




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/00, 12/50, 12/54	A1	(11) International Publication Number: WO 97/04552 (43) International Publication Date: 6 February 1997 (06.02.97)
(21) International Application Number: PCT/US96/11929 (22) International Filing Date: 18 July 1996 (18.07.96) (30) Priority Data: 60/001,498 19 July 1995 (19.07.95) US (71) Applicant (for all designated States except US): ASCOM NEXION INC. [US/US]; 289 Great Road, Acton, MA 01720 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): CALDARA, Stephen, A. [US/US]; 220 Horsepond Road, Sudbury, MA 01776 (US). HAUSER, Stephen, A. [US/US]; 106 Farms Drive, Burlington, MA 01803 (US). MANNING, Thomas, A. [US/US]; 26 Summer Street, Northboro, MA 01532 (US). (74) Agents: LEBOVICI, Victor, B. et al.; Weingarten, Schurgin, Gagnebin & Hayes, Ten Post Office Square, Boston, MA 02109 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published With international search report. With amended claims.

(54) Title: POINT-TO-MULTIPOINT TRANSMISSION USING SUBQUEUES**(57) Abstract**

An Asynchronous Transfer Mode switch and method which facilitate point-to-multipoint transmission are disclosed. Each input port (24) within the switch includes a switch allocation table ("SAT") (20) which manages bandwidth allocation and delay. Each SAT (20) includes a plurality of sequentially ordered cell time slots (50), a subqueue (60) and a synchronized pointer (52), which is directed to one of the slots (50) such that at any given point in time each of the pointers is directed to the same slot location in the respective SAT with which the pointer is associated. To execute point-to-multipoint operation where output port conflicts are present, the switch transmits copies of the cell to available the output ports (30) at different points in time.

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.

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

- 1 -

POINT-TO-MULTIPOINT TRANSMISSION USING SUBQUEUEUESFIELD OF THE INVENTION

5 The present invention is generally related to telecommunications networks, and more particularly to bandwidth allocation and delay management within an asynchronous transfer mode switch.

RELATED APPLICATION

10 A claim of priority is made to provisional application 60/001,498 entitled COMMUNICATION METHOD AND APPARATUS, filed July 19, 1995.

BACKGROUND OF THE INVENTION

15 Telecommunications networks such as asynchronous transfer mode ("ATM") networks are used for transfer of audio, video and other data. ATM networks deliver data by routing data units such as ATM cells from source to destination through switches. Switches include input/output ("I/O") ports through which ATM cells are received and transmitted. The appropriate output port for transmission of the cell is determined based on the cell header.

20 One problem associated with ATM networks is loss of cells. Cells are buffered within each switch before being routed and transmitted from the switch. More particularly, switches typically have buffers at either the inputs or outputs of the switch for temporarily storing cells prior to transmission. As network traffic increases, there is an increasing possibility that buffer space may be inadequate and data lost. If the buffer size is insufficient, cells are lost. Cell loss causes undesirable interruptions in audio and video data transmissions, and may cause more serious damage to other types of data transmissions.

30 One particular situation where cell loss occurs is in point-to-multipoint transmission. In point-to-multipoint transmission a cell is transmitted from a single input to

35

- 2 -

multiple outputs across the switch fabric. In order to execute such a transmission, each of the designated outputs must be available to receive the cell from the transmitting input, i.e., have adequate buffer space. However, the likelihood that each of the designated outputs will be simultaneously prepared to receive the cell when the cell is enqueued decreases as traffic within the switch increases. In some circumstances this may result in delayed transmission. In the worst case, cells will be delayed indefinitely and discarded. It would therefore be desirable to facilitate point-to-multipoint transmission by reducing or eliminating delays and cell loss.

