Cello—An ATM transport and control platform

Jonas Reinius

Ericsson's WCDMA products for third-generation cellular systems are based on Cello, a generic platform for small- to medium-scale telecom applications. Cello has a robust, distributed, real-time telecom control system, and low-cost ATM, TDM or IP transport. The transport system is built on AMAX, an ATM switch designed especially to give a good cost-to-performance ratio.

The author describes the Cello system concept—that is, the Cello platform and development environments—Cello node topology, the execution platform, and key new hardware.

Background

When Ericsson first began designing wide-band code-division multiple access (WCDMA) products for third-generation cellular systems, it lacked a system platform that could support the multimedia services that will characterize these systems. Designers also perceived that the new radio-network products would have to be cost-effective enough to compete with mature GSM products of the future. The Cello platform was selected, mostly because it demonstrated the best cost-to-performance ratio. The present-day design of Cello emphasizes WCDMA products, but it is nonetheless a generic platform that can be used for several different applications, such as media gateways, Internet protocol (IP) routers, private branch exchanges (PBX), and so on.

Figure 1
The main interfaces in Cello and the relationship of Cello nodes in a network.

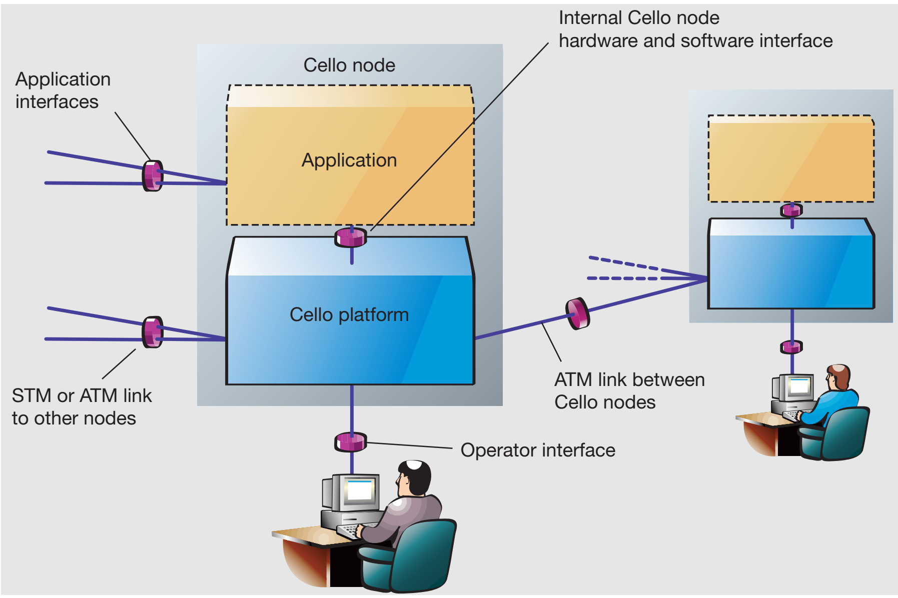

System concept

Cello is a platform product from which it is possible to develop a switching network node, such as a small- to medium-sized asynchronous transfer mode (ATM) switch, a radio base station, or a radio network control node. It comprises a distributed, real-time telecom control system, an ATM-based transport system, and an element-management system built from Web technology, the common object request broker architecture (CORBA), and the IP-based inter-ORB protocol (IIOP).

The ATM transport system was designed to yield a good cost-to-performance ratio. The wide range of interfaces, from 1.5 to 622 Mbit/s, can accommodate everything from low bit-rate to broadband connections. Cello also provides tools and methods for developing custom node software and hardware—such as radio algorithms, call-control programs and device cards. Cello is made up of two main parts:

- the Cello platform; and
- the Cello development environment.

These two parts comprise all the physical equipment and associated software needed to create a switching network node. Cello's modular design facilitates the creation of nodes with different configurations, functionality, capacity, reliability and performance levels.

The Cello platform

The Cello platform is built up of hardware and software modules. The operating system provides software-development and program-execution environments for Cello platform software and application software. It uses the ATM switch to interconnect processors on all types of devices, thus offering a flexible and scalable operating platform for network products.

