INTEGRATED COMMUNICATIONS MANAGEMENT OF BROADBAND NETWORKS

Crete University Press, Heraklio, Greece
ISBN 960 524 006 8

Edited by David Griffin

First published in 1996 by
CRETE UNIVERSITY PRESS
Foundation for Research and Technology
P.O. Box 1527, Heraklio, Crete, Greece 711 10
Tel: +30 81 394235, Fax: +30 81 394236
email: pek@iesl.forth.gr

The ICM consortium consists of the following companies:

Alcatel ISR, France
Alpha SAI, Greece
Ascom Monetel, France
Ascom Tech, Switzerland
Centro de Estudos de Telecommunicações, Portugal
Cray Communications Ltd., United Kingdom (Prime contractor)
Danish Electronics, Light & Acoustics, Denmark
De Nouvelles Architectures pour les Communications, France
Foundation for Research and Technology - Hellas, Institute of Computer Science, Greece
GN Nettest AS, Denmark
National Technical University of Athens, Greece
Nokia Corporation, Finland
Queen Mary and Westfield College, United Kingdom
Unipro Ltd., United Kingdom
University College London, United Kingdom
University of Durham, United Kingdom
VTT - Technical Research Centre of Finland

Chapter 2

ATM background

Editor: David Griffin
Authors: Eric Scharf, David Griffin

Much of the work of the project addresses the management of Asynchronous Transfer Mode (ATM) networks, and so some background knowledge of ATM is required by the reader. The introduction presented here together with the references listed will provide the reader who has not studied ATM in detail with an introduction to the network issues addressed by the management systems described in the later chapters.

2.1 What is ATM?

ATM [2.1] is a means of communicating digitally-encoded information between two or more nodes. Imagine that between a transmitter and receiver node there is a conveyor belt whose sections or slots can be filled with cells, each of which has a header and a payload. When an information-carrying cell (*assigned* cell) arrives at the transmitter node, it is put into the next available slot on the conveyor belt. If assigned cells arrive more quickly at the transmitter node than the conveyor belt can remove them, then the assigned cells will have to be queued in order of arrival until a slot is ready to transport them. If assigned cells arrive more slowly at the transmitter node than the conveyor belt is able to remove them, then the spare slots on the conveyor belt are filled with *empty* cells. The conceptual conveyor belt is the *transmission link* between the transmitter and receiver nodes.

Each cell consists of 53 octets, each octet having eight bits, five octets are used for header information and 48 octets for the payload or information field.

Within a link, assigned cells belong to uni-directional virtual communications channels, or connections. Assigned cells belonging to a particular connection will normally be interspersed with empty cells and assigned cells belonging to other connections. Each assigned ATM cell belonging to an individual connection will have the same header information, although, of course, the payload will usually differ from cell to cell.

ATM virtual connections usually traverse more than one link, and are of two types: *Virtual Channel Connections* (VCCs) and *Virtual Path Connections* (VPCs). A VCC is a one-way, end-to-end route through the network and is made up of a sequence of *Virtual Channel Links* (VCLs). Cells belonging to a VCC are routed from one VCL to another at intermediate nodes called *Virtual Channel Switches*. Collections of VCLs interconnecting the same two Virtual Channel Switches and traversing the same path through the network are grouped together in a VPC. A VPC is formed from a sequence of *Virtual Path Links* (VPLs) - sometimes referred to as *Virtual Paths* (VPs) - spanning a single transmission link. Cells belonging to a VPC are routed from one VPL to another at *Virtual Path Cross-Connects*.

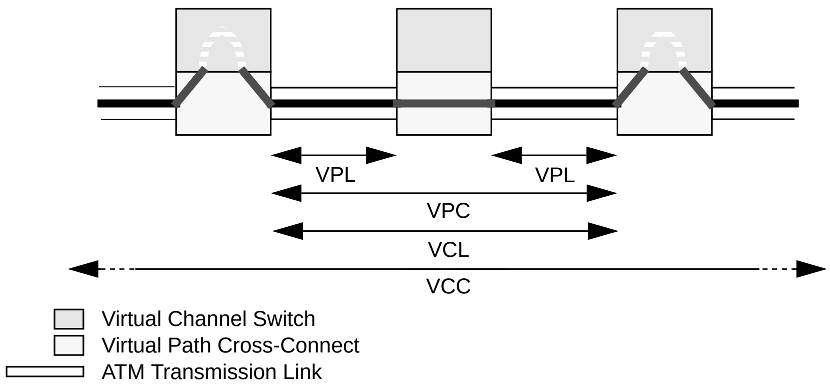

Figure 2.1 Span of VPLs, VPCs, VCLs and VCCs

So there is a two level hierarchy of connections in ATM: Virtual Channel Connections are multiplexed together over sequences of Virtual Path Connections; in turn, VPCs are multiplexed together on the ATM transmission links (Figure 2.2).

