Design and Implementation of a Multicast-Buffer ATM Switch

Jin Li and Chuan-lin Wu
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, Texas 78712
U.S.A

Abstract

In this paper, we propose a multicast-buffer ATM switch, where a dedicated buffer is allocated to store all multicast cells in an ATM switch. We describe the design of the dedicated multicast buffer in an output-buffer ATM switch and in a shared-buffer ATM switch. No additional copy circuit or complicated control circuit is needed to implement the multicast-buffer ATM switch. Our performance evaluation shows that the proposed switch has a low cell-loss ratio. We are implementing an eight-by-eight output buffer ATM switch with the multicast buffer. We also present a window control mechanism to further improve the utilization of the multicast buffer. Therefore, we believe the multicast-buffer ATM switch is a good candidate to support the multicast function for future broadband ISDN applications.

1. Introduction

Asynchronous Transfer Mode (ATM) is the standard interface proposed by CCITT to support broadband Integrated Service Digital Network (B-ISDN) applications [1]. One of the key technologies for an ATM network is the ATM switch. The multicast function is an important feature of an ATM switch. The function enables an ATM switch to transfer an incoming cell to several output ports simultaneously. Broadcasting is a special case of the multicast function where an incoming cell is sent to all output ports. The early applications of ATM will be ATM-oriented Local Area Network or multimedia services such as video conferencing or data broadcasting, all of which require handling not only bursty traffic but also multicast connections. Therefore, a good multicast scheme is crucial for an ATM switch that is expected to achieve a low cell-loss ratio with a simple control mechanism.

Figure 1 illustrates the general structure of an ATM switch. The switch has input interface modules (IIM), a switch unit, and output interface modules (OIM). An IIM supports the required preprocessing of all incoming cells. For example, each IIM aligns all incoming cells so that they arrive at the switch unit in a time-slotted manner. The IIM also supports the header transformation of each incoming cell. The switch unit, a core part of an ATM switch, transfers incoming ATM cells to the outputs destinations. An OIM supports the required post processing of all outgoing cells. In this paper, we focus on realizing the multicast function inside the switch unit. We propose a multicast-buffer ATM switch to support the multicast function efficiently. As illustrated in Figure 2, a multicast buffer is allocated in an ATM switch for storing all multicast cells. The proposed switch has a better buffer utilization since each multicast cell occupies only one buffer space in the switch. Therefore, the cell-loss ratio of the multicast-buffer ATM switch is reduced.

Various types of ATM switch architecture have been proposed [2]. Based on the buffering strategy, an ATM switch unit can be classified as an input-buffer ATM switch, an output-buffer ATM switch, or a shared-buffer ATM switch. Among these switches, an output-buffer switch and a shared-buffer switch have the best delay and throughput, and thus are more commonly used. In this paper, we describe the implementation of the multicast buffer to an output-buffer ATM switch and to a shared-buffer ATM switch. For an output-buffer ATM switch, we dedicate a shared First-In and First-Out (FIFO) buffer to all multicast cells. For a shared-buffer ATM switch, we dedicate a shared FIFO address queue for all multicast cells. A multicast cell has a higher priority than a point-to-point cell. The resulting new switch does not require complicated control circuitry. Since each multicast cell occupies only one buffer space, a multicast-buffer ATM switch achieves a better cell-loss performance under multicast traffic loads.

0-8186-7216-1/95 $04.00 © 1995 IEEE

The rest of this paper describes the new switch architecture. Section two gives an overview of the current multicast functions in an output-buffer ATM switch and a shared buffer ATM switch. Section three presents the design of the multicast buffer in an output-buffer ATM switch and in a shared-buffer ATM switch. Section four gives the computer simulation results of the cell-loss ratio performance of the multicast-buffer ATM switch. Section five briefly describes the implementation of the multicast buffer in an eight-by-eight 155 Mbs output-buffer ATM switch. Section six concludes the paper.

2. Previous work

As various types of ATM switch are proposed. Many multicast schemes are studied depending on the switch architecture. Before presenting our new multicast scheme, we first illustrate existing multicast schemes for an output-buffer ATM switch [3] and a shared-buffer ATM switch [4][5].

