AAL2 switching in the WCDMA radio access network

Bo Karlander, Szilveszter Nádas, Sandor Rácz and Jonas Reinius

New switching technologies are needed in the access network of third-generation mobile networks to provide a cost-effective transmission of different kinds of services, such as AMR-coded voice and large bandwidth data. The use of ATM and AAL2 switching techniques in the traffic concentration nodes can significantly reduce the need for link capacity in the access network. The most important sources of these savings are the statistical fluctuation of the number of AAL2 connections, the fluctuation of the number of users at a base station due to mobility, and the granularity of the ATM virtual channel cell rate. Ericsson's offering of AAL2 switching technology enables operators to maximize these gains in the WCDMA radio access network. Indeed, in large networks, this technology has the potential to triple the capacity of the transmission link.

The authors describe the advantages of AAL2 switching in the WCDMA radio access network. They also include results from a study comparing bandwidth requirements when traffic between node B and the RNC is aggregated either at the ATM layer, by means of switching AAL2 paths using ATM cross-connect, or at the AAL2 layer, by means of dynamically switching AAL2 connections. The authors also describe the general architecture for terminating and switching AAL2 connections in Ericsson's WCDMA RAN node products.

Introduction

Third-generation mobile networks offer a wide variety of services, such as voice, circuit-switched data, and packet-switched data at bit rates ranging from a few kbit/s up to 384 kbit/s (eventually up to 2 Mbit/s). In this environment, highly adaptive transfer and switching methods are needed to deliver different kinds of services in a cost-effective and high-quality fashion. This applies to both the

- radio network layer, which is responsible for transferring data over the WCDMA air interface; and
- transport network layer, which is responsible for transferring data between the nodes (node B and RNC) of the radio access network.

As specified by the Third-generation Partnership Project (3GPP), the transport network layer for the WCDMA radio access network (RAN) is to use ATM transport network technology and protocols. ATM adaptation layer type 2 (AAL2) technology is used for the dominant part of data transfer.

The AAL2 protocol enables several user (radio network layer) connections to be multiplexed flexibly and efficiently on a common ATM virtual channel connection (VCC) between two nodes. The use of AAL2 switching in intermediate nodes has the potential to yield significant statistical multiplexing gains on transmission links that carry aggregated traffic for multiple nodes without loss of control over the quality of service of individual connections.

BOX A, TERMS AND ABBREVIATIONS

3GPP	Third-generation Partnership Project	MAC	Medium access control
AAL2	ATM adaptation layer type 2	Node B	Radio base station, RBS
AMR	Adaptive multirate (voice codec)	O&M	Operation and maintenance
ATM	Asynchronous transfer mode	PS 64	Packet-switched data at 64 kbit/s
CAC	Connection admission control (algorithm)	PS 384	Packet-switched data at 384 kbit/s
CBR	Constant bit rate	PSTN	Public switched telephone network
CCH	Common transport channel	Q.2630	ITU AAL2 signaling protocol
CID	Connection identifier	QoS	Quality of service
CN	Core network	RAB	Radio access bearer
CNA	Concentration node area	RAN	Radio access network
CPP	Connectivity packet platform (formerly called Cello packet platform)	RLC	Radio link control
CPS	Common-part sub-layer (packet)	RNC	Radio network controller
CS 64	Circuit-switched data at 64 kbit/s	RXI	CPP-based aggregation node
DCCH	Dedicated control channel	SDH	Synchronous digital hierarchy
DCH	Dedicated transport channel	SL1, 2, 3	Switching level one, two, three
DRNC	Drift RNC	SRNC	Serving RNC
E1	ETSI 2 Mbit/s line interface	STM	Synchronous transfer mode
GoS	Grade of service	TTI	Transmission time interval
Iu	Core network-to-radio network interface (3GPP)	UE	User equipment
Iub	Node B-RNC interface (3GPP)	UP	User plane
Iur	RNC-RNC interface (3GPP)	VC	Virtual channel
LAC	Link admission control	VCC	Virtual channel connection
		VCI	Virtual channel identifier
		VPI	Virtual path identifier
		WCDMA	Wideband code-division multiple access

WCDMA RAN—transport network architecture and protocols

WCDMA RAN

Figure 1 gives a schematic view of a WCDMA network, which consists of user equipment (UE), the WCDMA terrestrial radio access network (WCDMA RAN), and the core network.

The WCDMA RAN handles all tasks that relate to radio access control, such as radio resource management and handover control. The core network, which is the backbone of WCDMA, connects the access network to external networks (PSTN, Internet). The user equipment (mobile terminal or station) is connected to radio base stations (node B) over the WCDMA air interface (*Iu*). During soft handover, one UE can communicate with several node Bs simultaneously.

