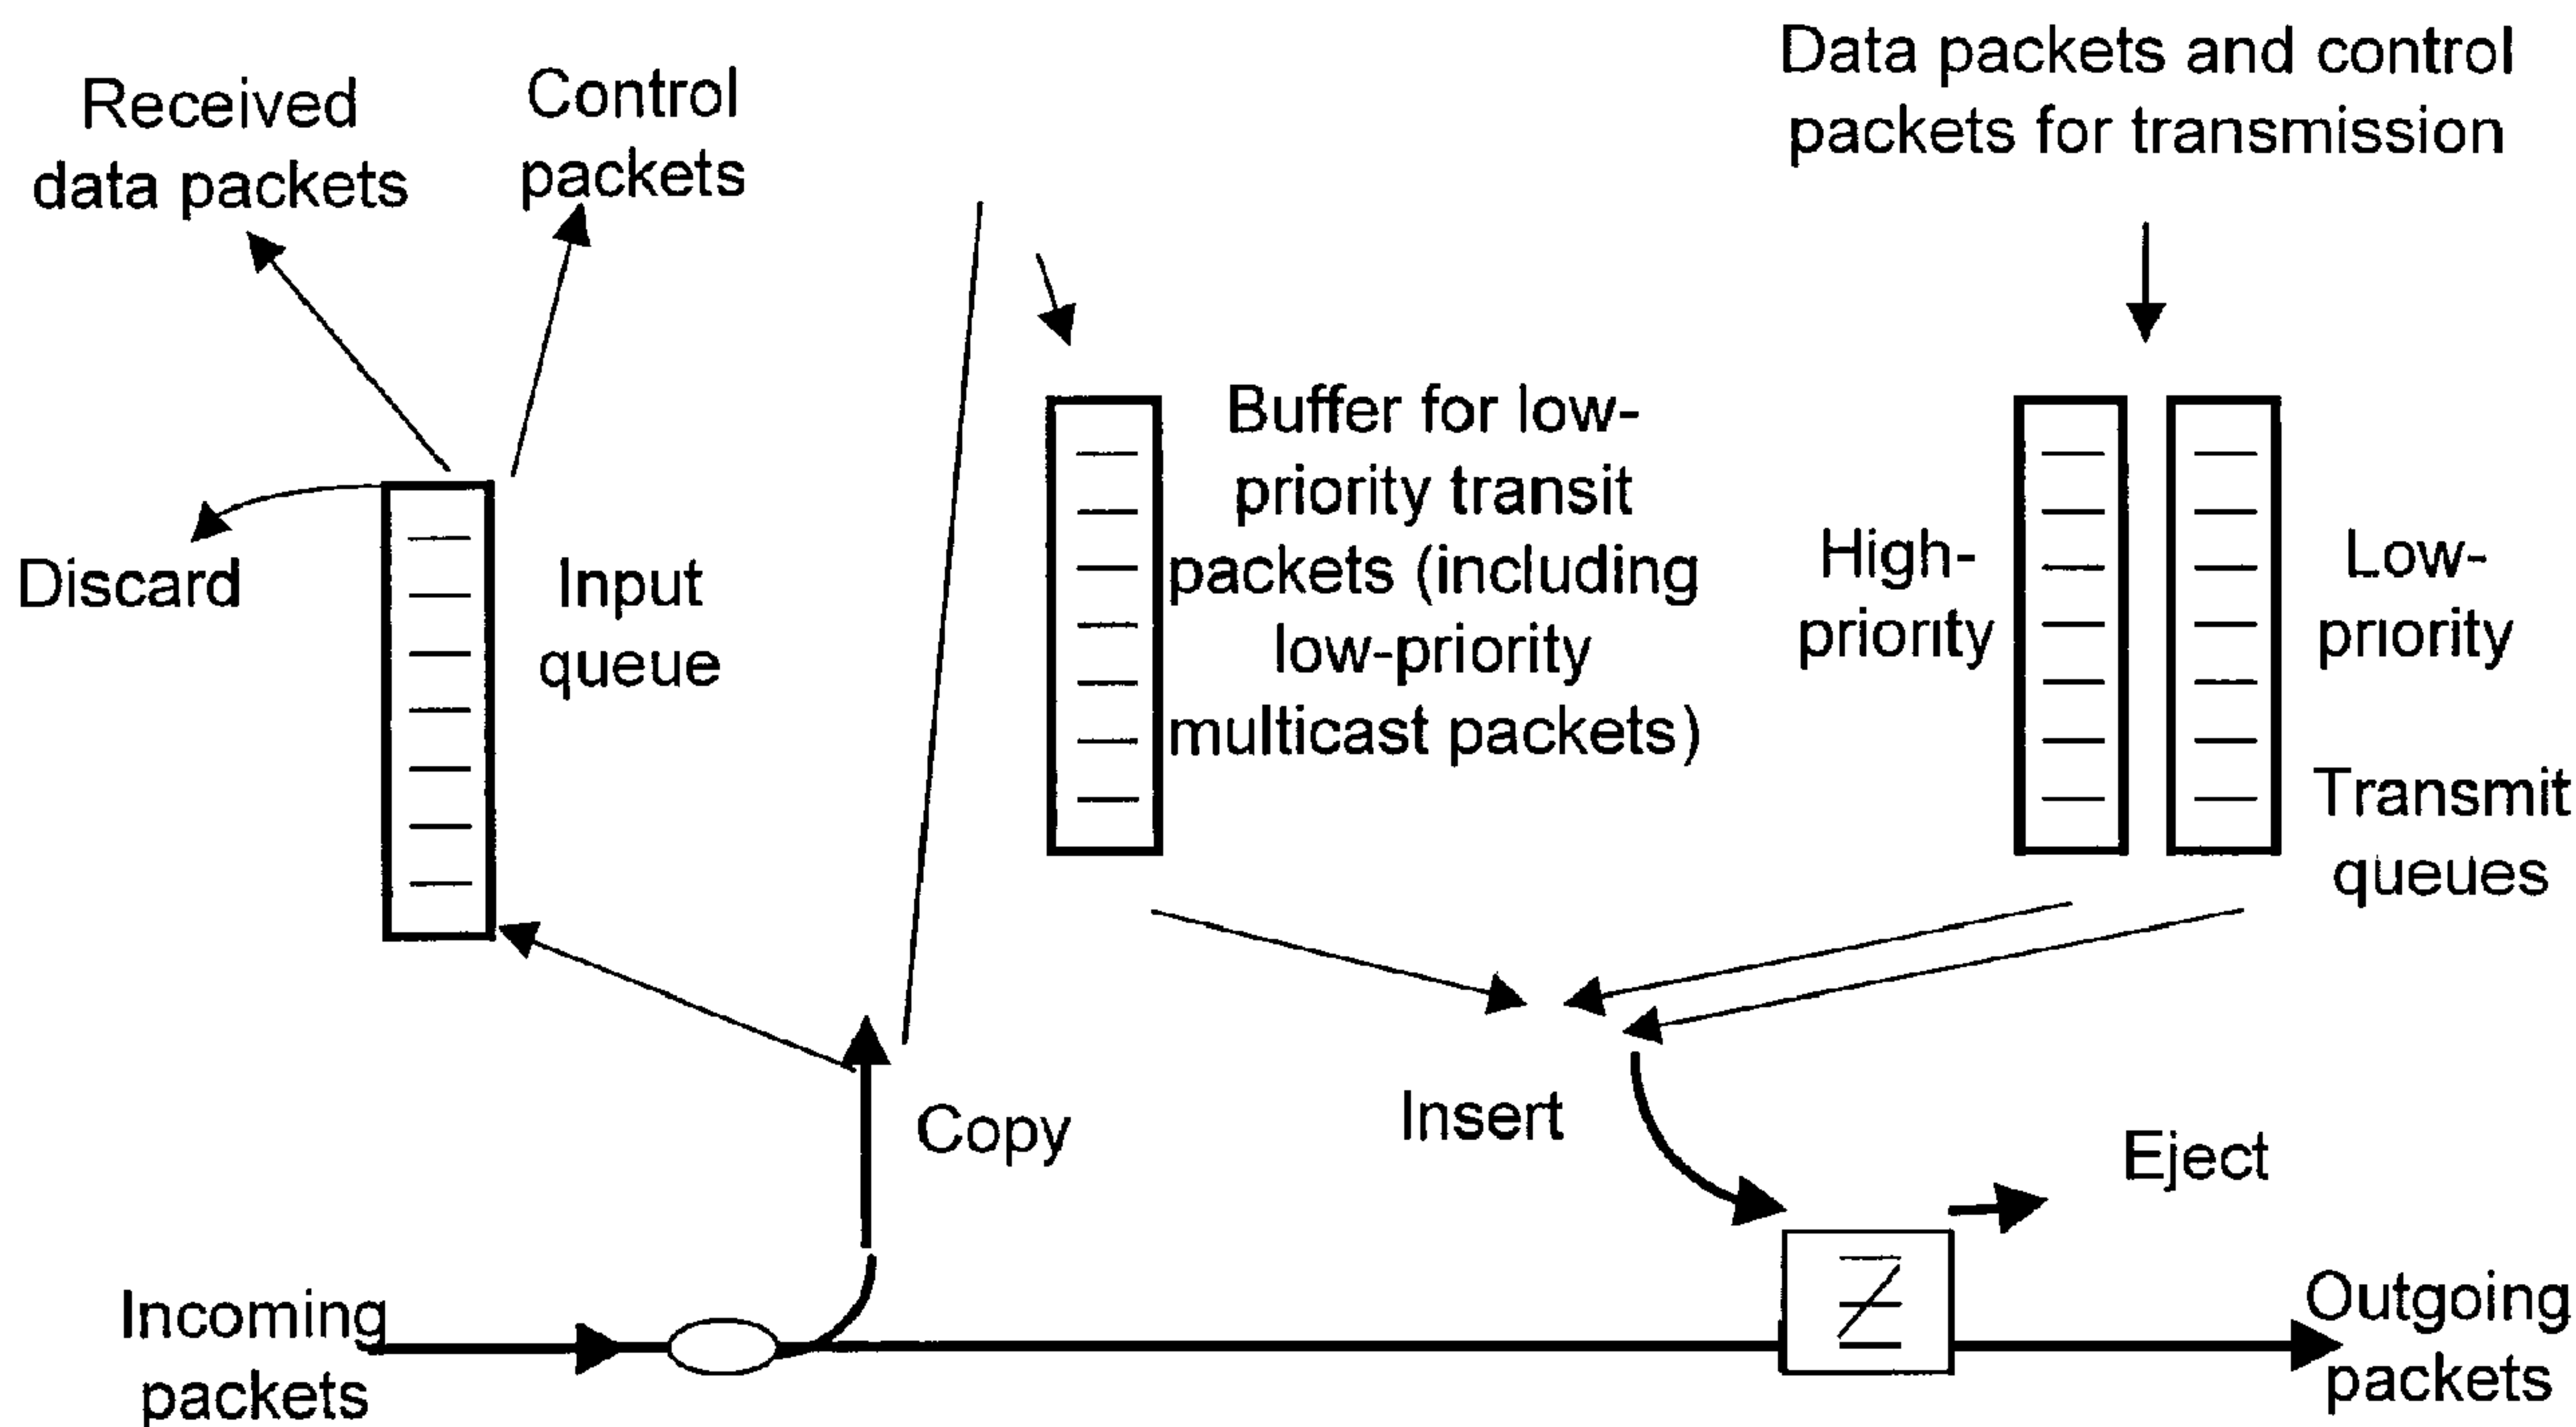




(86) Date de dépôt PCT/PCT Filing Date: 2000/05/05  
 (87) Date publication PCT/PCT Publication Date: 2000/11/16  
 (45) Date de délivrance/Issue Date: 2007/11/06  
 (85) Entrée phase nationale/National Entry: 2001/11/06  
 (86) N° demande PCT/PCT Application No.: GB 2000/001744  
 (87) N° publication PCT/PCT Publication No.: 2000/069126  
 (30) Priorités/Priorities: 1999/05/11 (EP99303672.2);  
 1999/08/12 (GB9919095.1); 1999/12/22 (GB9930373.7)

(51) Cl.Int./Int.Cl. *H04L 12/43* (2006.01),  
*H04J 14/02* (2006.01), *H04L 12/42* (2006.01),  
*H04L 12/56* (2006.01), *H04Q 11/04* (2006.01)  
 (72) Inventeurs/Inventors:  
 COTTER, DAVID, GB;  
 LUCEK, JULIAN KAZIMIERZ, GB  
 (73) Propriétaire/Owner:  
 BRITISH TELECOMMUNICATIONS PUBLIC LIMITED  
 COMPANY, GB  
 (74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : NOEUD ET METHODE DE TRAITEMENT DES PAQUETS DE TRANSMISSION A PRIORITE ELEVEE DANS UN RESEAU DE COMMUNICATIONS OPTIQUES  
 (54) Title: A NODE AND METHOD FOR PROCESSING HIGH PRIORITY TRANSIT PACKETS IN AN OPTICAL COMMUNICATIONS NETWORK



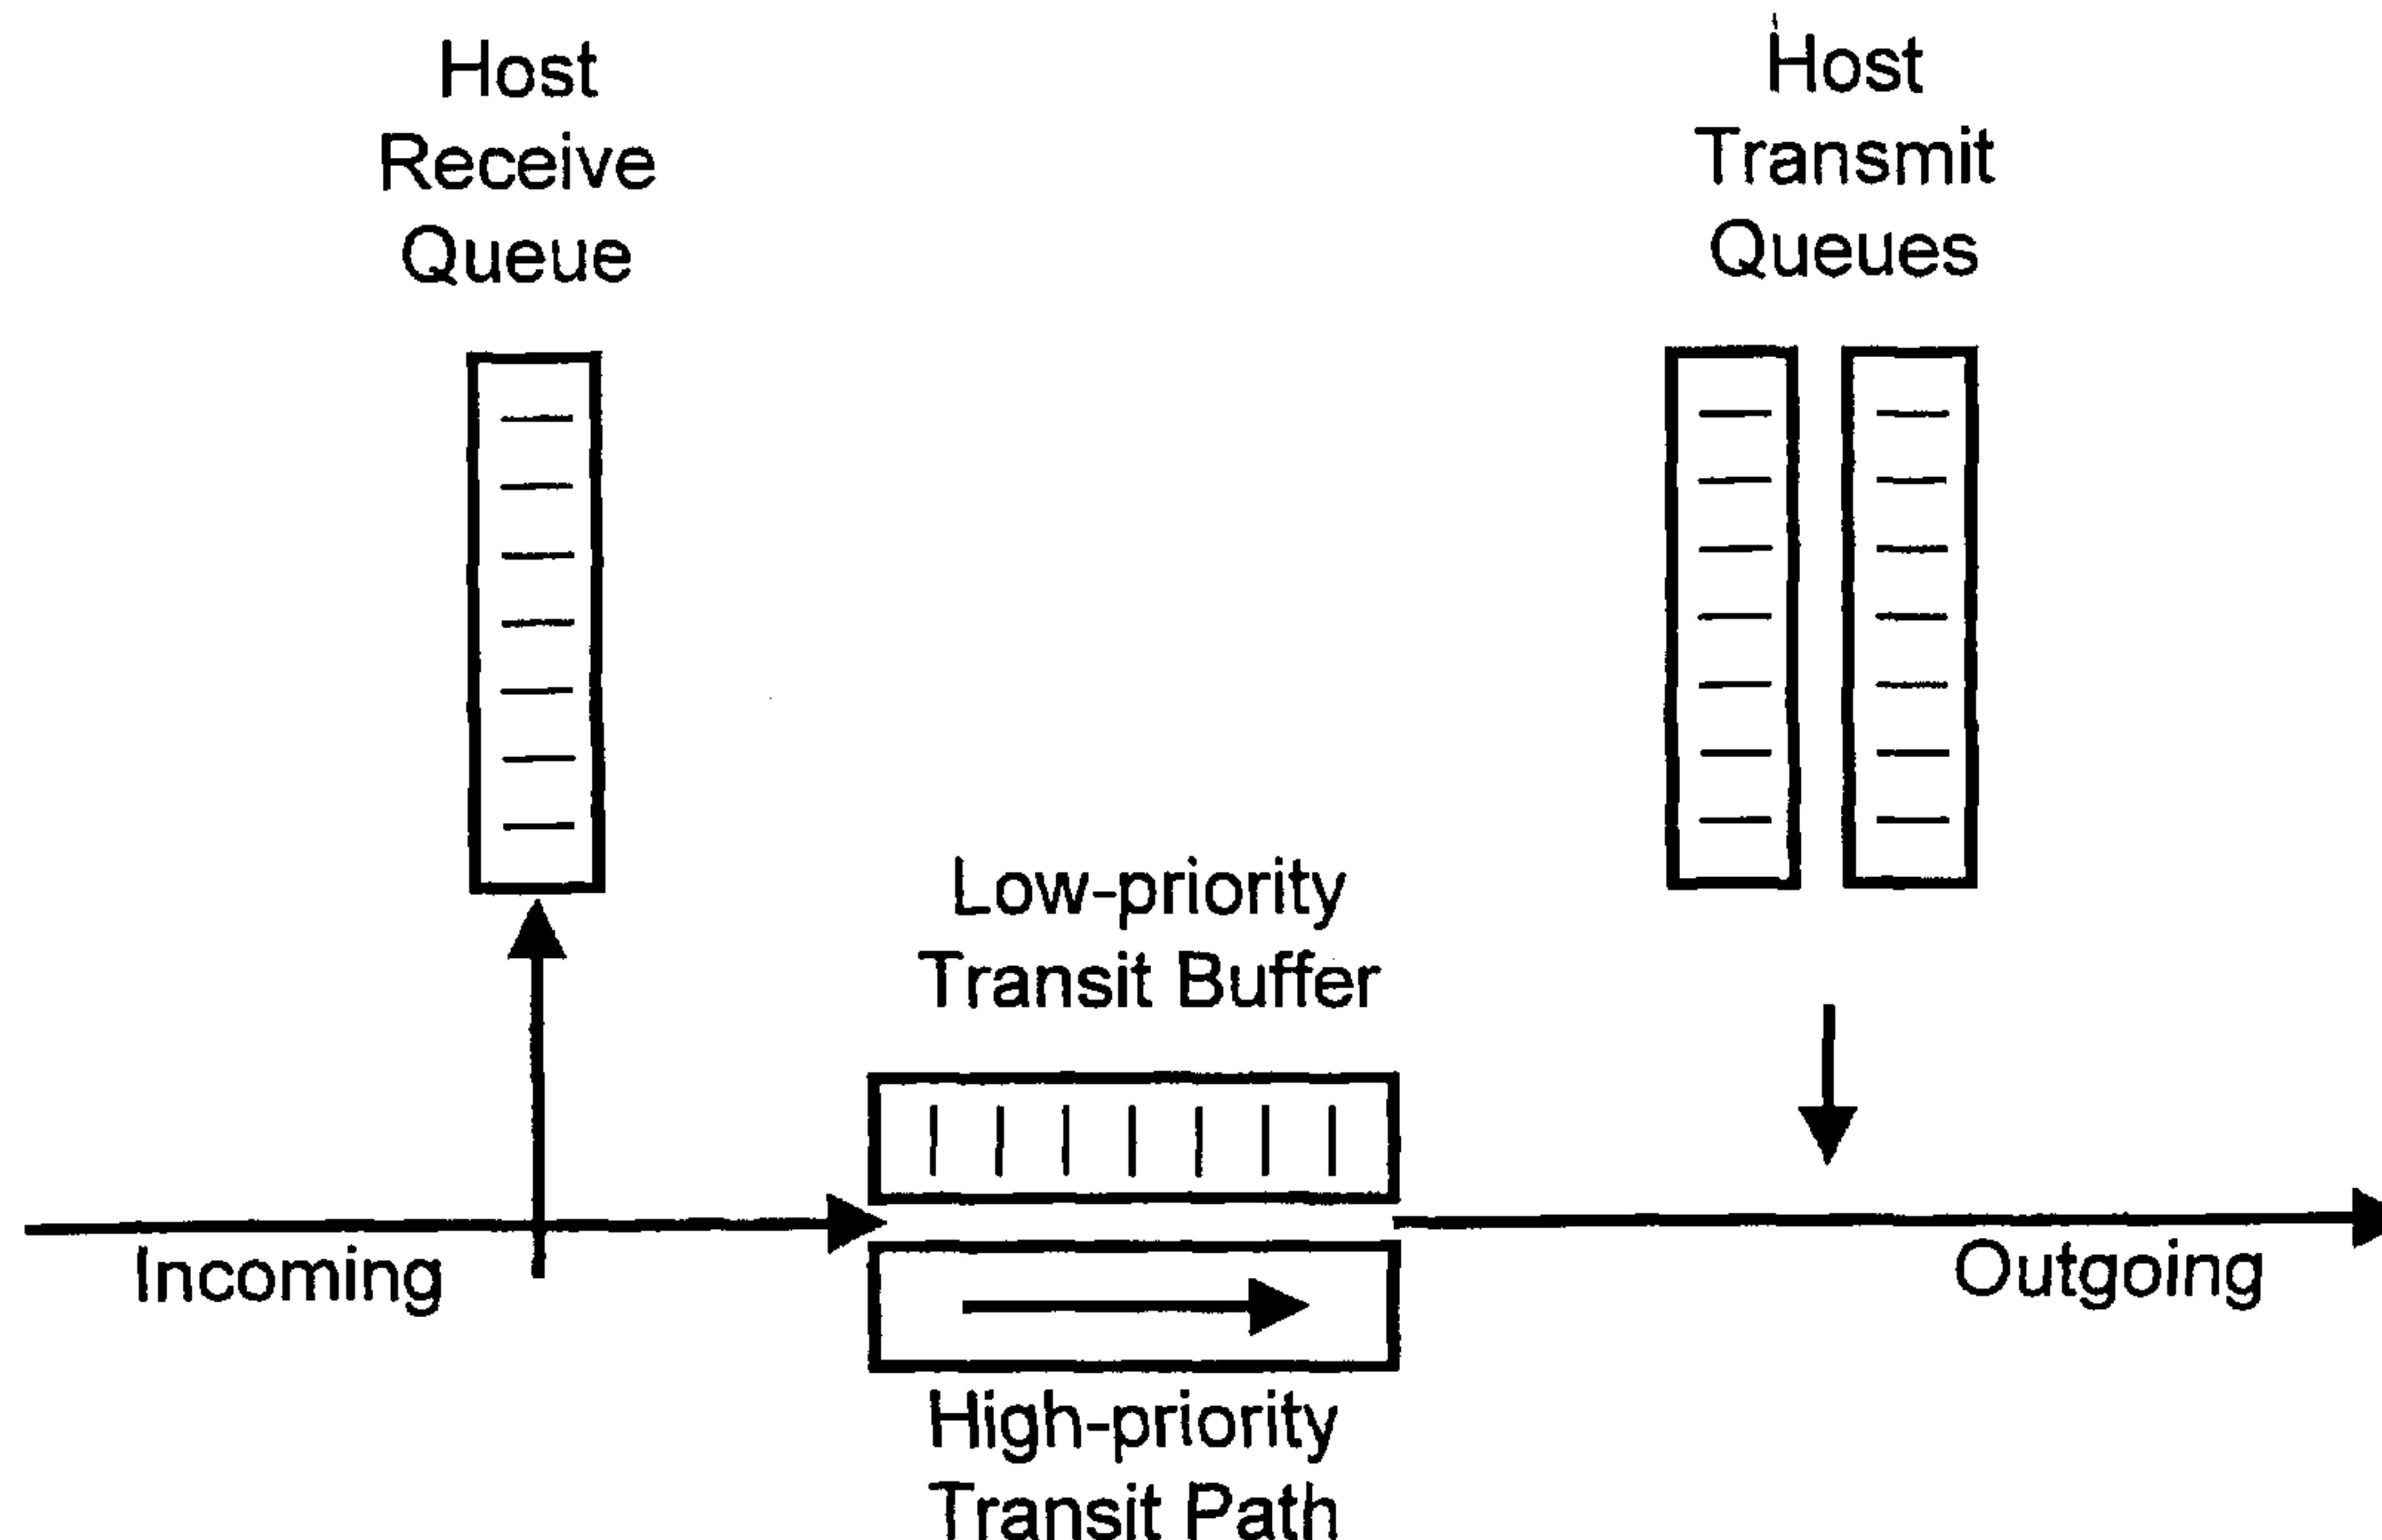
(57) Abrégé/Abstract:

A node in an optical communications network receives optical packets at its input. It determines a priority assigned to each optical packet. When the optical packet is a transit packet destined for another node then, if the packet is determined to have a relatively higher priority, it is output from the node via a continuous-flow transmission path. Other transit packets having relatively lower priorities are sent via an alternative transmission path that may include queues or buffers resulting in a variable delay.

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>7</sup> :</b> <b>H04L 12/43, H04Q 11/00, H04L 12/56</b>	<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 00/69126</b> <b>(43) International Publication Date:</b> 16 November 2000 (16.11.00)
<b>(21) International Application Number:</b> PCT/GB00/01744 <b>(22) International Filing Date:</b> 5 May 2000 (05.05.00) <b>(30) Priority Data:</b> 99303672.2      11 May 1999 (11.05.99)      EP 9919095.1      12 August 1999 (12.08.99)      GB 9930373.7      22 December 1999 (22.12.99)      GB <b>(71) Applicant (for all designated States except US):</b> BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB). <b>(72) Inventors; and</b> <b>(75) Inventors/Applicants (for US only):</b> COTTER, David [GB/GB]; 23 Moorfield Road, Woodbridge, Suffolk IP12 4JN (GB). LUCEK, Julian, Kazimierz [GB/GB]; 52 Wellesley Road, Ipswich, Suffolk IP4 1PL (GB). <b>(74) Agent:</b> WELLS, David; BT Group Legal Services, Intellectual Property Department, 8th Floor, Holborn Centre, 120 Holborn, London EC1N 2TE (GB).	<b>(81) Designated States:</b> AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).  <b>Published</b> <i>With international search report.</i>	

**(54) Title:** OPTICAL COMMUNICATIONS NETWORK**(57) Abstract**

A node in an optical communications network receives optical packets at its input. It determines a priority assigned to each optical packet. When the optical packet is a transit packet destined for another node then, if the packet is determined to have a relatively higher priority, it is output from the node via a continuous-flow transmission path. Other transit packets having relatively lower priorities are sent via an alternative transmission path that may include queues or buffers resulting in a variable delay.

## A NODE AND METHOD FOR PROCESSING HIGH PRIORITY TRANSIT PACKETS IN AN OPTICAL COMMUNICATIONS NETWORK

The present invention relates to a communications network, and in particular to a broadband optical network carrying optical packets.

5       Optical broadband networks have been deployed to meet the rapid increase in demand for bandwidth arising from the growth in the Internet and related data services. Conventionally, packets have been encapsulated for transmission using broadband network protocols such as ATM (Asynchronous Transfer Mode) or SDH (Synchronous Digital Hierarchy). However, it has been recognised that this  
10 introduces a redundant layer of network processing and management and that it would be desirable to transmit packets directly over the optical transmission medium. For example, Cisco have proposed an "Optical Internetworking Roadmap" which describes five stages of evolution from data traffic being carried over SDH/SONET equipment through to full optical internetworking. The penultimate stage of this road  
15 map has been reached with Cisco's introduction of architecture known as a "Dynamic Packet Transport" (DPT) technology. This uses a ring architecture in which packet routers are placed on two concentric counter-directional rings. The rings consist of optical fibre links, each of which terminate at the input and output ports of the routers. DPT operates by sending data packets in one direction and by sending  
20 the corresponding control packets in the opposite direction on the other fibre.

According to a first aspect of the present invention, there is provided a method of operating a node in a communications network, comprising:

- (a) receiving an optical packet at the node,
- (b) determining a priority assigned to the optical packet,
- 25 (c) when the optical packet is determined to be a transit packet having a relatively higher priority, then outputting the packet from the node via a continuous flow transmission path,
- (d) for at least some packets determined to have a relatively lower priority, transmitting the said packets through the node via an  
30 alternative transmission path including one or more queues.

The present invention provides a method of operating a node, which allows express packets, that is high priority transit packets, to flow past any intermediate routers between the source and destination. This gives the advantages of reduced

- 2 -

traffic load at the router, greater throughput, reduced delay, and reduced delay variation. At the same time, the nodes remain capable of implementing queuing and statistical multiplexing of packets where required. Using the invention, express multicast control packets may be transmitted to implement time-critical network management operations, such as support recovery. The fixed speed-of-light latency of express packet transfer also benefits time-critical applications such as distributed database concurrency. It also enables advanced latency-management techniques.

A continuous-flow transmission path is one which transmits packets in a steady, continuous fashion without using queues or other variable delays. The continuous-flow transmission path may be entirely in the optical domain. Alternatively, part or all of the continuous-flow transmission path may be in the electrical domain. In this case, electrical-optical conversion will be carried out at the input and output to the node.

According to a second aspect of the present invention, there is provided a node for connection in a communications network, the node comprising a continuous-flow transmission path connecting the input to the node to the output from the node, an alternative path connecting the input to the node to the output from the node and including one or more packet queues, and

control means arranged to determine the priority assigned to a transit packet to direct a transit packet having a relatively higher priority via the continuous-flow transmission path, and to direct a transit packet having a relatively lower priority via the alternative transmission path.

According to a third aspect of the present invention, there is provided a method of operating a node in a communications network, comprising

- (a) receiving a packet in the optical domain at an input to the node,
- (b) creating a copy of the packet,
- (c) directing at least some of the transit packets received at the input of the node via a continuous-flow transmission path, to the output of the node.