SUMMARY OF THE INVENTION

An Asynchronous Transfer Mode ("ATM") switch and method which facilitate point-to-multipoint transmission is disclosed. The ATM switch includes a plurality of input ports and a plurality of output ports. Each input port within the switch includes a switch allocation table ("SAT") which manages bandwidth allocation and delay. Each SAT includes a plurality of sequentially ordered cell time slots, a subqueue associated with each cell time slot and a pointer which is directed to one of the slots. The SAT pointers at each input port are synchronized such that, at any given point in time, each of the pointers is directed to the same slot location in the respective SAT with which the pointer is associated.

To execute point-to-multipoint operation where output port collisions are present the switch transmits copies of the cell to the output ports at different points in time. More particularly, the switch transmits copies of the cell to available output ports and tracks such transmission for managing future transmission to unserved output ports. To track transmission the switch includes a map and scoreboard. The map indicates which subqueues, and hence which output ports, are initially eligible to receive a copy of the cell.

- 3 -

The scoreboard indicates which subqueues, and hence which output ports, have received a copy of the cell. Tracking is realized by updating the scoreboard, i.e., toggling bits corresponding to serviced output ports, when copies of cells are transmitted. When subsequent opportunities to transmit arise as the SAT pointers are advanced, copies of the cell are not sent to already serviced output ports. When each designated output port has been serviced, the cell is dequeued.

Point-to-multipoint transmission is facilitated by transmitting copies of the cell at different points in time when output port collisions are detected. Due to cell traffic, a situation will sometimes arise where only a subset of the cells designated for receipt of a copy of the cell are available. In such a case, simultaneous transmission to each of the designated output ports is not possible. By transmitting at different points in time as the output ports become available, buffer space is more rapidly freed and the delay before entering the switch fabric is reduced for some or all of the cells. Traffic and delays are further reduced by tracking serviced output ports and transmitting only to unserviced output ports.

BRIEF DESCRIPTION OF THE DRAWING

These and other features and advantages of the present invention will become apparent from the following detailed description of the drawing in which:

Fig. 1 is a block diagram of a switch which facilitates point-to-multipoint operation;

Fig. 2 is a block diagram which illustrates operation of the switch allocation tables of Fig. 1;

Figs. 3A and 3B are block diagrams which further illustrates operation of the switch allocation tables; and

Fig. 4 is a flow diagram which illustrates a method for point-to-multipoint operation.

- 4 -

DETAILED DESCRIPTION OF THE DRAWING

Referring now to Fig. 1, the switch includes an NxN switch fabric 10, a bandwidth arbiter 12, a plurality of to switch port processors ("TSPP") 14, a plurality of from switch port processors ("FSPP") 16, a plurality of multipoint topology controllers ("MTC") 18 and a plurality of switch allocation tables ("SAT") 20. The NxN switch fabric, which may be an ECL crosspoint switch fabric, is used for cell data transport, and yields Nx670 Mbps throughput. The bandwidth arbiter controls switch fabric interconnection dynamically schedules momentarily unused bandwidth and resolves multipoint-to-point bandwidth contention. Each TSPP schedules transmission of cells 22 to the switch fabric from multiple connections. The FSPP receives cells from the switch fabric and transmits those cells onto output links. The switch allocation table controls crossbar input to output mapping, connection bandwidth and the maximum delay through the switch fabric.

In order to traverse the switch, a cell 22 first enters the switch through an input port 24 and is buffered in a queue 26 of input buffers. The cell is then transmitted from the input buffers to a queue 28 of output buffers in an output port. From the output port 30, the cell is transmitted outside of the switch, for example, to another switch. To facilitate traversal of the switch, each input port 24 includes a TSPP 14, and each output port 30 includes an FSPP 16. The TSPPs and FSPPs each include cell buffer RAM 32 which is organized into queues 26, 28. All cells in a connection pass through a single queue at each port, one at the TSPP and one at the FSPP, for the life of the connection. The queues thus preserve cell ordering by handling only one connection per queue. This strategy also allows quality of service ("QoS") guarantees on a per connection basis.