Figure 3 gives an example of a multicast function in a four-by-four output-buffer ATM switch, where a multicast cell is arriving at input port 1 and going out at output ports 1, 2, and 4. The multicast cell is multiplexed into a high-speed bus, and broadcast to all output ports. The cell only passes the address filters (AF) of its multicast destinations, and is stored in the dedicated buffer of output ports 1, 2 and 4. Although such multicast function in an output-buffer ATM switch offers the advantage of simple control, it results in a poor buffer utilization. A large amount of buffer is needed for each output port to prepare for the cell concentration at that output port.

A shared-buffer ATM switch achieves higher buffer utilization by sharing the memory among all output ports. Each output port maintains a FIFO queue to record the shared-buffer addresses of all cells going to that output. Every output cycle, each output port can fetch the corresponding cell correctly from the shared buffer based on the address information in the FIFO. With the same amount of buffers, the cell-loss ratio of a shared buffer switch is much smaller than that of an output-buffer switch.

Additional hardware is required to realize the multicast function in a shared-buffer ATM switch. Two multicast schemes exist for a shared-buffer ATM switch. Scheme one is the cell-copy method [6], where a copy network is added before the switch unit. A multicast cell is duplicated L times before being sent to the switch unit, where L is the number of output destinations of the multicast cell. The total size of the switch becomes very large due to the copy network. Figure 4 illustrates this scheme with a four-by-four ATM switch. A multicast cell arrives at input port 1 and exists at output ports 1, 2, and 4.

Scheme two is the address-copy method [7]. As shown in Figure 5, where a multicast cell arrives at input port 1 and exists at output ports 1, 2, and 4, the cell is stored in the shared buffer only once, but the location of the cell in the shared buffer is copied and queued into address queues for outputs 1, 2 and 4. As the length of each address queue may be different, the switch needs additional logic to detect the release time of each multicast cell. Inside the control logic, a multicast cell counter (MCC) is associated with each shared buffer space. Each MCC holds the number of destinations of an incoming cell. The value held by the MCC is decreased when the corresponding cell is read out from the shared memory by an output. When the value of the MCC reaches zero, the shared buffer address of this cell can be released. Although this solution realizes the multicast function within the shared-buffer switch without an additional copy network, it requires rather complicated control circuitry.

In summary, the current multicast schemes have drawbacks. To realize the multicast function, an output-buffer switch does not need extra hardware, but a large amount of buffer is required for achieving a low cell-loss ratio performance. On the other hand, a shared-buffer switch needs a smaller amount of buffer due to the sharing effect, but its hardware complexity becomes great in either the cell-copy circuit or the control circuit. Since the multicast function is essential for an ATM switch, a new scheme to implement the function in an output-buffer switch and in a shared-buffer ATM switch is presented in the following section.

3. Proposed multicast-buffer ATM switch

3.1 The multicast output-buffer ATM switch

The multicast output-buffer ATM switch is derived by injecting a multicast FIFO buffer into an output-buffer ATM switch. This multicast FIFO buffer is called a multicast shared buffer (MSB), while the other FIFO buffer in the original output-buffer ATM switch is called a dedicated buffer. An incoming multicast cell of the switch bypasses the dedicated FIFO buffer and is stored in the MSB directly, which is depicted in Figure 6. In a later time slot, the multicast cell is broadcast to all output port destinations. Buffer utilization increases because each multicast cell occupies only one-cell space inside the switch.

The control mechanism of the multicast scheme is described as follows. Each incoming cell has a physical address, a multicast cell indicator, and a multicast channel number (MCN) [8]. These items of information are obtained in the header translation process at the incoming interface module, which is shown in Figure 1. When the multicast cell indicator is disabled, the physical address is used to identify which output port the cell should go to. On the other hand, when the multicast cell indicator is activated, the physical address information is ignored. The obtained MCN is sent to the multicast control table (MCT), which stores a bit-map vector to indicate whether an output port should receive the cell or not. A multicast cell on the time-division bus is copied to the MSB. The multicast cell has a higher priority than a point-to-point cell. Therefore, when the MCT is nonempty, the header cell in the MSB is read out and sent to each output port. At the same time, the MCN of the header cell is also broadcast to the enable line of each selector at the output port. An output port receives the multicast cell if the corresponding enable bit of the output selector is 1; otherwise, it fetches the point-to-point cell from its own dedicated buffer. Finally, individual virtual-channel identifiers are obtained at outgoing interface modules through header re-translations using the MCN. This mechanism achieves a better buffer utilization without an extra cell-copying mechanism or a complicated control logic.