The Cello platform contains all the functionality necessary to switch ATM and variable-length cells (ATM adaptation layer 2, AAL2) within an ATM-based telecommunications network. It also supports signaling system no. 7 (SS7) and functionality for network signaling on the broadband

ISDN user part (B-ISUP) for ATM connections as well as Q.2630 for AAL2 connections. An IP router network has been integrated for network management traffic. The physical infrastructure consists of

- a 19-inch subrack with 25 generic slots—for processors or device boards. Each node may contain up to 50 subracks;
- cell switching capacity of up to 16 Gbit/s per subrack, with a single or redundant switch core;
- clustered processors with scalable capacity and robustness;
- device processors on each device board;
- a timing unit for network synchronization;
- special-purpose processors (for instance, for radio-link user plane protocols); and
- exchange terminals, plesiochronous or synchronous digital hierarchy (PDH/SDH) for ATM and time-division multiplexing (TDM, currently 1.5 to 155 Mbit/s; 622 Mbit/s will be added when the demand calls for it).

Figure 1 shows the main interfaces in Cello and how Cello nodes work together in a network.

Node topology

Cello nodes can be scaled both in terms of size and capacity. The smallest node consists of a single subrack. Each part of the node that can cause a single point of failure can be configured to be redundant. That is, the switch core, the inter-subrack link and the power feed can be duplicated (1+1 redundancy). Other resources, such as processors and exchange terminals are pooled in an *N+1* configuration.

Subracks comprise several plug-in units, such as processor boards, switch core boards, various exchange terminal boards, and device boards, each of which is attached to the subrack backplane. They also include a cable shelf and fan.

The Cello subrack has 27 slots of which two are reserved for switch core boards, which also include power distribution units. Two of the device slots are specially equipped for timing units. Depending on the processing power and the level of redundancy needed, a subrack may or may not contain main processor boards. In the configuration of large nodes, one subrack serves

BOX A, ABBREVIATIONS

AAL	ATM adaptation layer
AAL2	ATM adaptation layer 2
AAL5	ATM adaptation layer 5
AESA	ATM end-system address
API	Application program interface
ASCC	ATM switch-core circuit
ATM	Asynchronous transfer mode
B-ISUP	Broadband ISDN user part
BP	Board processor
CID	Connection identity
CORBA	Common object request broker architecture
DB	Device board
DHCP	Dynamic host configuration protocol
DNS	Domain name server
E1	ETSI 2 Mbit/s interface
E3	ETSI 34 Mbit/s interface
ET	Exchange terminal
ETB	Exchange terminal board
F4, F5	Management signaling for ATM connections
FTP	File transfer protocol
GPB	General processor board
GSM	Global system for mobile communications
GUI	Graphical user interface
HTML	Hypertext markup language
HTTP	Hypertext transfer protocol
ICP	Internal communication path
IDL	Interface definition language
IIOP	IP-based inter-ORB protocol
IP	Internet protocol
ISDN	Integrated services digital network
ISL	Inter-subrack link
J1	TTC (Japanese) 1.5 Mbit/s interface
J2	TTC (Japanese) 45 Mbit/s interface
MIB	Management information base
MO	Managed object
MOM	Managed object model
MPB	Main processor board
MPC	Main processor cluster
MTP	Message transfer part
NNI	Network-to-network interface
O&M	Operation and maintenance
OMG	Object Management Group
ORB	Object request broker
OSPF	Open shortest path first
PB	Processor board
PBX	Private branch exchange
PDH	Plesiochronous digital hierarchy
PID	Product identity
PPB	Parts per billion (10E-9)
PVC	Permanent virtual connection
Q.2630	Signaling protocol for AAL2 connections
RAM	Random access memory
SAAL	Signaling AAL
SAP	Service access point
SCB	Switch core board
SCCP	Signaling connection control protocol
SDH	Synchronous digital hierarchy
SS7	Signaling system no. 7
STM	Synchronous transfer module
SXB	Switch extension board
SNMP	Simple network management protocol
SPAS	Space switching system
SPIC	Space switch interface circuit
STS	Signaling transport signal level (for instance, STS-3c)
T1	ANSI 1.5 Mbit/s interface
T3	ANSI 45 Mbit/s interface
TCP	Transmission control protocol
TDM	Time division multiplexing
TMN	Telecommunications Management Network
TUB	Timing unit board
UDP	User datagram protocol
UNI	User network interface
VCC	Virtual channel connection
VCI	Virtual channel identity
VPI	Virtual path identity
WCDMA	Wideband code-division multiple access

