A Slotted Ring Copy Fabric for a Multicast Fast Packet Switch

Peter Newman and Matthew Doar

The University of Cambridge Computer Laboratory
New Museums Site, Pembroke Street
Cambridge CB2 3QG
U.K.

December 1989

ABSTRACT

Asynchronous transfer mode (ATM) and fast packet switching offer the ability to support many diverse forms of telecommunications services across a single integrated network. Some applications require the communications traffic from a single source to be received concurrently by many destinations. One approach to satisfy this requirement is to use a copy fabric to replicate the necessary number of copies of the incident traffic and then to route each copy to the required destination in a conventional point-to-point ATM fast packet switch. This paper presents the design and performance of a copy fabric based upon a slotted ring that permits a very simple implementation. It offers a performance which for many applications is comparable to that of much more complex designs.

1 Introduction

Asynchronous transfer mode (ATM) has been proposed as an integrated switching mechanism for use in the public broadband ISDN and within broadband private networks. Other areas of application include metropolitan area networks (MANs) and high speed local area networks (LANs). ATM is a transfer mode that uses fast packet switching with short fixed length packets to reduce the delay and the variance of delay for delay sensitive services such as voice and video [1]. These short packets are called cells and currently broadband ISDN proposes a cell length of 48 octets of information with 5 octets of header [2]. Many ATM fast packet switch designs only support unicast operation, i.e. point-to-point operation, where one incoming cell produces one outgoing cell routed to a single specific destination. There are a number of applications that require the switches of the network to

[1]To appear in Proc. Int. Switching Symposium, Stockholm, Sweden, May 1990.

be capable of multicast operation where a single incoming cell produces multiple outgoing cells. Each outgoing copy of a multicast cell contains the same information but is routed to a different destination. Such applications include video conferencing, audio conferencing, entertainment video distribution, the interconnection of local area networks and some aspects of distributed processing.

There are a number of possible approaches to the design of a multicast ATM fast packet switch but many of the existing designs lead to a very complex implementation. The approach proposed in this paper yields a very simple implementation that may easily be implemented in current gate array technology. It may be prefixed to any existing ATM switch design to offer multicast operation without requiring any modifications to the switch. The design was developed to serve both the private broadband networks market and within the high speed local area backbone network where switch designs that require vast areas of custom silicon are at present inappropriate.

2 Design of a Multicast Fast Packet Switch

The simplest means to achieve multicast operation is for the source to send multiple copies of each multicast cell to each destination, one after the other. This simple technique has several major drawbacks. The delay between the generation of the first and last copies of each multicast cell will be large. The source must waste valuable time individually transmitting many copies of the same cell. Also, network resources would be more efficiently utilised if copies of multicast cells were not generated until later nodes in the network wherever possible. A slightly more efficient approach uses a dedicated server at each switching node to make the required number of copies of every incident multicast cell and to route each copy individually to its required destination [3]. The copying and routing is still performed in series, however, thus a large delay may still exist between the generation of the first copy and the last. This approach is therefore not suitable for the multicast operation of real-time traffic such as voice or video and may not be fast enough for some distributed computing applications.

For such applications the operations of copying and routing must be performed in parallel, i.e. all of the copies of each incident multicast cell must be copied and routed concurrently. Also, many different multicast cells from separate switch ports should be capable of being handled at the same time. For ATM fast packet switch designs that use a single path switch fabric the operations of copying and routing the multicast cells may be implemented in the same switch fabric that handles the switching of the unicast traffic, but such switches may not be constructed with a total switch capacity above a few Gbits/sec in current technology [4, 5]. ATM fast packet switch designs that use a multi-stage switch fabric with output buffered switching elements may also implement multicast operation within the same switch fabric that performs the routing function but this makes the design of the switching element very complex [6, 7]. Also, non-buffered multi-stage switch fabrics exist that

Figure 1: General structure of a multicast ATM fast packet switch.

can handle both unicast and multicast traffic within a single switch fabric, e.g. the Richards network [8] and also Lea [9]. In general, however, these tend to be more suited to synchronous transfer mode (STM) with centralised control than for the distributed control that is essential for a high capacity ATM fast packet switch.

A more general solution is shown in fig. 1. A unicast (point-to-point) ATM fast packet switch is preceded by a copy fabric and a set of multicast controllers. The multicast controllers add a copy tag to each of the incoming cells and store them in an input buffer. The copy tag defines the number of copies of the cell that are required. The copy fabric then generates the required number of copies of each cell in parallel and each copy exits from the copy fabric on a separate output port at approximately the same time. All cells that exit from the copy fabric, whether originally unicast or multicast, may now be handled in the same manner by the unicast switch and routed to their respective destinations.