3.2 The multicast shared-buffer ATM switch

If we apply the same multicast scheme as mentioned for a shared-buffer ATM switch, a dedicated multicast buffer should be added inside the shared-buffer switch unit to store all multicast cells. However, due to the inherent sharing feature of the shared-buffer ATM switch, we can store each multicast cell into any empty space of the shared buffer. Only an additional shared FIFO multicast address queue (MAQ) is added inside the control unit to keep track of all multicast cells in the shared buffer. An incoming multicast cell is stored into the shared buffer, and its address is sent to MAQ in the control unit. The corresponding MCN is also sent to the read address by the input interface module. Therefore, the multicast shared-buffer ATM switch is derived by injecting a FIFO MAQ into a shared-buffer ATM switch. The FIFO MAQ implements a virtual multicast shared buffer. The basic idea, which is illustrated in Figure 7, is the same as the multicast scheme proposed for an output-buffer ATM switch.

For a shared-buffer ATM switch, as many as N outgoing cells can be fetched from the shared buffer during one time slot. Since multicast cells have higher priority than point-to-point cells, the control unit checks MAQ first during each time slot. If MAQ is nonempty, the control unit scans the MCN of the multicast cell to find out those output ports with bit information 0. A point-to-point cell can be fetched for each of these outputs along with the multicast cell. For example, when N equals 8, if the MCN is 10010111 (output ports 2, 3, and 5 have bit information 0), then the control unit fetches the multicast cell, as well as the point-to-point cell for output ports 2, 3, 5 during the same time slot. The multicast cell is sent to output ports 1, 4, 6, 7, 8 according to the MCN.

4. Cell-loss performance evaluation

For an ATM switch, a cell is simply lost upon buffer overflow. Cell loss is a main drawback of ATM switching system. Therefore, a switch architecture with low cell-loss ratio is desired. To provide a reliable communication service to users, the cell-loss ratio should be negligibly small, probably around 10^{-6} to 10^{-9} [1], where 10^{-9} is approximately the bit error rate for an optical fiber transmission line. The design of an ATM switch should allocate enough buffer to ensure such a required cell-loss ratio. In this section, we compare the cell-loss performance of the proposed multicast-buffer ATM switch with the current output-buffer ATM switch and shared-buffer ATM switch.

4.1 Multicast function in an Output-buffer ATM switch

The switch models are listed below:

(1) an output-buffer ATM switch with only dedicated buffers, OB-D;
(2) a multicast output-buffer ATM switch, OB-M.

The traffic conditions are given below:

(1) the average number of incoming multicast cells is S, and the average number of destinations of each multicast cell is D.
(2) each incoming cell has random destinations with uniform probability.
(3) the normalized throughput ρ of the switch is kept constant at 99%. That is, ρ_p, which is the point-to-point traffic load, varies with the multicast traffic load ρ_m.

$$\rho_m = \frac{S \times D}{N} \times \rho$$

$$\rho_p = \frac{N - S \times D}{N} \times \rho$$

$$\rho = \rho_m + \rho_p = 99\%$$

where N is the number of input/output ports.

For an N-by-N OB-M, to achieve a required cell-loss ratio with multicast traffic load ρ_m, the total amount of buffer being saved compared with the OB-D is

$$(TD\rho - TD(\rho - \rho_m)) * N - TS\rho_m$$

where $TD\rho$ is the required buffer size for each dedicated FIFO when no multicast traffic is download to the MSB, $TD(\rho - \rho_m)$ is the required buffer size for each dedicated FIFO when ρ_p is reduced from ρ to $(\rho - \rho_m)$, and $TS\rho_m$ is the required buffer capacity of the MSB with multicast traffic load ρ_m. Therefore, to evaluate the buffer being saved for the OB-M, we should evaluate the required buffer for the dedicated traffic and the multicast traffic respectively when achieving a desired cell-loss ratio under multicast traffic load ρ_m.