According to the WCDMA RAN specifications drafted by the 3GPP, all radio network functions and protocols are separate from the functions and protocols in the transport network layer. The transport network layer provides data and signaling bear-

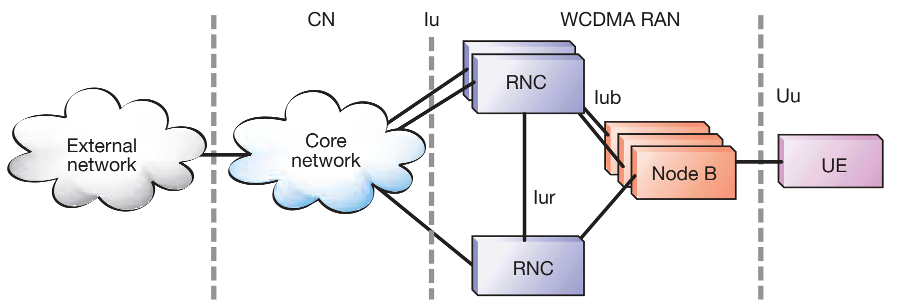

Figure 1
Schematic view of a WCDMA network.

ers for the radio network application protocols between RAN nodes, and includes transport network control-plane functions for establishing and releasing such bearers when instructed to do so by the radio network layer.

The initial WCDMA RAN specifications stipulate that the transport network layer must be based on ATM and AAL2 technology. However, Release 5 of the 3GPP specifications also includes the option of an IP-based transport network.

The focus of this article is on the design of ATM- and AAL2-based transport networks with particular emphasis on the *Iub* interface between the RNC and node B.

AAL2 at the *Iub* interface

Figure 2 shows the ATM and AAL2-based protocol stack at the *Iub* interface for transferring data streams on common transport channels (CCH) and dedicated transport channels (DCH) to the air interface.

The retransmission mechanism of the radio link control (RLC) protocol ensures reliable transmission of loss-sensitive traffic over the air interface. The RLC protocol is used by signaling radio bearers and by radio bearers for packet-switched data services, but not by radio bearers for circuit-switched services.

The medium access control (MAC) protocol forms sets of transport blocks in the air interface and schedules them according to the timing requirements of WCDMA. Each scheduled period, called a transmission time interval (TTI), is 10 ms in length or multiples thereof.

WCDMA radio connections, or radio access bearers (RAB), have bit rate values between 8 and 384 kbit/s. The size of the MAC transport block sets and length of the TTI are RAB-specific.

For data transfer over the *Iub* interface, the MAC transport block sets are encapsulated into *Iub* frames according to the *Iub* user-plane (UP) protocol for CCH or DCH data streams. Each *Iub* user-plane data stream needs a separate transport network connection between the RNC and node B. The transport network thus establishes one AAL2 connection for each data stream. In Figure 2, the AAL2 switch (optional) is used for building aggregating transport networks.

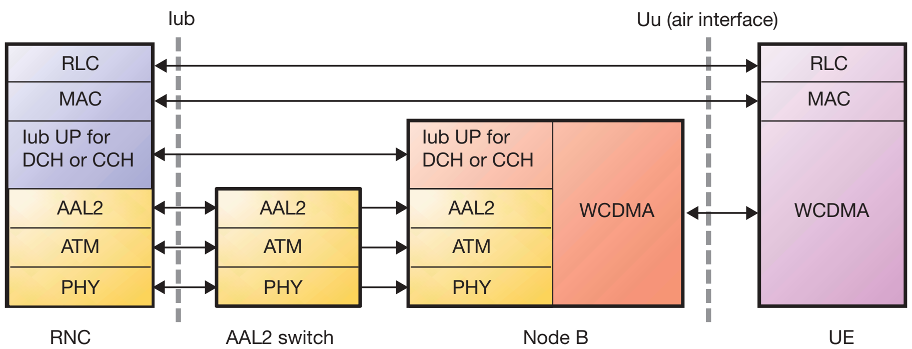

Figure 2
User-plane data transfer between the RNC and node B.

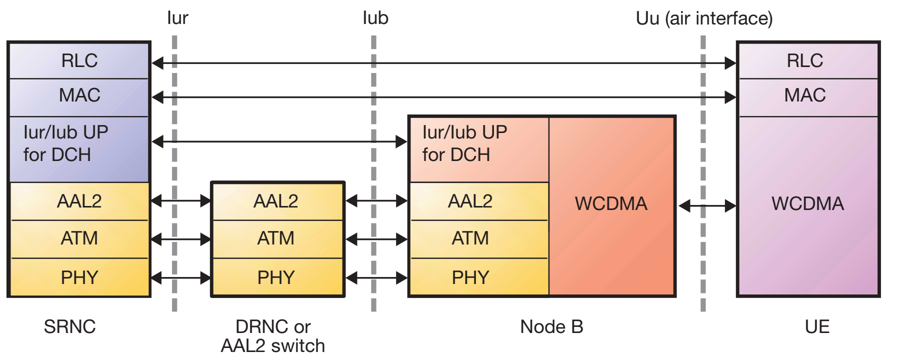

Figure 3
User-plane data transfer between the SRNC and node B.

