

(51) International Patent Classification ⁶ :

H04J 5/00, 5/22, 3/02

A1

(11) International Publication Number:

WO 96/04729

(43) International Publication Date:

15 February 1996 (15.02.96)

(21) International Application Number: PCT/US95/09665

(22) International Filing Date: 1 August 1995 (01.08.95)

(30) Priority Data:

110537

1 August 1994 (01.08.94)

IL

(71) Applicant: 3COM CORPORATION [US/US]; 5400 Bayfront Plaza, P.O. Box 58145, Santa Clara, CA 95052-8145 (US).

(72) Inventor: BRONSTEIN, Zvika; 5a George Elliot Street, 65235 Tel Aviv (IL).

(74) Agents: HAYNES, Mark, A. et al.; Haynes & Davis, Suite 310, 2180 Sand Hill Road, Menlo Park, CA 94025-6935 (US).

(81) Designated States: AU, CA, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).


Published

With international search report.

(54) Title: A NETWORK SWITCH

(57) Abstract

An asynchronous transfer mode (ATM) workstation (20), an ATM switch (12), legacy LAN workstations (14) and at least two network switches (10) are provided in a network. The ATM switch (12) is connectable to a multiplicity of ATM workstations (20) and the network switches (10) are each connectable to a different multiplicity of the legacy LAN workstations (14) and also to the ATM switch (12). Each network switch (10) includes apparatus for converting (fig. 5) back and forth between a legacy LAN protocol and an ATM protocol and apparatus for creating at least one emulated LAN from some of the legacy LAN and ATM workstations (20). The apparatus for converting (fig. 5) includes a cell memory (54) and a cell memory manager (50). The cell memory (54) stores fixed length cells of message data and the cell memory manager (50) for inputting and outputting queues of cells which form a message. The apparatus for converting (fig. 5) also includes legacy LAN and ATM interface units (40, 68) for communicating respectively with the legacy LAN and the ATM workstations. On input, the legacy LAN interface unit (40) breaks a legacy LAN message into a plurality of cells and, on output, provides a plurality of cells as a single message. The apparatus for creating (fig. 5) includes an address translation unit (44) having a map of the network and a listing of which legacy LAN and ATM workstations (14, 20) belong to which of the emulated LANs. The listing includes a listing of which legacy LAN workstations (14) require that their messages be converted to the ATM protocol in order to pass through the ATM switch (10).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo			SE	Sweden
CH	Switzerland	KR	Republic of Korea	SI	Slovenia
CI	Côte d'Ivoire	KZ	Kazakhstan	SK	Slovakia
CM	Cameroon	LI	Liechtenstein	SN	Senegal
CN	China	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TG	Togo
CZ	Czech Republic	LV	Latvia	TJ	Tajikistan
DE	Germany	MC	Monaco	TT	Trinidad and Tobago
DK	Denmark	MD	Republic of Moldova	UA	Ukraine
ES	Spain	MG	Madagascar	US	United States of America
FI	Finland	ML	Mali	UZ	Uzbekistan
FR	France	MN	Mongolia	VN	Viet Nam
GA	Gabon				

A NETWORK SWITCH

FIELD OF THE INVENTION

The present invention relates to network switches in general.

BACKGROUND OF THE INVENTION

There are many types of data communication networks, such as local and wide area networks (LANs and WANs), which connect together a multiplicity of computers. LANs connect computers which are close together and WANs connect computers which are far apart, respectively. LANs are implemented in a number of protocols, such as ETHERNET, Token Ring and Fiber Distributed Data Interface (FDDI). A more recent type of network protocol is the asynchronous transfer mode (ATM) protocol which is useful for both LANs and WANs. ATM networks, which are networks based on the ATM protocol, are connection-oriented in a manner similar to the telephone system.

Unfortunately, the various networks are incompatible with each other and therefore, if a user wants to take advantage of the qualities of the ATM networks, he must remove his old network and install a new ATM one.

Furthermore, most networks are not easily reconfigured. Typically, the workstation at which one person works is fixed in his room and is part of only one group of computers. For example, the groups might be the accounting and the engineering departments. If the user moves to a different room, the network has to be reconfigured, typically through bridges and routers, in order to access his new location. Similarly, if he changes groups but stays in the same room, the network has to be reconfigured. Any change typically also requires extensive management involvement.

SUMMARY OF THE PRESENT INVENTION

It is therefore an object of the present invention to provide a network switch which eases the problems described hereinabove.

There is therefore provided, in accordance with a preferred embodiment of the present invention, a network including asynchronous transfer mode (ATM) workstations, an ATM switch, legacy LAN workstations and at least two network switches. The ATM switch is connectable to a multiplicity of ATM workstations and the network switches are each connectable to a different multiplicity of the legacy LAN workstations and also to the ATM switch. Each network switch includes apparatus for converting back and forth between a legacy LAN protocol and an ATM protocol and apparatus for creating at least one emulated LAN from some of the legacy LAN and ATM workstations.