With the same buffer capacity, computer simulation shows that the cell-loss ratio of the OB-M reduces dramatically with ρ_m; while the cell-loss ratio for the OB-D does not vary with ρ_m at all. The cell-loss ratio performance corresponds to the behavior of a M/D/a model [9]. The impact of N on the cell-loss ratio is rather limited, where N is the number of input ports or output ports. Figure 8 plots cell-loss ratio of the dedicated traffic in the OB-M versus $TD\rho$ when ρ_p varies. We find that a large number of buffers can be saved even if a very small percentage of the multicast traffic is downloaded to the MSB. For example, if buffer capacity for each dedicated FIFO is 40, then the cell-loss ratio is higher than $10^{-0.5}$ when ρ_p is 99% (thus, ρ_m=0); whereas the cell-loss ratio is reduced to 10^{-6} when ρ_p is 84% (thus, ρ_m=15%). When ρ_m equals 0, the cell-loss ratio of the OB-M is equal to the cell-loss ratio of the OB-D.

As mentioned, as long as ρ remains the same, the cell-loss ratio of the OB-D does not vary with ρ_m. Therefore, we only need to evaluate the cell-loss ratio versus $TS\rho_m$ under different multicast traffic load ρ_m in the OB-M. Considering an ATM switch with eight input ports and eight output ports, S is set to 1, 2, and 4, D is set to 2, 4, and 8. The resulting multicast traffic load ρ_m is shown in Table 1. Under these traffic conditions, Figure 9 illustrates the cell-loss ratio of the multicast traffic versus the buffer capacity of the MSB. We can see that as long as the value of S is fixed, the cell-loss ratio for the multicast traffic does not vary when the destination number D changes. Therefore, the buffer utilization of the MSB becomes the most efficient for broadcasting, when D is equal to N.

	S=1			S=2		S=4
D	8	4	2	4	2	2
ρ_m	1	0.5	0.25	1	0.5	1

Table 1 Multicast Traffic with Various S and D

Comparing the simulation results in Figure 8 and Figure 9, we can see that a large amount of buffer can be reduced in the OB-M with multicast traffic. Since poor buffer utilization is a major drawback of an output-buffer ATM switch, this mechanism improves the buffer utilization under multicast traffic loads. Therefore, to achieve the same cell-loss ratio, fewer buffers are required for the OB-M than for the OB-D. Since each buffer space is 53 bytes for an ATM cell, and buffer memory consumes a dominant area of an ATM switch, the OB-M scheme is attractive for designing an output-buffer ATM switch with small chip size.

4.2 Multicast function in a shared-buffer ATM switch

For the shared-buffer ATM switch, we examine a switch with eight input ports and eight output ports. The switch models are listed below:

(1) a shared-buffer switch with cell-copy circuit, SB-C;
(2) a shared-buffer switch with address-copy circuit, SB-A;
(3) a multicast shared-buffer switch, which has the shared address FIFO buffer, SB-M.

The traffic load is as follows:

(1) each incoming cell has a constant number of destinations, D;
(2) each cell has random destinations with uniform probability;
(3) the normalized throughput of the switch is kept constant. Therefore, the number of incoming cells changes as D changes.

In the following simulation, the number of destinations is set to 1, 2, 4, and 8; the normalized throughput is kept at 90%; and 10% of the offered load is point-to-point traffic. We investigate the cell-loss performance by computer simulation. Figure 10 shows the cell-loss ratio performance versus buffer size with various destinations. The figure indicates that the cell-loss ratio of a shared-buffer switch with multicast address queue (SB-M) has the best performance. For example, when buffer size per port is 4 and the number of destinations is eight (D=8), SB-M achieves less than 10^{-7} of the cell-loss ratio, whereas a shared-buffer switch with cell-copy circuit (SB-C) has more than 10^{-1} cell-loss ratio, and a shared-buffer switch with address-

copy circuit (SB-A) has a cell-loss ratio of more than 10^{-5}.

In the SB-C, the cell-loss ratio increases as the number of destinations increases. On the other hand, in the SB-A and SB-M, the cell-loss ratio decreases as the number of destinations increases. This is because for the SB-A and SB-M, each multicast cell only has a single copy inside the shared-buffer, and the rate of incoming cell is in inverse proportion to the number of destinations. To understand why the SB-M has a better performance than the SB-A, we need to remember that in the SB-M, a multicast cell reaches all output port destinations when it is read out only once. On the other hand, in the SB-A, a multicast cell in the shared buffer can not be released until its multicast cell counter reaches zero. The shared buffer address for a multicast cell is released earlier in the SB-M than in the SB-A, and the released address is available for a new incoming cell. As a result, the cell-loss ratio of the SB-M is better than that of the SB-A even though each multicast cell occupies a single buffer space in both approaches.

