(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 March 2009 (05.03.2009)

PCT

(10) International Publication Number
WO 2009/029817 A1

(51) International Patent Classification:
H04L 12/56 (2006.01)

(21) International Application Number:
PCT/US2008/074828

(22) International Filing Date: 29 August 2008 (29.08.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/967,175 31 August 2007 (31.08.2007) US

(71) Applicant (*for all designated States except US*): **NEW JERSEY INSTITUTE OF TECHNOLOGY** [US/US]; 340 Fenster Hall, University Heights, Newark, New Jersey 07102 (US).

(72) Inventors; and
(75) Inventors/Applicants (*for US only*): **DONG, Ziqian** [CN/US]; 17 E. Midland Avenue, Apt. 1B, Kearny, New Jersey 07032 (US). **ROJAS-CESSA, Roberto** [MX/US]; 936 70th Street, 1st Floor, Brooklyn, New York 11228 (US).

(74) Agents: **HUNTER, Paul, S.** et al.; Foley & Lardner LLP, 150 East Gilman Street, P.O. Box 1497, Madison, Wisconsin 53701-1497 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL,

[Continued on next page]

(54) Title: REPLICATING AND SWITCHING MULTICAST INTERNET PACKETS IN ROUTERS USING CROSSPOINT MEMORY SHARED BY OUTPUT PORTS

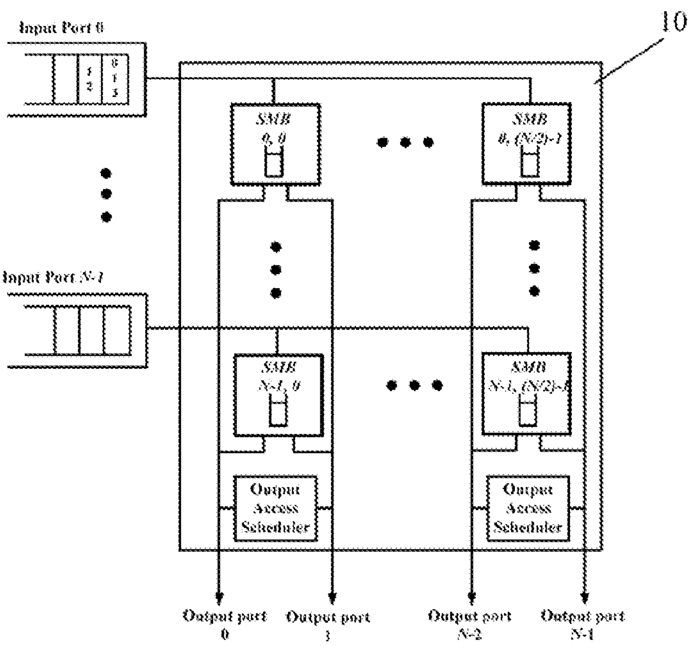

Figure 1

(57) Abstract: Multicast traffic is expected to increase in packet networks, and therefore in switches and routers, by including broadcast and multimedia-on-demand services. Combined input-crosspoint buffered (CICB) switches can provide high performance under uniform multicast traffic. However this is often at the expense of N2 crosspoint buffers. An output-based shared-memory crosspoint-buffered (O-SMCB) packet switch is used where the crosspoint buffers are shared by two outputs and use no speedup. An embodiment of the proposed switch provides high performance under admissible uniform and non-uniform multicast traffic models while using 50% of the memory used in CICB switches that has dedicated buffers. Furthermore, the O-SMCB switch provides higher throughput than an existing SMCB switch where the buffers are shared by inputs.

WO 2009/029817 A1

NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

REPLICATING AND SWITCHING MULTICAST INTERNET PACKETS IN ROUTERS USING CROSSPOINT MEMORY SHARED BY OUTPUT PORTS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority to U.S. Provisional Patent Application No. 60/967,175 filed August 31, 2007, the contents of which are incorporated herein by reference.

FIELD

[0002] The present disclosure relates to the field of packet-oriented networks. The present disclosure provides a method for replicating and switching multicast Internet packets in routers using cross-point memory shared by output ports.

BACKGROUND

[0003] This background section provides a context for the disclosure. The description herein may include concepts that could be pursued but are not necessarily ones that have been previously conceived or pursued. The description is not intended to be limiting and unless otherwise stated, nothing in this section is admitted as prior art simply by inclusion in this section.

