

(12) **UK Patent Application** (19) **GB** (11) **2 368 228** (13) **A**

(43) Date of A Publication **24.04.2002**

(21) Application No **0025507.5**

(22) Date of Filing **18.10.2000**

(71) Applicant(s)
3Com Corporation
(Incorporated in USA - Delaware)
5400 Bayfront Plaza, Santa Clara,
California 95052-8145, United States of America

(72) Inventor(s)
Kevin Jennings
Jerome Nolan
Paul O'Keefe
David Nolan
Sorcha O'Callaghan

(74) Agent and/or Address for Service
Bowles Horton
Felden House, Dower Mews, High Street,
BERKHAMSTED, Herts, HP4 2BL, United Kingdom

(51) INT CL⁷
H04Q 11/04 , H04L 12/56 29/06 , H04Q 3/00

(52) UK CL (Edition T)
H4K KTKX

(56) Documents Cited
GB 2338147 A **GB 2317542 A**
GB 2316589 A **EP 0594196 A**
EP 0551243 A **WO 96/00945 A**
US 6041053 A **US 4587610 A**

(58) Field of Search
UK CL (Edition S) H4K KTKX
INT CL⁷ H04L 12/56 29/06 , H04Q 3/00 11/04
ONLINE : WPI ; EPODOC ; PAJ

(54) Abstract Title
Network unit with address cache

(57) A network unit includes a cache memory for storing address pairs of frequently-occurring packet conversations through the unit. On a look-up request, a cache controller checks whether an address pair is held in the cache before allowing a normal look-up in a layer database. The controller has recourse to RMON statistics to update the cache by inserting address pairs of more frequent conversations in place of those for less frequent conversations.