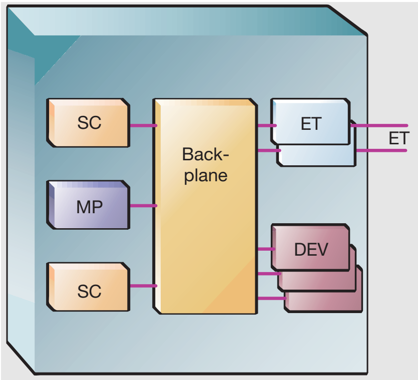

Figure 2
Single subrack node.

as a hub that interconnects the other device subracks.

Figure 2 shows an example of the minimum configuration: a node with a single switch core and a single processor. Apart from the main processor board (MPB) and the switch core board (SCB), the subrack can also be equipped with application device boards or exchange terminal boards (ETB) for connecting other nodes.

If greater capacity is needed, the node can be extended with device subracks. Each device subrack is connected to the hub subrack via internal high-speed transmission links. These links do not conform to standard ATM links. Instead, the switch's internal cell format is used, which yields rapid setup of internal node connections. Four internal links are included on each switch core board. These links also distribute the node's system clock, which is extracted from incoming physical links (for example, an E1). Multiple subrack nodes are configured as a star with a hub. Connections are thus set up through a maximum of three switches.

Local switching is also supported in each device subrack. A node can be expanded without disturbing traffic in the original configuration by connecting the device subracks—using the internal links on the switch core board. Configurations of nodes with more than five subracks require switch extension boards (SXB) in the hub subrack to interconnect the device subracks (Figure 3).

Figure 3
Large multiple subrack node.

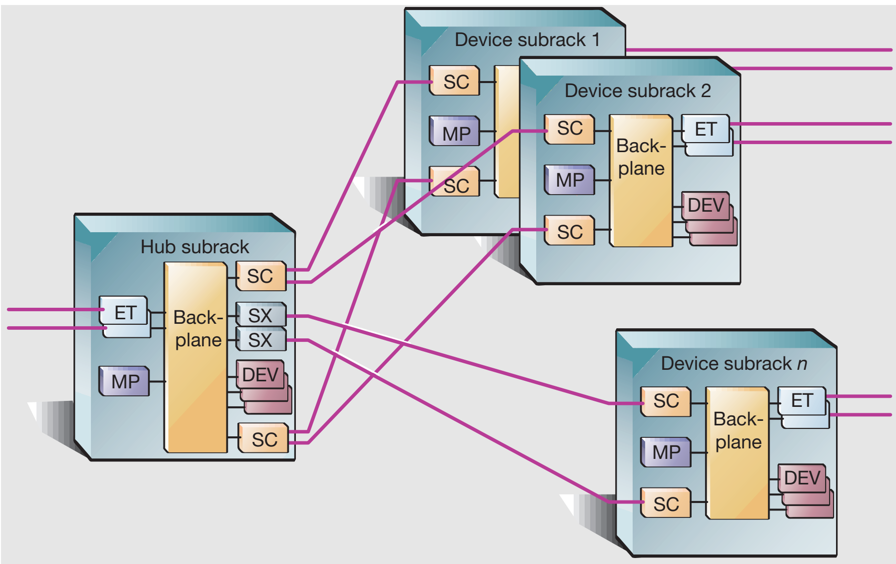

Execution platform

The Cello execution platform gives applications

- a scalable cluster of intercommunicating processors;
- a distributed, real-time operating system that supports robust application design and load balancing;
- a distributed, real-time database; and
- operation and maintenance (O&M) support.

The platform has a logical topology with a main processor cluster (MPC), built of multiple main processors. Each board in Cello contains a board processor. The MPC is the hub in a star topology with board processors at its edges. Each board processor may have one or more subordinate special processors connected to it (Figure 4).

Each processor in the main processor cluster communicates with all other processors in the cluster. Each board processor communicates with the main processor cluster but not with other board processors. The special-purpose processors (SP) solely communicate with the board processor with which they are co-located.