Preferably, the node includes a switch connected to its output, and the method includes selectively ejecting at least some of the packets from the continuous-flow transmission path. This aspect of the invention provides a new node architecture termed by the inventors the "copy-eject-and-insert" (CEI) architecture.

- 3 -

This architecture facilitates the transmission of high-speed multicast packets, and is also particularly suitable for use with the packet prioritisation scheme of the preceding aspects of the invention. In a conventional "drop-and-insert" (D&I) architecture, packets are extracted from the transmission path for the purpose of receiving them at a node. In the case of a CEI node embodying the invention, by contrast, packets are copied but not extracted from the transmission path when they are to be received by the node. All of the incoming traffic may be copied to the node, for example by means of a passive optical coupler. If an incoming packet has reached its destination or is otherwise intended to terminate at the node, then the packet may also be ejected from the transmission path, for example by means of a crossbar switch at the output of the node. When a multicast packet is received at the node, it may both be read from the copied packet and concurrently transmitted onwards via the continuous transmission path.

Systems embodying the present invention will now be described in further details, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic showing a prior art network;

Figure 2 is a schematic showing a first example of a network embodying the present invention;

Figure 3 shows the data flows through a node in the network of Figure 2;

Figure 4 is a schematic showing the architecture of a prior art node;

Figure 5 is a schematic showing the architecture of a node embodying the present invention;

Figure 6 is a schematic showing an alternative node architecture;

Figures 7i to 7iv show alternative configurations for the node of Figure 5;

Figure 8 illustrates the substitution of a high priority packet or a low priority packet;

Figure 9 shows data files within a node for different packet types;

Figure 10 shows the format of an optical packet;

Figure 11 shows the use of dual addresses;

Figure 12 shows a ring-wrap operation;

Figure 13 shows contention at a node output;

Figure 14 shows logic and synchronisation circuits;

- 4 -

Figure 15a shows the structure of a node used to switch wavelength division multiplexed optical packets;

Figure 15b shows the format of a packet handled by the node of 15a; and

Figures 16a to 16e show crossbar switch structures.

5 Figure 1 shows a prior art architecture developed by Cisco and known as the DPT (Dynamic Packet Transport) ring architecture. In this scheme, Internet routers are placed on two concentric counter-directional rings. The rings consist of a sequence of optical fibre links, each of which terminates at the input and output ports of the routers. One of the rings is known as the inner ring and the other as the  
10 outer ring. Data packets are sent in one direction and corresponding control packets are sent in the opposite direction on the other fibre. This IP (Internet Protocol) ring network is designed to provide a number of features. In particular, it supports statistical multiplexing of packets with no provision of point-to-point connections or use of dedicated bandwidth for interconnection of routers or for protection. It is also  
15 designed to support packet prioritisation and to offer multiple levels of queuing and scheduling and to support both unicast and multicast transmissions. In the network of Figure 1, by contrast with the network described below with reference to Figure 2, the signal transmission path is interrupted at each node.

Figure 2 shows a network embodying the present invention. A  
20 communications network 1 comprises concentric outer and inner rings 2.1, 2.2. The rings carry optical packets. A number of nodes 3a to 3g are connected to the rings. In this example, each of the nodes 3a to 3g is an Internet Protocol router. In operation, one of the routers, for example 3a, outputs an optical packet addressed to another of the routers, for example 3d, onto one of the dual concentric rings. The  
25 optical packet output by node 3a carries the network address of node 3d in an address field in the packet header. The packet passes around the ring. At each intermediate node, the packet address is compared with the address of the respective node. If the packet is not addressed to a given intermediate node, then the packet passes on to the next node, and so on. At the destination node, the address is read  
30 and the packet is recognised as being intended for receipt at that node. The packet, in the case of a unicast transmission is removed from the network at that node. It may then, for example, be converted to the electrical domain for onwards

- 5 -

transmission to a customer terminal on an electronic network connected to the respective node.

Figure 3 shows the flow of received and transmitted data packets at one of the nodes on one of the fibre rings. An incoming data packet is (i) sent to host receive queue (if it has reached its destination) or (ii) is passed over for onward transmission. A multicast packet may both be sent to the host receive queue and also be forwarded for onward transmission. Packets for onward transmission are treated according to their priority: a low-priority packet is passed to a transit queue. An express high-priority, packet is passed immediately to the outgoing path without queuing. When capacity on the ring allows, a packet from one of the transmit queues may be passed to the outgoing path. Multiple transmit queues may be used to manage prioritisation of outgoing traffic from the node.

In this example, the optical fibre rings and the nodes are configured to provide a continuous optical path for express packets.

In a continuous transmission path signals stream into and out from the path at a continuous and uniform rate. The path may be a continuous optical transmission path. The continuous optical transmission path may contain a substantially fixed delay as in Figure 3. In the case of a continuous transmission path in the electrical domain the path may contain storage elements such as an FIFO (first-in first-out) buffer. This differs from conventional systems, however, in this case, in that in normal operation, the output from the buffer is continuous and regular, so that the input and output of the buffer occur at substantially the same rate. In conventional prior systems, a buffer is used to hold transit packets in a queue which is serviced at rates and/or times which are dependent on factors such as packet priorities, queue length, status of contending queues etc.

A further function carried out by the control circuits is to determine whether the time-to-live (TTL) of a packet has expired. If the time-to-live is determined to have expired, then again a control signal is sent to the switch to cause such a packet to be ejected from the onwards transmission path. When one of the nodes generates and outputs an optical packet onto the ring, a value is written in a phase field "PH" in the packet header. One of the nodes functions as a master node. The master node outputs multicast express control packets that travel around each ring to inform each node of the current phase value. The phase value is regularly updated, at a rate of at

- 6 -

least once per round-trip time for speed-of-light travel around the ring. The updating of the phase value is carried out by the master node incrementing the value. For example, in the case of one control packet per round trip of the ring, the phase value is incremented by 1 (modulo 3). Each ring periodically receives and stores the  
5 current phase value as indicated by one of the multicast express control packets from the master node. When any other express packet is received at the node the control logic in the node determines whether the value in the PH field of the received express packet is such that:

$$(\text{current\_phase} - \text{PH}) \bmod 3 > 1,$$

10 If this inequality is satisfied, then the packet is expired and should be purged from the ring. This may be the case, for example, because the packet is addressed to a node which is currently malfunctioning or because there is an error in the address of the packet.

Figure 4 shows the traditional 'drop-and-insert' (D&I) architecture that has  
15 been widely considered previously for optical packet networks (e.g. J R Sauer, M N Islam and S P Dijaili, "A soliton ring network", Journal of Lightwave Technology, vol. 11, no. 12, December 1993, pp. 2182-2190). The node contains a 2x2 optical crossbar switch. When the switch is in the bar (straight-through) position: (i) an optical packet on the incoming fibre is 'dropped' (i.e. switched to the node receiver);  
20 or (ii) an optical packet created by the host transmitter is 'inserted' (i.e. switched to the outgoing fibre); or both (i) and (ii) occur simultaneously. A limitation of the D&I node architecture is that when a packet is received by the node it is physically removed from the optical path between the incoming and outgoing fibres. The only method for multicasting a packet is, at each node, to drop the packet and then  
25 subsequently reinsert it when sufficient vacant capacity on the ring becomes available. Therefore this architecture cannot support multicasting of express (high-priority) transit packets.

Figure 5 shows a preferred form of the new node architecture that we have invented. Instead of 'drop-and-insert' (D&I), this new architecture provides the  
30 function of 'copy-eject-and-insert' (CEI). In the case of a D&I node, packets are extracted from the transmission path for the purpose of receiving them at a node. In the case of a CEI node, packets are copied but not extracted from the transmission path for the purpose of receiving them at a node. All incoming traffic is copied to the

- 7 -

node, for example by means of a passive optical coupler as shown in Figure 5. In use, an incoming fibre from one of the optical rings passes to a coupler 5.1. From the coupler 5.1, one optical path passes through a further length of fibre 5.3 providing a fixed delay to an optical crossbar switch 5.4. An outgoing fibre  
5 connected in the fibre ring passes from one of the output ports of the crossbar switch. The other output from the coupler 5.1 is split again, for example, using a further coupler 5.2. One branch of the output from this further coupler is used to provide a copy of the optical packet on the incoming fibre. This copy may be passed, for example, to the host receive queue. The other branch of the output of  
10 the further coupler passes to control circuits. These control circuits may, for example, read a header carried with the optical packet, and carry out associated logic operations. A control output is generated by the control circuits and passes to the crossbar switch. If the control circuits determine that the packet is addressed to the respective node, then the crossbar switch is set to the cross state in order to eject  
15 the packet (in the case of a unicast transmission). If an incoming packet has reached its destination or is otherwise intended to terminate at the node, the packet is also ejected from the ring, for example by means of the crossbar switch as shown in Figure 5. Simultaneously, or whenever sufficient vacant capacity on the ring become available, a packet may be inserted onto the ring by the node. The control logic and  
20 synchronisation sub-systems use information contained in the packet (for example in the packet header) and other information (for example the status of transmit queues and the status of control flags) to operate the crossbar switch. Figure 5 shows a fixed optical delay, which allows sufficient time for the operation of these control logic, synchronisation sub-systems and switch before the arrival of the packet at the  
25 switch. As in the D&I architecture, an express (high priority) transit packet is passed directly to the outgoing fibre. However, unlike D&I, the CEI architecture allows express transit packets to be multicast, because they are simultaneously copied without delaying their onward passage.

Figure 14 shows the control logic subsystem in more detail. It includes  
30 optical logic stages 141 and combinatorial electronic logic gates 142. The four optical blocks (synch, address recognition PH reader and Unicast/Multicast) each have copies of the optical packet to the input. The output from the synchronisation block is an optical path to each of the other three optical blocks. The synch block

- 8 -

may be based on one of the self-synchronisation techniques we have described in our patent EP-B-687370 (eg. using pulses separated by 1.5 bit periods input to an optical AND gate). The address recognition block may be based on the technique we have described before in that patent. Addresses are coded using specially selected binary  
5 words, and recognised by inputting address and target word to an optical AND gate. The output from the AND gate is converted to give an output from the block that is an electrical binary signal that says packet ADDRESSEE matches/does not match local address. The PH may be two optical AND gates, each having the optical packet as one input and a synch pulse as the other input - this synch pulse timed to overlap  
10 with one of the PH bits. The output from PH reader block is two parallel electrical binary signals - each denoting one of the PH bits. The UNI/MULTI reader may be one optical AND gate, having the optical packet as one input and a synch pulse as the other input - this synch pulse timed to overlap with the UM bit. The output from UNI/MULTI reader block is one electrical binary signal - denoting unicast/multicast.  
15 These electrical signals, together with signals from the queue status then pass to the high-speed electronic logic part. This part performs the logic which is set out in the table below describing the action of the node for various types of packet (e.g. multicast low-priority transit packet, etc). The output from this electronic logic is a binary electrical signal to set the 2x2 cross bar optical switch configuration.

20 Each stage of this logic has to operate within a time shorter than the shortest packet. e.g. for a 50 byte packet at 100 Gbit/s = 4 ns. Since the optical stage and electronic logic stage are arranged in pipeline, each stage would have to take less than the minimum time (eg 4 ns). The optical stage is ultrafast - e.g. using four-wave mixing in semiconductor optical amplifier, or TOAD device, etc. to implement the  
25 optical AND gate. Since the electronic logic required is simple it is possible to construct suitable fast circuit using hard-wired combinatorial logic.

Figure 16a shows the design of a simple crossbar switch. This includes a region 161 of lithium niobate. When an electrical control signal is applied to a gate on the lithium niobate, then the refractive index of the lithium niobate changes, so  
30 that the optical signal is coupled across from one optical path to the other, the "cross" state of the switch.

Figure 16b shows an alternative design for a crossbar switch. In this case, the switch uses a Mach Zehnder interferometer configuration with a control element

- 9 -

162 in one branch of the interferometer. A phase change is generated by the control element 162 when the control signal is applied, to switch the output from one optical output port to the other. Both of these designs suffer significant cross talk: that is, when an optical signal is switched to one of the ports, a significant level of optical power, for example 20 or 30 percent of the input power, is present at the other output port.

To overcome this problem, each crossbar switch in the nodes described above may be replaced by a network of four such switches in the configuration shown in Figure 16c. The switches are cross-connected as shown. The control signal  $c$  is applied to the first pair of switches in the signal path, and the logical compliment of the control signal,  $\bar{c}$  is applied to each of the second pair of switches in the signal path. Logically, the switching effect of such a network is equivalent to that of each the single switches of Figure 16a or 16b. However, the crosstalk is reduced by an order of magnitude. Since the crossbar switch required only two inputs, one for the insert path and one for the straight through path, and only two outputs, one for the eject path and one for the output path, it can be shown that the full network shown in figure c is not required. Instead, the full crossbar switching function with the benefit of crosstalk reduction can be achieved using a pair of cross bar switches connected in series as shown in Figure 16d.

Figure 16e shows a further alternative design also offering improved cross talk performance. This uses a pair of passive optical couplers 163, 164 with an electro-absorption modulator (EAM) 165 connected between them. The copy optical signal is taken from the first of the optical couplers 163 and the insert optical signal is added at the second of the optical couplers 164. In the absence of a control signal  $c$ , the EAM 165 is in its transparent state and the optical packet passes straight through from the input path to the output path. When a packet is to be ejected then the control signal  $c$  switches on the EAM 165. The ejected packet may then be replaced by further packet added on the INSERT optical path of the coupler 164.

Figure 6 shows an example of the CEI architecture in an alternative form. Here there are two crossbar switches, SW1 and SW2. (For clarity, the control logic and synchronisation sub-systems are not shown.) In this case not all incoming packets are copied; instead only those that arrive whilst SW1 is in the cross state will be copied. In this case an express multicast packet may be copied and

- 10 -

immediately reinserted onto the ring if switches SW1 and SW2 are both placed in the cross state. However in this case the express multicast packet is delayed by the transit time of the feedback loop from SW1 to SW2 and back again to SW1. The alternative form of the CEI architecture shown in Figure 6 has certain drawbacks:

5 First, the forced additional delay for express multicast packets increases the probability of contention at SW1 (described further below) unless special preventative measures are used (such as, for example, providing a sufficient time guard band following an express multicast packet, with the resulting penalty of reduced network throughput and higher complexity). Second, the architecture shown in Figure 6

10 requires two optical switches, rather than one.

In the case of the D&I architecture (Figure 4), the extraction of packets from the transmission path for the purpose of receiving them at a node is necessarily a selective process in the optical domain, and may be performed for example by an optical crossbar switch. In the preferred form of the CEI architecture (Figure 5), the

15 copying of packets from the transmission path for the purpose of receiving them at a node is non-selective in the optical domain, and may be performed for example by a passive optical coupler.

In the CEI architecture the continuous transmission path for express transit packets is preferably a continuous optical transmission path, but not necessarily so,

20 as shown in Figures 7 (i)-(iv). Figure 7(i) shows a preferred arrangement with a continuous optical transmission path for express transit packets. Figure 7(ii) shows a transponder (a regenerator using optical-electrical and electrical-optical conversion stages). Figure 7(iii) shows a case where the copy function is performed in the electrical domain. Figure 7(iv) shows a case where the copy, eject and insert

25 functions are all performed in the electrical domain. In this case the switch can be a 1x2 electrical switch, and as in the case of optical switching, the channel left open circuit is in effect 'ejected'.

Figure 7(iv) is distinguished from conventional nodes by a 'continuous transmission path'. In the continuous transmission path signals stream into and out

30 from the path at a continuous and uniform rate.

An important feature of an optical packet network with different priority classes is the ability to reallocate the time occupied on the ring by a transit packet to another packet with higher priority. The reuse of certain time slots using a D&I

- 11 -

optical node architecture is previously known. For example, the previously cited paper of Sauer et al describes a D&I optical node in which a slot made vacant by dropping a packet (i.e. removing from the optical ring a packet when the packet destination address matches the node address) can be immediately reused by inserting a transmit packet. Here, distinctively, we reallocate any time slot by ejecting an incoming packet from the ring so as to vacate the time slot, and immediately or simultaneously insert into the same time slot another packet taken from one of the transit or transmit queues. Using the CEI architecture, the steps required to reallocate to a packet of higher priority a time slot currently occupied by an incoming transit packet are: copy the incoming transit packet; determine its priority; if appropriate for reallocation, then do {eject the packet from the ring; deliver the packet copy to the transit buffer; transmit another packet of higher priority reusing the same time slot}. This is illustrated in Figure 8. Notice that the CEI architecture is not essential for this process. For example, using a D&I architecture the steps required would be: determine the priority of an incoming transit packet; if appropriate for reallocation, then do {drop the packet; store the packet for later transmission; and transmit another packet of higher priority reusing the same time slot}. Although reference is made here to 'slot', this feature is equally applicable to unslotted systems in which packets arrive asynchronously, and may be of variable length.

Figure 9 shows the flow of incoming and outgoing packets in the optical CEI node of Figure 5, so as to realise the architecture of Figure 3. All packets are assumed to be in one of two priority classes: low and high (express). Packets may also be unicast or multicast, and they may be data packets (with a payload of user data) or control packets (uniquely for network control purposes). All incoming packets are copied, whereupon they may be allocated to various input queues or else discarded. The operations at the node to handle different types of legitimate (e.g. non-expired) incoming packet are as follows:

- Unicast
  - Express transit packet: The crossbar switch is put in the bar (straight through) state so that the packet is passed directly to the outgoing fibre. The packet copy is discarded.

- 12 -

- 5

  - Express packet addressed to node: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to the input queue, and subsequently sorted according to whether it is a data or control packet.
  - Low-priority transit packet: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to the low-priority transit buffer.
  - Low-priority packet addressed to node: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to the input queue, and is subsequently sorted according to whether it is a data or control packet.
- 10

  - Multicast
    - Express packet (different source): The crossbar switch is put in the bar state so that the packet is passed directly to the outgoing fibre. The packet copy is delivered to the input queue, and is subsequently sorted according to whether it is a data or control packet intended for delivery to the node, otherwise it is discarded.
    - Express packet (node is the source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is discarded.
    - Low-priority packet (different source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to low-priority transit buffer. The packet copy is also delivered to the input queue, where it is subsequently sorted according to whether it is a data or control packet intended for delivery to the node, otherwise it is discarded.
    - Low-priority packet (node is the source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is discarded.
- 15
- 20
- 25
- 30

Packets for transmission from the node are selected from the low-priority transit buffer or the transmit queues according to the current status of queue depths, appropriate rate controls, fairness algorithm, etc. Subject to these controls, packets are transmitted as vacant capacity becomes available on the ring (either vacant

- 13 -

capacity on the incoming fibre or new vacant capacity created by the ejection of packets from the ring by the node itself).

Each packet carries an appropriate MAC protocol header, designed according to the principles set out earlier (Figure 10 shows a suggested example). This MAC header can consist of two parts: one part (which we call the 'optical MAC header') contains the minimum information needed for the processing functions that must be carried out at high speed 'on the fly' to enable the operations listed above. The second part of the MAC header contains information needed for other Layer 2 functions.

10 In the example shown in Figure 10, SS (self-synchronisation) is a 2-bit field used for timing recovery. ADDRESSEE is a 10-bit field used to indicate which node should strip the packet from the ring. In the 'Spatial Reuse Protocol' (SRP) described by Cisco (part of the Dynamic Packet Transport), packets are stripped from the ring by the destination node in the case of unicast packets. It is proposed here that for  
15 the purpose of routing in the high-speed network, the optical MAC header will not include both the source and destination addresses, but instead will contain a single address field ADDRESSEE. For the transmission of a unicast packet, ADDRESSEE will be set to the destination address. For the transmission of a multicast packet, ADDRESSEE will be set to the source address. A node is required to eject a packet  
20 from the ring if ADDRESSEE matches the address of the node. UM (unicast/multicast) is a 1-bit field used to indicate whether the packet is unicast or multicast. PH (phase) is a 2-bit field with a dual purpose; as described more fully below, PH indicates the priority of the packet and also contains a simple time-to-live mechanism. P (parity) is a 1-bit field used to set the parity of the optical MAC header  
25 (approximately 2 bytes), and will be used in combination with 'on the fly' parity checking to provide some header integrity.

The reading of the optical MAC header must be performed 'on the fly' using very simple high-speed logic operations. For ultra-high-speed optical implementation, processing based on simple optical logic gates, such as AND, will be used. Various  
30 methods of timing recovery based on self-synchronisation are described in our patent EP-B-687370. For example, the field SS could consist of 2 return-to-zero format optical pulses separated by 1.5 bit periods, and self-synchronisation would be based on the output from an optical AND gate whose first input is a copy of the packet and

-14-

the second input is a further copy of the packet delayed by 1.5 bit periods with respect to the first. The output from the AND gate is an optical pulse in precise  
5 synchronism with the start of the packet. This pulse (which we term here the "timing pulse") can then be used for a variety of operations on the fly. The timing pulse may be used in the process of determining whether there is a match between ADDRESSEE and the address of the node, as described in our patent EP-B-687370. The timing pulse may also be used in combination with optical AND gates to read the  
10 fields UM and PH.

The dual ring structure is resilient because, in the event of an isolated fibre break or node failure, the nodes can perform the "ring wrap" operation illustrated in Figure 12 [as described, for example in "Dynamic Packet Transport Technology and Applications Overview, published by Cisco Solutions Inc. in 1999]. For proper  
15 operation of the control mechanisms which allow the network to reconfigure and re-establish after ring wrap, it is clearly necessary for the nodes to distinguish between those incoming packets which are travelling on their "correct" ring and those that have been wrapped onto their "wrong" ring. Here we propose the technique of "dual addressing" to avoid the need for an additional header field to indicate the "correct"  
20 ring for each packet. Dual addressing is illustrated in Figure 11. Rather than provide each node with a single address, it is proposed to provide each node with two addresses, one for each of the rings. In this case, the transmitter of a packet uses the appropriate value of ADDRESSEE corresponding to the ring that is being used. In the event of ring wrap, the usual rule for packet stripping is simply followed (a node  
25 is required to eject a packet from the ring if ADDRESSEE matches the address of the node), without the need to read further header fields to check whether the packet has been wrapped onto the "wrong ring".

Some network control operations require point-to-point signalling between adjacent nodes. This can be done in a number of ways: creation of an independent  
30 control packet; "piggy back" technique such as overwriting certain fields in a newly created data packet or in a low-priority transit packet; or out-of-band signalling. Out-of-band signalling could be performed in time guard bands between packets.

A potential limitation of the CEI node architecture is the possibility of contention at the eject-and-insert switch. This is illustrated in Figure 13. Contention  
35 may arise when an incoming packet B arrives when the switch is in the cross position

- 15 -

and the node is in the process of inserting a packet A. In the absence of an optical buffering mechanism, various possibilities for contention resolution are available including switch over to the bar state to allow B to pass, thus forcing the ejection of part of A. Resend A later. The recipients of A will recognise that the packet has  
5 been truncated, and discard it.

A node embodying the present invention may also be used in a network employing wavelength division multiplexing. Figure 15a shows one example of such a node. A wavelength division multiplexed packet on the incoming fibre passes to a wavelength division demultiplexer 15. Suitable wavelength division multiplexers are  
10 available commercially and may be formed, for example, from wavelength-selective optical couplers or from a fibre Bragg grating. The demultiplexer 151 outputs different wavelength channels on different output fibres. For ease of illustration, only four wavelength channels,  $\lambda_1 - \lambda_4$  are shown, but in practice, a greater number of wavelength channels, for example 8, 16 or 32 channels may be used. The optical  
15 signal on each wavelength channel is passed through an opto-electronic converter, such as a photodiode, and is converted to a signal in the electrical domain. The outputs from the opto-electronic converters are tapped to provide the copy of the packet. The copy may be passed to the host receive queue and to control/synchronisation circuits, as described previously in relation to the other  
20 embodiments. The outputs of the opto-electronic converters are also connected via an electrical continuous-flow path to the input of an electronic switch 152. A parallel datastream of signals to be inserted into the flow path are also input to the switch 152. The parallel outputs of the switch 152 pass through electro-optic converters and are converted back into the optical domain at different respective wavelengths  
25  $\lambda_1 - \lambda_4$ . The signals at the different wavelengths are combined onto the outgoing fibre by an optical multiplexer 153. Figure 15b shows the format of the optical packets as received on the incoming fibre and as output on the outgoing fibre. In this example, 32 wavelengths are used and the packets are 32bits wide in the wavelength domain. Bit positions 1, 33, 65,... are received at wavelength  $\lambda_1$ , bit  
30 positions 2, 34, 66,... at wavelength  $\lambda_2$ , and so on. In the node, in the input fibre prior to demultiplexing of the WDM packet, the optical signals may be passed through a dispersion-compensating device, such as a length of dispersion-

- 16 -

compensating fibre, to realign the different wavelength channels in the time domain so as to correct for the effects of group-velocity dispersion in the optical network.

**CLAIMS**

- 5 1. A method of operating a node in a communications network, the node having an input and an output with a continuous flow transmission path therebetween, the continuous flow transmission path including a transmission medium, the method including the steps of:
- 10 (a) receiving an optical packet at the node;
- (b) copying the packet such that there is a first packet and a second packet, the first and second packets being duplicates of one another;
- (c) reading at least part of the first packet to determine if (i) the packet is a transit packet and if (ii) the packet is a high priority packet;
- (d) if the first packet is determined to be a transit packet having a high priority, then outputting the second packet from the node via the continuous flow transmission path; and,
- 15 (e) if the first packet is determined to be a transit packet not having a high priority, placing the first packet in a queue for later transmission and ejecting the second packet from the transmission medium.
- 20 2. A method as claimed in claim 1, including the further step of:
- (f) if the first packet is determined to be addressed to the node, ejecting the second packet from the transmission medium, and receiving the first packet for storage and/or processing at the node.
- 25 3. A method as claimed in claim 1 or claim 2, wherein the continuous flow path is an optical path.
4. A method as claimed in any one of claims 1 to 3, wherein the priority of a packet is determined from a flag carried with the packet.
- 30 5. A node for use in a communications network, the node having:
- (a) an input and an output with a continuous flow transmission path therebetween, the continuous flow transmission path including a transmission medium, the input being adapted for receiving an incoming optical packet;
- 35

-18-

- (b) copying means for copying the incoming packet such that there is a first packet and a second packet, the first and second packets being duplicates of one another;
- 5 (c) processing means arranged to (i) read the first packet, (ii) determine whether a packet is a transit packet, and (iii) determined whether a packet is a high priority packet; and,
- 10 (d) control means which, in response to the processing means determining that the first packet is a transit packet having a high priority, are arranged to direct the second packet to the output of the node via the continuous flow transmission path, the control means being further arranged to, in response to the processing means determining that the first packet is a transit packet not having a high priority, (i) eject the second packet from the transmission medium and
- 15 (ii) to place the first packet in a queue for later transmission.

6. A node according to claim 5, further comprising a wavelength demultiplexer connected to the input to the node and arranged to demultiplex a wavelength-multiplexed packet received at the input to the node, and a

20 wavelength multiplexer connected to the output of the node and arranged to multiplex a plurality of wavelength channels to form a wavelength division multiplexed optical packet for transmission from the output of the node.

7. An optical communications network including one or more nodes

25 according to claim 5.