Three communication paths are used to implement flow control within the switch via probe and feedback messages:

- 5 -

a Probe Crossbar 42, an XOFF Crossbar 44 and an XON Crossbar 46. The Probe Crossbar 42 is an NxN crosspoint switch fabric which is used to transmit an FSPP multiQueue Number to each FSPP. The multiQueue number identifies a plurality of destination queues for the cell for use in point-to-multipoint connections. The FSPP uses the queue number or multiqueue number to direct a probe to the appropriate output queue or queues and thereby determine if there are enough output buffers available in the destination queues for receipt of the cell or cells.

The XOFF Crossbar 44 is an NxN serial crosspoint switch fabric which is used to communicate "Don't Send" type messages from the FSPP 16 to the TSPP 14. Each TSPP includes a scheduling list 47 of cells to be transmitted. The first bit of the feedback message, namely XOFF, is asserted to halt transmission of request message probes from a particular TSPP, and is thus a state control bit which puts the receiving TSPP scheduling list in an XOFF state, meaning that TSPP Scheduling List 47 will not use bandwidth. XOFF is asserted if the FSPP destination queue is full. The TSPP Scheduling List then remains in the XOFF state until receiving an XON message. The second bit, namely REJECT, is asserted when the entire pool of output buffers at one or more FSPP has been exhausted. The FSPPs maintain multiple buffer pools for different connection groups. The TSPP responds to an asserted REJECT feedback message by not dequeuing the cell through the data crossbar 10. A null cell denoted by a complemented CRC, is transmitted instead. The TSPP responds to an asserted XOFF feedback message by modifying the TSPP XOFF state bits. The XOFF state bits prevent the TSPP from attempting to send a request message from any queue on that Scheduling List until notified by the FSPPs that cell buffers are available.

The XON Crossbar 46 is an NxN serial contention-based switch which is used to communicate "Enable Send" type messages. More particularly, the XON Crossbar is employed

- 6 -

to communicate the XON message from the FSPP to the TSPP. When the number of cells in the FSPP destination queue increases above a predetermined threshold the FSPP asserts an XOFF and sets a state bit in a queue descriptor indicating that XOFF is asserted. When the number of cells in the FSPP queue drops below an XON threshold, the XON message is sent from the FSPP to the TSPP. The XON message enables the TSPP Scheduling List to resume sending request messages.

The Probe & XOFF communication paths operate in a pipeline fashion. First, the TSPP 14 selects a scheduling list, and information associated with that scheduling list is used to determine output ports 30 for transmission, i.e., a destination output queue. The bandwidth arbitrator 12 reduces this information to a TSPP to FSPP connectivity map which is employed to control the Probe, XOFF, and data cross-points in sequence. More particularly, the FSPP Queue number or multiQueue number is transmitted to the FSPP using the Probe crossbar 42. The FSPP then tests for buffer availability, and asserts REJECT and/or XOFF on the XOFF crossbar 44 if sufficient buffers are not available. The TSPP then transmits a null cell if REJECT was asserted. If XOFF was asserted, the TSPP puts the Scheduling List into the XOFF state. If sufficient buffers are available, the TSPP transmits the cell to the FSPP output queue through the data crossbar 47.

Referring now to Figs. 1 & 2, each TSPP within the switch includes an SAT 20 which manages bandwidth allocation and delay. The SAT is the basic mechanism behind cell scheduling. Each SAT 20 includes a plurality of sequentially ordered cell time slots 50 and a pointer 52 which is directed to one of the slots. All of the pointers in the switch are synchronized such that at any given point in time each of the pointers is directed to the same slot location in the respective SAT with which the pointer is associated, e.g., the first slot. In operation, the pointers are advanced in lock-step, each slot being active for 32 clock cycles at 50

- 7 -

MHz. When the pointer is directed toward a slot, the TSPP uses the corresponding entry 51 in the SAT to obtain a cell for launching into the switch fabric 10. For example, SAT 1 will transmit a cell from connection "a" to output ports "2" and "3" with the pointer in the position illustrated in Fig. 2.

