



(86) Date de dépôt PCT/PCT Filing Date: 2004/09/10
(87) Date publication PCT/PCT Publication Date: 2005/03/24
(85) Entrée phase nationale/National Entry: 2006/03/02
(86) N° demande PCT/PCT Application No.: GB 2004/003869
(87) N° publication PCT/PCT Publication No.: 2005/027462
(30) Priorité/Priority: 2003/09/11 (EP03255686.2)

(51) Cl.Int./Int.Cl. *H04L 29/06* (2006.01),
H04L 12/56 (2006.01)
(71) Demandeur/Applicant:
DETICA LIMITED, GB
(72) Inventeurs/Inventors:
BENNETT, MARK ARWYN, GB;
PIGGOTT, ALEXANDER COLIN, GB;
GARFIELD, DAVID JOHN MICHAEL, GB;
MORRIS, PHILIP, GB
(74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : CONTROLE ET SECURITE DE RESEAU EN TEMPS REEL
(54) Title: REAL-TIME NETWORK MONITORING AND SECURITY

(57) **Abrégé/Abstract:**

There is provided a hardware device for monitoring and intercepting data packetized data traffic at full line rate. In preferred high bandwidth embodiments, full line rate corresponds to rates that exceed 100 Mbytes/s and in some cases 1000 Mbytes/s. Monitoring and intercepting software, alone, is not able to operate on such volumes of data in real-time. A preferred embodiment comprises: a data delay buffer (208) with multiple delay outputs (216); a search engine logic (210) for implementing a set of basic search tools that operate in real-time on the data traffic; a programmable gate array (206); an interface (212) for passing data quickly to software sub systems; and control means for implementing software control of the operation of the search tools. The programmable gate array (206) inserts the data packets into the delay buffer (208), extracts them for searching at the delay outputs and formats and schedules the operation of the search engine logic (210). One preferred embodiment uses an IP co-processor as the search engine logic.



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

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
24 March 2005 (24.03.2005)

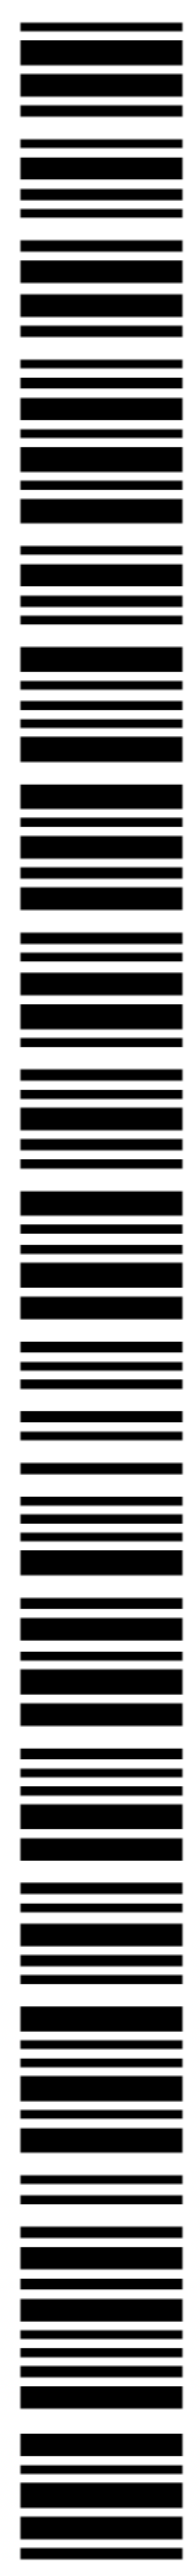
PCT

(10) International Publication Number
WO 2005/027462 A1

- (51) International Patent Classification⁷: **H04L 29/06**, (74) Agent: **GILL JENNINGS & EVERY**; Broadgate House, 7 Eldon Street, London EC2M 7LH (GB).
12/56
- (21) International Application Number: PCT/GB2004/003869 (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: 10 September 2004 (10.09.2004) (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, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 03255686.2 11 September 2003 (11.09.2003) EP
- (71) Applicant (*for all designated States except US*): **DETICA LIMITED** [GB/GB]; Surrey Research Park, Guilford, Surrey GU2 7YP (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (*for US only*): **BENNETT, Mark, Arwyn** [GB/GB]; 22 Grantley Close, Shalford, Surrey GU4 8DL (GB). **PIGGOTT, Alexander, Colin** [GB/GB]; 24 Ellingham, Midhope Road, Woking, Surrey GU22 7UQ (GB). **GARFIELD, David, John, Michael** [GB/GB]; Ground Floor Flat, 170 Hartfield Road, Wimbledon, London SW19 3TQ (GB). **MORRIS, Philip** [GB/GB]; Flat 3, Strathmore House, 22 Jenner Road, Guildford, Surrey GU1 3PP (GB).
- Published:**
— with international search report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: REAL-TIME NETWORK MONITORING AND SECURITY

(57) Abstract: There is provided a hardware device for monitoring and intercepting data packetized data traffic at full line rate. In preferred high bandwidth embodiments, full line rate corresponds to rates that exceed 100 Mbytes/s and in some cases 1000 Mbytes/s. Monitoring and intercepting software, alone, is not able to operate on such volumes of data in real-time. A preferred embodiment comprises: a data delay buffer (208) with multiple delay outputs (216); a search engine logic (210) for implementing a set of basic search tools that operate in real-time on the data traffic; a programmable gate array (206); an interface (212) for passing data quickly to software sub systems; and control means for implementing software control of the operation of the search tools. The programmable gate array (206) inserts the data packets into the delay buffer (208), extracts them for searching at the delay outputs and formats and schedules the operation of the search engine logic (210). One preferred embodiment uses an IP co-processor as the search engine logic.



WO 2005/027462 A1

REAL-TIME NETWORK MONITORING AND SECURITY

Field of the Invention

The present invention relates to the monitoring of network data traffic in real-time. In particular, the present invention relates to apparatus for the real-time monitoring and processing of internet protocol (IP) packets at very high data rates to enable a range of data monitoring and real-time network security applications.

Background to the Invention

The relentless increase in data traffic over distributed networks has engendered a concomitant requirement for monitoring and filtering the data passing over such networks. Whether the distributed network is a local area network (LAN), a metropolitan area network (MAN), or a wide area network (WAN) that spans an entire continent, the monitoring and filtering tasks are conventionally carried out either in hardware, using dedicated application specific integrated circuits (ASICs), or in software, running on processing devices hosting server software. In the latter case, the processing devices may in fact be network servers - computers one of whose programmed tasks is to operate server software within a particular LAN or WAN. Alternatively, they may be dedicated devices that run only network software, so-called "network appliances". It will be understood that the term "server" is used in the literature to denote both the physical processing devices and the server software they host. Where a distinction is not made between these meanings in the following text, the term refers to both simultaneously.

Typically, a range of monitoring and filtering tasks are performed at various points throughout distributed networks in order to protect servers and/or their clients from unwanted data. Indeed, the routing strategies adopted by routers may be considered as fulfilling a basic filtering task. There are many software applications, executed by servers themselves, that implement monitoring and/or filtering tasks. An important subset of these applications deals primarily with potential threats to the integrity of the network or the

information transmitted over it: they are generally referred to as security software applications or simply "security elements".

Examples of existing internet security software applications for execution on servers include: antivirus (AV) detection software; spam detection software; 5 content monitoring software; firewall software; traffic sniffing software, anti-fraud software; and, intrusion detection and prevention software.

Servers can either be owned and operated by business or residential users, or provided as "managed security services" by a service provider, for example a telecommunications operator, a communications service provider 10 (CSP) or an internet service provider (ISP). In the case of a managed security service, the service provider may incorporate one or more security elements in their service offering. This may be done by placing appropriately configured servers on the customer premises or by providing the security element in their own servers at a data centre. Service providers typically supply either an all- 15 purpose network server or a dedicated network appliance for each group of one or more customers; a separate piece of server software processing data traffic for each security element offered. In addition, the service provider will generally deploy additional servers within their own network to provide internal security and to prevent attacks on their own network infrastructure.

20 Software security systems that enable the monitoring and controlling of data traffic in packet-oriented networks are known. Examples of security systems used with the ubiquitous TCP/IP protocol suite include firewall applications, AV software, and intrusion detection software, such as so-called "sniffer" software.

25 Firewall applications include a packet inspection functionality which allows content-aware filtering of the packets passing into a sub-network, for example allowing only certain FTP (file transfer protocol) commands to be used or blocking encrypted data traffic over a given port.

30 An ISP may provide AV software configured to be applied to email traffic in their mail service offering. AV software monitors the email traffic for data patterns that correspond to a regularly updated list of "virus signatures". Virus

signatures are patterns of data that identify potentially malicious executable code.

Sniffer software can be used to monitor and log the activity of a predetermined network user, thereby allowing a network administrator to detect
5 abusive or suspicious network activity.

Such software solutions are fundamentally limited by the capacity of servers to receive, process in software, and retransmit monitored data onto a network quickly enough, while avoiding the introduction of large delays (latency) or missing data packets entirely. Where large numbers of subscribers are
10 offered managed security services (as they would be in, for example, a residential service or a service aimed at small/medium sized enterprises) or where high traffic loads (input bit rates of 100Mbytes/s or more) are carried on large networks, the software approach runs up against these constraints. Conventionally, this problem can be addressed by deploying additional servers
15 in conjunction with (hardware) network switches that distribute the traffic load between available servers. This solution comes at cost in terms of: complexity; administration and management overheads; physical infrastructure space required for the network and attached devices; and, often crucially, in terms of financial outlay.

20 Conventional processing and monitoring devices, such as those used for processing IP data traffic, include a hardware network switch and a processing device. In operation, the processing device executes software servers; each software server offering the functionality of a security element to one or more users.

25 Real-time interception of network traffic is distributed amongst the software servers so that each software server hosts a security element functionality on behalf of a predetermined group of users.

Prior art systems do not cope effectively with high bandwidth data traffic, as might be found in the core networks of ISPs and CSPs. Examples of
30 communications standards exhibiting input rates that are considered to be high bandwidth include: STM-64 (1244 Mbytes/s); STM-16 (311 Mbytes/s); 10

Gigabit Ethernet (1250 Mbytes/s LAN mode, else as STM-64); and Gigabit Ethernet (125 Mbytes/s).

A hardware platform is able to handle these high bandwidth input rates. Conventional hardware platforms, however, replace each required security
5 element with corresponding, dedicated (often custom built) ASIC components.

Such devices have typically been built to perform specific protocol processing tasks. The major disadvantage of known custom built hardware devices is their static functionality. This arrangement, while capable of handling the data traffic at full line rate, is not easily adaptable or upgradable. This
10 means that they are complex, risky and time-consuming to build, update and maintain. Furthermore, the inherent difficulty of constructing firmware or hardware analogous to "real-time" security software applications means that there are relatively few appropriate developers.

A more recent approach has been to provide an ASIC that includes a
15 field programmable gate array (FPGA) with dynamically programmable logic. To achieve the dynamic programming either a language such as VHDL is used or a programmable finite state machine (FSM) is instigated.

