access to slots by assigning specific "transmit opportunities" to ranges of minislots (a transmit opportunity can be contention- or reservation-based). A reserved slot is a timeslot that is reserved for a particular cable modem—that is, no other cable modem may transmit in that timeslot. The CMTS allocates timeslots through a bandwidth-allocation algorithm. The algorithm is vendor-specific and might differ considerably from vendor to vendor. Reserved slots are generally suited for longer data transmissions. After synchronization is complete, the cable modem receives the upstream parameters it needs to inform the CMTS of its presence on the network. The cable modem receives the upstream allocation information—which it uses to start the ranging process.

Due to the physical distance between the head-end and cable modem, the time delay (in milliseconds) can vary significantly. To compensate for delay, each cable modem employs a ranging protocol that effectively adjusts its internal clock. To do so, a number of consecutive timeslots (normally three) is reserved for each ranging process. The cable modem is instructed to transmit in the second timeslot. The CMTS measures the transmission and instructs the modem to adjust its clock as necessary. The two timeslots before and after create a gap to ensure that the ranging burst does not collide with other traffic.

Ranging is also used to coordinate the transmission power level of all cable modems, so that the upstream bursts arrive at the CMTS at the same level. A balanced transmission power level is essential for maintaining optimum performance of the upstream demodulator in the head-end. The attenuation from the cable modem to the head-end can vary by more than 15 dB.

Contention slots are open for all cable modems to transmit in. If two cable modems transmit in the same timeslot, their packets collide and the data is lost. When this occurs, the lack of positive ACK from the CMTS serves to indicate that the CMTS did not receive any data, and the cable modems retransmit at another, randomly selected time. Contention slots are normally used for very short data transmissions, such as a request for reserved slots, in order to transmit more data.

Ranging registration allows the cable modem to identify itself to the CMTS. It also allows the head-end

- to assign downstream and upstream frequencies;
- to set power levels; and
- to distribute other administration information necessary to manage the network.

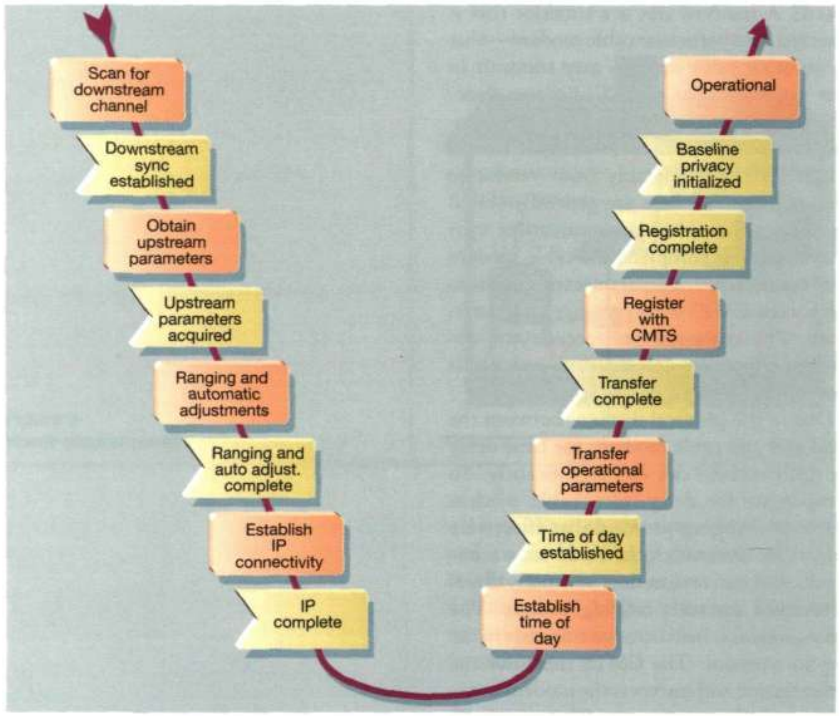

Figure 8
Initialization process.

The time-division multiple access (TDMA) transmissions from cable modems at various distances from the head-end must be coordinated so that all transmissions align with the boundaries of the head-end minislots. Discrepancies between individual modems and the head-end are caused by propagation delays in the CATV plant, the FEC interleaving function (a variable-depth interleaver supports latency-sensitive and latency-insensitive data), and processing time.

The CMTS informs the modem of the propagation delay after it receives the ranging request. Upstream frequency assignments can change at any time. The head-end ensures that the cable modem receives the new frequency assignment before listening for the modem's transmission on the frequency. After ranging is complete, the cable modem must invoke dynamic host configuration protocol (DHCP) mechanisms to obtain an IP address and continue the registration/configuration process (provisioning process).

The DHCP server responds with IP addresses, the name of the configuration file, and server addresses. After establishing a security association, the cable modem must download (via TFTP) a file with configuration parameters from the configuration server. The CMTS then checks that the configuration was obtained from a legitimate configuration server. Finally, the cable modem registers (Figure 8) with the CMTS and be-

Figure 9
Generic frame format.