[0004] The migration of broadcasting services, such as cable television and multimedia-on-demand, to packet-oriented networks as well as the embracing of emerging applications, such as teleconferencing and storage networks by the Internet, place significant traffic demands on the Internet. To keep up with such bandwidth demand, packet switches and routers need to provide efficient multicast switching and packet replication.

[0005] The forwarding of packets in the Internet depends on the switching efficiency of routers and switches. Presently, there are different buffering strategies to build packet switches. Input buffered (IB) switches provide limited throughput and require complex scheduling schemes. Output buffered (OB) switches offer high throughput but require infeasible memory speedup. Combined input crosspoint buffered (CICB) switches deliver better switching performance than input buffered switches, but the memory amount in the buffered crossbar is large (the number of crosspoint buffers equals N^2, where N is the number of ports).

[0006] In addition to the high performance of CICB packet switches under unicast traffic, the crosspoint buffers in these switches help to provide call splitting intrinsically. Different from IB switches, CICB switches do not require cell transmission after inputs and outputs have been matched. In CICB switches, one input can send up to one (multicast) cell to the crossbar, and one or more cells destined to a single output port can be forwarded from multiple inputs to the crossbar at the same time slot. Therefore, CICB switches have natural properties favorable for multicast switching as contending copies for a single output can be sent to the crosspoint buffers from several inputs at the same time slot without blocking each other.

[0007] Existing packet switches mostly target unicast traffic. It is expected that in the near future multicast and broadcast services will demand most of the available bandwidth of the Internet. Current packet switches provide limited services for multicast traffic. CICB switches are seen as the promising architecture for building efficient switches. However, multicast packets can easily exhaust the available internal buffers in current CICB switches. Handling and management of multicast traffic by switches and routers can become very expensive as the sheer amount and speed of memory necessary often becomes cost prohibitive if not simply infeasible. To lower the cost of memory implementation, the required memory amount should be reduced or kept to a minimum while keeping high efficiency in packet switching and replication.

SUMMARY

[0008] Various embodiments relate to switch architecture and selection schemes for efficient replication and switching of multicast Internet packets. In at least one embodiment, switch buffered crosspoints are utilized to allow for sharing of memory by different outputs.

[0009] At least one embodiment is based on a CICB switch that uses crosspoint buffers shared by two or more output ports. Buffer sharing by two output ports is used to achieve high performance and simplify the implementation.

[0010] At least one embodiment includes a method of scheduling access to the shared crosspoint buffers by output ports and a flow control mechanism to avoid buffer underflow and overflow. The resulting switch architecture is referred to as shared-memory crosspoint-buffered switch with output-based sharing (O-SMCB).

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] To assist those of ordinary skill in the relevant art in making and using the subject matter hereof, reference is made to the appended drawings, wherein:

[0012] Fig. 1 is a schematic drawing of the $N \times N$ output-based shared-memory crosspoint buffered (O-SMCB) switch with shared-memory crosspoints by outputs.

[0013] Fig. 2 is a schematic drawing of the $N \times N$ input-based shared-memory crosspoint buffered (I-SMCB) switch with shared-memory crosspoints by inputs.

DETAILED DESCRIPTION

[0014] Fig. 1 illustrates a representative output-based shared-memory crosspoint buffered (O-SMCB) switch 10. The O-SMCB switch 10 requires less memory than a CICB switch to achieve comparable performance under multicast traffic and no speedup. Furthermore, the O-SMCB switch provides higher throughput under uniform and non-uniform multicast traffic models than an input-based SMCB (I-SMCB) switch (Fig. 2), where two inputs share crosspoint buffers. The O-SMCB switch 10 includes input ports 0 to N-1; shared memory buffers (SMB); output access schedulers; and output ports 0, 1, ..., N-2, and N-1. The input ports receive input from a first-in first-out (FIFO) queue at each input. The switch 10 has N^2 crosspoints and $\frac{N^2}{2}$ crosspoint buffers in the crossbar. A crosspoint in the buffered crossbar connects input port i to output j is denoted as $CP(i,j)$. The buffer shared by $CP(i,j)$ and $CP(i, j')$ that stores cells for output ports j or j', where $j \neq j'$, is denoted as $SMB\ (i,q)$, where $0 \leq q \leq \frac{N}{2} - 1$.