8. An optical communications network as claimed in claim 7, wherein the network has a ring topology.

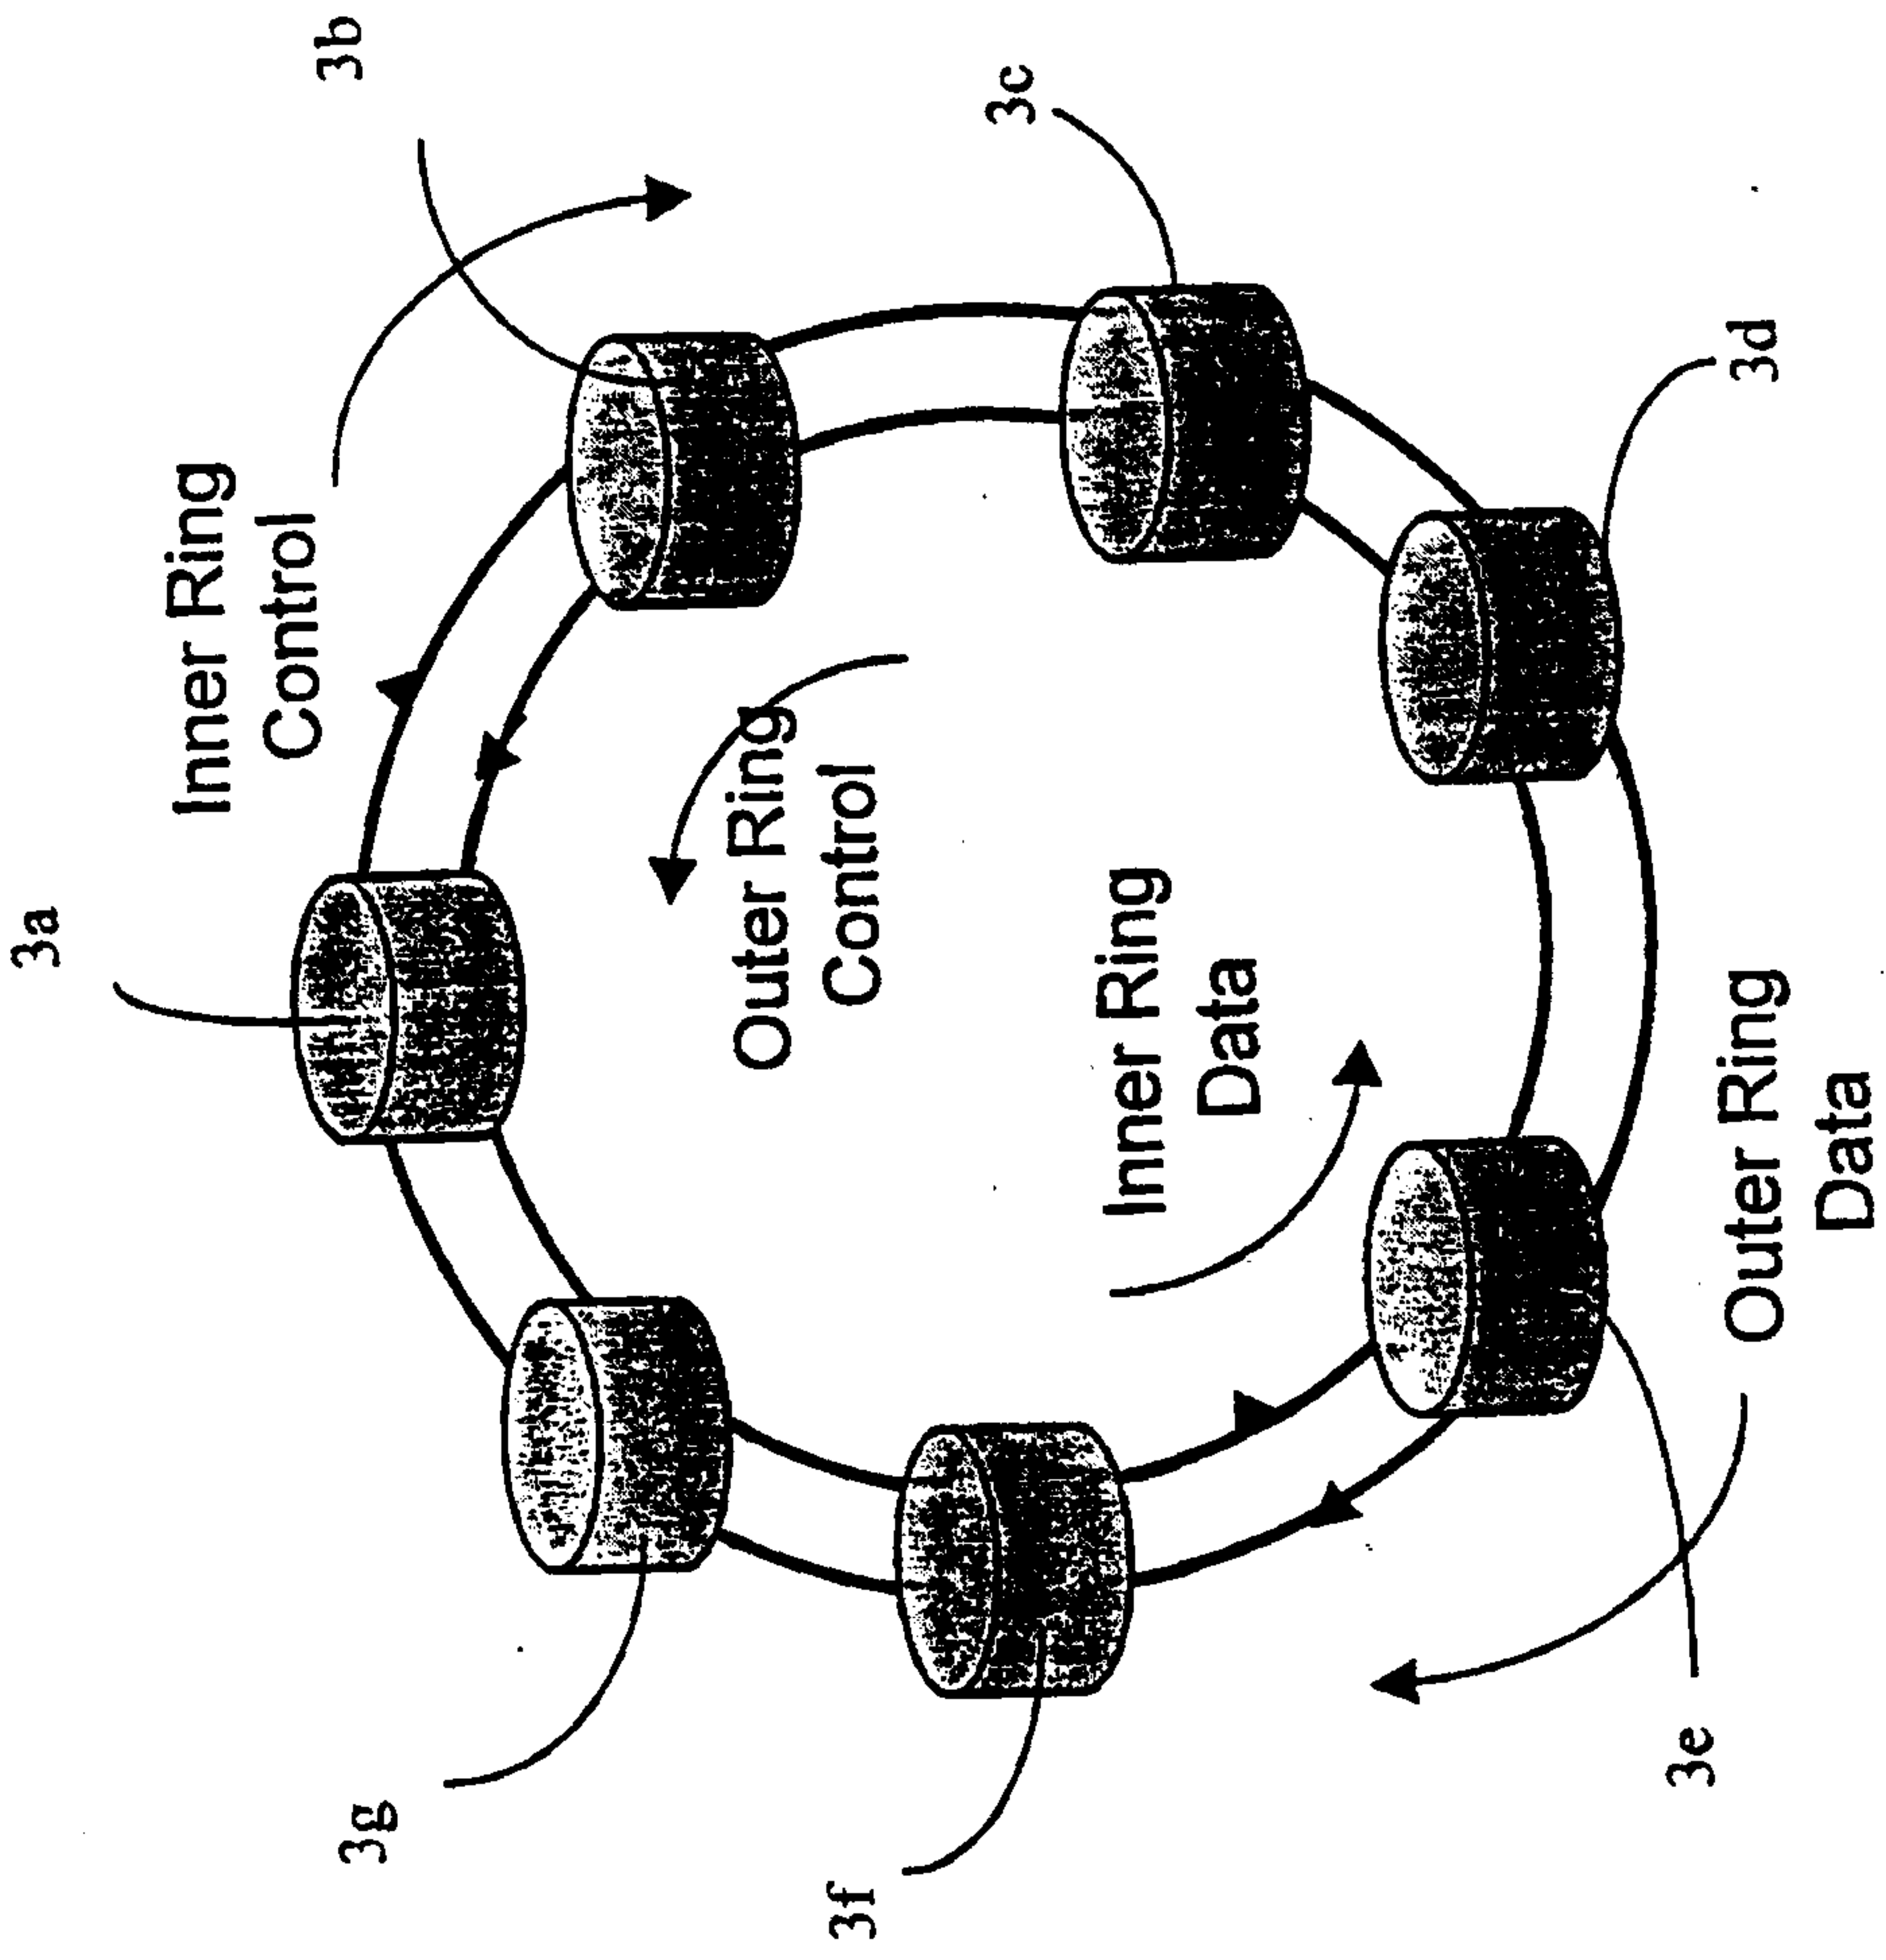
9. A method of operating a node in a communications network, the node having an input and an output with a continuous flow transmission path therebetween, the continuous flow transmission path including a transmission medium, the method including the steps of:

30

- (a) receiving an optical packet at the node;
- 35 (b) copying the packet such that there is an original packet and a copy packet, the original and the copy packets being duplicates of one another;

-19-

- 5
- (c) copying the copy packet such that there is a first copy packet and a second copy packet;
  - (d) reading at least part of the first copy packet to determine if (i) the packet is a transit packet and if (ii) the packet is a high priority packet;
  - (e) if the first copy packet is determined to be a transit packet having a high priority, then outputting the original packet from the node via the continuous flow transmission path; and,
  - (f) if the first copy packet is determined to be a transit packet not having a high priority, placing the second copy packet in a queue for later transmission and ejecting the original packet from the transmission medium.
- 10
10. A method of operating a node as claimed in claim 9, including the further steps of:
- 15
- (g) if the first copy packet is determined to be addressed to the node, ejecting the original packet from the transmission medium, and receiving the second copy packet for storage and/or processing at the node.
- 20