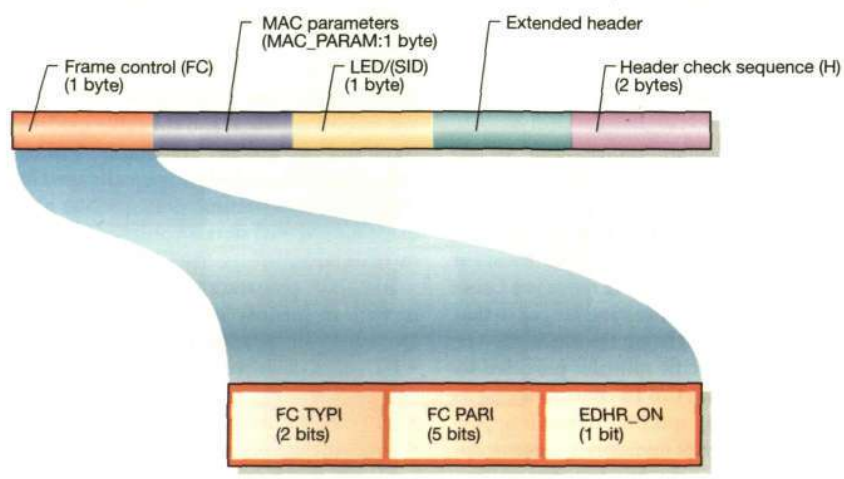

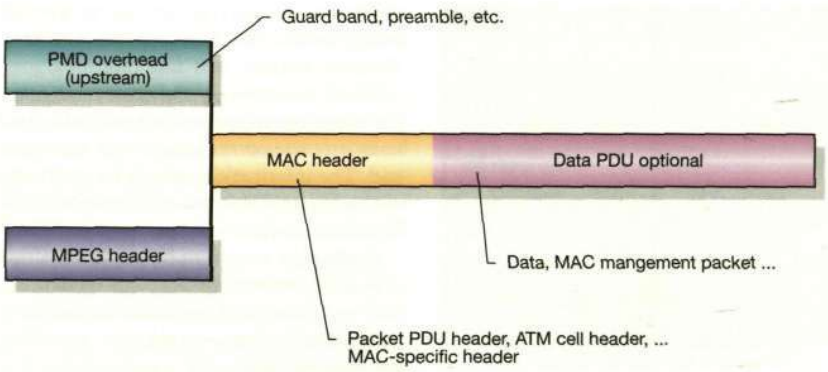

Figure 10
Generic MAC frame format.

comes operational. The CMTS must authorize the cable modem before it can send traffic into the network.

Each cable modem should contain a unique 48-bit MAC address (IEEE 802) that is assigned during the manufacturing process, and security information necessary to authenticate the cable modem.

When the initialization process has been completed, the customer premises equipment can communicate with the outside world through the CATV installation. The communication is bidirectional—downstream and upstream.

Upstream traffic

The term upstream is used to indicate the signal transmitted by the cable modem. Upstream traffic always occurs in bursts, so that many cable modems can transmit in the same frequency. Of the two modulation schemes (QPSK and 16 QAM), 16 QAM (four bits per symbol) is the fastest, but it is also the most sensitive to ingress noise. The upstream direction is characterized by

- a flexible and programmable cable modem under the control of the CMTS;
- frequency agility;
- time-division multiple access;
- QPSK and 16 QAM modulation formats;
- support for fixed-frame and variable-length packet data unit (PDU) formats;
- multiple symbol rates;
- programmable Reed-Solomon block coding; and
- programmable preambles.

When an Ethernet packet arrives at the cable modem from customer premises equipment, it is encapsulated in a MAC packet with a PDU header. The cable modem assesses bandwidth allocation for transmission opportunities and sends the MAC frame when allowed. Transmission opportunities have two basic components: number of mini-slots, and physical layer characteristics.

At the other end, the CMTS receives the packet, removes the header, and then forwards the packet to another cable modem through the radio frequency interface or to the WAN through the network side interface. Figure 9 shows a generic frame format; Figure 10 shows a generic MAC frame format. There are three types of MAC header: packet PDU MAC header, ATM cell MAC header, and MAC-specific header.

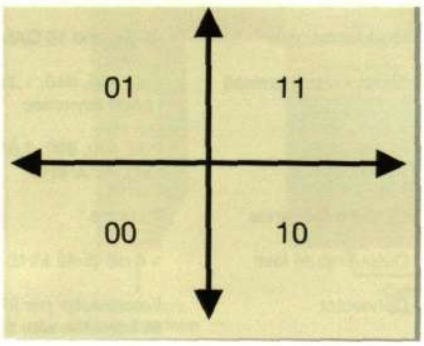

Figure 11
QPSK symbol mapping.