[0015] According to at least one embodiment, the O-SMCB switch 10 uses round-robin selection in its arbitration schemes and achieves a high throughput under multicast traffic.

This selection scheme was adopted for its simplicity and as an example. Other selection schemes can also be used. In another embodiment, the O-SMCB switch 10 is provisioned with one multicast first-in first-out (FIFO) queue at each input. For this embodiment, an even N value is assumed for the sake of clarity. However, an odd N value can be used with one input port using dedicated buffers of size 0.5 to 1.0 the size of an SMB. The size of an SMB, in number of cells that can be stored, is k_s. The case of minimum amount of memory, or when $k_s = 1$ (equivalent to having 50% of the memory in the crossbar of a CICB switch), is considered. Therefore, *SMB(i,q)* with $k_s = 1$ can store a cell that can be directed to either j or j'. The SMB has two egress lines, one per output.

[0016] To avoid the need for speedup at SMBs, only one output is allowed to access an SMB at a time. The access to one of the N SMBs by each output is decided by an output-access scheduler. A scheduler performs a match between SMBs and the outputs that share them by using round-robin selection. There are $\frac{N}{2}$ output-access schedulers in the buffered crossbar, one for each pair of outputs. Multicast cells at the inputs have an N-bit multicast bitmap to indicate the destination of the multicast cells. Each bit of the bitmap is denoted as D_j, where $D_j = \begin{cases} 1 & \text{if output } j \text{ is one of the cell destination,} \\ 0 & \text{otherwise.} \end{cases}$

[0017] Each time a multicast copy is forwarded to the SMB for the cell's destination, the corresponding bit in the bitmap is reset. When all bits of a multicast bitmap are zero, the multicast cell is considered completely served. Call splitting is used by this switch to allow effective replication and to alleviate a possible head-of-line blocking. A credit-based flow control mechanism is used to notify the inputs about which output replicates a multicast copy and to avoid buffer overflow.

[0018] Much research has focused on unicast traffic, where each packet has a single destination. It has been shown that unicast switches achieve 100% throughput under admissible conditions, $\sum_i \lambda_{i,j} < 1$ and $\sum_j \lambda_{i,j} < 1$ where i is the index of inputs ($0 \le i \le N$-1), j is the index of outputs ($0 \le i \le N$- 1) for an $N \times N$ port switch, and $\lambda_{i,j}$ is the data rate from input i to output j, in a plethora of switch architectures and switch configuration schemes.

[0019] Although it is difficult to describe actual multicast traffic models, switches of this type also should provide 100% throughput under admissible multicast traffic. In multicast switches, the admissibility conditions are similar to those for unicast traffic, however, with the consideration of the fanout of multicast packets. The fanout of a multicast packet is the number of different destinations that expect copies of the packet. This implies that the average fanout of multicast traffic increases the average output load of a switch. Therefore, the average output load in a multicast switch is proportional to the product of the average input load and the average fanout for a given multicast traffic model.

[0020] In at least one embodiment, incoming variable-size packets are segmented into fixed-length packets, also called cells, at the ingress side of a switch and being re-assembled at the egress side, before the packets leave the switch. Therefore, the time to transmit a cell from an input to an output takes a fixed amount of time, or time slot. Also, cell replication is performed at the switch fabric by exploiting its space capabilities. One embodiment focuses on crossbar-based switches in order to allow multicast cells to be stored in a single queue at the input.

[0021] Multicast switching has been largely considered for input buffered (IB) switches. In these switches, matching is performed between inputs and outputs to define the configuration on a time-slot basis. This matching process can be complex when considering multicast traffic. Combined input crosspoint-buffered (CICB) packet switches have higher performance than IB switches at the expense of having crosspoint buffers, which run at the same speed as the input buffers in an IB switch, under unicast traffic. In these switches, an input might have up to N buffers where each one stores cells destined to a particular output to avoid head-of-line blocking. The crosspoint buffers in CICB switches can be used to provide call splitting (or fanout splitting) intrinsically. Different from IB switches, CICB switches are not required to have cell transmission after inputs and output have been matched. This feature makes CICB switches attractive for implementation. In CICB switches, one input can send up to one (multicast) cell to the crossbar, and one or more cells destined to a single output port can be forwarded from multiple inputs to the crossbar at the same time slot.

