Management and Control Functions in ATM Switching Systems

As research has progressed, it has become clear that the main difficulties in ATM pertain to its operational details rather than the concept. And it seems likely that these control issues will be much more complicated and costly for ATM switches when compared with current telephone circuit switches.

Thomas M. Chen and Steve S. Liu

Asynchronous transfer mode (ATM) is the target switching technique for the future public Broadband Integrated Services Digital Network (B-ISDN) [1, 2]. A great deal of early ATM research has focused on switch fabrics to demonstrate the feasibility of the ATM concept. Many switch fabric architectures have been proposed and prototyped [3], and a number of broadband switching systems with ATM fabrics are commercially available.

The switch fabric provides the essential routing and buffering functions, but ATM switching systems will also require management and control functions necessary for the efficient operation of the network. As research has progressed, it has become clear that the main difficulties in ATM pertain to its operational details (e.g., control of multiple types of traffic with different service requirements) rather than the concept. Relatively little attention has been devoted to the control functions required in ATM switching systems. Part of the reason is that standards on operations and maintenance (OAM) and traffic control principles have not yet been finalized. While the problems are widely understood, agreement within the industry on operations and traffic control schemes has been somewhat slow.

There are reasons to expect that the control aspects will be much more complicated and costly for ATM switches than current telephone circuit switches. First, processing and control must be exercised on the levels of individual ATM cells (53-byte packets consisting of 5-byte headers and 48-byte payloads) as well as virtual connections. Second, various types of traffic (e.g., voice, data, video) with different quality of service (QOS) requirements will be mixed within the network. Some of these services will be variable bit-rate which implies that traffic flows within the network will fluctuate randomly. Finally, at the high speeds of ATM, even brief deteriorations in service may result in serious losses of user information. Therefore it will be important to maintain the desired level of network performance. For these reasons, we believe that the consideration of control aspects is an important factor in the design of ATM switching systems.

The purpose of this article is to examine the management and control functions in ATM switching systems implied by current industry standards and agreements on OAM and traffic control. Until now, ATM research in the areas of switch design and traffic control have progressed essentially independently. First, we briefly review the B-ISDN Protocol Reference Model and its representation of the different information flows in ATM. Network management and traffic control principles in ATM, and in particular OAM, are overviewed. With this information as background, we attempt to infer their implications on the functional blocks of an ATM switching system. An example switch architecture model with distributed management and control functions is outlined, and some design issues are discussed.

Information Flows in ATM

ATM switching systems can be viewed as network elements that store, process, and relay the flows of information through the network. The information flows in ATM are defined by the B-ISDN Protocol Reference Model which consists of three separate planes: User Plane, Control Plane, and Management Plane [4]. These planes represent the protocol layers associated with the three types of traffic flows: user information; signaling information for call/connection control; and management information for the efficient operation of the ATM network.

The User Plane has a hierarchical structure consisting of the Physical Layer, ATM Layer, ATM Adaptation Layer (AAL), and Services Layer. The Physical Layer will be unframed (i.e., cell-based [5]) or framed using SONET [6] or SDH [7]. The AAL is subdivided into two sublayers: the Segmentation and Reassembly Sublayer and the Convergence Sublayer [2, 8]. The Segmentation and Reassembly Sublayer maps user information into the 48-byte payloads of ATM cells before entry into the ATM network and reassembles the cell payloads into user information after delivery. The Convergence Sublayer is responsible for provid-

THOMAS M. CHEN is involved in research in ATM switching and network control at GTE Laboratories.

STEPHEN S. LIU is a principal member of technical staff in the Network Transport Systems Department at GTE Laboratories.

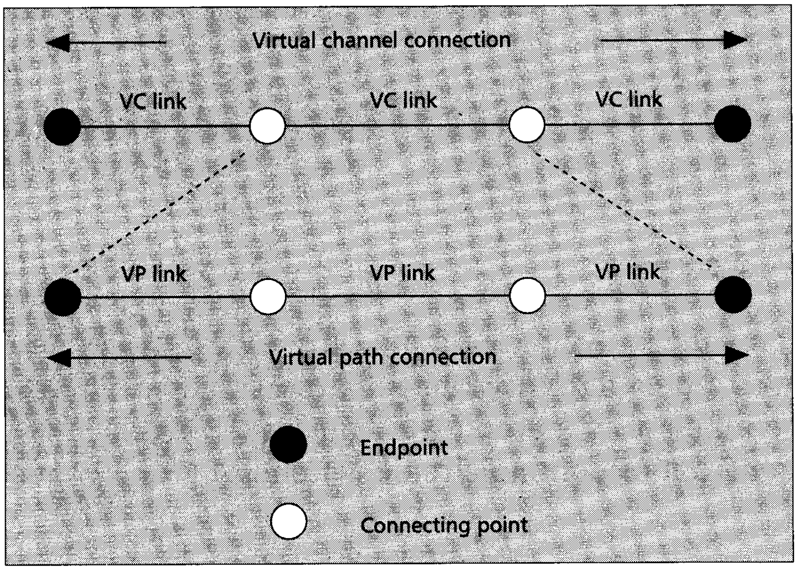

■ **Figure 1.** *Connections at the VP and VC levels.*

ing adaptation functions to the Services Layer that are service-specific.

Between the Physical Layer and AAL, the ATM Layer is responsible for the end-to-end sequence-preserving transfer of ATM cell streams [2, 9]. ATM is connection-oriented at the two levels of virtual channels (VCs) and virtual paths (VPs), which will be relevant to the later discussion of OAM and traffic control principles. At each level, a virtual connection consists of links, endpoints, and connecting points as shown in Fig. 1 [10]. A virtual channel connection (VCC) or virtual path connection (VPC) is characterized by its required quality of service (QOS) and throughput parameters. QOS is specified mainly in terms of cell delays and cell loss rate [11]. The ATM Layer is intended to support a number of different QOS classes. The QOS class of each VCC is specified during connection setup and can be changed only by renegotiation between the user and network. Specific QOS parameter values may be explicitly requested or implicitly associated with a service. A VPC will carry VCCs of possibly different QOS classes, and the QOS of the VPC must satisfy the most demanding QOS of the carried VCCs. Hence VCCs with similar QOS requirements will probably be grouped in the same VPC. More details of the ATM Layer can be found in the extensive literature on ATM (e.g. [12, 13]).

The Control Plane representing the protocol layers for transport of signaling information shares the same layered structure as the User Plane. Thus signaling information is exchanged among users and network nodes by means of ATM cells called signaling cells. Signaling procedures use a signaling virtual channel to establish, change, and release VPCs/VCCs for user information transfer. The default point-to-point signaling VC is identified by the header field codes VPI=0, VCI=5 [11]. To establish and control other signaling virtual channels, "meta-signaling" procedures (in the Management Plane) use a permanent VC called the meta-signaling virtual channel identified by the standardized field code VPI=0, VCI=1 [11]. Specific meta-signaling procedures are under study.

Finally, the Management Plane represents the transfer of control information used for maintaining the efficient operation of the ATM network. It is subdivided into Plane Management and Layer Management. The Plane Management is responsible for functions related to the network as a whole and coordination between all planes, and hence it is not layered. Specific functions of Plane Management are as follows.

- Fault management: to dynamically detect, isolate, and correct failures.
- Performance management: to continually monitor, report, and evaluate the behavior of network elements.
- Configuration management: to initialize installed equipment into service, and check or change their service status.
- Accounting management: to collect, process, and report information on resource usage for customer billing.
- Security management: to regulate the access to and control of network elements' databases.

They are consistent with the conventional domains of network management [14, 15]. ATM standards to date have addressed only fault management and performance management [16].

The Layer Management is responsible for layer-specific management functions. We will be concerned primarily with the ATM Layer Management which includes ATM Layer OAM, resource management, and meta-signaling.

While further details of the Management Plane continue to be studied, the ATM Forum has specified an Interim Local Management Interface (ILMI) [11], which allows bidirectional exchanges of messages between network management entities residing in ATM devices situated across the UNI. The ILMI is described further in the next section.

Network Management and Traffic Control in ATM

In ATM, the objective of network management is to enable the successful completion of as many B-ISDN service calls as possible [17]. Naturally, an important task is monitoring the performance of network facilities to detect failures and malfunctions, and responding with appropriate actions in order to minimize the effect on offered services. This is the responsibility of operations and maintenance (OAM) in ATM. Fault management is concerned with the detection, isolation, and correction of acute failures that interrupt the availability of network resources. Besides acute failures, some failures may be manifested intermittently, or malfunctions may subtly degrade network performance while network resources remain available (e.g., a corruption of VPI/VCI translation tables may cause misrouting for certain VCCs). It is the role of performance management to continually monitor the behavior of network facilities to detect degradations in performance caused by these conditions.

In the absence of faults and malfunctions, efficient operation of the network must be maintained by the functions of traffic control and resource management. Without traffic controls, the network would exhibit the well known phenomenon of congestion where the actual throughput will begin to decrease with the offered load. This deterioration in network throughput caused by congestion is characteristic of both circuit-switched and packet-switched networks without the protection of traffic controls. In circuit-switched voice networks, network efficiency begins to suffer from "trunk congestion" as the offered traffic increases and alternate routes become heavily loaded. Then "switch congestion" occurs when call processors become increasingly occupied with unsuccessful call attempts [18, 19]. In packet-switched data networks, excessive offered traffic results in overflows of network buffers and long packet transfer delays. Consequently, more packets are retransmitted, which increases the network load even further [20]. The role of traffic control is to regulate the flows of traffic and manage the efficient allocation of network resources in order to prevent and react to congestion.

Because ATM is connection-oriented and supports real-time services, traffic control at the level of call admissions will be important. As in the circuit-switched telephone network, call admission control consists of selectively accepting or rejecting requests for new connections. When indications of congestion are detected, the network response may be a combination of expansive or restrictive controls depending on the nature of the congestion [18, 19]. For peak-day overloads, expansive controls can make more alternate routes available and change the normal traffic patterns to alleviate congestion in any particular area. For focused overloads (e.g., during an emergency in an area), restrictive controls such as selectively blocking calls to that area can prevent excessive call attempts from overwhelming that local switch.