The configuration of the main processor cluster is not dependent on the location of processor boards in the node's subracks. For example, the main processor cluster may or may not have processors mounted in each subrack.

To achieve a high degree of robustness for critical application functions, Cello execution support is able to detect software and hardware failures and to take measures to preserve the execution of the software involved. If one of the MPC's processors fails or crashes, critical processes on that processor are automatically restarted on another of the MPC's processors. The services handled by the faulty processor are taken over by other execution resources of the MPC.

The main processor cluster also permits scalable processing; that is, by adding main processors to it, operators can increase the capacity of the MPC. Possible MPC configurations range from a non-redundant, single-processor MPC to a redundant, high-performance MPC with many main processors.

The initial distribution of processes to the MPC's processors is determined by configuration data for the MPC. If a fault occurs, the distribution scheme may temporarily differ from the default settings. Once the faulty processor is replaced, however, the initial distribution scheme is revived.

The term *main processor* denotes the role played by a member of the main processor cluster. Every main processor is equal—that is, there are no master-slave relationships between main processors.

Board processor denotes the role of a processor that executes software on a device board. Board processors are typically used as a control plane for user data channels and for controlling software. These processors, which are subordinate to the main processor cluster, are less robust than the MPC. For example, it is not possible to upgrade a board processor without disturbances, and if a board processor must be restarted, no actions are taken to preserve the execution of tasks on it. Software on board processors can access central core functions, such as the file system, the database, and the name server.

Cello also includes special-purpose processors and a certain amount of support for application software executing on them.

The channels between the processors, through which the processors communicate with one another, are called internal communication paths (ICP). The ICPs, which use the same physical path through the switch as user data, use AAL5 on ATM channels set up by Cello at startup. The paths are supervised by control system software. The transmitting processor identifies itself by means of a specific virtual path or virtual channel identity (VPI/VCI). The distribution of the execution platform is maintained by means of the ICPs. Thus, the application software can send the same kinds of signal between processes on the same processor as between different processors.

Cello contains a name server that is used for publishing the MPC's tasks and services. The name server enables software on any processor to find a process in the main process cluster.

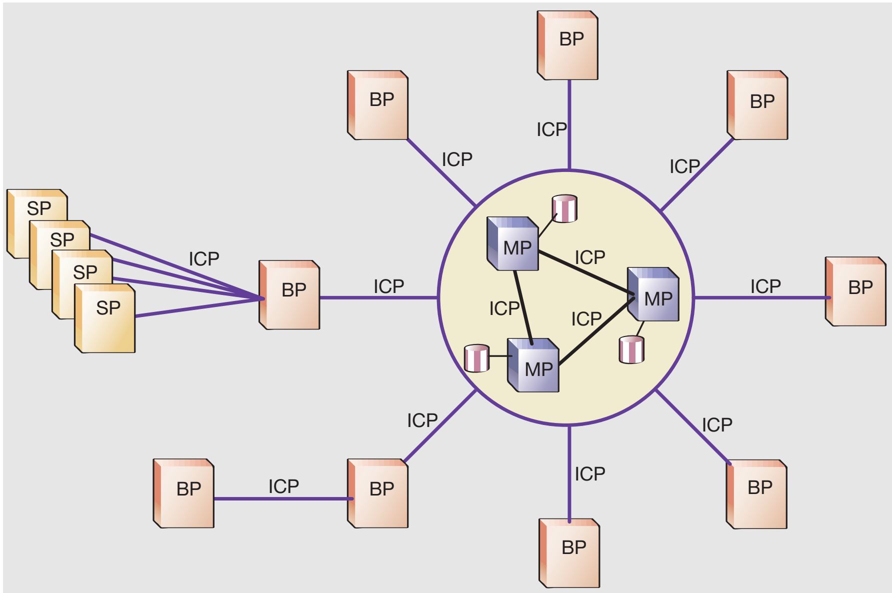

Figure 4
Control system structure of a Cello configuration with eight board processors (some with special processors) and a main processor cluster with three processors.

Upgrading software

Operators may upgrade software without interrupting service, although some degradation of service is expected. Basically, the new software is downloaded and automatically updated with current data, the "old" software is stopped, and the new software takes over. Processors need not be restarted except when certain of Cello's core functions are being upgraded.