[0022] Therefore, CICB switches have natural properties favorable for multicast switching as contending copies for a single output can be sent to the crosspoint buffers from several inputs at the same time slot without blocking each other. In general, CICB switches have dedicated crosspoint buffers for each input-output pair, for a total of N^2 crosspoint buffers. Since memory used in the crosspoint buffers has to be fast, it is desirable to minimize the total amount this fast, expensive memory.

[0023] Fig. 2 shows an I-SMCB switch. The performance of the O-SMCB switch described with reference to Fig. 1 is compared to that of the I-SMCB switch. Models of a 16 × 16 O-SMCB switch and a 16 × 16 I-SMCB switch can be implemented in discrete-event simulation programs. The I-SMCB switch has an architecture similar to the O-SMCB switch, in that the SMBs are shared by (two) inputs. For a fair comparison, the I-SMCB also uses round-robin selections for SMB-access by inputs and for output arbitration. Simulation results are obtained with a 95% confidence interval and a standard error not greater than 5%.

[0024] Multicast traffic models are considered with uniform and nonuniform distributions and Bernoulli arrivals: multicast uniform, multicast diagonal with fanouts of 2 and 4, and broadcast. In the uniform multicast traffic model, multicast cells are generated with a uniformly distributed fanout among N outputs. For this traffic model, the average fanout is $\frac{1+N}{2}=\frac{17}{2}$ and a maximum admissible input load of $\frac{1}{fanout}=\frac{1}{8.5}$.

[0025] The diagonal multicast traffic model with a fanout of 2 has a destination distribution to $j=i$ and $j=(i+1)\%N$ for each multicast cell, and a maximum admissible input load of 0.5. The diagonal multicast traffic model with a fanout of 4 has the copies of a multicast cell destined to $j=\{i, (i+1)\%N, (i+2)\%N$, and $(i+3)\%N\}$\$ for each multicast cell, and its admissible input load is 0.25. A broadcast multicast cell generates copies for all N different outputs and has a maximum admissible load is 1/16=0.0625.

[0026] Under uniform traffic, the I-SMCB and O-SMCB switches deliver 100% throughput. Under multicast diagonal traffic with fanout of 2, the throughputs observed are 100% for the O-SMCB switch and 96% for the I-SMCB switch. Under multicast diagonal traffic with a fanout of 4, the performance of the I-SMCB switch decreases to 67%, while the performance of the O-SMCB switch remains high, close to 100%. Under broadcast traffic (fanout equal to N), the throughput of the O-SMCB switch is 99% while the throughput of the I-SMCB switch is 95%.

[0027] Multicast is a traffic type difficult to police for admissibility. Furthermore, the performance of switches under inadmissible traffic (produced by larger fanouts than the expected average) might change. In cases of unicast traffic, the maximum throughput of a switch can remain high with a fair scheduler. However, this might not be the case under multicast traffic. In this experiment, the input load is increased beyond the maximum admissible values in the considered traffic models to observe throughput changes of the O-SMCB and I-SMCB switches under these overload conditions.

[0028] Under uniform multicast traffic, the throughput of both switches degrades to 93% when the input load is larger than 0.117 (the output load is larger than 1.0). This throughput degradation occurs because of the increased number of contentions for SMB access as the traffic load increases. Under multicast diagonal traffic with fanout of 2, the throughput of the I-SMCB switch is reduced to 90% while the throughput of the O-SMCB switch remains close to 100%. Under multicast diagonal traffic with a fanout of 4, the throughput of the I-SMCB switch falls to 68% while the throughput of the O-SMCB switch also remains close to 100%. Under broadcast traffic, the throughput of the I-SMCB switch decreases to 79%. However, the throughput of the O-SMCB switch remains close to 100%.