AAL2 at the *Iur* interface

Mobile users sometimes move from radio cells controlled by a serving RNC (SRNC) to radio cells controlled by another RNC, designated drift RNC (DRNC). The *Iur* interface between the RNCs allows the SRNC to maintain contact with mobile users connected via the air interface to one or more cells controlled by another RNC. The user-plane protocol for the DCH data streams is established between the SRNC and node B. If the DRNC incorporates AAL2 switching, an AAL2 layer connection can be established from the SRNC to the node B—that is, without AAL2 connection termination in the DRNC—minimizing the transfer delays via the DRNC. This is particularly important for DCH data streams, which have strict timing requirements for soft handover. If an AAL2 switching network is built to interconnect multiple RNCs and node Bs, then every AAL2 connection for DCH data streams can be set up directly between the SRNC and node B without passing the DRNC. This configuration further reduces transmission costs. The AAL2 control plane (not shown in Figure 3) is terminated in every AAL2 switching node.

User frames are segmented and packed into AAL2 common-part sub-layer (CPS) packets, which are multiplexed into ATM cells (Figure 4). The AAL2 payload can vary in length (up to 45 bytes). The AAL2 header is 3 bytes in length. All ATM cells are 53 bytes in length, including a 5-byte header.

Thanks to AAL2 multiplexing, the AAL2 packets from several AAL2 connections can be transported on one ATM virtual channel connection (VCC). Each ATM cell on the VCC can carry AAL2 packets from different AAL2 connections. The connection identifier (CID) field in each AAL2 packet header identifies the AAL2 connection to which the packet belongs, much the same as the virtual path identifier (VPI) and virtual channel identifier (VCI) fields in the ATM cell header identify the ATM virtual channel connection. [1, 2]

QoS requirements and admission control

The WCDMA RAN transport layer services must meet stringent quality of service (QoS) requirements. The most important measure of QoS performance in the *Iub* and *Iur* interfaces is maximum packet delay.[3]

To satisfy QoS requirements, AAL2 connection admission control (CAC) is executed before a new AAL2 connection is set up in the system. Connection admission decisions are based on the traffic descriptors and QoS requirements assigned to the connections. The AAL2 admission control procedure allocates bandwidth resources (from available virtual channel and path identifier resources) to AAL2 connections in the transport network. If the requisite amount of resources is not available to accommodate a new connection, it is rejected.

If AAL2 connections are transported in end-to-end virtual circuits with resource allocation, AAL2 connection administration control is only executed at the end points of the virtual circuits (node B and RNC). How-

Figure 4
Assembly of ATM cells.

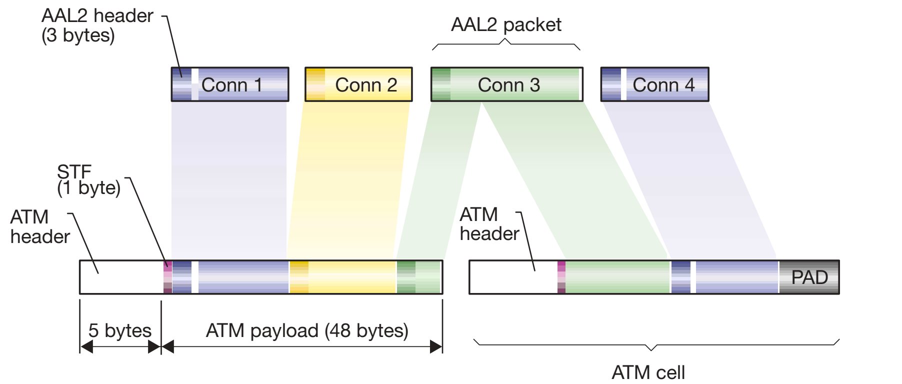

ever, if the resources along the path of an AAL2 connection are not allocated end-to-end, the CAC decisions are replaced by or based on hop-by-hop link admission control (LAC) decisions at every AAL2 switch along the path.