Additionally, in accordance with a preferred embodiment of the present invention, the apparatus for converting includes a cell memory and a cell memory manager. The cell memory stores fixed length cells of message data and the cell memory manager for inputting and outputting queues of cells which form a message.

Moreover, in accordance with a preferred embodiment of the present invention, the apparatus for converting also includes legacy LAN and ATM interface units for communicating respectively with the legacy LAN and the ATM workstations. On input, the legacy LAN interface unit breaks a legacy LAN message into a plurality of cells and, on output, provides a plurality of cells as a single message.

Furthermore, in accordance with a preferred embodiment of the present invention, the apparatus for creating includes an address translation unit having a map of the network and a listing of which legacy LAN and ATM workstations belong to which of the emulated LANs. The listing includes a listing of which legacy LAN workstations require that their messages be converted to the ATM protocol in order to pass through the ATM switch.

Additionally, in accordance with a preferred embodiment of the present invention, the cell memory manager receives translation information from the address translation unit.

Finally, in accordance with a preferred embodiment of the present invention, the legacy LAN interface unit includes a flow control unit having a queue of messages to be sent. The flow control unit indicates to an ATM interface on a remote ATM workstation whenever the legacy LAN interface is about to start sending a message for which not all of its cells are stored in the queue.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the drawings in which:

Fig. 1 is a schematic illustration of a network having a network switch constructed and operative in accordance with a preferred embodiment of the present invention;

Fig. 2 is a schematic illustration of the emulated network which the network switch of Fig. 1 creates;

Figs. 3A, 3B and 3C are schematic illustrations of data flow through the network switch of Fig. 1, where Fig. 1 illustrates regular communication, Fig. 3B illustrates point to multi-point communication and Fig. 3C illustrates Ethernet bridging;

Fig. 4 is a block diagram illustration of the network switch shown in Fig. 1;

Fig. 5 is a more detailed illustration of the network switch of Fig. 4;

Fig. 6 is a schematic illustration of Ethernet and ATM messages and their relationship, useful in understanding the network switch of Fig. 4;

Fig. 7A is a schematic illustration of the structure of a cell memory utilized in the network switch of Fig. 5; and

Fig. 7B is a schematic detailed illustration of portions of the cell memory of Fig. 7A.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Reference is now made to Figs. 1, 2, 3A and 3B which illustrate a network having a network switch, constructed and operative in accordance with a preferred embodiment of the present invention, and the operation of the switch.

5 The network of Fig. 1 typically comprises a plurality of network switches 10, constructed and operative in accordance with a preferred embodiment of the present invention, and an asynchronous transfer mode (ATM) switch 12 connected to each of the network switches 10. For example, the ATM switch can be any of the ATM switches commercially available from SynOptics of Santa Clara,
10 California, USA.

 Each network switch 10 is also connected to a multiplicity of "Legacy LAN" workstations 14, where the term "legacy LAN" denotes an Ethernet, Token Ring or FDDI network. Each legacy LAN workstation 14 has a legacy LAN controller card therein and is connected either directly (workstations 14a) or
15 indirectly through a hub 16 (workstations 14b) to the relevant network switch 10. Workstations 14c are servers and are shown directly connected to a network switch 10. Switches 10 can also be connected to a router 18 which connects the network shown in Fig. 1 to another network, marked WAN. The other network can be any type of network and does not have to be a WAN as marked.

20 The ATM switch 12 is also connected to a multiplicity of ATM workstations 20, each having an ATM controller card therein.

 The switches 10 of the present invention perform a number of operations. They provide communication among the legacy LAN workstations 14 and between legacy LAN workstations 14 and ATM workstations 20. As shown in Fig. 2 and
25 in accordance with the present invention, switches 10 create one or more interconnected emulated networks 22 to which any of legacy LAN workstations 14a - 14c and ATM workstations 20 can be attached. All traffic within an emulated network 22 is switched while traffic between emulated networks 22 is routed.

Typically, each emulated network 22 is provided to one working group in a company. If someone from one group changes rooms, no routing or recabling is required; his new workstation is just redefined to belong to his old emulated network 22. Similarly, if someone changes groups, his old workstation is just redefined to belong to his new emulated network 22. Thus, the network described in Fig. 2 is "location independent", meaning that the user is connected to his network no matter where he is. Furthermore, with an emulated network 22, the changes of location or group do not entail performance degradations.