Traffic type	*Ta(I)*	*Ta(O)*	*Ti(I)*	*Ti(O)*
Uniform	100%	100%	93%	93%
Diagonal 2	96%	100%	90%	100%
Diagonal 4	67%	100%	68%	100%
Broadcast	95%	99%	79%	100%

Table I

[0029] Table I summarizes the obtained throughput for all tested traffic models. In this table, *Ta* stands for the measured throughput under admissible traffic and *Ti* for the measured throughput under inadmissible traffic. The *I* in parenthesis indicates that the result is related to the I-SMCB switch and the *O* indicates that the result is related to the O-SMCB switch.

[0030] The embodiment described provides a switch architecture to support multicast traffic using a shared-memory switch that uses crosspoint buffers shared by outputs to use 50% of the memory amount in the crossbar fabric that CICB switches require. One embodiment of the proposed O-SMCB switch delivers high performance under multicast traffic while using no speedup. The switch has buffers that are shared by outputs instead of the inputs. This has the effect of facilitating call splitting by allowing inputs directly access to the crosspoint buffers.

[0031] This improvement has a significant impact on switching performance. As a result, the O-SMCB provides 100% throughput under both uniform multicast traffic and diagonal multicast traffic with fanouts of 2 and 4, all with Bernoulli arrivals. Furthermore, an embodiment of the proposed switch keeps the throughput high under nonuniform traffic with overloading conditions. The disadvantage of SMCB switches is that time relaxation of CICB switches is minimized because of the matching process considered. However, the matching is performed in chip and among a moderate number of outputs. Furthermore, the matching process is simpler in the SMCB switches than those used in IB switches for multicast traffic.

[0032] The representative method offers significant advantages relative to prior art. The advantageous properties and/or characteristics of the disclosed method include, but are not limited to, scalability, effectiveness, robustness, and efficiency.

[0033] The above-described embodiments may be implemented within the context of methods, computer readable media and computer program processes. As such, it is contemplated that some of the steps discussed herein as methods, algorithms and/or software processes may be implemented within hardware (e.g., circuitry that cooperates with a processor to perform various steps), software or a combination of hardware and software. The representative embodiments may be implemented as a computer program product wherein computer instructions, when processed by a computer, adapt the operation of the computer such that the methods and/or techniques are invoked or otherwise provided. Instructions for invoking the methods may be stored in fixed or removable media, transmitted via a data stream in a signal bearing medium such as a broadcast medium, and/or stored within a working memory or mass storage device associated with a computing device operating according to the instructions. Generally speaking, a computing device including a processor, memory and input/output means may be used to process software instructions, store software instructions and/or propagate software instructions to or from a communications channel, storage device or other computer/system.

[0034] Applicant has attempted to disclose all embodiments and applications of the disclosed subject matter that could be reasonably foreseen. However, there may be unforeseeable, insubstantial modifications that remain as equivalents. While the present invention has been described in conjunction with specific, exemplary embodiments thereof, it is evident that many alterations, modifications, and variations will be apparent to those skilled in the art in light of the foregoing description without departing from the spirit or scope of the present disclosure. Accordingly, the present disclosure is intended to embrace all such alterations, modifications, and variations of the above detailed description.

CLAIMS

What is claimed is:

1. A switch that supports multi-cast traffic in a network, the switch comprising:

input ports configured to receive data packets;

output ports configured to provide data packets to the network;

crosspoint memory buffers coupling the input ports and the output ports; and

at least one scheduler that matches crosspoint memory buffers and output ports, the at least one scheduler further configured to permit only one output port access to a particular crosspoint memory buffer at a time.

2. The switch of claim 1, wherein the data packets received by the input ports are variable size data packets.

3. The switch of claim 1, wherein there are $N/2$ schedulers where N is a number of output ports.

4. The switch of claim 1, wherein there are $\frac{N^2}{2}$ crosspoint memory buffers where N is a number of output ports.

5. The switch of claim 1, wherein the received data packets have a multi-cast fanout number of 2.

6. The switch of claim 1, wherein the received data packets are uniform multi-cast traffic data packets.

7. A method of replicating and switching multicast network packets using crosspoint memory shared by output ports, the method comprising:

segmenting incoming variable-size packets into fixed length packets at an ingress side of a switch;

matching shared memory buffers and output ports; and

re-assembling the fixed-length packets into outgoing variable-size packets at an egress side of the switch;

8. The method of claim 7, wherein the incoming variable-size packets have a multi-cast fanout number of 4.

9. The method of claim 7, wherein matching shared memory buffers and outputs that share them is done using a round-robin selection process.