Each of the counters is incremented once for each cell time, and the pointer returns to the first slot after reaching the last slot. Given an SAT depth of 8k, which defines a frame, the pointers scan the SATs every 6msec, thereby providing a maximum delay for transmission opportunity of 6msec. The delay can be decreased by duplicating a given entry at a plurality of slots within the SAT. The maximum delay that an incoming cell will experience corresponds to the number of slots 54 between the pointer and the slot containing the entry which specifies the destination of the cell. When multiple entries are made in order to decrease the maximum possible number of separating slots, the duplicate entries are preferably spaced equidistantly within the SAT, as illustrated. Hence, maximum delay for transmission opportunity corresponds to the frequency and spacing of duplicate entries within the SAT.

The amount of bandwidth allocated to a particular connection corresponds to the frequency at which a given entry appears in the SAT. Each slot 50 provides 64Kbps of bandwidth. Since the pointers cycle through the SATs at a constant rate, the total bandwidth allocated to a particular connection is equal to the product of 64Kbps and the number of occurrences of that entry. For example, entry (a,2,3), which occurs in three slots, is allocated 192Kbps of bandwidth.

Referring now to Figs. 1, 3A & 3B, in point-to-multipoint operation a situation will sometimes arise wherein it is not possible to transmit the cell to each output port designated by the entry, typically because of output port collisions due to other traffic within the switch. To

- 8 -

execute point-to-multipoint operation in such a situation, the switch transmits the cells to the designated output ports at different points in time. More particularly, the switch transmits the cell to the designated output ports as those ports become available. For example, connection "a" is designated for ports 2, 3 & 4, but the transmission to port 2 is in collision with connection "b" from SAT 2 at a first cell time and the transmission to port 3 is in collision with connection "c" from SAT 3 even though the transmission to port 4 is not in collision. To resolve this situation the transmission to port 4 is executed and the transmission to ports 2 and 3 is delayed. Repeat transmission to an already serviced port is prevented by tracking the serviced ports. Tracking also provides an indication of when all designated output ports have been serviced.

The lookup table is employed to determine the appropriate output ports associated with the entry, e.g., "a (2, 3)." The lookup table includes an area corresponding to each input port (here quadrants), and each quadrant includes words stored therein. The input port number is used to select the appropriate quadrant in the lookup table. The scheduling list number is then used to select a subgroup of eight 16-bit words (for a 16 port switch). Finally, the subqueue number is used to select one word from the subgroup; the word which indicates the output ports for transmission.

In order to track point-to-multipoint transmission executed at different points in time, the switch includes a map 56 and scoreboard 58, and the SAT 20 includes a three bit subqueue 60 field associated with each slot. The map and scoreboard are each 8 bits long and form bitmasks. The map is initialized according to the subqueues which act as offsets into the map. For example, subqueue "110" indicates an offset of 6 and subqueue "100" indicates an offset of 4. Hence, bits 4 and 6 in the map are initialized to logic "1" and the other bits are initialized to logic "0." The scoreboard is initialized to all logic "0" bits.

- 9 -

Tracking is realized by updating the scoreboard when copies of cells are transmitted. When the pointer reaches the slot corresponding to the connection for the enqueued point-to-multipoint cell, e.g., "a (4)," the offset indicated by the subqueue (110 indicating 6) is employed to set bit 6 of the scoreboard. The scoreboard is then compared to the map and if a complete match is determined then each cell has been transmitted and the scoreboard is cleared. If the map and scoreboard do not match completely then the point-to-multipoint transmission has not been made to each designated output queue. The updated scoreboard and map thus track which designated output ports have received copies of the cell.

Fig. 4 illustrates a method for point-to-multipoint operation. In an initial step 70 the need for a point-to-multipoint connection with allocated bandwidth is identified. In a following step 72 output port collisions are identified. If there are no collisions, the copies of the cell can be transmitted to each designated output port simultaneously. If there are collisions, the point-to-multipoint transmission is broken into pieces to avoid the output port collisions. To do so, the map and scoreboard are initialized 74. Flow then continues to a transmission portion 76 to transmit each copy and determine when each designated output port has been serviced.