However, the situation in ATM is much more complicated because the range of B-ISDN services will involve calls with widely different characteristics and service requirements. It is not clear at this time how to determine acceptance or rejection of B-ISDN calls on an equitable basis. For example, acceptance of a high-rate connection might result in the blocking of several low-rate connections affecting many users. Fairness might require some form of bandwidth partitioning or priorities between different services. Another issue arises with variable bit-rate (VBR) connections. A conservative admission policy for VBR connections based on allocation of network resources according to their peak rates may result in low network utilization. If a high network utilization is desired by taking advantage of statistical multiplexing, connection admission control in ATM will be based on resource allocation at less than the peak rates for VBR traffic. In this case, there will be a probability that bursts at peak rate could coincide and cause buffer overflows or excessive cell delays. Traffic control at the cell level will be important to protect the quality of services from random peak fluctuations of traffic.

Because there is similarity between the cell-based statistical nature of ATM and conventional packet-switched data networks, the traffic control mechanisms used in data networks [20] might be expected to be applicable in ATM. However, the principles of feedback and delays that are characteristic of traffic controls in data networks (e.g., choke packets, credits, sliding windows) are not appropriate for real-time ATM traffic. Also, the effectiveness of schemes using feedback is fundamentally limited by the propagation delay; i.e., a high-speed source is capable of sending too many cells before feedback can propagate backward to control it. Hence ATM will likely require a combination of various traffic control mechanisms operating on different levels and timescales with sensitivity to different classes of traffic [21]. These are discussed later in this section.

5 bytes	48 bytes			
Cell header	OAM cell type	OAM function type	Function specific field	EDC
bits: 40	4	4	366	10

■ **Figure 2.** *OAM cell format.*

OAM Cell Type	Value	OAM Function Type	Value
Fault management	0001	Alarm indication signal (AIS)	0000
		Remote defect indicator (RDI)	0001
		Cell loopback	1000
		Continuity check	0100
Performance management	0010	Forward monitoring	0000
		Backward reporting	0001
		Monitoring and reporting	0010
Activation/deactivation	1000	Performance monitoring	0000
		Continuity check	0001

■ **Table 1.** *Field codes of OAM cells.*

OAM in the ATM Layer

The flows of OAM information in the Physical Layer are designated as F1-F3 OAM flows (corresponding to the levels of regenerator section, digital section, and transmission path). In SONET, OAM information is carried in the SONET section, line, and path overhead. The overhead allows section and line error monitoring, payload error monitoring, path tracing, and includes embedded data channels for alarm gathering, remote provisioning, and other communication needs.

Standards for OAM in the ATM Layer have mainly addressed performance management and fault management [16, 22]. ATM Layer OAM is based on the bidirectional exchanges of cells carrying OAM information related to specific VPCs and VCCs. At the VP and VC levels, they are referred to as F4 and F5 OAM flows, respectively. F4 OAM cells have the same VPI field value as user cells of the same VPC but are distinguished by pre-assigned VCI values: VCI=3 for "segment" OAM cells (i.e., cells communicated within only a portion of a VPC) and VCI=4 for "connection" OAM cells (communicated end-to-end in a VPC). F5 OAM cells have the same VPI/VCI values as the user cells of the same VCC but are identified by pre-assigned PT values: PT=4 for segment OAM cells and PT=5 for connection OAM cells.

All OAM cells will have the format shown in Fig. 2. The OAM Cell Type field indicates the type of management function as listed in Table 1. Only fault management, performance management, and activation/deactivation cells have been

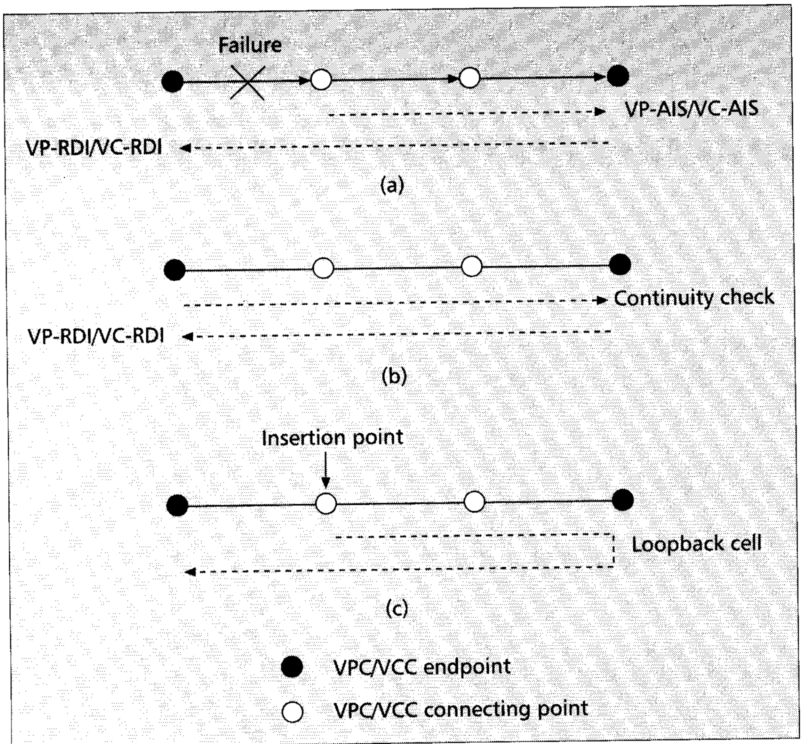

■ **Figure 3.** *Fault management procedures: a) alarm surveillance; b) continuity checking; c) OAM cell loopback.*

defined so far [16, 22]. The OAM Function Type field indicates the specific function performed by that cell. The EDC field is a CRC-10 error detection code computed over the information field of the OAM cell using the generator polynomial $x^{10} + x^9 + x^5 + x^4 + x + 1$.

F4 and F5 OAM flows are initiated during or after connection set-up. OAM cells for both directions of the F4 and F5 flows must follow the same physical route so that any connecting points can monitor the OAM information from both directions. Both endpoints and connecting points can generate, insert, monitor (non-intrusively), and process OAM cells for that virtual connection; only endpoints of a flow can extract OAM cells.

Fault Management — The objective of fault management in the ATM Layer is to monitor and test the availability of VPCs/VCCs. This is performed by three methods: surveillance of alarms from the Physical Layer; continuity checking; and on-demand testing of connectivity by cell loopback. These are illustrated in Fig. 3.

If the Physical Layer indicates a failure (e.g., loss of signal, loss of cell synchronization), VPC/VCC failure will be reported in the ATM Layer with two types of cells: VP/VC-AIS (Alarm Indication Signal) and VP/VC-RDI (Remote Defect Indicator). Both AIS and RDI cells include fields for failure type (1 byte) and failure location (15 bytes). Upon receiving a failure indication from the Physical Layer, the detecting node will wait a short time for the Physical Layer automatic protection switching. If the failure indication persists, the node will issue a VP/VC-AIS cell to notify downstream nodes of connection unavailability. VP/VC-AIS cells will continue to be generated periodically until the fault is corrected. After receiving a certain number of VP/VC-AIS cells, the VPC/VCC endpoint will begin to send VP/VC-RDI cells upstream to notify the source VPC/VCC endpoint of the downstream failure. Fault localization and recovery actions are then initiated.

Continuity checking is not supported in [11] but it is described as a future possible option in [22]. If a VPC/VCC failure has not been indicated from the Physical Layer but no user cells have been sent downstream by a source VPC/VCC endpoint for some time T_s, it may send a Continuity Check cell downstream. Its purpose is to confirm that an inactive connection is still alive. If the destination VPC/VCC endpoint does not receive any cell within T_r ($T_r > 2T_s$) time, it implies that connectivity has been lost and the VPC/VCC endpoint will send a VP/VC-RDI cell to the source VPC/VCC endpoint. The same procedure is carried out in the reverse direction as well.

Although alarms and continuity checks are useful for fault detection, a means to test VPC/VCC connectivity on-demand will be useful to locate faults. The OAM cell loopback capability will allow an OAM cell to be inserted into a VPC/VCC and looped back (i.e., returned in the reverse OAM flow). Fields in the OAM loopback cell include: Loopback Indication (1 byte) to signify whether loopback has occurred; Correlation Tag (4 bytes) to uniquely identify the cell; Loopback Location (15 bytes) to indicate the point for loopback; and Source ID (15 bytes) to allow the originating node to recognize its own cells upon their return.

Performance Management — While acute failure conditions will be detected by fault monitoring, intermittent error conditions may cause a gradual deterioration in QOS. The continuous collection of performance measurements is needed to detect such deterioration. OAM performance management cells provide a mechanism to measure the performance of VPCs/VCCs and report the collected performance data. The procedure can be activated at a VPC/VCC endpoint by a user or operations system (OS) or at a connecting point by an OS.

The procedure consists of inserting OAM performance management cells between blocks of user cells. The OAM cells carry information about the preceding block to the downstream VPC/VCC endpoint. The forward monitoring and backward reporting procedures are shown in Fig. 4. The source VPC/VCC node performs an error check calculation over a block of N user cells. The block size N is nominally 128, 256, 512, or 1,024, but the actual block size may differ somewhat from these values because the OAM cell must wait for the next unassigned cell slot (the OAM cell may eventually be forced in). The error check, with other information, is carried in the OAM performance monitoring cell. At the destination node, the same error check is calculated over the received block of user cells and compared with the contents of the following OAM cell. The results of the comparison are reported back to the source node in an OAM cell using the reverse OAM flow. Intermediate connecting points have the option of monitoring the procedure and results.

The OAM performance management cell includes fields for: Monitoring Cell Sequence Number (1 byte) to indicate the sequential identity of the cell; Total User Cell Number (2 bytes) to indicate the user block size N; even parity BIP-16 error detection code (2 bytes) computed over the information field of the block; Timestamp (4 bytes) to indicate the time of origination; Block Error Results (1 byte) for backward reporting of the number of errored BIP-16 parity bits observed at the destination; and Lost/Misinserted Cell Count (2 bytes) for backward reporting of the number of lost/misinserted cells observed at the destination.

