Monitoring and Control of ATM Networks Using Special Cells

Thomas M. Chen, Steve S. Liu, David Wang, Vijay K. Samalam, Michael J. Procanik, and Dinyar Kavouspour, GTE

Abstract

This article describes a framework for monitoring and controlling ATM networks based on the use of management cells. By management cells we refer to a broad class of special cells that circulate around the network to perform a variety of useful functions for monitoring and optimizing the operation of the network. The proposed framework includes current ATM-layer OAM (operations and maintenance) methods as well as new mechanisms for more sophisticated long-term network management. This article explores various existing and new uses of management cells for performance monitoring, traffic control, fault management, and network administration.

Asynchronous transfer mode (ATM) is a connection-oriented fast packet-switching protocol that has been standardized for the broadband integrated services digital network (B-ISDN). All types of user information are carried in 53-byte cells (packets) consisting of a 5-byte header and 48-byte information field. An ATM network provides end-to-end transport of user cells along virtual connections with a specified quality of service (QoS). The QoS is the end-to-end network performance perceived by users for a connection, characterized mainly in terms of the cell loss ratio and cell delays. A range of voice, video, and data services are supported above the ATM layer through a number of service-specific ATM adaptation layers.

ATM exemplifies the fast packet-switching idea of streamlining the network protocol to enable high-throughput, low-delay cell switching. For high-speed switching, the cell header is simple and consists mainly of the virtual path/virtual channel identifier (VPI/VCI) fields to indicate the connection. There are two additional header fields for control purposes: a cell loss priority (CLP) bit to enable low loss priority (CLP = 1) cells to be discarded before high loss priority (CLP = 0) cells, and a payload type (PT) field to identify different types of cells. The PT field also allows explicit forward congestion indication (EFCI), whereby a cell can indicate the presence or absence of congestion somewhere along its route.

A number of common fields are notably absent from the ATM cell header, such as sequence numbers, acknowledgments, and error protection for the information field. The simplicity is intentional in order to relay user cells through the network with minimal processing delays. However, certain header fields would be useful for the purposes of network control and management. For example, cell sequence numbers would allow easy detection of cell loss, and time stamps would allow measurements of cell delays. The absence of these types of header fields has implications for network operations, which is becoming more important as ATM is being deployed more widely. In the absence of useful header fields, other means must be devised to support monitoring and control capabilities in the ATM network.

This article explores the uses of special cells for a wide range of monitoring and control functions. The basic idea is to use the ATM cell format as a flexible vehicle to carry all types of management information. These special cells have the usual cell header (identified by a particular VCI or PT field code) for fast transport through the ATM network, but their payloads contain measurement or control data that may require processing by network nodes. The ATM cell is designed to move information quickly, and hence is a fast and convenient vehicle for conveying control information as well as user data. The special cells take advantage of the ATM transport facilities at the cost of additional bandwidth and processing at network nodes. We refer to this general class of special cells as *management cells* in order to distinguish the broad class from some of its particular members that have been standardized already, such as operations and maintenance (OAM) cells [1], resource management (RM) cell [2, 3], and test cell [4]. They are called management cells to reflect their use for general management and control functions. OAM cells and other examples of special-purpose cells have been defined to serve particular needs without consideration of their membership in a broader class. By studying management cells as a general class, it may be possible to anticipate uses for these special cells before their needs are encountered in practice.

For discussion, management cells can be classified into certain categories. First, *regular* management cells are generated periodically and circulate on a connection, for example, for in-service performance monitoring. On the other hand, *on-demand* management cells are generated only as needed, for instance, to test the integrity of a path. Second, management cells are said to be *associated* (with a user VP or VC connection) if they are embedded within a user cell stream; they

0890-8044/96/$05.00 © 1996 IEEE

Performance aspect	QoS parameter	Definition
Speed	Max. cell transfer delay	α-percentile of end-to-end cell delays (worst-case cell delay)
	Peak-to-peak cell delay variation	Difference between maximum cell delay (above) and (fixed) minimum cell delay
Dependability	Cell loss ratio	Ratio of lost cells to total cells transmitted, specified for CLP = 0 + 1 or only CLP = 0 cells
Accuracy	Cell error ratio	Ratio of cells with payload errors to total cells transmitted (excluding severely errored cell blocks)
	Severely errored cell block ratio	Ratio of cell blocks with more than *M* errored cells to total cell blocks transmitted
	Cell misinsertion rate	Number of misinserted cells divided by the measurement time interval

■ Table 1. *QoS parameters.*

have the same VPI or VPI/VCI value as the user cells and follow the same route. Because they may experience the same QoS as the user cells, they are very useful for in-service performance monitoring. An example is the OAM performance management cell. In contrast, *nonassociated* management cells are not constrained to follow the same route as a user cell stream and will use an exclusive VP or VC connection (i.e., not used by user cells). They are appropriate for functions such as out-of-service testing. Third, like OAM cells, management cells can be referred to as *F4* or *F5 type*, depending on whether they are routed at the virtual path or virtual channel level, respectively. F5 type management cells are carried within the same virtual channel as user cells (with the same VPI/VCI but different PT); F4 type management cells are carried within the same virtual path but in a separate virtual channel from user cells (with the same VPI but different VCI). Finally, they may be categorized according to their scope, defined by the two location points where the cells are generated and finally terminated. A management cell may be *user-to-user* (exchanged between two users transparent to the network), *user-to-network*, *network-to-user*, *node-to-node* (carried and processed within a single network), or *network-to-network* (carried and processed by multiple networks).

Because management cells may involve specific processing procedures in the network, they might be viewed as an adjunct protocol to ATM. The idea of an adjunct protocol defined by a special class of control packets has been practiced for other network protocols, for example, the Internet Control Message Protocol (ICMP) messages in the Internet Protocol (IP). Within our broad definition, management cells may conceivably carry many different types of information, such as:

- Measurement data (e.g., time stamps, error checks, cell counts)
- Network state information (e.g., congestion level, alarms, available bandwidth)
- Control information (e.g., control parameters, routing changes, bandwidth reallocation)
- Administrative information (e.g., resource usage, passwords, encryption keys)

Since the payload can possibly be any type of management information, the only limitation on the utility of management cells is the processing capabilities of each node. Specific examples of management cells and procedures will be covered throughout the article. The following section discusses procedures to use management cells for performance monitoring beyond current OAM capabilities. The third section explores ideas for using management cells in traffic control. The fourth section suggests possible uses of management cells in fault detection and recovery. Finally, the fifth section outlines potential uses in network administration, such as security and customer network management.

Performance Monitoring

Quality of Service

Since some B-ISDN services depend on a minimum level of network performance (e.g., real-time services depend on bounded end-to-end cell delays), the QoS is specified and guaranteed for each ATM connection. The QoS is characterized by the parameters listed in Table 1 which quantify the end-to-end impairments experienced by cells along a connection [3]. Cell delays and cell loss ratio are usually the main parameters of concern because cell errors depend on the physical transmission layer. The particular set of QoS parameters specified for a connection will depend on the type of service. The QoS specification is part of the so-called traffic contract established between the network and users during connection admission control (CAC) [2, 3]. For their part, users specify the desired QoS and their traffic rate parameters (e.g., peak cell rate) and agree to conform to them for the duration of the connection. Based on these parameters, the network estimates the resources required for the connection and determines whether the necessary resources are available. In return for users conforming to their declared traffic parameters, the network agrees to maintain the specified QoS for that connection (if the QoS is unattainable, the requested connection will be rejected by CAC).

CAC is the network's primary means of preventing congestion and sustaining QoS because traffic sources are generally not controllable by feedback after a connection has been accepted (with the exception of the new available bit rate, or ABR, service [3]). CAC depends on a resource allocation algorithm that attempts to estimate the effect of accepting new connections. Numerous approaches have been proposed, but their effectiveness depends on the accuracy of the source traffic models and choice of optimization criteria. Accurate source models have been difficult to develop, and even with valid source models, QoS estimation is usually a very difficult problem due to the random interactions among different traffic streams.