The transmission portion is looped until each output port has been serviced. In step 78 the scheduling list number and subqueue number are retrieved from the SAT. A determination 80 is then made as to whether a cell is available to transmit. If no cell is available, other processing takes place. If a cell is available, the output ports are determined by indexing into the bit vector lookup table with the entry in the subqueue in step 82. Copies of the cell are then transmitted to the available output ports in step 84. The subqueue bit is then set in step 86 and a determination 88 is made as to whether the subqueue

- 10 -

scoreboard is equal to the subqueue map, i.e., whether all of the designated output ports have been serviced. If the subqueue scoreboard is equal to the subqueue map, then all output ports have been serviced and the scoreboard is cleared 5 90 and the cell is dequeued 92. If the subqueue scoreboard is not equal to the subqueue map, i.e., some designated output ports remain unserved, flow returns to step 78.

Having described the preferred embodiments of the invention, it will now become apparent to one of skill in 10 the art that other embodiments incorporating their concepts may be used. It is felt therefore that the invention should not be limited to disclosed embodiments, but rather should be limited only by the spirit and scope of the appended claims.

- 11 -

CLAIMS

What is claimed is:

1. A switch for receiving a data unit, and transmitting the data unit to a plurality of designated destinations, comprising:

a plurality of input ports operative to receive the data unit, the data unit being received in one of said plurality of input ports;

a data unit buffer memory associated with said plurality of input ports for buffering the data unit received in said one of said plurality of input ports;

a plurality of output ports operative to receive copies of the data unit from said one of said plurality of input ports and transmit the copies of the data unit to at least one of the plurality of designated destinations;

a data unit buffer memory associated with said plurality of output ports for buffering the copy of the data unit received in said output ports; and

a tracking device operative to track receipt of the copies of the data unit by said output ports,

whereby the data unit received in said input port may be transmitted to selected output ports at different points in time without executing redundant transmission to any given selected output port.

2. The switch of claim 1 further including a scoreboard bitmask having a plurality of bits, each bit corresponding to one of said plurality of subqueues, wherein each subqueue is associated with one or more output ports.

3. The switch of claim 2 further including a map bitmask having a plurality of bits, each bit corresponding to one of said plurality of subqueues, wherein each subqueue is associated with one or more output ports and said bits are set corresponding to eligible output ports.

- 12 -

4. The switch of claim 3 further including a switch allocation table associated with each input port of said plurality of input ports, said switch allocation table having a plurality of entries and a pointer directed toward one of said plurality of entries, each switch allocation table entry being associated with a connection identifier and at least one output port identifier, said pointer being incremented through said entries at a predetermined rate.

5. The switch of claim 4 wherein said switch allocation table includes a subqueue field having an entry associated with each input port of said plurality of input ports, said subqueue entries including an index to a table of unserved output ports.

6. The switch of claim 5 further including a bit vector lookup table including output port identifiers, said bit vector lookup table providing identifiers of unserved output ports for a given connection when indexed by said subqueue field entry.

7. A method for point-to-multipoint transmission of a data unit from a telecommunications switch having a plurality of input ports and a plurality of output ports comprising the steps of:

receiving the data unit in the switch;
designating a plurality of the output ports for receipt of copies of the data unit;

determining which output ports of the plurality of output ports are available for receipt of a copy of the data unit;

transmitting a copy of the data unit to each output port identified as being available for receipt of a copy of the data unit; and

delaying transmission of a copy of the data unit to each output port identified as not being available for receipt of

- 13 -

a copy of the data unit until a later point in time.

8. The method of claim 7 including the further step of initializing a scoreboard of unserviced output ports designated in said designating step.