Each workstation 14 is connected to the network switch 10 in its vicinity, regardless of the emulated network 22 to which the workstation 14 belongs. Each network switch 10 stores within it the emulated LAN 22 to which each of its connected workstations 14 belongs and enables the communication within and between emulated LANs 22.

It is noted that the ATM workstations 20 are not connected to network switches 10 but rather, are directly connected to the ATM switch 12. Because of this, for the ATM workstations 20 to be part of an emulated LAN, their ATM controllers must include in them units for producing emulated LANs. Otherwise, they can only communicate with other ATM workstations 20.

It will be appreciated that the type of workstation the user has is of little importance, since the switches 10 provide the necessary transformations between the legacy LAN and ATM communication protocols. It will be appreciated that the network switches 10 can convert between any legacy LAN protocol and the ATM protocol. Since the FDDI protocol is similar to the Ethernet one, FDDI and Ethernet workstations can be part of the same emulated LAN. Token Ring workstations, however, must either form a separate emulated LAN or incorporate an external mechanism by which to convert the Token Ring protocol to the Ethernet protocol.

Figs. 3A, 3B and 3C illustrate the ATM switch 12 with three network switches 10, 10b and 10c. In Fig. 3A, regular communication between switch 10a and switch 10b is shown. As can be seen, the data packet passes through the ATM

switch 12, since the switches 10 are connected to each other through the ATM switch 12.

In Fig. 3B, point to multi-point communication is shown in which data packets from one network switch are provided to the other two network switches. Fig. 3B shows three such communications. Once again, the data packets pass through the ATM switch 12.

In Fig. 3C, "switching" or communication between legacy LAN workstations 14 connected to a single network switch 10, is shown. In this case, communication between the workstations does not pass through the ATM switch 12.

If desired, the network switches 10 can perform the simple network management protocol (SNMP) which enables network management units to communicate with the network switches 10 to ascertain and change operation parameters.

Reference is now made to Figs. 4 and 5 which illustrate the elements of a network switch 10, to Fig. 6 which is useful in understanding the operation of the network switch 10 of Figs. 4 and 5, and to Fig. 7 which illustrates the structure of a cell memory of switch 10. Figs. 4 - 6 illustrate the switch 10 in the context of switching between the Ethernet and ATM protocols. It will be appreciated that similar constructions are required for the FDDI and Token Ring protocols.

The network switch 10 typically comprises Ethernet and ATM interface units 40 and 42 (Fig. 4), respectively, an address translation unit 44, a data control unit 46 and a host computer 48. The Ethernet and ATM interface units 40 and 42, respectively, convert messages to and from the Ethernet and ATM protocols from and to the internal protocol of the network switch 10, which, in the switch of Figs. 4 - 6, is the AAL5 ATM frame format. The Ethernet interface unit 40 communicates with the Ethernet workstations 14 to which it is connected and the ATM interface unit 42 communicates with the ATM switch 12.

Fig. 6 illustrates the relationship between an AAL5 frame 51, ATM messages 67 and an Ethernet frame 53. In accordance with the Ethernet format,

each Ethernet frame comprises destination and source media access control (MAC) addresses 61 and 64 and a variable length data portion 55. The length of data portion 55 is variable up to a maximum, standardized length. FDDI frames are organized similar to Ethernet frames.

5 The AAL5 frame 51 takes the Ethernet frame 53 as its payload 65 and adds to it a LAN emulation header (LECID) and some additional information 57 which includes the length of the payload 65, an error code and padding to reach a desired cell length.

10 Each ATM cell 67 comprises a header 69 and a fixed length cell data 71 forming part of the AAL5 frame 51. As shown in Fig. 6, the payload 65 of the AAL5 frame 51 may be divided among a number of ATM cells 67 and the additional information 57 may form the payload of another ATM cell 67. It is noted that the first cell data 71 corresponding to an AAL5 frame 51 includes in it the source and destination MACs 61 and 63 of the Ethernet frame 53.

15 The address translation unit 44 performs the LAN emulation, forwards and filters the messages passed through the switch. It also learns about and ages the topology of the network. The learning and aging activities are standard operations performed by standard bridges and routers. In learning, the network device, in this case the network switch 10, records the port at which each station resides and
20 utilizes this information when forwarding messages to that station. In aging, if a station has not been seen after a predetermined length of time, it is removed.

 The address translation unit 44 knows to which emulated network 22 each workstation belongs. Unit 44 receives each message and monitors the MAC address portions. It then determines what type of conversion, if any, must be
25 performed on the message. If the incoming message is from an Ethernet workstation 14 and is to be sent to another Ethernet workstation attached to the same network switch 10, the message will initially be converted to the AAL5 internal protocol and will be converted back to the Ethernet protocol on transmission. Otherwise (if it is to go to an ATM workstation 20 or to an Ethernet
30 workstation attached to a different network switch 10), the message will eventually

be converted from the internal AAL5 protocol to the ATM protocol. Conversely for a message from an ATM workstation 20.