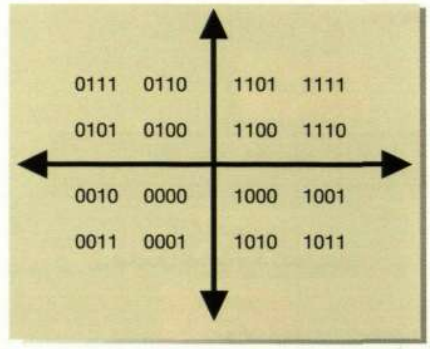

Figure 12
16 QAM symbol mapping.

In the upstream direction, the digital signal is encoded to QPSK or 16 QAM, converted into an analog signal, and sent to the tuner for transmission in the 5 to 65 MHz bandwidth. Figures 11 and 12 show the upstream symbol mappings. Figure 12 represents differential-coded symbol mapping.

Table 4 summarizes the characteristics of the upstream packets generated by the cable modem.

Downstream traffic

The RF signals from the CMTS to the cable modem are encoded in either 64 or 256 QAM. The symbol mappings resemble those shown in Figure 12, but with six-bit encoding for 64 QAM with 16 codes per quadrant, and eight-bit encoding for 256 QAM with 64 codes per quadrant.

There are several stages in the downstream transmission of packets from the CMTS to the cable modem.

When powering up, the cable modem extracts information that regulates when and how it is to communicate with the head-end. The information needed for the initial contact is available in the frames that are broadcast on the network.

During the ranging process, information is transmitted to help the cable modem adjust its timing and upstream transmission levels. After the modem has been registered, the downstream traffic will contain the download data and administrative information that the cable modem requested or information that the CMTS wants to distribute. Every cable modem can listen to the signals broadcast from the CMTS, but only the modem with the correct destination address can access the information contained in the payload section of the packet.

Ordinarily, one downstream channel is paired with multiple upstream channels to achieve the requisite balance in data bandwidth. Each modem transmits bursts in timeslots (reserved, contention or ranging). The cable modem must accept a modulated RF signal (Table 4).

Basic operation of the CMTS

A head-end cable modem termination system communicates with cable modems located in subscribers' homes, to create a virtual local area network (LAN) connection.

TABLE 4, ELECTRICAL OUTPUT FROM THE CABLE MODEM

Parameter	North America	Europe
Frequency	5 to 42 MHz edge to edge	5 to 65 MHz edge to edge
Level range	+8 to 55 dBmV (16 QAM) +8 to +58 dBmV (QPSK)	+68 to 115 dBµV (16 QAM) +68 to +118 dBµV (QPSK)
Modulation type	QPSK and 16 QAM	QPSK and 16 QAM
Symbol rate (nominal)	160, 320, 640, 1,280 and 2,560 ksym/sec	160, 320, 640, 1,280 and 2,560 ksym/sec
Bandwidth	200, 400, 800, 1,600 and 3,200 kHz	200, 400, 800, 1,600 and 3,200 kHz
Output impedance	75 ohms	75 ohms
Output return loss	> 6 dB (5-42 MHz)	> 6 dB (5-65 MHz)
Connector	F-connector per [IPS-SP-406] (in common with the input)	F-connector per [IPS-SP-406] (in common with the input)

Signals from various sources, including broadcast transmissions, satellite-delivered programming, and local television broadcasts, are received and processed in the head-end. Each television signal travels on a different frequency that acts as a self-contained spectrum inside the cable.

Network-layer requirements for the CMTS extend beyond transparency to IP traffic. The CMTS must also support

- variable-length subnet masks;
- classless addressing;
- IP multicast addressing and forwarding;
- Internet group management protocol (IGMP);
- proxy ARP; and
- the filtering of DHCP downstream-bound broadcast packets to protect against BOOTP server spoofing.

The data-over-cable protocol relies heavily on the CMTS for its implementation. Each node on the head-end is capable of supporting between 1 and 2,000 cable modems. The average number of cable modems per node is expected to be around 500. The CMTS is responsible for the initialization, ranging and maintenance of the network formed by the cable modems. The initialization process, which is managed by the CMTS, can be divided into the following phases:

- synchronization—the CMTS sends timing and frequency information to the cable modem, to establish synchronization;
- ranging—the CMTS guides the cable modem through the ranging process;
- IP connectivity—the CMTS establishes IP connectivity;
- time—the CMTS establishes the time of day;
- settings—the CMTS transfers operational parameters; and
- baseline privacy initialization (BPI)—the CMTS initializes baseline privacy if the cable modem can run it.

The CMTS and the cable modem operate as forwarding agents and as hosts. Figure 13 shows the protocol stack used by these components. The data forwarded through the cable modem is *link-layer transparent bridging* and supports multiple network layers.

The main function of the CMTS is to transmit IP packets transparently between the head-end and end-user. Management functions, such as support for spectrum management and software downloads, are sent as IP packets as demonstrated in Figure 13. Both the CMTS and the cable modem operate as IP and logical link control (LLC) hosts according to the IEEE 802 standard for communication over the cable network. The CMTS must support the transport of IP traffic and must be able to restrict the network layer to a single protocol, such as IP.