Interim Local Management

While further details of the Management Plane continue to be studied, the ATM Forum has specified an ILMI [11]. Based on the Simple Network Management Protocol (SNMP) and a standard ATM UNI Management Information Base (MIB), it allows the ATM user to obtain status and control information about VPCs/VCCs at its UNI.

Across the UNI, each ATM device has an UNI Management Entity (UME) that supports ILMI functions. ILMI communications take place between adjacent ATM UMEs through the exchange of SNMP messages which are encapsulated into ATM (ILMI) cells using AAL5. ILMI cells are identified by the default code values VPI=0, VCI=16. By means of SNMP messages, an UME can access the UNI MIB information associated with an adjacent UME. The ATM UNI MIB contains information about the Physical Layer, ATM Layer, ATM Layer statistics, and VPCs/VCCs (for details, refer to [11]).

Traffic Control and Resource Management

The objective of traffic control and resource management is the protection of the network in order to sustain a desired level of network performance; a secondary goal is the efficient utilization of network resources [21]. Traffic control consists of a set of mechanisms that regulate the flows of traffic in order to prevent random fluctuations from adversely effecting the quality of offered services. Traffic control will be a combination of preventive and reactive mechanisms operating on different levels and timescales. Some traffic control functions have been identified as: resource management; connection admission control (CAC); usage/network parameter control (UPC/NPC); and congestion control.

This section reviews the ideas offered in [11, 21]. It should be noted that the standards in this area have not been finalized. Functions such as feedback control, fast resource management [23], cell scheduling [24-27], and traffic shaping [28, 29] are still under study; there is a great deal of ongoing research that can be readily found in the extensive literature, e.g., refer to [28, 30].

Traffic Parameters — A traffic descriptor is a set of traffic parameters that can be used to characterize an ATM connection, e.g., peak rate and average rate. A source traffic descriptor is a subset of the ATM traffic descriptor that is used during connection set-up. Based on the source traffic descriptor, the network will make CAC decisions, allocate network resources, and derive the appropriate parameters for UPC/NPC.

Traffic descriptors, at the minimum, must include the peak rate [11]. Other parameters (e.g., sustainable cell rate or average rate, if known) may be optionally provided by the user in the interests of improving network efficiency. All parameters must be simple enough to be understood and calculable by the user, useful to CAC, and enforcable by UPC/NPC.

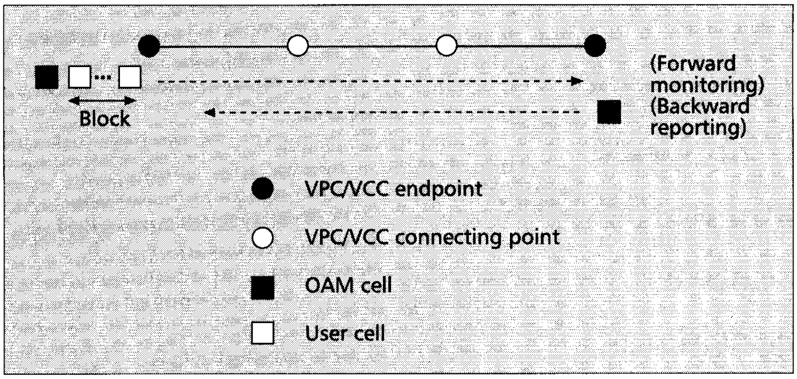

■ **Figure 4.** *Performance management procedure.*

ATM services are provided on the basis of a traffic contract negotiated between the user and network during connection set-up. It consists of the source traffic descriptor, requested QOS, cell delay variation tolerance, definition of conforming cells (for the UPC), and definition of a compliant connection (e.g., how many non-conforming cells may be allowed) [11].

Constant bit-rate (CBR) connections may be described simply by their rate. It is expected that cells of the same CBR connection will have the same cell loss priority under normal circumstances. However, the situation with variable bit-rate connections is much more complicated. VBR connections will require more than the peak rate to indicate their degree of burstiness; and two source traffic descriptors may be used to separately specify the high loss priority CLP=0 and low loss priority CLP=1 flows.

Resource Management — The initial use of virtual paths is expected to be an important part of resource management in ATM networks for a number of reasons [11, 31]. First, by reserving capacity on VPCs in anticipation of future VCCs, the processing required to establish individual VCCs can be reduced. VCCs can be established by making simple connection control decisions at VPC endpoints, and call processing is not required at connecting points. Second, VPCs allow a way to logically segregate traffic types requiring different QOS, while allowing VCCs to be statistically multiplexed. VCCs with similar QOS requirements can be grouped in the same VPCs; the QOS for a VPC must satisfy the most demanding QOS requirements for VCCs in that VPC. Third, VPCs allow a group of VCCs to be managed and policed more simply. Finally, dynamic routing control at the level of virtual paths allows a simple method for adaptive network reconfiguration. Path routing can be changed simply by modifying routing information at VPC connecting points.

Connection Admission Control — A B-ISDN call may involve one or more connections (e.g., for multimedia or multiparty services). If multiple connections are involved, connection admission is determined for each VCC/VPC. A new connection request is accepted only if it is determined that sufficient network resources are available to establish the connection with the required QOS and

> It is evident from OAM and traffic control considerations that ATM switching systems are much more than a switch fabric that simply routes and buffers cells.

maintain the agreed QOS of existing connections. The parameters of an established connection may be changed only through renegotiation between the user and network.

Signaling messages sent by the user will contain source traffic descriptors and the required QOS. The signaling protocol is defined by the Q.93B standard [32], and Q.93B signaling messages are encapsulated in ATM cells according to the signaling AAL (SAAL) [33, 34]. From the signaling information, the network determines: acceptance or rejection of new connections; traffic parameters for UPC/NPC; and allocation of network resources.

A specific connection admission policy will not be standardized and will be determined by the network provider. Any number of different approaches are possible; some proposals are summarized in [35, 36]. Many follow the general approach of using queueing analysis with a stochastic traffic model. New connections are admitted if the resulting QOS, as determined from the queueing model, is maintained above the desired level. Another general approach estimates an "equivalent capacity" for each connection and its required QOS, then the admission policy reduces to simply checking that the equivalent capacity is available [37].

Usage/Network Parameter Control — Since network resources are allocated on the basis of traffic contracts, UPC and NPC are necessary to monitor and regulate incoming traffic flows across the user-network interface (UNI) and network node interface (NNI), respectively. The objective is protection of the network from (intentional or unintentional) deviations from the given source parameters that can adversely affect the QOS of other connections. Specific tasks are verification of the validity of VPI/VCI values, and checking that the rates of incoming traffic from active VPCs/VCCs are conforming to agreed parameters.

For cells conforming to the negotiated traffic contract, the UPC/NPC may pass the cells or reschedule them if traffic shaping is implemented. The purpose of traffic shaping (at the option of the network provider) is to reduce the peak cell rate, burst length, or cell delay variation by spacing cells in time. For non-conforming cells, the UPC/NPC might discard them, or tag high loss priority cells by changing the header bit CLP=0 to CLP=1. A final possible option might be the release of the connection (presumably as a last resort).

A specific UPC algorithm called generic cell rate algorithm (GCRA), essentially a continuous-state leaky bucket, is supported by the ATM Forum for the UNI [11]. The leaky bucket algorithm is a well known policing method that allows some tolerance for burstiness, depending on its parameters [38-40]. It is also amenable to simple implementation (not prescribed by [11]). Multiple GCRAs can be combined together to monitor various traffic parameters. UPC/NPC is revisited in the section on input and output modules.

Congestion Control — The purpose of congestion control is to detect the onset of congestion and react to minimize the speed, effects, and duration of congestion. Specific control actions will include: selectively discarding cells according to the CLP bit; and Explicit Forward Congestion Indication (EFCI). Other actions are under study.

In EFCI, congested nodes along a VPC/VCC can inform downstream nodes and the destination endpoint by changing the PT field in the cells that pass through. Congestion indications are collected at VPC/VCC endpoints that inform the appropriate sources to adjust their rates (for traffic that can be flow controlled or adjusted). Due to the propagation delays involved, this mechanism may not be expected to prevent congestion but may help to mitigate the amount of cell loss during periods of persistent congestion. The procedure by which a network node monitors and classifies its internal congestion state is considered to be dependent on implementation.

ATM Switching Systems

It is evident from OAM and traffic control considerations that ATM switching systems are much more than a switch fabric that simply routes and buffers cells. In addition to relaying cells, ATM switching systems must contain the functions related to the Control Plane and Management Plane discussed earlier, and support a set of traffic control functions. We refer to the distribution of all these functions within the system as the switch architecture.

In this section, we apply the preceding background information towards the development of a general functional architecture model of an ATM switching system. The architecture model specifies only the *functional* requirements of a switching system and their distribution, and is not intended to prescribe any particular implementation. We omit interworking functions that would be necessary in practice to support various service-specific interfaces for narrowband circuits, frame relay, Switched Multimegabit Data Service (SMDS), and other services. It is assumed here that the switch interfaces are the standardized SONET-framed ATM UNI or NNI.

Functional Architecture Model

As represented by the ATM Layer in the User Plane, the main function of an ATM switch is to relay user cells from its input ports to the appropriate output ports. The flow of user traffic through three functional blocks is shown in Fig. 5a. First, incoming cells are extracted from the Physical Layer signals by input modules (IMs) and prepared for routing by translating the VPI/VCI fields. During the translation, a tag with internal routing information may be attached. The routing and buffering is performed by a cell switch fabric (CSF). Finally, the cells are prepared for physical transmission by output modules (OMs) at the switch output ports. The input modules, cell switch fabric, and output modules together perform the basic cell routing and buffering functions required in ATM switching.

As described earlier, the Control Plane depicts signaling information carried in ATM cells. Unlike user cells, the information within signaling cells is not transparent to the network; the signaling information must be processed and interpreted by the switch. Therefore the switch must identify incoming signaling cells, separate them from user cells, and process the signaling information.

If the switch generates control information, it must encapsulate the information into signaling cells that are merged with the outgoing user cell traffic.