PRIOR ART

Figure 1

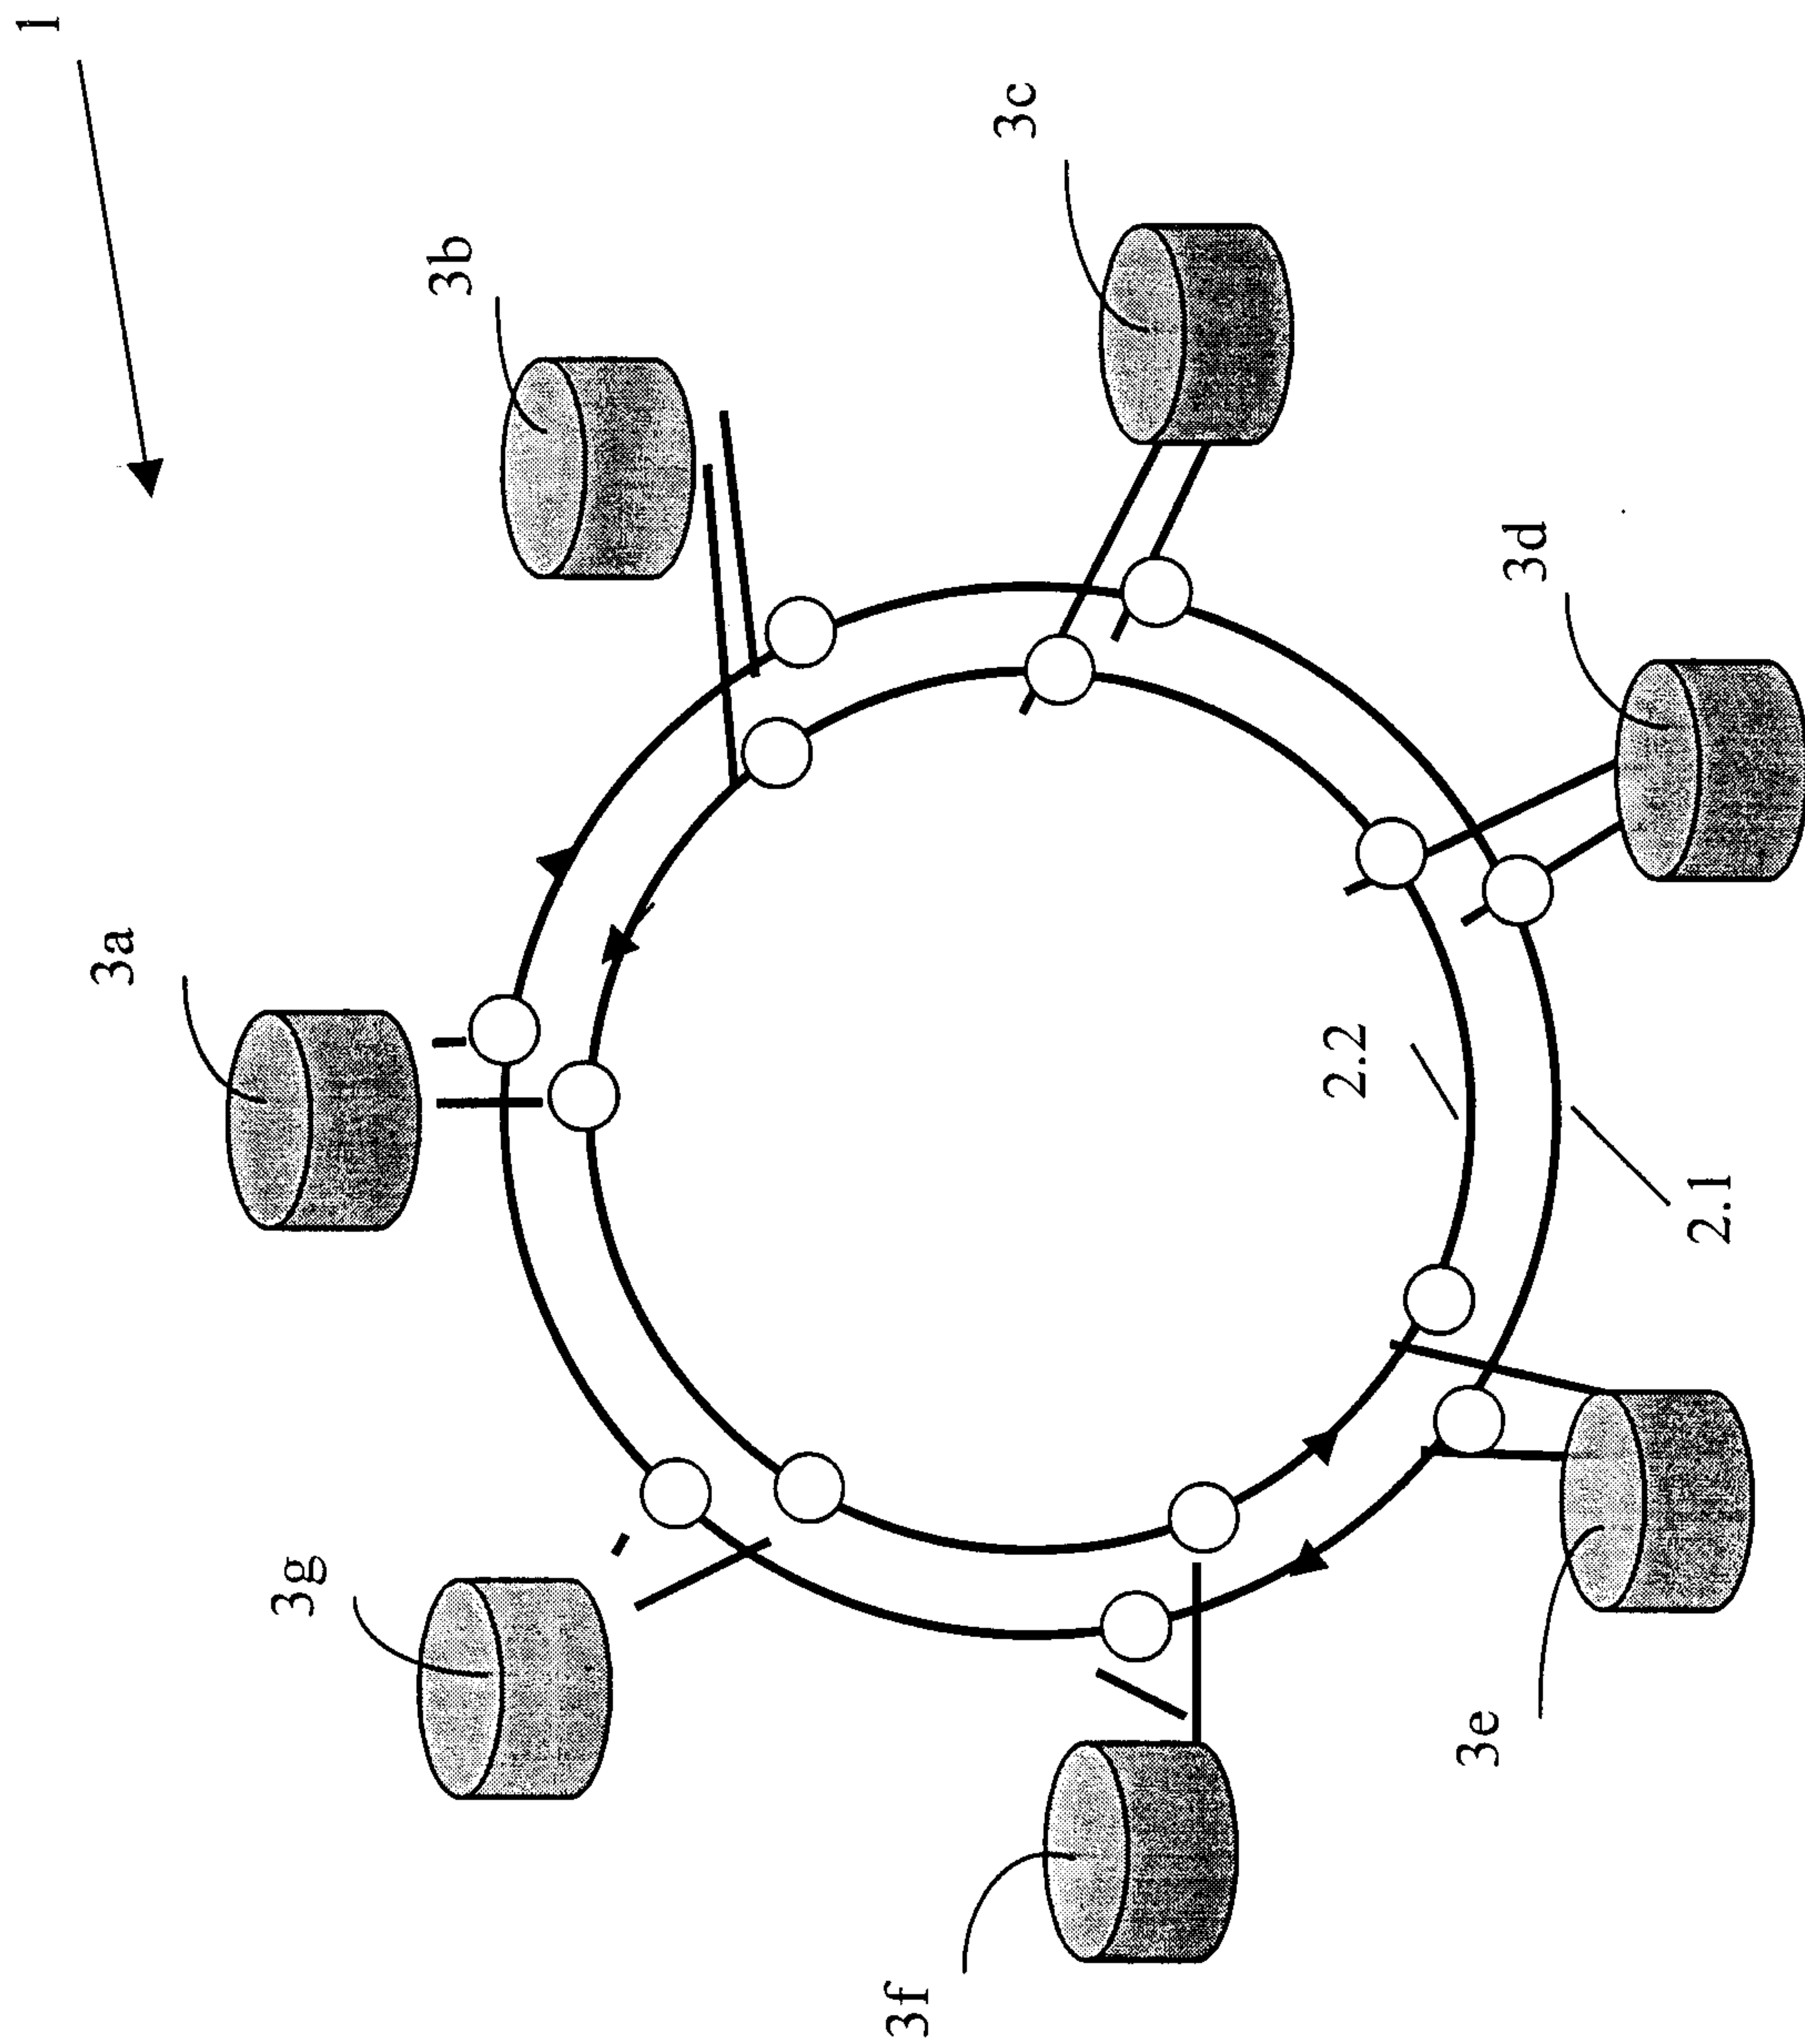


Figure 2

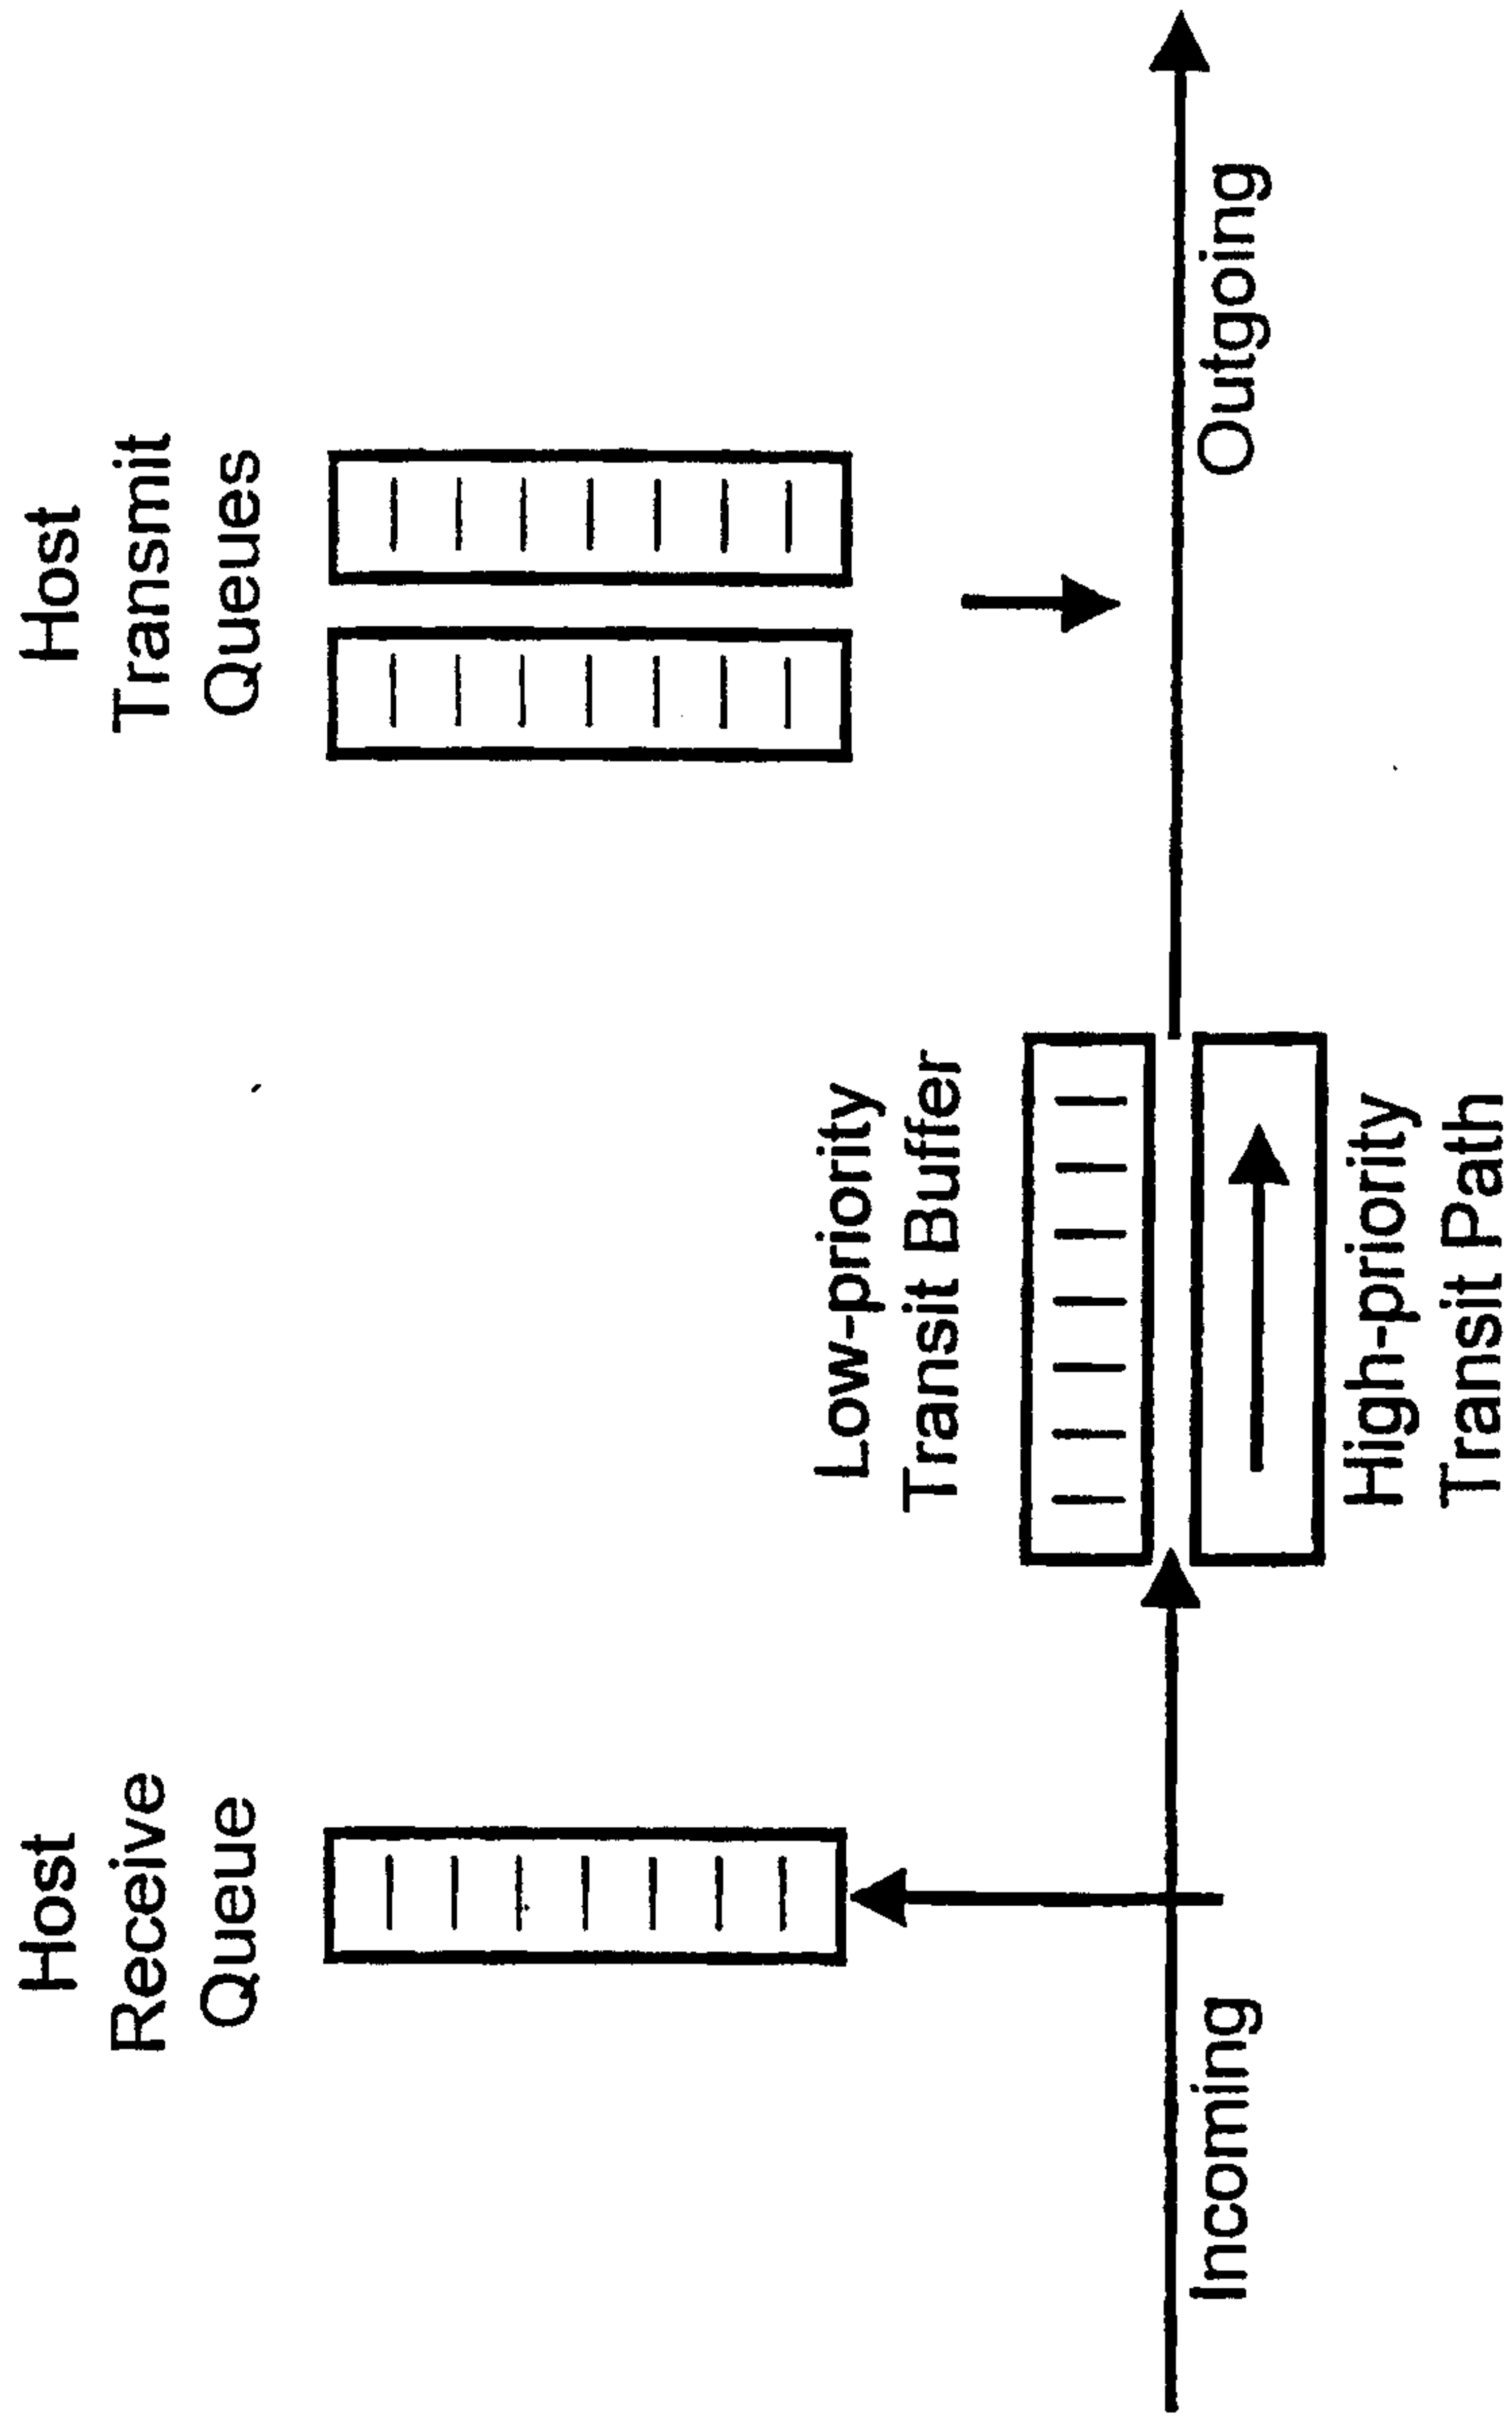


Figure 3

PRIOR ART

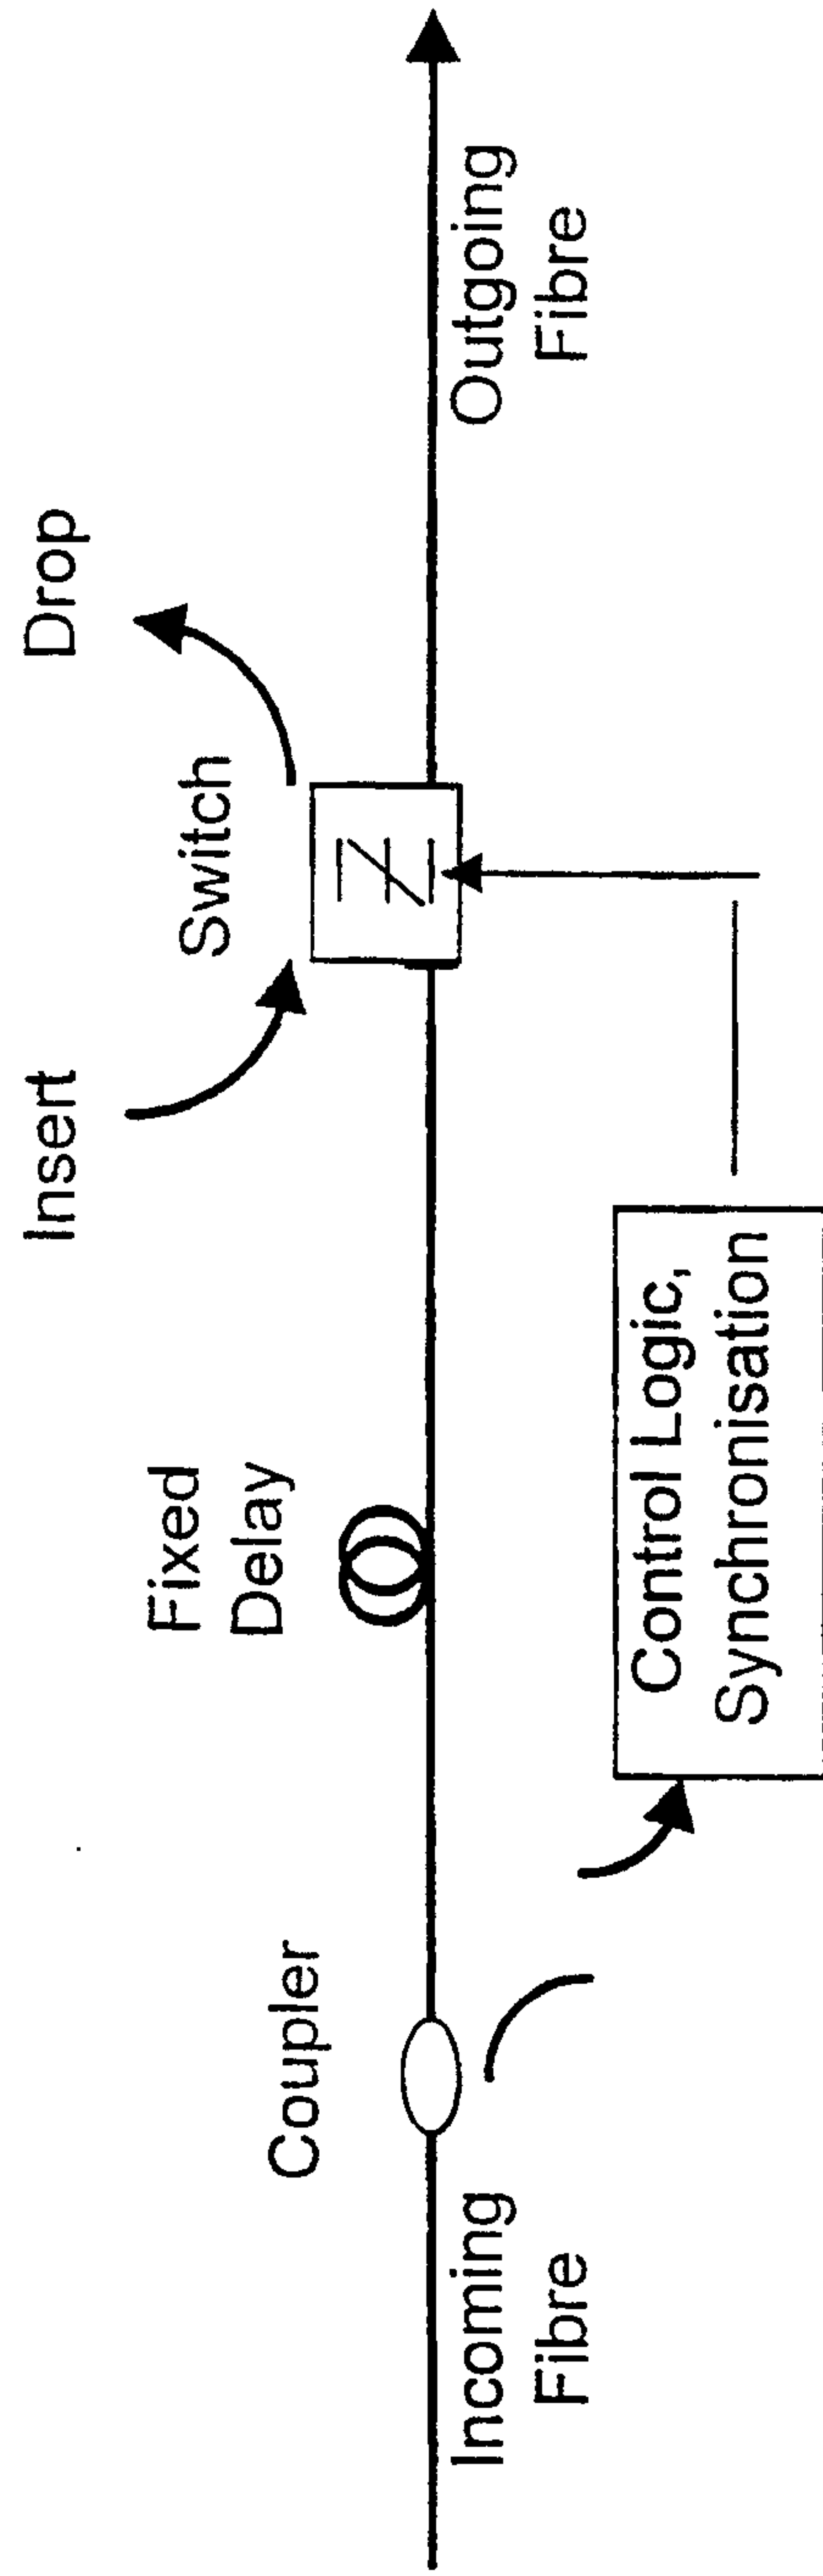