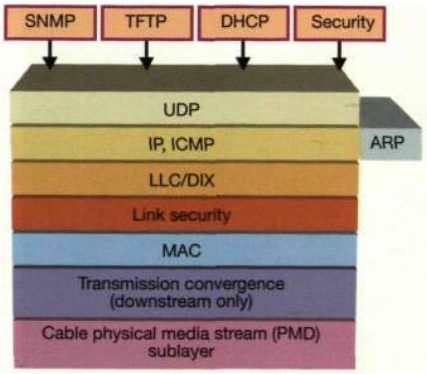

Figure 13
Protocol stack on the RF interface.

Conceptually, the CMTS forwards data packets at two interfaces between the CMTS-RFI and the CMTS-NSI, and between the upstream and downstream channels (Figure 14).

The cable modem access network operates at layers 1 (physical) and 2 (media access control/logical link control) of the open systems interconnection (OSI) reference model. Thus, layer 3 (network) protocols, such as

Figure 14
Data forwarding through the CMTS and cable modem.

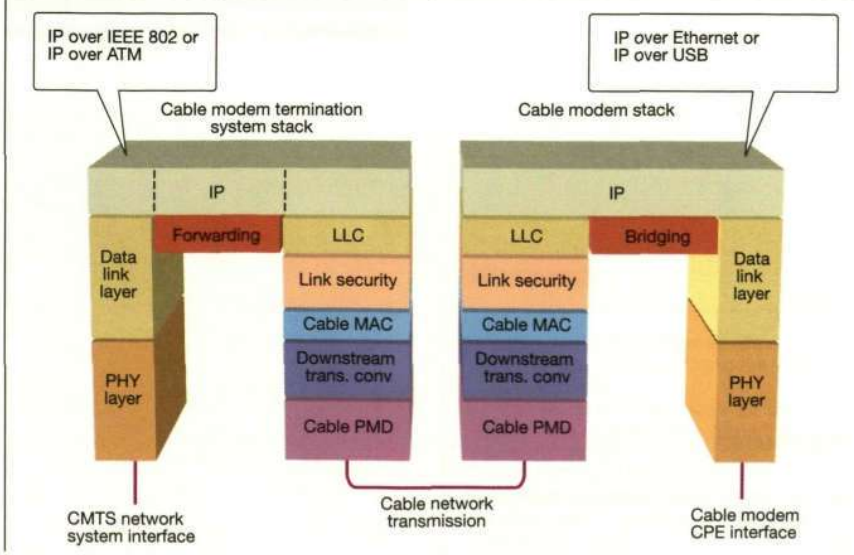

Figure 15
Ericsson currently offers two cable modem products: the PipeRider HM200c and HM201c. The HM200c is DOCSIS 1.0 certified and the HM201c is Euro-DOCSIS certified. The HM200c has received conditional @Home certification. The PipeRiders connect to customer premises equipment via USB and 10 Mbit/s Ethernet interfaces. A PipeLock security button allows the user to isolate the LAN and WAN sides of the cable modem when it is not in use, but keeps the modem attached to the system and available for instant access. In accordance with a specification in the DOCSIS standard, the modem software can be upgraded remotely by the cable operator. The HM200c/HM201c has been certified by regulatory bodies for use in the US, Canada, the EU, Japan, Korea, Australia, and China.

IP traffic, can be seamlessly delivered to end-users over the cable modem platform.

Standardization

In the US, cable operators formed a consortium called CableLabs to accelerate the development of standards within the cable industry. In Europe, tComLabs performs a similar role. Other efforts are currently underway in Europe and North America to standardize voice over IP (VoIP) over cable and home networking.

One of CableLabs' recent standards regulates the sending of data over cable. This effort resulted in the development of the data-over-cable service interface specification (DOCSIS) which was adopted as IEEE 802.14. An appendix to this standard applies to European television standards. The standard was also adopted by tComLabs, a consortium of European operators who coordinated it as Euro-DOCSIS 1.0. (ETSI standard, ES 201 488 V1.1.1). It is based on the 1.1 RFI Version IO6, which is included in tComLabs' appendix to Euro-DOCSIS.

Every cable modem has to go through DOCSIS certification by CableLabs. To date, DOCSIS has gone through seventeen certification waves. Each wave consists of a set of lab tests and documentation to ensure cable modem design compliance with established specifications and standards. Euro-DOCSIS requirements are certified by tComLabs. They are currently in their third wave of compliance testing.

VoIP

Many cable operators are eyeing voice over IP as a major source of revenue and have asked CableLabs to develop a standard, known as PacketCable, which includes voice as well as provisions for other kinds of media. The standard prompted several enhancements to the DOCSIS standard (DOCSIS 1.1), including quality of service (QoS) and additional remote-management features.

Some of the enhancements are

- configuration and registration for cable-modem-based QoS;
- fragmentation of upstream packet data;
- payload header suppression; and
- dynamic establishment of QoS-enabled service flows by the CMTS.

The PacketCable standard stipulates how quality of service should be implemented for different kinds of multimedia flow. The security part of the standard also incorporates provisions for baseline privacy and encryption protocols, for encrypting packet data across the cable network.