Where VHDL and similar languages are used the drawback has been that the developer community is small, and that programming and particularly
20 debugging applications has been found very difficult. By contrast, the problem with FSMs is that, while they might be easy to program, the functionality it is possible to implement using them is very restricted compared to the versatility of programming languages: attempts to extend the functionality have resulted in unmaintainable code.

25

Statement of Invention

According to one aspect of the present invention there is provided an apparatus for detecting a predetermined bit pattern in data traffic comprising a plurality of data packets, the data packets having a header portion and a
30 payload portion, the apparatus comprising:

a processing logic for receiving the data traffic in real-time and for dividing the data packets in the data traffic into bit sequences; and

an Internet Protocol (IP) co-processor unit having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the IP co-processor unit comprising means for identifying, from within at least the payload portion, each data packet that contains a bit sequence that matches a predetermined bit pattern entry within the memory array.

Rather than use dedicated ASICs, the invention makes use of more general purpose hardware components for performing searches of data traffic for occurrences of specific bit patterns: and in particular, an IP co-processor. In this aspect of the invention, one or more IP co-processors are used to provide content addressable memory (CAM) arrays that are suitably fast at performing basic comparison tasks. Not only are IP co-processors capable of high comparison rates, they are also arranged so that the "selectors", i.e. the terms for which the device searches, can be altered, added, and deleted under software control.

The bit pattern storage memory array may store bit patterns as ternary data. As a result, the IP co-processor can be arranged to match bit patterns of variable lengths. The matching of alphabetic characters can be made 'case insensitive': for instance, ASCII differentiates between cases by changing one bit of the ASCII character code and when searching for a bit pattern in any fixed length ASCII character the 'case' bit can be represented by an 'X' for "don't care". Furthermore, by using the 'X' state to represent a 'wildcard', many bit patterns can be searched in one comparison: so-called "selector compression".

In certain circumstances, the bit sequences supplied to the IP co-processor unit may be consecutive sequences of bits from a single data packet. To match patterns across packet boundaries, at least one of the predetermined bit patterns would preferably comprise two subpatterns, the at least one bit pattern having a length L bits, each subpattern being stored as a separate bit pattern entry in the memory array; the identifying means would compare the subpatterns with bit sequences from an end portion of each data packet, the portion being within L bits of the beginning or of the end of the data packet; and the identifying means would identify each data packet that contains a bit sequence that matches one or other of the two subpatterns.

The apparatus preferably further comprises a control interface, the control interface being in communication with the processing logic, the interface conveying software instructions to the processing logic to alter the bit pattern entries in the IP co-processor unit. The alteration may be an addition or a modification of a given entry (selector) in the bit pattern storage memory array and the IP co-processor unit is thereby configured to search for a new bit pattern. Alternatively, the alteration is a deletion of one or more pattern storage array entries, with the result that the search engine logic no longer searches for the or each deleted bit pattern.

Advantageously, the control interface is integral with the processing logic. Through this control interface, software applications can adapt the bit pattern search functionality of the IP co-processor unit.

The bit pattern storage memory array of the IP co-processor unit may include a plurality of bit offset entries corresponding to one or more predetermined bit patterns. Provided there are sufficient unused locations in the bit pattern storage memory array, multiple bit-offset entries will permit a higher rate of data packet throughput in the IP co-processor unit. A special case of this would be an eight-bit offset (or byte offset).

Another way to extend the number of selectors is to provide additional IP co-processor units, each with their respective bit pattern storage memory arrays.

The apparatus preferably further comprises at least one further IP co-processor unit, the or each further IP co-processor unit operating upon bit sequences having a predetermined offset relative to the bit sequences searched by the IP co-processor unit. Successively offset versions of the same data packet are thereby compared against each selector in the selector table stored on each bit pattern storage memory array.

Advantageously, the or each further IP co-processor unit is provided with an identical copy of the contents of the bit pattern storage memory array of the IP co-processor unit.

It is preferred that the apparatus further comprises a delay buffer for storing the data traffic at full line rate, wherein data streams that contain a data

packet that is identified as having a bit sequence that matches the predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for further processing.

In preferred high bandwidth embodiments, full line rate corresponds to bit rates that exceed 100 Mbytes/s (in fact, often by a factor of 10 or more).

In accordance with a further aspect of the invention, there is provided an apparatus for detecting a predetermined bit pattern in data traffic comprising a plurality of data streams of data packets, the apparatus comprising:

a processing logic for receiving the data traffic in real-time and for dividing the data traffic into bit sequences;

a search engine logic having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the search engine logic comprising means for identifying each data packet that contains a bit sequence that matches a predetermined bit pattern entry within the memory array; and

a delay buffer for storing the data traffic at full line rate, the delay buffer having one or more outputs,

wherein data streams that contain a data packet that is identified as having a bit sequence that matches a predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for further processing.

Again, the invention makes use of more general purpose hardware components for performing searches of data traffic for occurrences of specific bit patterns. In this aspect of the invention, a delay buffer creates a delay path that holds all data traffic for a length of time determined by the operation of the, or each, delay output and by the total capacity of the buffer.

In this case, it is advantageous that the memory capacity of the delay buffer is sufficient to retain the full bandwidth of the data traffic for a period of seconds. In preferred high bandwidth embodiments, full line rate corresponds to bit rates that exceed 100 Mbytes/s (in fact, often by a factor of 10 or more).

The provision of such a significant amount of memory for use as a buffer adds desirable functionality, including allowing further searches and combination searches of data traffic. The presence of several delay outputs of

different durations enables software modules to trade-off latency vs prior amount of data available for processing.

An actionable data packet may occur at any point in a data stream. The delay buffer permits the storage of a significant proportion of the packets in an data stream so that if a later data packet is found to be actionable, data packets belonging to the same data stream can be identified by further searches of the data exiting at one or more of the delay outputs. Consequently, once a data packet belonging to a data stream is identified as possibly actionable (by virtue of a match of a bit sequence in the data packet with the predetermined bit pattern), any data packets belonging to the same data stream, even if they were transmitted before the actionable data packet, can be identified (by virtue of a subsequent match of a further bit sequence, for instance an address in the header field) and forwarded, along with the actionable packet and any further packets belonging to the same data stream, to software for further processing.

The search engine logic advantageously comprises means for performing a protocol based search, in which target data packets are identified by virtue of the bit sequence of at least one data field within each data packet. The data field may be a header field, the header field including data relating to one of the group including: address information; to/from port number information; and, packet type identifier information.

The protocol based search is advantageously a search on IP header fields. Data packets having IP address information that matches the predetermined bit pattern can therefore be flagged and software security applications can then act upon sets of data packets representing whole email messages or data files containing computer viruses.

Search engine logic can make use of hash-table technologies but it is preferred that the device use CAM instead. It is consequently preferred that the search engine logic is an IP co-processor unit: IP co-processor units incorporate CAM.

In accordance with yet another aspect of the invention, there is provided a system for detecting a predetermined bit pattern in input data traffic

comprising a plurality of data packets, the system comprising a hardware component and a software component, the hardware component including:

a processing logic for receiving the data traffic in real-time and for dividing the data traffic into bit sequences; and

5 a search engine logic having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the search engine logic comprising means for identifying each data packet that contains a bit sequence that matches the predetermined bit pattern entry within the memory array,

10 wherein data packets processed by the search engine logic of the hardware component and identified as containing a matching bit sequence are passed to an executable software application of the software component for further processing, the bandwidth of data traffic passed on to the executable software application of the software component being substantially less than the bandwidth of the data traffic input into the hardware component.

15 Hardware components are thus used to narrow down the bandwidth of data traffic so that software components need only operate on a restricted subset of the data packets in the overall data traffic. Consequently, the search functionality is split into repetitive tasks that can be defined in simple terms (and performed fast by hardware components) and more computationally intensive
20 tasks that involve more complex definition (and therefore more conveniently executed by software at a slower rate).

The hardware component is capable of performing processing and data packet filtering at input rates too high to be achieved purely in software. It is also adaptable to changes in operational environment and relatively
25 straightforward to upgrade. In particular, data packet selection may be dynamically programmed from software to implement an arbitrary number of filters, including, for example, combinations of TCP/IP payload bit pattern searches and TCP/IP header look-ups.

30 The comparison performed in the hardware search engine logic replaces similar comparison functionality previously performed in software. Software security applications are thereby freed to act upon matches found by the hardware search engine. The present invention therefore allows a wide range

of security elements to be provided on a single network device, even when the number of subscribers is large or the traffic loads are very high. By monitoring and intercepting the data traffic in real-time using hardware, it substantially obviates the scalability and latency problems of the conventional software-only, server-based approach.

By permitting the real-time processing and interception of data traffic, the system allows data traffic to be processed to remove unwanted, dangerous or illicit content before this content reaches any customer's premises.

Preferably, the search engine logic of the hardware component is an IP co-processor unit.

It is preferred that the hardware component of the system further comprises a delay buffer for storing the data traffic at full line rate, wherein data streams that contain a data packet that is identified as having a bit sequence that matches the predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for further processing.

The processing logic is preferably programmable in a higher level programming language.

Higher level programming languages are generally easier to programme than languages like ASSEMBLY, FSM or firmware. They are also easier to maintain.

As the hardware component can be arranged to output data in a standardised, application independent format (for example using a small suite of API calls), high speed monitoring applications (e.g. anti-virus applications) can be developed without requiring extra code to be written to handle the output. In combination with the hardware component, these monitoring applications can operate on data at bandwidths too high for software applications on their own to handle. Furthermore, as the software developer will appreciate, by providing a suitable 'front end' interface, the system can be arranged to use off-the-shelf monitoring applications.

The hardware component can thereby detect the presence of a signature bit pattern, such as a virus signature, wherever in the data packet that signature might appear, generate application independent output, and pass that output

to the software component. The software component can then retask the hardware component to search for, and optionally intercept, related data packets.

The apparatus can therefore implement bit pattern detection functionality
5 for the full range of security elements found in software-based systems. However, since it operates in real-time, analysing data packets as they travel over the network, the system can provide additional security elements traditionally not found in server-based systems.

The combination of a free form search with a subsequent header field
10 search is particularly advantageous. In the free form search, the search engine detects instances of data packets that contain a match for a first bit pattern. These instances are analysed to discover information that identifies the network transaction to which the instances belong. The identification information is then used to generate a second bit pattern that corresponds to the contents of the
15 header field of any data packets belonging to the identified network transaction. In the subsequent header field search, the search engine is then instructed to search through data output from the delay buffer for data packets containing the second bit pattern. All packets containing the second bit pattern are then fed to software applications for further processing.

20 In this combination of searches, the search engine logic of the hardware component, performing a free form search for the first bit pattern, operates as a "coarse filter" such that only a small fraction of network transactions (sessions comprising one or more data packets related by their content or the fields in their protocol headers) that may contain actionable content is passed to the
25 software component.

The combination of signature and header will usually define both the threat and the service/protocol. Therefore, many software modules covering a wide range of threats and protocols may run concurrently without impacting the overall system performance (because only a small number, eg 1, will be
30 operating on a given filtered data packet at any one time).