5. Implementation

5.1 Hardware complexity of a multicast shared-buffer switch

As mentioned in section three, the implementation of a multicast shared-buffer ATM switch is straightforward by just adding a multicast address FIFO queue. Therefore, we only briefly discuss the hardware complexity of the multicast shared-buffer ATM switch in this subsection.

In the shared-buffer ATM switch with cell-copy circuit, a multicast cell occupies several shared buffer spaces. As a result, data rate should be faster in proportion to the number of copies, so that multicast cells can be copied and written in buffers at the same time slot. If the number of copies is large, the speed of the data processing unit becomes a bottleneck, especially in high-speed switch design.

On the other hand, in the shared-buffer ATM switch with address-copy circuit, a multicast cell occupies several address spaces. As a result, the speed of the control circuit must be faster in proportion to the number of copies, so that each multicast cell address can be copied and written into the address queue at the same time slot. If the number of copies is large, the speed of the control processing unit becomes a bottleneck. Furthermore, in the shared-buffer ATM switch with address-copy circuit, a $(\log_2 N+1)$-bit counter is added for every buffer space to detect the release time of each multicast cell. The total number of bits consumed by the multicast cell counter is $N \times B \times (\log_2 N+1)$, where N is number of output ports and B is the buffer size per output port. The complexity increases when N and B become large.

It is well known that the required memory-access and control-processing speed is a major bottleneck for a shared-buffer ATM switch [10]. For the multicast shared-buffer ATM switch, every multicast cell occupies only one buffer space and one address space. Therefore, it does not require the same high-speed data processing or the high-speed control processing as in the other two switches. Furthermore, the proposed switch is free from the cell-copy circuit or the address-copy circuit, and no multicast cell counter is associated with each shared buffer space. Performance evaluation shows that to achieve the same cell-loss ratio, the multicast share-buffer ATM switch requires the least amount of buffers. Therefore, the proposed shared-buffer ATM switch is the architecture with the smallest hardware size of the three.

5.2 Implementation of an eight-by-eight multicast output-buffer ATM switch

In this sub section, we describe the implementation of a multicast output-buffer ATM switch. Currently, we are implementing an eight-by-eight ATM switch with an input speed of 155 Mbs. Figure 11 shows the implementation that integrates the multicast buffer into an output-buffer switch by the bus structure. A 150 Mbs bus can be implemented as eight lines, each operating at 18.75 Mbs. This rate is compatible with the current CMOS technology. A bit-sliced structure has been adopted to reduce the internal circuit speed for memory. Information cells on incoming ports are converted into eight-bit parallel signals, and individual signals are connected to individual input ports of bit-sliced sub switches. Each sub switch has a sliced time-division multiple bus and a sliced FIFO memory unit. They operate in a synchronous way. Address filtering (AF) operations for each outgoing port in individual sub-switches are all the same. To simplify the hardware design, no AF function is implemented within each sub switch. The AF operations and the multicast control functions for plural sub-switches can be centralized in an address controller. All incoming cells are multiplexed on the high-speed bus, and their header information is sent to the address controller directly by IIMs. The address controller generates write signals to the corresponding FIFO correctly based on the multicast indicator bit and the physical address information.

Instead of designing a new multicast output-buffer ATM switch, an existing output-buffer ATM switch can also be easily changed into an multicast output-buffer

ATM switch. Figure 12 shows the implementation block diagram by simply adding a separate memory chip into an existing output-buffer ATM switch. The input interface distributes an incoming cell either to a dedicated buffer or to a multicast buffer, and the output interface fetches a cell from the correct buffer. When the multicast traffic load varies, the switch can be upgraded easily by replacing a different memory chip for the multicast buffer without modifying the existing output-buffer ATM switch. Therefore, the proposed multicast output-buffer ATM switch not only coexists with the current output-buffer ATM switch very well, but also has a good scalability.