GB 2 368 228 A




FIG.1




FIG.2

	SA	SB	SC	SD	
DA	*	T1	T2	T3	
DB	*	*	T4	T5	
DC	*	*	*	T6	

21

FIG.3

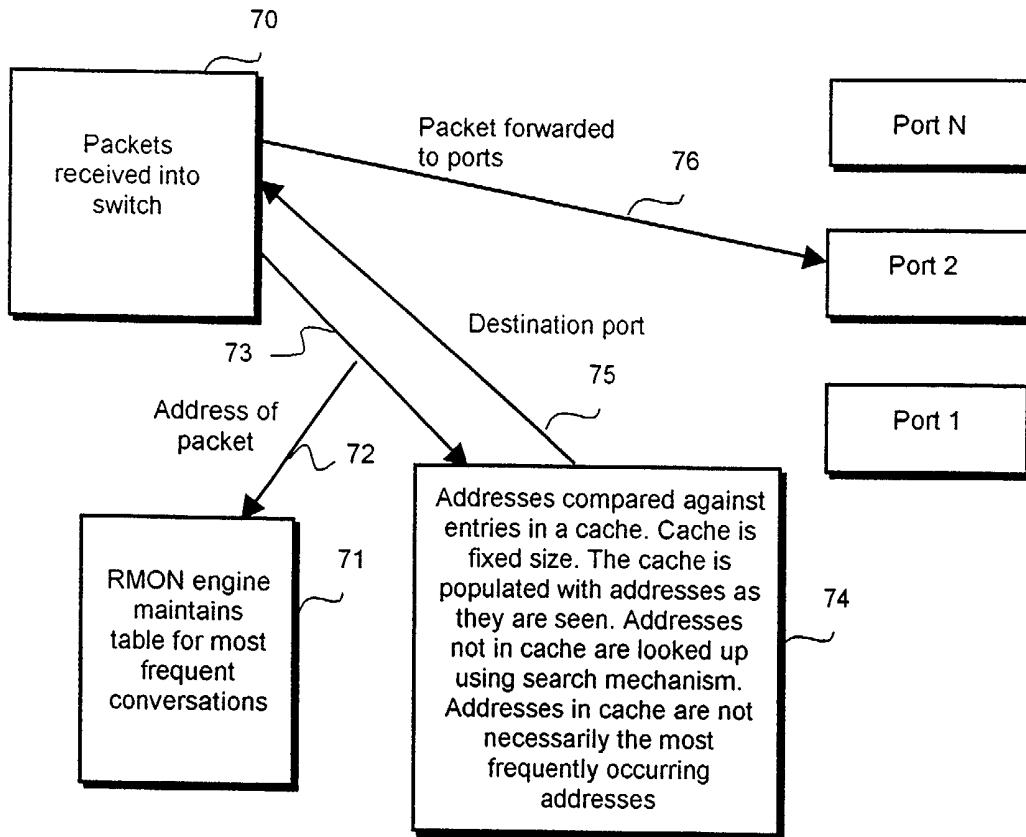
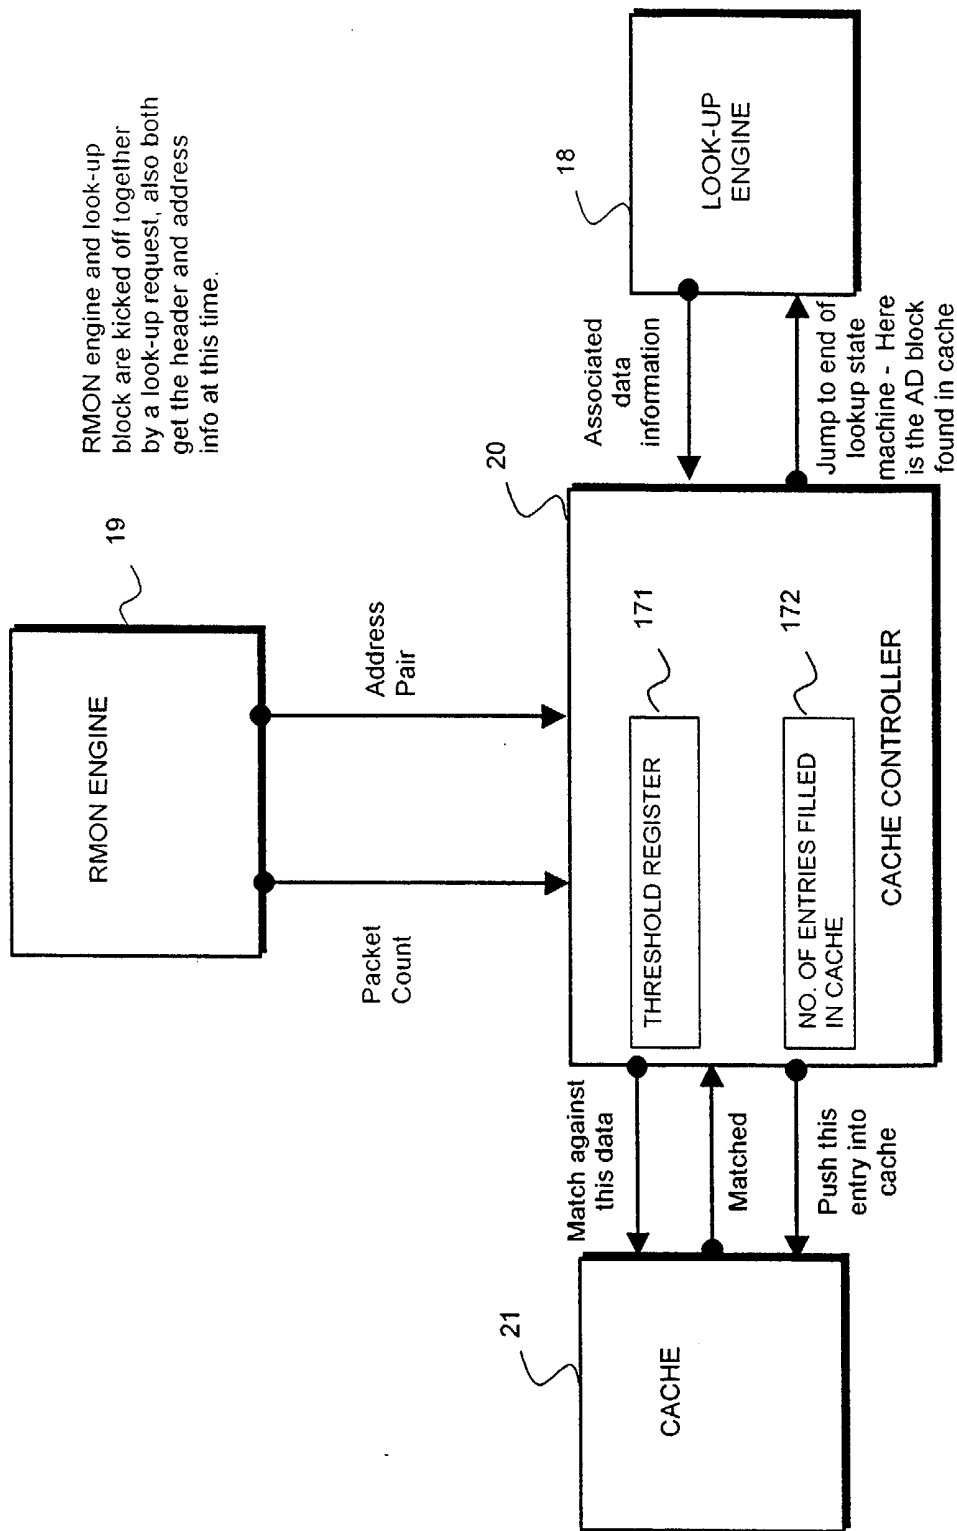


FIG.4



RMON engine and look-up block are kicked off together by a look-up request, also both get the header and address info at this time.