9. The method of claim 8 including the further step of retrieving a subqueue index entry from a switch allocation table.

10. The method of claim 9 including the further step of indexing a bit vector lookup table with the retrieved subqueue index entry to obtain at least one output port identifier.

11. The method of claim 11 including the further step of setting the subqueue index entry based upon the copies of the data unit transmitted in said transmitting step.

12. The method of claim 11 including the further step of comparing the subqueue to the scoreboard to determine whether each designated output port has been serviced.

13. The method of claim 12 including the further step of updating the scoreboard to indicate which output ports have been serviced.

14. The method of claim 13 including the further step of clearing the scoreboard if each designated output port has been serviced.

- 14 -

15. The method of claim 14 including the further step of dequeuing the data unit when each designated output port has been serviced.

AMENDED CLAIMS

[received by the International Bureau on 19 November 1996 (19.11.96);
original claim 7 amended; new claims 16-27 added;
remaining claims unchanged (6 pages)]

1. A switch for receiving a data unit, and transmitting the
5 data unit to a plurality of designated destinations,
comprising:

a plurality of input ports operative to receive the data
unit, the data unit being received in one of said plurality
of input ports;

10 a data unit buffer memory associated with said plurality
of input ports for buffering the data unit received in said
one of said plurality of input ports;

a plurality of output ports operative to receive copies
of the data unit from said one of said plurality of input
15 ports and transmit the copies of the data unit to at least
one of the plurality of designated destinations;

a data unit buffer memory associated with said plurality
of output ports for buffering the copy of the data unit
received in said output ports; and

20 a tracking device operative to track receipt of the
copies of the data unit by said output ports,

whereby the data unit received in said input port may
be transmitted to selected output ports at different points
in time without executing redundant transmission to any given
25 selected output port.

2. The switch of claim 1 further including a scoreboard
bitmask having a plurality of bits, each bit corresponding
to one of said plurality of subqueues, wherein each subqueue
30 is associated with one or more output ports.

3. The switch of claim 2 further including a map bitmask
having a plurality of bits, each bit corresponding to one of
said plurality of subqueues, wherein each subqueue is
35 associated with one or more output ports and said bits are
set corresponding to eligible output ports.

4. The switch of claim 3 further including a switch allocation table associated with each input port of said plurality of input ports, said switch allocation table having a plurality of entries and a pointer directed toward one of said plurality of entries, each switch allocation table entry being associated with a connection identifier and at least one output port identifier, said pointer being incremented through said entries at a predetermined rate.

5. The switch of claim 4 wherein said switch allocation table includes a subqueue field having an entry associated with each input port of said plurality of input ports, said subqueue entries including an index to a table of unserved output ports.

6. The switch of claim 5 further including a bit vector lookup table including output port identifiers, said bit vector lookup table providing identifiers of unserved output ports for a given connection when indexed by said subqueue field entry.

7. A method for point-to-multipoint transmission of a data unit from a telecommunications switch having a plurality of input ports and a plurality of output ports comprising the steps of:

receiving the data unit in the switch;
designating a plurality of the output ports for receipt of copies of the data unit;

determining which output ports of the plurality of output ports are available for receipt of a copy of the data unit;

transmitting a copy of the data unit to each output port identified in said determining step as being available for receipt of a copy of the data unit; and

delaying transmission of a copy of the data unit to each output port identified in said determining step as not being

- 17 -

available for receipt of a copy of the data unit until a later point in time.

5 8. The method of claim 7 including the further step of
initializing a scoreboard of unserviced output ports
designated in said designating step.

10 9. The method of claim 8 including the further step of
retrieving a subqueue index entry from a switch allocation
table.

15 10. The method of claim 9 including the further step of
indexing a bit vector lookup table with the retrieved
subqueue index entry to obtain at least one output port
identifier.

20 11. The method of claim 11 including the further step of
setting the subqueue index entry based upon the copies of the
data unit transmitted in said transmitting step.