3 Banyan Copy Fabric Designs

Two designs of copy fabric have been proposed in the literature based upon a banyan interconnection network constructed from 2×2 broadcast switching elements (nodes). Each broadcast switching element has the ability to route an incident cell to either one of its two outputs, or to generate two copies of the cell, one on each output. Turner's switch [10, 11] uses buffered nodes in the copy fabric. A fanout field, appended to the front of each cell, specifies the required number of copies and each node in the interconnection network uses this field to determine whether to produce a single copy of each incident cell or two copies. The copy fabric is blocking in that incident cells may contend for the same resources (nodes and links) within the copy fabric therefore each node must have at least a single cell buffer on each input. Busyback signals are used between the nodes in each stage of the network to indicate a full cell buffer. Due to the buffering, each copy of the cell is not guaranteed to exit the copy fabric at the same time and there is some randomness built into the network as to the outputs of the network from which each copy will emerge.

Lee's copy fabric [12] is also based upon a banyan network but is non-blocking and also non-buffered. The fanout field at the head of each incident multicast cell is presented to a running adder network which computes a range of contiguous output ports from which the required number of copies will emerge. This address range is prefixed to the cell as a minimum and a maximum output address. All cells are applied to the copy fabric in alignment and from the minimum and maximum address fields each node may decide on which output to route the cell or whether to produce two copies. In this copy fabric all copies will exit the fabric at exactly the same time but the ports upon which they will exit are dependent upon the other traffic within the network. Furthermore, for Lee's copy fabric to work, active ports may not be interspersed with inactive ports, thus a concentration fabric is required prior to the copy fabric.

The above examples of copy fabrics do not offer simple implementation in current gate array technology. They are also difficult to partition into a single component that may be replicated in order to implement copy fabrics of any size. There are many applications that do not require the extremes of performance that a banyan copy fabric may be able to offer. Also, many applications cannot justify the investment in custom silicon required to implement such copy fabrics. The following is a very simple design of copy fabric that may easily be implemented in current gate array technology yet it offers a multicast capacity that for many applications may not be greatly inferior to the above examples.

4 A Slotted Ring Copy Fabric

4.1 Design

The proposed copy fabric is based upon a slotted ring and is illustrated in fig. 2. Each node on the ring supports a single input port and a single output port of the copy fabric. A multicast controller is connected to each input port of the copy fabric and the output ports of the copy fabric feed into the unicast switch as illustrated in fig. 1. Sufficient storage (typically one or two octets) is located in each node of the ring for an integer number of slots to circulate on the ring. Each slot is long enough to contain a complete cell. The slot structure is illustrated in fig. 3 and consists of the copy tag followed by the cell. The copy tag specifies the required number of copies of the cell. The slot structure may be maintained by a single timing source and distributed to each node of the ring as the copy fabric is likely to be implemented within a single cabinet.

ATM fast packet switching requires that all cells have a header that contains a label. The label defines which connection each cell belongs to. Each multicast controller performs a table look-up on the label of every incident cell and extracts from the table the copy tag associated with the connection to which that cell belongs. It prefixes the copy tag to the front of each incoming cell and stores the cells in the input buffer which is a first in first out (FIFO) queue.

Figure 2: The slotted ring copy fabric.

Figure 3: Slot structure.

When a node on the copy fabric has a cell waiting in its input buffer it waits for the beginning of the next slot to arrive on the ring. If that slot is free the cell is written into that slot else if the slot is full the node waits for the next free slot. A slot is considered free if the copy tag at the head of the slot is zero. As a full slot passes each node on the ring a copy of the cell is transmitted to the output port of that node and the copy tag at the head of the slot is decremented by one. Thus the required number of copies of the cell will exit from the copy fabric, the first copy from the output port of the node at which the cell entered and further copies from consecutive downstream ports. As soon as the required number of copies of the cell have been produced the copy tag becomes zero hence the slot is free for use by the next node on the ring that has a cell waiting. The copy tag may need to be removed from each cell on exit from the copy fabric to restore the cell to the format expected by the unicast switch.

Many proposals exist for the use of a slotted ring as a switching fabric, e.g. [13, 14]. In a slotted ring used as a switching fabric each cell must exit on a specific output port of the ring depending upon its required destination. For traffic with a random distribution of destinations each cell will on average have to travel half way round the ring to reach the destination it requires. Thus, on average, for a single slot ring two cells may be serviced for each complete rotation of the slot.