FIG.5

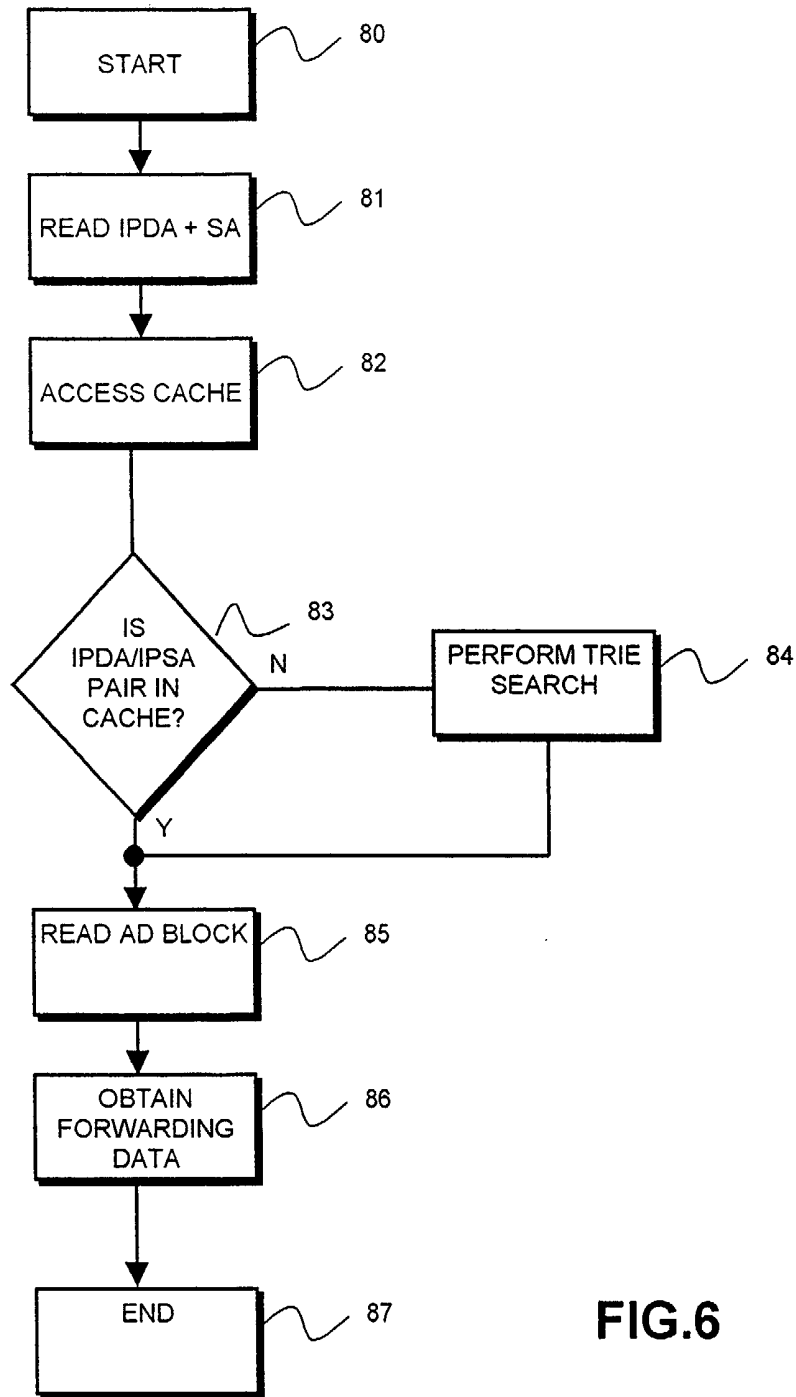


FIG.6

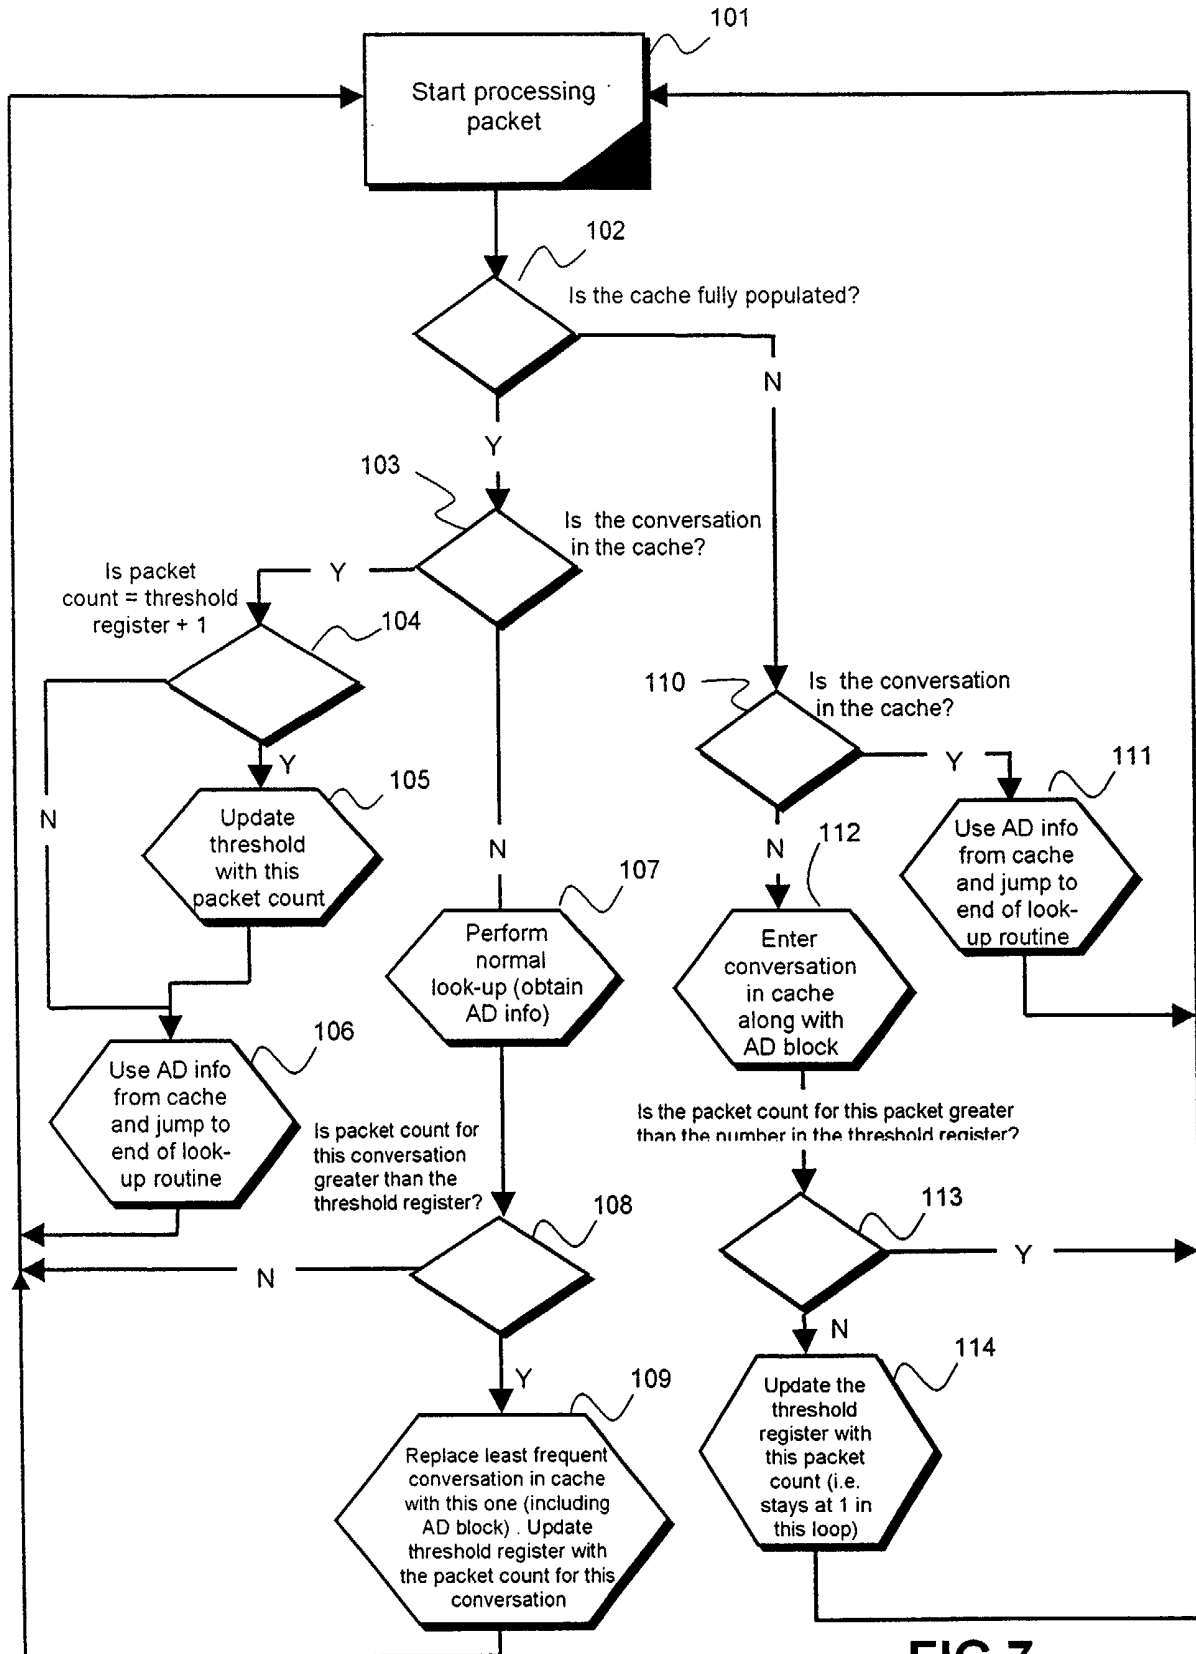


FIG.7

NETWORK UNIT WITH ADDRESS CACHE FOR FREQUENTLY OCCURRING NETWORK CONVERSATIONS

Field of the Invention

5

This invention relates to network units, that is to say multi-port units which can receive and transmit addressed data packets, such as, for example, Ethernet data packets and which need to perform address look-ups so as to relate address data, which maybe for example media access control address data or network address data or both in order to determine by means of a database relating address data to forwarding data (such as a numerical identification of a respective port or ports) which port or ports are required to send a data packet previously received by the unit. Examples of such network units are "switches", "bridges" and "routers". The invention more particularly relates to making a look-up process more efficient, at least in respect of addresses which appear in more frequently occurring conversations (i.e. message traffic between specific representative by sources and destinations).

15

Background to the Invention

20

25

30

Various architectures and associated software are available for network units which are required to perform address look-ups to determine the physical destination or destinations of a packet received by a network unit of the kind to which the invention relates. A common feature of such units is a look-up table or "forwarding database" which, despite having a variety of forms and different manners of organization dependant on circumstances, in essence relates address data in a packet to forwarding data. Forwarding databases are generally more complicated than the simple description indicates. For example, they may cope with different layer addresses, have a variety of means of restricting