The future of cable modems

In the future, cable modems will probably be bundled with other networking features and LAN technologies to facilitate the distribution of broadband IP into the home. At the low end, simple versions of cable modems will be incorporated into personal computers. This evolution into a home gateway is already starting to take place. For convenience and security, the gateway will provide basic software features, such as network address translation (NAT), and function as a DHCP server for local devices, allowing users to set up an extensive home network. Firewalls and content filtering can also be incorporated into the modem to enhance the security of the family computing environment. Operators can simplify the implementation and generate additional revenue by offering these features in the form of a managed service.

Universal plug and play is being investigated to make implementation easy. Other wired and wireless LAN interfaces are being considered, including Bluetooth, HiperLAN/2, HomePNA and powerline networks. Specialty devices, such as set-top boxes and game machines, might also come with built-in cable modems in the future. These devices might even serve as home gateways.

Regardless of the form these devices take, the most interesting and exciting part will be the additional capabilities, services, and entertainment that broadband access will bring into our homes.

Conclusion

A tremendous opportunity exists with broadband over cable. Cable industry standardization is driving exponential growth as evidenced by sales of DOCSIS-standard modems. Cable modems with wireless interfaces will form a foundation step to increased opportunities for mobile Internet services in the home. Soon Internet protocol solutions for various broadband-to-home methods (cable, DSL, fiber) will converge to similar systems. These solutions will not only include infrastructure, but also a suite of additional revenue-producing services for operators.

The full set of services available from Ericsson to cable operators will be a key differentiator among equipment providers. Much of the work for fixed and mobile IP access already developed by Ericsson can be applied to the cable end-to-end solution.

Ericsson has the right mix of capabilities to make broadband over cable a success—cable modems, infrastructure and services. This combination constitutes the bridge that links the worldwide Internet to the home.

REFERENCES

1. High Speed Cable Modems; Albert Azzam, McGraw Hill Series on Computer Communications, 1997
2. Cable Modems: The Answer for Increased Bandwidth or Just a Stop Along the Way?, Analyst Richard Wawrzyniak, Semico Research Corporation, April 1999
3. Data Over Cable Service Interface Specifications; Radio Frequency Interface Specification; SP-RFI-105-991105

The R380s—The first smartphone from the Ericsson–Symbian partnership

Steve Bridges

The Ericsson R380s is the first GSM mobile phone to use the Symbian EPOC32 operating system, which gives the phone the functions that make it a smartphone. The original concept for the R380s was to provide mobile phone and personal organizer functionality in a single phone-sized device. The device was characterized as a business tool rather than a lifestyle accessory. As development progressed, the specification was modified to include a WAP browser, unified messaging (e-mail and SMS), and secure access to restricted-access systems, such as a corporate intranet.

The author explains why EPOC was chosen, and describes the work required to develop the R380s software from EPOC32.

What is a smartphone?

The term *smartphone* has been coined to describe devices which, while being primarily a mobile phone, incorporate elements of functions found in paper-based personal organizer systems or in modern electronic personal digital assistants (PDA). Typically, a smartphone contains a calendar, an address book, e-mail and messaging functions, and a browser for the wireless application protocol (WAP), together with a range of ancillary functions. It is also expected to work simply as a mobile telephone. A smartphone is roughly the same size as a standard business mobile phone—for example, the Ericsson R320—and like a standard phone, it allows users to operate its basic functions in a one-handed fashion. Extra functionality is accessed via the large touchscreen, which differentiates it from standard phones.

What is the R380s?

The Ericsson R380s (Figure 1) is the first GSM mobile phone to use the Symbian EPOC32 operating system, which gives the phone the functions that make it a smartphone.

Why EPOC?

EPOC32 is an ideal operating system for smartphone devices. This real-time, multitasking, multithreaded operating system was created specifically for use in handheld devices, and has been optimized for the constraints of such an environment, the chief of which are

- low power consumption;
- small display;
- limited input and output methods;
- slow processor speed; and
- stand-alone operation.

The creation of Symbian—a joint venture between Ericsson, Matsushita, Motorola, Nokia and Psion—meant that a suitable in-house operating system became available as the R380s project was starting up. However, the fit between EPOC32 and the R380s was not a perfect one. This article describes the work required to develop the R380s software from EPOC32.

The R380s concept

The original concept for the R380s was to provide mobile phone and personal organizer functionality in a single phone-sized device. The device was characterized as a business tool rather than a lifestyle accessory and was specified accordingly. As development progressed and the concept of the mobile Internet took shape, the specification was modified to include features that would take advantage of this mobile business environment:

- a WAP browser would take advantage of the mobile services expected to be offered for this new wireless standard;
- a unified messaging feature would allow standardized control of all available text-transmission services, such as e-mail and short message service (SMS); and
- secure access with built-in security algorithms would allow users to dial up a secure link with restricted-access systems—for example, a corporate intranet.

To provide good functionality both as a smartphone and as a basic phone, two modes of operation were specified. This design affected the hardware of the phone and the user interface (UI).