The flow of signaling information is shown in Fig. 5b. Since signaling cells use the same ATM Layer transport as user cells, they flow through the input modules, cell switch fabric, and output modules. However, in addition, the signaling information is processed by a functional block shown as connection admission control (CAC). Here it is assumed that the cell switch fabric routes the signaling information to and from the CAC, but it is not necessary to go through the CSF. Alternatively, the signaling information might be passed directly from the input modules to the CAC, and from the CAC to the output modules.

Although the Management Plane has not been fully defined yet, the management functions required in ATM switching systems can be broadly identified as: fault management; performance management; configuration management; resource management; security management; accounting management; and meta-signaling. These functions are the responsibility of a functional block shown as system management (SM) in Fig. 5c.

A major responsibility of the SM is support of the ATM Layer OAM procedures defined for fault management and performance management. Like signaling cells, the information contained in OAM cells must be recognized and processed by the ATM switch. The ATM switch must identify OAM cells, separate them from other cells, and perform the necessary processing functions to support the ATM Layer OAM procedures. The flow of OAM information is shown in Fig. 5c. Along with the other cells, OAM cells flow through the input modules, cell switch fabric, and output modules. It is assumed here that the cell switch fabric routes the OAM cells to and from the SM, but it is not necessary to use the CSF. Alternatively, the OAM cells might be passed directly from the input modules to the SM, and from the SM to the output modules.

Another responsibility of the SM is support of the ILMI for each UNI. The SM contains a UME for each UNI that monitors the objects defined in a standardized ATM UNI MIB. The UME at the user device can communicate with their adjacent UMEs across the UNI by means of SNMP mesages carried in ILMI cells. Again, the ILMI cells follow the flow through input modules, CSF, SM, and output modules in Fig. 5c.

Finally, the SM should support network-wide operations, administration, maintenance and provisioning (OAM&P) functions which will probably be proprietary to the network provider. The SM may exchange network management information with an OS through a direct interface or via a separate telecommunications management network (TMN) [14].

In summary, the general switch architecture model in Fig. 5c consists of these basic functional blocks:

- input modules: to receive incoming cells and prepare them for routing through the CSF;
- output modules: to prepare outgoing cells for transmission;
- cell switch fabric: to route user-data cells from input to output ports, and possibly route signaling and management cells between the other functional blocks;

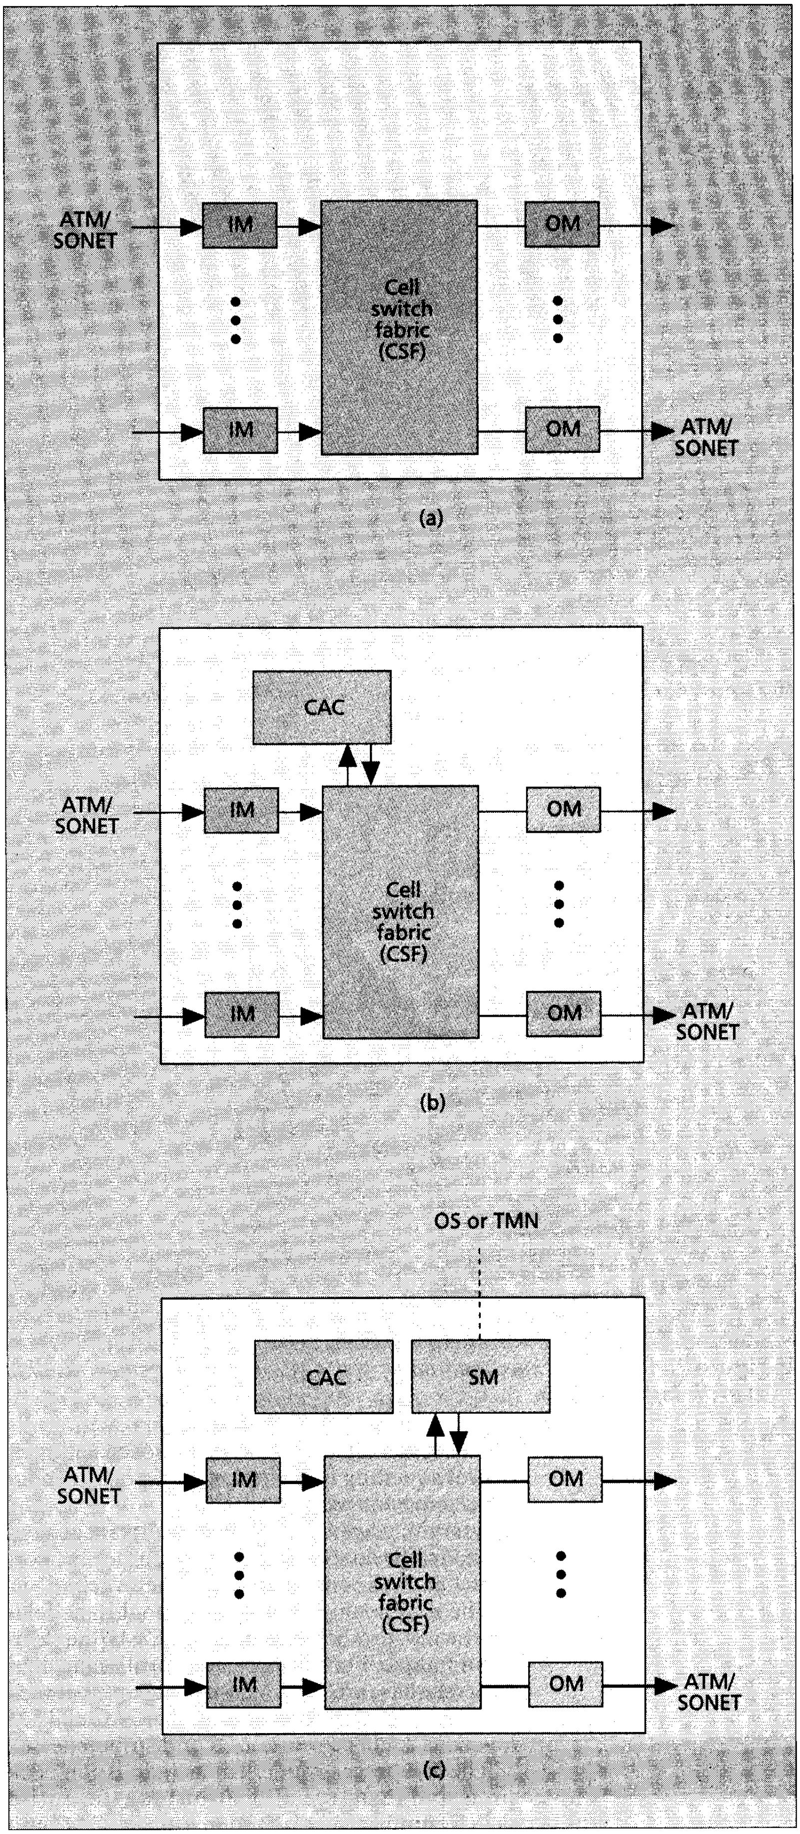

■ **Figure 5.** *The flows of a) user-data cells; b) signaling information and c) management information.*

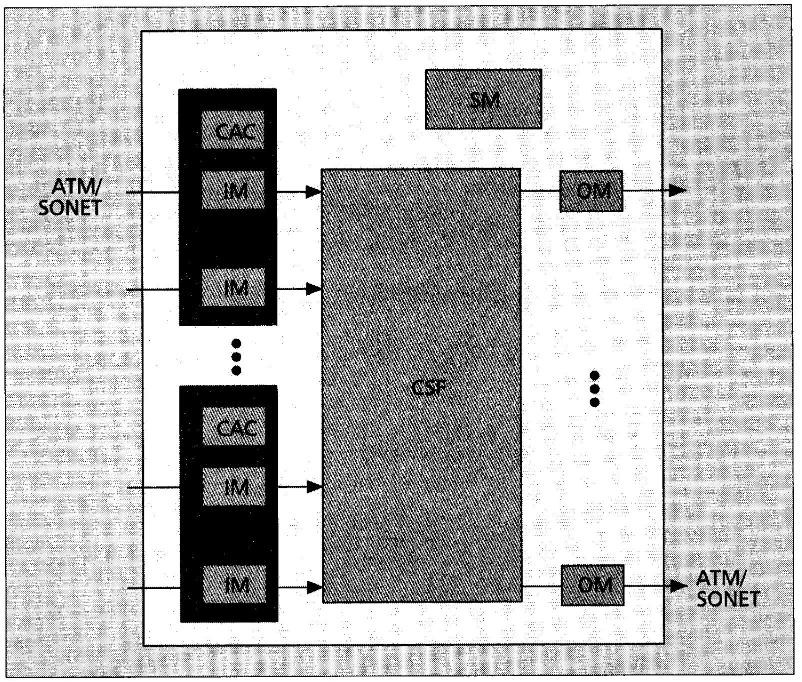

■ **Figure 6.** *Call processing functions distributed to blocks of IMs.*

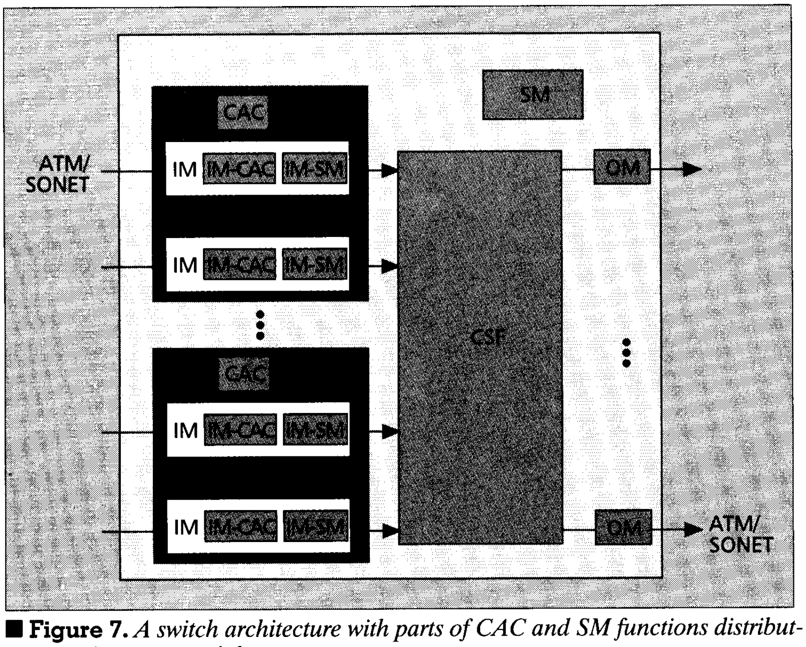

■ **Figure 7.** *A switch architecture with parts of CAC and SM functions distributed to the input modules.*

- connection admission control: to process and interpret signaling information, and perform connection admission or rejection;
- system management: to perform all management and traffic control functions to ensure the correct and efficient operation of the switch.