Cell header subfields include the *Virtual Channel Identifier* (VCI) and the *Virtual Path Identifier* (VPI). The VCI remains constant within a VCL (the span of a VPC) and the VPI remains constant over a VPL (the span of a single transmission link), neither have end-to-end significance. It is the responsibility Virtual Path Cross-Connects and the Virtual Channel Switches to alter the VPI and VCI fields in the cell header according to predefined routing table entries that were created at connection set-up time.

The network deals with uni-directional connections implemented by VCCs. Connections may be of one of a number of pre-defined types called *Connection Types* or

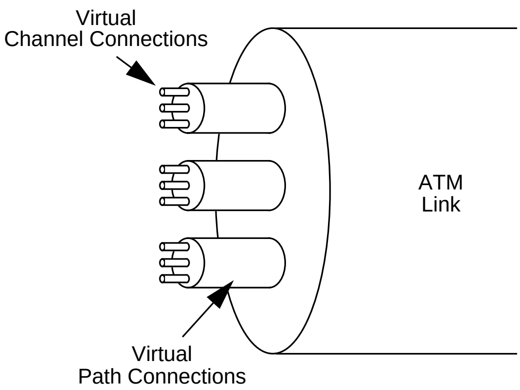

Figure 2.2 Links, VPCs and VCCs

Classes of Service (CoSs). Each Cos is defined in terms of its bandwidth and performance requirements. A CoS identifier is used to identify the type of connection during the connection set-up procedure.

A *call* is made up from one or more related VCCs usually sharing the same route through the ATM network. Thus a traditional telephone call can be considered as comprising two connections, one from the initiator of the call to the receiver, the other to accommodate replies made by the receiver to the initiator. A video telephone call comprises four connections, two for audio and two for video.

ATM makes use of both electronic and optical technology. ATM links can be copper or fibre-optic; typical link rates are 155.55 and 622.20 Mbit/s, corresponding to 365863 and 1467453 cells/s respectively. Because of technology limitations, ATM switches are predominately based on semiconductor technology; CMOS offers speeds up to about 250 Mbit/s with low power dissipation and high packing density; Emitter Coupled Logic achieves speeds approaching 10 Gbit/s but compromises on power dissipation and packing density.

The techniques underlying ATM communications were already being considered in 1983 in Europe and the United States, and in 1990, the CCITT produced thirteen recommendations defining ATM principles and parameters.

2.2 General characteristics of ATM

In comparison with many legacy systems, ATM has the following features:

- statistical multiplexing,
- connection-oriented operation,
- no error or flow control on a link,
- limited header functionality, and,
- a small information field.

Statistical multiplexing is a particularly important feature of ATM. It allows many connections to share ATM links and VPCs, even though the sum of their peak bandwidths exceeds the overall bandwidth offered by the link or VPC; this is done on the assumption that the sum of their instantaneous bandwidths exceeds the overall link capacity with a rarity sufficient to give an acceptable cell loss for each of the connections in that link. Statistical multiplexing therefore fosters efficiency in the use of a link.

The *connection-oriented* aspect ties in with that adopted by legacy networks catering for voice and video services. The connection-oriented approach has a considerable influence on how an ATM network is managed.

Simplicity of design, manufacture, operation and maintenance is fostered by the lack of link error and flow control mechanisms, as well as by the limited cell header functionality and small payload or information field.

2.3 Loss, delay and priority issues

Cell loss is concerned with the concept of "semantic transparency." Some services such the transfer of computer data and files, can be sensitive to bit errors and cell loss. Synchronisation information in video services should also not be lost, although otherwise video and audio services are usually tolerant of some information loss. The bit error rate in an ATM network is determined by transmission errors, buffer overflow and packet misrouting. *Transmission errors* are independent of the use of ATM and will give rise to the loss or incorrect arrival of bits in the cell payload (information field). *Buffer overflow* can occur in the switching systems when a cell queue expands, possibly as a result of traffic congestion, beyond the size of the buffer concerned; the buffer size should be dimensioned to accommodate the expected traffic conditions, but may also be influenced by technological constraints. *Packet misrouting* is caused by the misinterpretation of the cell header in a switching system and will give rise to the loss or incorrect arrival of cells.

Video and audio services can be particularly susceptible to variable *delay.* Interactive services can, in addition, be susceptible to absolute delay, if the two or more communicating parties have a separation such as to cause a delay approaching 0.5 seconds. The overall network delay is made up of packetisation, fixed switching, queueing, transmission and depacketisation delays. *Packetisation delay* occurs where a real time service such as video or voice is converted to an ATM cell stream. *Fixed switching delay* depends on the architecture of the given ATM switch and is normally kept low by suitable switch design. *Queueing delay* occurs in the switch buffers and will depend on the traffic load in the network. The *transmission delay* for fibre-optic links is between 4 and 5 μs/km. *Depacketisation delay* is introduced at the receiver to smooth out delay variation (delay jitter) experienced by real time services such as video or voice.