unnecessary look-ups, and employ "aging" to remove entries which have not been used for a particular time. Furthermore, there is a variety of ways of organizing the look-ups in such a database. Owing to the considerable lengths of addresses and the high cost both financially and in temporal terms, of database searching, various schemes have been employed to reduce the memory space required and to accelerate the search process. One method of reducing the space required is to employ "hashing" of address data. The search process may be organised on a

binary tree basis but other forms of searching, such as "trie" searching, are now commonplace.

5 More complex databases may require more than the destination address of a packet for a successful look-up. For example when the network unit is used in conjunction with a 'farm' of servers and all the servers in the farm have the same 'virtual' IP (internet protocol) address, the unit, in this example a router, may store traffic among the servers according to the IP source addresses of the traffic. A look-up table for this purpose can be based on a trie search using a key identifying both the network destination address (i.e. the servers' virtual IP
10 address) and the IP source address, and if appropriate the identification of the virtual LAN of the source.

Such a key may be used to search a trie data structure until a leaf node is found, i.e the look-up is performed by traversing the trie structure until an associated data block containing the
15 'forwarding data' is reached or the search terminates without success. Each stage in the trie uses a few more bits from the key. This means that many addresses with a common prefix can refer to the same target address.

By way of further example, load balancing is a method of sharing the load between a