Brief Description of the Drawings

Examples of the present invention will now be described in detail with reference to the accompanying drawings, in which:

5 Figure 1A shows an active configuration of a device in accordance with the invention;

Figure 1B shows a passive configuration of a device in accordance with the invention;

Figure 2 shows a more detailed schematic diagram of the components in an apparatus according to the present invention;

10 Figure 3 shows the header fields of a TCP/IP data packet;

Figure 4 illustrates the lookup operation implemented upon a search engine logic;

Figure 5 illustrates a CAM array containing four versions of the same selector with successive byte offsets;

15 Figure 6 illustrates four different CAM arrays (belonging to four different search engine logics) each containing one version of the same selector;

Figure 7 illustrates an accelerated lookup operation implemented upon either a single search engine logic using a CAM array as shown in Figure 7 or a set of search engine logics each containing one version of the same selector with successive byte offsets;

20 Figure 8 is a flowchart illustrating how the apparatus operates to detect virus signatures; and

Figure 9 shows an example of a system architecture including two interception cards.

25

Detailed Description

The invention can be implemented in a network in two configurations: an active configuration and a passive configuration. The active configuration (shown in Figure 1A) not only monitors and filters data but performs all the tasks of a conventional router - storing and forwarding data packets. The passive configuration (shown in Figure 1B) performs no routing functions, but may be

30

arranged to forward a copy of any suspect packages to a specific node on the network.

The passive configuration can be characterised as providing data packet sniffing applications, where wanted 'target' data packets are output from software. The active configuration allows 'safe' data packets to be streamed to
5 their destination, and is suited to streaming applications, such as real-time AV.

In an embodiment of the present invention, the apparatus is implemented as hardware logic units in a network device that also includes switching apparatus, for conventional network switching tasks. The hardware logic units
10 are arranged to provide output that is substantially independent of the software applications that use that output. Consequently, the hardware platform can provide data for a range of different applications without need for format conversion.

The hardware logic units thus perform, in real-time and at full line rate, the monitoring and processing functions conventionally carried out by software
15 servers. More precisely, the real-time approach involves providing, in a dedicated hardware search engine logic, the bit pattern detection functionality that has conventionally formed part of security software applications implemented in server software (e.g. looking for virus signatures). The
20 hardware search engine logic is provided as a component of the hardware logic units. The real-time checking path of the apparatus is implemented by the passage of the incoming data packets through the search engine logic.

A suitable arrangement of hardware logic units in accordance with this embodiment of the invention is illustrated schematically in Figure 2. Data enters
25 the device at full line rate via a standard line interface unit 202 and a framer/deframer 204. The hardware logic units comprise: a search engine logic 210 for implementing a set of basic search tools that operate in real-time on data packetized data traffic; an interface component 212 for passing data quickly to software sub-systems; and a processing logic 206 for implementing
30 software control of the operation of the search tools.

The search engine logic 210 can be implemented: as a hash-table search; using a hardware correlation engine; or by using a CAM (content addressable memory).

5 In a preferred implementation of the invention, an IP co-processor is used as the search engine logic. To understand why IP co-processors are suited to this task, one must consider the task for which IP co-processors are intended: routing in data packet-oriented networks.

10 An intrinsic component of any data packet-oriented network is the switch or router. This ensures that data packets are correctly sent to their destinations. Data packets within data packet-oriented networks generally have both header and payload portions. The header portion includes fields for storing: package number, destination address information, port numbers and protocol fields. The payload portion contains a block of data for transmission across the network.

15 Routers store each data packet temporarily (in a queue), inspect the destination address field, then forward the data packet to the destination (or another router) according to a routing strategy. No network can carry data traffic at a higher rate than the rate at which the router can repeat this process. In order to maximise data throughput, the hardware components of routers are continually improved: to increase the volume of data stored; to accelerate the
20 address inspection and forwarding steps; and to streamline the routing strategy. A particularly important component in this respect is the IP co-processor, which is specifically designed for use in routers for WANs that implement the TCP/IP protocol suite.

25 The fact that a data packet-switched network can not function at line rates that exceed the capabilities of the routers in that network means that there will always be a router component capable of storing, processing and forwarding data packets at the maximum line rate. IP co-processor chipsets currently fulfil this criterion. Furthermore, suitable IP co-processor chipsets are readily available as off-the-shelf items.

30 IP co-processor chipsets, such as the IDT [RTM] Network Search Engine devices, integrate suitably large arrays of CAM. These chipsets can be instructed to compare a data structure (temporarily stored in a register) with one

or more entries in the CAM array ("lookup"). In addition, new entries can be written into an unused location in the array ("learn"), thereby facilitating searches for the new entries.

In this preferred implementation, the IP co-processor provides the search functionality necessary to facilitate a range of security elements found in software-based systems, for example, firewall, intrusion detection & prevention, anti-virus for Internet email and FTP file downloading, spam filtering, and URL filtering. However, since it operates in real-time, analysing data packets as they travel over the network, apparatus incorporating IP co-processors can provide additional security elements traditionally not found in server-based systems: -

Anti-virus for all email not just email downloaded from the Internet service provider e.g. web mail, mail services offer by other ISPs etc.

Anti-virus protection for real-time network services such as HTTP, instant messaging, voice-over-IP (VoIP).

Chat-room watch guard as a child protection feature: watching the information typed into a chat room service in real-time and preventing personal data (real name, phone number, home address etc.) being sent or logging that personal data has been entered and alerting the child's parents or guardians.

As a parental control feature removal illicit material e.g. blanking out swear words in a web page or email.

The apparatus can also be used to monitor the activities of employees or subscribers on the network. In the case of employees, this might involve identifying which web sites they are visiting and for how long; the content of the emails, downloads, instant messaging, chat rooms and peer-to-peer applications they are using and whether it involves exchanging or downloading illicit or copyright materials; and the details of the other individuals with whom they are communicating. Similarly, for subscribers to an Internet service, this may involve monitoring their activities to make sure they are not violating the terms of their service agreement. The system, if it has identified a potential target, could be used to collect all data packets related to the target as evidence for employee dismissal case or service termination.

As noted previously, IP co-processors are constantly being developed to increase their capacity and IP co-processor capacity fundamentally limits network line rates. IP co-processors are attractive because, whatever the network, there is bound to be an IP co-processor designed to operate at the full line for that network. In other words, IP co-processors are good candidates for implementing a CAM that operates at full line rate because they are available at all achievable line rates.

While initial checking is performed at full line rate and substantially in real-time, subsequent software processing of intercepted content (in the separate sub-system) may require more time. By ensuring that all data traffic remains accessible for as long as it is needed, processing can be carried out at a rate that the software can handle efficiently, even if that rate is no longer in real-time.

Thus, in another preferred implementation of the invention, the accessibility of data traffic is assured by providing a delay path through the hardware logic units, i.e. a temporary memory, in parallel to the (full line rate) checking path: the duration of storage being sufficiently long to allow the software processing to complete. Meanwhile, data that has not been intercepted can be re-transmitted, from the checking path, with no significant latency. The invention therefore allows data traffic to be processed to remove unwanted, dangerous or illicit content before this content reaches any customer's premises.

In Figure 2, the hardware logic units also include a data delay buffer 208 with multiple tap points 216. The data delay buffer 208 can be implemented as volatile memory, for example DDR RAM or SRAM. The delay buffer operates as first-in-first-out (FIFO) memory so that for an input data rate of 0.3 Gbytes/s and a requirement of 1.5 seconds of delay, it will therefore suffice to provide 0.5 Gbytes of DDR RAM.

The processing logic 206 may be implemented as one or more network processing units (NPU), ASIC units, field programmable gate arrays (FPGAs) or a combination of such components. The processing logic 206 is arranged to manage the interaction between the other hardware components, between

hardware components and software control instructions (thereby mediating control of the operation of the search engine logic 210), and, via the interface component, between the apparatus and connected software sub-systems.

In the light of the discussion of the previous implementation, it is also preferred that IP co-processor chipsets are used to implement CAMs for the present implementation.

The search engine logic 210 and the processing logic 206 cooperate to implement basic search tools of two types:

a) a bit sequence search, which allows bit patterns to be matched anywhere within the data flowing through the apparatus, even across data packet boundaries; and

b) a header field search, which allows data packets to be identified based on the contents of data packet header fields, for instance: IP addresses; to/from port numbers; and protocol fields.

The bit sequence searches are also referred to as 'freeform' (FF) searches. FF searches are significantly more demanding in terms of utilisation of search engine logic, the key requirement in executing FF searches being to perform a bit-wise comparison of the incoming data against a dictionary of wanted bit patterns. IP co-processor chipsets are thus particularly suitable. As a rule, the longer the bit pattern and the larger the dictionary, the better: again, IP co-processor chipsets are seen to provide the flexibility to search on longer bit patterns (for example, allowing searches for 64*1024 (64K) pattern entries of up to 18 bytes in length).

With these search tools, the delay buffer 208 and software interaction, sophisticated protocol (data packet content) processing can be achieved.

In one specific embodiment of the invention, the processing and monitoring apparatus is adapted to operate on IP data traffic. This apparatus is now described in order to illustrate further features of the invention particular to the context of IP network communications. The apparatus of the invention is conveniently manufactured from a combination of custom built hardware and standard computing platforms.

The IP data traffic embodiment follows the architecture illustrated in Figure 2 and incorporates a combination of hardware components suitable for handling IP data traffic, including the following elements:

5 a) a standard line interface unit (optical transceiver) 202 and framer/deframer 204 for performing basic input and output (I/O). These are OTS chips configured for router applications, and suitable for handling large bandwidths at full line rate: in one possible implementation, the line interface unit 202 is suitable for SONET/SDH operation. Raw IP data packets (e.g. with the POS - Data packet Over SONET overhead removed) are output by the
10 framer/deframer 204.

b) a plurality of field programmable gate arrays (FPGAs) - the favoured implementation of the processing logic 206 in Figure 2. The hardware IP processing is handled by two FPGA modules: a memory management FPGA module for inserting the data packets from the framer/deframer 204 into the
15 main delay buffer 208, and extracting them for searching at one or more tap points 216 along the delay path; and, a data packet processing FPGA module, responsible for formatting and scheduling the searching. The data packet processing FPGA module is associated with a so-called "scratch pad" 218, for example a volatile memory device such as an SRAM chip. Data packets are
20 held in the "scratch pad" 218 during the search, and then can be forwarded in a suitable format, discarded, or re-submitted for searching.

c) the delay buffer 208 is provided by at least one volatile memory chip, for example a DDR RAM SODIMM module such as might be used in a laptop computer. Four 1Gbyte DDR RAM modules would be sufficient for storing three
25 seconds of data traffic at full STM-64 rate, and potentially more in practice (since data traffic is often bursty and 50% or less utilisation is common).

d) the search engine logic is implemented as one or more IP co-processors, for example IDT [RTM] Network Search Engine devices, for performing the actual searching of IP data packets. In one instance, the IP
30 co-processor can support search entries (also referred to as "selectors") 144 bits (18 bytes) in length and each IP co-processor has a selector table of 65,536 entries, and can perform 100 million matches per second across all of these.