Transport network functions

Node types

The Ericsson WCDMA RAN system is composed of three kinds of traffic-handling nodes:

- Different versions of node B are available for indoor and outdoor placement and to satisfy different needs for the air interface and transport network capacity. One node B can be configured to serve as a transport network hub using ATM cross-connect and AAL2 switching techniques to aggregate traffic to and from other node Bs.
- The radio network controller (RNC) is a modular multi-subrack node. It is available in various sizes to satisfy different capacity needs.
- The RXI is a transport network node that provides ATM cross-connect, AAL2 switching, and IP router services. This fault-tolerant node is a single-subrack design with numerous interface options and a high-capacity switch core.

The WCDMA RAN system also includes various other nodes not described in this article for operation and maintenance (O&M) support.

The system platform

All of Ericsson's WCDMA RAN nodes are based on the same carrier-class technology—the connectivity packet platform (CPP, formerly called Cello packet platform). Modular and robust in design, CPP is is characterized by a multiprocessor control system with multiple processor levels, and its use of cell-switching technology to internally interconnect all types of processor boards, external interface boards and application-specific boards. CPP also includes functionality for terminating and switching ATM and AAL2 connections and for terminating and routing IP traffic.[4]

Switching and termination of ATM and AAL2 connections in CPP

The CPP solution for terminating and switching ATM and AAL2 traffic is at the heart of Ericsson's ATM and AAL2 transport solutions. The node function that provides the through-connection of an ATM virtual channel connection is called the ATM VC cross-connect. Similarly, the node function that provides the through-connection of an AAL2 connection is denoted AAL2 switching.

The 3GPP does not specify how ATM layer connections are to be established and released in the WCDMA radio access network. Ordinarily, permanent virtual channel or virtual path connections are configured using network management actions. In the 3GPP specifications, the establishment and release of AAL2 connections is to be controlled dynamically and in real time by user requests—that is, by the functions of the radio network layer. The AAL2 connections in the WCDMA radio access network are controlled by Q.2630 signaling between the RNC and node B. This signaling can thus be used for setting up connections in a switching transport network between the node B and RNC. The network can be built up of multiple CPP nodes—for example, a tree structure of node Bs combined with pure transport aggregating nodes, such as the Ericsson RXI820 (Figure 5).

The internal cell-switching architecture of the node allows the terminating function of each ATM VCC to be distributed to the processor board or application-specific de-

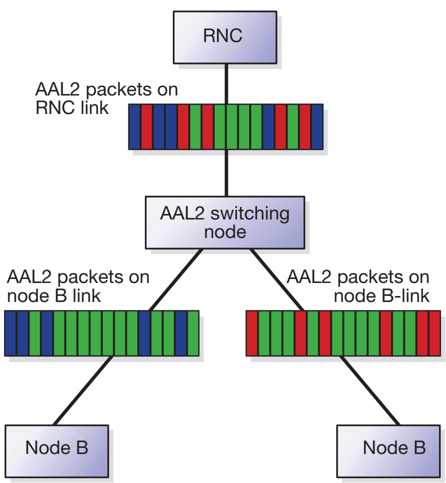

Figure 5
AAL2 switching network.

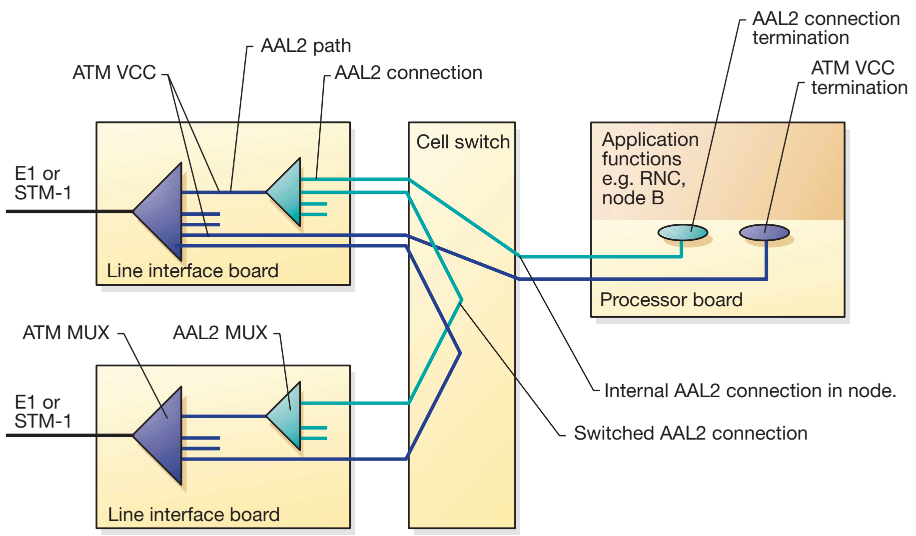

Figure 6
Switching and termination of ATM virtual channel and AAL2 connections in CPP nodes.

vice board on which the application function that uses the connection resides. The ATM line interface boards thus forward ATM cells in either direction between the external line interface and cell switch. The cell switch in CPP transfers ATM cells to and from connection-terminating boards in the node, or to and from external interface boards for cross-connected ATM VCCs.