20 multiplicity of servers for higher performance and for the protection of redundancy. All the servers in a server "farm" may have the same virtual internet protocol address. The router, such as the network unit 10 shown in Figure 1, shares traffic destined for this virtual internet protocol address between the servers based on some selective algorithm. One typical algorithm employs the internet protocol source address of a packet to determine which server
25 will be selected for handling such a packet.

Currently therefore in server load balancing (and also IP multicasting) the entire internet protocol destination address and internet protocol source address need to be stored in a trie look-up table. This means that the search of a very long key may be required. Very
30 typically, in a trie search the key is read in slices, and typically ten "slices" are required. This consumes substantial time and it is desirable to improve the performance of the switch by rendering more common reads are necessary.

It is known, for a variety of purposes, to provide a "cache" memory, for example constituted by a content addressable memory (CAM). Such a memory can be accessed by direct application of address data, for example as described in European patent application EP-
5 0594196-A1, wherein the content addressable memory is employed in conjunction with hashing of address data to provide a search key and particularly to deal with cases where
more than a particular number of addresses would hash to the same key.

Content addressable memories, for example as described in US patent 4587610, require only
10 one read operation to compare all stored addresses with an incoming address. However, this speed of operation is obtained at the cost of great complexity and cost. The present invention will employ a content addressable memory only for a small number of address pairs.