In practice, CAC and other network controls will work in combination with performance monitoring. Performance monitoring involves the continuous collection of measurement data for evaluation of network behavior. Performance measurement data is essential input for correct and optimal control decisions, and serves to evaluate the consequences of control actions. In addition, performance data is indispensable for network planning, provisioning, maintenance, and accounting. Performance monitoring is especially important in ATM networks because ATM services are more complicated than simple dial-tone service in telephone networks or best-effort service in packet-switched networks. Monitoring is also more challenging because it is insufficient to measure the general overall performance of the ATM network; performance monitoring must be exercised *per connection* because a different QoS is guaranteed on each connection.

A distinction is usually drawn between in-service performance monitoring and out-of-service testing, which are car-

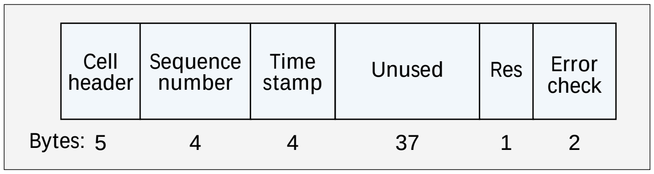

Figure 1. *Test cell.*

ried out with different types of management cells. Out-of-service testing consists of performance measurements of a system (one or more switches) under controlled traffic conditions without real user traffic. The objective is to characterize the system performance, verify connectivity, confirm the interoperability between switches, or validate that the performance conforms to design specifications. The type of management cell under consideration in ATM standards is the *test cell* [4]. As shown in Fig. 1, the test cell payload contains a number of fields to measure performance:

- Sequence number to detect test cell loss or misinsertion
- Time stamp to measure cell delay and delay variation
- Error check field to detect bit errors in the payload

Typically, test cells are generated and injected into the system under test; some may be used as unobserved "background" traffic, while others are captured and analyzed upon exit.

Out-of-service testing is necessary to certify equipment before deployment and to diagnose trouble in faulty equipment. However, a tested system is not guaranteed to perform adequately under real network conditions, and diagnostic testing is done only after trouble has been detected and the system has been configured out of service. In contrast, in-service performance monitoring is carried out continuously on connections with real user traffic. It can accurately measure the end-to-end QoS seen by users and detect troubles quickly before they become serious. For these reasons, in-service monitoring is currently used to measure the bit error performance of digital transmission links [5] and synchronous optical network/synchronous digital hierarchy (SONET/SDH) links [6, 7].

In-Service Performance Monitoring by OAM Cells

For in-service monitoring, the type of management cell prescribed in current ATM standards is the *OAM performance management cell* [1, 8–10]. They are inserted between blocks of user cells and carried in the user cell stream without change in relative position or modification of payload fields. The OAM cells contain information about the preceding user cell block. At the destination, this information is compared with the received user cell block. As shown in Fig. 2, the OAM performance management cell is generated with these fields for performance monitoring:

- Monitoring cell sequence number (MCSN) to detect the loss of OAM cells and ensure their proper sequence
- Total number of CLP = 0 + 1 user cells (TUCN_{0+1}) and CLP = 0 user cells (TUCN_0) that have been transmitted up to the OAM cell
- Block error detection code (BEDC) calculated over the preceding block of user cells
- An optional time stamp (TS)

The destination endpoint records the total number of received CLP = 0+1 user cells (TRCC_{0+1}) and CLP = 0 user cells (TRCC_0), and writes the recalculated error code in the block error result (BLER). The destination endpoint uses TUCN_{0+1} to infer cell loss or misinsertion (or TUCN_0 for only CLP = 0 cells). From the TUCN_{0+1} field in the previous OAM cell, the destination can determine the number of user cells (say N) transmitted in the cell block between the previous OAM cell and the current one. If $M < N$ user cells were received in the preceding block, then it can be inferred that $N - M$ user cells were lost. Likewise, if $M > N$ user cells were received, it can be inferred that $M - N$ user cells were misinserted. The optional time stamp field is intended to enable cell delay measurements, but no procedure is currently defined. Conceivably, the destination may use the time stamp to infer the end-to-end cell delay, but this requires the source and destination to be synchronized to the same time of day. Alternatively, the destination may loop back the time stamp to the source endpoint to measure the round-trip delay, but this may not accurately estimate the one-way delay [11].

The current OAM performance management procedure has the advantage of requiring no special functions in ATM switches (unless the switch has the role of a source or destination endpoint). Intermediate switches simply relay OAM cells without modifying or adding information. Thus, ATM switch implementations can be simple, but the capabilities of the procedure are limited. For example, the procedure cannot determine the exact number of lost or misinserted cells. In the case $M = N$, for instance, it will be inferred that no cells were lost or misinserted, but it is possible that an equal number of cells were lost and misinserted. In addition, the procedure cannot determine where cells were lost or misinserted along a connection. For cell delay measurements, clocks at the source and destination endpoints must be synchronized to the same time of day; otherwise, only a round-trip delay can be measured unambiguously.

Improved Monitoring by Management Cells

The current OAM performance monitoring procedure can be improved by allowing ATM switches to modify management cells with information that each switch may collect about its own internal performance (we now discuss management cells instead of OAM cells to distinguish this procedure from the current OAM standards). Like OAM cells, management cells may be inserted between blocks of user cells and carried in the user cell stream without changing their relative positions. However, to improve the monitoring procedure, management cells will carry different information in the payload, which will require specific processing capabilities at each switch [12].

To accurately measure cell delays, each switch could measure the ingress-to-egress delay experienced by a management cell through that switch and write this delay amount into the management cell upon its departure. This can be done by timestamping a cell when it is received at a switch input, and reading the time stamp just before the cell is transmitted from

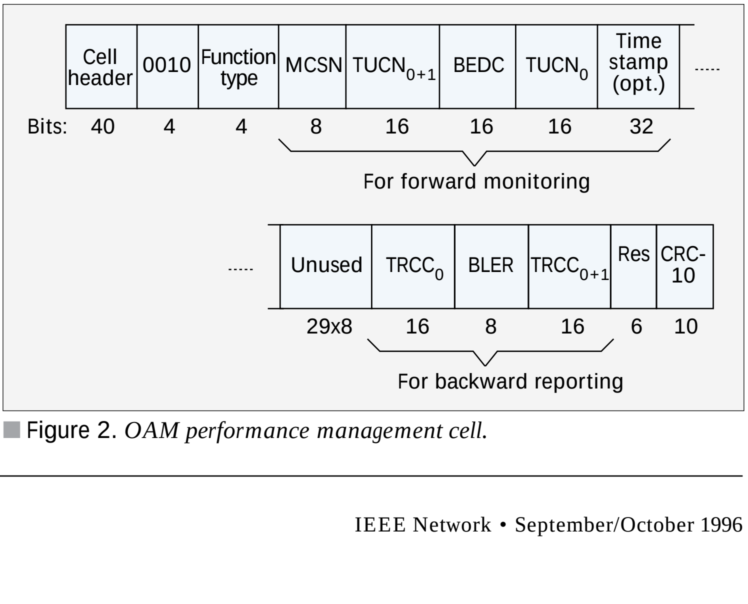

Figure 2. *OAM performance management cell.*

the switch. The difference between the egress and ingress times can be added to a field in the cell when it is transmitted. As shown in Fig. 3a, this "delay-stamp" field in the management cell would be initialized to zero and accumulate the delays measured at every switch along a connection (fixed propagation delays can be precomputed and added by each switch). At the destination, the delay-stamp field would reveal the end-to-end cell delay experienced by the management cell along the connection. Since (associated) management cells experience the same QoS as user cells, this cell delay measurement would be a random sample of the cell delay seen by the user cells in that connection. Alternatively, a management cell could have multiple delay-stamp fields, as shown in Fig. 3b. As it passes through each switch, the switch writes the measured ingress-to-egress delay into the next empty delay-stamp field. Thus, the management cell could record the delays experienced at individual switches along the connection, which may be useful for diagnostic purposes. This procedure can accurately measure end-to-end cell delays without the need for clock synchronization between connection endpoints as required in the OAM procedure.