To upgrade board processor software, the board processor must first be reconfigured and then restarted. Consequently, this upgrade disturbs service (on the board). Processors are taken out of service in a controlled manner, which allows the traffic load to cease before restart.

Startup in a Cello node is divided into phases. When each phase is complete, a particular set of services becomes available. A checkpoint at the end of each phase indicates which services are to be provided at that time.

During initial startup, the system usually has a primitive Cello configuration on disk that establishes connection to an operation and maintenance (O&M) function for loading software and configuration data. Thus, only a few configuration parameters (for example, the node identifier) need to be set locally.

When the software has been loaded, the node restarts, after which each board described in the configuration data is loaded with software from the disk and started, beginning with the Cello software and then the applications.

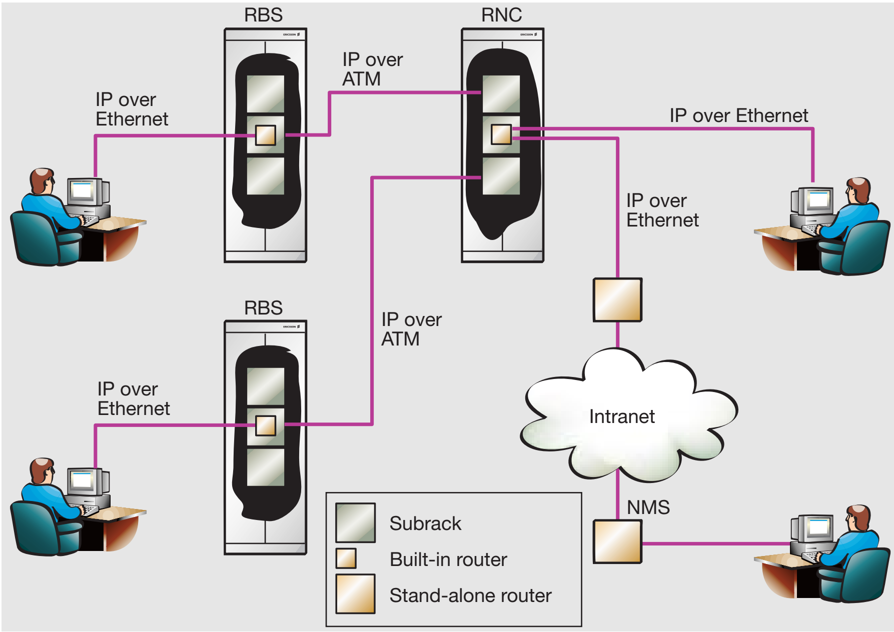

Figure 5
The operation and maintenance IP network.

Network management

A Cello-based node contains every function needed to manage the network element. Cello provides basic element-management services and a foundation for building node-specific element-management applications. The IP-based protocols and services to external managing systems are open and compatible with platforms from Microsoft, Sun Microsystems, Hewlett-Packard and others. Web-based technology is used for implementing element-management applications. The applications can be stored as Java applets in the network element and accessed from virtually any common personal computer or workstation using a standard Web browser.

Online documentation is supplied in hypertext markup language (HTML) and thus can also be read via a Web browser. CORBA is used as a foundation for the management interfaces. This gives high-level support for application programmers on both the network element side and on the side of the network managing system.

Remote management uses "in-band" communication (IP over ATM) between nodes; on-site management uses IP over Ethernet. For management purposes, operators can access every node from any node in a Cello-based network (Figure 5). The logical management interface, known as the managed object model (MOM), is defined in a formal standard notation called the interface definition language (IDL), which was developed for handling complex management interfaces. When designing client/server systems, designers adopted a "thin client" architecture, which enables easy adaptations of client applications for specific customer demands. To a large extent, the O&M functions in the network element have been implemented in Java.

The software fault-management function includes a fault repair escalation ladder, which allows several levels of fault detection and related actions. Faulty devices and software can automatically be isolated and taken out of service. A list handler and alarm-filtering functions, together with a mass alarm-suppression function, reduce the need for manual alarm analysis. Configurations are managed via an application-independent managed-object browser, via "command line" entry or command files used for the initial configuration. Moreover, Cello provides a real-time performance viewer, which is an element-management application with a graphical user interface (GUI) that presents performance data in real time. To ensure security, every user who accesses the management network is properly authenticated by means of user ID, password, and authentication tokens.