5.3 Multicast output-buffer ATM switch with window control

In the multicast shared-buffer ATM switch, the FIFO MAQ implements a virtual multicast shared buffer. Depending on the traffic condition, the multicast shared-buffer switch has the advantage of dynamic partition of the multicast buffer and the dedicated buffer. Therefore, the multicast buffer has a higher utilization in the proposed shared-buffer ATM switch that in a proposed output-buffer ATM switch. A window control mechanism can be used to improve the utilization of the multicast buffer in an output-buffer ATM switch .

For the multicast-buffer ATM switch described in section 5.2, only one multicast cell is cleared from the MSB during each time slot (the transmission time for one ATM cell). For the multicast buffer, the worst case occurs when each multicast cell has only two destinations, and a total of $\frac{N}{2}$ cells are directed to the MSB during each time slot, where N be the number of input or output ports. With the window control mechanism, several multicast cells in the MSB can be sent out during the same time slot as long as each output port is free from contention. If W is the window size ($W \leq \frac{N}{2}$), then up to W cells can be read out from the MSB and sent to all output ports in one time slot. A demultiplexer is employed at the read port of the MSB as shown in Figure 13. The MCN of each individual multicast cell is also sent to the demultiplexer, so that each multicast cell can be directed into the corresponding output port correctly.

The window control mechanism operates as follows. Let B be an N-bit signal of 0's or 1's. During each time slot, B is assigned the value of the first MCN (denoted as S1) in the multicast control table, therefore, B = S1. If the second cell in the head of the MSB is fetched, its MCN (denoted as S2) is compared with the value of B. As long as the destinations of these cells do not overlap, these cells can be cleared from the MSB during the same time slot. For example, suppose N equals 8, the MCN values of the first two cells in the MSB are listed as follows:

S1 = 1001000 S2 = 0010001

Since S1 AND S2 = 0, these cell go to different output ports. They can be read out at the same time slot, and the value of B is the result of bit-wise OR operation of S1 and S2:

B = S1 OR S2 = 1011001

The third MCN (denoted as S3) is compared with B again. This process continues until up to W cells in the head of the MSB are read out or an output contention occurs. The final bit value of B is sent to each output selector. This mechanism can be used to increase the utilization of the MSB if the average number of destinations of each multicast cell is low.

6. Conclusion

One of the key roles of an ATM switch is to support the multicast function with a low cell-loss ratio and a simple control method. We have proposed a new multicast-buffer ATM switch architecture. We describe a new design for an output-buffer ATM switch with the improved buffer utilization, and a new design for a shared-buffer ATM switch with less complexity of the control circuitry. Both of the switches can realize an efficient multicast scheme. We have discussed the cell-loss ratio of ATM switches under multicast traffic conditions. The evaluation by computer simulation shows that this approach has a better cell-loss ratio and requires less hardware of buffer size or control logic. We are developing an 155 Mbs eight-by-eight switch to implement the multicast scheme.

References

[1] CCITT, I-Series Recommendation on B-ISDN, Geneva, 1991.

[2] A. Pattavina, "Nonblocking architectures for ATM switching", IEEE Commun. Magazine, P.13, Feb. 1993.

[3] H. Suzuki, et al., "Output-buffer switch architecture for asynchronous transfer mode", ICC'89, P.99, June 1989.

[4] A. Thomas, J. P. Coudreuse and M. Servel: "Asynchronous time-division techniques: An experimental packet network integrating video communications," ISS'84, P.32, C.2 May, 1984.

[5] H. Kuwaharea, et al., "A shared buffer memory switch for ATM exchange,", ICC'89, P.118, June, 1989.

[6] T. Koinuma, et al, "An ATM switching system based on a distributed control architecture", ISS'90, Vol.5, P.21, May, 1990.

[7] S. Miura, et al, "Multiple shared buffer type ATM switch with multicast function", IEICE'90 Autumn National Convention Record, B-360 in Japanese, Sep.,1990.

[8] J. S. Turner, "Design of a broadcast packet switching network", IEEE Trans. Commun., Com-36, P.734-743. June, 1988.
[9] G. Fross and C. M. Harris, "Fundamentals of queuing theory", New York, Wiley, 1974.
[10] Martine de. Prycher, "Asynchronous Transfer Mode: Solution for Broadband ISDN," New York, 1993.