At higher rates, four IP co-processors would then be required to handle the 400 million searches/second expected. Each of the four co-processors would, in this case, be provided with the same selector table (i.e. the CAM array stores the same table of entries for which to search); the higher throughput being
5 achieved by running the chipsets in parallel for speed.

e) finally, the interface component 212 is conveniently provided by an Peripheral Component Interconnect (PCI) interface. The PCI interface component 212 provides high bandwidth direct memory access (DMA) scatter-gather egress to software sub-systems, via a PCI backplane 214. It is
10 also used for notifying of events (via interrupts), and accepting new tasking. The PCI interface may either be provided as an integrated component within the firmware of the processing logic 206 or as a separate hardware "bridge" unit with direct communication to an interface controller component of the processing logic. A preferred interface standard is the Compact Peripheral
15 Component Interconnect (cPCI), a hardware interconnection standard that has been specially developed for telecommunications and telephony applications. cPCI devices are manufactured to operate at high data rates, to have high reliability and to permit hot-swapping. The apparatus of the present invention may suitably be implemented as one or more cPCI cards, thereby providing
20 high-speed IP interception and filtering in a low form factor standard interface format (i.e. single slot cPCI).

In this IP-specific embodiment, the inventive processing apparatus facilitates real-time application layer processing of high bandwidth IP streams. Furthermore, the apparatus can screen this high bandwidth IP data traffic for
25 content-based threats at full line rate.

Apart from the benefits of the searching operations, the invention benefits from the use of a delay buffer with multiple taps. The delay buffer allows for a number of latencies, including: the time taken for the hardware to interrupt the software; the time required for the software to process the incoming
30 data packet, and retask the hardware; as well as LAN or WAN delays when remote tasking is performed.

Interrupt delays will almost always be less than 5 milliseconds, even on general purpose operating systems. With real-time extensions, the latency will be less than 50 microseconds.

Network delays will obviously depend on the characteristics of the link:
5 from less than a millisecond for a fast point-to-point connection, to over 100 milliseconds for a WAN connection over a satellite link.

The size of the delay buffer and its tap points can be chosen to minimise latency (if required) while ensuring that there is sufficient time for the hardware components to interact with the software components.

10 The simplest example of this interaction between software and hardware is discussed above: a free form search followed by the detection of other data packets in an IP stream by retasking the search engine logic to search the output of the delay buffer for an IPQ. Even this example conveys some of the power and versatility of the apparatus of the invention.

15 In another preferred embodiment, the presence of two or more outputs from the delay buffer allows more complex search strategies to be constructed in parallel and/or in series with the application of successive search processes. For example, free form searches for different predetermined bit patterns (different virus signatures, say) can be combined, giving "qualified" search
20 strategies of a type "<signature_1> AND NOT (<signature_2> AND <signature_3>". Only data packets that match the combination will cause the software to take the output from the delay buffer to retrieve the rest of an IP stream.

Searches can also be constructed within an IP stream by applying the
25 successive search criteria as the data emerges from the successive taps in the tapped delay line. The search terms used in successive searches need not be predetermined: "<signature_n>" can be changed by software 'on the fly', based on the information found within the data stream from an earlier tap (that is to say, the search engine logic is retasked in dependence upon the result of an
30 earlier search while the data is still present in the delay buffer). Later searches benefit from the results of earlier searches of the same data packets.

To illustrate, consider a suitable delay buffer with three outputs (or "taps") a 'hit' occurs in traffic at tap 1 and at least some of the matched packet(s) are passed to software. The software analyses the packet(s) and decides to place an alternative/additional search on the data in time before it re-emerges at tap 2. The additional search criteria used (for instance, a bit pattern or a header field) can be designated by the software, based on the contents it finds in the analysed packet(s). A new (but related) hit occurs at tap 2 as a result of the new search terms and some further matching packet(s) are passed to software. The software performs further analysis and, if required, sets up yet another search sufficiently early for the data to be scanned yet again (and with a different criteria again) at tap 3.

It is noted that while delay buffers with three outputs as described above are suitable, the delay buffer is not limited to this number of outputs. Different numbers of tap points are provided in further embodiments of the invention. The number of taps does limit the number of searches that can be performed on the same data in series with intervening retasking of the search engine logic.

In a preferred embodiment, the invention may be implemented as a system comprising a hardware component and a software component. Data traffic is received at full line rate (high bandwidth) by the hardware component. The hardware component acts as a coarse filter upon the data traffic, only passing data that fulfils certain criteria to the software component. The software component (a security element, for example) further restricts the data upon which an action needs to be performed.

The data traffic being monitored and filtered comprises a plurality of data packets. Each data packet crossing the apparatus is checked against a set of criteria defined for each security element (for example virus signatures, intrusion/attack signatures, spam signatures, content fields, or IP socket numbers). Whenever a data packet matches the specific criteria, it is passed to a separate sub-system for more detailed analysis, for removal, for quarantine and/or for correction.

The separate sub-system includes a processor that executes the software component of the system. In many instances, the software component

is arranged to process the matched data packets and, in consequence, to feedback further bit patterns which may be of interest. Such a combination search facility is of particular use in the detection of streams of data that contain matching bit patterns.

5 Consequently, the hardware logic units need to perform a number of tasks: adding incoming data packets to the delay path in real-time; screening (filtering) incoming data packets in real-time; passing 'safe' data packets onward with low latency; blocking 'dangerous' data packets and passing them to software; and, updating and managing filtering parameters. Some of these
10 tasks need to be done in hardware, for example filtering, since a bandwidth of, say, 311 Mbytes/s (for STM-16) is too fast for software executing on general purpose computers. Whereas other tasks are better done in software, for example, the update and management of filtering parameters, and deeper inspection of suspicious data packets.

15 In a further enhancement of the invention, data packets that match one of the criteria are then subjected to a post-selection matching process. The matching data packets are analysed to locate certain suitable fields, these fields can then be 'tagged' for easier recovery by software applications. An example
20 might be the identification of data packets as potential threats followed by a step in which the "to:" field of the data packet is marked for attention. Software applications analysing the potentially threatening data packets can thereby easily locate the 'to:' field in that data packet without requiring any further processing.

25 Basic hardware search operations are now explained in greater depth, in the context of high bandwidth IP data traffic. For the purposes of this discussion, the apparatus used is the IP-specific apparatus described above, and illustrated with reference to the architecture shown in Figure 2.

30 The firmware processing logic directs the IP co-processor 210 to perform two generic search operations: matching one or more header field contents of an IP data packet with a selector; and matching a bit pattern at any position within an entire IP data packet (Figures 4 and 7). In carrying out either search

operation, the IP co-processors allow the storage of selectors in ternary (0,1 and X, where X is effectively a single bit wildcard indicator).

In one particular instance of the first type of search operation, the IP co-processor matches a class of selector, referred to as an IP "quint", or simply an IPQ, with the contents of one or more IP header fields. This IP quint comprises the contents of five distinct IP header fields: source and destination IP address (2 x 32 bits), IP protocol type (e.g. UDP or TCP, 1 x 8 bits), and UDP/TCP source and destination ports (2 x 16 bits), giving a total of 104 bits (as shown in the shaded fields in Figure 3). The IP quint is then compared to each entry in the selector table. The co-processor's ternary storage capability permits a "mask" to be applied across chosen bit values in a given selector: thus, for instance, an IPQ search can easily be directed to data packets that have source and destination IP addresses in the same domain, even if the full IP addresses do not match. Only 104 bits of the available 144 bits need be used in an IPQ search, however other header fields can be placed in the spare 40 bits (for instance, TCP flags).

Consider an IP co-processor that can perform 100 million look-ups per second, using all 65,536 144-bit selectors. At STM-64, the maximum UDP or TCP data packet rate is 36Mdata packets/s (ie with a minimum data packet size of 28 bytes, plus 8 bytes of POS frame overhead). Therefore, a single IP co-processor can handle IP quint searching at rates characteristic of STM-64 almost three times over, even in the worst case: the "worst case" corresponding to the smallest data packet size.

Matching a bit pattern that may appear at any position within an entire IP data packet (payload or header) or even across a data packet boundary has different requirements. For the same IP co-processor, storing entries 144 bits long, each entry corresponds to 18 bytes of data, i.e. 18 characters represented in ASCII (or 9 characters of a language represented in Unicode, such as Chinese). Longer words or e-mail addresses need to be truncated so that the hardware reduces the hit rate to the level where it can be handled easily in software. Text matching involves heavier loading on the co-processor, and some trade-offs are required to achieve the required volumetrics. As will be

seen, there is an important trade-off between speed of searching and selector table size. The issues are further described below.

Free text searches require far heavier utilisation than IP quint searching. In the simplest case, an 18-byte window is slid over the data packet, with a comparison every byte against all entries in the selector table (e.g. all e-mail addresses), as shown in Figure 4. In the IP co-processor used above, 100 million of these comparisons are possible every second; even at STM-16, the byte rate is 311 Mbytes/s, therefore the co-processor is too slow by a factor of three.

One solution is to perform several searches at once, with successive byte offsets (i.e. eight bit offset), by using spare selectors, with the search term shifted along by a byte each time (at the expense of the available selector length). This is shown in Figure 5.

Using four selectors for each search term allows a single IP co-processor to handle an STM-16 ($311/4 = 77.75 < 100$), with enough spare capacity for 2-3 searches on IP quint in addition, whilst still providing a large enough selector table size and individual selector length. Here, the speed of searching is maintained at the expense of the number of possible selectors.

More IP co-processors can be brought in to extend the number of selectors that could be searched: each additional co-processor maintaining a different selector table. The speed of searching and the number of possible selectors are both then maintained at the expense of added complexity and the need for additional components.

Additional co-processors can also be used in an alternative solution. Figure 6 illustrates a scenario where four IP co-processors are provided, each with their own copy of the same selector table, and where each of the four co-processors is supplied with data packets having differing offsets relative to each other. As noted above, this amounts to the substantially simultaneous processing of the same data structure against four byte-offset selectors.

Pattern matching across data packet boundaries presents a particular problem. Where cross-packet matching is considered necessary, initial free form patterns are sent to software to be divided into two shorter subpatterns.

The subpatterns are both written to the memory array of the IP co-processor as selectors. The predetermined bit patterns are generally of a length, L bits, that is shorter than the data packets being searched. Since the predetermined bit pattern will be split across no more than two data packets, at least one of the two subpatterns will be intact in one or other of the two data packets. Both shorter patterns must therefore be searched on to ensure that instances of the full bit pattern split over two data packets are always detected. Where no match is found for the whole predetermined bit pattern and either subpattern matches a bit pattern from the middle of a data packet, it can safely be assumed that the subpattern match is a false match. The IP co-processor thus restricts the pattern searching for the subpatterns of the predetermined bit pattern to bit patterns at either the start or the end of the searched packet: i.e. less than L bits from the first or last bit in each packet.