10. The method of claim 7, further comprising scheduling access to shared memory buffers by output ports such that only one output port accesses a particular crosspoint memory buffer at a time.

11. The method of claim 7, further comprising providing call splitting with the incoming variable-size packets.

12. The method of claim 7, wherein multicast cells at the ingress side of the switch have an N-bit multicast bitmap to indicate a destination of the multicast cells.

13. The method of claim 12, wherein each bit of the N-bit bitmap is denoted as D_j, where $D_j = \begin{cases} 1 & \text{if output } j \text{ is one of the cell destination,} \\ 0 & \text{otherwise.} \end{cases}$

14. A computer program product wherein computer instructions, when processed by a computer, adapt the operation of the computer such that the computer:

segments incoming variable-size packets into fixed length packets at an ingress side of a switch;

matches shared memory buffers and output ports for communication of fixed length packets to output ports;

permits only one output port access to a particular shared memory buffer at a time; and

re-assembles the fixed-length packets into outgoing variable-size packets at an egress side of the switch.

15. The computer program product of claim 14, wherein the computer is further configured to match shared memory buffers and output ports using a round-robin selection process.

16. The computer program product of claim 14, wherein there are $\frac{N^2}{2}$ shared memory buffers where N is a number of output ports.

17. The computer program product of claim 16, wherein N is an even number.

18. The computer program product of claim 14, wherein multicast cells at the ingress side of the switch have an N-bit multicast bitmap to indicate a destination of the multicast cells.

19. The computer program product of claim 18, wherein each bit of the N-bit bitmap is denoted as D_j, where $D_j = \begin{cases} 1 & \text{if output } j \text{ is one of the cell destination,} \\ 0 & \text{otherwise.} \end{cases}$

20. The computer program product of claim 18, wherein each time a multicast copy is sent to a particular shared memory buffer, a corresponding bit in the N-bit bitmap is reset.