The solution to the product requirements

The product would include features and functions not previously found in mobile devices. In hardware, the device would have to operate in two significantly different modes: normal phone mode, and smartphone mode. The device must also control two different basic functions with maximum efficiency:

- GSM functionality and the air interface; and
- PDA-type operations.

The requirement to support two different modes of operation while having the largest screen possible for the size of device resulted in the design of a hinged, passive flip. The flip holds the normal phone keypad and covers 60% of the touchscreen. In flip-closed (FC) mode, the R380s looks and operates

BOX A, TERMS AND ABBREVIATIONS

API	Application program interface
ARM	Advanced Research Machines (the manufacturers of the processor used in EPOC devices
ECK	Ericsson component kit
ETSI	European Telecommunications Standards Institute
FC	Flip closed
FO	Flip open
GSM	Global system for mobile communication
OS	Operating system
PC	Personal computer
PDA	Personal digital assistant
QWERTY	Standard typewriter keyboard: first six letters, top-row, left
ROM	Read-only memory
UI	User interface
VGA	Video graphics array
WAP	Wireless application protocol

like an ordinary GSM mobile phone. Key presses are physically transmitted through the flip via plungers attached to the keys. The resultant contact with the underlying touchscreen is interpreted by the operating software. In this mode, the exposed part of the screen does not respond to touch, and the phone looks and feels like an ordinary mobile device, with some added display functionality to help users. The flip is a purely mechanical device that contains no trouble-prone electronic parts or flexible electrical connections. In the event of damage, it can easily be replaced by an approved Ericsson service center.

When the flip is opened, a sensor signals the operating software, which alters the display. In flip-open (FO) mode, the entire touchscreen is exposed, the touch input is enabled, and the screen orientation is rotated through 90 degrees to give a landscape-style display. This mode maximizes the potential for data display and input with a look and feel that more closely resembles established PDA design than that of a mobile phone. Many PDA-type functions are solely available in the flip-open mode, since they are best served by static, two-handed user operation.

Given the requirements for data-input methods, and that in software the device had no proven user interface for the size and form-factor of its screen, the user interface had to be designed from the ground up. A two-processor design was implemented to provide efficient execution of the two main groups of functions in the smartphone, maximizing the possible reuse of software and hardware from earlier successful development work.

The GSM (phone-side) functionality is controlled by a processor similar to that used in standard Ericsson GSM phones. This processor, which runs the OSE real-time operating system from Enea Data, handles all operations related to the air interface and controls audio operation.

The organizer (PDA-side) functionality is controlled by an ARM-based processor design, which manages PDA-side functions and controls the phone-side processor at a high level, initiating and terminating calls on user command. The ARM processor runs the EPOC32 operating system, which has been optimized for this processor family.

The need to close-couple the two different sets of functionality without causing loss of performance on either side led to the design of a specialized communications channel. The two processors communicate via a high-speed software backplane, which provides an efficient interchange of data and commands. Because standard GSM voice calls are handled entirely by the phone-side processor, the PDA-side processor is free to manage other activities during a call.

Figure 1
Image of the R380s.

EPOC32—the starting point

Before Symbian was formed, EPOC32 was associated with the popular Psion 5 organizer (Figure 2), a PDA with a half-VGA-sized touchscreen, a full QWERTY key-

Figure 2
Image of the Psion 5 personal digital assistant.