The address translation unit 44 typically manages a multiplicity of ongoing one-way communications by giving each its own tag which is to be placed into the header 69 of the ATM message. The tag is known herein as the "virtual path/channel identification tag". Some of the communications can be flow control communications which are useful, though not necessary, in controlling the traffic between network switches 10.

The data control unit 46 performs the conversion between the Ethernet and ATM protocols. Unit 46 breaks the incoming messages, which may be of any length, into at least two fixed length cells of data and stores the cells in a managed queue. For the outgoing message, it provides the relevant cell or cells of the message in their proper order. The Ethernet protocol requires a plurality of cells to produce an Ethernet frame while the ATM protocol requires just the cells. The virtual tag maintained by the address translation unit 44 is used to identify which data belongs to which frame of a message.

The host computer 48 typically manages the operation of the network switch 10. It also provides output to network managers as to the activity and efficiency of the network switch 10.

Fig. 5 details one implementation of the network switch 10 and its connection to the Ethernet workstations 14. In this implementation, the switch 10 comprises an application specific integrated circuit (ASIC) 50, a translation memory 52, a cell memory 54, an arbiter 56, an ATM interface 58 and the host computer 48.

The ASIC 50 comprises a timing synchronizer 60, an Ethernet interface 62, a host interface 64, a translator 66, an ATM protocol module 68, a cell memory address generator 70, a cell memory data interface unit 72 and a queue control block 74.

The translation memory 52 and translator 66 form the address translation unit 44. The ATM protocol module 68 and ATM interface 58 form the ATM

interface unit 42. The cell memory 54, the address generator 70, data interface unit 72 and queue control block 74 form the data control unit 46. The Ethernet interface 62 and the arbiter 56 form the interface unit 40. Finally, the host interface 64 interfaces between the host 48 and the elements of the ASIC 50.

5 The ATM module 68 communicates with a single ATM controller 80 and, on input, breaks an ATM cell 67 into its header 69 and cell data 71 portions. If the incoming cell is the first of a series of cells, ATM module 68 removes the source and destination MAC addresses 61 and 63 and provides them to the address translation unit 44. At the same time, ATM module 68 provides the cell data 71
10 (which includes the MAC addresses 61 and 63) to the data control unit 46. If the incoming cell is not a first cell, ATM module 68 only provides the cell data 71 to the data control unit 46. On output, ATM module 68 receives the relevant data from the data control block 46 and the header information from the translation block 66. The ATM module 68 combines the data and header information together
15 to produce cells 67, where the first cell 67 has all of the address information and some data and subsequent cells 67 only have data.

 The translation block 66 analyzes the received addresses, translates them into locations in one of the emulated LANs 22 and determines whether or not the source and destination addresses are permitted to communicate (i.e. are part of the
20 same emulated LAN 22). The topologies of the emulated LANs 22 (e.g. the workstations belonging to each LAN 22) are stored in translation memory 52.

 Block 66 determines the direction the message has to go, either to one or more Ethernet workstation connected to the switch 10 and/or to the ATM switch 12 to be sent to one or more ATM workstation. To that end, it controls the queue
25 control block 74, as will be described in more detail hereinbelow.

 The translation block 66 performs bridging and aging of addresses, as described hereinabove. Whenever the forwarding address of a workstation changes, the translation block 66 provides this information to the host computer 48.

 The arbiter 56 determines which of Ethernet controllers 82, which are
30 connected together along a bus 84, is to communicate with Ethernet interface 62.

To do so, arbiter 56 performs a time multiplexing, or round robin, scan of the bus 84, connecting each Ethernet controller 82 with the interface 62, one after the other, for a short period of time.

5 In accordance with the present invention, rather than communicating an entire frame at once, the Ethernet interface 62 only receives or provides (depending on the direction of the frame) a cell-sized portion of the frame 53 at one time. Therefore, the Ethernet interface 62 must receive or provide a number of times, in accordance with the round robin scan of arbiter 56, before interface 62 finishes communicating the entire frame.

10 So that the Ethernet controllers 82 can operate according to the Ethernet standard and can receive and send Ethernet frames, rather than cells of frames, arbiter 56 operates at least 20 times faster than the speed of the controllers 82.

Furthermore, to ensure that all the cells which make up an Ethernet frame can be sent one right after the other, the Ethernet and ATM interfaces 62 and 68, respectively, perform flow control through the ATM switch 12. The Ethernet
15 interface 62 has a queue of frames, each consisting of a list of cells, to be sent out. Whenever a frame is read from a queue and the Ethernet interface 62 is about to start sending a frame for which it does not yet have all the cells, it notifies the ATM interface 68 to send a flow control cell to the source network access device.
20 The ATM interface of the source network access device, in turn, raises the priority of the missing cells so that they pass through the ATM switch 12 to the data control unit 46 of the destination network switch 10.