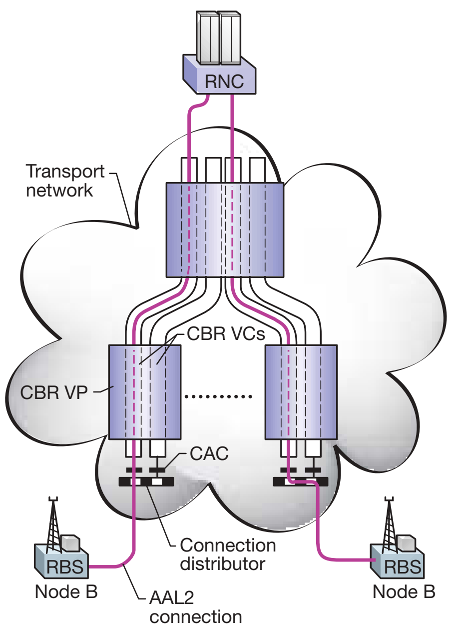

Figure 7
VC-CC model.

Each ATM VCC used by an AAL2 path is terminated by an AAL2 multiplexer function (located on the line interface board) that handles the AAL2 CPS layer, multiplexing and demultiplexing the AAL2 packets to and from ATM cells on the AAL2 path (an ATM VCC, see Figure 6).

For AAL2 connections that terminate within the node, each AAL2 packet is carried via the cell switch directly between the line interface board and the application board that terminates the AAL2 connection.

By contrast, for AAL2 connections that are switched through the node, each AAL2 packet is forwarded via the cell switch directly between the incoming and outgoing line interface boards.

AAL2 packets are switched by hardware on the line interface boards and via the high-capacity cell switch in the same way as ATM cells are switched. Therefore, every traffic-handling RAN node based on CPP technology can efficiently perform switching within the AAL2 layer.

WCDMA RAN transport network topologies and switching alternatives

Although it is possible to connect a node B to its RNC via a direct, physical point-to-point connection, more commonly an *Iub* connection is established via one or more intermediate aggregation nodes. Connections from multiple node Bs can, in this way, be multiplexed onto the same physical link interface to the RNC.

The basic methods of multiplexing and concentrating traffic in the WCDMA RAN transport network are

- physical-layer multiplexing, using SDH network technology—one STM-1 interface (155 Mbit/s) to the RNC can multiplex, for example, 63 E1 (2 Mbit/s) physical-layer connections to as many node Bs;
- ATM VP or virtual channel cross-connect, using ATM network technology—one STM-1 interface to the RNC can multiplex several virtual channels representing ATM layer connections to many node Bs. The number of VCCs is the same at the RNC interface as the aggregated number of VCCs at the node B interfaces; and
- AAL2 switching, using AAL2 network technology—one STM-1 interface to the RNC can multiplex AAL2 connections to multiple node Bs on a common group of AAL2 paths between the RNC and an intermediate node with AAL2 switching capability. Other AAL2 paths are established between the intermediate node and each node B. The AAL2 switching method is combined with ATM VC cross-connect in the intermediate node of end-to-end ATM VCCs between the RNC and each node B; for example, for signaling and O&M access.

These methods are typically combined—that is, AAL2 switching is introduced on top of ATM cross-connect, which is introduced on top of SDH multiplexing.

To accommodate AAL2 switching in intermediate nodes, resources must be allocated to process signaling for the set-up and release of AAL2 connection and for multiplexing and demultiplexing the AAL2 user-plane.

Compared to an intermediate node, which merely cross-connects AAL2 path VCCs, each AAL2 switching node introduces some delay during connection set-up and transport of the user-plane. The network designer must weigh these costs against possible savings in transmission capacity and network management, and decide to what extent AAL2 switching can be employed.

We have studied the gains in transmission capacity from AAL2 switching at different aggregation levels for *Iub* traffic, comparing two cases of network dimensioning:

- VC-CC—use of ATM VC cross-connect of AAL2 path VCCs in aggregation nodes (Figure 7).
- AAL2—dynamic switching of AAL2 connections between multiple downlink AAL2 paths and fewer uplink AAL2 paths in aggregation nodes (Figure 8).

All AAL2 paths are assumed to be configured as constant bit rate (CBR) virtual channels with a fixed, guaranteed capacity. The connection admission control function, which operates on each AAL2 path, needs to know the total capacity of the AAL2 path to calculate the number of different kinds of AAL2 connections that can be allowed on each AAL2 path while maintaining quality of service for every connection.