25 12. The method of claim 11 including the further step of
comparing the subqueue to the scoreboard to determine whether
each designated output port has been serviced.

30 13. The method of claim 12 including the further step of
updating the scoreboard to indicate which output ports have
been serviced.

35 14. The method of claim 13 including the further step of
clearing the scoreboard if each designated output port has
been serviced.

- 18 -

15. The method of claim 14 including the further step of dequeuing the data unit when each designated output port has been serviced.

5

16. An asynchronous transfer mode ("ATM") switch for transmitting copies of a cell received at an input port of said switch to a plurality of output ports of said switch through a switch fabric, comprising:

10

a clock generator operative to generate cell time intervals; and

15

at least one cell transmission controller for controlling transmission of copies of the cell over a plurality of cell time intervals when output port conflicts prevent transmission of copies of the cell to said plurality of output ports during a single time interval.

20

17. The ATM switch of Claim 16 further including a tracking device for tracking which output ports of said plurality of output ports have received a copy of the cell, thereby avoiding transmission of redundant copies of the cell to already-serviced output ports.

25

18. The ATM switch of Claim 17 further including an output port map indicating which output ports are initially eligible to receive a copy of the cell.

30

19. The ATM switch of Claim 18 further including a scoreboard bitmask indicating which output ports have received a copy of the cell.

35

20. The ATM switch of Claim 19 wherein said scoreboard is updated for each cell time interval in which a copy of the cell is transmitted.

- 19 -

21. The ATM switch of Claim 20 further including a plurality of switch allocation tables ("SATs"), each SAT including a plurality of sequentially ordered cell time slots, a subqueue associated with each cell time slot, and a synchronized pointer directed toward one of said cell time slots such that at any given point in time each said synchronized pointer is directed to a corresponding slot in the respective SAT with which said pointer is associated.

22. The ATM switch of Claim 21 further including a lookup table having a respective area associated with each said input port, said input ports being identified by input port numbers.

23. The ATM switch of Claim 22 further including a plurality of scheduling lists of cells to be transmitted through said switch.

24. The ATM switch of Claim 23 wherein said transmission control device

a) utilizes said input port number to select an area of said lookup table associated with the input port transmitting the cell,

b) employs the scheduling list number to select a subgroup of words from said area selected with the input port number, and

c) employs the subqueue number to select a single word from said subgroup, wherein said word indicates output ports for transmission of copies of the cell.

25. The ATM switch of Claim 24 wherein said subqueue field contains a three bit number.

26. The ATM switch of Claim 25 wherein said output port map and said scoreboard are each eight bits in length, respectively.

- 20 -

27. The ATM switch of Claim 26 wherein said subqueues provide an offset into said output port map for initialization thereof.

5

AMENDED SHEET (ARTICLE 19)