Figure 4

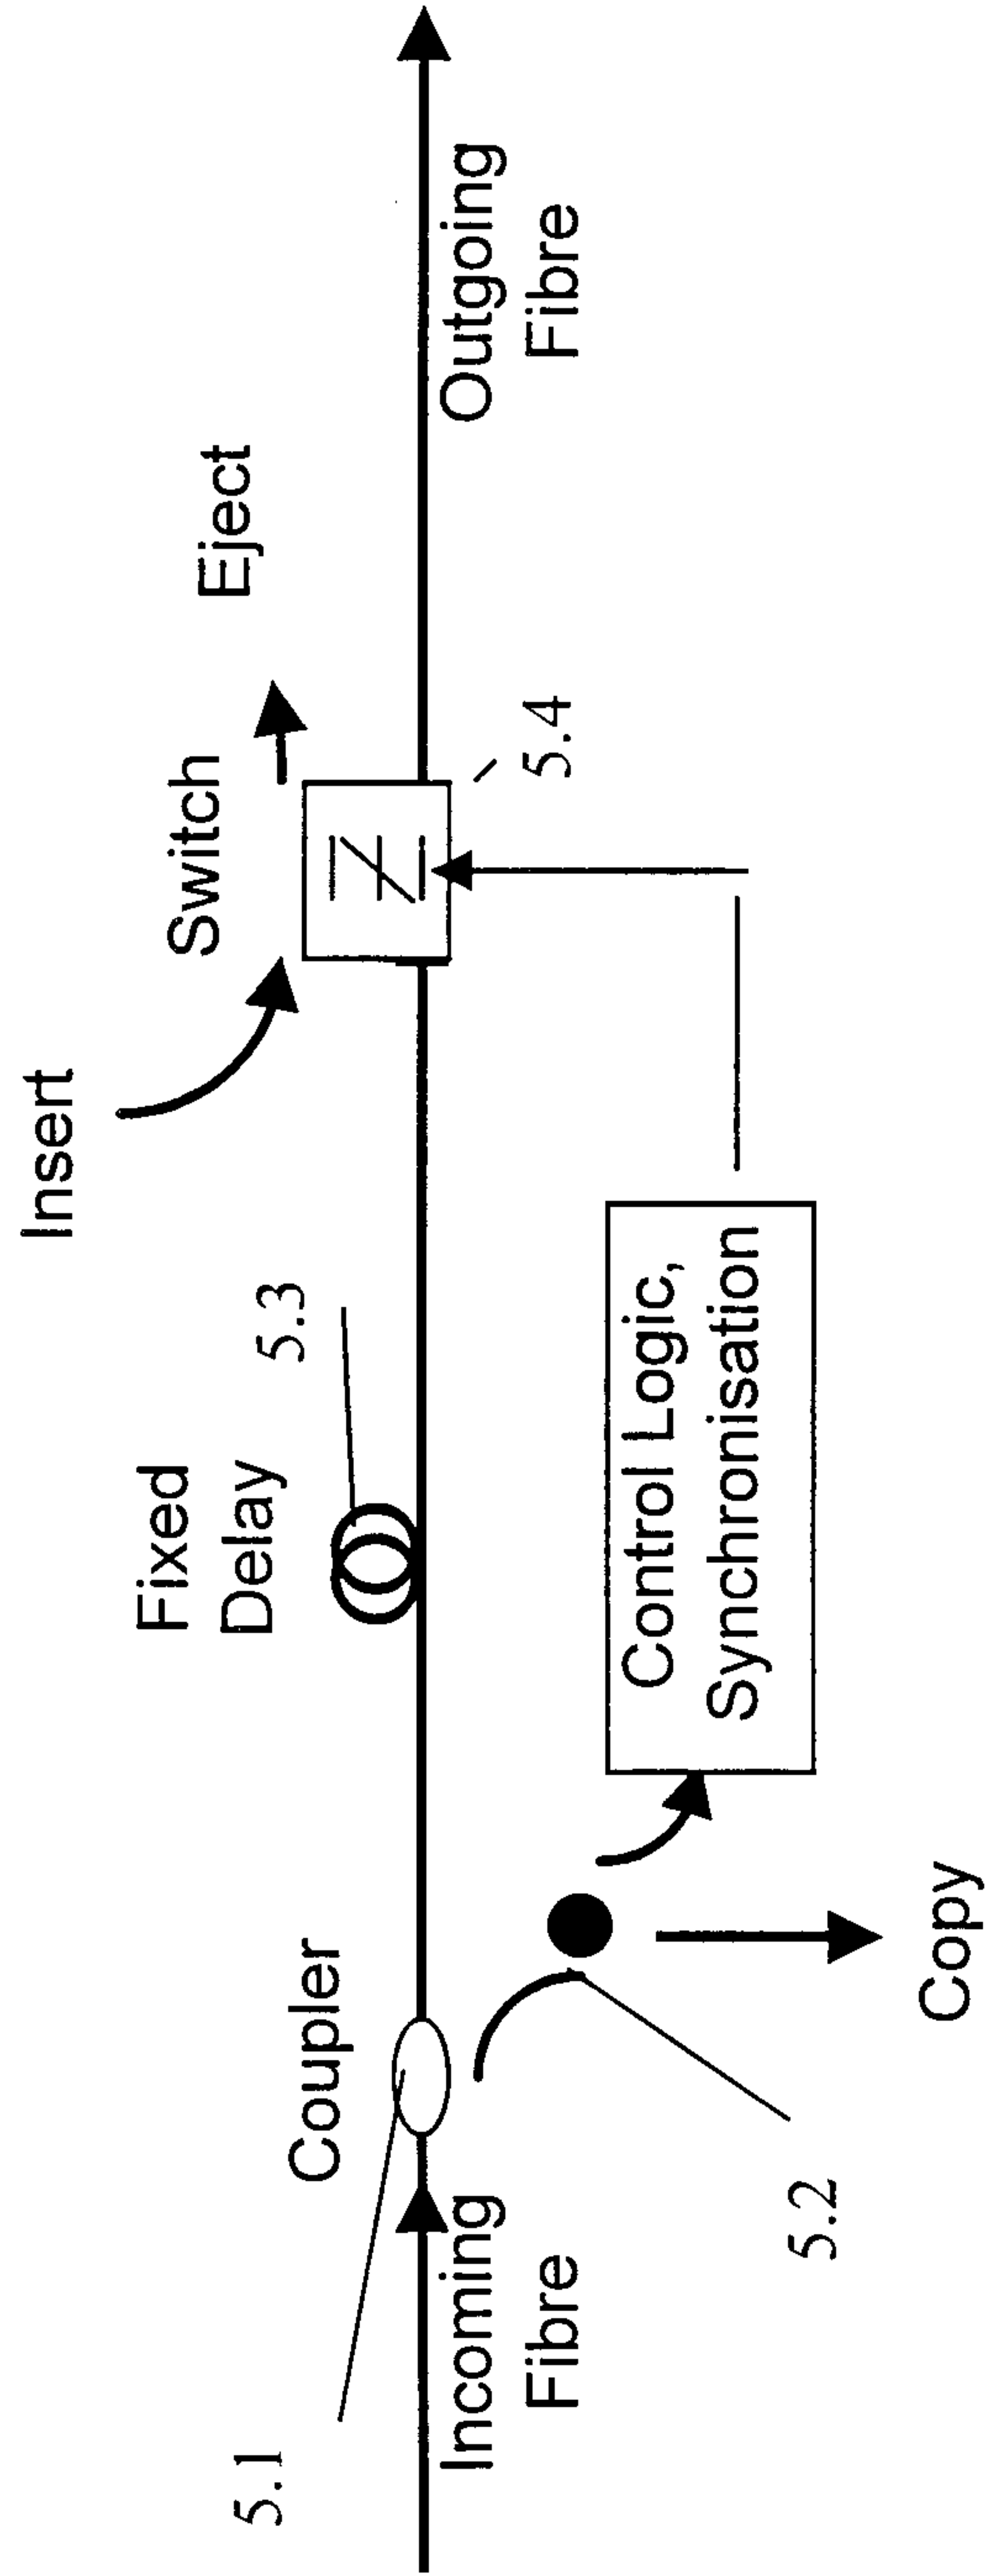


Figure 5

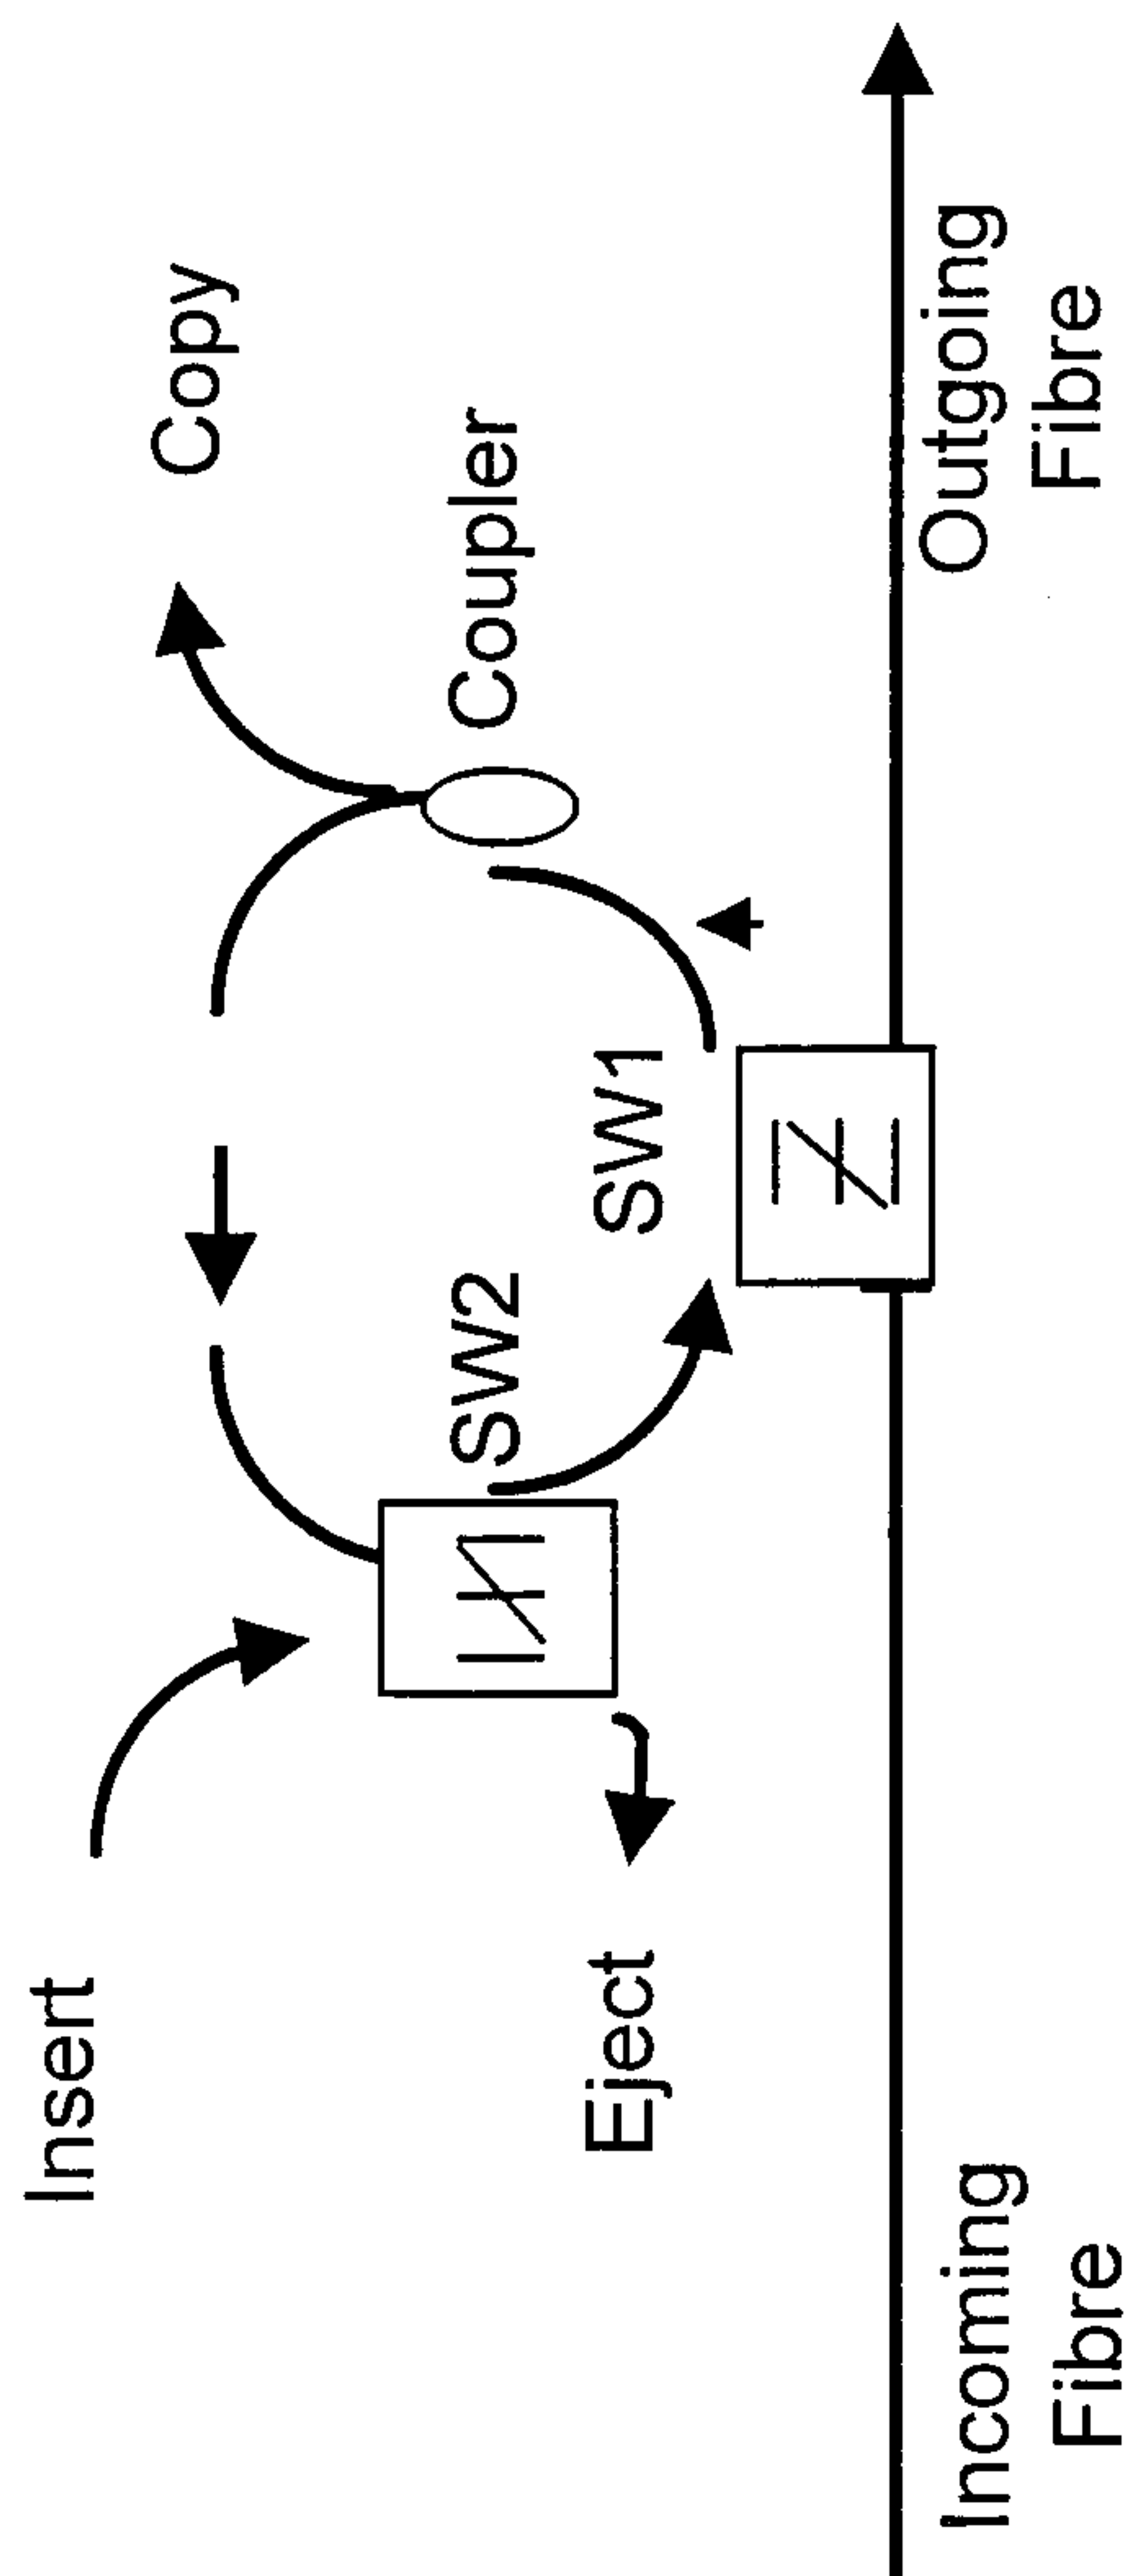


Figure 6

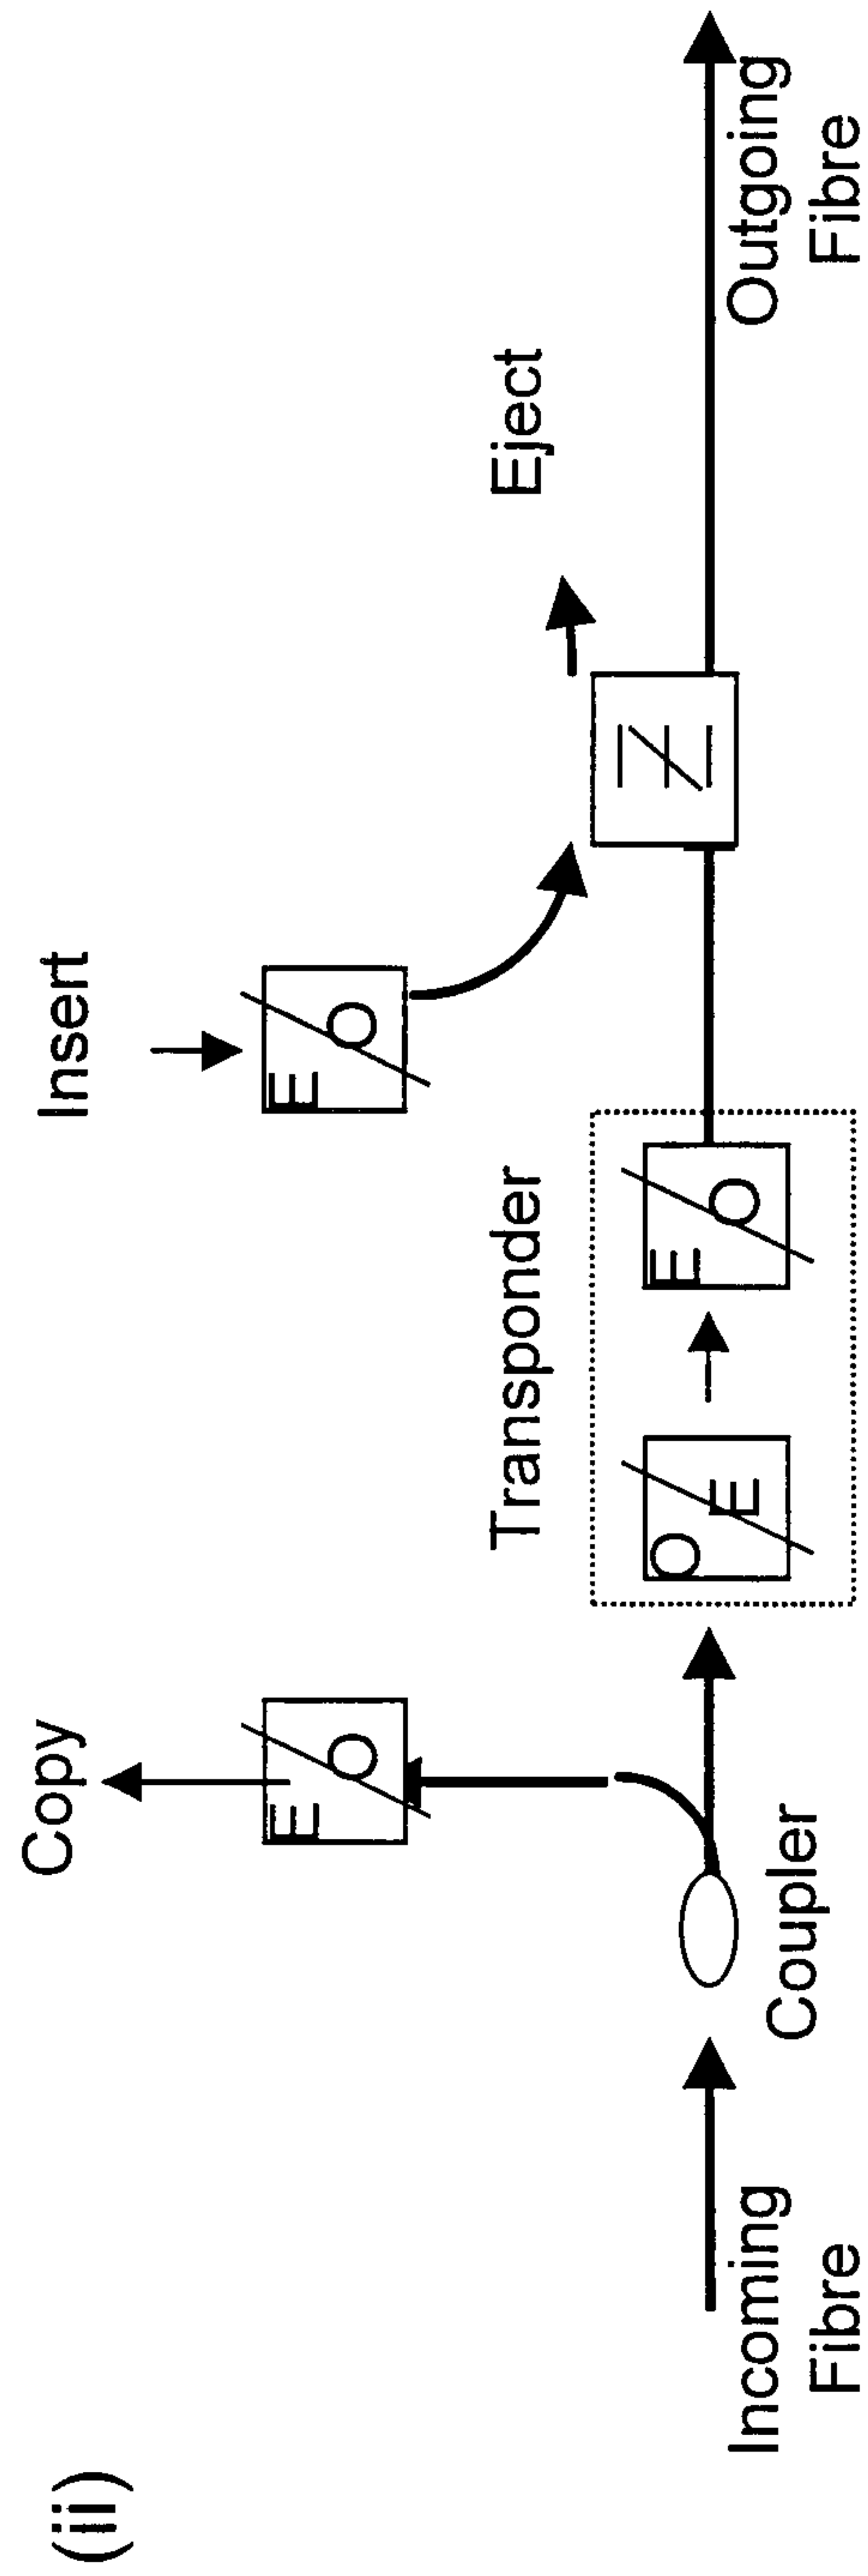
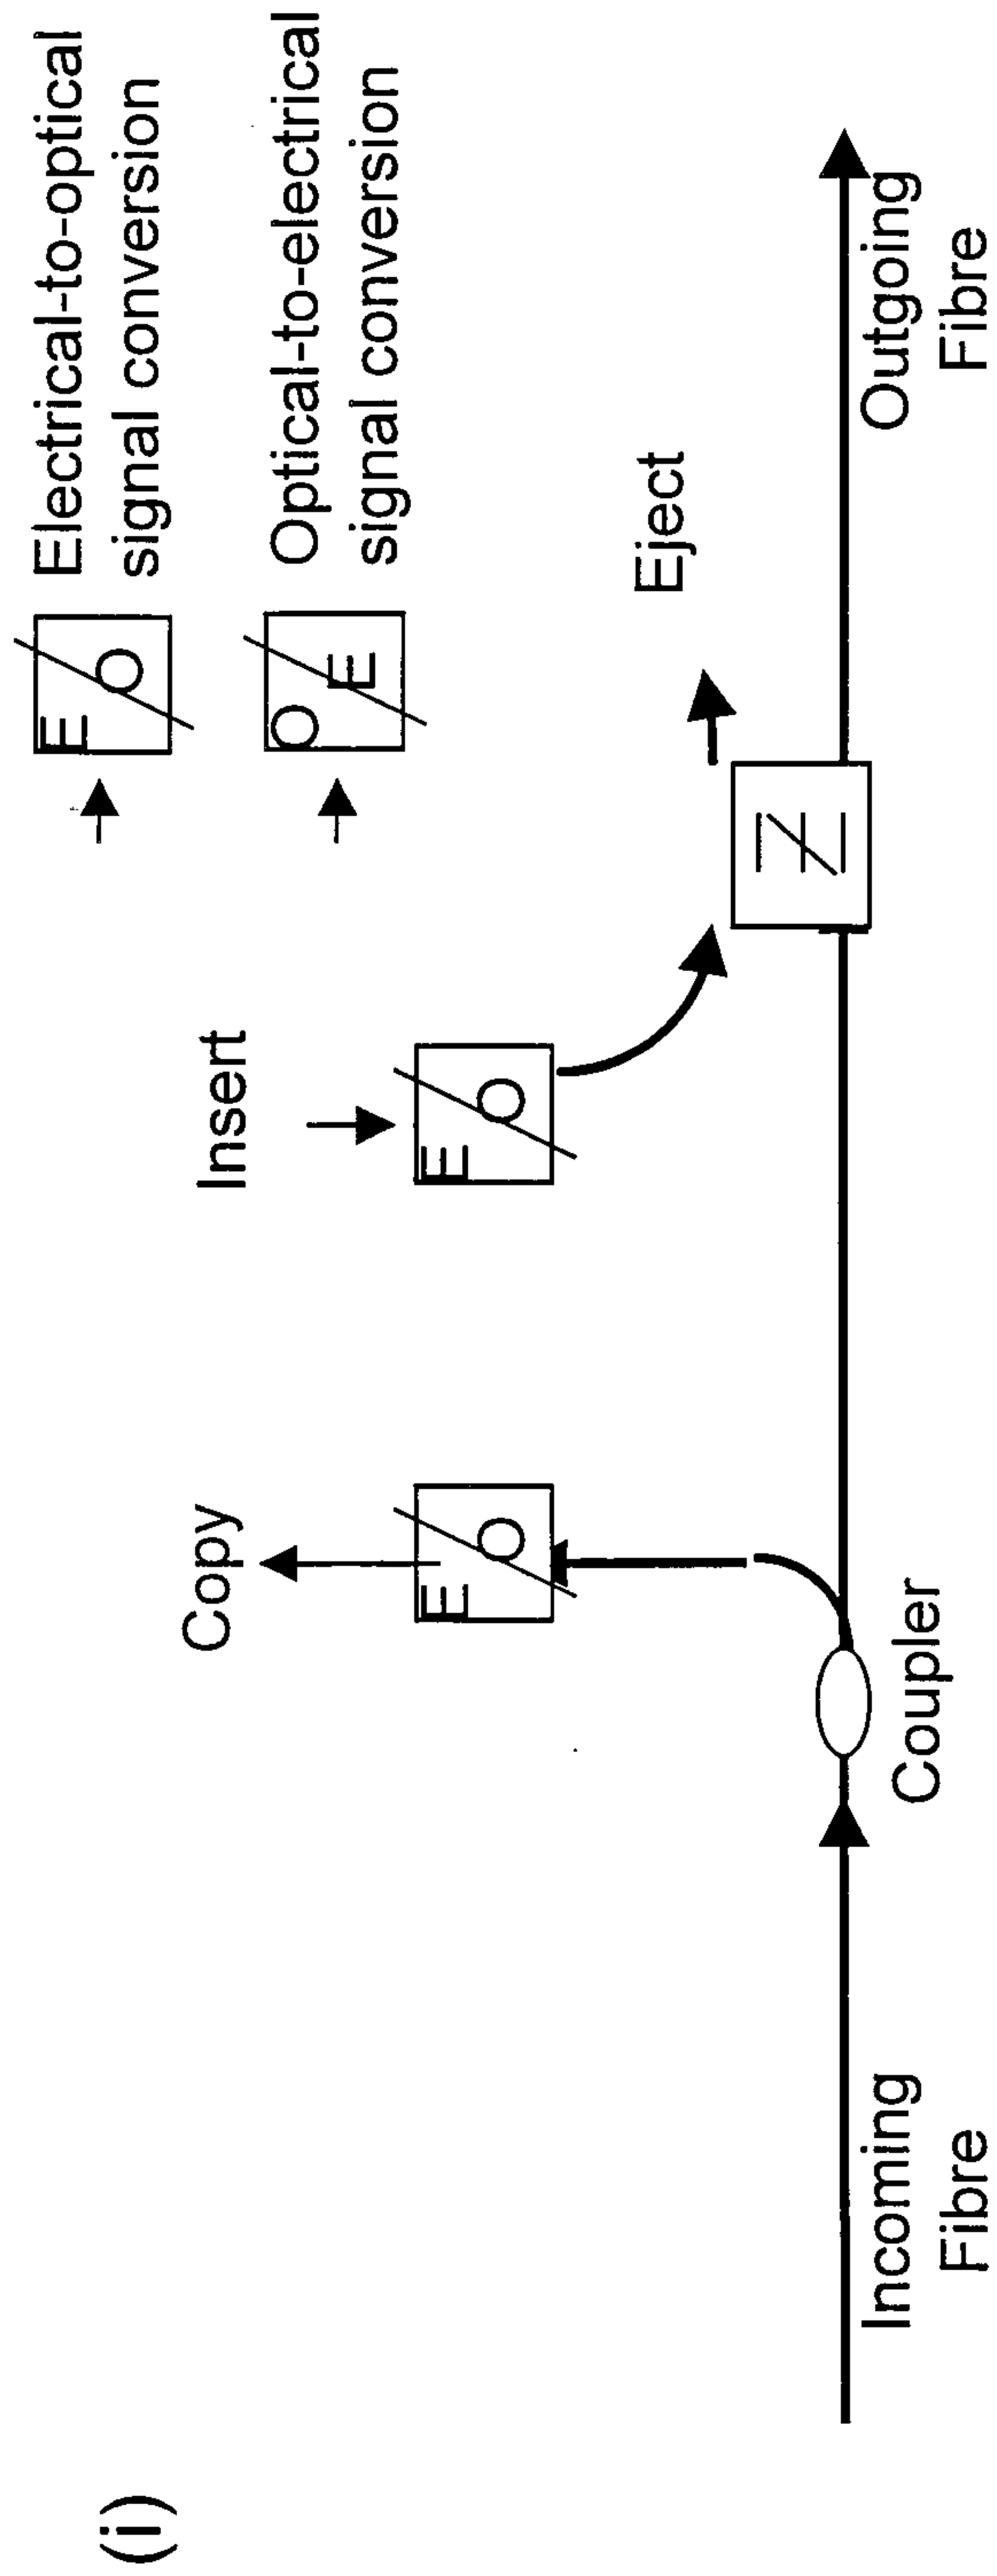