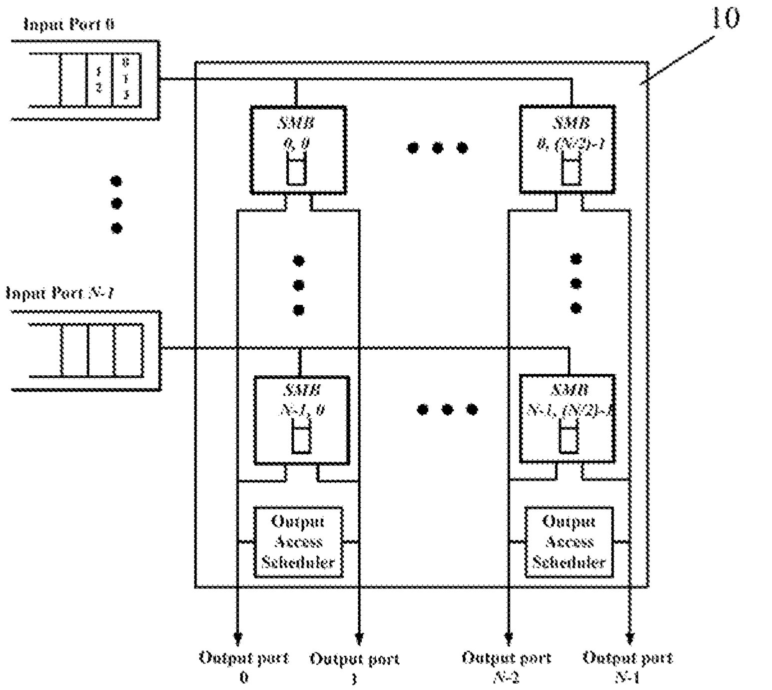

Figure 1

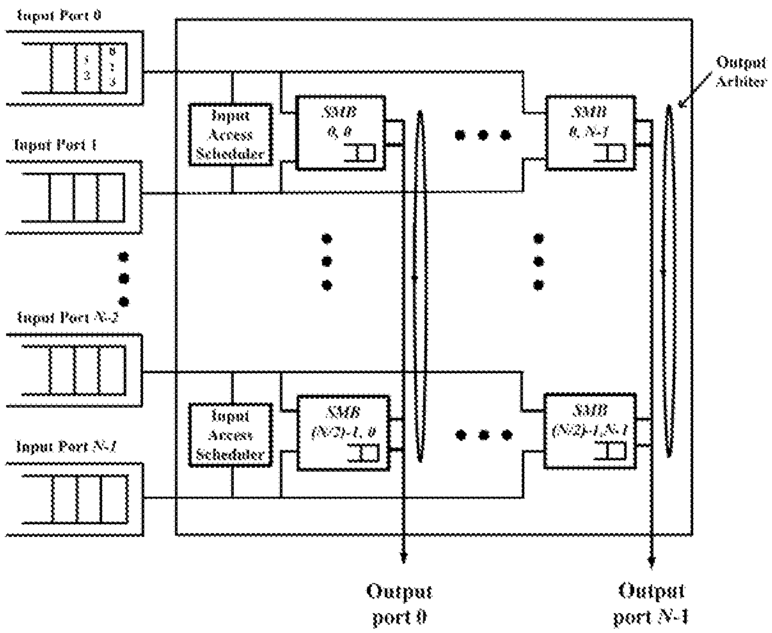

Figure 2

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2008/074828

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/56

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)
H04L

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 practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	US 2004/081184 A1 (MAGILL ROBERT B [US] ET AL) 29 April 2004 (2004-04-29) paragraphs [0025] - [0063]; figures 1,2	1,2,7,9,10,14,15 3-6,8,11-13,16-20
X Y	WO 94/17617 A (ERICSSON TELEFON AB L M [SE]) 4 August 1994 (1994-08-04) page 24, line 6 - page 26, line 10; figures 14-18 page 27, line 23 - page 28, line 11; figure 20	1,7,14 3,4,16,17

-/--

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"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)

"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

"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

"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

"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.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
5 December 2008	12/12/2008

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Kreppel, Jan

International application No
PCT/US2008/074828

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	US 6 151 301 A (HOLDEN BRIAN D [US]) 21 November 2000 (2000-11-21) column 5, line 52 - column 7, line 51; figure 5 column 11, line 21 - column 12, line 34	1,7,14 5,6,8, 11-13, 18-20
A	ROJAS-CESSA R ET AL: "CIXOB-K: COMBINED INPUT-CROSSPOINT-OUTPUT BUFFERED PACKET SWITCH" GLOBECOM'01. 2001 IEEE GLOBAL TELECOMMUNICATIONS CONFERENCE. SAN ANTONIO, TX, NOV. 25 - 29, 2001; [IEEE GLOBAL TELECOMMUNICATIONS CONFERENCE], NEW YORK, NY : IEEE, US, 25 November 2001 (2001-11-25), pages 2654-2660, XP001060622 ISBN: 978-0-7803-7206-1 the whole document	1-20

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2008/074828

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
   because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
   because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
   because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers allsearchable claims.

2. ☒ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search reportcovers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-6,14-20

   switching device adapted to provide a crosspoint memory buffer shared by multiple output ports

   ---

2. claims: 7-13,14-20

   method and computer program for segmenting / re-assembling variable-size packets into / from fixed-length packets in a shared memory switching device

   ---

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2008/074828

Patent document cited in search report		Publication date	Patent family member(s)			Publication date
US 2004081184	A1	29-04-2004	NONE			
WO 9417617	A	04-08-1994	AU	693084	B2	25-06-1998
			AU	5982494	A	15-08-1994
			AU	6381798	A	18-06-1998
			BR	9406142	A	12-12-1995
			CA	2153172	A1	04-08-1994
			CN	1386021	A	18-12-2002
			CN	1097535	A	18-01-1995
			DE	69428107	D1	04-10-2001
			DE	69428107	T2	02-05-2002
			DK	681770	T3	10-12-2001
			EP	0681770	A1	15-11-1995
			ES	2163431	T3	01-02-2002
			FI	953594	A	27-07-1995
			JP	3476194	B2	10-12-2003
			JP	8505991	T	25-06-1996
			JP	3589660	B2	17-11-2004
			JP	2003224582	A	08-08-2003
			MX	9308193	A1	31-01-1995
			NO	952980	A	21-09-1995
			US	5361257	A	01-11-1994
			US	5467347	A	14-11-1995
US 6151301	A	21-11-2000	NONE			