Benefits of AAL2 switching on the *Iub* interface

Evaluating the performance of AAL2 switching

To evaluate the performance of AAL2 switching, we used a fast and accurate CAC algorithm that takes into account the properties of traffic across the *Iub* interface.[5] The transport network layer grade-of-service (GoS) requirement of each RAB type is met if the blocking probability of its connection remains below a given threshold, for example, 0.3% for voice. Besides the offered traffic (measured in Erlang), the blocking prob-

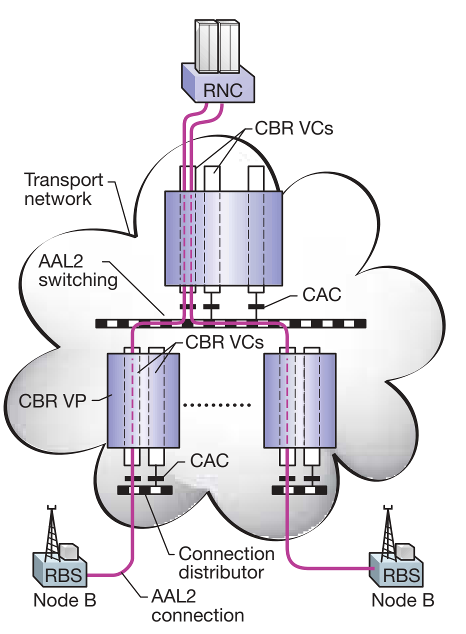

Figure 8
AAL2 switching model.

TABLE 1, CONSIDERED TRAFFIC MIX, USER LOAD IN MERLANG AND GOS REQUIREMENT

mErlang per user	Standard	Data-oriented	Contains PS 384	Target GoS (%)
Voice	20.25	17.00	25.00	0.3
CS 64	2.25	1.00	-	0.3
PS 64	3.30	5.00	1.00	0.7
PS 384	-	-	0.04	4.0

abilities also depend on the admission control algorithm, and on the number of available AAL2 paths (ATM VPs) between the two nodes.

The connections were generated randomly in a simulator, and the packet-level traffic descriptors were attributes of the generated connections. When a new connection was generated in the simulator, the CAC algorithms executed in the appropriate nodes in the access network, and we measured the blocking probabilities for each service.

Traffic parameters

During the simulation, different RAB types were considered for the following services: AMR-coded voice, circuit-switched data at 64 kbit/s (CS 64), packet-switched data at 64 kbit/s (PS 64) and packet-switched data at 384 kbit/s (PS 384). The evaluation also took into account all common channels in a cell and a dedicated control channel (DCCH) connection for each RAB.

We ran simulations for different kinds of traffic mixes (Table 1). The results represent a standard mix of voice (12.2 kbit/s AMR-coded), CS 64 and PS 64. Since many operators do not intend to provide PS 384 bearers when network load is heavy, it was relevant to study traffic mixes that excluded this rate. The results from the other traffic mixes also favored AAL2 switching over VC cross-connect.

For each *Iub* interface, about 200 kbit/s was allocated for non-AAL2 control-plane and O&M signaling.

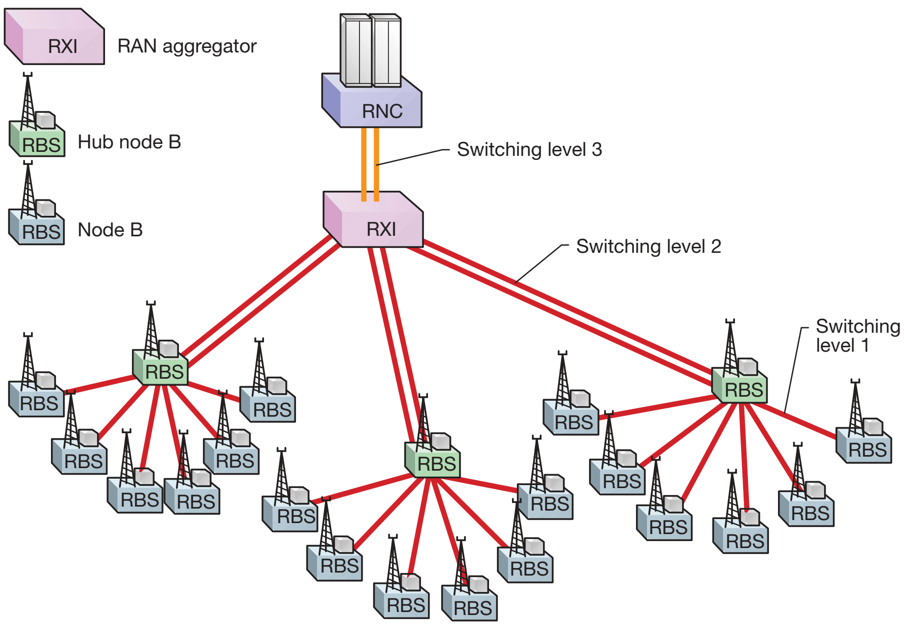

Figure 9
Tree topology of the WCDMA RAN.