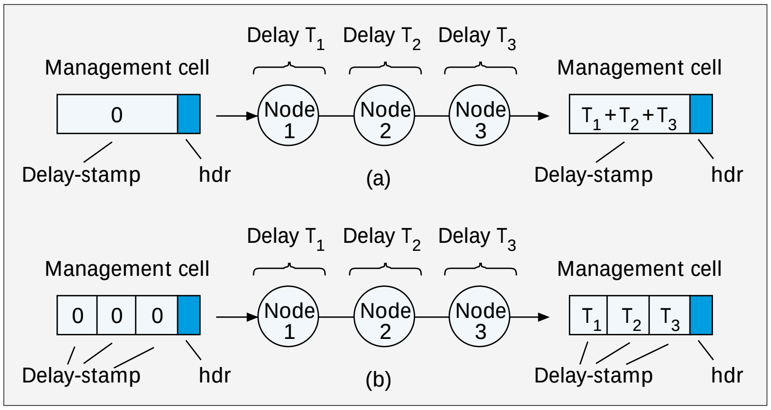

■ Figure 3. *Cell delay measurements using a management cell with a) one delay-stamp field; b) multiple delay-stamp fields.*

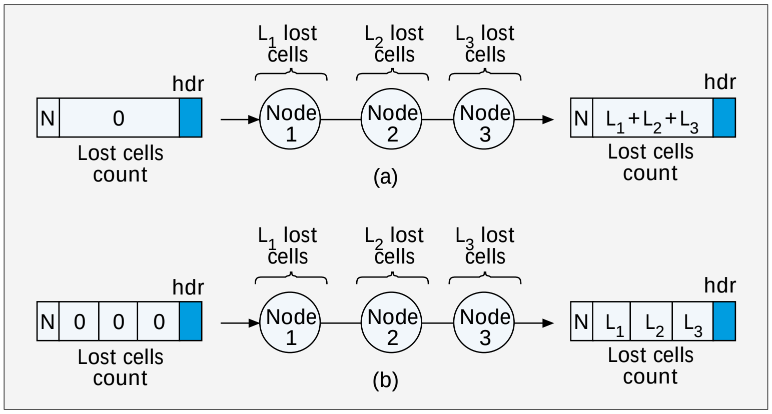

■ Figure 4. *Cell loss measurements using a management cell with a) a lost cells count field; b) multiple lost cells count fields.*

Because each measurement is a random sample of cell delay, an important issue is the statistical meaningfulness of a set of cell delay measurements. Clearly, a larger number of measurement samples is desirable for more accuracy, but there is a tradeoff between accuracy and sampling rate. Greater statistical accuracy is achieved by more samples, but the number of samples taken over a monitoring interval depends on the bandwidth allocated to the management cells. Unfortunately, the bandwidth for management cells must be practically constrained to a small fraction of the total bandwidth (e.g., 1 percent or less). This constraint will prevent management cells from significantly affecting the connection they are monitoring. Furthermore, it is recognized that the main responsibility of the network is transport of user data; hence, management cells should not impose a significant burden on the network that could detract from this responsibility. As a consequence, management cells will be generated relatively infrequently on a connection. The worst case is therefore a brief, low-rate connection; it may be impossible to collect a sufficient number of cell delay measurements for a desired level of accuracy. A lengthy high-rate connection would be ideal for accurate monitoring. For example, to measure the 99-percentile of cell delay with an accuracy of 99 ± 0.1 with 95 percent confidence, about 38,416 samples are required. For a high-rate 150-Mb/s video connection, the monitoring interval required would be about 1 s (assuming that 1 percent of the total bandwidth is dedicated to management cells). In contrast, the necessary monitoring interval for a low-rate 64-kb/s speech connection would be about 40 min. Naturally, less time would be required if the desired accuracy is relaxed.

In a similar manner, cell loss can be monitored accurately using (associated) management cells if each switch can keep count of the number of cells discarded at that switch. To accurately measure cell loss, a management cell may carry the number of transmitted cells, say N, and a field for the counted lost cells (initially zero). As shown in Fig. 4a, each switch adds the number of cells discarded at that switch since the previous management cell. At the destination, the management cell would reveal the total number of cells lost along the connection, say L, and a simple calculation L/N would yield the cell loss ratio. For more detail, the management cell might carry additional fields to record the causes of cell loss, such as header errors, usage parameter control (UPC), or buffer overflow. Alternatively, a management cell might have multiple fields, as shown in Fig. 4b, for example, a separate field for the number of cells discarded at each switch. When the management cell passes through each switch, the switch writes the number of lost cells into the next empty field. Thus, the management cell could reveal the cell loss ratio at individual switches along the connection.

As a cumulative count, the cell loss measurement is different than the cell delay measurement, which is an instantaneous sample. The number of cell loss measurements made over a monitoring interval is not an issue for accuracy as it is for cell delay measurements. However, there is another consideration that argues for frequent cell loss measurements. The cell loss ratio should be very low in practice; essentially, cell loss measurements should always be expected to be close to zero. If the cell loss becomes measurable, it probably indicates an unusually serious congestion condition. The frequency of cell loss measurements should be high to detect such an alarming condition as quickly as possible or within a target response time (if given). Again, the rate will be limited in practice because the bandwidth for management cells must be constrained to a small fraction of the total bandwidth. By implication, the response time can be faster for high-rate connections than for low-rate connections. For example, assuming 1 percent of the bandwidth is allocated to management cells, the response time for a high-rate 150-Mb/s video connection can be about 0.3 ms (the time between management cells). For a low-rate 64-kb/s speech connection, the response time can be only about 0.6 s.

Traffic Management

Congestion Detection and Notification

Direct measurements of cell delay and cell loss, as described earlier, are useful to verify that users are receiving their guaranteed QoS. For optimizing network operations, additional measures of the network state are also informative. For example, additional congestion measures might be queue lengths or buffer occupancies at each switch, or utilization (load factor) on each link. Generally, a combination of different measures will provide the most accurate information about the severity and nature of congestion. Management cells can be used for these measurements and for distribution of congestion information.

Clearly, ATM switches must be responsible for monitoring the state of congestion of its buffers. Congestion monitoring is straightforward in early switch implementations which have simple single-priority FIFO (first-in-first-out) queues. In this case, cell delays and queue lengths are directly correlated, and congestion can be measured by queue length threshold crossings. However, later-generation ATM switch implementations have or will have more sophisticated buffer systems in order to provide more flexible service classes. Generally, an ATM switch buffer system may consist of separate virtual queues for different traffic classes with dynamic space allocation between the queues. Cell scheduling may be governed by complex scheduling algorithms based on time (service) priorities, and access to buffer space may be controlled by space (loss) priorities. In a general buffer system, congestion might be distinguished into two types:

- Time congestion, for example, low-service priority traffic suffering excessive delays due to high-service priority traffic
- Space congestion caused by buffers filling up and overflowing

Because time and space congestion might occur separately in a complex buffer system due to different causes, it is desirable to identify the nature of congestion as well as its severity.