Figure 8 illustrates how the apparatus operates in the detection of a virus signature in IP data traffic. The signature of a virus could be anywhere within the IP payloads of one or more data packets passing through the network. The virus signature is therefore searched for in a free form bit pattern search. The bit pattern corresponding to the virus signature is written as a selector onto the CAM array of the IP co-processor.

The incoming data traffic is written at full line rate into both the delay buffer 208 and the register of the IP co-processor 210. The data structure stored in the register is compared with the stored virus signature (step 802): if a match is made, the data packet containing the match (and only that data packet) is passed quickly to software (steps marked by dashed outline, 810). Further inspection then takes place in software.

Since the signature is likely to be only a part of a virus, the software initiates an IPQ search to select other data packets associated with the signature-containing IP data packet, whereby a full IP stream may be isolated from the data path. The IPQ search is initiated by deducing the IP header information that is associated with the IP stream containing the virus signature (step 814), and then writing an entry corresponding to that IP header information to the CAM array of the IP co-processor 210 (step 816). The IP co-

processor 210 will then perform an IPQ search of data stored on the delay buffer 208 (step 804) as that data exits at one of the delay outputs 216.

What then happens to IPQ search hits (step 818) would depend on the service being run – in an AV role, the IP stream would be suppressed or discarded – in a network monitoring role, the IP stream might be passed to the software for further analysis (e.g. passed back to the sender) and storage.

The virus signature may not be detected in a given data packet. In which case the data packet will be allowed to pass (step 820). Such data packets, though they do not themselves contain the virus signature, may yet be part of an IP stream that does contain a virus signature. The copy of such a data packet will still be present on the delay buffer 208, so the second (IPQ) search (step 804) will be able to intercept it at one of the delay outputs 216. No further action is needed if the data packet is determined to be free of the virus signature and unrelated to any IP stream that does contain a signature (the data packet will simply be discarded (step 830)). Of course, many further searches may be applied to that data packet (step 820) by the IP co-processor.

Note that the architecture proposed has allowed an 'in-data packet' virus detection to be performed at full line rate, without attempting to process every byte of data traffic in software (which is prohibitive at high data rates).

The FF and IPQ search operations can thus be combined to give highly flexible functionality. The hardware component of the system performs combinations of these searches, at various points on the delay path (216, in Figure 2), as tasked by the control software component. The control software component is also solely responsible for adding and subtracting selectors to/from the IP co-processor.

The preferred operation of the control software component with respect to selectors is now set out in more detail. Initially, a list of bit patterns and/or TCP/IP headers is read in to a software application, either from a file (that could be in binary, hex, or ASCII format), or from a user interface.

Adding a selector to the IP co-processor, as might occur once a particular IP stream was identified as containing a matched bit pattern for instance, is relatively simple. The software application maintains a table of

unused entries in the IP co-processor. For each new bit pattern, the next available free entry in the CAM array is chosen. The bit pattern together with the entry is written across the cPCI bus to the line card, which writes the bit pattern to the IP co-processor (this can be done in two write cycles of 9 bytes each, say), and marks it as active in the IP co-processor.

When a match occurs, the IP co-processor returns the index of the match, which enables the software application to correlate the match with the bit pattern that caused it.

Subtracting a selector from the IP co-processor is also straightforward. To detask a bit pattern, effectively deleting it, the software application simply identifies the CAM entry corresponding to the pattern to be detasked, and (via the cPCI bus and line card firmware) marks that entry as inactive in the IP co-processor. Again, this can be initiated from a user interface or from a tasking file.

Placing all of the complexity in the software has a number of advantages. Firstly, new or enhanced protocol support can be implemented and integrated far more easily. Furthermore, for scenarios where several units need to be deployed (either across the cPCI backplane, or with LAN/WAN links), and where relevant data packets can arrive at any of the interfaces, the processing needs be distributed and hence done in software. VoIP "go and return" paths is an example of this.

Preferred embodiments combine many aspects of the invention: header/data packet searches, delay buffer, software control, and remote tasking. In consequence, they enable support for any IP-based protocol to be added easily: for example, VoIP (multiple session reconstruction, text and binary searching); IP usage statistics (survey functionality); and e-mail (intensive text searching).

The firmware processing logic will provide a few auxiliary services, such as data packet counting.

Although the preceding discussion deals mainly with an embodiment of the apparatus for detecting bit patterns in IP traffic, the present invention is not limited to handling IP data traffic. Indeed, a similar architecture to that

illustrated in Figure 2 allows real-time processing of a variety of complex data packetized protocols without the need for development of complex hardware or firmware. Examples of complex protocols to which this architecture can be applied include: the TCP/IP protocol family (including RTP, UDP and ICMP);
5 ATM; X.25; NetBEUI; and DECNET.

In the example shown in Figure 9, the hardware interception and filtering apparatus of the present invention is implemented as a cPCI processing card 902,904 for slotting into a cPCI backbone (bus) 920. The separate software sub-system, where matching data packets are sent for more detailed analysis,
10 is implemented as a single board computer 910 that also slots into the cPCI backbone 920. The computer may conveniently be based upon the x86 family of processors and may run an off-the-shelf operating system, such as Microsoft Windows [RTM]. The illustrated system has two cPCI processing cards 902,904 that are controlled from the single board computer 910. In practice,
15 there can be an arbitrary, many-to-many relationship. For example, if a high percentage of the raw traffic must be processed in software, then several PCs can service a single interception card, with the on-board firmware (the processing logic) balancing the load between the PCs. Conversely, if the hit rate of 'dangerous' data packets in hardware is low, then a single PC can task
20 and control many processing cards, reducing the total system cost.

For passive, data packet sniffing applications, the wanted 'target' data packets are output from the software running on the single board computer. For active, streaming applications, such as real-time AV, the 'safe' data packets are forwarded directly from the processing card - this path is not, however, shown
25 in the Figure.

In one embodiment of the present invention, the apparatus is installed centrally, in a network core. When applied centrally, the ability to monitor data packets in real-time in order to identify undesirable, dangerous or illicit content (for instance, unsolicited advertisements, computer viruses, or pornography) for
30 a large number of subscribers allows service providers to offer a brand new service to their subscribers: a "clean internet service". Unlike prior art internet services where the data traffic has to be "cleaned" *after* data has entered the

customer's premises, a clean internet service could be configured to remove all the undesirable, dangerous or illicit content *before* it reaches the customer's premises.

CLAIMS

1. An apparatus for detecting a predetermined bit pattern in data traffic comprising a plurality of data packets, the data packets having a header portion and a payload portion, the apparatus comprising:
- 5 a processing logic for receiving the data traffic in real-time and for dividing the data packets in the data traffic into bit sequences; and
- an Internet Protocol (IP) co-processor unit having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the
- 10 IP co-processor unit comprising means for identifying, from within at least the payload portion, each data packet that contains a bit sequence that matches a predetermined bit pattern entry within the memory array.
2. An apparatus as claimed in claim 1, wherein the bit pattern storage
- 15 memory array stores bit patterns as ternary data.
3. An apparatus as claimed in claim 1 or claim 2, wherein the bit sequences supplied to the IP co-processor unit are consecutive sequences of bits from a
- 20 single data packet.
4. An apparatus as claimed in claim 3, wherein at least one of the predetermined bit patterns comprises two subpatterns, the at least one bit pattern having a length L bits, each subpattern being stored as a separate bit pattern entry in the memory array;
- 25 wherein the identifying means compares the subpatterns with bit sequences from an end portion of each data packet, the portion being within L bits of the beginning or of the end of the data packet; and
- wherein the identifying means identifies each data packet that contains a bit sequence that matches one or other of the two subpatterns.
- 30
5. An apparatus as claimed in any one of the preceding claims, further comprising a control interface, the control interface being in communication with

the processing logic, the interface conveying software instructions to the processing logic to alter the bit pattern entries in the IP co-processor unit.

5 6. An apparatus as claimed in claim 5, wherein the control interface is integral with the processing logic.

7. An apparatus as claimed in any one of the preceding claims, wherein the bit pattern storage memory array of the IP co-processor unit includes a plurality of bit offset entries corresponding to one or more predetermined bit patterns.

10

8. An apparatus as claimed in any one of the preceding claims, further comprising at least one further IP co-processor unit, the or each further IP co-processor unit operating upon bit sequences having a predetermined offset relative to the bit sequences searched by the IP co-processor unit.

15

9. An apparatus as claimed in claim 8, wherein the or each further IP co-processor unit is provided with an identical copy of the contents of the bit pattern storage memory array of the IP co-processor unit.

20 10. An apparatus as claimed in any one of the preceding claims, the apparatus further comprising a delay buffer for storing the data traffic at full line rate,

25 wherein data streams that contain a data packet that is identified as having a bit sequence that matches the predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for further processing.

30 11. An apparatus for detecting a predetermined bit pattern in data traffic comprising a plurality of data streams of data packets, the apparatus comprising:

a processing logic for receiving the data traffic in real-time and for dividing the data traffic into bit sequences;

a search engine logic having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the search engine logic comprising means for identifying each data packet that contains a bit sequence that matches a predetermined bit pattern entry within the memory array; and

5 a delay buffer for storing the data traffic at full line rate, the delay buffer having one or more outputs,

wherein data streams that contain a data packet that is identified as having a bit sequence that matches a predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for
10 further processing.

12. An apparatus as claimed in claim 11, wherein the bit sequences supplied to the search engine logic are consecutive sequences of bits from a single data packet.

15

13. An apparatus as claimed in claim 12, wherein at least one of the predetermined bit patterns comprises two subpatterns, each subpattern being stored as a separate bit pattern entry in the memory array; and

20 wherein the identifying means compares the subpatterns with the formatted bit sequences and identifies each data packet that contains a bit sequence that matches one or other of the two subpatterns.

14. An apparatus as claimed in claims 11, 12 or 13, wherein the search engine logic comprises means for performing a protocol based search, in which
25 target data packets are identified by virtue of the bit sequence of at least one data field within each data packet.

15. An apparatus as claimed in claim 14, wherein the data field is a header field, the header field including data relating to one of the group including:

30 address information;
to/from port number information; and,
data packet type identifier information.

16. An apparatus as claimed in any one claims 11 to 15, further comprising a control interface, the control interface being in communication with the processing logic, the interface conveying software instructions to the processing logic to alter the bit pattern entries in the search engine logic.
- 5
17. An apparatus as claimed in claim 16, wherein the control interface is integral with the processing logic.
18. An apparatus as claimed in any one claims 11 to 17, wherein the bit
10 pattern storage memory array of the search engine logic includes a plurality of bit offset entries corresponding to one or more predetermined bit patterns.
19. An apparatus as claimed in any one claims 11 to 18, further comprising
15 at least one further search engine logic, the or each further search engine logic operating upon data sequences having a predetermined offset relative to the data sequences searched by the search engine logic.
20. An apparatus as claimed in claim 19, wherein the or each further search
20 engine logic is provided with an identical copy of the contents of the bit pattern storage memory array of the search engine logic.
21. An apparatus as claimed in any one of claims 11 to 20, wherein the search engine logic is an IP co-processor unit.
- 25
22. A system for detecting a predetermined bit pattern in input data traffic comprising a plurality of data packets, the system comprising a hardware component and a software component, the hardware component including:
- 30 a processing logic for receiving the data traffic in real-time and for dividing the data traffic into bit sequences; and
- a search engine logic having a bit pattern storage memory array, in which one or more predetermined bit patterns are stored, the search engine logic

comprising means for identifying each data packet that contains a bit sequence that matches the predetermined bit pattern entry within the memory array,

wherein data packets processed by the search engine logic of the hardware component and identified as containing a matching bit sequence are passed to an executable software application of the software component for further processing, the bandwidth of data traffic passed on to the executable software application of the software component being substantially less than the bandwidth of the data traffic input into the hardware component.