The partitioning of functions into this model is consistent with other related studies [22, 41] but clearly not unique. Furthermore, the partitioning does not always have precisely defined boundaries between functional blocks. For instance, the CAC and SM functions could be centralized or distributed.

The call processing functions (admission decisions and resource allocation) for the entire switch could be centralized in a single CAC unit as shown in Fig. 5b, or alternatively, call processing could be distributed to blocks of input ports as in Fig. 6 (cf. [42, 43]). The centralized approach is simpler to implement but the single call processor could become a bottleneck for large switch sizes as it handles signaling information from all switch inputs. In the distributed approach in Fig. 6, each CAC block processes the signaling information received by a small number of input ports. Thus large switches could be constructed without the bottleneck caused by centralized call processing, but this approach is more complex. Each distributed CAC block uses global state information which must be updated, distributed, and coordinated.

Figure 7 shows that part of the CAC functions may be distributed further among the input modules. The distributed portion of the CAC functions within the IMs will be designated as IM-CACs. The IM-CACs could handle extraction, interpretation, and insertion of signaling cells. The IM-CACs within the input modules may operate in parallel and assume some of the processing burden from the CAC. The IM-CACs are discussed further in the section on connection admission control.

In Fig. 5c, the SM functions are centralized and OAM cells must be diverted from the incoming cell streams to the SM for processing. As a consequence, OAM cells will not maintain their relative positions within the user cell streams. This is important to note because OAM performance management cells (in forward monitoring) should maintain their positions between blocks of user cells in each VPC/VCC. It implies the necessity to distribute a portion of the SM functions to the input modules such that OAM cells may be processed within the input modules without extracting them from the user cell streams, as shown in Fig. 7. Thus OAM cells may maintain their relative positions among user cells within a VPC/VCC. Furthermore, this distributed approach could alleviate the processing bottleneck that could occur with a centralized SM. The distributed parts of the SM functions will be referred to as IM-SMs and described further in the section on system management.

Input and Output Modules

A functional diagram of an input module is shown in Fig. 8a. The IM-CAC and IM-SM blocks show that parts of the CAC and SM functions can reside in each input module. At the minimum, the IMs must include the SONET functions for the termination of the Physical Layer. More details are shown in Fig. 8b. The optical signal is converted to electrical, and the digital bitstream is recovered. The payload and overhead are separated. The overhead is processed and Physical Layer maintenance is performed. Cells are delineated from the payload and unassigned (empty) cells are dropped.

After this, error control is performed for each cell header using the polynomial $x^8 + x^2 + x + 1$. In the default mode, single bit errors are detected and corrected. If an error is detected and corrected, the mode is switched to multiple bit error detection (but no correction) for the next cell header. It remains in this error detection mode until a header without detected errors is encountered; then it switches back to the default mode. In error detection mode, cells with detected errors are discarded. Discarded cells should be recorded, and counts of errored/corrected cells should be maintained.

Next the IMs translate the VPI/VCI fields and

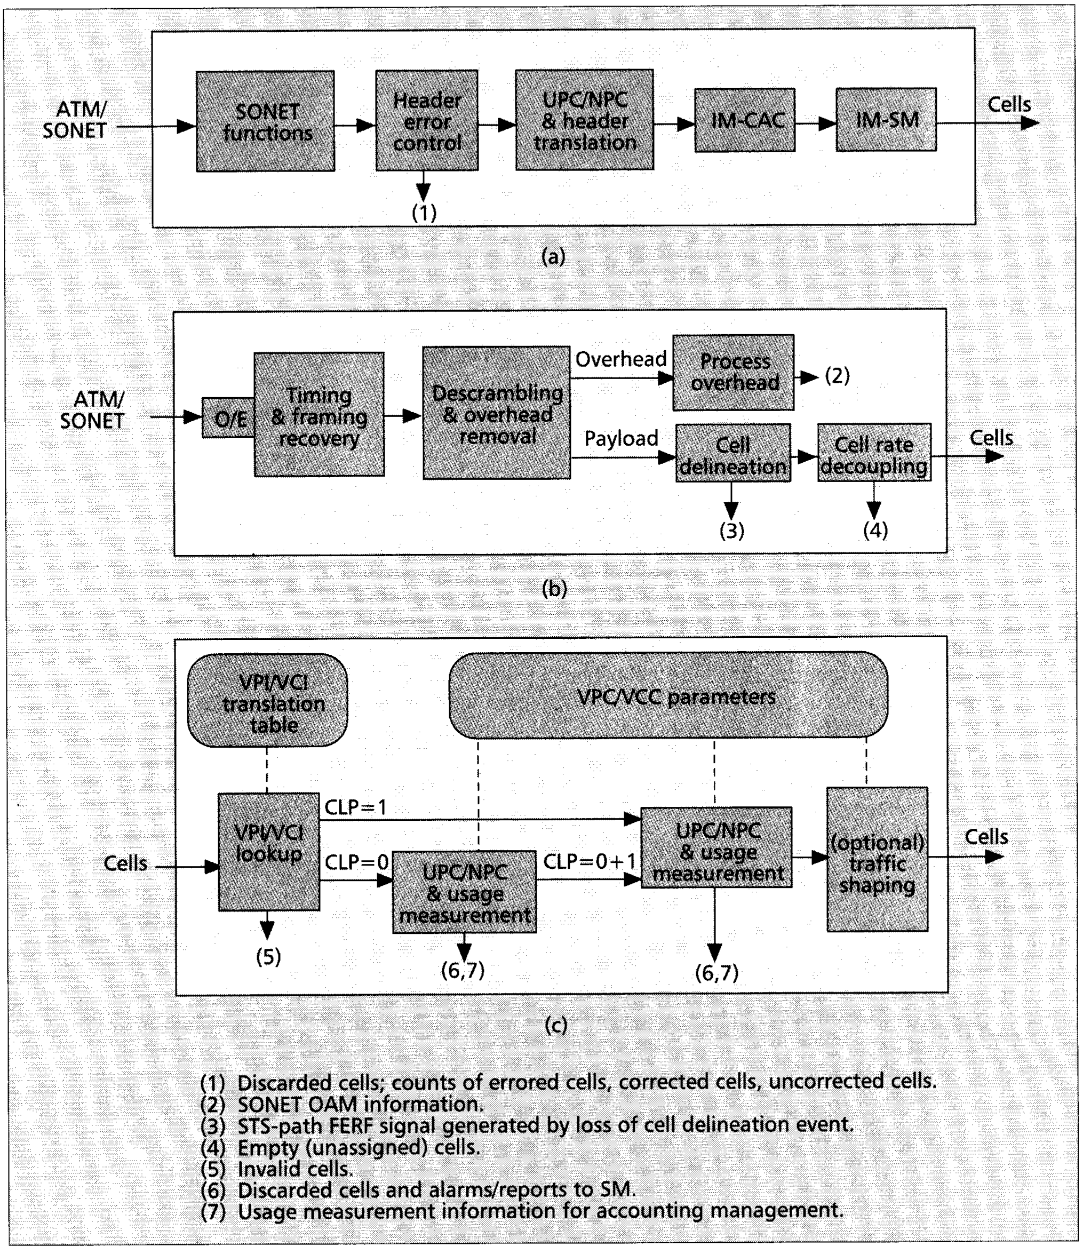

■ **Figure 8.** *Functional diagram of a) an IM; b) SONET functions; and c) UPC/NPC and header translation.*

The basic function of the CSF is to route cells from its inputs to the appropriate outputs according to the attached internal routing tag.

discard cells with invalid fields. In preparation for routing through the CSF, an internal routing tag may be attached to each cell (which will be removed later by the OMs). The internal routing tag might include destination output port; source input port; loss tolerance; delay priority; timestamp; maximum delay bound; delay variation tolerance; or other housekeeping information for internal monitoring and control purposes. The tag might also indicate whether the cell contains user, signaling, or management information. Since the tag exists only within the switch, its contents are determined by the switch designer.

UPC/NPC is performed for each VPC/VCC according to the set of traffic parameters specified in the traffic contract. The parameters would include at least the peak rate and cell delay variation tolerance, and might include the sustainable cell rate and burst tolerance as defined in [11]. The GCRA, essentially a continuous-state leaky bucket algorithm, is specified for UPC [11]. One or more UPC/NPC mechanisms could be configured, depending on the given parameters and whether the parameters refer to the CLP=0 flow or total CLP=0+1 flow. An example showing dual UPC/NPC mechanisms to enforce the peak rates for CLP=0 and CLP=0+1 flows is shown in Fig. 8c. First, UPC/NPC is applied to the peak rate of the high priority CLP=0 cells. Non-conforming cells may be dropped, or if tagging is requested and implemented, they may be tagged by changing them to low priority CLP=1. Tagging may be advantageous if tagged cells are likely to be delivered successfully without detriment to the QOS to other users, e.g., under light network loads. The second UPC/NPC enforces the peak rate for the total CLP=0+1 flow. A record of discarded cells and loss statistics should be maintained.