Time congestion is manifested as long delays seen by low-service-priority traffic. Hence, it can be measured directly by monitoring cell delays per connection, as described earlier. The severity of congestion may be judged by the proximity of the measured cell delays to the agreed cell delay bound. The delays for low- and high-service-priority traffic should be compared because low-service-priority traffic will be affected first before congestion becomes serious enough to affect high-service-priority traffic. Note that queue length is not a reliable measure of time congestion; for example, low-service-priority traffic may experience long delays even when its queue is short. Also, link utilization (or load) is not an accurate indication of time congestion because it is a measure of the average activity on a link, including low- and high-service-priority traffic. It would be more revealing to measure the amounts of load contributed separately by the high- and low-priority traffic, because the load contributed by the high-priority traffic has an effect on the time congestion seen by the low-priority traffic. Burstiness of the high-priority traffic may also be an informative measurement because greater burstiness implies more delay for low-priority traffic.

Field	Name	Length	Use for ABR
ID	Protocol identifier	8 bits	1 = ABR service
DIR	Direction	1 bit	0 = forward 1 = backward
BN	Backward explicit congestion notification	1 bit	0 = source generated cell 1 = switch generated cell
CI	Congestion indication	1 bit	0 = no congestion 1 = congestion
NI	No increase	1 bit	0 = additive increase allowed 1 = no additive increase allowed
R/A	Request/acknowledge	1 bit	Not used
E/R	Elastic/rigid	1 bit	Not used
RES	Reserved	2 bits	Not used
ER	Explicit cell rate	16 bits	Explicit rate feedback
CCR	Current cell rate	16 bits	Source rate limit when this RM cell was generated
MCR	Minimum cell rate	16 bits	Guaranteed cell rate
QL	Queue length	32 bits	Optional measure of maximum number of cells queued among the nodes along this connection
SN	Sequence number	32 bits	Optionally incremented for each forward RM cell, otherwise set to 0
RES	Reserved	246 bits	Not used
CRC-10	Cyclic redundancy check	10 bits	Detect errors in payload

■ Table 2. *RM cell payload fields for ABR service.*

Space congestion is manifested by increasing cell loss, first for low-loss-priority (CLP = 1) traffic and then for high-loss-priority (CLP = 0) traffic if it becomes serious. Cell loss is an obvious measure of space congestion, but this is difficult when the cell loss ratio is expected to be extremely low (e.g., 10^{-6} or less). Hence, cell loss is virtually unmeasurable (except over very long times), and it would not be informative to continually measure zero cell loss. A more useful measure is possible if the cell loss ratio is specified only for CLP = 0 cells (cell loss ratio may be specified for CLP = 0 + 1 or only CLP = 0 cells [3]). In this case, no cell loss bound is imposed for CLP = 1 cells, and loss of CLP = 1 cells may be measurable under normal network conditions to reveal the occurrence of buffer overflows (assuming that switches are capable of selectively discarding CLP = 1 cells before CLP = 0 cells).

Buffer occupancy or queue length is another obvious candidate measure of space congestion. Buffer occupancy can be related to very low cell loss ratios by the theory of large deviations, which establishes a linear relationship between buffer size and the logarithm of the cell loss ratio under certain conditions [13–15]. In this approach, the incoming traffic is monitored and copied into several conceptual "pseudo-buffers" with various sizes smaller than the real physical buffer. Since the pseudo-buffers are smaller, their cell loss ratios may be significant and measurable even over a short monitoring period. Extrapolating from these measured cell loss ratios can yield an estimate of the cell loss ratio for the physical buffer [15].

If the cell loss ratio is specified only for CLP = 0 cells, the buffer contents of CLP = 0 and CLP = 1 cells should be measured separately. A buffer filled up mostly with CLP = 1 cells has a lower risk of violating the guaranteed CLP = 0 cell loss bound than a buffer filled up mostly with CLP = 0 cells. It would be useful to know the proportion of CLP = 1 and CLP = 0 cells in the buffer in addition to the overall buffer occupancy.

Congestion information can be exchanged between nodes by the use of management cells. Note that EFCI already allows congestion notification in the forward direction using only user cells. However, management cells can carry more detailed congestion information in both the backward and forward directions. Backward notification may be useful, for example, to inform the UPC mechanism located at the source of the connection about the level of congestion. If the connection is experiencing congestion, UPC should adjust by discarding excessive cells instead of tagging them because these cells will be discarded eventually with high probability. Cell tagging makes sense when there is a high chance that CLP = 1 cells will be delivered successfully. In more severe congestion, UPC might automatically begin to discard incoming CLP = 1 cells in addition to excessive cells [16].

Congestion notification directed to the users has questionable usefulness for constant and variable bit rate (CBR/VBR) connections because their rate parameters are established during connection setup. CBR/VBR sources are not obligated to adjust their rates in response to congestion notification. On the other hand, available bit rate (ABR) service relies on feedback information from the network via management cells called RM cells (discussed below). RM cells take a round-trip between the ABR users and accumulate congestion state information from the network. In the backward direction, the RM cells inform the source about the proper adjustment of its transmission rate; they can also be used to dynamically adjust the UPC mechanism.

A drawback to the use of management cells for explicit congestion notification is the extra load imposed on the network which is certainly not desirable if the network is congested [8]. But management cells such as RM cells or OAM cells will be circulating in the network anyway. Congestion information might be simply carried in the unused portions of these cells.

Another issue concerning the effectiveness of congestion monitoring is the time between management cells. As mentioned before, the bandwidth for management cells will be constrained to a small fraction of the total bandwidth, which means that management cells may be generated infrequently especially for low-rate connections. Recalling an earlier example, the time between management cells will be 0.6 seconds on a low-rate 64-kb/s connection (assuming one percent of the bandwidth is dedicated to management cells). If congestion can develop faster than that, the network may be too slow to detect and react to it. In this case, there must be significant overdimensioning to prevent congestion or additional mechanisms to predict and avoid congestion as early as possible.

Feedback Rate Control by RM Cells

Congestion monitoring and notification have a central role in the ABR service specified recently [3]. The nature of the ABR service is different from the usual CBR/VBR services [17]. Network resources are allocated to CBR/VBR connections during connection establishment; sources are not controllable by feedback after a connection is accepted. Sufficient resources must allow for the simultaneous peak-rate bursts of VBR connections, but VBR sources may not be transmitting at their peak rates much of the time. In this case, some bandwidth will be left unused. ABR connections are allowed to share this time-varying excess capacity in the network without effecting the QoS of CBR/VBR connections. ABR sources attempt to adapt their rates to their proper share of the available bandwidth by obtaining feedback information from the network. The feedback information is carried in management cells, specifically RM cells [2, 3]. The RM cell fields are listed in Table 2.

The feedback control operates in two modes, explicit binary or explicit rate indication. In both cases, RM cells are generated by the source between blocks of user cells and returned in the backward direction by the receiver. The explicit binary indication mode requires only that network nodes are capable of EFCI (whereby a congested node may set EFCI = 1 in forward-direction user cells). The receiver monitors the EFCI state of the received cells and, if congestion is indicated, returns CI = 1 or NI = 1 in the backward RM cell. The source examines the CI and NI bits in the returned RM cell. If CI = 1, it must reduce its rate by a relative amount; if NI = 1, it is prevented from increasing its rate; and if CI = 0 and NI = 0, it is allowed to increase its rate by a fixed amount.

The explicit rate indication mode requires that network nodes are capable of calculating the proper share of available bandwidth for each source and writing this amount into the ER field in passing RM cells. The receiver will loop back the RM cell with the same ER value but has the option to reduce it. When the source receives the returned RM cell, it first calculates its rate adjustment according to the CI and NI bits as described earlier, and adjusts to the minimum of the calculated new rate or the ER field value.

Bandwidth Renegotiation by RM Cells

RM cells are also used for the new ATM block transfer (ABT) capability under standardization [2]. Both ABT and ABR services may serve similar data applications, and both will use the RM cell for control of time-varying bandwidth (although the RM cell for ABT will have the slightly different format in Table 3). However, they are fundamentally different from each other. The ABR service attempts to divide the unreserved excess bandwidth, whereas ABT is a method of bandwidth reservation (ABT has been known for years as the fast reservation protocol). The objective of ABT is to allow a data source to renegotiate its bandwidth reservation on the basis of blocks of cells.