During receipt of a message, the Ethernet interface 62 typically breaks the received partial message into cell-sized data portions and provides them to
25 interface 62. The translation block 66 converts the MAC address portion into an internal address. The interface 62 also provides the cell-sized data to the cell memory interface unit 72 and information regarding the frame to which it belongs to the queue control block 74.

For outgoing messages, the queue control block 74 causes the cell memory
30 interface 72 to provide the cells of the frame, in order. The interface 62 then

creates partial messages and provides them to the bus 82. At the same time, it indicates to arbiter 56 to activate the relevant controller 82.

It will be appreciated that, because the interface 62 communicates cells rather than frames, the network switch 10 can begin switching the message to its destination as soon as the first cell arrives. This is known as a "cut-through" mode. When operating in the cut-through mode, the switching latency (defined as the time it takes for a message to pass through a switch) is very low. The network switch 10 can alternatively operate in a "store and forward" mode in which a frame is transmitted only after it has been completely received.

Fig. 7A illustrates the structure and operation of the cell memory 54 and Fig. 7B is a blowup of a portion of Fig. 7A. The cell memory 54 typically comprises a frame pointer array 90, a queue head array 91, a multiplicity of cell data fragments 92 and a VPCI table 94.

Each data fragment 92 contains a cell of data 94, an element listing the length of useful data in cell 96, a copy index 98 defining how many times the message is to be sent (useful for sending point to multi-point messages) and a pointer 100 to a next fragment. Each queue head, labeled 122, points to a chained series of data fragments 92 available for output. Each frame pointer, labeled 120, points to the beginning of a series of data fragments 92 currently being put into the cell memory 54.

The cell memory is managed by the cell RAM address generator 70 which comprises the following pointers: the Ethernet read (ER) pointer 102, the Ethernet write (EW) pointer 108, the ATM write (AW) pointer 110 and the empty list pointer 112.

An Ethernet frame is entered into the cell memory 54 in cell-sized portions. When the first cell-sized portion of the Ethernet frame arrives, the location of an empty data fragment 92 is taken from the empty list pointer 112 and placed into the appropriate location of both frame pointer array 90 (for input) and the queue head array 92 (for output). This location becomes the frame pointer 120 for the entire incoming frame and the value of the frame pointer 120 is also placed into the ER

pointer 102. The first portion of the Ethernet frame is stored into the empty data fragment 92. At the same time, the location of the next empty data fragment is taken from the empty list pointer 112 and stored into data fragment 92 as the pointer 100 to the next data fragment. The value of the pointer 100 is then used to
5 update ER pointer 102.

When the next cell-sized portion of the Ethernet frame arrives, it is stored into the data fragment 92 indicated by ER pointer 102 and the location of the next empty data fragment is taken from the empty list pointer 112 and stored in both pointer 100 of the data fragment 92 and in the ER pointer 102. The process is
10 repeated until all of the portions of the Ethernet frame have arrived.

Since the location of the first cell-sized portion of the frame was put into the queue head array 92, the chained series of cells can be put out (either for ATM or Ethernet use) almost immediately, even while portions of it are still being entered.

On output, the relevant queue head 122, as indicated by the queue control 74 block 66, is read and copied to either the EW pointer 108 (for Ethernet output) or to the AW pointer 110 (for ATM output). The data fragment 92 to which it points is read out via memory interface 72 and the copy index 98 is reduced every time a cell is sent to a destination. Thus, if the message has to be sent five times,
15 to five different destinations, the cell will be removed from memory only after the fifth copy has been sent. Once the last copy has been sent, the location of the data fragment 92 is placed into the empty list pointer 112 since the data fragment is now empty.

Each time the data fragment is read, the value of the pointer 100 in the data
25 fragment 92 is provided to the relevant EW pointer 108, for Ethernet output, and to the AW pointer 110 for ATM output. The relevant pointers then access the next data fragment 92 in the chain. The process is repeated until the entire frame or cell series have been accessed.

Typically, only a small number, such as 12, of Ethernet frames are input at
30 once. In contrast, a large number, such as 100 - 300, ATM frames are processed

at one time. Therefore, there is no ATM read pointer. Instead, the VPCI table includes both the initial cell location pointer and the current cell location pointer. Like for the Ethernet process, the VPCI table receives the location of the next available data fragment from the empty list pointer 112 and, on input, stores the incoming ATM cell in the indicated data fragment. If a series of ATM cells have to be stored, then they are chained together as described for the Ethernet frames with the exception that there is no AR pointer. Instead, the VPCI table continually updates the read pointers.