It would be natural to measure and record the incoming traffic flow at the point where it is controlled. Hence there should be a mechanism com-

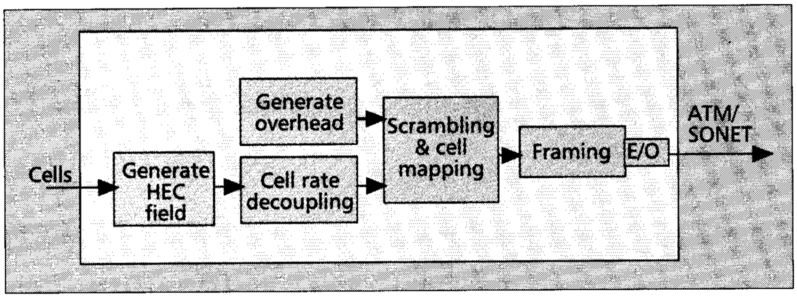

■ **Figure 9.** *Functional diagram of OM.*

bined with the UPC/NPC to collect measurements about usage per VCC/VPC for billing purposes. This usage information is provided to the SM which may receive additional usage measurements from the CSF (where cells may be discarded). If customer billing is handled by centralized network management, the SM processes the collected information and prepares reports for network management.

The output modules are considerably simpler than the input modules, particularly if the IM-CAC and IM-SM functions are present in the IMs. The output modules do not involve UPC/NPC and VPI/VCI lookup tables, nor do they contain buffers which we have included in the CSF. A functional diagram of an output module is shown in Fig. 9. An OM accepts a stream of cells from the cell switch fabric, removes the internal routing tag, generates the HEC field for each header, maps the cells into SONET payloads, generates the SONET overhead, and converts the digital bitstream into optical form. It is assumed that user cells, signaling cells, and management cells all go through the CSF. Alternatively, it might be possible that the OMs receive signaling cells from the CAC and management cells from the SM directly; these cells are mixed with user cells in the OMs before transmission.

Cell Switch Fabric

A diagram of the CSF is shown in Fig. 10. The basic function of the CSF is to route cells from its inputs to the appropriate outputs according to the attached internal routing tag. First, cells may be aligned in time by means of single-cell buffers. Some buffering is necessary within the routing fabric due to inevitable contention for output ports (typically output buffers are on the order of hundreds of cells or less per port). The favored approach is buffering at the outputs, although the buffering may be situated at the inputs, internally, or externally (as a recirculating buffer) [3, 44].

The routing operation should recognize delay priorities (which could be explicit in the internal routing tag), and the buffering operation must be capable of selective cell discarding according to the CLP bit. Both operations should preserve the sequential order of cells belonging to the same VCC, which will all have the same delay priority but not necessarily the same loss priorities. Implementation of the specific cell scheduling policy is determined by the network provider. Some proposed scheduling policies for ATM with delay/loss priorities are described in [24, 25, 45, 46].

The routing fabric may be implemented following any number of approaches such as shared memory buffer, shared medium, or space-division network [3]. It appears that almost any approach is capable of realizing small switch modules, e.g., 64 x 64 size with 155 Mb/s port speeds. Larger fabrics can be constructed by interconnection of small modules in multiple stages. In addition, the CSF

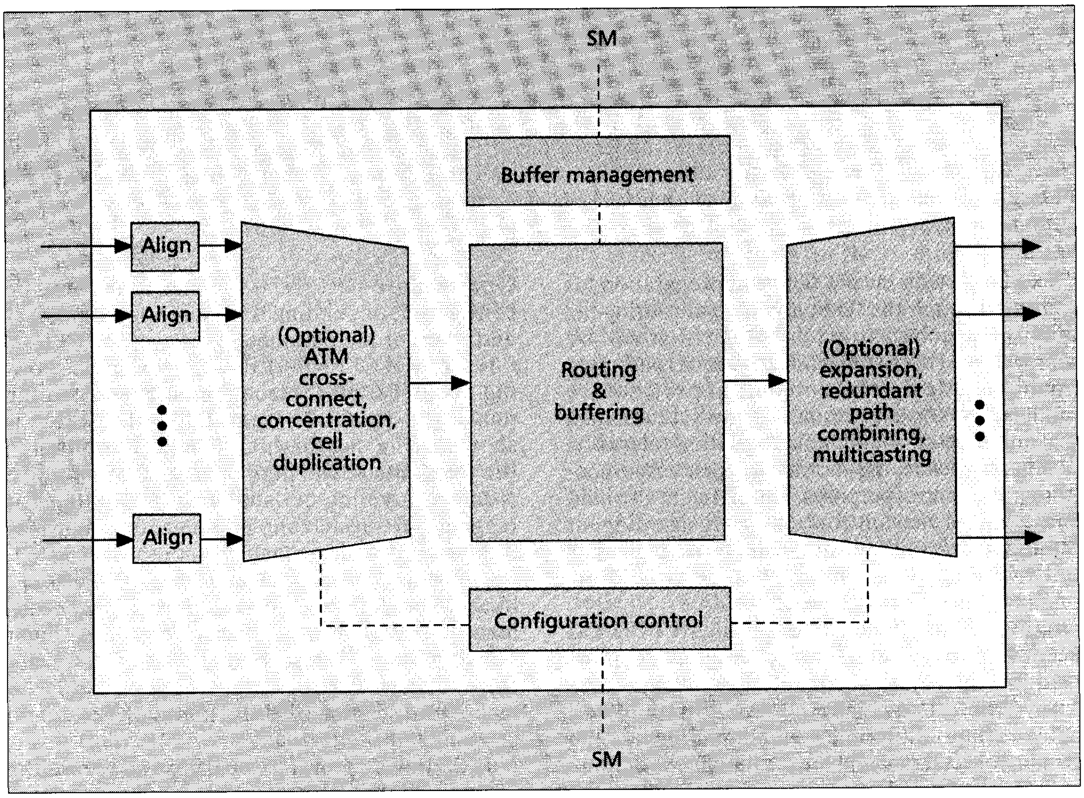

■ **Figure 10.** *Functional diagram of the CSF.*

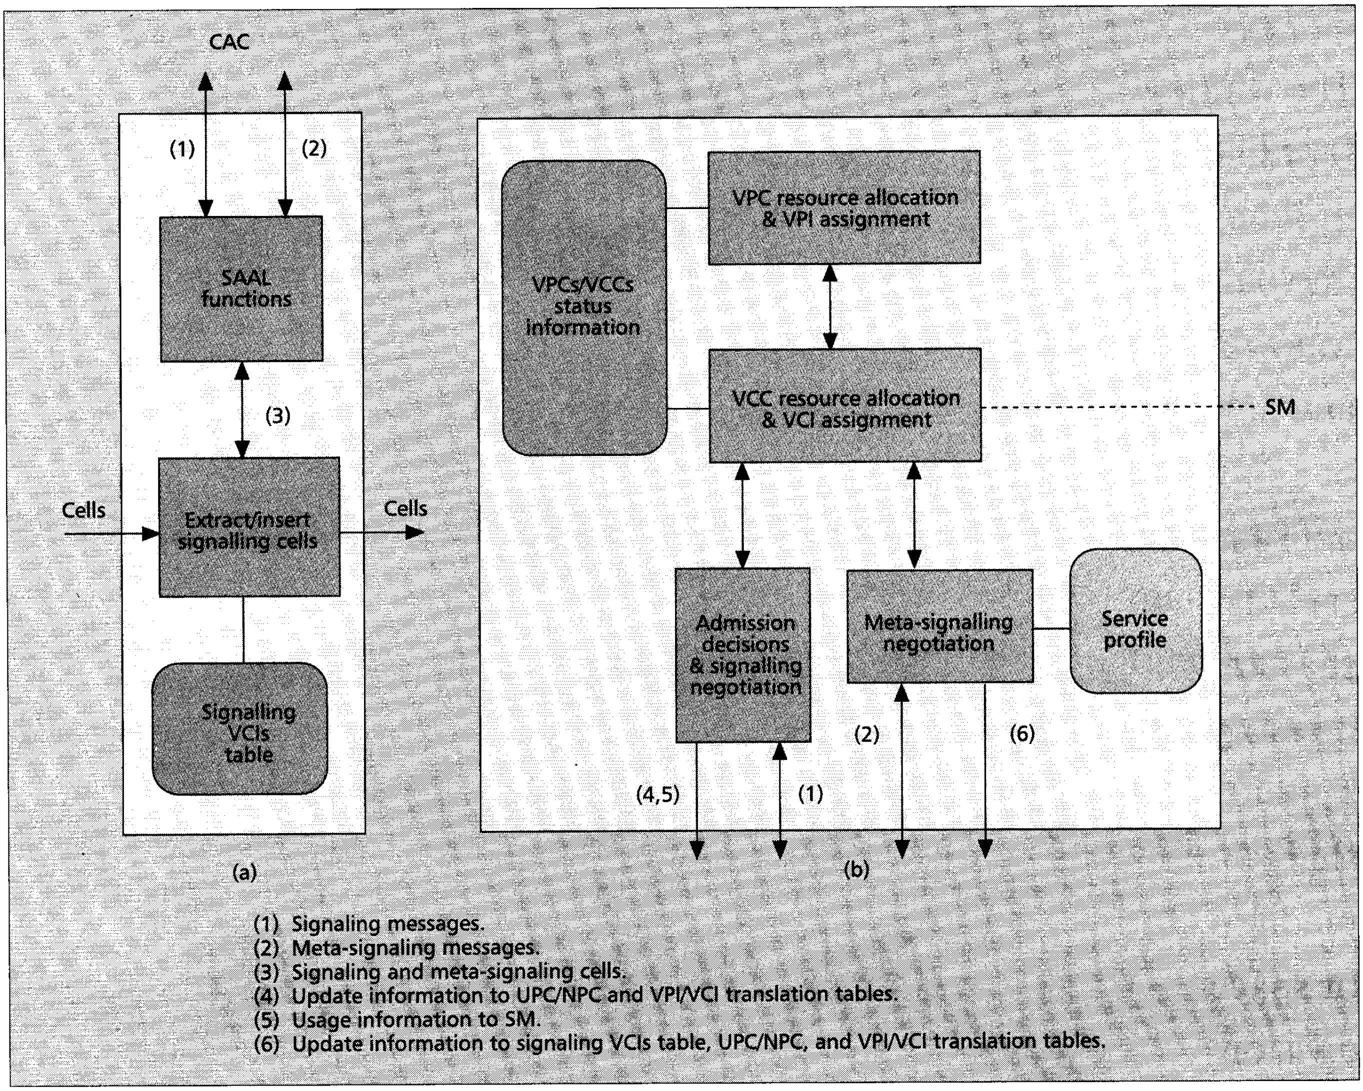

■ **Figure 11.** *Functional diagram of a) an IM-CAC and b) CAC.*

should provide some amount of redundancy for reliability. This may be achieved by redundant parallel fabric planes (e.g., [47]) or redundant paths in a single fabric plane (e.g., [48]).