A known technique in association with network units is known as "RMON" (remote
15 monitoring). This in essence comprises remotely controlled "snooping" on packets which pass through the switch in order to obtain information from the header portions of the packets. Remote monitoring is employed for a variety of purposes, particularly the analysis of traffic flows in the system to which the network unit is connected. It is described in a variety of publications, such as for example in GB-2317542, GB-2316589, 2338147 and so on. A
20 known feature of remote monitoring is the identification of each "conversation" (identified by a pair of addresses) which is handled by the unit at any particular time.

Summary of the Invention

25 The present invention has the object of avoiding or at least reducing the number of look-ups in a forwarding database that may be necessary by maintaining in a cache entries identifying frequently used 'conversations' occurring in packets handled by the unit together with the relevant forwarding data, or at least a means of accessing the relevant forwarding data for such packets. The cache memory may be constituted by a content addressable memory which
30 can rapidly provide access to the relevant forwarding data.

In a preferred form of the invention a look-up process comprises accessing the cache to determine whether a conversation, defined by an address pair, is the subject of an entry in the cache and resorting to a larger database if the conversation is not the subject of an entry in the cache.

5

The cache is preferably established and maintained by recourse to the table or list of conversations, preferably established by RMON techniques. In particular, as part of the look-up process, if a measure of traffic flow (such as a packet count) relevant to a first conversation represented by a packet is greater than that for a second conversation which is the subject of an entry in the cache, the first conversation (if not in the cache) may be put in the cache in place of the second. The cache may be continually updated in this manner.

10

Further features of the invention will be apparent from a description of a more detailed example, with reference to the drawings.

15

Brief Description of the Drawings

Figure 1 is a schematic diagram of a network unit which would typically host the present invention.

20

Figure 2 illustrates a trie search.

Figure 3 illustrates a conversation table.

25

Figure 4 illustrates a search process according to the invention.

Figure 5 illustrates the interrelationship between the main functional parts participating in a search process according to the invention.

30

Figure 6 is a flow diagram illustrating in simplified form a search process according to the invention.

Figure 7 is a flow diagram illustrating the search process more particularly in relation to the access to and control of the cache memory.

Description of a Preferred Example

5

Figure 1 of the drawings illustrates in a schematic manner a network unit 10 in accordance with the invention. The unit is intended to represent one example of a wide variety of units which can host the present invention.

10

Merely by way of example, the unit shown in figure 1 has a multiplicity of ports, represented in the example by four ports 11, 12, 13 and 14. These ports are adapted for connection to an external network by any suitable medium, for example twisted pair, coaxial line, fibre optic line or wireless link according to circumstances. The ports 11 to 14 contain as is usual physical layer control and media access control, for example in accordance with the relevant transmission standard.

15

The unit includes a bus system which is shown as the bus 15 for the conveyance of control data, packet data and such like around the various parts of the unit. Although the ports 11 to 14 may have temporary buffers associated with them, in accordance with known practice, storage of packets that have been received and before they are transmitted is effected in a central packet memory 16. In order to determine where a given packet should go, a look-up engine 18 is provided. In essence this includes a database (e.g. as shown in Figure 2) which contains entries relating address data to forwarding data (a number or other data identifying a port) so that for any given destination (which may be single or multiple) the accessing of the look-up table from the relevant data will yield the forwarding data which is used by a forwarding engine, in this example represented by the CPU 17, which will cause extraction of the relevant packet from memory 16 and the forwarding of the packet to the relevant port or ports.

20

25

30

As thus so far described the unit is commonplace and various different architectures are known and available commercially. The remaining blocks in Figure 1 will be described later.

As is more particularly described later, for a variety of purposes, such as load balancing in a server farm, it is desirable or necessary to perform look-ups using both source and destination data of a packet, and for this purpose a trie search facility may be provided.

5 Figure 2 illustrates a known manner of organising a search table according to a trie look-up scheme. The look-up table is organised in blocks to two types. One type, exemplified by blocks 53, 54 and 55, contain pointers which are either of a trie type, such as pointer 56, that point to another block of trie type or, as exemplified by pointer 60, are of an 'associated data' type which point to the second type of block, the AD block exemplified by blocks 59 and 61.
10 The AD or associated data blocks contain the forwarding information for the relevant address key.

The first block, 53, in the trie search is accessed by the first set of bits from the trie key. These bits identify a pointer within the block 53. In the case of a pointer such as pointer 56,
15 which points to block 54, the pointer within block 54 will be accessed by the second set of bits from the trie key and so on until an associated data block, such as block 59, is reached.

It may be in some cases that a group of addresses share the same route; then the first part of the trie key will immediately yield, as in the case of pointer 60, a pointer to an associated data
20 block 61 and it is then unnecessary to traverse the whole key in order to obtain the forwarding information.