23. A system as claimed in claim 22, wherein the processing logic is programmable in a higher level programming language.

24. A system as claimed in claim 22 or claim 23, wherein the bit sequences supplied to the search engine logic of the hardware component are consecutive sequences of bits from a single data packet.

25. A system as claimed in claim 24, wherein at least one of the predetermined bit patterns comprises two subpatterns, each subpattern being stored as a separate bit pattern entry in the memory array; and

wherein the identifying means compares the subpatterns with the bit sequences and identifies each data packet that contains a bit sequence that matches one or other of the two subpatterns.

26. A system as claimed in any one of claims 22 to 25, wherein the search engine logic of the hardware component comprises means for performing a protocol based search, in which target data packets are identified by virtue of the bit sequence of at least one data field within each data packet.

27. A system as claimed in claim 26, wherein the data field is a header field, the header field including data relating to one of the group including:

address information;

to/from port number information; and,

data packet type identifier information.

28. A system as claimed in any one of claims 22 to 27, wherein the hardware component further comprises a control interface, the control interface being in communication with the processing logic, the interface conveying software instructions from the software component to the processing logic to alter the bit pattern entries in the search engine logic.

29. A system as claimed in claim 28, wherein the control interface is integral with the processing logic.

30. A system as claimed in any one of claims 22 to 29, wherein the bit pattern storage memory array of the search engine logic includes a plurality of bit offset entries corresponding to one or more predetermined bit patterns.

31. A system as claimed in any one of claims 22 to 30, wherein the hardware component further comprises at least one further search engine logic, the or each further search engine logic operating upon bit sequences having a predetermined offset relative to the bit sequences searched by the search engine logic.

32. A system as claimed in claim 31, wherein the or each further search engine logic is provided with an identical copy of the contents of the bit pattern storage memory array of the search engine logic.

33. A system as claimed in any one of claims 22 to 32, wherein the hardware component of the system further comprises a delay buffer for storing the data traffic at full line rate, wherein data streams that contain a data packet that is identified as having a bit sequence that matches the predetermined bit pattern are extracted from an output of the delay buffer and passed to a software application for further processing.

34. A system as claimed in any one of claims 22 to 33, wherein the or each search engine logic of the hardware component is an IP co-processor unit.
35. An apparatus for detecting predetermined bit patterns in data traffic comprising a plurality of data packets, the apparatus comprising:
5 a processing logic for receiving the data traffic in real-time and for dividing the data traffic into bit sequences;
a search engine logic having a bit pattern storage memory array, in which a plurality of predetermined bit patterns are stored, the search engine logic
10 comprising means for identifying each data packet that contains a bit sequence that matches a predetermined bit pattern entry within the memory array; and
a delay buffer for storing the data traffic at full line rate, the delay buffer having an input and a plurality of outputs, each output corresponding to a delay of a respective duration,
15 wherein the search engine logic is adapted to execute different bit pattern search strategies across the data traffic, where data packets which match a particular search strategy are output at a respective output of the delay buffer for further processing.
- 20 36. An apparatus as claimed in claim 35, wherein the search engine logic is adapted to execute a series of bit pattern search strategies, with each successive strategy being dependent upon at least one preceding search strategy.
- 25 37. An apparatus as claimed in claim 36, wherein successive strategies are determined in dependence upon the results of at least one preceding search strategy.