The use of a slotted ring as a copy fabric differs from its use as a switching fabric because the copy fabric places no restrictions upon which output ports of the copy fabric the cells must exit. The only constraint is that each copy of a cell exits the copy fabric on a separate output port. Each copy of a cell may therefore exit the ring on the nearest available output port. This occurs when each copy of a cell exits the ring on consecutive downstream output ports. Thus under saturation, during a single rotation of a single slot ring, every output port of the ring may produce a copy of a cell. This compares with a maximum of two ports producing output cells for a slotted ring used as a switching fabric. This will remain true whatever the distribution of copy requests in the copy tags. Thus a slotted ring used as a copy fabric in the above manner has a far greater capacity than a slotted ring used as a switching fabric.

Fig. 4 illustrates a single slot ring copy fabric under saturation for one complete rotation of the slot. Although every input is assumed to have a multicast cell waiting to access the ring only five have gained access to the ring during a single rotation of the slot. Every output, however, is busy producing one copy of a cell for every rotation of the slot therefore the copy fabric is operating at saturation. (The illustration of fig. 4 suggests that all cells exit the copy fabric with their headers aligned, whereas in reality there will be a small delay of a few bit times between the start of each cell due to the storage in each node of the ring.)

4.2 Label Allocation

In a unicast ATM fast packet switch the connection to which a cell belongs is identified by the label, or virtual channel identifier (VCI), in the cell header. In general each input port of a unicast switch allocates its own labels. A multicast cell will be replicated in the copy fabric and identical copies of the cell with identical labels will appear at a number of input ports to the unicast switch. Care must therefore be taken to ensure that there is no conflict between the labels of unicast and multicast traffic. A simple method of doing so is to allocate unicast labels from the bottom end of the address space and multicast labels from the top of the address space. A more efficient method is to translate the labels of multicast cells as they pass through the multicast controllers. Thus each port may select both its unicast and multicast labels independently of all other ports and the translation of multicast labels in the multicast controllers permits any conflict between labels to

Figure 4: An example of copy fabric operation at saturation.

be avoided. All copies of multicast cells emerge from the copy fabric on fixed output ports assigned at call setup. Thus the translated multicast labels need not be unique across all input ports of the unicast switch but only across those input ports that belong to the same multicast connection. So the same translated multicast label may be used for any number of multicast connections that do not have any input ports of the unicast switch in common.

In this simple slotted ring copy fabric the output port of the copy fabric at which each copy of a multicast cell will exit is fixed and known at the time the connection is established. This removes the requirement for label translation at the output of the copy fabric and permits the copy fabric to be used with any design of unicast switch without requiring that switch to be modified. In more complex schemes, e.g. Lee [12], any copy of a multicast cell may emerge at any output port of the copy fabric depending upon the current traffic load within the copy fabric. In this case every output port of the copy fabric requires a translation table that contains a label translation for every copy of each of the multicast labels. This rapidly leads to a very large translation table on every output of the copy fabric even for moderate sizes of switch.

4.3 Enhancements

Unfortunately, at high loads the copy fabric may become unfair in that some busy input ports may 'hog' the ring preventing downstream ports from gaining a fair share of access to the ring. If the bandwidth of the ring can be made sufficiently high this effect may not become important, otherwise a mechanism that ensures fairness may need to be implemented. One possible mechanism is to use a counter-rotating ring containing a single bit per ring node which is used to indicate which downstream nodes have cells to send. Upstream nodes may therefore take into account the traffic of downstream nodes before inserting their own cells onto the ring in a similar manner to that of the DQDB metropolitan area network proposal [15]. This mechanism might also be used to include priority information into the ring access algorithm to ensure that high priority traffic is served before traffic of lower priority. This is likely to be important for delay sensitive multicast traffic such as conference voice and video in the presence of delay insensitive traffic [16].

It is possible that in some applications most multicast cells will only require a few copies to be made. In this case, if the mean multicast traffic load is fairly low, a ring much smaller than the number of input ports may be used as a copy fabric. Each node of the copy fabric would handle the multicast traffic from several input ports and any multicast cells that require more copies than the size of the copy fabric could circulate around the ring more than once. The information in the copy tag attached to each cell would have to be used to distinguish between multiple copies of the same cell arriving at the same output port of the copy fabric.