Figure 2 shows two keys 51 and 52 which are long keys composed of the internet protocol destination address and the internet protocol source address. They may also contain VLAN
25 information.

For further information on the organisation of trie searches, reference may be made to European patent application EP-0551243-A2, US patent 6041053 (Douceur et al) and international patent application publication number WO96/00945.

30

As indicated in the foregoing, currently an entire address pair, the internet protocol destination address and internet protocol source address needs to be stored in a trie look-up

table preferably with the number of the source VLAN. This means a search of the entire key may be needed. Typically, the key is read in slices of bits, for example an initial segment of eleven bits, subsequent segments of seven bits and a final segment of four bits, requiring up to ten reads from the trie search table.

5

A further known feature in switches of this kind is a remote monitoring facility, shown in Figure 1 as "RMON" engine 19. In this example the remote monitoring system is coupled to the internal bus 15, as described in British patent GB-2316589 and elsewhere.

10

Remote monitoring is known to provide a "conversation table" which is essentially a table indexed according to address pairs together with an indication of the number of times that address pair has occurred.

15

There is a variety of ways in which the remote monitoring can be performed. It may be done randomly, or at regular intervals. The measure of traffic flow for any given address pair maybe in terms of bytes or packets.

20

The RMON engine 19 monitors traffic that is bridged and/or routed by the unit. For each packet (or each of a succession of sample packets at regular or pseudorandom intervals) entering the switch it checks whether an entry for a conversation pair has been entered in a table in static random access memory. If it has not been so entered, the RMON engine creates a new entry. If an entry does exist it updates the statistics for the conversation.

25

The address pairs for the most active conversations as determined by the RMON engine are put into a cache memory, namely the CAM (content addressable memory) 21. Along with each entry will be a pointer to the associated data block for this conversation. The access to the CAM 21 and its updating is controlled by a cache controller 20, preferably constituted by a state machine which will be described with reference to Figure 7. For each packet entering the switch the relevant destination and source address data will be entered into the cache if this conversation has more frequently occurred than the least frequent in the cache and if it has not been entered before. Entries may also be aged out if they are inactive for a specified

30

time. Each time the associated data block is added or removed or a branch of layer 3 trie search is modified, the AD block pointer in the cache will be updated accordingly.

5 When a layer 3 packet is about to be looked up, its destination address and source address will be compared with the entries in the cache. If the entry is found the pointer will be used to go straight to the AD block. This saves many reads and enables routing of these packets at the wire rate.

10 A variation on this method is to have the associated data block located in the cache also. In this case each time the information and the associated data block is modified the corresponding entry in the cache needs to be updated.

15 Figure 3 illustrates a known form of conversation table 21 which can be constructed using a known RMON engine 19. The monitoring system is, for example, coupled to the main data bus of the switch so that it can snoop on packets (or a regular or random selection of packets) in order to obtain, at each sample, the IP destination and source addresses of the sampled packet. The conversation table also needs a measure of the traffic; this may be represented by a count of packets or may be represented by a byte count according to preference. The conversation table thus has entries schematically represented in Figure 3, each of which identifies an IP source and destination pair and contains an indication of the packet count flowing between that pair. The source addresses are SA, SB etc and the destination addresses are DA, DB etc. The values T1, T2 etc represent the volume of traffic containing the respective pair of addresses.

25 Figure 4 illustrates in a generally schematic form a search process according to the invention.

30 Block 70 represents the known process by which packets are received into the switch, temporarily stored in packet memory and so on. Block 71 represents the operation of the RMON engine, which maintains a table for the most frequent conversations, as obtained by examination of the packet headers as more fully described later. Line 72 represents the updating of the RMON engine as each new packet is received by the switch. Arrow 73

represents a look-up request performed in respect of a packet while it is stored in the switch in order to obtain the forwarding data for the particular packet.

Block 74 represents the search process performed according to the invention. The pair of
5 addresses (IPSA and IPDA) are compared with the entries in the cache, which is normally a fixed size. The cache is populated with addresses as they are seen (Figure 7). If the address pair is not in the cache, the 'ordinary' look-up process will be initiated, as denoted by arrow 75. That process will normally yield the forwarding data, so that the packet is forwarded (arrow 76) to one or other of the ports.

10 Owing to the nature of the RMON engine, the entries in the cache are not necessarily the most frequently occurring conversations at any given time, though statistically they will constitute most of the frequent conversations.