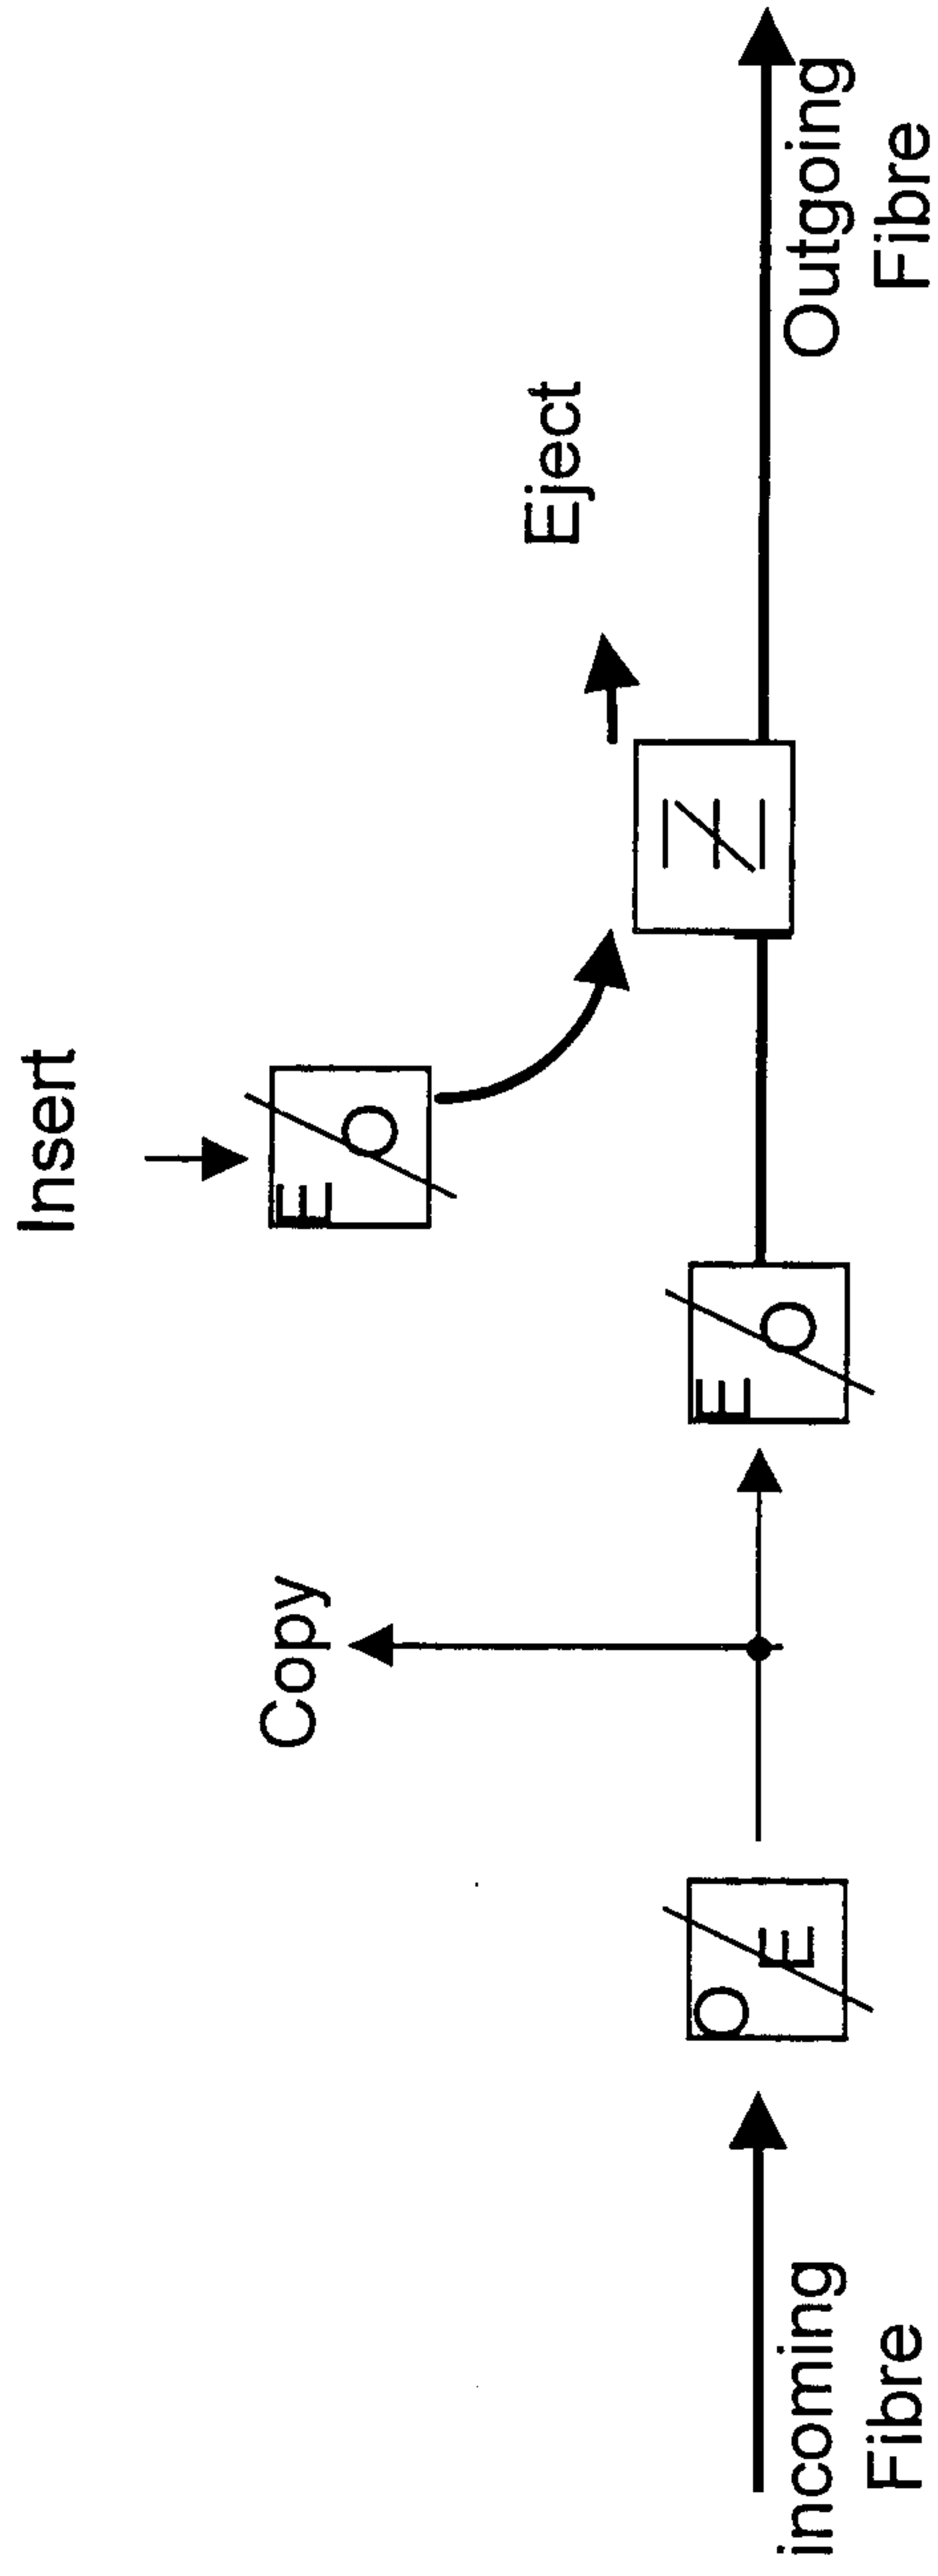


Figure 7(i), 7(ii)

(iii)



(iv)

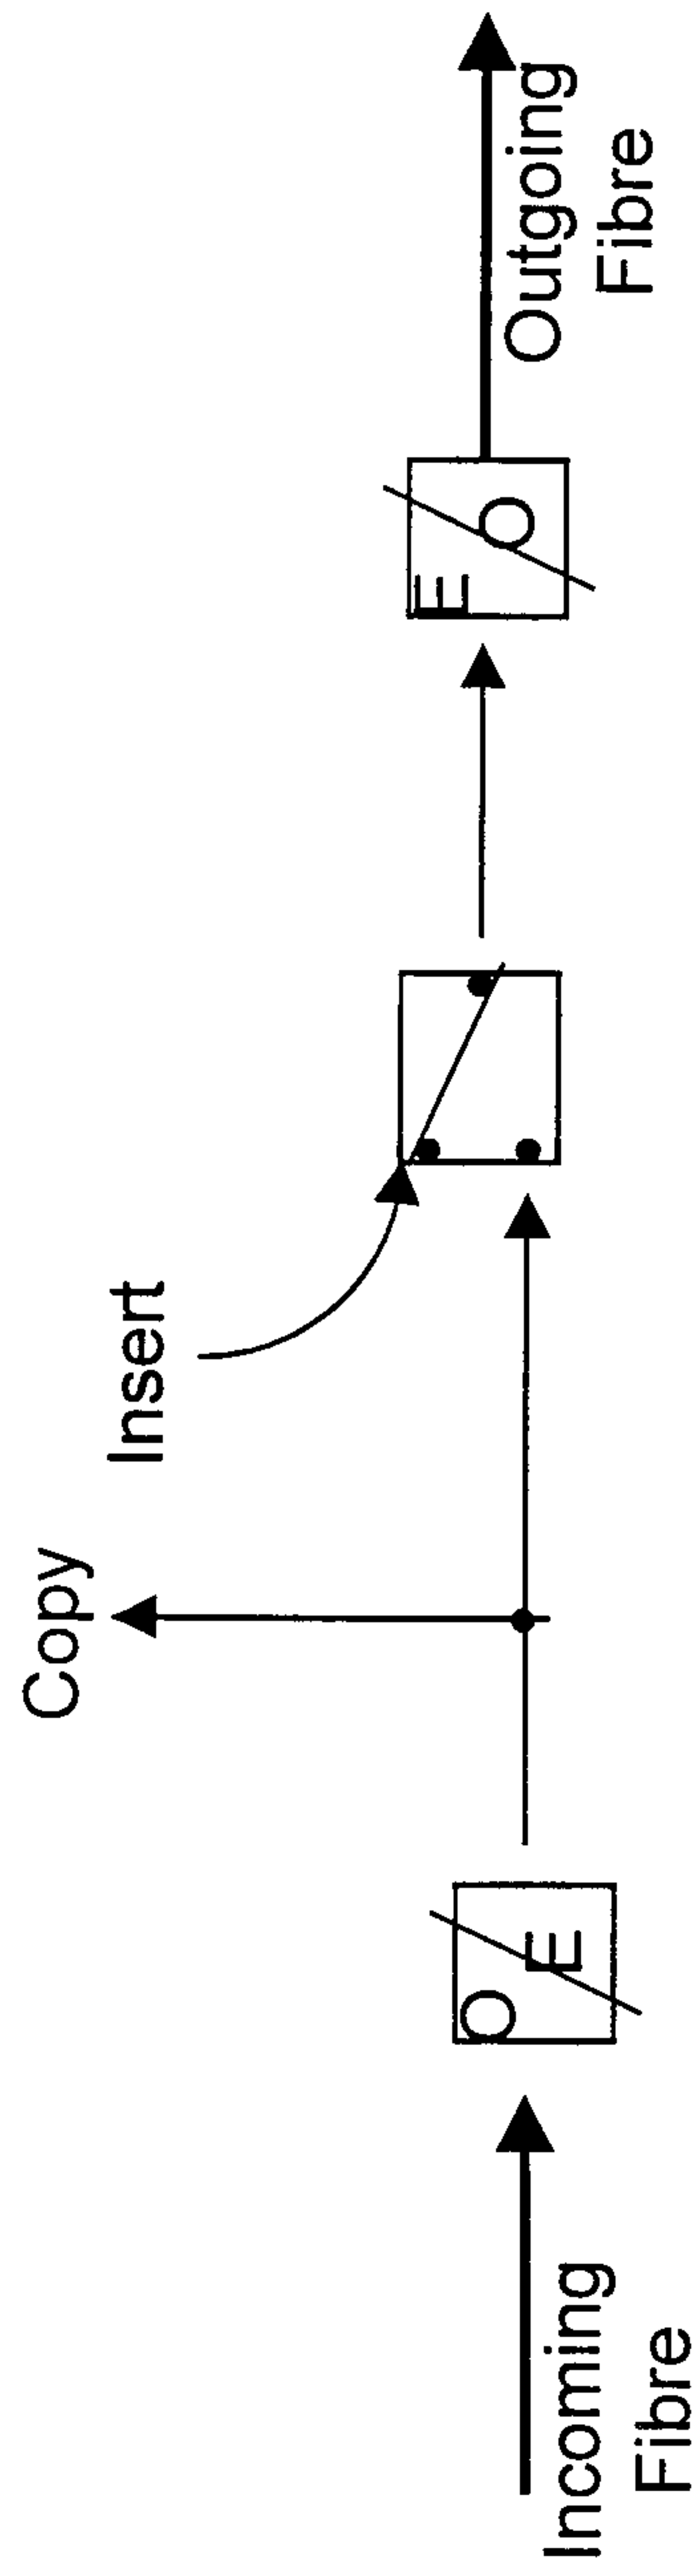
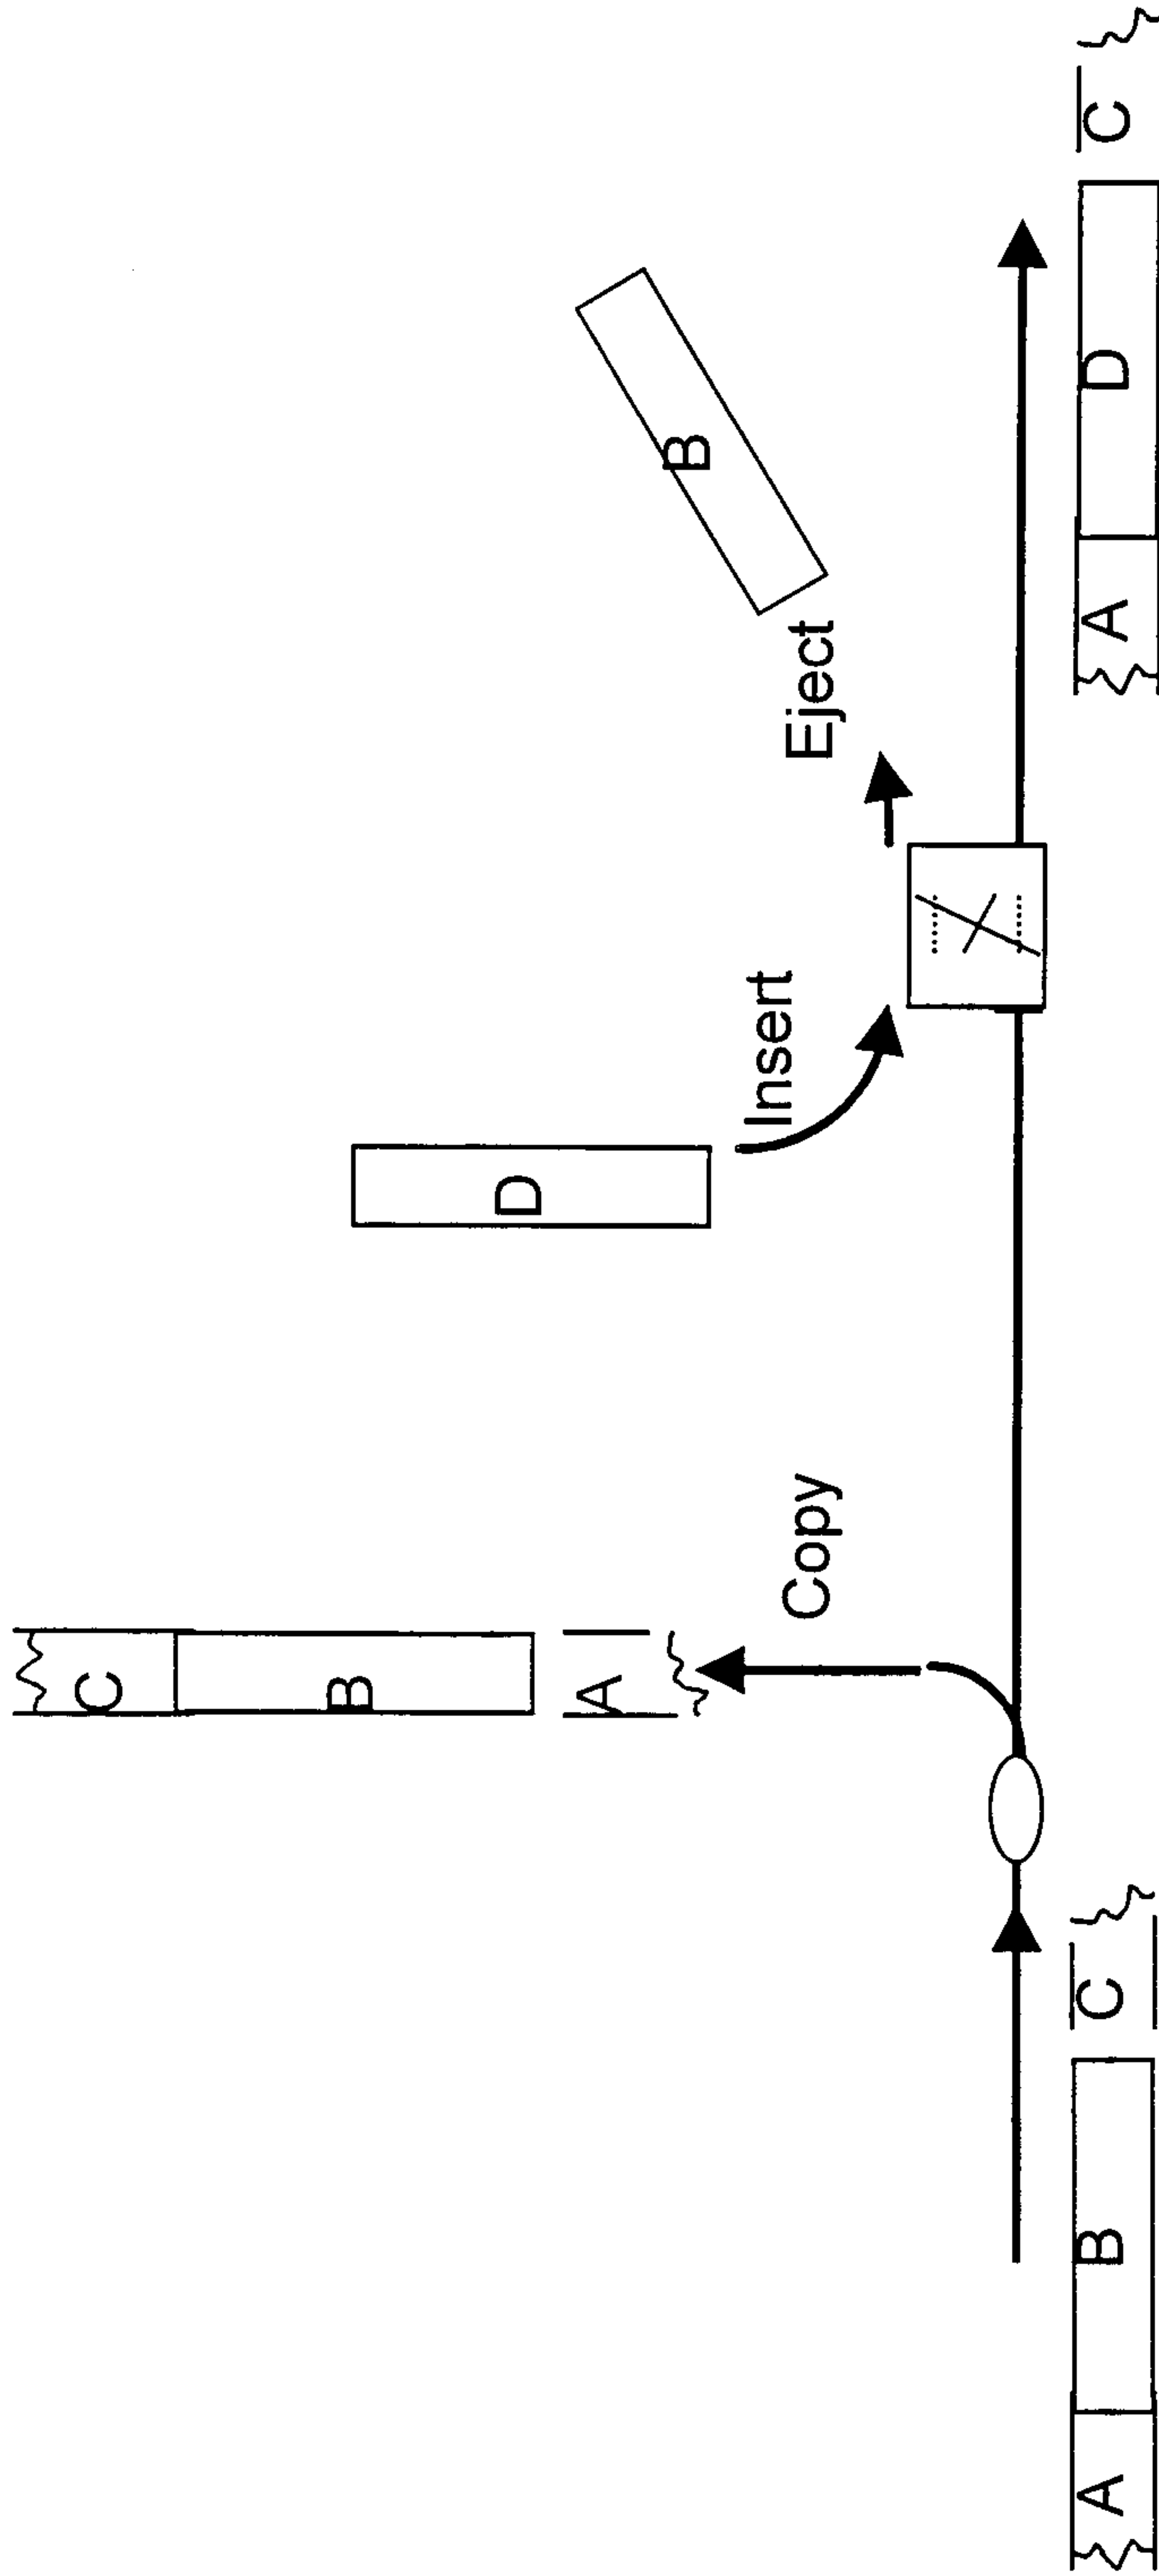