1/6

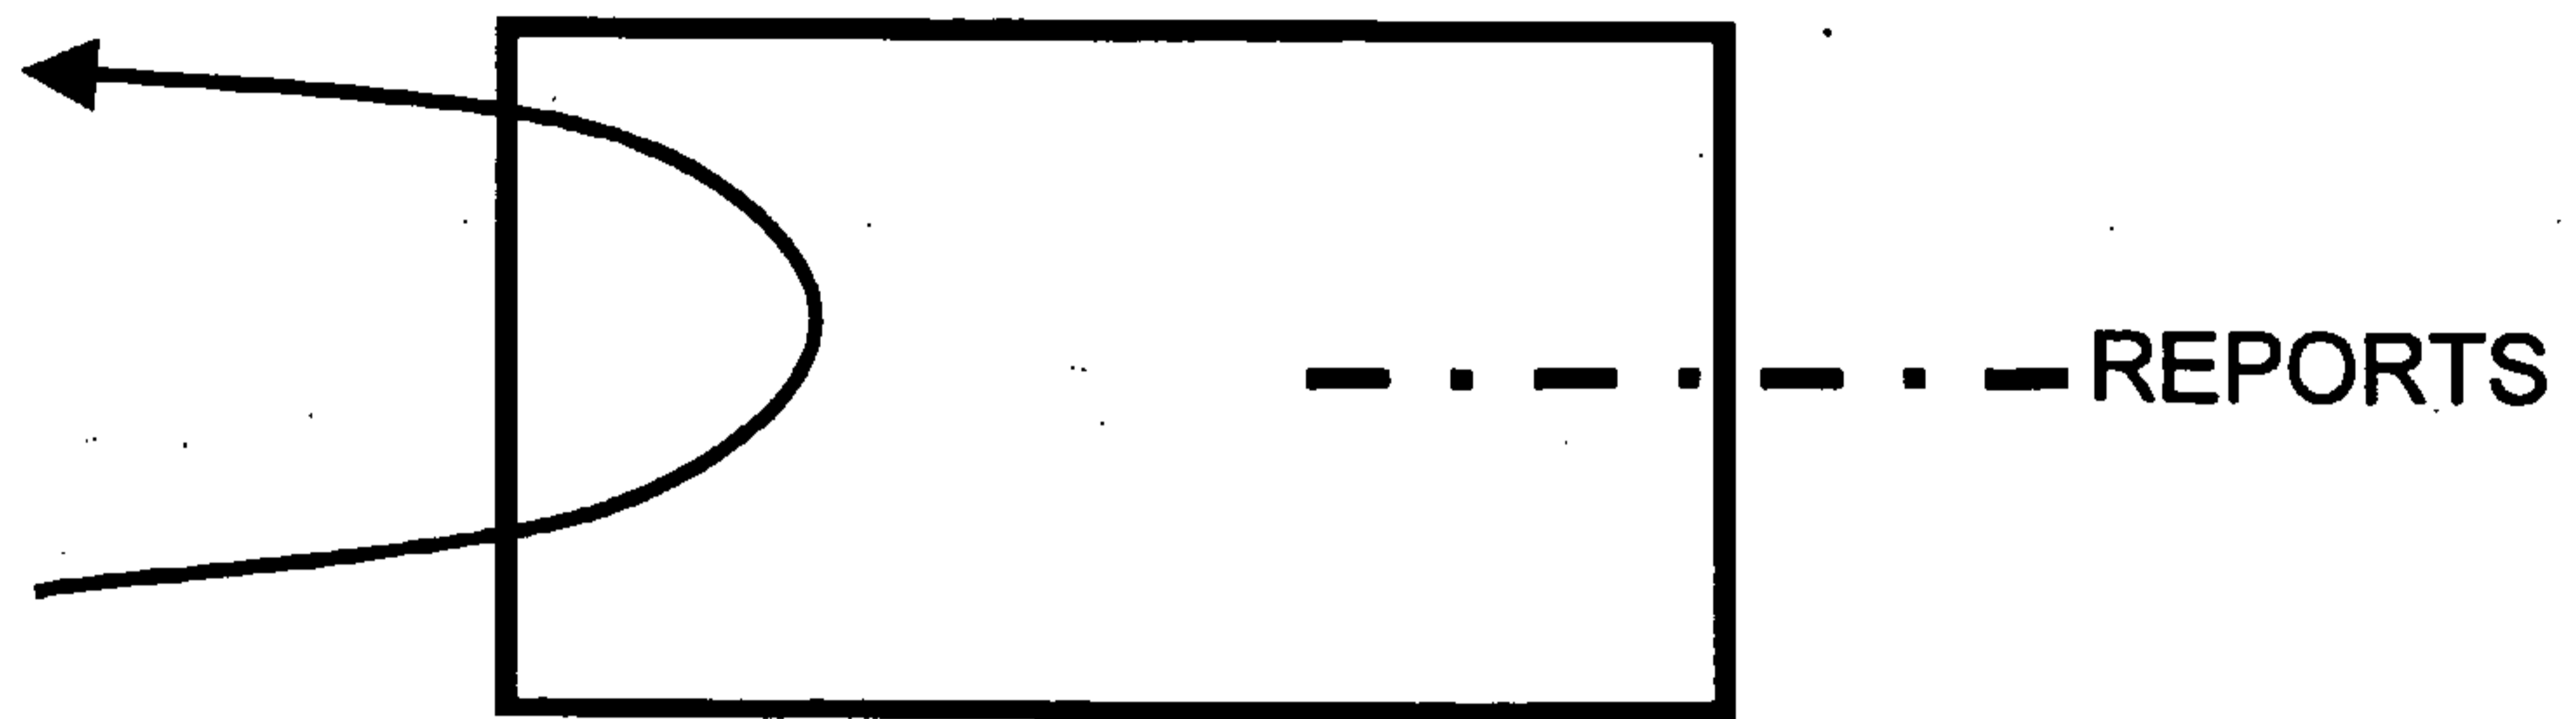


Fig 1A

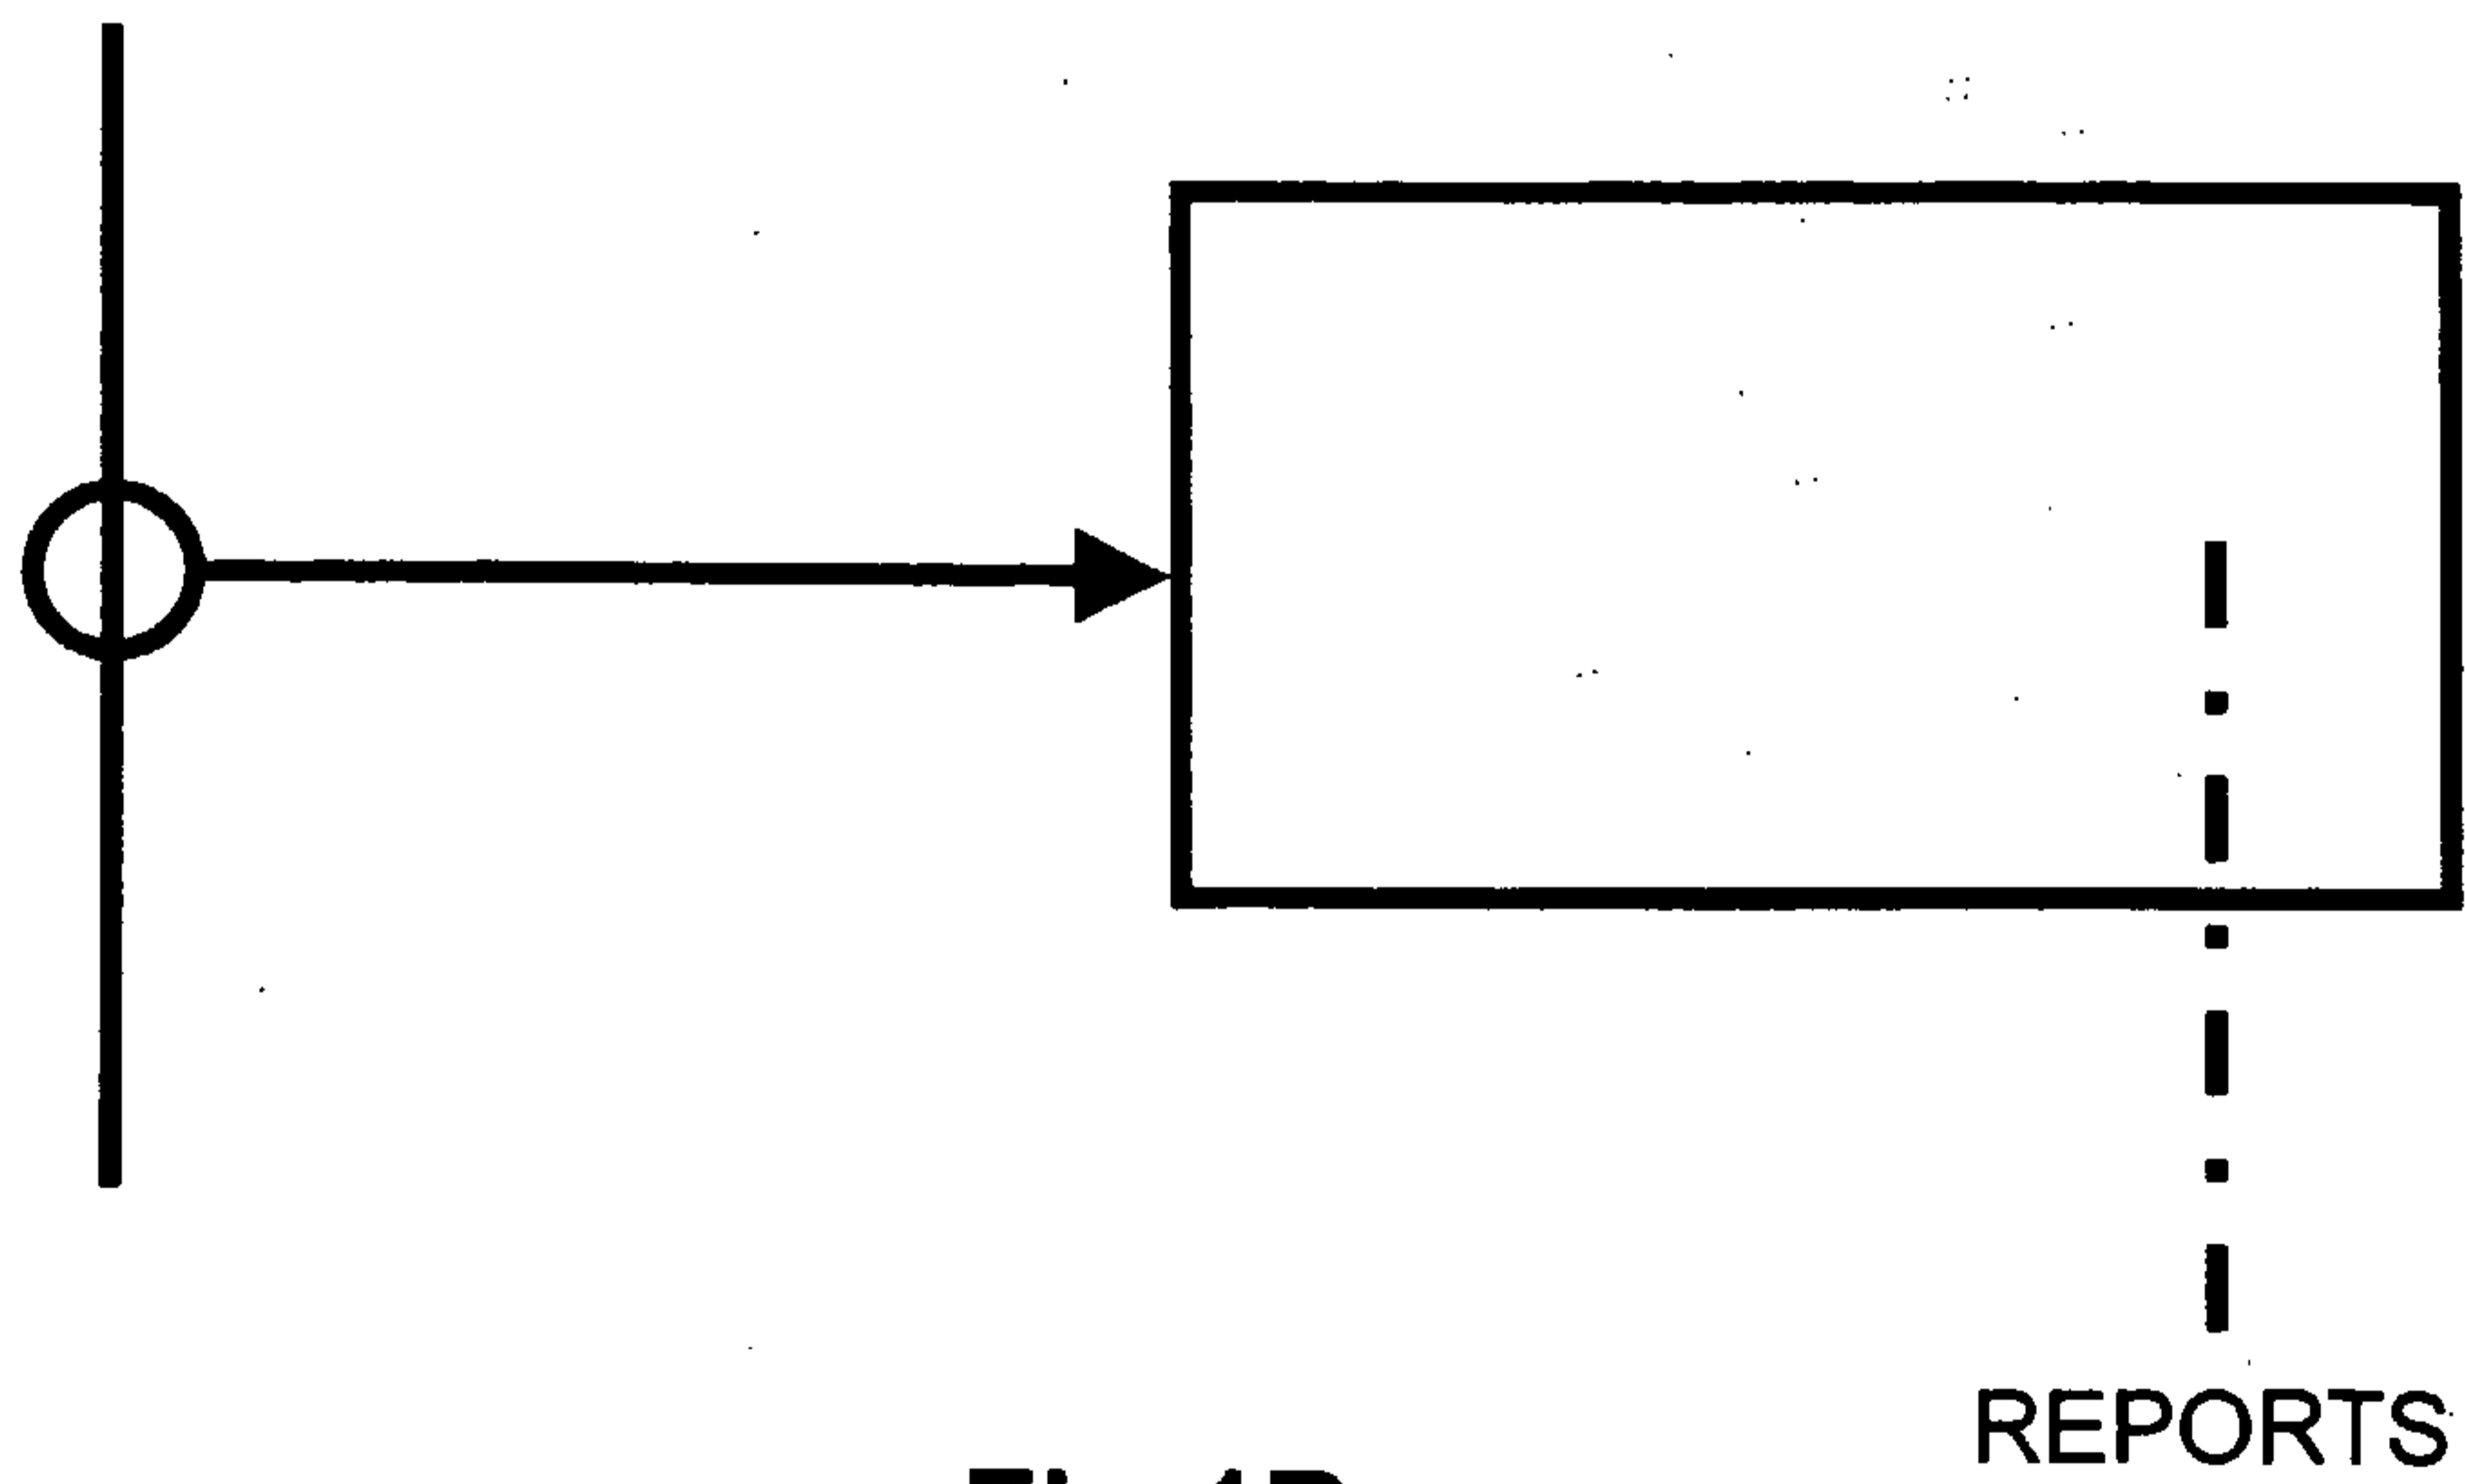


Fig 1B

2/6

Fig.2.