Data sources are often very bursty, that is, mostly inactive except when higher protocol layers pass down packets to the ATM layer. The higher-layer packets are often fragmented into many ATM cells. This fragmentation causes a problem where even a low cell loss rate can cause a significant loss rate of higher-layer packets. The goodput of a connection (the amount of higher-layer packets delivered successfully) may be much less than the throughput of the connection measured in delivered ATM cells because many delivered cells can belong to a corrupted packet.

Hence, the ABT method is based on reserving bandwidth for bursts of cells, referred to as *blocks*. A block may consist of a portion of a higher-layer packet, a single packet, or several packets; a direct relationship between blocks and packets is not mandated. A block is bracketed by two RM cells. A preceding RM cell requests the bandwidth required for the block. Another RM cell follows at the end of the block to release the reserved resources. By treating cells in units of blocks, the objective is to avoid the situation where cell losses are spread over many packets and thus impair the goodput. Blocks are either discarded entirely or delivered with a high level of network performance (which could include a low residual cell loss ratio).

Initially, an ABT connection is set up with a peak cell rate but no allocated bandwidth. A peak cell rate is also established for RM cells which limits the frequency of bandwidth renegotiations. Each block negotiates for a change in the bandwidth allocation by one of two methods:

Field	Name	Length	Use for ABT
ID	Protocol identifier	8 bits	2 = ABT/DT 3 = ABT/IT
DIR	Direction	1 bit	0 = forward 1 = backward
TM	Traffic management cell	1 bit	0 = source generated cell 1 = network generated TM cell
CI	Congestion indication	1 bit	0 = bandwidth renegotiation succeeded 1 = bandwidth renegotiation failed
M	Maintenance	1 bit	0 = bandwidth renegotiation cell 1 = maintenance cell
R/A	Request/acknowledge	1 bit	0 = bandwidth renegotiation request 1 = acknowledgement
E/R	Elastic/rigid	1 bit	0 = network may overwrite cell fields 1 = network may not overwrite cell fields
RES	Reserved	2 bits	Not used
CLP = 0 + 1 BCR	CLP = 0 + 1 block cell rate	16 bits	Requested or allocated bandwidth for CLP = 0 + 1 user cells + OAM cells
OAM BCR	OAM block cell rate	16 bits	Requested or allocated bandwidth for OAM cells only
MCR	Minimum cell rate	16 bits	Not used
BS	Block size	32 bits	Number of cells in following block (ABT/IT only)
SN	Sequence number	32 bits	Optionally incremented for each RM cell (ABT/IT only)
RES	Reserved	246 bits	Not used
CRC-10	Cyclic redundancy check	10 bits	Detect errors in payload

■ Table 3. *RM cell payload fields for ABT service.*

ABT with Delayed Transmission (ABT/DT) — The RM cell preceding the block contains a request for a peak cell rate and cell delay variation tolerance. The block waits for acceptance by the network (via a backward RM cell) before it is sent. Acceptance by the network implies a guaranteed QoS equivalent to that for a CBR connection, as long as the source conforms to its given traffic parameters.

ABT with Immediate Transmission (ABT/IT) — The block is sent immediately after the preceding RM cell which contains the requested peak cell rate and cell delay variation tolerance. Because the required bandwidth may not be available, the block proceeds on a node-by-node basis. Each node has the opportunity to forward the block with the required QoS or, if unable to do so, discard the entire block.

Dynamic VP Bandwidth Management

For the purpose of efficient bandwidth management, the network may attempt to dynamically adjust the bandwidth allocated to virtual paths [18]. VPs should capture only as much bandwidth as they need to satisfy their throughput and QoS requirements so that the unused bandwidth will be available to other VPs which may need the bandwidth. As more VC connections are accepted within a VP, the VP may need to request more bandwidth to accommodate the new connections. The changes in VP bandwidth allocation may be managed centrally, but a distributed approach may be more scalable to large networks [19].

The distributed approach requires the exchange of management cells along a VP. These management cells have a role similar to signaling messages. A management cell is generated at the source endpoint and propagated along the VP with the requested bandwidth. It is examined by each node, which decides whether the request can be accommodated and writes the result into the management cell. The destination endpoint returns the management cell in the backward direction, which confirms or cancels the bandwidth change at each node.

Routing

As mentioned earlier, CAC is the network's primary means of preventing congestion and sustaining QoS because CBR/VBR sources are generally not controllable by feedback after a connection has been established. In response to a connection request, the network attempts to select a route among a set of candidate routes (routing) and makes a decision to accept or reject the connection request based on the available and required resources (resource allocation). The routing problem consists of discovering the network topology, selecting feasible routes, and distributing updated routing information. Management cells are well suited to the exchange of routing information between network nodes. Naturally, routing algorithms have been developed in both circuit- and packet-switched networks. Since ATM is connection-oriented, the routing problem is similar to that in circuit-switched networks, but more complicated due to different QoS requirements for each ATM connection. In addition, VBR traffic may mean more dynamic traffic patterns than typical telephone traffic; hence, it is useful to borrow ideas from routing in data networks.

The private network-network interface (PNNI) specification is based on routing procedures proven in packet-switched data networks [20]. Nodes learn about their neighbors by exchanging *Hello* messages. The network topology is constructed and updated based on the well-known link-state routing protocol. Each node monitors the status of its links and periodically floods this state information throughout the network in a type of management cells called *PNNI topology state packets* (PTSPs). By this process, each node maintains a complete and consistent topology map of the network. Feasible routes can be selected by the source node (source routing) or independently by each node along the route (hop-by-hop routing). To avoid loops and inconsistency, source routing is preferred [20].

Management cells can also accumulate useful measurements to compare candidate routes. As described earlier, the QoS and congestion can be continually monitored on each VC or VP connection by regular management cells. In addition, it is easy to simultaneously measure other selection criteria such as link utilization, length (number of hops), queue lengths, or available bandwidth [14, 21–25]. If source routing is used, management cells can continuously collect and report this information to the source node of each connection. Thus, in response to a new connection request, a source node will have already collected end-to-end information about every active connection (from that starting point) and can make a well-informed selection of a new route. Alternatively, a source node can send "probe" management cells to collect end-to-end information along every candidate route [26]. The probe method will collect the most current status information about candidate routes, but will increase the delay in the connection setup process. If hop-by-hop routing is used, management cells need to make round-trips so that all nodes along a connection will gain consistent information from backward-traveling cells. Each node will have to read backward-traveling management cells and keep track of information pertaining to all connections.

After a candidate route is selected, the network attempts to estimate the effect on QoS of accepting the new connection. This is generally a difficult problem because ATM sources may have widely different and unpredictable characteristics, minimal traffic parameters (perhaps only peak cell rate) may be known in advance, and traffic flow characteristics will change within the network because of random interaction with other traffic streams. QoS estimation usually relies on an assumed source model and calculates QoS from a queuing analysis, approximation, or simulation. In any case, there may be inaccuracy between the estimated QoS and the actual resulting QoS, and this difference must be evaluated through QoS measurements by management cells. The measured differences may be useful to dynamically adjust the parameters of the CAC algorithm.