1 / 5




FIG. 1




FIG. 2

3 / 5




FIG. 3A

4 / 5




FIG. 3B

5 / 5

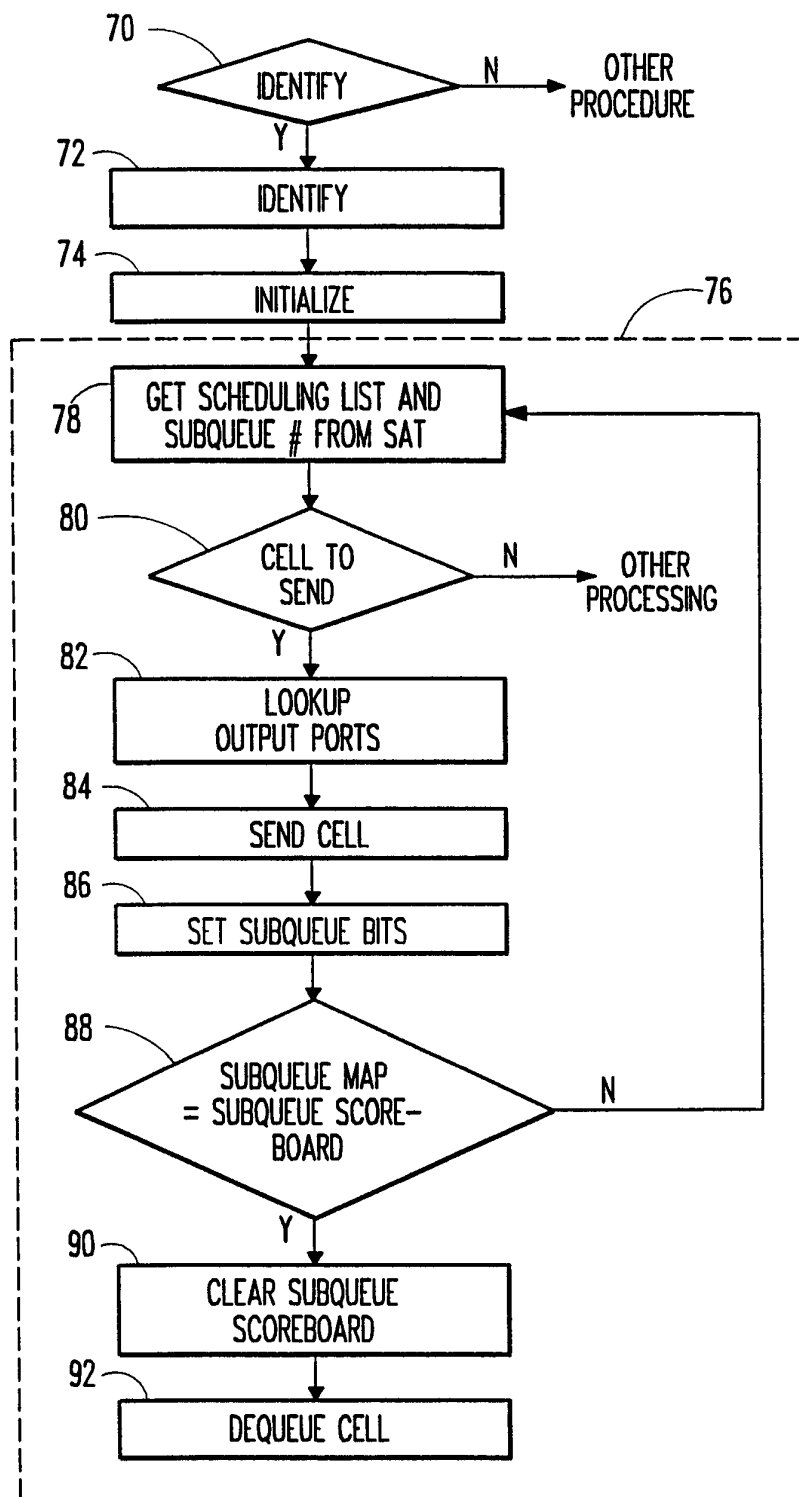


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/11929

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : HO4L 12/00 12/50 12/54

US CL : 370/60.1, 61

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. : 370/58.1-58.3, 60, 60.1, 61, 94.1, 94.2; 340/825.89

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)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	US, A, 5,448,559 (HAYTER) 05 SEPTEMBER 1995, col.2, lines 10-68, col.3, lines 24-66, col.4, lines 22-55.	1-14
X,P	US, A, 5,455,820 (YAMADA) 03 OCTOBER 1995, col. 2, lines 23-61, col.3, lines 17-34, col.3, line 46 - col.4, line 38.	1-14
X,P	US, A, 5,475,679 (MUNTER) 12 DECEMBER 1995, col.2, line 41 - col.3, line 13, col.4, line 23 - col.5, line 9.	1-14



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)	"&" 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

26 SEPTEMBER 1996

Date of mailing of the international search report

15 OCT 1996

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

Facsimile No. (703) 305-3230

Authorized officer

RICKY Q. NGO

Telephone No. (703) 305-4798