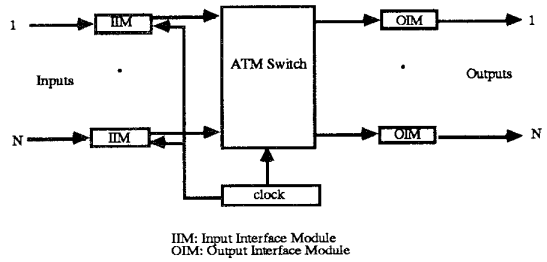

Figure 1 General structure of an ATM switch

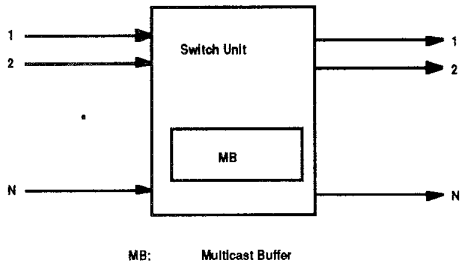

Figure 2 Proposed multicast-buffer ATM switch

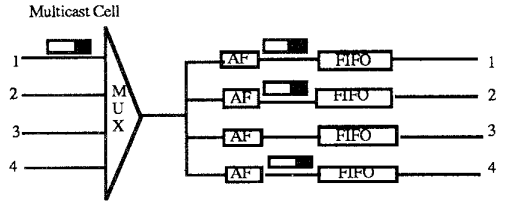

Figure 3 An existing output-buffer ATM switch

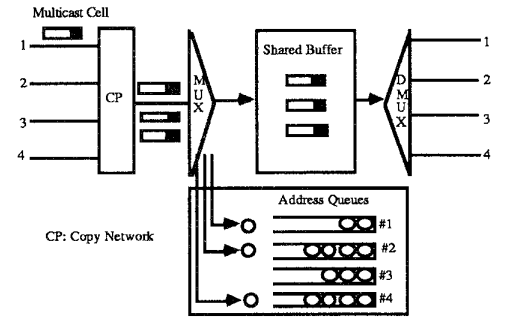

Figure 4 A existing shared-buffer ATM switch with cell-copy circuit

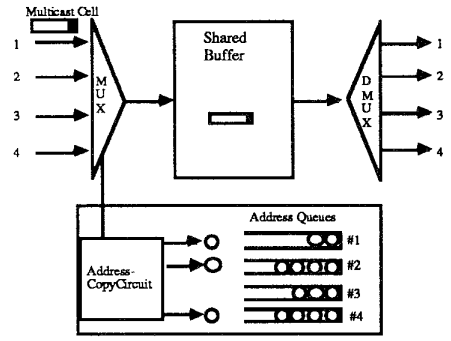

Figure 5 A shared-buffer ATM switch with address-copy circuit

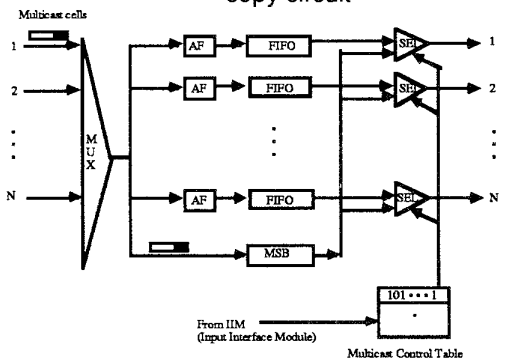

MSB: Multicast shared buffer

(a)

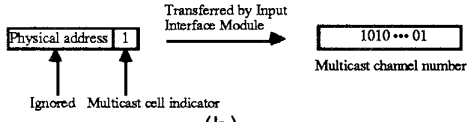

(b)