Figure 7(iii), 7(iv)



B = received packet or lo-priority transit packet  
D = packet from transmit queue or lo-priority transit buffer

Figure 8

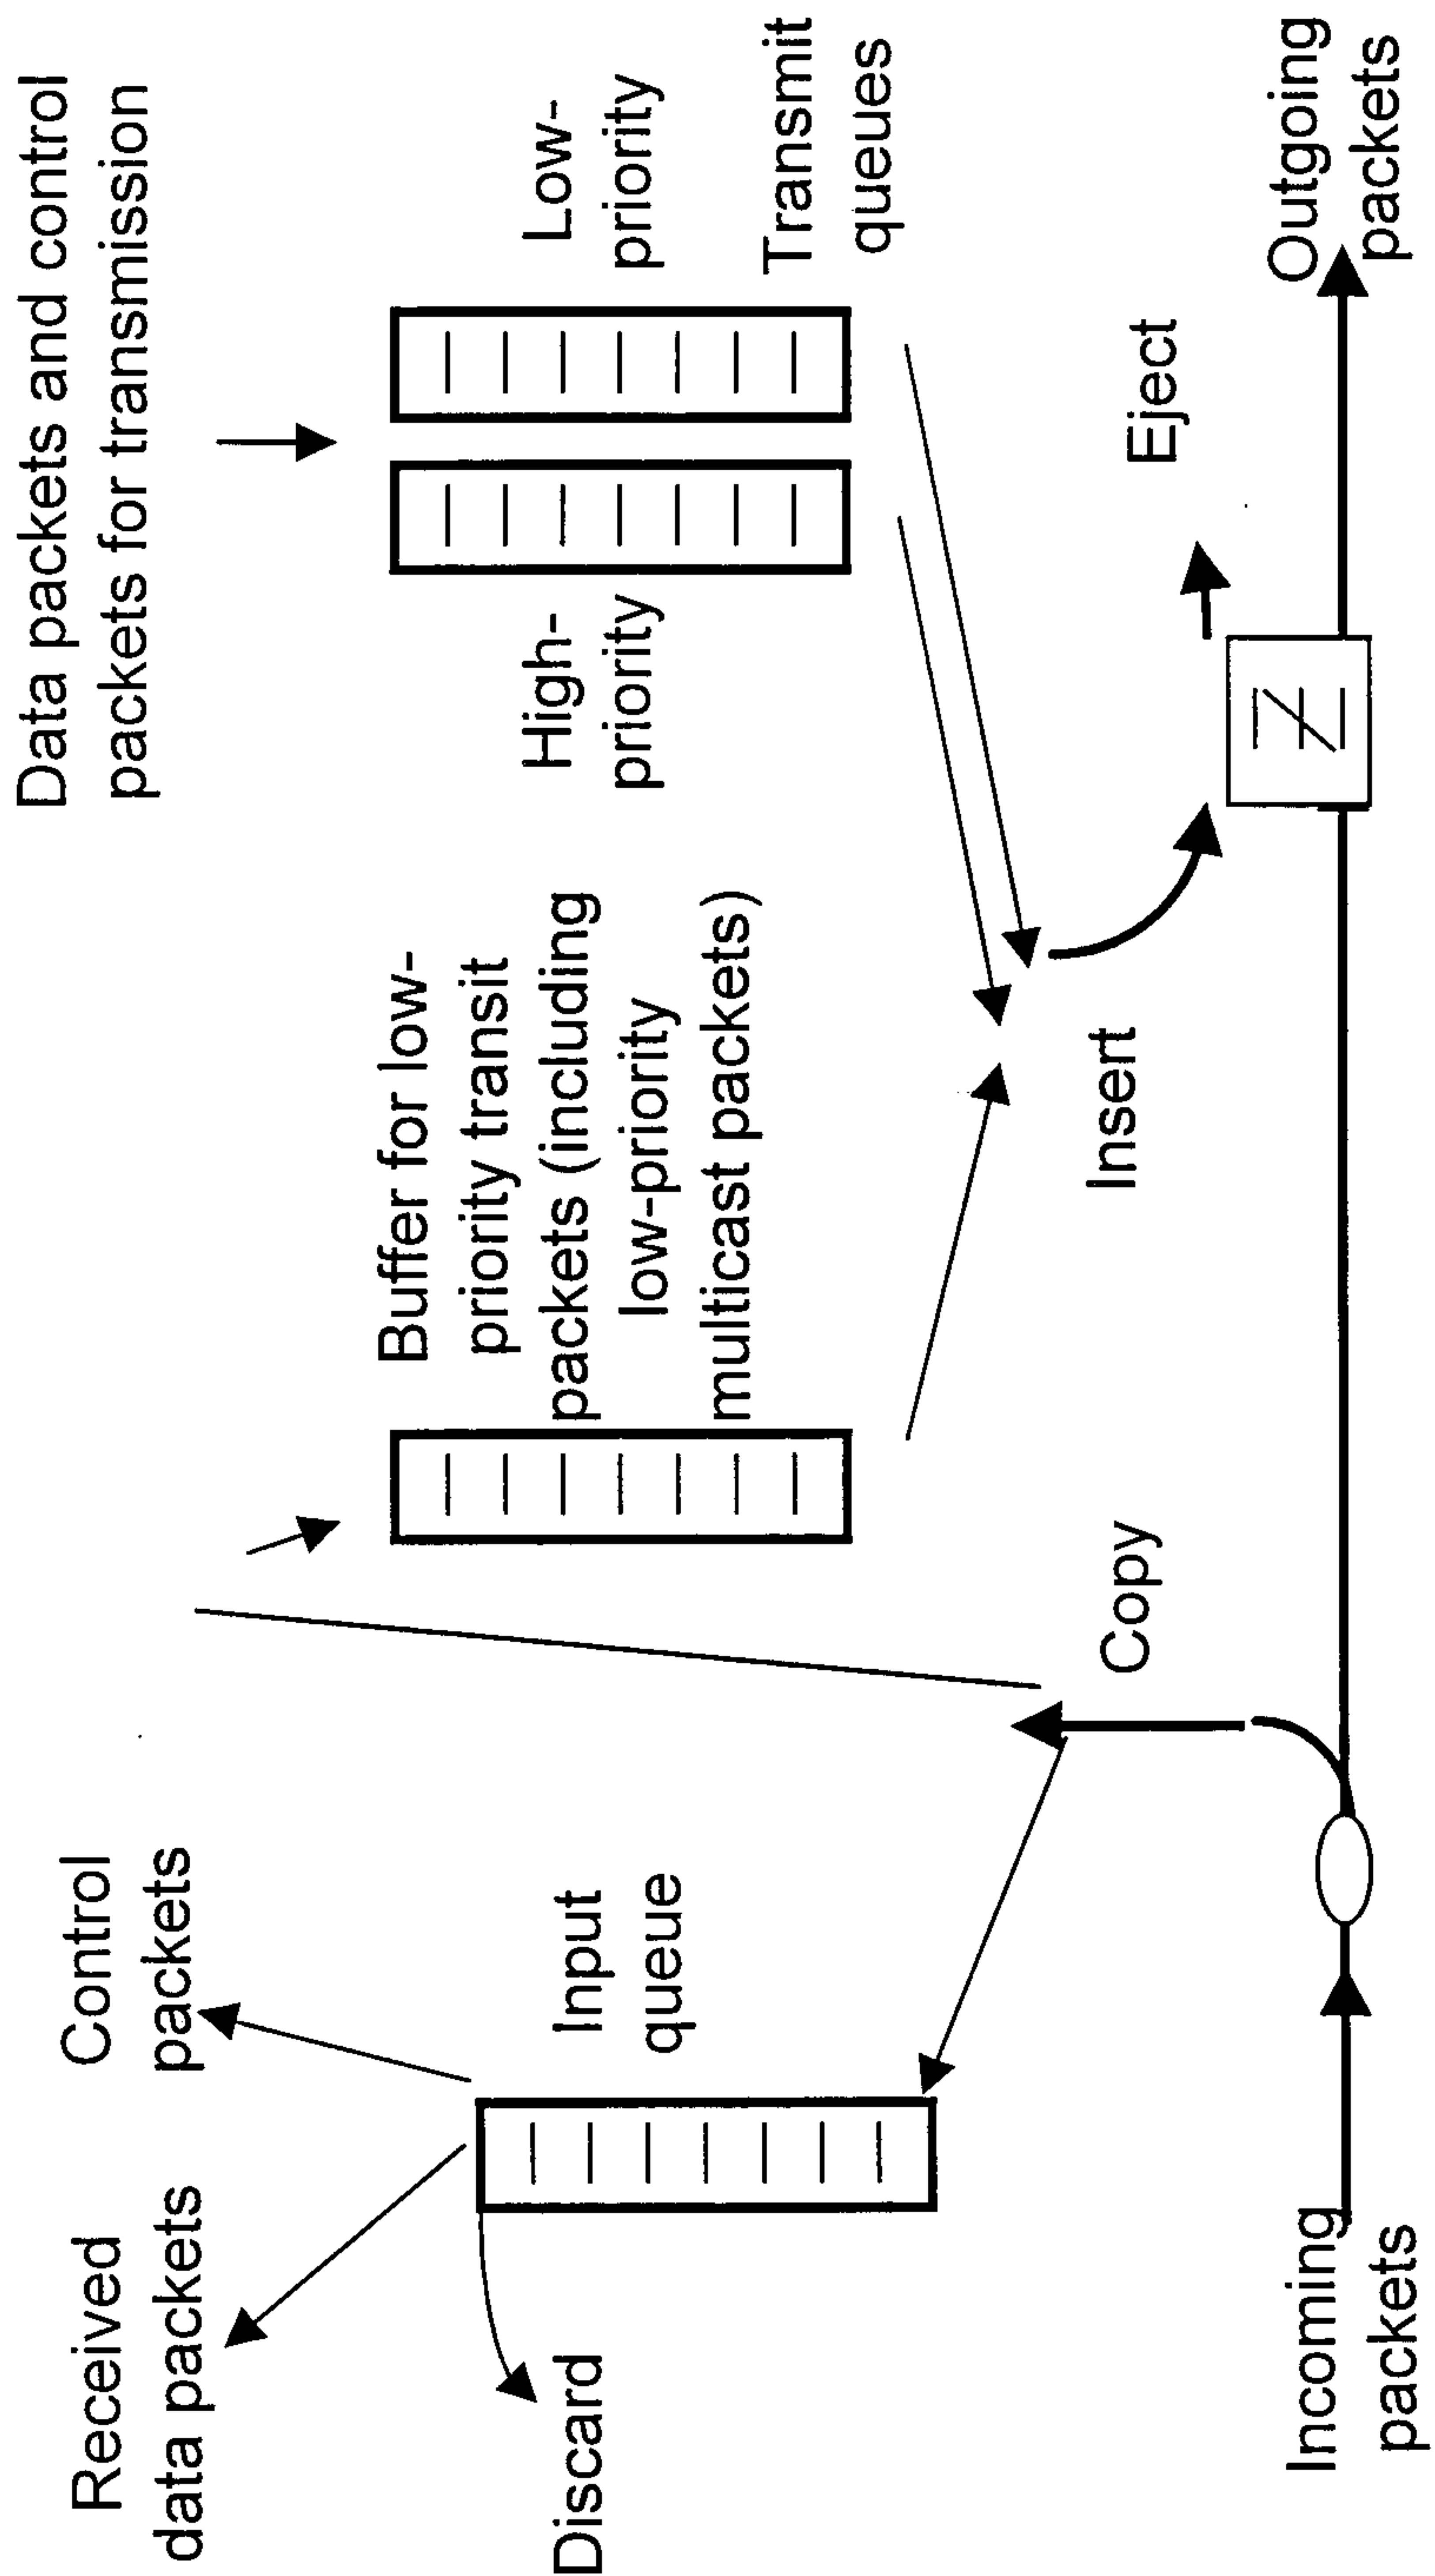


Figure 9

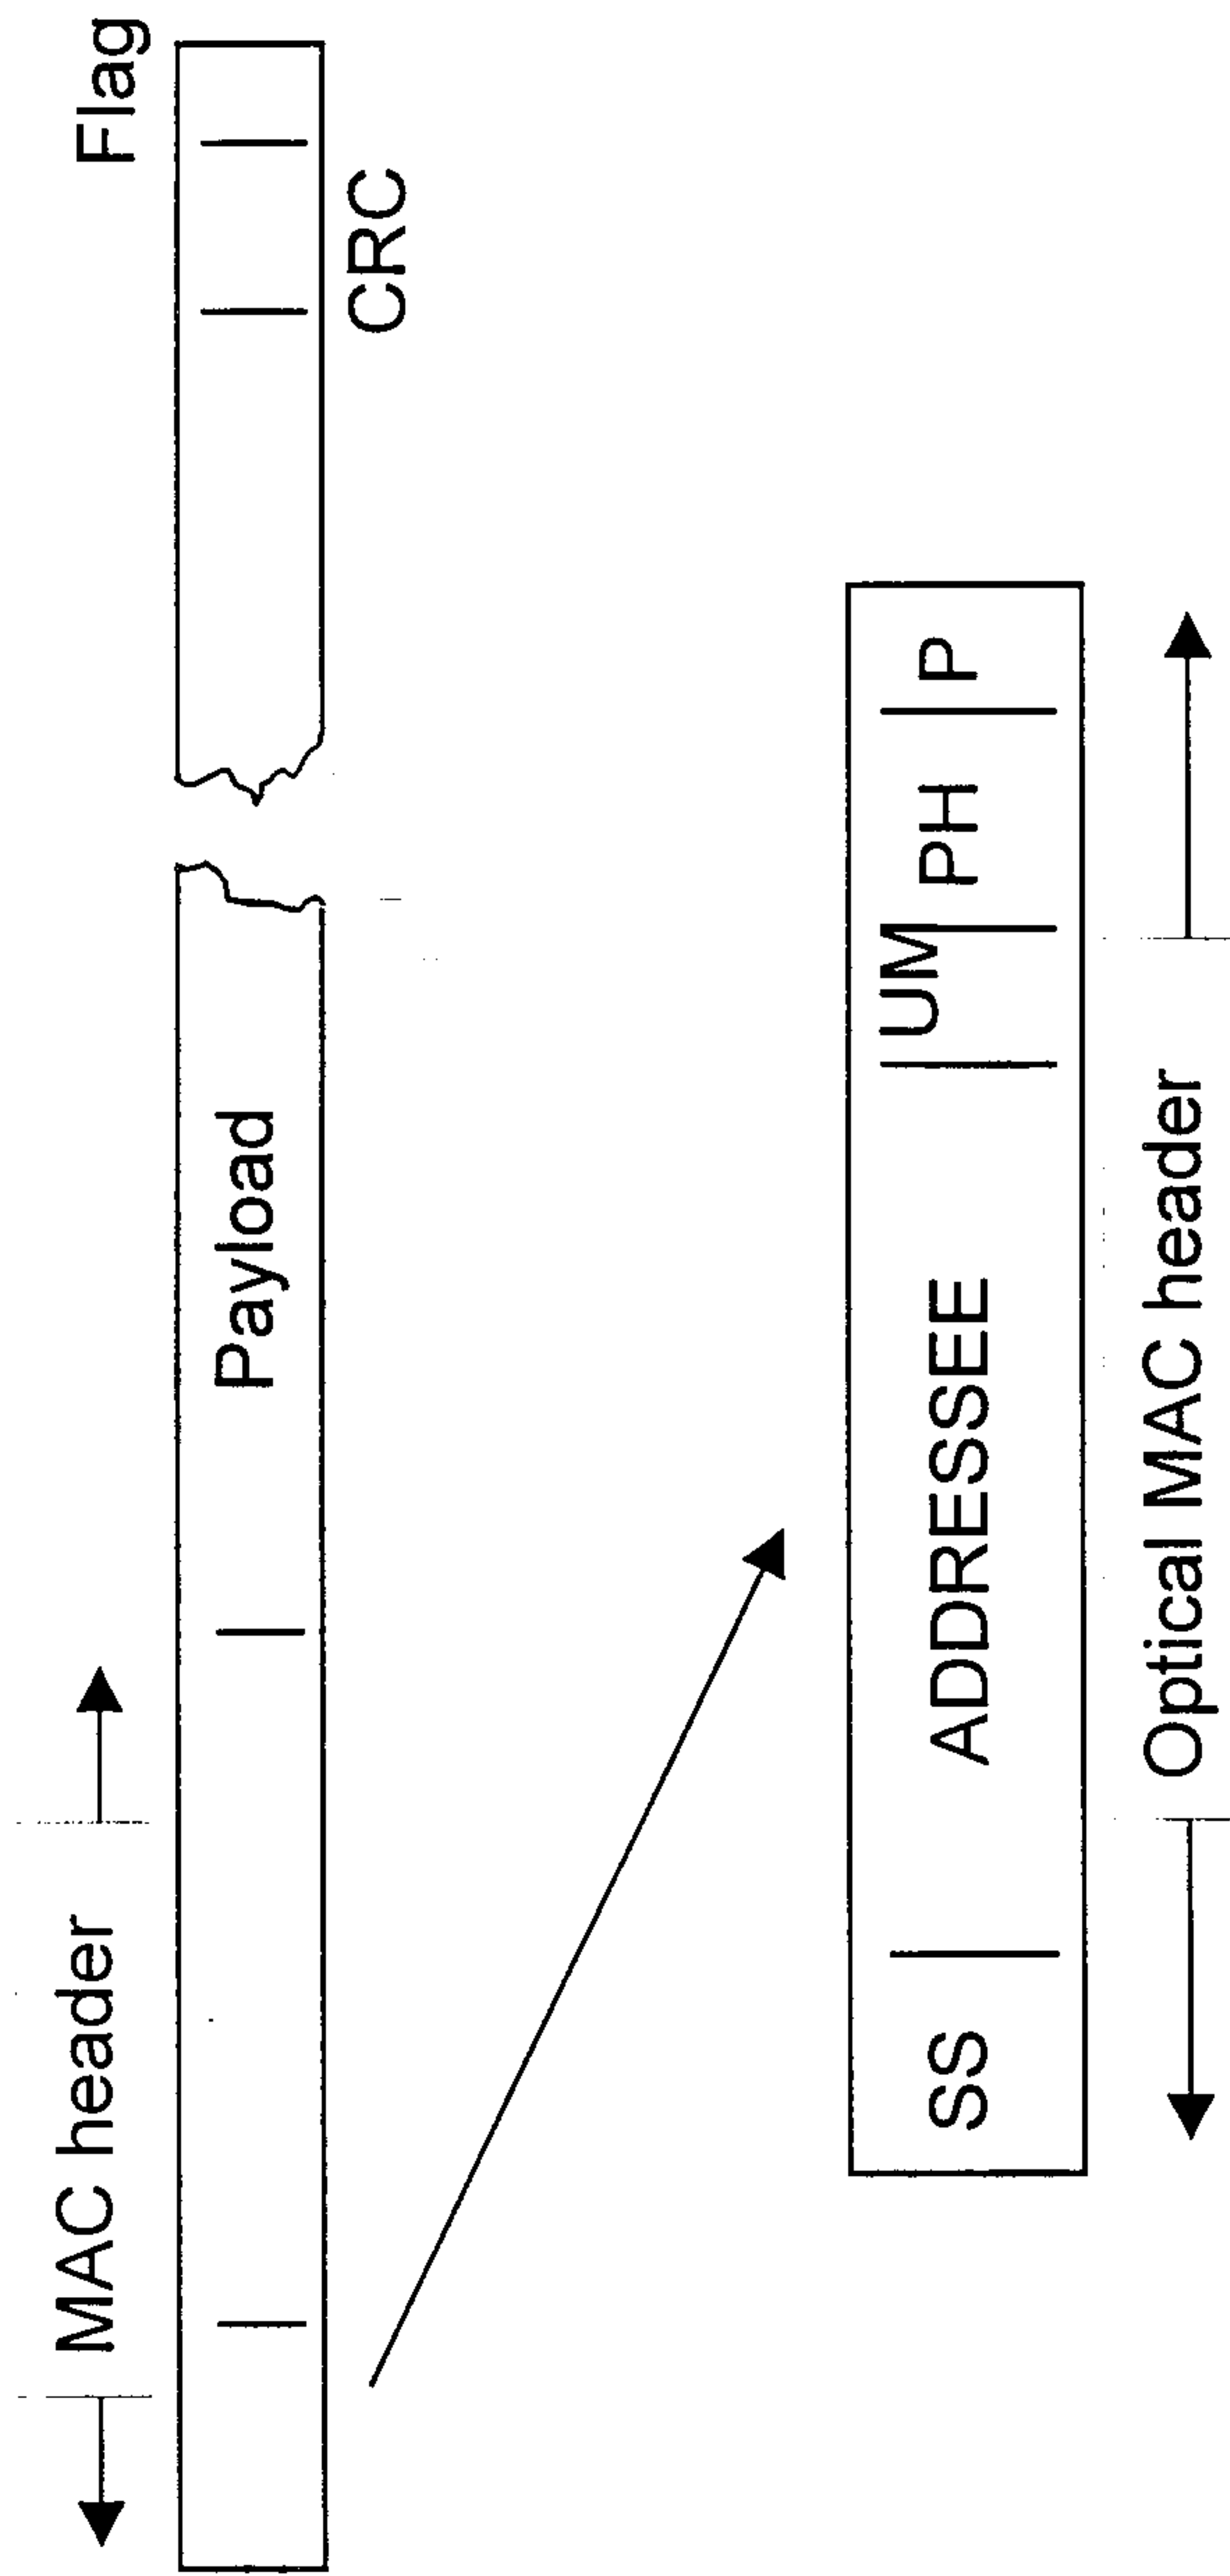


Figure 10

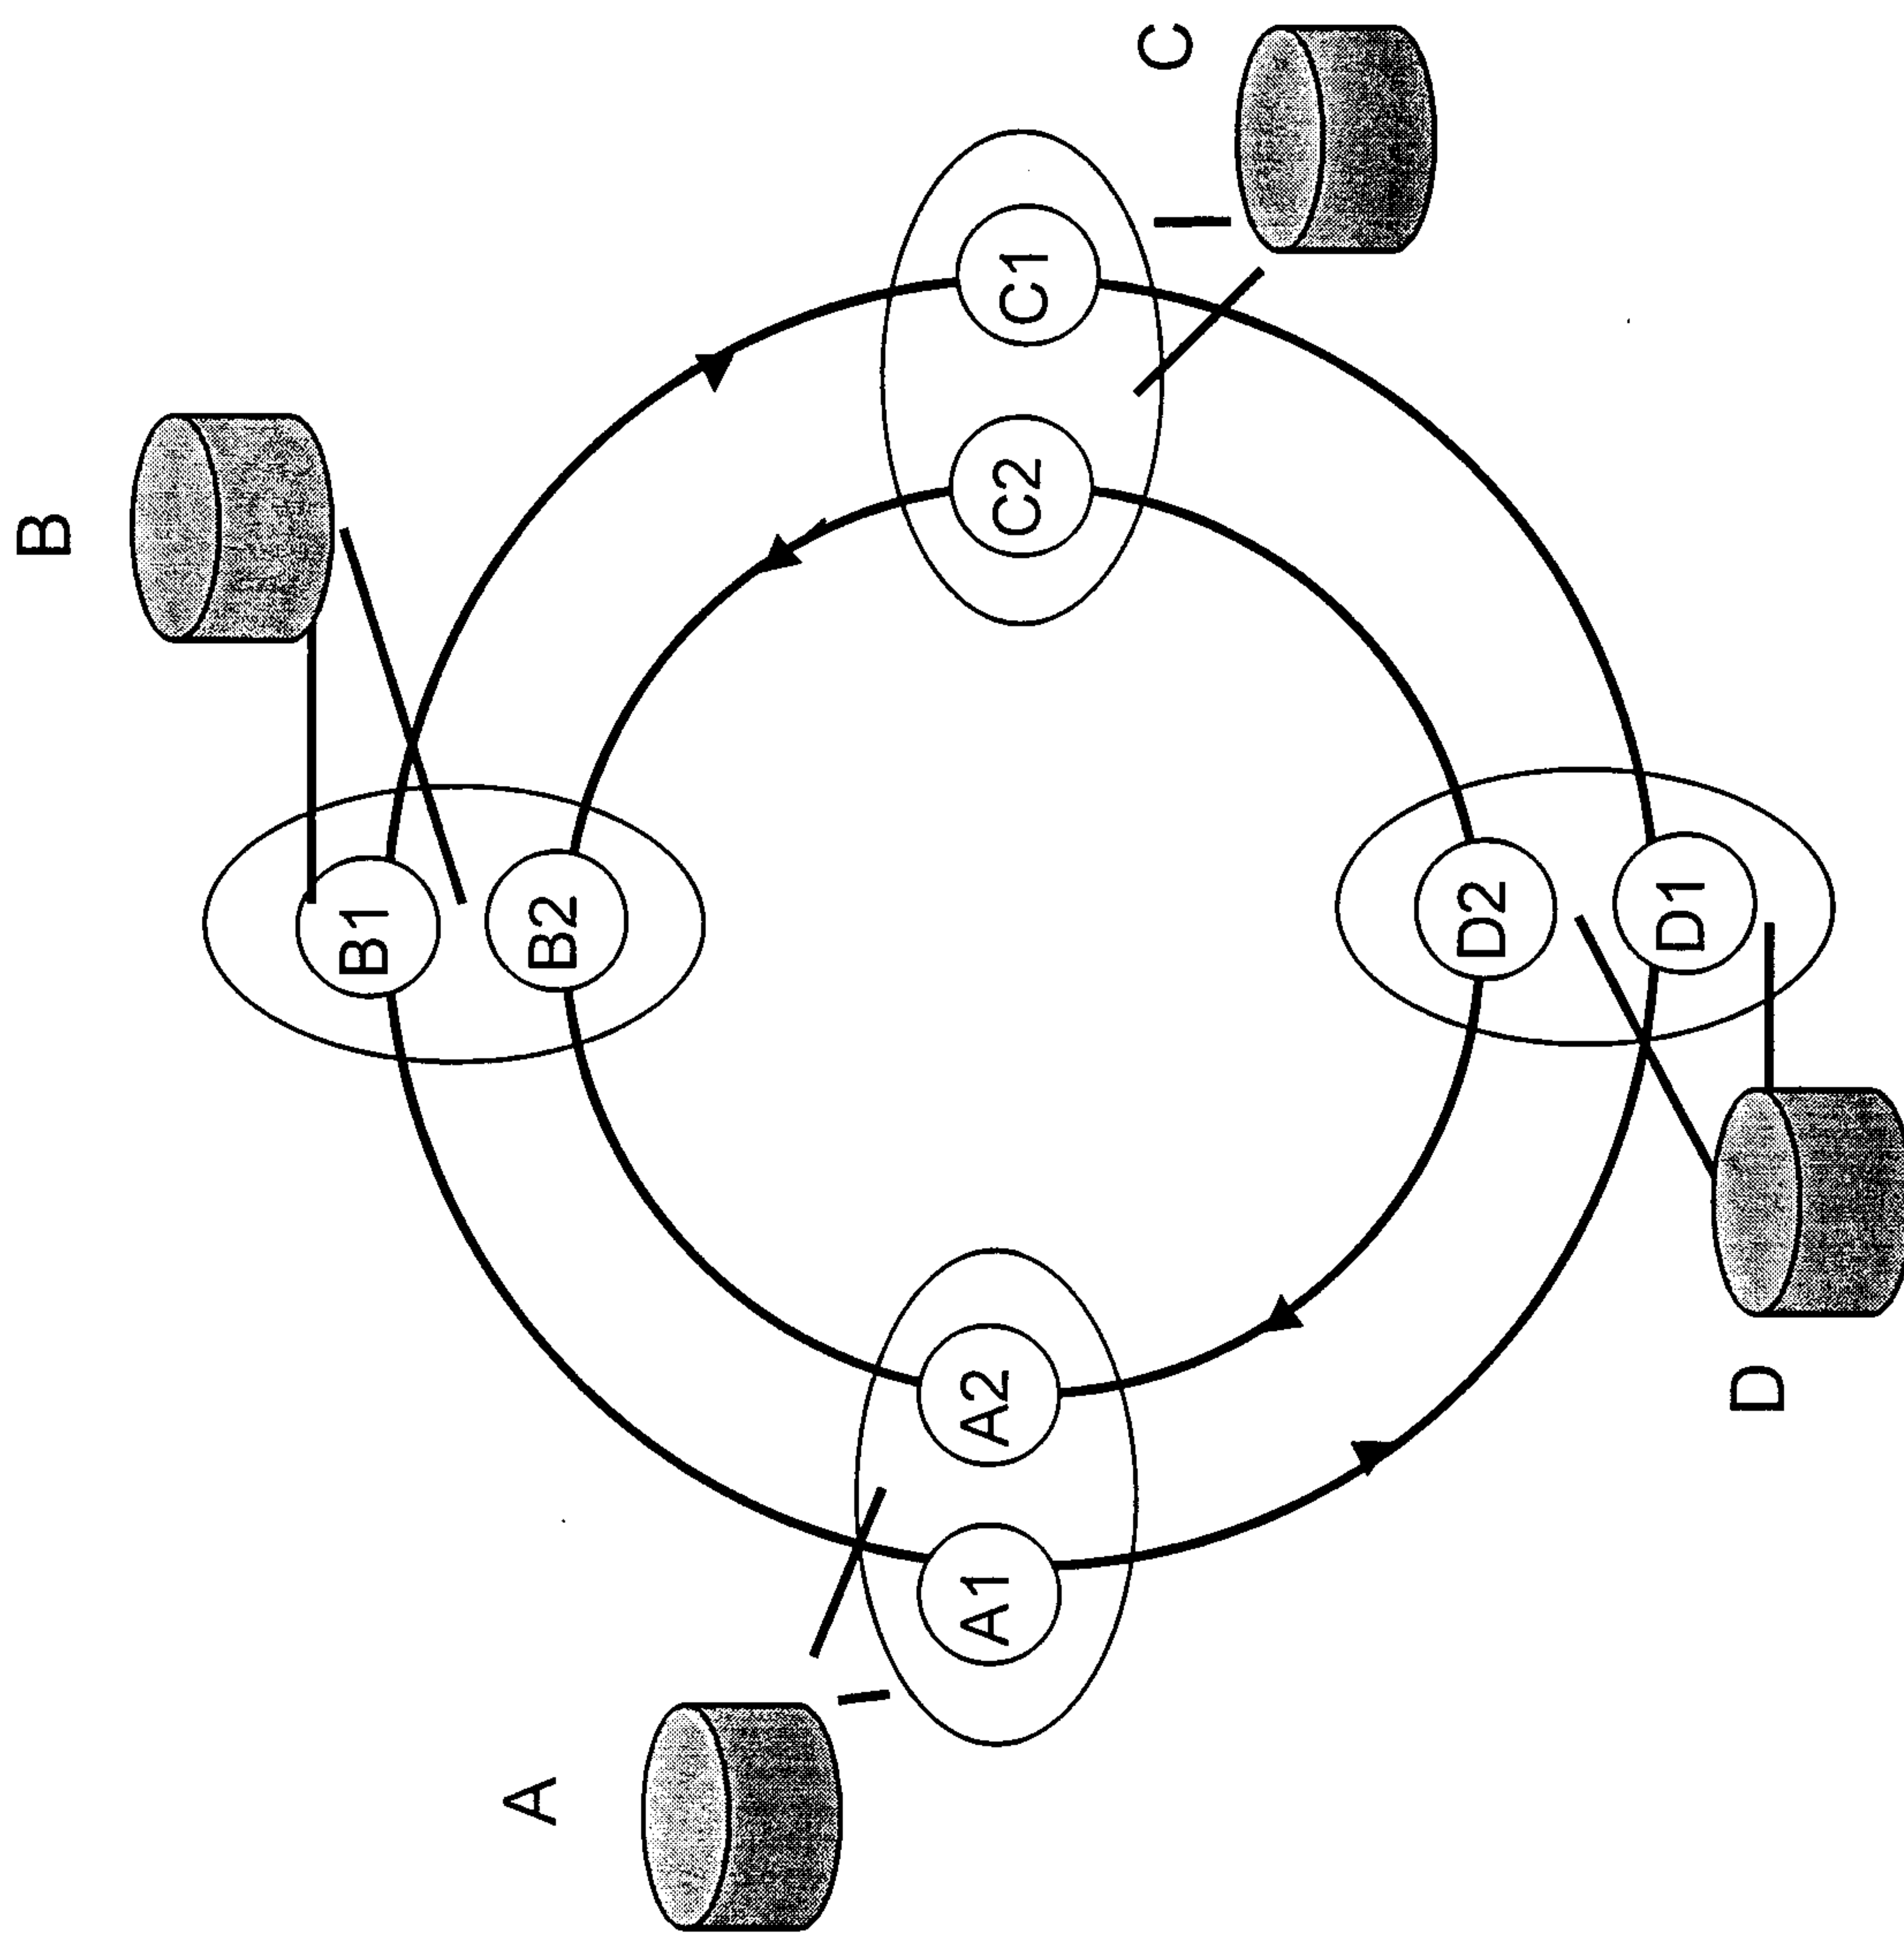