Concentration before the routing fabric could reduce the size of the routing fabric (e.g., a 10:1 concentration ratio is suggested in [49]). Alternatively, this concentration could occur in a remote multiplexer outside of the switch. An optional expansion stage at the output side could combine redundant cells from the parallel fabric planes and perhaps perform multicasting or broadcasting if required. Some cross-connect capability in the CSF could be useful for protection switching and load balancing purposes.

The buffer management must keep track of queue statistics and indicate alarms to the SM if congestion is detected. To monitor performance and resource usage in the routing fabric, the buffer management may use housekeeping or timing information contained in the internal routing tag attached to each cell. It must collect sufficient information to determine the state and nature of congestion of the CSF, e.g., whether congestion is increasing or receding, and whether it is focused or general. Depending on the state and nature of congestion, it may adjust the cell scheduling scheme and selective cell discarding in the routing fabric. Upon request from the SM, the buffer management should be able to provide performance data, congestion information, records of discarded cells for analysis, and usage measurement data for accounting management. The SM may decide to activate the Explicit Forward Congestion Indication, in which case the CSF will provide the capability to change the PT codes in cell headers to indicate congestion.

Connection Admission Control

The basic function of the CAC is negotiation of a traffic contract with the user through signaling and allocation of resources to support the negotiated QOS. For discussion we assume that a part of the CAC functions are distributed to the input modules, referred to as IM-CACs in Fig. 7.

As shown in Fig. 11a, the IM-CAC recognizes signaling cells (either by their VCIs or an explicit field in their internal routing tags) within the incoming cell streams. As appropriate, it will pass, insert, or extract signaling cells. It would be convenient to handle meta-signaling cells as well, although meta-signaling is considered to be part of the Management Plane. After signaling cells are extracted, they require SAAL processing to convert them from cells into signaling messages that

The routing fabric may be implemented following any number of approaches such as shared memory buffer, shared medium, or space-division network.

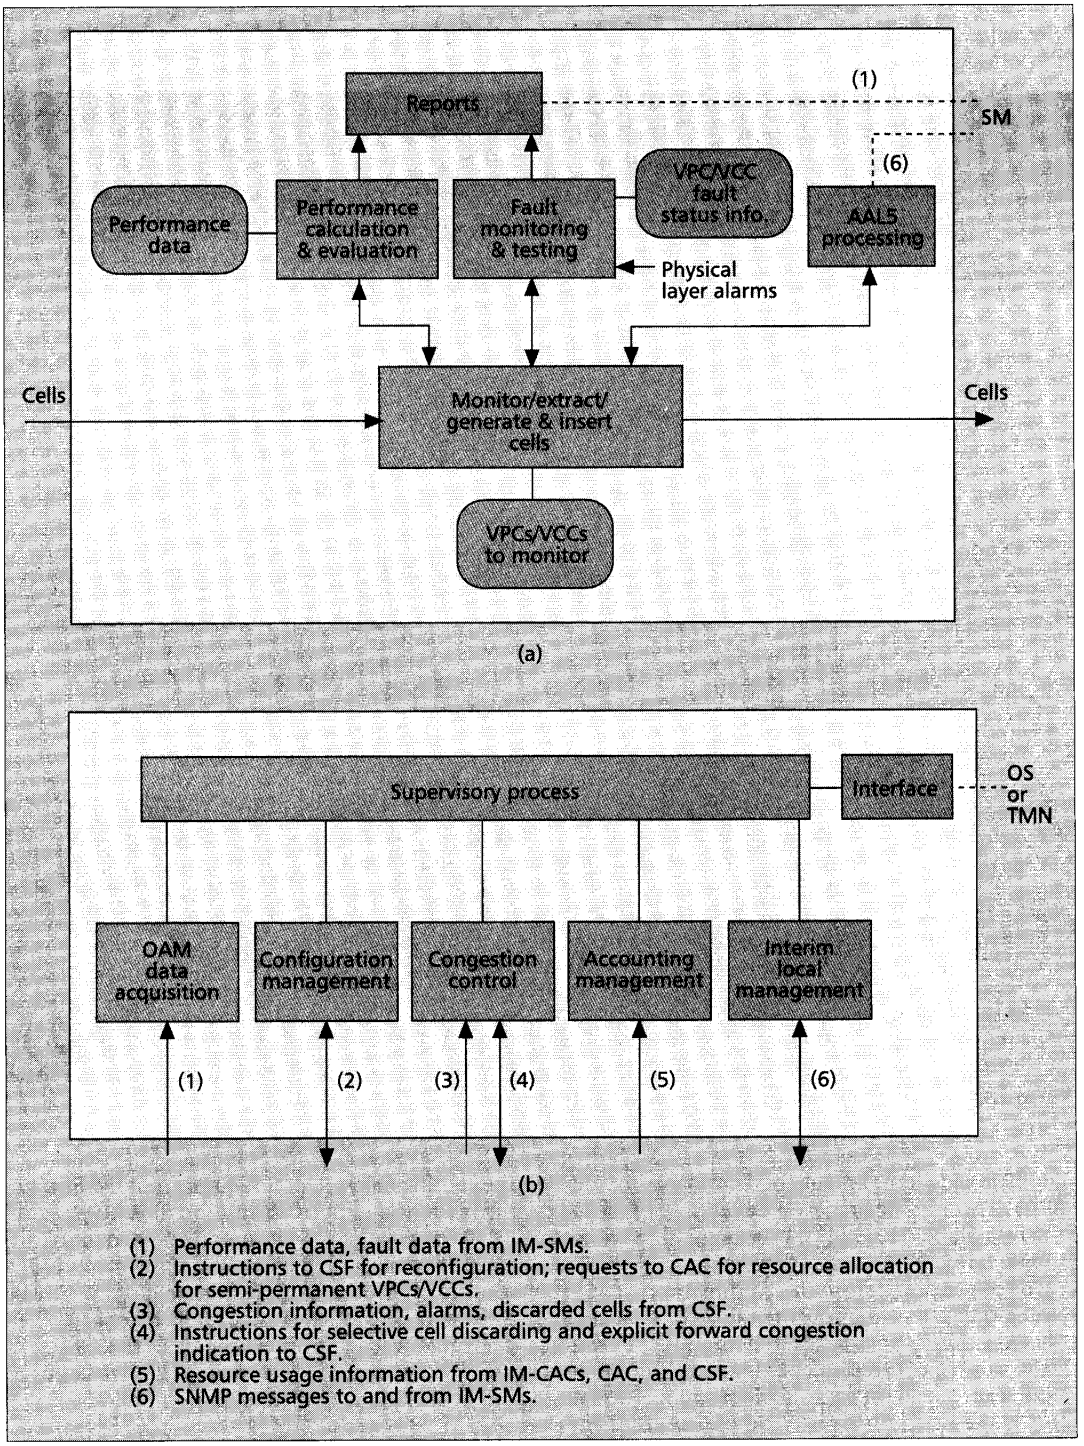

■ **Figure 12.** *Functional diagram of a) an IM-SM and b) SM.*

can be understood by the CAC. Likewise, signaling messages from the CAC must be encapsulated into ATM cells before insertion into the cell stream flow. Without the IM-CACs, the CSF could be used to route signaling cells to and from the CAC. In that case, the CAC would be responsible for SAAL functions.

A functional diagram of the CAC is shown in Fig. 11b. Meta-signaling messages are processed in the meta-signaling negotiation block. When a user requests a signaling VC through meta-signaling, its service profile identified by the service profile identifier (SPID) specifies its level of service. A signaling VC is allocated resources and assigned a VCI value which is communicated to the user. A notification is sent to the signaling VCIs table (to indicate that a signaling VC has been activated); UPC/NPC parameters are generated; and the VPI/VCI translation tables are updated.

Through the signaling VC, the user may request a new VCC for the exchange of user cells. Based on the requested QOS and source traffic descriptor, an appropriate VPC will be chosen, or if one does not exist, a new VPC will be established. If the VPC has sufficient capacity for the new VCC, then the VCC is accepted. The method for determining sufficiency will be decided by the network provider. If the VPC has insufficient capacity for the requested VCC, then its capacity is increased if possible. If the capacity cannot be increased sufficiently, the requested VCC is not accepted. Whenever a VCC is established or changed, the appropriate tables (UPC/NPC parameters, VPI/VCI translation, VPC/VCC status information) are updated and the changes are reported to the SM for accounting management. Upon request from the SM, the CAC will allocate resources for semi-permanent VPCs/VCCs.

System Management

The SM is complex because it has many responsibilities. It handles Management Plane functions and communications with the OS or TMN. For discussion we assume that a part of the SM functions are distributed to the input modules, designated as IM-SMs in Fig. 7. An example of an IM-SM is shown in Fig. 12a. The IM-SMs perform some processing for ATM Layer OAM. The VPCs/VCCs to monitor are identified in a table. For performance management, user cells are monitored as they pass, and performance data (e.g., error check) is calculated. The performance monitoring cell following the user cells is inspected (and extracted if appropriate), and its contents are compared with the calculated performance data. The comparison results may be reported to other nodes via performance management cells and reported to the SM.

The IM-SM also performs ATM Layer OAM fault management. When alarms from the Physical Layer are detected, it will generate and receive VP/VC-AIS and VP/VC-RDI cells according to the procedures described earlier. If continuity checking is exercised, it keeps track of inactive VPCs/VCCs and generates Continuity Check cells periodically, or generates VP/VC-RDI if any VPC/VCC is inactive for too long. In addition, it monitors and exchanges loopback cells, or upon request from the SM, it initiates the loopback procedure. All fault information is reported to the SM.