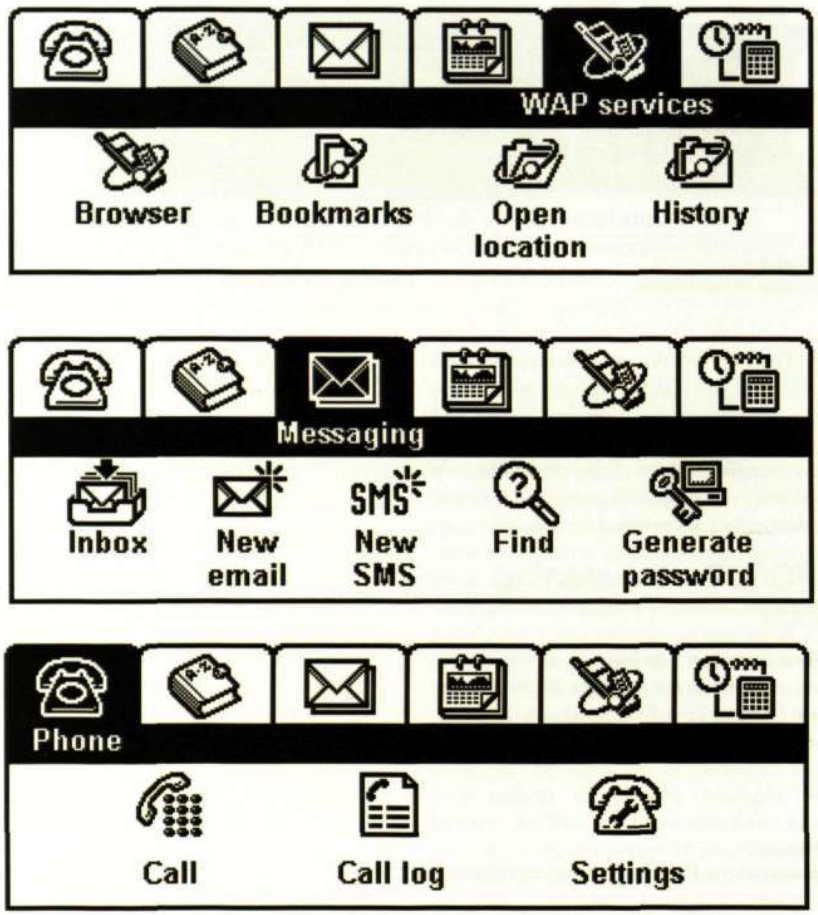

Figure 3
Screen shots of the R380s user interface.

board, and the ability to add extra software packages via PC download or plug-in cards.

EPOC32 offered several benefits to R380s developers, which simplified the task of producing the PDA-side software:

- EPOC is a real-time operating system that permits the support of time-dependent operations, such as telephony;
- the user interface is effectively decoupled from the underlying operating system and application engines—it can thus be modified to accommodate the look and feel of a new device; and
- the software is modular, which simplifies parallel development at multiple sites.

New features needed in EPOC32

Depite its obvious suitability for use in the R380s project, the version of EPOC32 used in the Psion 5 had certain drawbacks for Smartphone use:

- Psion devices adhere to an *always-on* concept, with both a main and a backup battery. The expectation is that the device is always powered and that data in RAM is persistent. However, if both batteries are allowed to run down simultaneously, all user data is lost. Mobile phones adhere to an *often-off* concept. Consequently, to be persistent, user data must be backed up to permanent storage;
- Psion PDAs have no concept of telephony. No engine or user interface was available in EPOC32 to support GSM telephony;
- EPOC32 had been designed to support half-VGA devices with real keyboards and multiple on-screen windows. The user interface and engine software had been optimized for this sort of design;
- the concept of dual-mode operation (FO, FC) was completely alien to EPOC32; and
- software internationalization had been considered a ROM-swap function, but a mobile phone must be able to change languages on the fly.

To overcome these obstacles, and numerous other minor issues, the software was modified as follows:

- A flash file system was introduced, to accommodate persistent data. The challenge was to produce a system that made regular updates—to guarantee an acceptable level of security—without noticeably affecting performance in the rest of the device.
- A telephony support module, known as ETEL, was produced to support telephony functions (from the processor on the PDA side). This module serves as an interface between the phone-side software and PDA-side applications and services that require telephony access.
- The entire user interface was deconstructed and redesigned to function appropriately for the R380 (Figure 3). Instead of being file-based, the concept for the new user interface was task-based.
- Application engines were modified where necessary, and an entire unified messaging application was constructed.
- The operating system was modified to support dual-mode (FO and FC) operation, with
  - seamless switching between defined equivalent views in both modes; and
  - persistence of viewable data.
- The operating system was modified to allow users to control language/locale without the need for external hardware or

without having to modify the hardware. Not all available language/locale packages can be contained in the memory of a single R380s, so provision has been made for downloading the language population not shipped with a particular phone.

- The Symbian software component kit, known as EIKON, was unsuitable for the new user interface. A new component kit—the Ericsson Component Kit (ECK)—was developed to give the correct functionality of components, such as on-screen buttons, icons, dialog boxes, and so on.

Each of these modifications was required to make a truly effective smartphone. The omission of any of them would have resulted in a significantly less attractive, less useful device. However, the primary motivation for this work was usability.

The EPOC32 architecture on the R380s

A detailed description of the R380s software architecture is beyond the scope of this article, but a glance at Figure 4 gives some idea of its complexity.

Usability and the user interface

Since no proven designs existed for this type of device, the usability team had to create a new user interface from the requirements. The team took this opportunity to incorporate continuous usability testing into the design process, to ensure that users would end up with a worthwhile and desirable product. As its guideline, the team took the European Telecommunications Standards Institute (ETSI) definition of usability—in particular, the requirement which stipulates that at least 75% of novice users should be able to perform a given task correctly on their first attempt. All major use cases on the device were tested in this manner. Reiterative testing applied to every aspect of the design, from the on-screen appearance of details, such as buttons, to the overall navigation concept of the operating system.

In this respect, the R380s software can be said to have been designed by users rather than engineers—the activities of the usability team, which acted as users by proxy, ensured that the user interface is the best fit for the size, form factor, and functionality of the R380s. The software started life as Psion 5 software, but as Figure 2 shows, the changes to the user interface were significant.