In the above discussion it has been assumed that unicast cells (which do not require multiple copies) are handled in the same manner as multicast cells. The copy tag for a unicast cell is set to one and it passes through the copy fabric in exactly the same manner as multicast traffic. It emerges on the output port of the node at which it entered without further copies being made. It therefore suffers the same delay through the copy fabric as multicast traffic. If the delay performance for unicast traffic is to be optimised it would be better if unicast traffic were routed directly to the unicast switch without passing through the copy fabric. Fig. 5 illustrates one method of achieving this. A larger unicast switch is used and the outputs of the copy fabric feed into separate input ports of the unicast switch. The copy fabric need not be the same size as the number of input ports to the multicast controllers as several multicast controllers may share the same input port of the copy fabric. This design has the advantage that all multicast traffic is handled on separate ports of the unicast switch from the unicast traffic. This prevents the multicast traffic of one port from interfering with the unicast traffic of a downstream port which is a disadvantage of the former design. A concentrating interconnection network may be used as the switch fabric of the unicast switch if a large number of multicast ports are required.

Figure 5: Switch structure for enhanced performance.

The performance of the copy fabric at high loads will be impaired if the incident traffic is not evenly distributed across all inputs of the copy fabric. For a high performance copy fabric an even distribution of traffic across the copy fabric may be ensured by placing a distribution fabric before the copy fabric. Cells arriving at any of the input ports of the switch may be switched through the distribution fabric to any of the input ports of the copy fabric. All cells belonging to the same connection must follow the same path through the the distribution and copy fabrics to avoid out of sequence errors between cells on the same connection. The distribution of the traffic must therefore be performed on a per call basis at call setup and may be based upon an estimate of the current load on each section of the ring. One possibility for the implementation of the distribution fabric is to use another slotted ring since the load on the distribution fabric will be lower than that on the copy fabric.

5 Performance

5.1 The Simulation Model

The performance of the copy fabric was investigated using a simulation model. Copy fabrics of size 16 ports to 256 ports were modelled. It was assumed that each node on the ring possessed the required number of bits of storage in order that an integer number of slots fitted exactly onto the ring. The results are presented for the case in which exactly one slot fits onto the ring as the performance does not vary with the number of slots on the ring. The slot is assumed to be exactly one cell in length and all delay results are normalised to the length of a cell. A cell is assumed to be available for transmission as soon as it begins to arrive, i.e. the copy fabric operates in cut-through mode. Delays are measured from the time a cell begins to arrive until it commences transmission across the copy fabric. Each input port is assumed to carry slotted traffic and the incident load on any port gives the probability of any slot containing a valid cell. All input ports are assumed to be equally loaded and the phase relationship between the slot structure arriving at each input port is assumed to be random.

Each multicast cell will request multiple copies to be generated from a minimum of two to a maximum equal to the size of the copy fabric. The number of copies requested is called the fanout. Two fanout distributions were modelled: fixed and geometric. In the fixed fanout model all fanouts are the same fixed value. In the geometric case the fanout of each multicast cell is given by a truncated geometric distribution with parameter p similar to that used in [11]. The probability that the fanout is k is given by

$$\Pr\ (\text{fanout} = k) = \begin{cases} p(1-p)^{k-2} & 2 \leq k < N \\ (1-p)^{N-2} & k = N \end{cases}$$

where $0 \leq p \leq 1$ and N is the size of the copy fabric. The mean fanout obtained from this distribution is given by

$$E\ (\text{fanout}) = \begin{cases} (1/p)[1-(1-p)^{N-1}] + 1 & 0 < p \leq 1 \\ N & p = 0 \end{cases}$$

As each incident multicast cell gives rise to multiple outgoing cells the applied load is given by the product of the incident load and the mean fanout and is normalised to the throughput per port at saturation of the copy fabric. Each simulation run was allowed to reach stability and then measurements were taken for a total of 200,000 cells. This yielded measurements of mean delay with a standard deviation of about 6% of the mean for low loads and about 3% of the mean for medium to high loads.

5.2 Delay

The mean delay of a copy fabric of size 64 ports is given in fig. 6 for a geometric fanout distribution with mean fanouts from 2 to 32. The 99^{th} percentile of the delay distribution under the same conditions is given in fig. 7. The delay is at a maximum for a mean fanout of about half of the size of the ring (a fanout of 32 in this case) and beyond this fanout the mean delay decreases.

For a fixed fanout distribution the mean delay for fanouts up to 8 does not differ significantly from the geometric fanout distribution. For higher mean fanouts the fixed distribution gives a mean delay of up to about 30% greater at loads above 0.5 but a reduced 99^{th} percentile of delay.