The ASIC synchronizer 60 provides the basic timing for the ASIC 50. Typically, the operations which the network switch 10 performs are repeated after four clocks. On the first clock, a partial message (typically of four bytes long) is read from one Ethernet controller 82. On the second clock, a partial message is written to one Ethernet controller 82, which is a different controller than that from which a partial message was read. On the third and fourth clocks, messages are received from and written to, respectively, the ATM controller 80.

The host computer 48 manages the operation of ASIC 50. This includes:

- a) learning the topology of the emulated LANs;
- b) aging of addresses;
- c) gathering of statistics;
- d) recovering errors;
- e) SNMP management; and
- f) other bridging operations.

It will be appreciated that the present invention enables one or more emulated LANs to be produced from a plurality of existing LANs, whether legacy LAN or ATM or both.

It will further be appreciated that the principles of operation of the network switch 10 provided herein are applicable to other types of legacy LANs. For example, to create a network switch for FDDI networks, the Ethernet controller 82 (Fig. 5) and arbiter 56 must be replaced by an FDDI interface and an FDDI arbiter, respectively. The remaining elements need not be changed. For a Token Ring

network, in addition to replacing the Ethernet controller 82 and arbiter 56, the VPCI table 94, routing tables of the translation memory 52 and some of the internals of the ASIC itself, must be changed. However, the basic operation of the data control unit 46 remains the same for all legacy LANs.

- 5 It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather the scope of the present invention is defined by the claims which follow:

CLAIMS

1. A network having at least one emulated local area network (LAN) therein, said network comprising:

5 asynchronous transfer mode (ATM) workstations;
an ATM switch connectable to a first multiplicity of said ATM workstations;

legacy LAN workstations; and

10 at least two network switches, each connectable to a second multiplicity of said legacy LAN workstations and to said ATM switch;

wherein each network switch comprises:

means for converting back and forth between a legacy LAN protocol and an ATM protocol; and

15 means for creating said at least one emulated LAN from some of said legacy LAN and ATM workstations.

2. A network according to claim 1 and wherein said means for converting comprises:

a cell memory for storing fixed length cells of message data; and

20 a cell memory manager for inputting and outputting queues of cells which form a message.

3. A network according to claim 2 and wherein said means for converting additionally comprises legacy LAN and ATM interface units for communicating
25 respectively with said legacy LAN and said ATM workstations and wherein, on input, said legacy LAN interface unit breaks a legacy LAN message into a plurality of cells and, on output, provides a plurality of cells as a single message.

4. A network according to claim 1 wherein said means for creating comprises
30 an address translation unit having stored therein a map of said network and a listing

of which legacy LAN and ATM workstations belong to which of said at least one emulated LAN.

5 5. A network according to claim 4 and wherein said listing includes a listing of which legacy LAN workstations require that their messages be converted to the ATM protocol in order to pass through the ATM switch.

6. A network according to claims 2 and 4 and wherein said cell memory manager receives translation information from said address translation unit.

10

7. A network according to claim 3 and wherein said legacy LAN interface unit comprises a flow control unit having a queue of messages to be sent, wherein said flow control unit indicates to an ATM interface on a remote ATM workstation whenever said legacy LAN interface is about to start sending a message for which not all of its cells are stored in said queue.

15

8. A network switch connectable to a multiplicity of legacy LAN workstations and to an ATM switch, which, in turn, is connected to a different multiplicity of ATM workstations, and network switch comprising:

20 means for converting back and forth between a legacy LAN protocol and an ATM protocol; and

means for creating at least one emulated LAN from some of said legacy LAN and ATM workstations.

25 9. A network switch according to claim 8 and wherein said means for converting comprises:

a cell memory for storing fixed length cells of message data; and

a cell memory manager for inputting and outputting queues of cells which form a message.

30

10. A network switch according to claim 9 and wherein said means for converting additionally comprises legacy LAN and ATM interface units for communicating respectively with said legacy LAN and said ATM workstations and wherein, on input, said legacy LAN interface unit breaks a legacy LAN message into a plurality of cells and, on output, provides a plurality of cells as a single message.

11. A network switch according to any of claims 8 - 10 and wherein said means for creating comprises an address translation unit having stored therein a map of said network and a listing of which legacy LAN and ATM workstations belong to which of said at least one emulated LAN.

12. A network switch according to claim 11 and wherein said listing includes a listing of which legacy LAN workstations require that their messages be converted to the ATM protocol in order to pass through the ATM switch.

13. A network switch according to claim 9 wherein said cell memory manager receives translation information from said address translation unit.

14. A network switch according to claim 10 and wherein said legacy LAN interface unit comprises a flow control unit having a queue of messages to be sent, wherein said flow control unit indicates to an ATM interface on a remote ATM workstation whenever said legacy LAN interface is about to start sending a message for which not all of its cells are stored in said queue.