The main standards bodies considered for O&M were the Telecommunications Management Network (TMN), the Object Management Group (OMG), and the ATM Forum.

ATM network and connection handling

Cello supports the establishment and release of both permanent and switched connections in ATM and AAL2 transport networks. Permanent virtual connections (PVC) over a Cello node are set up by means of element-management procedures. Network connections that pass through several network elements require a network-management application. Cello's built-in functionality enables it to serve as an ATM switch in cooperation with any commercial ATM switch that uses

- B-ISUP (basic call); or
- a switched AAL2 connection (Q.2630).

Cello can also function as an SS7 signaling transfer point.

The application on the Cello platform uses an application program interface (API) to set up switched connections. Addresses in the API call can address ports within the same switching node or any other node in the network. The user can request different quality of service depending on the application, peak packet size, peak packet rate, and the service descriptor for voice, data, video, and so on.

Cello manages traffic by means of link admission control—the node's internal resources (switch cores and internal subrack links) are normally over-provisioned. A private ATM end-system address (AESA) identifies each connection point in a Cello network.

A signaling service interface is also provided for the applications. The signaling connection control part (SCCP) utilizes the signaling network (MTP3b), giving users the means of exchanging control messages between applications in different nodes. The following signaling services are supported:

- NNI-SAAL;
- MTP3b;
- SCCP;
- B-ISUP (basic call); and
- Q.2630 (AAL2 signaling).

Cello also contains a wealth of ATM transport functionality that is essential to the narrowband requirements of the cellular access network:

- circuit emulation of $n \cdot 64$ kbit/s ($n = 1$ to 30) with AAL1 on E1/J1/T1 links;
- fractional ATM, one fraction of $n \cdot 64$ kbit/s ($n = 1$ to 30) on E1/J1/T1 links;
- AAL2 multiplexing for E1/J1/T1, high-capacity E3/T3 and STM-1 links;
- AAL5 termination and support for the SAAL protocol;
- virtual path and virtual channel cross-connect;
- policing and shaping; and
- ATM O&M support for supervision of connections with F4/F5.

Thanks to its synchronization support, Cello can freely choose the reference clock of any ingress line on any multi-line exchange terminal.

ATM virtual channel or AAL connections that span several nodes are established on a "hop-by-hop" basis. That is, signaling is performed between adjacent nodes along the path of the connection, and routing is established in each node based on the destination ATM end-system address and the current availability of resources. The type of signaling used depends on the type of connection. The signaling service provides the signaling for the transport service. B-ISUP signaling is used for AAL0 and AAL5 connections, whereas Q.2630 signaling is for AAL2 connections (Figure 6).

AAL2 transport

Cello supports AAL2 transport, which guarantees narrowband ATM connections with short delay and good utilization of narrowband links (1.5 and 2 Mbit/s). The AAL2 connections, which are identified with a connection identity (CID), are multiplexed on established ATM virtual connections, called AAL2 paths, between adjacent AAL2 nodes.

TCP/IP services

Cello platform support for TCP/IP transport is a low-capacity solution intended for network management only. It enables communication with applications that are part of Cello nodes and use IP for transport as well as communication with external equipment that is attached to Cello nodes via external interfaces. The following functions in the Cello platform support IP communication:

- IP network layer with routing functionality;

Figure 6
Example of user access to the AAL bearer service in a three-node Cello network.