Figure 6: Mean delay for a 64 port copy fabric.

Figure 7: 99th Percentile of delay distribution for a 64 port copy fabric.

It may be seen that for an applied load of 0.7 the mean delay is about 4 cell lengths and the 99^{th} percentile of delay about 20 cell lengths. This is comparable with the mean delay performance for the buffered banyan copy fabric reported in [11]. It also corresponds very closely to the delay performance of a unicast input buffered ATM fast packet switch operating at 80% of its throughput at saturation [16, 17]. Thus assuming a maximum load of 0.7 the copy fabric gives a very similar delay performance to a unicast input buffered switch operating at its maximum load if the incident traffic is randomly distributed across all inputs.

5.3 Size of Copy Fabric

Fig. 8 shows that the above measurements of mean delay do not vary greatly with the size of the copy fabric for small values of mean fanout. Fig. 9 illustrates how the mean delay varies with the mean fanout for various sizes of copy fabric and a geometric fanout distribution at an applied load of 0.7. It may be seen that the mean delay reaches a maximum when the mean fanout is about 30% to 50% of the size of the copy fabric.

5.4 Mixed Traffic

The simplest implementation occurs when unicast traffic passes through the copy fabric together with the multicast traffic. In this case the unicast traffic will experience the same delay as the multicast traffic. Fig. 10 shows how the delay through the copy fabric is reduced as the ratio of unicast traffic to multicast traffic is increased for a copy fabric of size 64 with a mean fanout of 8. The probability of any incident cell being a unicast cell is given by the unicast traffic ratio. Unicast traffic has a fanout of unity whereas multicast traffic has a fanout given by the geometric distri-

Figure 8: Mean delay for various sizes of copy fabric.

Figure 9: Variation of mean delay with mean fanout for various sizes of copy fabric at an applied load of 0.7.

bution thus the effective mean fanout of the traffic mix depends upon the unicast traffic ratio. The applied load is adjusted to take into account the effective mean fanout of the traffic mix.

For a unicast traffic ratio of zero all traffic is multicast and the curve is the same as that of fig. 6. For a unicast ratio of one all traffic is unicast and the delay is 0.5 which represents the average time any cell must wait for the arrival of the head of the slot on the ring given a random phase difference between traffic arrivals on each input port. It may be seen that the presence of a relatively small amount of multicast traffic in the traffic mix rapidly increases the mean delay to approach that of pure multicast traffic.

5.5 Fairness

The basic design of the copy fabric exhibits sensitivity to the distribution of the incident traffic across the input ports. For example, if two consecutive input ports are heavily loaded the downstream port will receive much poorer service than the upstream port. A simple solution to the problem is to make the bandwidth of the ring much greater than that of the input ports. This is not difficult to implement as the ring may be constructed up to 8 or 16 bits wide and clocked at up to 50 MHz in current CMOS gate array technology. This would offer a basic throughput at saturation approaching 1 Gbit/sec on every port.

An alternative approach is to implement a fairness algorithm. A counter-rotating ring with a single bit per input port may be used to inform upstream ports when a cell is waiting in an input port further downstream. Each input port may use this information in deciding whether to transmit a cell or to defer to a downstream port with a cell waiting. Investigations have shown that while not offering perfect

Figure 10: Variation of mean delay with traffic mix for a 64 port copy fabric and a mean fanout of 8.

fairness this approach can ensure that all ports receive a reasonable share of the bandwidth even under the worst case conditions.

6 Implementation

The slotted ring copy fabric is currently being implemented as part of an enhanced design of the Cambridge Fast Packet Switch [16, 17]. The objective is to produce a low cost ATM fast packet switch that can be used in the experimental investigation of ATM networking. From this work a design will be produced that is relevant to private wide area networks and also as a high speed local area backbone network.

In order to reduce the experimental hardware required the copy tag will be inserted as a field in the cell header. This permits the function of the multicast controller to be replaced by a simple FIFO provided that the label translation function in the unicast switch also translates the copy tag. The copy fabric has been designed using an 8 bit wide ring clocked at 25 MHz which gives a throughput at saturation of 200 Mbits/sec for each port. Thus, in this implementation, a 64 port copy fabric operating at a maximum load of 0.7 will offer a total traffic capacity in the region of 10 Gbits/sec. Each node of the copy fabric may be constructed in less than 500 gates using current CMOS gate array technology. This allows the multicast function to be integrated into the input port controllers of the unicast switch.