15 The search process is summarised in Figure 6. From a start 80 the addresses obtained from a packet, stage 81, are read. The cache is accessed, stage 22. The cache will be populated with addresses as they are seen. If the address pair, the IP destination address and IP source address are in the cache (stage 83), then the associated data block (stored in the cache along with the address pair) is read (stage 85), the forwarding data obtained (stage 86) and the
20 process terminated. If the address pair is not available from the cache, a look-up will be performed by means of the trie search facility (stage 84) as previously described with reference to Figure 2.

25 Figure 5 shows the interrelationship between the principal functional parts of the search process, involving the cache 21, the RMON engine 19, the cache controller 20 and the look-up engine. The RMON engine provides packet counts and address pairs. The cache controller 20 maintains a threshold and the number of entries filled in the cache, in registers 171 and 172. The cache controller has recourse to the cache to determine whether the address pair is matched against data held in the cache and the cache will signal the cache controller to signify
30 a match. The cache controller will also push an entry on to the cache under certain conditions to be described later.

Figure 7 illustrates more specifically the operation of the cache controller. The start of the process shown by Figure 7 is the commencement 101 of the processing of the packet. At stage 102 the controller determines whether the cache is fully populated. Initially it will be assumed that the cache is fully populated so that the next stage 103 is a determination whether the conversation (the particular IP source and destination address pair) is in the cache. If the conversation is the cache, there is a determination 104 whether the packet count is equal to a value held in threshold register 171 incremented by unity. If it is, then the threshold register is updated (105) with the new packet count.

10 Then the AD information from the cache entry (or elsewhere) is obtained (106) and a jump is made to the end of the look-up routine.

If the conversation is not in the cache several actions will be required. One (107) is to perform the ordinary look-up (e.g. as shown in Figure 2). At stage 108 the controller determines whether the packet count for the conversation is greater than that in the threshold in register 171. If the packet count for the conversation is not greater than the threshold register there will be no need to enter this conversation in the cache and the process will terminate in readiness for the processing of the next packet. If however the packet count for the conversation is greater than the count in the threshold register then the least frequent conversation in the cache will be replaced (stage 109) by the conversation represented by the new packet, including the AD block when it is found. The threshold register will be updated with the packet count for this conversation.

If however the cache is not fully populated then the new conversation may be entered into the cache. A determination (110) will be made whether the conversation is in the cache already. If it is, then the AD information from the cache is employed (stage 111) and there is a jump to the end of the look-up routine. If the conversation is not in the cache, and the cache is not fully populated, the conversation must be entered in the cache (stage 112). This will be accompanied by the AD block. There needs also to be a determination (stage 113) whether the packet count for the packet is greater than the number in the threshold register. If it is not, then the threshold register will be updated (stage 114) with the packet count.

Claims

5 1. A network unit which includes:

a look-up engine for performing an address look-up in response to a key including a network address pair in a packet to obtain forwarding data for said packet;

10 means for monitoring conversations defined by network address pairs in packets received by the unit to obtain measures of traffic flow for said conversations;

a cache memory for storing entries accessible by network address pairs and enabling forwarding data to be obtained for entries in the cache; and

15

a cache controller operative:

(i) to cause in response to a look-up request a determination whether the address pair in a packet is held in the cache ;

20

(ii) to allow the look-up engine to perform the address look-up when the address pair in said packet is not held in the cache; and

25

(iii) to update said cache so as to displace entries associated with relatively low measures of traffic flow by entries associated with relatively high measures of traffic flow.

30

2. A network unit according to claim 1 wherein the cache controller is operative, when address pair of the packet is not held in the cache, to compare the measure of traffic flow with a threshold and to insert the address pair as a new entry in the cache if the measure exceeds the threshold and is operative when the address pair of the packet is held in the cache to determine whether the said measure exceeds said threshold and thereupon to increase said threshold.

3. A network unit according to claim 2 wherein the cache controller is operative to remove the entry associated with the least measure of traffic flow when inserting said new entry.
- 5 4. A network unit according to any foregoing claim wherein said look-up engine is organised to perform a trie search.



INVESTOR IN PEOPLE

Application No: GB 0025507.5
Claims searched: 1-4

Examiner: Richard Howe
Date of search: 22 June 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK Cl (Ed.S): H4K (KTKX)
Int Cl (Ed.7): H04L (12/56, 29/06) ; H04Q (3/00, 11/04)
Other: Online : wpi ; epodoc ; paj

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	GB 2 338 147 A (3Com) - see abstract	
A	GB 2 317 542 A (3Com) - see abstract	
A	GB 2 316 589 A (3Com) - see abstract	
A	EP 0 594 196 A (Digital Equipment Corporation) - see abstract	
A	EP 0 551 243 A (Digital Equipment Corporation) - see abstract	
A	WO 96/00945 A (International Business Machines) - see abstract	
A	US 6 041 053 (Douceur) - see abstract	
A	US 4 587 610 (Prime Computer) - see abstract	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.