Figure 11

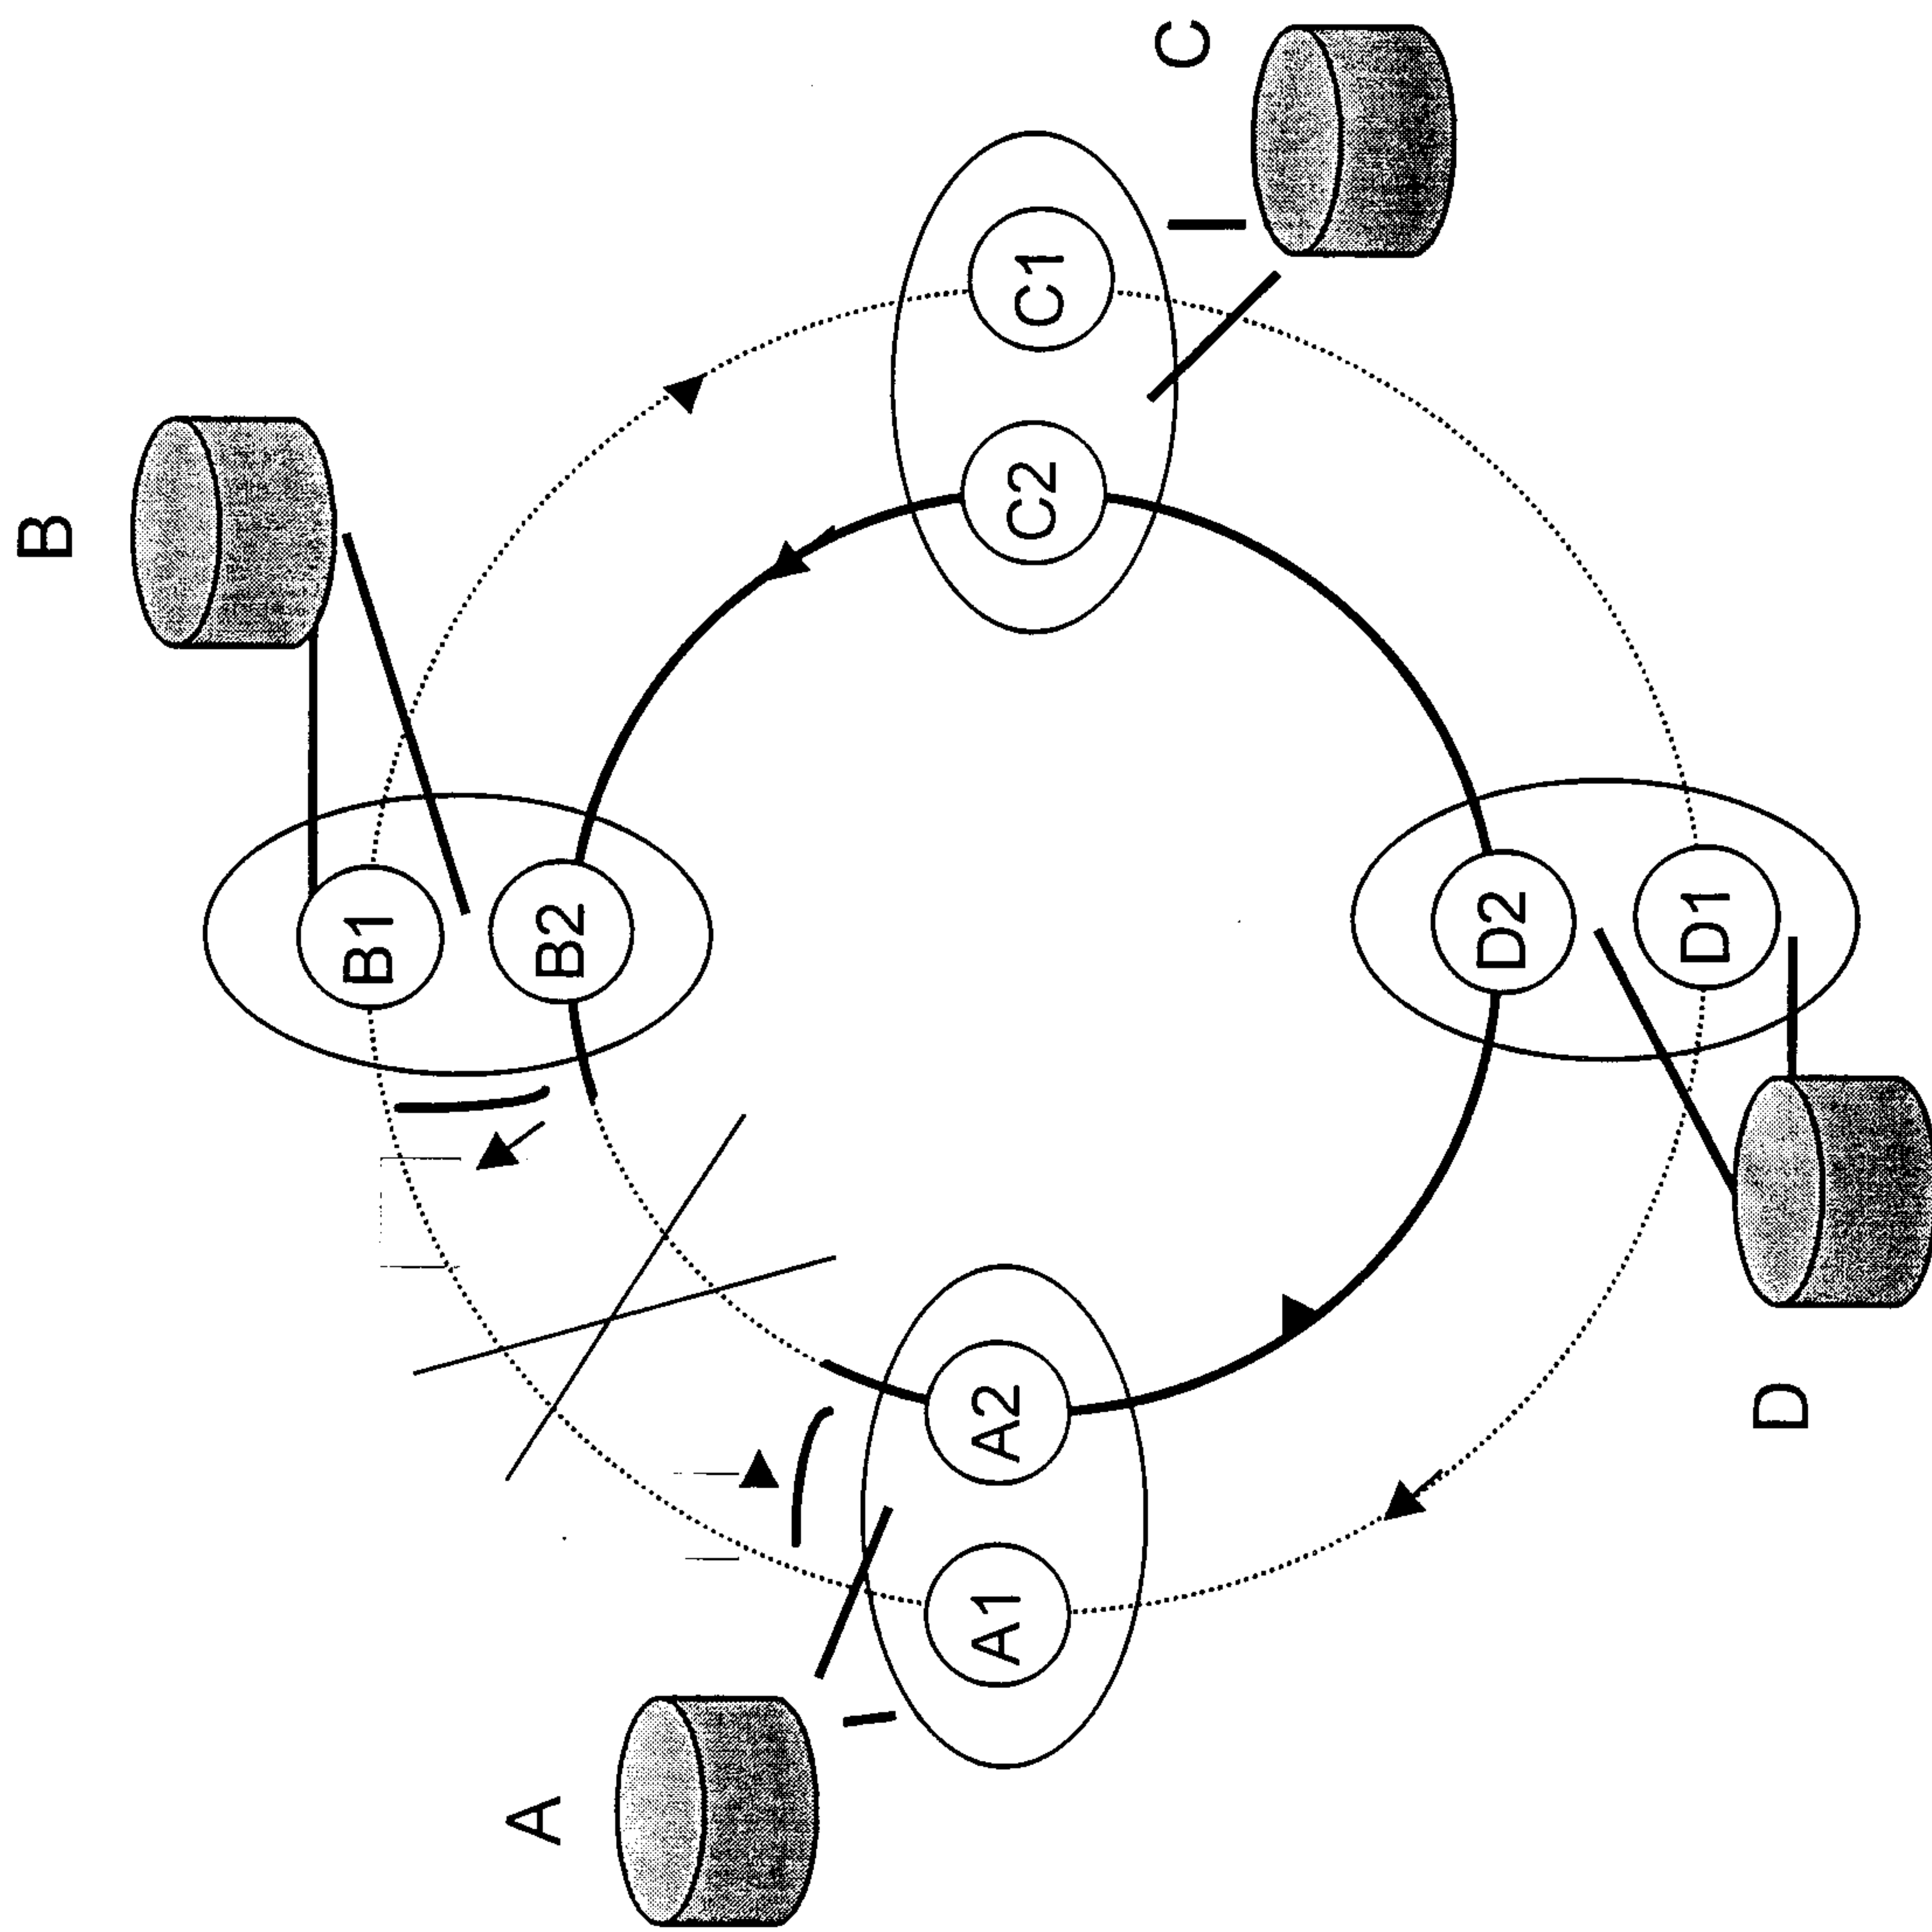


Figure 12

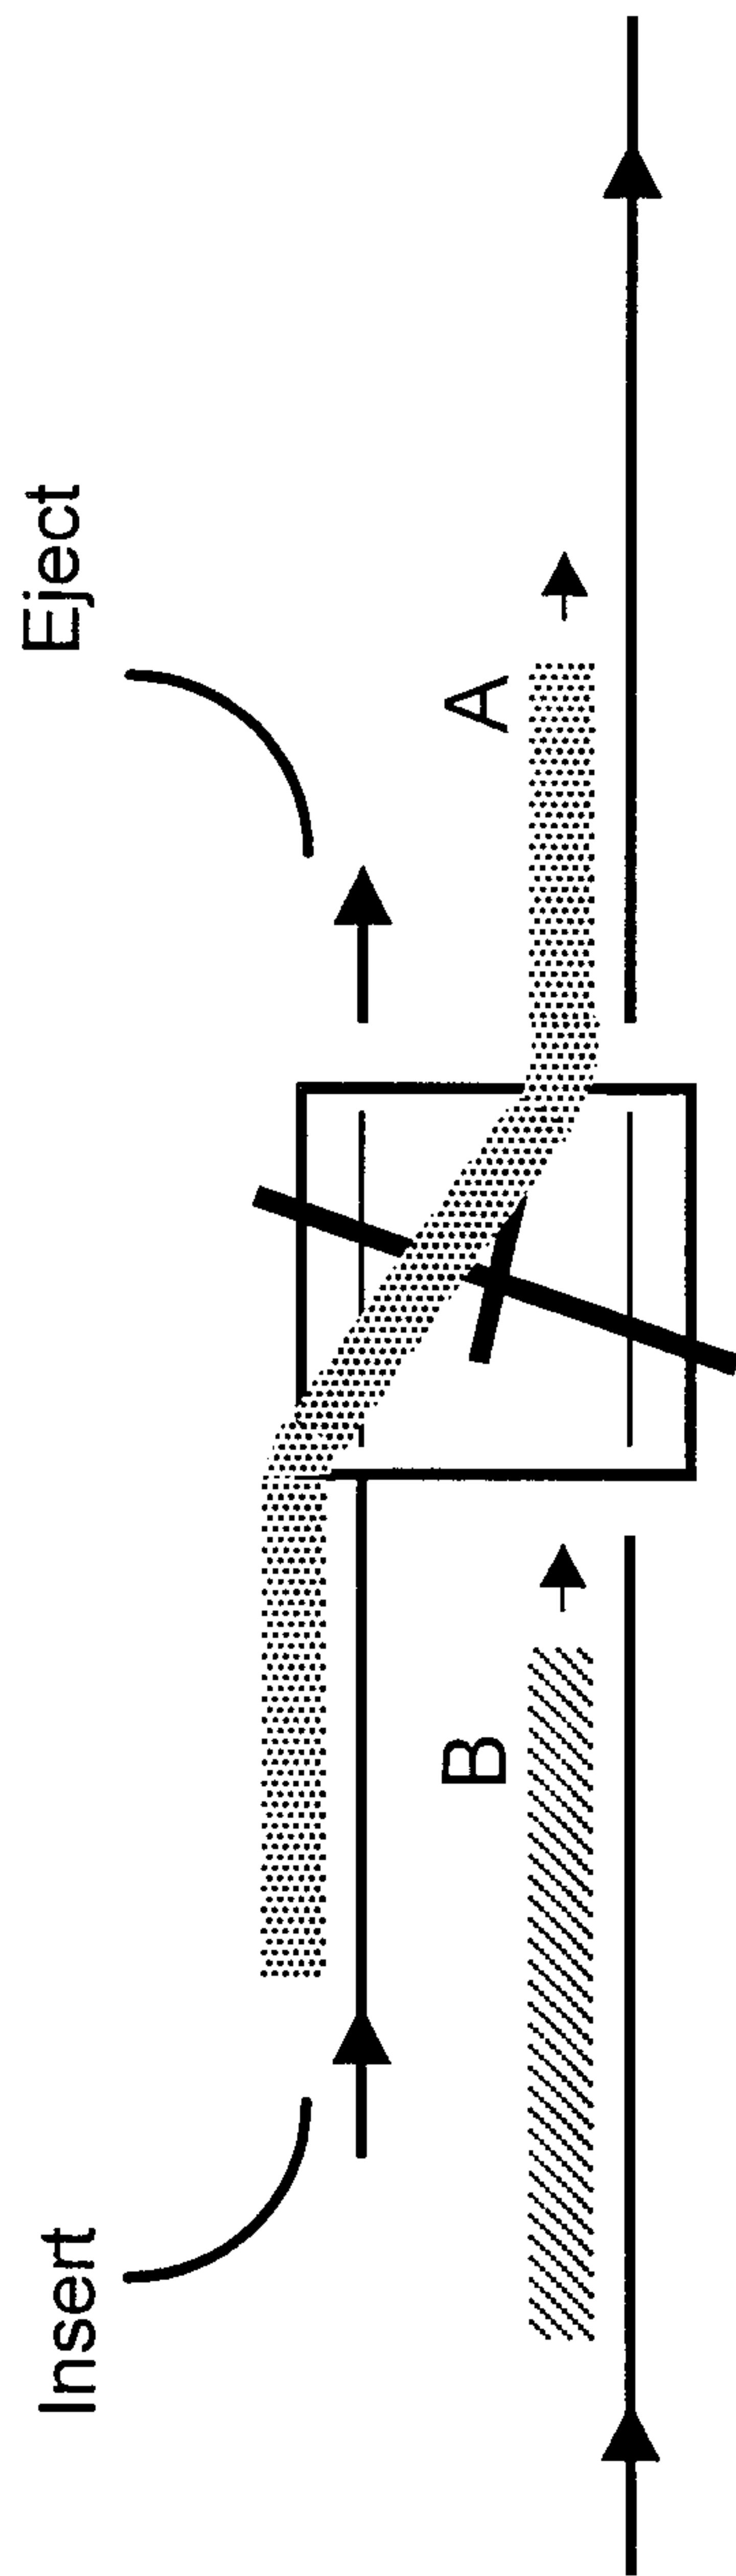


Figure 13

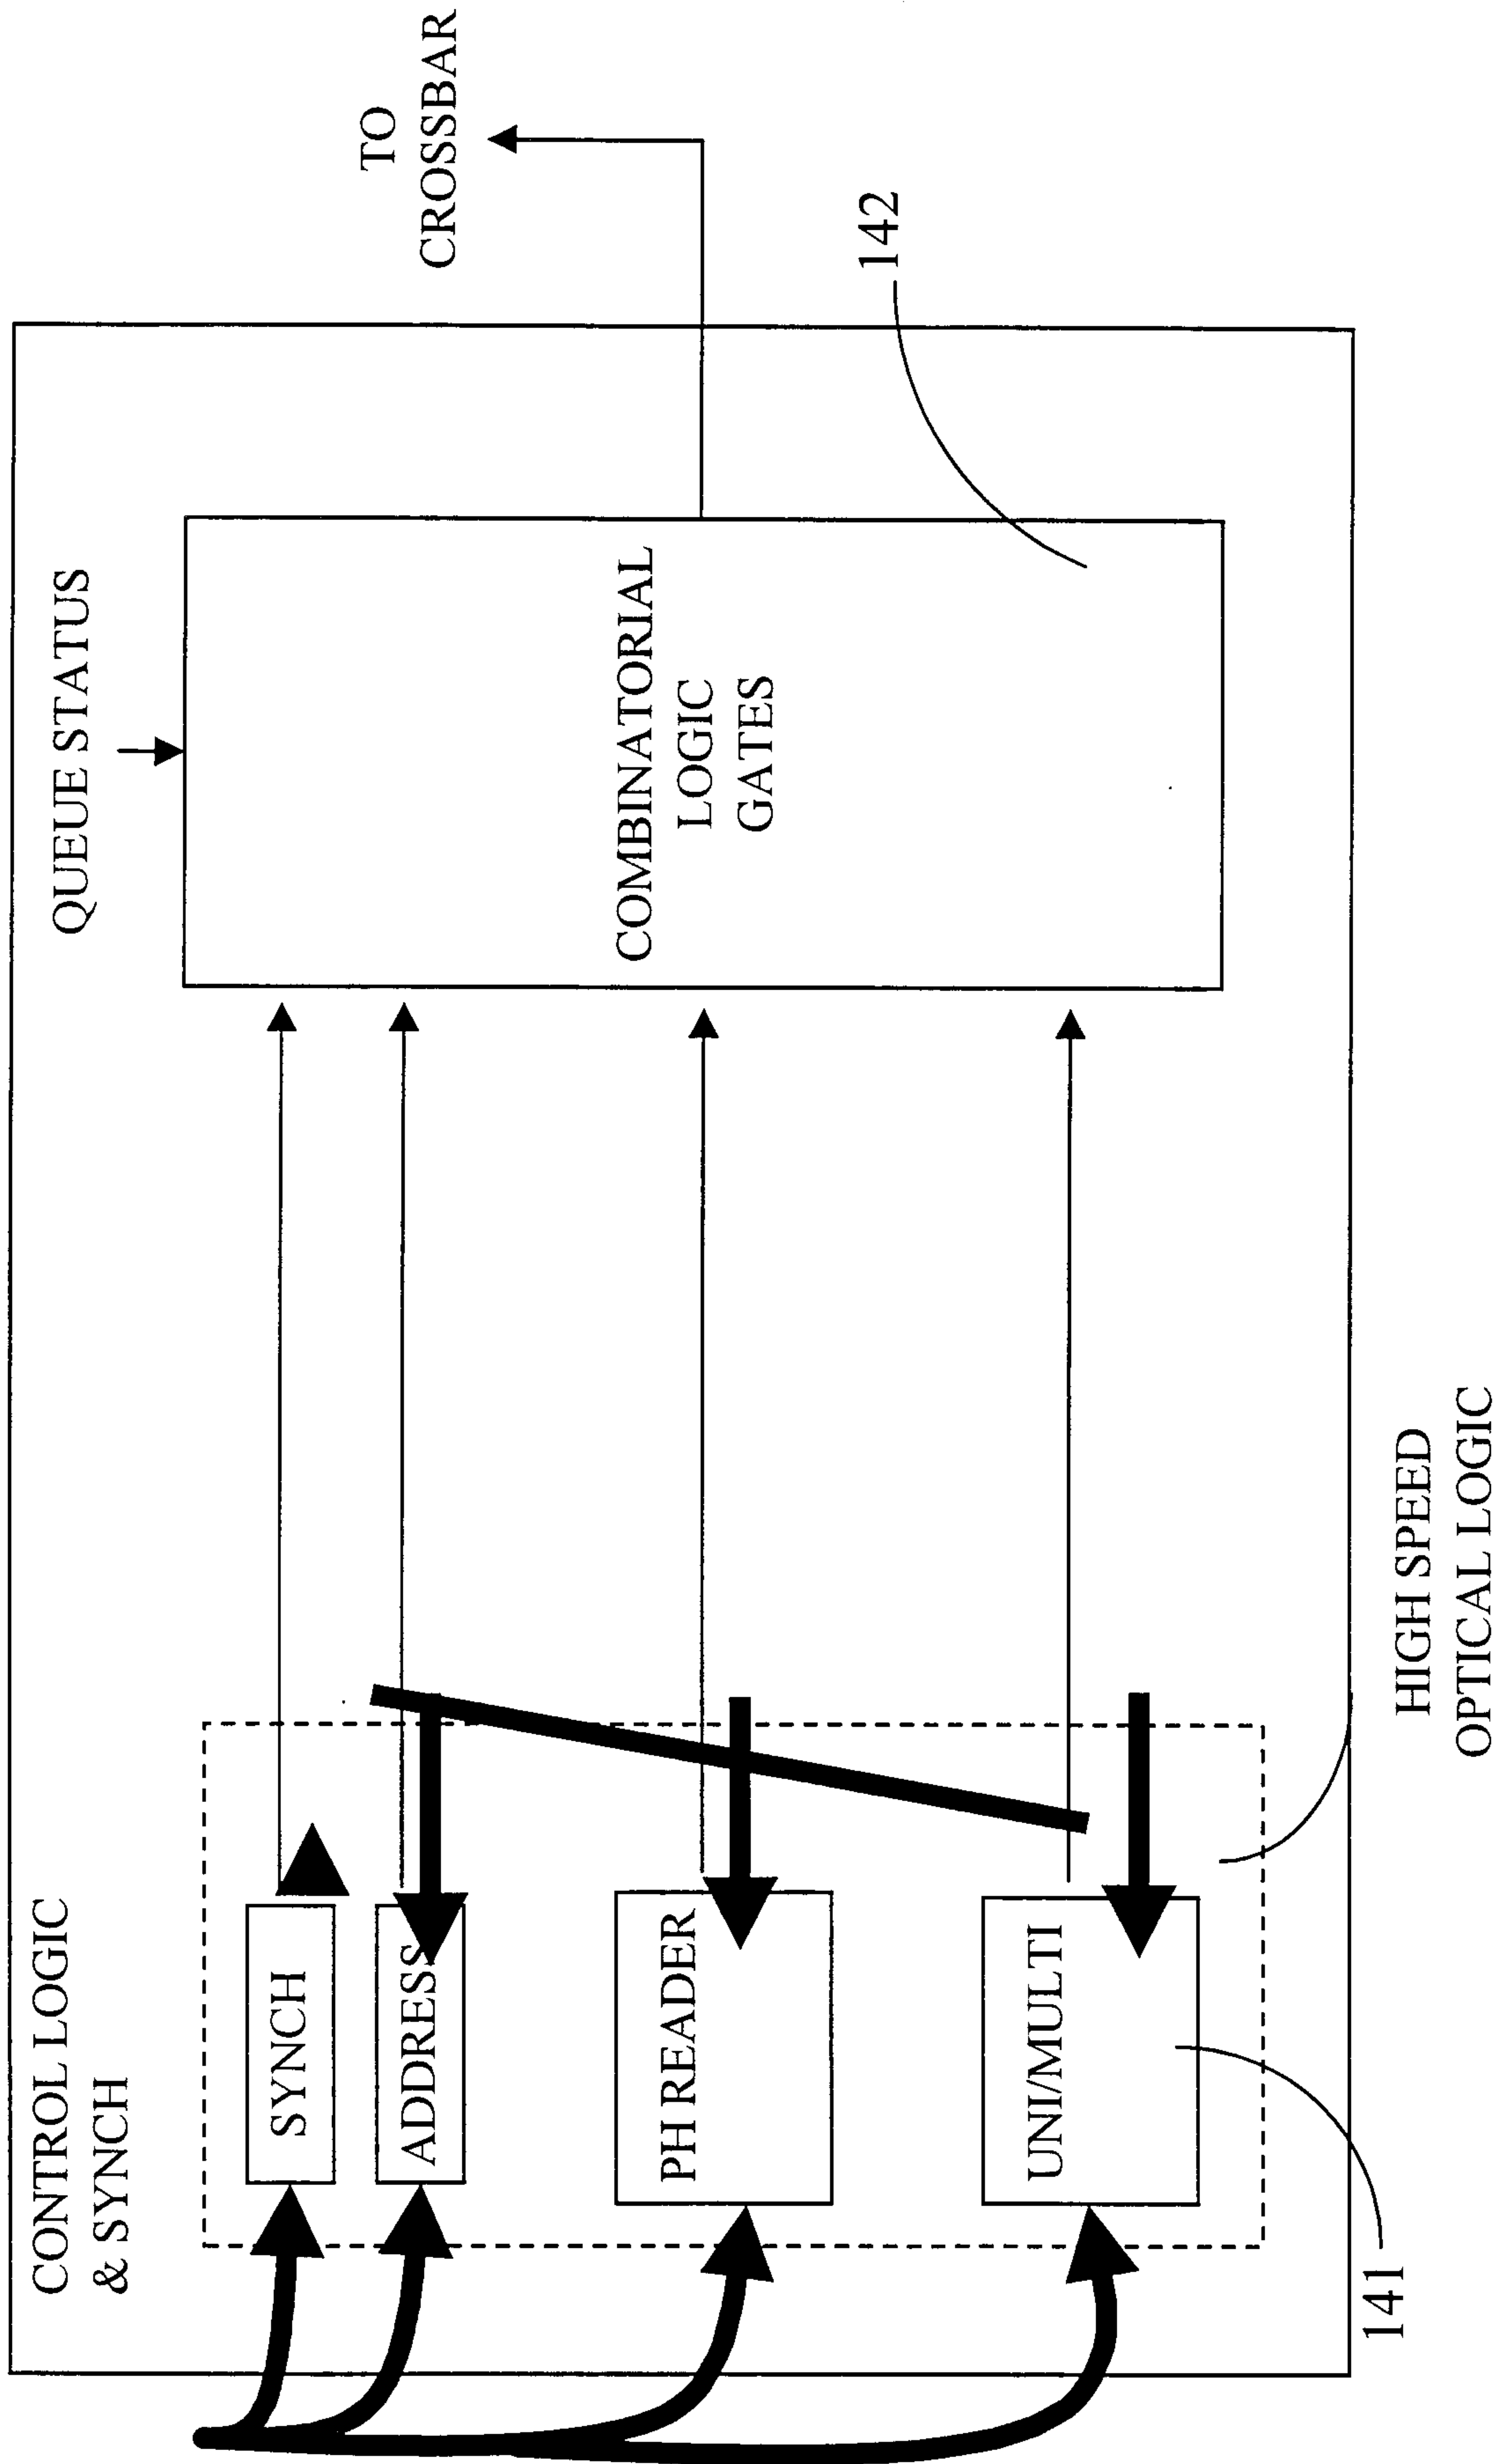


Figure 14

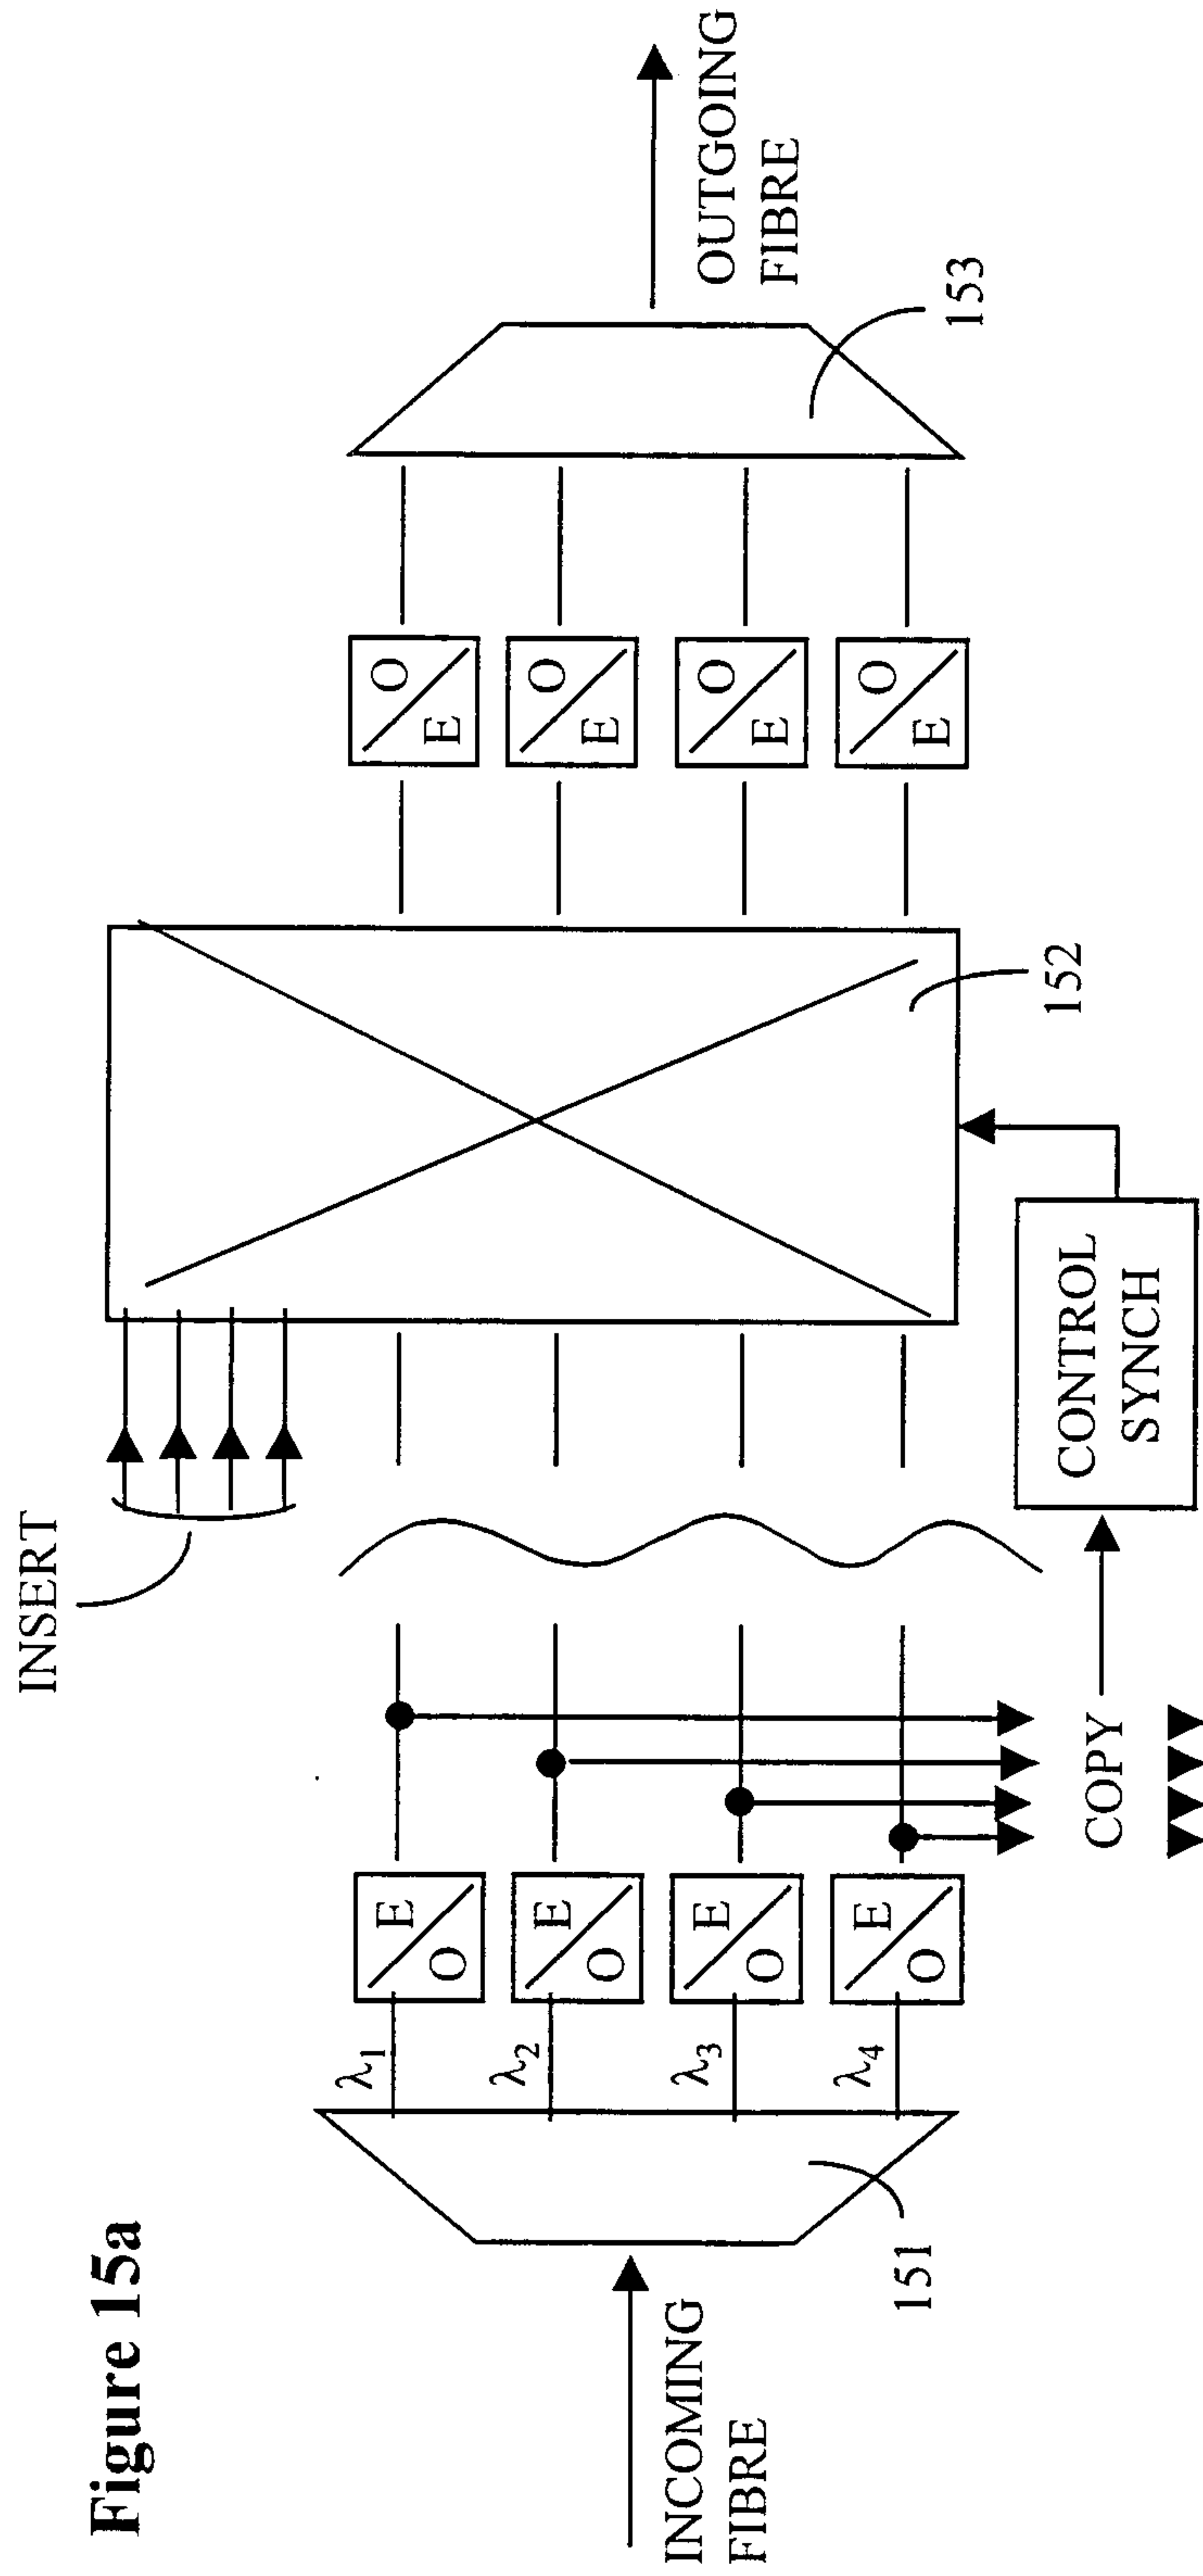


Figure 15a

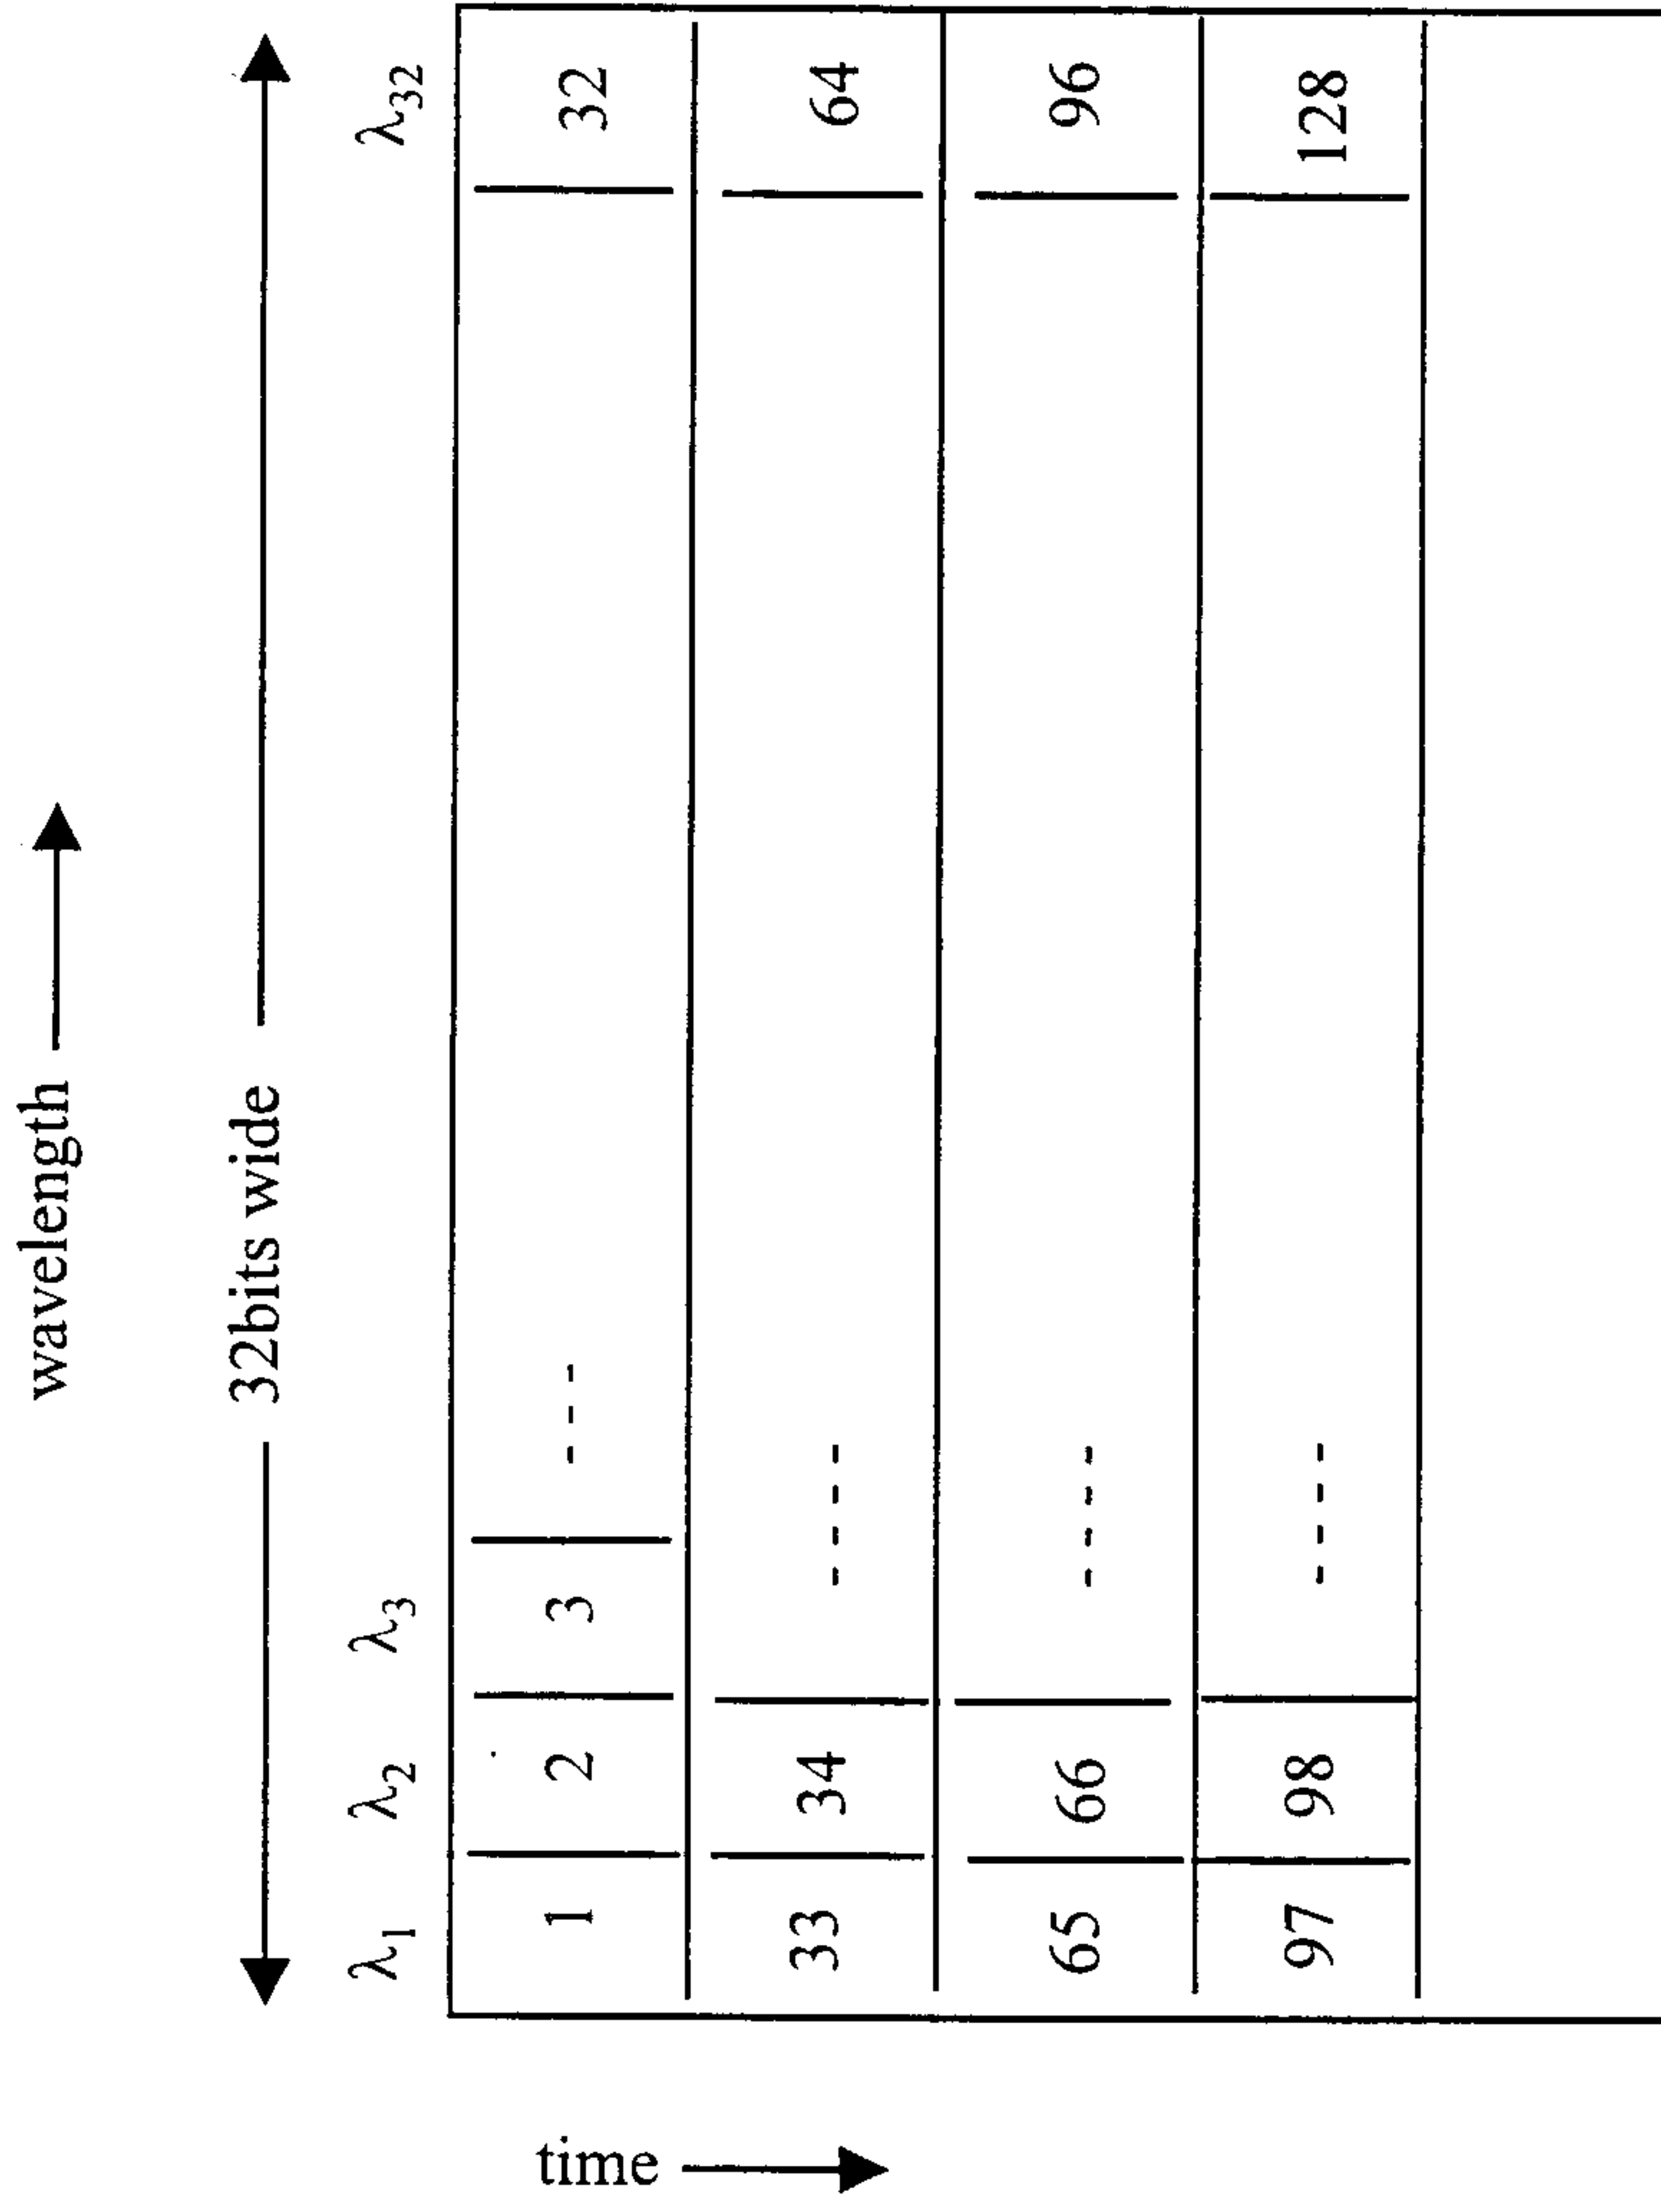


Figure 15b

Figure 16