15. A network switch according to claim 11 wherein said cell memory manager receives translation information from said address translation unit.

1 / 10




FIG. 1




FIG.2

3/ 10




FIG.3A

4 / 10

FIG. 3B







FIG. 3C





6 / 10

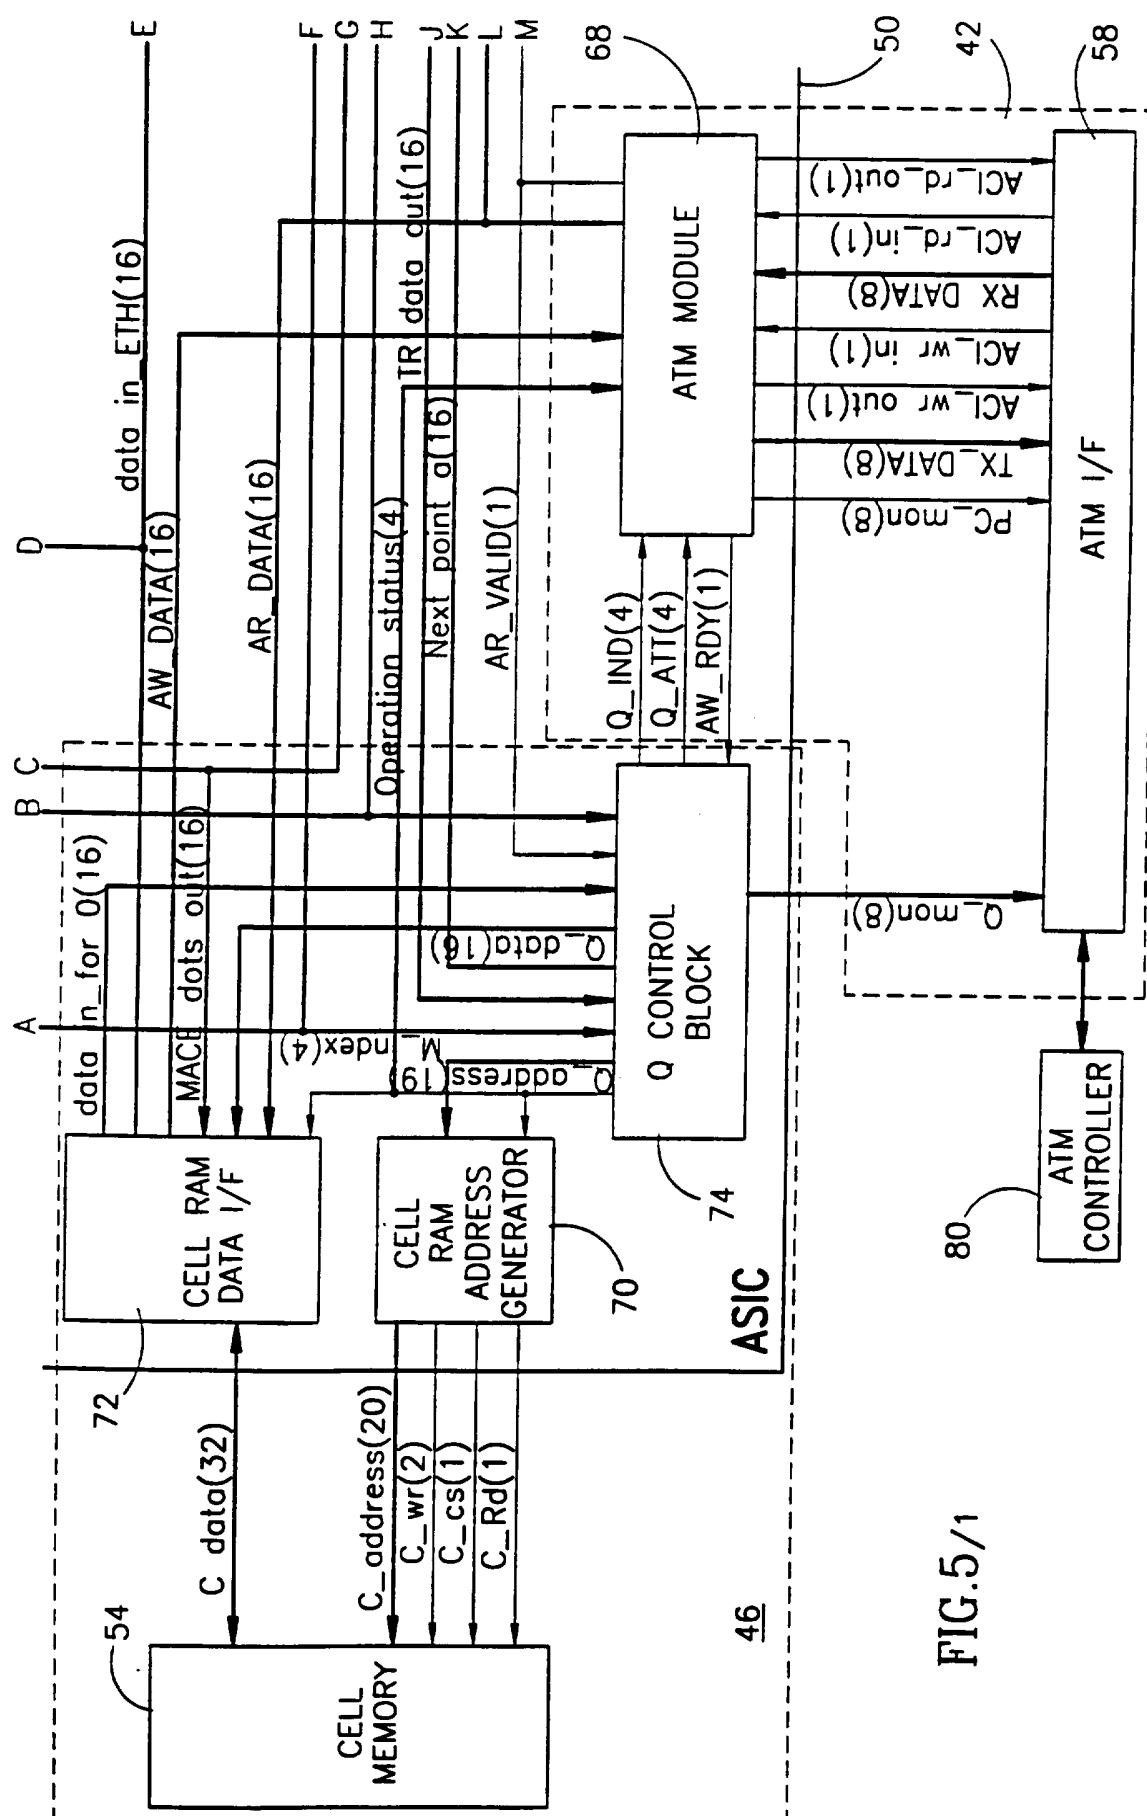


FIG.5/1





FIG.5/3






FIG.6

10 / 10




FIG. 7A

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/09665**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(6) : H04J 5/00, 5/22; H04J 3/02

US CL : 395/200.02; 370/85.13

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/200.02; 370/85.13

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Please See Extra Sheet.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 5,329,527 (UJIHASHI ET AL) 12 July 1994, Abstract, fig.1	1-15
Y	US, A, 4,975,906 (TAKIYASU ET AL) 04 December 1990, Abstract, fig.2.	1-15
Y,P	US, A, 5,425,026 (MORI) 13 June 1995, Abstract.	1-15
Y	US, A, 5,280,481 (CHANG ET AL) 18 January 1994, abstract.	
Y	US, A, 5,394,402 (ROSS) 28 February 1995, abstract.	1-15
Y,P	US, A, 5,363,369 (HEMMADY ET AL) 08 November 1994, abstract and fig.5.	1-15

☒ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	*T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
*A document defining the general state of the art which is not considered to be part of particular relevance	*X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