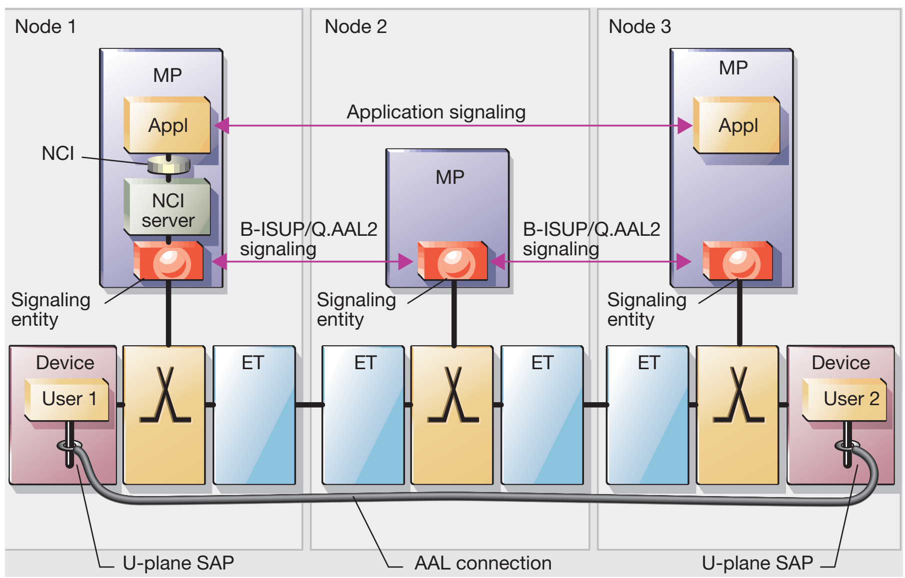

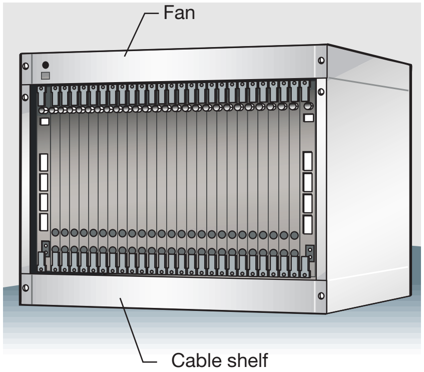

Figure 7
The Cello subrack.

- open shortest path first (OSPF) protocol for IP route identification;
- transmission control protocol (TCP) and user datagram protocol (UDP) transport services;
- dynamic host configuration protocol (DHCP); and
- domain name server (DNS).

Hardware

ATM switch

The ATM switch, called AMAX, was specially designed to derive a good cost-to-performance ratio from ATM networking products that

- need support for narrowband connections on narrowband transport links, T1/E1/J1, and fractions thereof; and
- put stringent requirements on delay and high utilization of the transport links. The AMAX provides this support, handling 25 high-capacity transport interface boards at data rates up to 622 Mbit/s.

The AMAX hardware consists of two specially designed hardware circuits: the ATM switch core circuit (ASCC), and the switch port interface circuit (SPIC).

The ASCC is contained on the switch-core board (SCB). The switch core, which has 28 ports operating at 800 Mbit/s, can handle 28 boards with user data rates of 622 Mbit/s, yielding approximately 16 Gbit/s duplex throughput. However, because there are only 27 boards in the subrack, one port is not used.

On each board the SPIC adapts the proprietary, eight-wire, backplane, switch-core interface to a standard UTOPIA interface, which is a common interface on commercial ATM layer parts. This makes it easy to build “any” type of device in Cello; for example, special protocol machines, special transmission-line interfaces, and so on.

Subrack equipment

The Cello subrack is designed to fit in BYB 501 and BYB 502 cabinets. It is also designed for mounting in industry-standard 19-inch racks. Forced-air cooling is normally required and provided by a low-height fan mounted above the Cello subrack. A cable shelf and air guide are placed under the Cello subrack. The total height of the complete subrack is 400 mm (Figure 7).

The external power supply is -48 V DC. A separate (redundant) distribution of -48 V power is provided to each board in the subrack. The -48 V power is converted on each board according to its specific voltage level and power requirements. Power distribution is dimensioned to feed 800 W per subrack.

The board size is 265 · 225 mm (height · depth) including backplane connectors. The minimum board spacing is 15 mm (one slot). Switch-core boards have 20 mm spacing—some boards may need 2 · 15 mm. The backplane has dedicated positions for switch-core and timing-unit boards. Device boards may be mounted in any position except those reserved for the switch-core boards. The switch core is located on the switch-core board at the edge of the subrack. For redundancy, the subrack can be equipped with two switch-core boards.

Each board must be equipped with a switch port interface circuit, which provides the interface to the switch core. For redundancy, the SPIC can support two core switches. For network synchronization, one timing unit is used per Cello node. If redundancy is required, two timing units may be used.