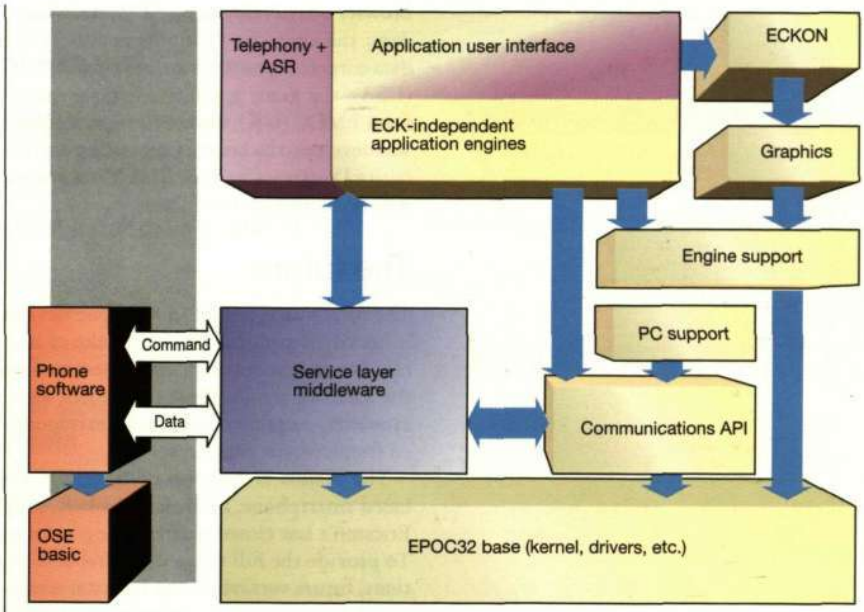

Figure 4
Software architecture diagram.

Multisite software development

The commercial requirement for rapid development dictated that the software would have to be developed at multiple sites. Obviously, this meant that the development team needed a robust system of configuration management and quality control, as well as a strong suite of common development and test tools. Three main sites were involved in software development:

- Ericsson in Stockholm produced the functional specifications and lower-layer PDA software, modified the phone software and verified the system;
- Ericsson Software Applications Laboratory in Warrington (England) produced the user interface specifications (scenario documents), the ECK, the telephony application, and the user interface code for most of the organizer applications. This team also performed the functional test on much of the software; and
- Symbian in London carried out the base port to the ARM processor, supplied software development kits, and modified some of the organizer applications.

In addition, some software development was subcontracted. For instance, the WAP

browser was developed by AU System (Sweden); the character-recognition part of the data entry component was developed by CIC (USA); the game application was acquired from EMCC (UK); authentication software was developed by Secure Computing and Security Dynamics; and the flash filing system was developed by Enea Data.

The future

The future direction of smartphone devices is heavily dependent on the uptake of mobile Internet services by consumers and on the kinds of service these consumers want. However, certain trends can be extrapolated from current data.

The R380s is Ericsson's first EPOC32-based smartphone, but it is also likely to be Ericsson's last closed smartphone platform. To provide the full range of desirable functions, future versions of the Ericsson smartphone platform will be open to third-party developers, in the same way as the Psion 5 or Palm Pilot devices are. The ability to add applications to a smartphone, either via the PC or directly over the air will be vital to the success of these devices in the next decade.

To support development, Ericsson has launched the Ericsson Developers' Zone (www.ericsson.com/developerszone)—a one-stop shop for developers looking to get information, support, training and tools related to Ericsson platforms. Test and certification services for externally developed software are also available via this site.

The modular and scalable nature of EPOC32 will make it the operating system of choice for a wide range of devices, thus ensuring a common software architecture. Future smartphones will support multimedia functions as well as traditional mobile data services.

Conclusion

The R380s, which is the first smartphone on the market to incorporate Symbian's EPOC32 operating system, opens up the future of devices by matching hardware and software to give a full range of desirable functions in a user-friendly, efficient, and portable package. The R380s points the way to a community of devices whose common architecture and open platform allow user requirements to be met by specialized developers.

Previous issues

No. 4, 2000

Experiences of the live WCDMA network in Stockholm, Sweden
Media Gateway for mobile networks
Ericsson's IP-based BSS and radio network server
Control servers in the core network
Ronja—A Java application platform
Wireless *ad hoc* networking—The art of networking without a network
Ericsson's WebOnAir Information Server

No. 3, 2000

ENGINE network server
ENGINE Access Ramp—The next-generation access architecture
Speech recognition technology for mobile phones
GPRS support nodes
Family of RBS 3000 products for WCDMA systems
Ericsson's network-based IP-VPN solutions
Fixed wireless access for broadband services

No. 2, 2000

Evolving from cdmaOne to third-generation systems
Third-generation TDMA
Business solutions for mobile e-commerce
The RBS 2206—A flexible ticket to third-generation wireless systems
Communications security in an all-IP world
HIPERLAN type 2 for broadband wireless communication

No. 1, 2000

IPv6—The new generation Internet
WAP—The catalyst of the mobile Internet
The challenges of voice over IP over wireless
Openness in AXE
Management solutions for IP networks

No. 4, 1999

Network evolution the Ericsson way
The AXI 540 router and the public IP network edge
Real-time routers for wireless networks
Ericsson's Bluetooth modules
Mobile Internet—An industry-wide paradigm shift?
Ericsson's mobile location solution

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

ISSN 0014-0171
Ljungföretagen, Örebro 2001