Figure 6 Proposed multicast scheme in an output-buffer ATM switch

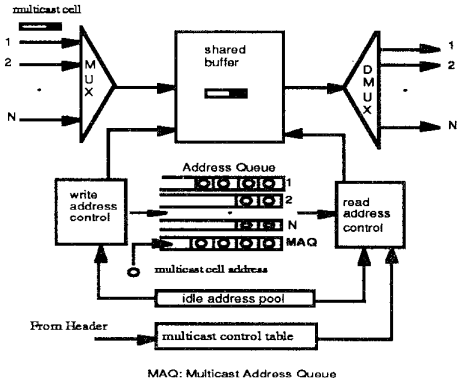

Figure 7 Proposed multicast scheme in a shared-buffer ATM switch

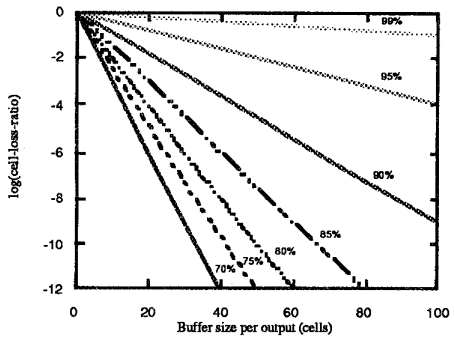

Figure 8 Cell-loss ratio of point-to-point cells in an OB-M switch

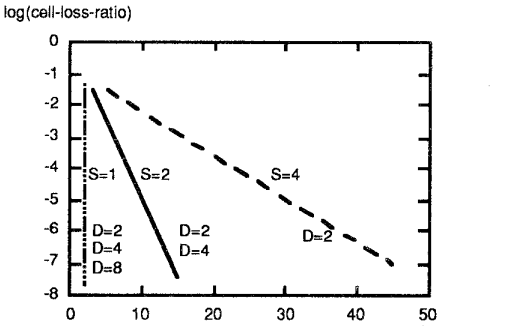

Figure 9 Cell-loss ratio versus buffer size for multicast cells in OB-M

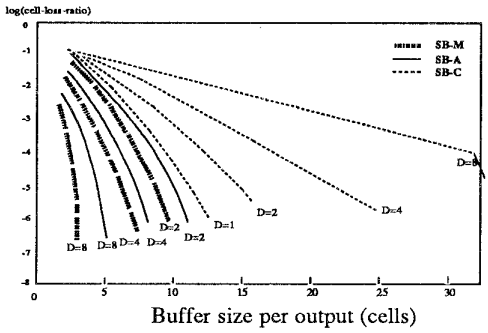

Figure 10 Cell-loss ratio of a shared-buffer ATM switch with different multicast scheme

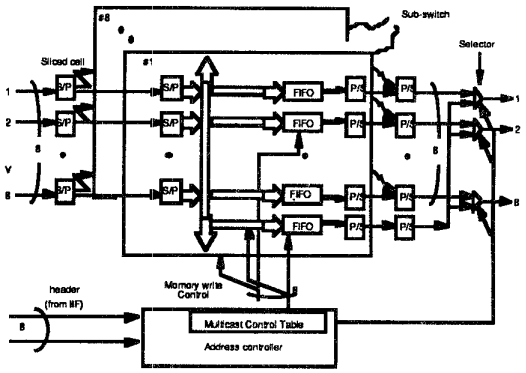

Figure 11 Implementation of an output-buffer ATM switch

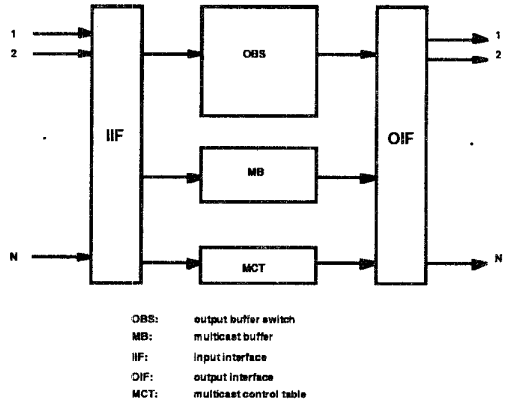

Figure 12 Add a separate memory chip to an existing output-buffer ATM switch

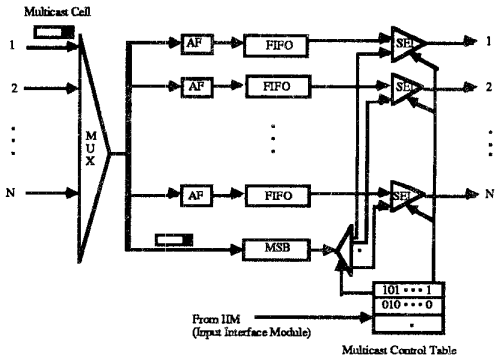

MSB: Multicast shared buffer

Figure 13 Multicast output-buffer ATM switch with window control mechanism