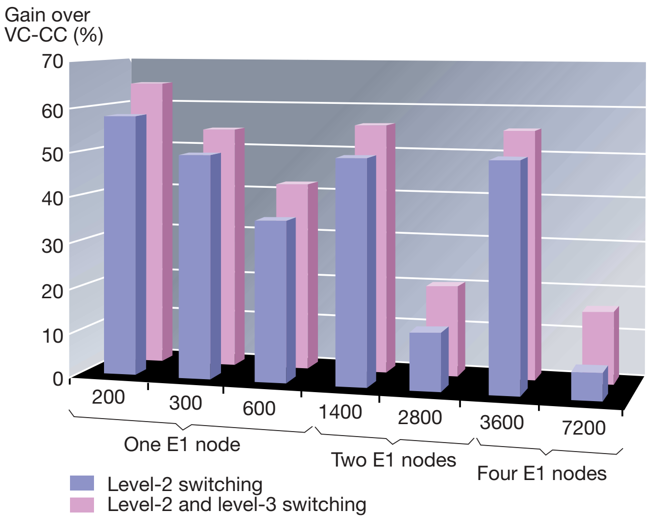

Figure 10
Reduced need for capacity for different number of users and switching layers in the 3x6 topology.

Note: While the gains from AAL2 switching are typical, the results apply to a specific traffic mix. The actual gains for other traffic volumes, mixes of voice and data, and other activity factors for voice and packet data might vary.

Statistical gain at different traffic loads

We obtained the numerical results by simulating the network at the connection level.

Low traffic volume

Compared to VC-CC switching, the greatest reduction in bandwidth using AAL2 switching can be expected when traffic is light—for example, when only a small part of the allocated bandwidth of the VCs is used.

Let us consider a tree network topology with three AAL2 switching nodes at switching level 2 (SL2), and six node Bs connected to each concentration node (Figure 9). We call this topology the 3x6 topology. Initially, the link capacity at the lowest level is one E1 per node B. The number of users is less than 600 per node B (600 is the maximum number, assuming a standard or data-oriented mix).

For VC-CC switching, the required bandwidth is 14 Mbit/s for a switching node in SL2, and 42 Mbit/s for a switching node in SL3.

We next increased the number of users served by a node B. Obviously, at some point, as the number of users grows, more capacity is needed to connect node B. Figure 10 shows how AAL2 switching reduces the need for capacity in the link compared to VC-CC switching at SL2 and SL3 for different numbers of users per node B using the standard traffic mix. When the number of users exceeds 600, two E1 access links are used. Likewise, when the number exceeds 2,800, four E1 access links are used.

The gain in capacity from using AAL2 switching is calculated as follows:

$$\text{Gain} = (C_{\text{VC-CC}} - C_{\text{AAL2}}) / C_{\text{VC-CC}}$$

where CVC-CC is the entire link capacity (aggregated capacity at the considered level) needed for VC-CC, and CAAL2 is the link capacity needed for AAL2 switching.

As can be seen, applying AAL2 switching in SL2 reduces the consumption of link capacity by 35-60%. The gain in capacity obtained from AAL2 switching is especially high (around 60%) when the use of SL1 links is low. The gain is even greater (up to 65%) at SL3.

Thanks to the large gain in capacity from AAL2 switching, the operator can install less capacity at SL2 and SL3 when first introducing service, and later upgrade the link capacity between switching nodes.

Fully loaded system

Operators generally try to avoid loading the network to its limit. Nevertheless, they do need to determine the capacity limit and in-