*E earlier document published on or after the international filing date	*Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
*L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z document member of the same patent family
*O document referring to an oral disclosure, use, exhibition or other means	
*P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

15 SEPTEMBER 1995

Date of mailing of the international search report

05 OCT 1995

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Authorized officer

THOMAS LEE

Facsimile No. (703) 305-3230

Telephone No. (703) 305-9646

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US95/09665

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US, A, 5,414,707 (JOHNSTON ET AL) 09 MAY 1995, abstract.	1-15
Y,P	US, A, 5,425,029 (HLUCHYJ ET AL) 13 June 1995, abstract.	1-15
Y	US, A, 5,287,535 (SAKAGAWA ET AL) 15 February 1994, abstract.	1-15
A	IEEE Network, March 1993 "Designing a Practical ATM LAN", Biagioni et al. pages 32-39.	1-15

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US95/09665

B. FIELDS SEARCHED

Electronic data bases consulted (Name of data base and where practicable terms used):

APS, JPABS, Proquest IEEE CD

L1	3 S EMULATED (3A) LAN
L2	6 S (4,787,082 OR 4,818,984 OR 4,868,866 OR 4,926,420 OR 4,9
30,	
L3	52 S ATM (P) LAN
L4	71 S LAN# (P) PROTOCOL# (P) (CONVERSION OR CONVERT###)
L5	7 S ATM AND L4