Cello platform hardware

A brief description of the functionality on each of the different boards follows below.

The switch-core board (SCB) contains the

BOX B, MANAGEMENT PROTOCOLS

IP-based protocols are used for element- and network-management communication. Consequently, each network element must have a unique IP address. Remote management uses “in-band” communication (IP over ATM); on-site management uses Ethernet.

Element managers and network managers use the IP-based inter-ORB protocol (IIOP)—the protocol used in CORBA for monitoring and configuring Cello nodes.

The hypertext transfer protocol (HTTP) is used for transferring HTML documents and Java applets from the file system in the Cello node to a managing system.

The simple network management protocol (SNMP) is used for monitoring and configuring IP routers and Internet protocols (MIB II). It can optionally be used for monitoring ATM resources.

The file transfer protocol (FTP) is used for transferring files to and from the Cello node’s file system, primarily for loading software, configuration data and other large volumes of data.

Telnet is used for local and remote access to the operating system shell in network elements. Telnet is also used for accessing managed objects.

switch core, clock distribution, fan supervision, -48 V power feed, and four intersubrack links (ISL), with a maximum user data bandwidth of 310 Mbit/s per link. The aggregate bandwidth of the four ISLs is limited to 622 Mbit/s. The switch extension board (SXB) includes four ISLs. These are used for building nodes with more than five device subracks.

The general-processor board (GPB) is a standard processor with its own RAM (128 MB) and hard disk (2 GB). It also contains Ethernet interfaces for managment access. The board, which uses one slot (15 mm), serves as a main processor in a main processor cluster.

The special-purpose processor board (SPB) is a generic, user-plane processing board used for protocol handlers—for example, for radio link protocols.

The timing-unit board (TUB) contains the system clock oscillator, which is very stable (50 parts per billion accuracy and more than one hour holding time). The system clock oscillator is synchronized with any of the exchange terminals.

The exchange-terminal boards (ETB) provide interfaces for ATM cell transport over various kinds of transmission line. Different ETB are needed for implementing adaptations to different physical media and different standards:

- ETM1—1.5 Mbit/s (T1, JTI.431-a) and 2 Mbit/s (E1). The ETM1 has eight ports.
- ETM23—6.3 Mbit/s (J2), 34 Mbit/s (E3) and 45 Mbit/s (DS3). The ETM23 has four ports.
- ETM4—155 Mbit/s (STM-1/STS-3c), optical or electrical link, 51 Mbit/s (STS-1). The ETM4 has two ports.

The ETM1 is one circuit board that can be configured, per port, for 1.5 or 2 Mbit/s links for ATM or STM. When configured for ATM, the ETM1 can use a fraction of the 1.5 to 2 Mbit/s capacity ($n \cdot 64$ kbit/s, where $n = 1$ to 30). This configuration also supports AAL2 multiplexing on the board.

When configured for STM, the ETM1 uses AAL1 circuit emulation ($n \cdot 64$ kbit/s combinations within the 1.5 or 2 Mbit/s, where $n = 1$ to 30).

Conclusion

Cello is a telecom platform that can be used for several different products. Its robust, real-time control system and efficient ATM transport system guarantee cost-effective applications with support for circuit-switched and packet-switched services alike.

The Cello execution platform gives applications a scalable cluster of intercommunicating processors, a distributed, real-time operating system, a distributed, real-time database, and O&M support.

The transport system is built on AMAX, an ATM switch designed especially to give a good cost-to-performance ratio, even in small systems.

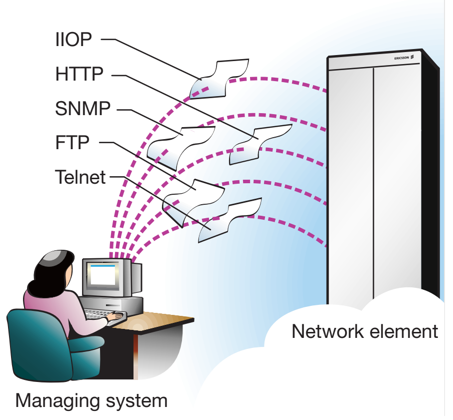

Figure 8
The protocols that a managing system uses to manage a Cello node (the managed system).