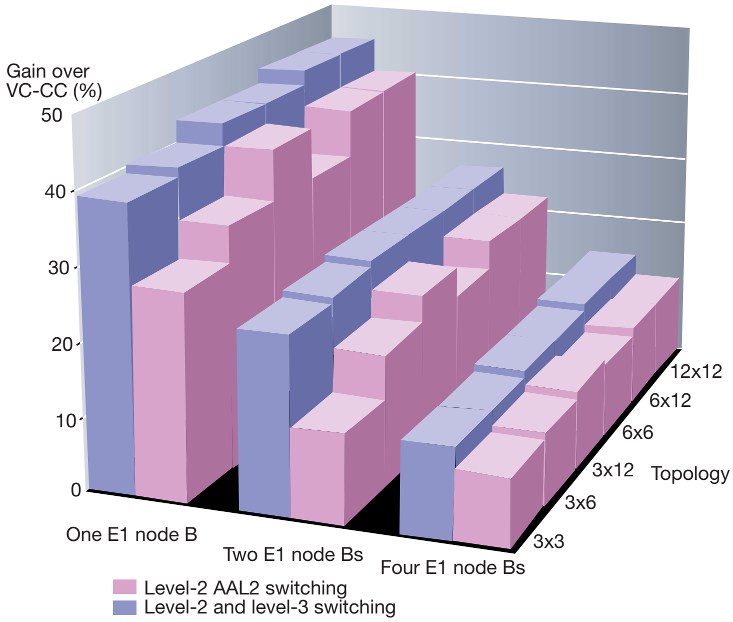

Figure 11
Fully loaded system. Gain in level-2 and level-3 link capacity from AAL2 switching.

vestigate traffic cases close to this limit. QoS problems are more likely to occur when load is heavy. Figure 11 shows the gain in capacity for the different network sizes with up to twelve switching nodes in the second level, twelve node Bs connected to each switching node (12x12 topology), and a fully loaded system.

Compared to VC-CC switching, the gain in capacity increases as more node Bs are added per SL2 switching node. Due to a homogenous distribution of traffic among node Bs, the gain from AAL2 switching at SL2 is the same for configurations with the same number of node Bs per switching node (for example, 3x12, 6x12 and 12x12). By

TABLE 2, USER DISTRIBUTION THROUGHOUT THE DAY

Users per node B	CNA1	CNA2	CNA3
Morning	1,400	2,800	2,800
Mid-day	2,800	1,400	2,800

applying AAL2 switching at L2 and L3, the gain increases slightly as a function of the number of switching nodes.

The 3x6 topology used two E1s per node B (Figure 11): An 18 percent gain in capacity was obtained when AAL2 switching was applied at level 2, and 25%, when applied at levels 2 and 3.

Gain from changing traffic distribution

In general, the distribution of traffic among node Bs is not homogenous. A concentrating switching node aggregates the traffic of several node Bs. Therefore each switching node has an associated serving area, which we call the concentration node area (CNA).

Let us assume that a CNA covers an office area and a residential area. Applying VC-CC switching, we attempt to establish as many VC connections between node Bs and the RNC as are needed to serve the sum of the maximum traffic of each node B. In this case, it does not matter that peak traffic occurs at different hours in the office and residential areas.

However, AAL2 switching allows the concentrated link capacity to be dimensioned for the sum of the actual traffic in the CNA.

Consider a 3x6 topology, in which the node Bs are connected by two E1 links. The maximum number of users in the switching node areas is 2,800 per node B assuming a standard mix of traffic. Let us also assume that due to user mobility, the number of users per node B changes over time (Table 2). Compared to VC-CC switching, the gain from AAL2 switching is 26% on SL1 and 32% on SL2.

TABLE 3, COMPARISON OF SWITCHING ALTERNATIVES

	VC-CC	AAL2
Statistical multiplexing capacity reduction	No	High (7 to 65%)
Reduction if traffic distribution changes	No	Yes (26 to 32%)
QoS guarantee	Yes, with CAC for each AAL2 path configuration	Yes, with CAC for each connection setup
Number of VCs minimized	No	Yes

Conclusion

The performance of AAL2 switching is superior to that of VC-CC switching. In several network scenarios that use AAL2 switching, the increased cost of processing AAL2 signaling is more than compensated for by significantly increased efficiency in the transmission network, and reduced costs of transmission.

The savings are greater when there is little traffic in the network—precisely when the network operator needs most to keep costs down. AAL2 switching also considerably reduces the need for link capacity in a fully loaded network with traffic concentration.

By applying AAL2 switching, operators can significantly reduce the need for link capacity since the network can adapt to changes in traffic levels.

REFERENCES

1 ITU-T Rec. I.363.2, "B-ISDN ATM adaptation layer type 2 specification"
2 ITU-T Rec. Q.2630.1 and Q.2630.2, "AAL Type 2 Signaling Protocol (Capability Set 1 and 2)"
3 3GPP TR 25.853, "Delay Budget within the Access Stratum"
4 Kling, L.-Ö., Lindholm, Å., Marklund, L. and Nilsson, G. B.: CPP—Cello packet platform. Ericsson Review Vol 79(2002): 2, pp. 68-75.
5 Sz. Malomsoky, S. Rácz and Sz. Nádas, "Connection Admission Control in UMTS Radio Access Networks", accepted for publication to Computer Communications, Elsevier Science, 2002.