7 Conclusions

A copy fabric based upon a slotted ring has been proposed that may be prefixed to any unicast ATM fast packet switch to support multicast operation. The performance of the copy fabric has been investigated and shown to be comparable with much more complex designs provided that the incident traffic is evenly distributed across the input ports. The even distribution of traffic across the copy fabric may be ensured by the use of a distribution fabric before the copy fabric but in general this should only be necessary for switches designed for high multicast loads. For a copy fabric operating at a maximum load of 0.7 the mean delay is about 4 cell times and the 99^{th} percentile of delay about 20 cell times. This translates to a mean delay of about 12 μsecs and a maximum delay for 99% of the traffic of about 60 μsecs for standard ATM cells at 150 Mbits/sec. This is very similar to the performance that would be expected from an input buffered unicast ATM switch operating at its maximum load of 80% of its throughput at saturation. The performance does not vary with the number of slots on the ring and varies only slightly with the size of the ring.

The slotted ring copy fabric offers a very simple implementation in current gate array technology. This makes it most appropriate for integration with a simple design of unicast ATM switch such as the Cambridge Fast Packet Switch [16, 17]. In this case it may be implemented within the input port controllers of the switch prior to the label translation function. The resulting switch design is most appropriate to the emerging ATM private networks market and also within high speed local and metropolitan area networks.

References

[1] P Newman. Fast packet switching for broadband isdn. In *Proc. IEE Second Nat. Conf. on Telecommun.*, pages 391–396, York, U.K., Apr. 1989.

[2] S E Minzer. Broadband isdn and asynchronous transfer mode (atm). *IEEE Commun. Mag.*, 17–57, Sep. 1989.

[3] K Y Eng, M G Hluchyj, and Y S Yeh. Multicast and broadcast services in a knockout packet switch. In *Proc. IEEE Infocom*, pages 29–34, New Orleans, Mar. 1988.

[4] J P Coudreuse and M Servel. Prelude: an asynchronous time-division switched network. In *Proc. IEEE Int. Conf. Commun. (ICC '87)*, pages 769–773, Seattle, June 1987.

[5] I Cidon, I S Gopal, and H Heleis. Paris: an approach to integrated private networks. In *Proc. IEEE Int. Conf. Commun. (ICC '87)*, pages 764–768, Seattle, June 1987.

[6] H Suzuki et al. Output-buffer switch architecture for asynchronous transfer mode. In *Proc. IEEE Int. Conf. Commun. (ICC '89) vol. 1*, pages 99–103, Boston, June 1989.

[7] H Kuwahara et al. A shared buffer memory switch for an atm exchange. In *Proc. IEEE Int. Conf. Commun. (ICC '89) vol. 1*, pages 118–122, Boston, June 1989.

[8] G W Richards and F K Hwang. A two-stage rearrangeable broadcast packet switching network. *IEEE Trans. Commun.*, 33(10):1025–1035, Oct. 1985.

[9] C T Lea. A new broadcast switching network. *IEEE Trans. Commun.*, 36(10):1128–1137, Oct. 1988.

[10] J S Turner. Design of a broadcast packet switching network. *IEEE Trans. Commun.*, 36(6):734–743, June 1988.

[11] R G Bubenik and J S Turner. Performance of a broadcast packet switch. *IEEE Trans. Commun.*, 37(1):60–69, Jan. 1989.

[12] T T Lee. Nonblocking copy networks for multicast packet switching. *IEEE J. Select Areas in Commun.*, 6(9):1455–1467, Dec. 1988.

[13] A Hopper and R M Needham. The cambridge fast ring networking system. *IEEE Trans. Computers*, 37(10):1214–1223, Oct. 1988.

[14] T Takeuchi et al. Synchronous composite packet switching — a switching architecture for broadband isdn. *IEEE J. Select. Areas Commun.*, SAC-5(8):1365–1376, Oct. 1987.

[15] R M Newman, Z L Budrikis, and J L Hullett. The qpsx man. *IEEE Commun. Mag.*, 26(4):20–28, Apr. 1988.

[16] P Newman. A fast packet switch for the integrated services backbone network. *IEEE J. Select. Areas in Commun.*, 6(9):1468–1479, Dec. 1988.

[17] P Newman. *Fast Packet Switching for Integrated Services*. PhD thesis, The Computer Laboratory, University of Cambridge, Cambridge, U.K., Dec. 1988.

[18] M Doar and P Newman. A novel design of multicast fast packet switch. Submitted to IEEE Infocom, June 1990.