Fault Management

Alarm Surveillance

Management cells have already found various uses in fault management, for example, ATM-layer fault alarms. Fault detection may occur in the physical layer or ATM layer. The physical layer (e.g., SONET/SDH) will probably have fault monitoring that will detect the loss of physical-layer transmission (e.g., loss of signal, loss of frame, loss of pointer) and generate physical-layer alarms. If a fault is detected in the physical layer which is not automatically protected, the ATM layer will be notified. ATM-layer alarms will be generated to notify other nodes of a detected fault. The type of management cells for fault alarms are the OAM alarm indication signal (AIS) cell and remote defect indication (RDI) cell [1]. The AIS cell is sent to notify downstream nodes, and repeated at 1 cell/s as long as the fault persists. In response to AIS, the receiver generates RDI cells in the backward direction to notify upstream nodes at the rate of 1 cell/s as long as AIS persists. As shown in Fig. 5, both AIS and RDI cells contain fields to optionally indicate the type of fault and fault location.

In addition, the ATM layer continually monitors idle connections for signs of a fault. It may not be obvious whether a connection is idle due to a fault or simply a lack of data to transmit. Hence, a management cell should be sent periodically on an idle connection as a sign of connectivity. The type of management cell prescribed for this purpose is the *OAM continuity cell* [1]. They are transmitted after 1 s of idleness and repeated at a rate of 1 cell/s. Currently, there are no specific fields standardized for the continuity check cell. The absence of both user cells and continuity check cells indicates a failure with either the connection or the source.

Diagnosis and Fault Isolation

Management cells may also be sent on demand to verify the continuity of a path or diagnose the extent of a fault. The management cell useful for this type of testing is the *OAM loopback cell* [1]. A loopback cell is transmitted and returned in the backward direction at an indicated node. As shown in Fig. 6, the loopback cell contains these fields:

- Loopback indication, which is decremented after loopback has been done
- Correlation tag, to uniquely identify the looped-back cell
- Loopback location, where the loopback should be performed
- Source identifier, to enable the originator to recognize its own returned cells

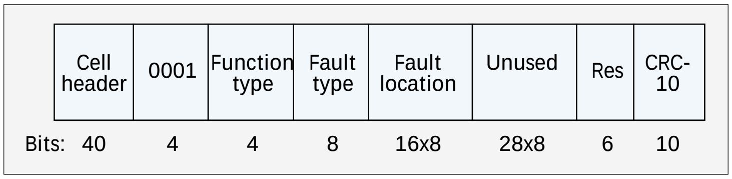

	Cell header	0001	Function type	Fault type	Fault location	Unused	Res	CRC-10
Bits:	40	4	4	8	16x8	28x8	6	10

■ Figure 5. *OAM AIS/RDI cells.*

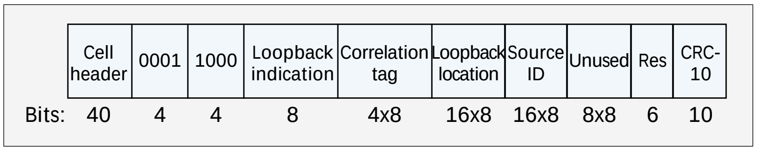

	Cell header	0001	1000	Loopback indication	Correlation tag	Loopback location	Source ID	Unused	Res	CRC-10
Bits:	40	4	4	8	4x8	16x8	16x8	8x8	6	10

■ Figure 6. *OAM loopback cell.*

As an additional aid in diagnosis, management cells can monitor the QoS per connection as described earlier. This may be useful to detect a soft failure which exhibits performance degradation over an extended time before eventually becoming a hard failure.

After diagnosis of a fault, the network can be reconfigured to isolate the trouble. Traditionally, reconfiguration has been managed centrally. After testing the faulty equipment, it is repaired and reconfigured back into service. However, the need for fast and automatic fault recovery is becoming increasingly critical for ATM networks. The response time to restore service must be minimal because high-speed ATM networks will carry large amounts of user data. Automation is also needed because ATM networks will be more complicated than telephone equipment to understand and diagnose. Trained experts will be valuable, and it will be costly to dispatch experts and complex testing equipment around the network. The need for fast and automated fault recovery has motivated research in distributed control, such as self-healing VP networks.

Self-Healing VP Restoration

Self-healing networks are capable of restoring failed paths using distributed control mechanisms. For instance, if a link fails, a self-healing SONET ring or mesh network would detect this link failure and automatically reroute traffic to backup links without the assistance of a centralized network manager [27]. Likewise, a self-healing ATM network would automatically restore failed VPs to backup VPs using distributed control. Restoration in the ATM layer has an advantage over that in the physical layer: backup VPs can be established without the need to reserve capacity, unlike physical-layer backup paths. Self-healing at the VP level can be implemented in ATM rings [28] or mesh networks [29–35]. Path restoration in VP mesh networks is accomplished through the exchange of management cells.

In VP mesh networks, the backup VPs are either pre-established or dynamically searched. In the dynamic search approach, the downstream endpoint of the failed VP (called the *sender node*) broadcasts management cells called *restoration messages* to seek an alternative path to the upstream endpoint (called the *chooser node*). A restoration message might contain these fields [35]:

- Message type

- Sender node ID
- Chooser node ID
- Failed link information (link ID, link capacity, failure time)
- Hop count limit
- Information from intermediate nodes (node ID, available VPs, available VP capacity)

When an intermediate node receives a restoration message, it decrements the hop count and records information about the available VP. The message is broadcast toward the chooser node. If multiple alternative paths are found, the chooser node selects one according to a given criterion such as shortest delay. The chosen backup VP is confirmed with a management cell called a *confirmation message* to the sender node, while *release messages* are sent along the other alternative paths. This dynamic search process can be accelerated by flooding restoration messages from both endpoints of the failed VP which meet at intermediate nodes [30].

Although robust, this dynamic search approach involves a large number of node-to-node management cells in both directions between the VP endpoints. By preassigning backup VPs with zero bandwidth, the number of messages and the restoration time can be reduced drastically [31–33]. The sender node transmits a single restoration message along the preassigned backup VP to the chooser node. Along the way, the message sets up the required bandwidth at each node (backup VPs must be chosen to ensure that sufficient restoration bandwidth will be available on demand). After receiving the management cell, the chooser node switches the traffic from the failed VP to the backup VP. If multiple failures make a pre-established backup VP unavailable, additional steps are needed to find an available backup VP dynamically.

We believe that management cells can be a basic mechanism underlying a multitude of management functions, and it is instrumental to study them as a single general class of special cells.

Network Administration and Management

Time Synchronization

The capability to synchronize the clocks at network nodes to a common reference time of day is valuable for various purposes in network administration. For example, time stamps can then be used for accurate performance monitoring; multiple alarms generated by a single fault can be correlated by examining their times; time stamps can be used to protect against replay security attacks; and the update of encryption keys can be coordinated (discussed below). A commonly proposed approach to synchronization to coordinated universal time (UTC) involves reception of radio signals from the 24-satellite Global Positioning System (GPS) which is accurate to about 10–100 ns [36, 37]. Not every node needs to be synchronized directly to GPS; selected nodes designated as primary time servers can be synchronized directly and distribute timing information to the remaining nodes through physical transmission signals or timestamped packets. The Internet network time protocol (NTP) follows the latter approach of exchanging timestamped UDP (user datagram protocol) packets between primary time servers and other time servers [38].

Similarly, time-of-day information can be disseminated in an ATM network through the bidirectional exchange of timestamped node-to-node management cells. Suppose that primary server A is being referenced by secondary server B. In the general procedure, server B will send a management cell to A containing the origination time $T1$. Server A will write the time of receipt $T2$ into the cell, and the time of transmission $T3$ back to B. Upon return, server B will note the time of receipt $T4$, and calculate the round-trip time $D = (T4 - T1) - (T3 - T2)$ and clock offset of A relative to B as $d = (T3 - T4 + T2 - T1)/2$. Since $T1$ and $T4$ are read from B's clock, and $T2$ and $T3$ from A's clock, the round-trip time is unambiguous. The clock offset assumes that the management cell takes an equal time to travel from A to B as from B to A. If the clocks were perfectly synchronized, the clock offset would turn out to be zero. The true offset has been shown to be bounded within the interval $d \pm D/2$, that is, it must be in an interval equal to the measured round-trip delay centered around the measured offset [38]. As expected, the procedure is more accurate when A and B are closer to each other. Also, multiple sets of measurements can perhaps improve the accuracy [11].