If the switch supports UNIs, the IM-SM recognizes ILMI cells and extracts them from the incoming cell stream. The ILMI cells go through AAL5 processing to recover the contained SNMP messages which are then passed to the SM.

A functional diagram of the SM is shown in Fig. 12b. The OAM data acquisition block collects information concerning performance and faults per VPC/VCC that is reported by the IM-SMs. It analyzes this data and decides on the appropriate course of action. For example, indications of acute failure may require deactivating the faulty equipment and reconfiguring the CSF. These changes may be controlled through the configuration management block. The configuration management block may also control semi-permanent virtual connections through negotiation with the CAC.

The congestion control block continually collects data about impending congestion and receives alarms triggered by severe congestion in the CSF. It may request additonal information from the CSF such as records of discarded cells for further analysis. It decides on the level and nature of congestion and the appropriate actions to be taken by the CSF such as selectively discarding cells or activating Explicit Forward Congestion Indication.

The accounting management block collects data about resource usage per VPC/VCC from the IM-CACs, CAC, and CSF. If customer billing is handled by centralized network management, the SM processes the collected information and sends billing reports to the OS or TMN.

Finally, the interim local management block is responsible for ILMI communications. It maintains an ATM UNI MIB and collects information about the objects defined in the MIB. For each UNI supported by the switch, it contains a UME which communicates with the user across the UNI by means of exchanging SNMP messages through the IM-SMs.

Conclusions

We have reviewed the status of industry agreements on OAM and traffic control in ATM. It is clear that they will have significant implications on the design of an ATM switching system because management and control functions will be a major part of an ATM switch in addition to the switch fabric.

We believe that the control functions will be significantly more challenging to implement than the cell switch fabric which simply routes and buffers cells. Because control must be exerted at the levels of every virtual connection and individual cells, we believe that the processing involved may be intensive. Hence the distribution of these control and management functions throughout the ATM switching system will be an important design consideration.

A functional architecture model of an ATM switching system with basic control and management functions has been developed in this paper. In the model, these control functions are partially distributed to the input modules. This distribution allows the control functions to be performed in parallel, and OAM cells may maintain their positions in the cell streams. However, it suggests that the input modules may be rather complex to be capable of the required processing.

The SM is complex because it has many responsibilities. It handles Management Plane functions and communications with the OS or TMN.

References

[1] ITU-T Rec. I.121, Broadband Aspects of ISDN, Melbourne, Nov. 1988.
[2] Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, TA-NWT-001113, issue 2, July 1993.
[3] F. Tobagi, "Fast packet switch architectures for Broadband Integrated Services Digital Network," Proc. IEEE, vol. 78, pp. 133-167, Jan. 1990.
[4] ITU-T Rec. I.321, B-ISDN Protocol Reference Model and Its Applications, Matsuyama, Nov. 1990.
[5] ITU-T Rec. I.432, B-ISDN User-Network Interface - Physical Layer Specification, Geneva, June 1992.
[6] ANSI T1.105, American National Standard for Telecommunications - Digital Hierarchy - Optical Interface Rates and Formats Specifications (SONET), 1988.
[7] ITU-T Rec. G.709, Synchronous Multiplexing Structure, Melbourne, Nov. 1988.
[8] ITU-T Rec. I.362, B-ISDN ATM Adaptation Layer (AAL) Functional Description, Geneva, June 1992.
[9] ITU-T Rec. I.361, B-ISDN ATM Layer Specification, Geneva, June 1992.
[10] ITU-T Rec. I.311, B-ISDN General Aspects, Geneva, June 1992.
[11] ATM Forum, ATM User-Network Interface Specification Version 3.0, Sept. 10, 1993.
[12] J-Y. LeBoudec, "The asynchronous transfer mode: a tutorial," *Comp. Networks and ISDN Sys.*, vol. 24, pp. 279-309, 1992.
[13] E. Sykas *et al.*, "Overview of ATM networks: functions and procedures," *Comp. Commun.*, vol. 14, pp. 615-626, Dec. 1991.
[14] ITU-T Rec. M.30, Principles for a Telecommunications Management Network, Melbourne, Nov. 1988.
[15] K. Terplan, Communication Networks Management. 2nd ed., (Prentice-Hall, 1992).
[16] ITU-T Rec. I.610, B-ISDN Operations and Maintenance Principles and Functions, Geneva, June 1992.
[17] S. Yoneda, "Broadband ISDN ATM layer management: operations, administration, and maintenance considerations," *IEEE Network*, vol. 4, pp. 31-35, May 1990.
[18] AT&T Bell Laboratories, Engineering and Operations in the Bell System. NJ, 1983.

■ ■ ■ ■ ■

Management and control functions will be a major part of an ATM switch in addition to the switch fabric.

[19] D. Tow, "Network management - recent advances and future trends," *IEEE J. on Sel. Areas in Commun.*, vol. 6, pp. 732-741, May 1988.
[20] M. Gerla, L. Kleinrock, "Flow control: a comparative survey," *IEEE Trans. on Commun.*, vol. COM-28, pp. 289-310, April 1980.
[21] ITU-T Rec. I.371, Traffic Control and Congestion Control in B-ISDN, TD 69 (XVIII), June 1992.
[22] Bellcore, Generic Requirements for Operations of Broadband Switching Systems, TA-NWT-001248, issue 2, Oct. 1993.
[23] P. Boyer, D. Tranchier, "A reservation principle with applications to the ATM traffic control," *Comp. Networks and ISDN Sys.*, vol. 24, pp. 321-334, 1992.
[24] J. Hyman *et al.*, "Real-time scheduling with quality of service constraints," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 1052-1063, Sept. 1991.
[25] Y. Takagi, *et al.*, "Priority assignment control of ATM line buffers with multiple QOS classes," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 1078-1092, Sept. 1991.
[26] A. Demers, *et al.*, "Analysis and simulation of a fair queueing algorithm," *Internet. Res. & Exper.*, vol. 1, pp. 3-26, Sept. 1990.
[27] A. Parekh, R. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single-node case," *IEEE/ACM Trans. on Networking*, vol. 1, pp. 344-357, June 1993.
[28] C. Partridge, Gigabit Networking. (Addison-Wesley, 1993).
[29] P. Boyer *et al.*, "Spacing cells protects and enhances utilization of ATM network links," *IEEE Network*, vol. 6, pp. 38-49, Sept. 1992.
[30] R. Onvural, Asynchronous Transfer Mode Networks: Performance Issues, (Artech House, 1993).
[31] K. Sato, *et al.*, "Broadband ATM network architecture based on virtual paths," *IEEE Trans. on Commun.*, vol. 38, pp. 1212-1222, Aug. 1990.
[32] ITU-T Draft Rec. Q.93B, B-ISDN User-Network Interface Layer 3 Specification for Basic Call/Bearer Control, Geneva, May 1993.
[33] ITU-T Draft Rec. Q.SAAL1, Service Specific Connection Oriented Protocol (SSCOP) Specification, Geneva, May 1993.
[34] ITU-T Draft Rec. Q.SAAL2, SSCF for Signaling at the UNI, Geneva, May 1993.
[35] F. Vakil and H. Saito, "On congestion control in ATM networks," *IEEE LTS*, pp. 55-65, Aug. 1991.
[36] H. Saito *et al.*, "Traffic control technologies in ATM networks," *IEICE Trans.*, vol. E74, pp. 761-771, April 1991.
[37] R. Guerin *et al.*, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 968-981, Sept. 1991.
[38] E. Rathgeb, "Modeling and performance comparison of policing mechanisms for ATM networks," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 325-334, April 1991.
[39] L. Dittmann *et al.*, "Flow enforcement algorithms for ATM networks," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 343-350, April 1991.
[40] M. Butto *et al.*, "Effectiveness of the 'leaky bucket' policing mechanism in ATM networks," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 335-342, April 1991.
[41] T. Helster and M. Izzo, "Functional architecture for a next generation switching system," *INFOCOM '90*, pp. 790-795.
[42] K. Suzuki *et al.*, "An ATM switching system — development and evaluation," *NEC Res. & Dev.*, vol. 32, pp. 242-251, April 1991.
[43] Y. Sakurai, *et al.*, "ATM switching system for B-ISDN," *Hitachi Rev.*, vol. 40, pp. 193-198, 1991.
[44] M. Karol, *et al.*, "Input versus output queueing on a space-division packet switch," *IEEE Trans. on Commun.*, vol. COM-35, pp. 1347-1356, Dec. 1987.
[45] G. Awater and F. Schoute, "Optimal queueing policies for fast packet switching of mixed traffic," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 458-467, April 1991.
[46] H. Kroner *et al.*, "Priority management in ATM switching nodes," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 418-427, April 1991.
[47] W. Fischer *et al.*, "A scalable ATM switching system architecture," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 1299-1307, Oct. 1991.
[48] A. Itoh, "A fault-tolerant switching network for BISDN," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 1218-1226, Oct. 1991.
[49] T. Banwell *et al.*, "Physical design issues for very large scale ATM switching systems," *IEEE J. on Sel. Areas in Commun.*, vol. 9, pp. 1227-1238, Oct. 1991.

Biographies

THOMAS M. CHEN [M '89] received a Ph.D. in electrical engineering from the University of California at Berkeley, and M.S. and B.S. degrees from the Massachusetts Institute of Technology. He joined GTE Laboratories, Waltham, Massachusetts, in 1989 where he is currently involved in research in ATM switching and network control. His e-mail address is: tmc0@gte.com.

STEPHEN S. LIU [SM '88] received M.S. and Ph.D. degrees in electrical engineering from Georgia Institute of Technology. He joined GTE Laboratories, Waltham, Massachusetts, in 1981. Since then he has been working on packet switching technology. He is a principal member of technical staff in the Network Transport Systems Department. His current research interest is traffic management and control of wide-area ATM networks for broadband multimedia services. His e-mail address is: ss10@gte.com.