Priority was not originally seen as a feature of ATM, but more recently, space and time priorities have been studied. *Space priority* ensures that in the event of buffer overflow, "precious" cells are retained in preference to other cells; the transmission of critical data such as computer files or video synchronisation information, benefits from space priority schemes. *Time priority* ensures that in the event of buffer overflow, cells carrying real time information (e.g. video or audio) are favoured over other cells.

2.4 ATM network control

The number of connections using the ATM network must be restricted in order to provide the existing customers with an acceptable Quality of Service (QoS) and the network operator with an acceptable Network Performance (NP). Important parameters for QoS and NP include mean and variance of cell loss and delay for a connection.

Flow control is achieved by controlling the traffic entering the network. This is usually achieved by Connection Admission Control (CAC) and Usage Parameter Control (UPC), whose operations are coordinated by Call Control (CC).

Connection Admission Control, formerly known as Connection Acceptance Control, decides whether or not a new connection can be set up in view of the traffic that such a connection will impose upon the link or VPC. The CAC algorithm controls link or VPC loading to ensure it is within the admissible region i.e. in the region where link buffer overflow is within the bounds of a pre-defined probability (cell loss target of the CAC) - this is the definition used in this publication. CAC also has a second definition, namely: CAC decides whether or not a new connection can be set up in view of the traffic that such a connection will impose upon the *network as a whole*; in this case, CAC is also assumed to include the *Route Selection Algorithm* (RSA) and the associated distributed CC functionality for routing a connection through the required succession of nodes. This book treats the CAC and RSA functionality separately. The purpose of the RSA at a given node (access or transit) is, at connection set-up time, to select the "best" outgoing route for the requested connection according to the required destination address and CoS. The role of the RSA is illustrated later in this chapter, and a taxonomy of RSA types is discussed in some detail in Chapter 5.

Usage Parameter Control (UPC), formerly known as Source Policing, regulates the number of cells injected into the network by a source. UPC monitors the traffic rate of a VCC (for switched, on-demand services) or VPC (in the case of leased lines or Virtual Private Networks - see Chapter 6), and, discards cells which violate the agreed traffic rate parameters. *Network Parameter Control* (NPC), provides similar functionality to UPC, but within the ATM network itself, in order to protect the network against the malfunctioning of its own network elements.

Call Control is a distributed process involving all of the Virtual Channel Switches over which the call will be routed. The CC mechanisms in each node communicate amongst themselves and with the user terminals via signalling protocols. In particular, CC deals with:

- call set-up which in turn invokes:
    - a RSA at each node to select a suitable outgoing route for each connection,
    - a CAC algorithm at each node for each connection of a call leaving that node,
    - selection of UPC parameters.
- call-in-progress monitoring, and,
- call release.

2.5 Further ATM concepts

These will only be referred to briefly here for the sake of completeness. The ATM layered model, covered by ITU-T Recommendation I.320 [2.3], considers the following four layers:

- the Physical Layer,
- ATM Layer,
- ATM Adaptation Layer (AAL), and,
- Higher Layer Protocols.

Cells visible at one or both of the physical and ATM layers include: Cells with an information payload, Unassigned cells, Idle Cells, Meta-signalling cells, Broadcast signalling cells and Physical Layer Operations and Maintenance Cells. Security issues include Cell Header Error Correction and Cell delineation and payload scrambling.

Furthermore, this Recommendation, together with I.321 [2.4] distinguishes between three planes in telecommunications networks (Figure 2.3):

- the Control Plane,
- the User Plane, and,
- the Management Plane.