Security

Like other networks, ATM networks are vulnerable to security attacks [39]. Some well-known examples of security threats include eavesdropping (when data cells are captured and read), alteration (when data cells are intercepted and changed), masquerading (when a source sends cells pretending to be another source), replay (capture and later retransmission of cells), service interruption (by sending signaling messages affecting another connection), and unauthorized access to network systems or resources. ATM networks will need a variety of security mechanisms which are beginning to be studied [40]. Basic security services include:

- Encryption of the user cell payload (perhaps including a digital signature for authentication)
- Authentication of signaling parties during connection establishment and termination
- Encryption of confidential information in signaling messages
- Agreement on initial encryption keys during connection setup

Some additional security services can be facilitated by management cells. For instance, management cells may serve to exchange in-band security information during a connection (because signaling is currently not allowed during a connection) [41]. In-band security messaging would be useful for key exchange without interrupting a connection. Encryption keys will have limited lifetimes to reduce the chance of a key being discovered. Hence, keys must be updated periodically, and the updates must be synchronized to prevent parties from using different keys at the same time. Management cells can carry key updates and help to synchronize the key management.

Management cells can also serve as the vehicles to distribute certificates. Key certification protects against attacks by an unauthorized party to change the public key of another party. Usually a trusted third party, called a *certificate authority*, is employed to store and distribute public keys in the form of certificates. A certificate contains the owner name, its public key, optional additional information (such as lifetime), and the digital signature of the certificate authority. Since the certificate authority's public key is known, any party can verify the genuineness of a certificate and recover the public key in the certificate. However, certificates cannot be forged without knowledge of the certificate authority's private key. Once signed, certificates may be stored anywhere and transmitted over unsecure channels via management cells.

Finally, management cells can help to detect and mitigate the effect of bit errors or lost cells on the decryption process. It is clearly undesirable to allow a bit error or lost cell to disrupt the proper operation of the decryptor for a long time. Chaining is a method to limit the effects to blocks of cells or cipher chains [39]. Management cells can be inserted between cipher chains to mark the boundaries between them. Separation into cipher chains allows the decryption process to reinitialize and resume normal operation when cells within a cipher chain are lost.

Customer Network Management

Customer network management is a service provided by the network to the users, to allow a user's network management system to interact with a management agent in the network. Through the exchange of user-to-network and network-to-user management cells, users can manage many aspects of their ATM services. The management cells may have specific formats and functions, or simply encapsulate messages from a higher application-layer management protocol such as SNMP (Simple Network Management Protocol).

As an example, users might inquire about the cell loss ratio or cell delays on their connections for the purposes of verifying their QoS or comparing the QoS on different connections. The network agent could respond with QoS measurements made with management cells in procedures described earlier. As another example, users might inquire about their current level of resource usage, especially if billing for services is usage-based. An end-to-end management cell might pick up usage measurements made at the network ingress and network egress points (it may be useful to know how much traffic has both entered the network and been successfully delivered).

Currently, limited customer network management capabilities are provided by the interim local management interface (ILMI) [42]. The ILMI allows the exchange of SNMP messages encapsulated in ATM cells across the UNI. SNMP messages are encapsulated by AAL5 (ATM adaptation layer 5), and ILMI cells are identified by the header fields VPI = 0, VCI = 16. Users can obtain status and control information about the UNI by accessing a standard ATM UNI management information base.

Conclusions

We have argued for the use of management cells as a flexible tool to monitor and optimize ATM network operations. Traditionally, when a new control problem has been identified, a new type of special-purpose cell is developed for that particular need. The result is a disparate collection of specific-purpose cells. The opposite approach has been advocated here; we have proposed a broad class of management cells and sought all potential uses for network monitoring and control. This may bring a more cohesive and broad conceptual framework over various disparate aspects of network operations.

We believe management cells can be a basic mechanism underlying a multitude of management functions, and it is instrumental to study them as a single general class of special cells. A number of useful functions have been identified here but more uses can probably be found with further investigation.

However, these capabilities are gained at a cost in additional requirements on the network. Network nodes must be able to monitor and record their own behavior; and identify, process, modify, and generate management cells according to specific procedures. The increased complexity may be justifiable in the long term when ATM is deployed more widely and more sophisticated network management capabilities are needed.

References

[1] ITU-T Rec. I.610, "B-ISDN Operation and Maintenance Principles and Functions," Geneva, Switzerland, July 1995.

[2] ITU-T Rec. I.371, "Traffic Control and Congestion Control in B-ISDN," Perth, Australia, Nov. 1995.

[3] ATM Forum, "Traffic Management Specification Version 4.0," ATM Forum/95-0013R10, Feb. 1996.

[4] ITU-T Draft Rec. 0.191, "Equipment to Assess ATM Layer Cell Transfer Performance," Mar. 1995.

[5] N. Sato, *et al.*, "In-Service Monitoring Methods — Better Ways to Assure Service Quality of Digital Transmission," *IEEE JSAC*, vol. 12, Feb. 1994, pp. 355–60.

[6] ANSI, "American National Standard for Telecommunications — Digital Hierarchy Optical Rates and Formats Specifications," T1.105-1988, 1988.

[7] ITU-T Rec. G.707, "Synchronous Digital Hierarchy Bit Rates," Melbourne, Australia, Nov. 14–25, 1988.

[8] S. Farkouh, "Managing ATM-Based Broadband Networks," *IEEE Commun. Mag.*, vol. 31, May 1993, pp. 82–86.

[9] H. Murakami *et al.*, "Monitoring Method for Cell Transfer Performance in ATM Networks," *NTT Rev.*, vol. 4, July 1992, pp. 38–44.

[10] T. Chen and S. Liu, "Management and Control Functions in ATM Switching Systems," *IEEE Network*, vol. 8, July 1994, pp. 27–40.

[11] C. Roppel, "Estimating Cell Transfer Delay in ATM Networks Using In-Service Monitoring Methods," *Proc. Globecom '95*, pp. 904–8.

[12] T. Chen and S. Liu, "Method and Apparatus for Performance Monitoring in Electronic Communications Networks," U.S. patent appl. serial no. 08/625,102, Apr. 1, 1996.

[13] N. Duffield *et al.*, "Entropy of ATM Traffic Streams: A Tool for Estimating QoS Parameters," *IEEE JSAC*, vol. 13, Aug. 1995, pp. 981–90.

[14] C. Courcoubetis *et al.*, "Admission Control and Routing in ATM Networks Using Inferences from Measured Buffer Occupancy," *IEEE Trans. Commun.*, vol. 43, Feb. 1995, pp. 1778–84.

[15] H. Zhu and V. Frost, "In-Service Monitoring for Cell Loss Quality of Service Violation in ATM Networks," *IEEE/ACM Trans. Networking*, vol. 4, Apr. 1996, pp. 240–48.

[16] D. Gaiti and G. Pujolle, "Performance Management Issues in ATM Networks: Traffic and Congestion Control," *IEEE/ACM Trans. Networking*, vol. 4, Apr. 1996, pp. 249–57.

[17] T. Chen *et al.*, "The Available Bit Rate Service for Data in ATM Networks," *IEEE Commun. Mag.*, vol. 34, May 1996, pp. 56–71.

[18] S. Ohta, "Dynamic Bandwidth Control of the Virtual Path in an Asynchronous Transfer Mode Network," *IEEE Trans. Commun.*, vol. 40, July 1992, pp. 1239–47.

[19] R. Kawamura *et al.*, "Fast VP-Bandwidth Management with Distributed Control in ATM Networks," *IEICE Trans. Commun.*, vol. E77-B, Jan. 1994, pp. 5–14.

[20] ATM Forum, "Private Network-Network Interface Specification Version 1.0," ATM Forum/af-pnni-0055.000, Mar. 1996.

[21] S. Bahk, M. El Zarki, "Congestion Control Based Dynamic Routing in ATM Networks," *Comp. Commun.*, vol. 17, Dec. 1994, pp. 826–35.

[22] W. Lee, *et al.*, "Routing Subject to Quality of Service Constraints in Integrated Communication Networks," *IEEE Network*, vol. 9, July 1995, pp. 46–55.

[23] T. Oser *et al.*, "Routing with Admission Control in ATM Networks," *J. Network and Sys. Mgmt.*, vol. 3, 1995, pp. 151–71.

[24] S. Plotkin, "Competitive Routing of Virtual Circuits in ATM Networks," *IEEE JSAC*, vol. 13, Aug. 1995, pp. 1128–36.

[25] H. Yokoi *et al.*, "Performance Evaluation of Routing Schemes in B-ISDN," *IEICE Trans. on Commun.*, vol. E78-B, Apr. 1995, pp. 514–22.

[26] N-F. Huang *et al.*, "Some routing problems in broadband ISDN," *Comp. Networks and ISDN Sys.*, vol. 27, Oct. 1994, pp. 29–43.

[27] S. Hasegawa *et al.*, "Control Algorithms of SONET Integrated Self-Healing Networks," *IEEE JSAC*, vol. 12, Jan. 1994, pp. 110–19.

[28] Y. Kajiyama *et al.*, "An ATM VP-Based Self-Healing Ring," *IEEE JSAC*, vol. 12, Jan. 1994, pp. 171–77.

[29] J. Anderson *et al.*, "Fast Restoration of ATM Networks," *IEEE JSAC*, vol. 12, Jan. 1994, pp. 128–38.

[30] H. Fujii and N. Yoshikai, "Restoration Message Transfer Mechanism and Restoration Characteristics of Double-Search Self-Healing ATM Network," *IEEE JSAC*, vol. 12, Jan. 1994, pp. 149–57.

[31] R. Kawamura *et al.*, "Self-Healing ATM Networks Based on Virtual Path Concept," *IEEE JSAC*, vol. 12, Jan. 1994, pp. 120–27.

[32] R. Kawamura and I. Tokizawa, "Self-Healing Virtual Path Architecture in ATM Networks," *IEEE Commun. Mag.*, vol. 33, Sept. 1995, pp. 72–79.

[33] R. Kawamura *et al.*, "Implementation of Self-Healing Function in ATM Networks," *J. Network and Sys. Mgmt.*, vol. 3, 1995, pp. 243–63.

[34] K. Murakami and H. Kim, "Virtual Path Routing for Survivable ATM Networks," *IEEE/ACM Trans. Networking*, vol. 4, Feb. 1996, pp. 22–39.

[35] N. Yoshikai and T-H. Wu, "Control Protocol and Its Performance Analysis for Distributed ATM Virtual Path Self-Healing Network," *IEEE JSAC*, vol. 12, Aug. 1994, pp. 1020–30.

[36] W. Klepczynski, "Modern Navigation Systems and Their Relation to Timekeeping," *Proc. IEEE*, vol. 71, Oct. 1983, pp. 1193–98.

[37] Bellcore, "Network Element Time-of-Day Transfer, Coordination, and Synchronization Generic Requirements," GR-2861-CORE, issue 1, Sept. 1994.

[38] D. Mills, "Internet Time Synchronization: The Network Time Protocol," *IEEE Trans. Commun.*, vol. 39, Oct. 1991, pp. 1482–93.

[39] D. Stevenson *et al.*, "Secure communications in ATM networks," *Commun. of ACM*, vol. 28, Feb. 1995, pp. 45–52.

[40] ATM Forum, "Phase I ATM Security Specification (Draft)," ATM Forum/95-1473R2, Apr. 1996.

[41] M. Peyravian and T. Tarman, "In-Band Security Messaging within the Data Channel," ATM Forum/96-0381, Apr. 1996.

[42] ATM Forum, "Interim Local Management Interface (ILMI) Specification Version 4.0," ATM Forum/95-0417R7, Apr. 1996.

Biographies

THOMAS M. CHEN [M] received the Ph.D. degree in electrical engineering from the University of California at Berkeley, and M.S. and B.S. degrees from Massachusetts Institute of Technology. He joined GTE Laboratories, Inc. in 1989 where he is currently a senior member of the technical staff involved in research in ATM traffic control and performance management. Dr. Chen is a technical editor of *IEEE Communications Magazine* and editor of the electronic journal *IEEE Communications Surveys*. He is the co-author of *ATM Switching Systems* (Artech House, 1995). E-mail: tchen@gte.com.

STEPHEN S. LIU [SM] received the B.S. degree from National Cheng-Kung University in Taiwan, and the M.S. and Ph.D. degrees from Georgia Institute of Technology in Atlanta, Georgia, all in electrical engineering. He joined GTE Laboratories, Inc. in 1981, where he is currently a principal member of technical staff. His research interests include wide-area ATM network traffic control and multimedia service provisions over ATM. He co-authored a book entitled *ATM Switching Systems*. E-mail: sliu@gte.com.

DAVID C. WANG [SM'91] received the B.S. from National Cheng-Kung University at Taiwan, and the M.S. and Ph.D. from Carnegie-Mellon University at Pittsburgh, all in electrical engineering. Currently, Dr. Wang is an architecture scientist at GTE Telephone Operations in Irving, Texas. He is responsible for the specification and design of broadband network architectures. He also serves as an adjunct professor at University of Texas at Arlington, where he teaches a course in advanced data networks. Prior to joining GTE, Dr. Wang was a member of technical staff at AT&T-Bell Labs for 10 years. He was involved in the performance specification, testing, and modeling of several networking technologies. E-mail: david.wang@telops.gte.com.

VIJAY K. SAMALAM received the M.S. degree from the Indian Institute of Technology, Kanpur, India, and the Ph.D. from the State University of New York at Stony Brook. He joined GTE Laboratories in 1984, where he worked on semiconductor device physics. Since then, he has worked on integrated circuit cooling, very large scale integration (VLSI) complexity theory, and neural networks, and was the principal firmware designer for GTE's project on gigabit networks and the broadband customer service module installed in the video service testbed in Cerritos, California. He is currently a principal member of technical staff and project leader on research into network performance for high-speed ATM networks. E-mail: vsamalam@gte.com.

MICHAEL J. PROCANIK graduated valedictorian from Devry Technical Institute, Woodbridge, NJ, in 1985. He is currently completing the B.S.E.E. degree at Northeastern University, Boston, Massachusetts. He joined AT&T Bell Laboratories, Murray Hill, NJ, in 1985, where his work focused on characterizing submicron high-speed digital NMOS integrated circuits. These ICs included a mux/demux chip set operating between 2 and 3 GHz. The devices were intended for AT&T's 1.7 Gb/s lightwave system. In 1988, he joined GTE Laboratories, Waltham, Massachusetts, where he has been involved in various optoelectronic SONET/ATM data links and systems. His current interests are traffic control, congestion control, and performance monitoring for high-speed ATM networks. E-mail: mprocanik@gte.com.

DINYAR KAVOUSPOUR is a standardization manager in GTE Telops' Customer Products and Services department. His responsibilities include defining and establishing requirements for broadband services in addition to selecting and approving the platforms for deployments. His other work areas include network management systems, presently leading the design effort on GTE' Customer Network Management Gateway. Mr. Kavouspour represents GTE in ATM Forum's Testing and Residential Broadband Groups. He received a B.S.E.E. degree from University of Southern California in 1982 and an M.S. degree in telecommunication engineering and management from Southern Methodists University in 1988. E-mail: dk@gte.net.