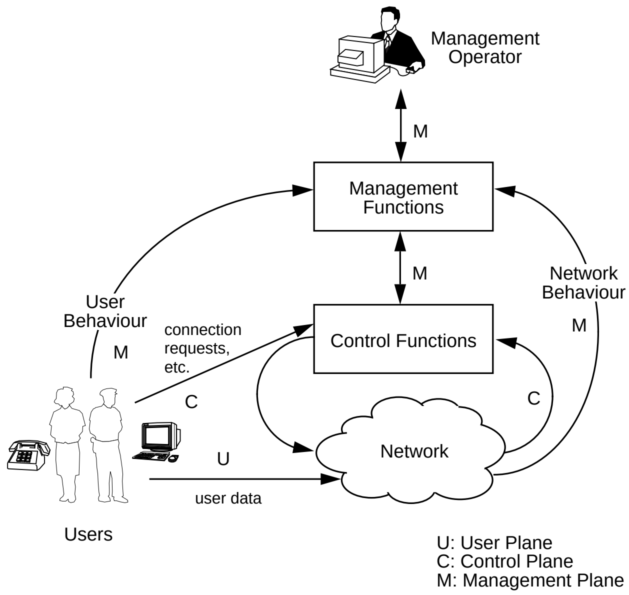

Figure 2.3 User, control and management planes

The relationship between these planes is an important issue in this book. The main focus of the ICM project was in the management plane [2.2], but the control and user planes had to be understood in order than the requirements and the environment of the management issues dealt with by the project could be put into perspective. The management plane should compliment and enhance the control plane functions by configuring operational parameters. Management should not replace the control functions and in general it has less stringent requirements on real-time response. The following section illustrates the operation of the control plane, the remainder of the book details the management issues relating to the control functions.

2.6 Illustration of the connection establishment procedure

A introduced previously, ATM networks are connection-oriented networks. Each node basically provides switching and CC functionality. Switching is done at two levels: VP Cross-Connects route cells within a VPC based on their VPI; VC Switches switch cells belonging to a particular VCC between VPCs based on their VCI. VPCs are created on a semi-permanent basis by management activity whereas VCCs are created dynamically by the control plane of the network via UNI (User-Network Interface) and NNI (Network-Network Interface) signalling. RSA and CAC mechanisms, are part of the control plane, however their behaviour is specified according to operational parameters which are set by the management plane.

Users, via UNI signalling, identify the number of connections they require for their calls; each connection is specified by the destination address and a CoS identifier. VCC establishment is achieved in a distributed way with each switch along the route making local decisions. In order to accomplish routing, all feasible routes towards a given destination for a particular CoS are stored locally in a route selection table - the entries in this table are initially set and subsequently modified by the management plane. For a number of reasons (increased availability, reduced vulnerability to failures, adaptivity) more than one route may exist to the destination for the CoS. At connection set-up time, the RSA in each switch is responsible for selecting the most appropriate VPC to be used for the next hop of the connection from VPCs in its route selection table associated with the selected CoS and the required destination node. The RSA is based on route selection parameters associated with the available routes. These parameters reflect the preference of selecting one route over another. Several types of RSA can be distinguished, according to the method they employ, the information they utilise and the degree of adaptivity they offer.

The selected route identifies an outgoing VPC and the request is passed to the CAC function responsible for that VPC. Based on the bandwidth parameters of the requested CoS and those of the existing connections, the CAC function employs its algorithm to determine whether there is capacity on the VPC for the new connection. If there is capacity the connection request is forwarded, via NNI signalling, to the switch at the other end of the VPC where the process is repeated until the destination is reached. If CAC determines that capacity is not available on a VPC the RSA tries the alternative routes until all routes are exhausted and the connection request is rejected at that node.

In order to take full advantage of all the alternative routes in the network the CC may provide the capability of route back-tracking whereby the previous switch is signalled to attempt its alternative routes. If all the routes at the access node have been exhausted the user is informed that the connection has been rejected.

When a connection is accepted the CAC function at the access switch invokes a UPC policing function to ensure that the user traffic does not exceed the bandwidth parameters specified for that CoS.

2.7 Summary

This chapter briefly introduced the important concepts associated with ATM networks. The knowledge of which will aid the understanding of the ICM work presented in the following chapters. In particular, the chapters on VPC and Routing management and Virtual Private Networks management require an appreciation of general ATM principles. Specifically, the following should be taken into account:

- that ATM networks are connection-oriented,
- the distinction between control and management plane activities,
- the fact that VCCs are multiplexed onto VPCs which in turn are multiplexed onto ATM transmission links,
- the differences between VC switching functions and VP cross-connect functions,
- the fact that VPCs are configured by management activity prior to the establishment of VCCs through signalling and control plane actions,
- the relationship between calls and connections,
- the concepts of CAC and RSA, and,
- the notion of Connection Types, or CoSs.

2.8 References

[2.1] Martin de Prycker, “Asynchronous Transfer Mode,” Ellis Horwood, Chichester, U.K., 1991.

[2.2] ITU-T Recommendation M.3010 (November 1991). “Principles for a Telecommunications Management Network,” Version R5.

[2.3] ITU-T Recommendation I.320, “ISDN Protocol Reference Model.”

[2.4] ITU-T Recommendation I.321, “B-ISDN Protocol Reference Model and its Application.”

[2.5] ITU-T Recommendation I.150, “B-ISDN Asynchronous Transfer Mode functional characteristics.”

[2.6] ITU-T Recommendation I.362, “B-ISDN ATM Adaptation Layer (AAL) functional description.”