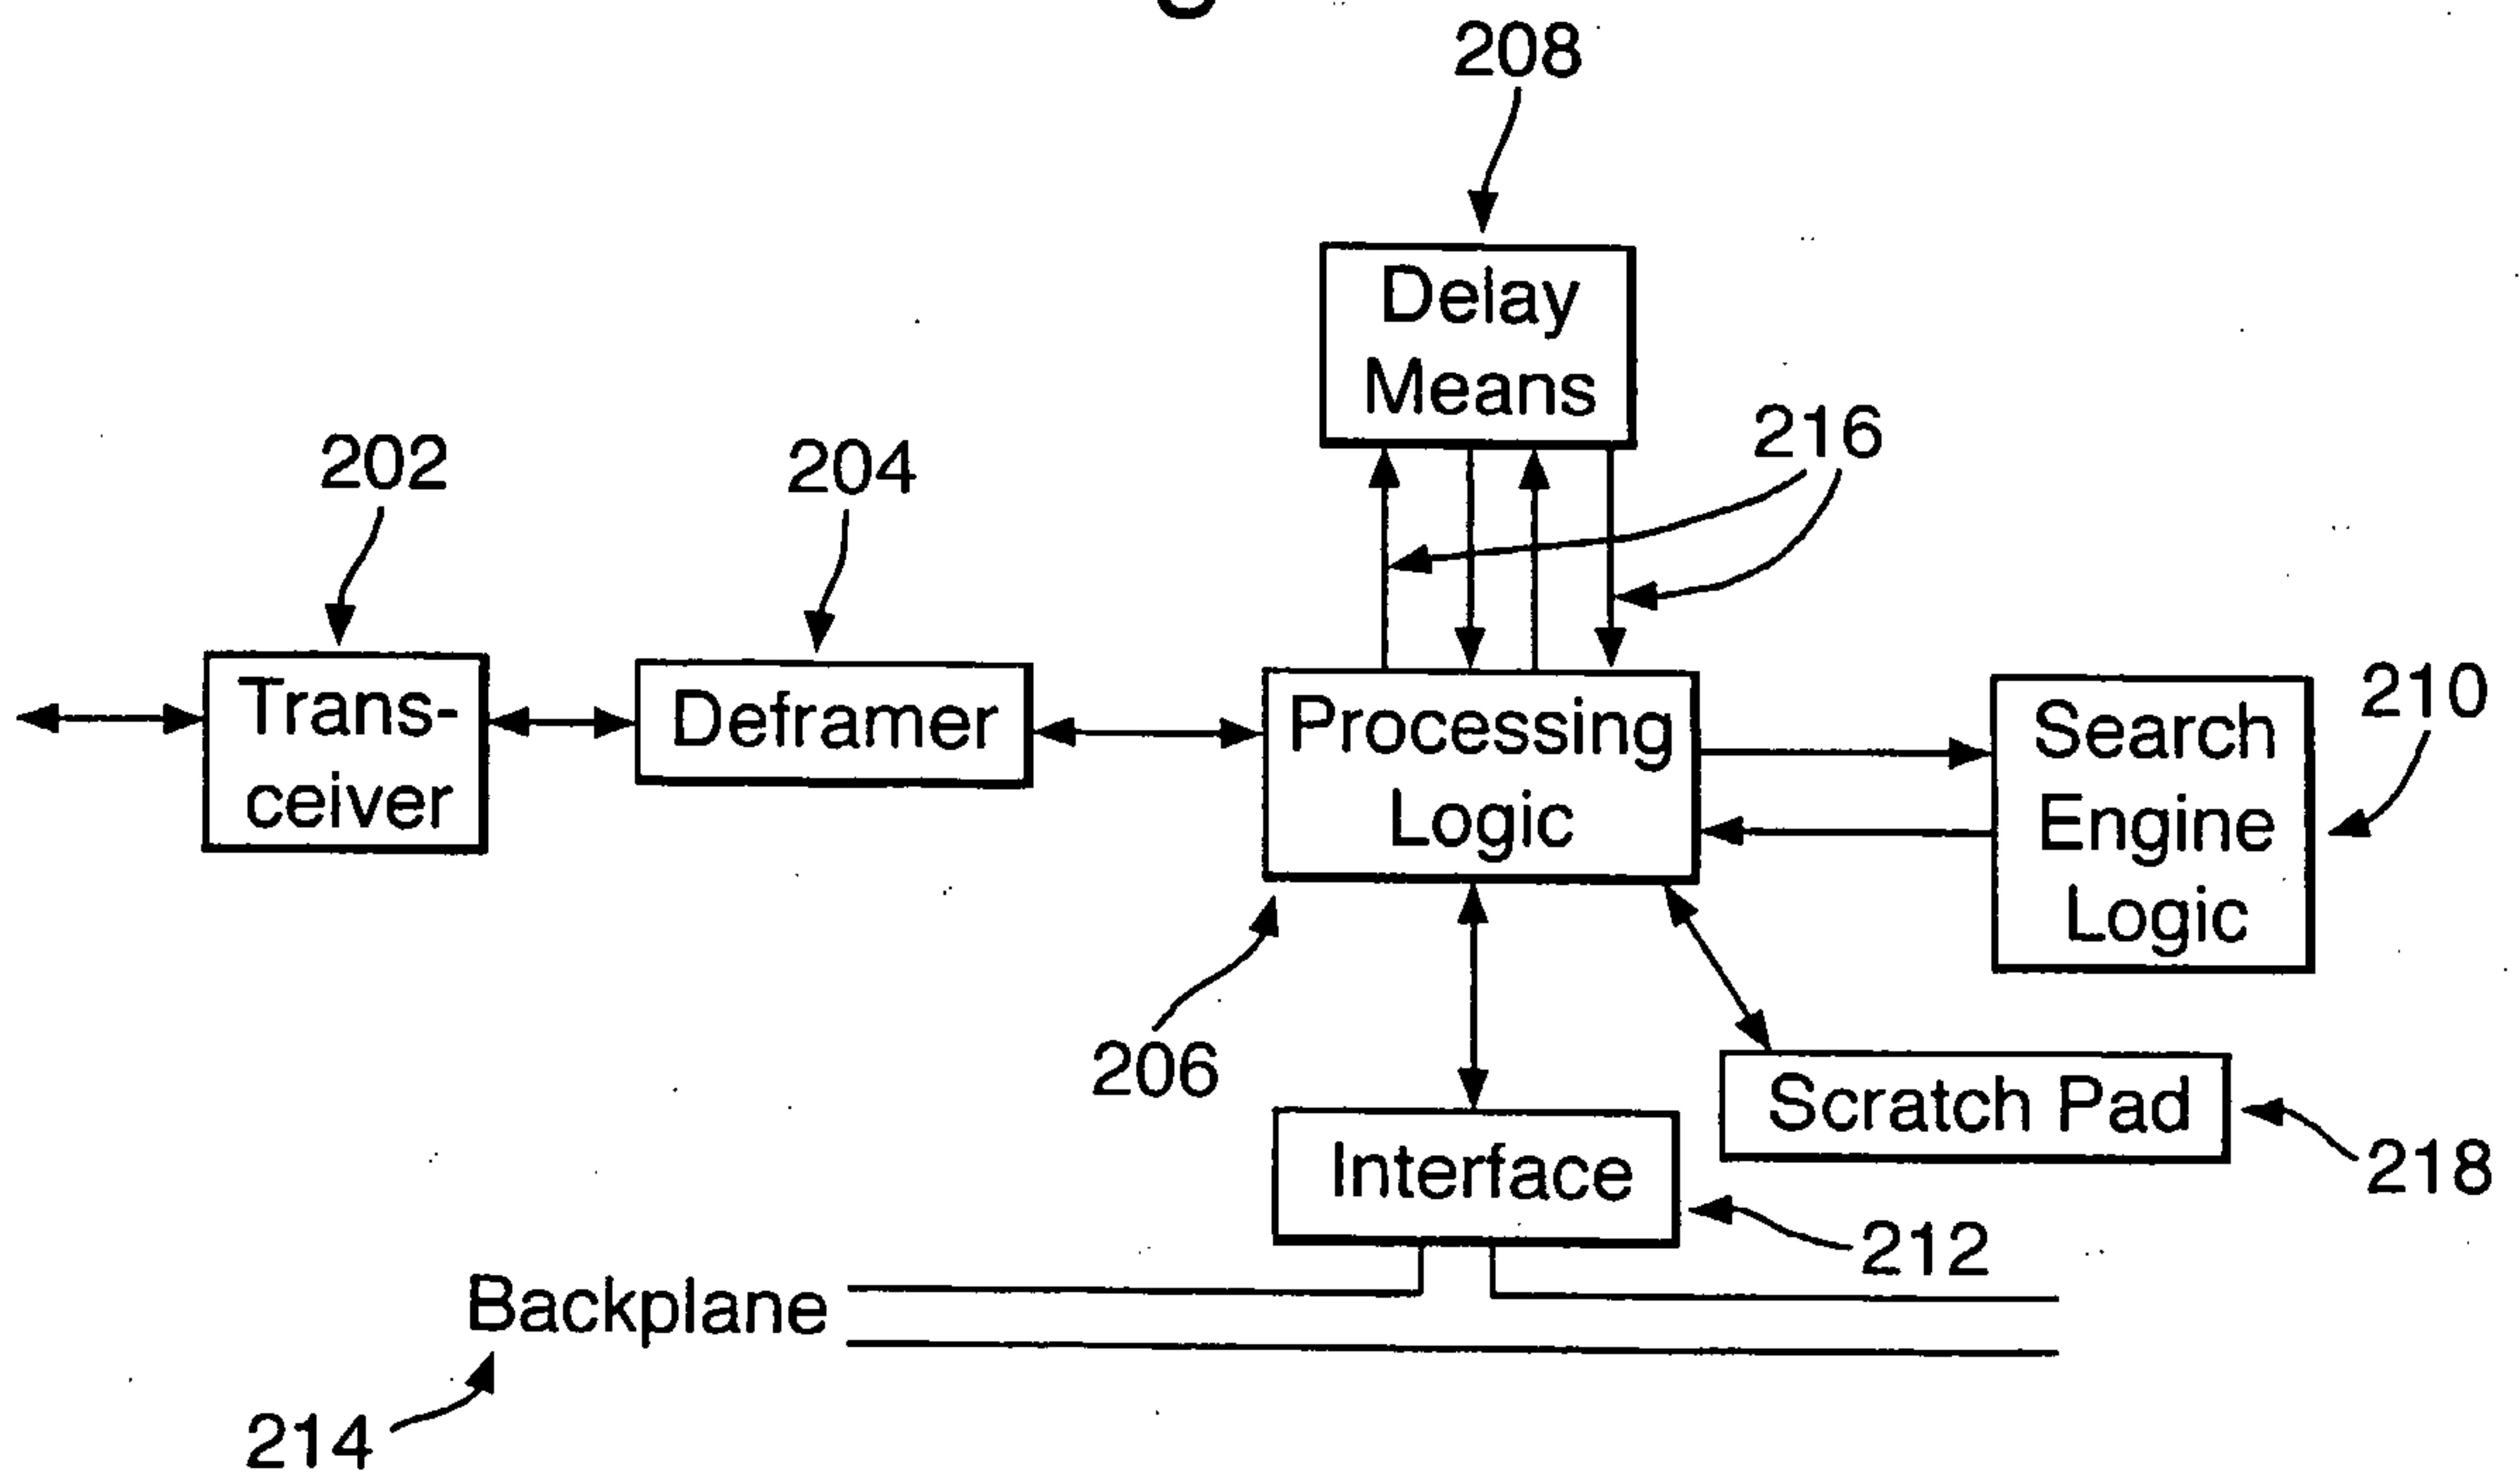
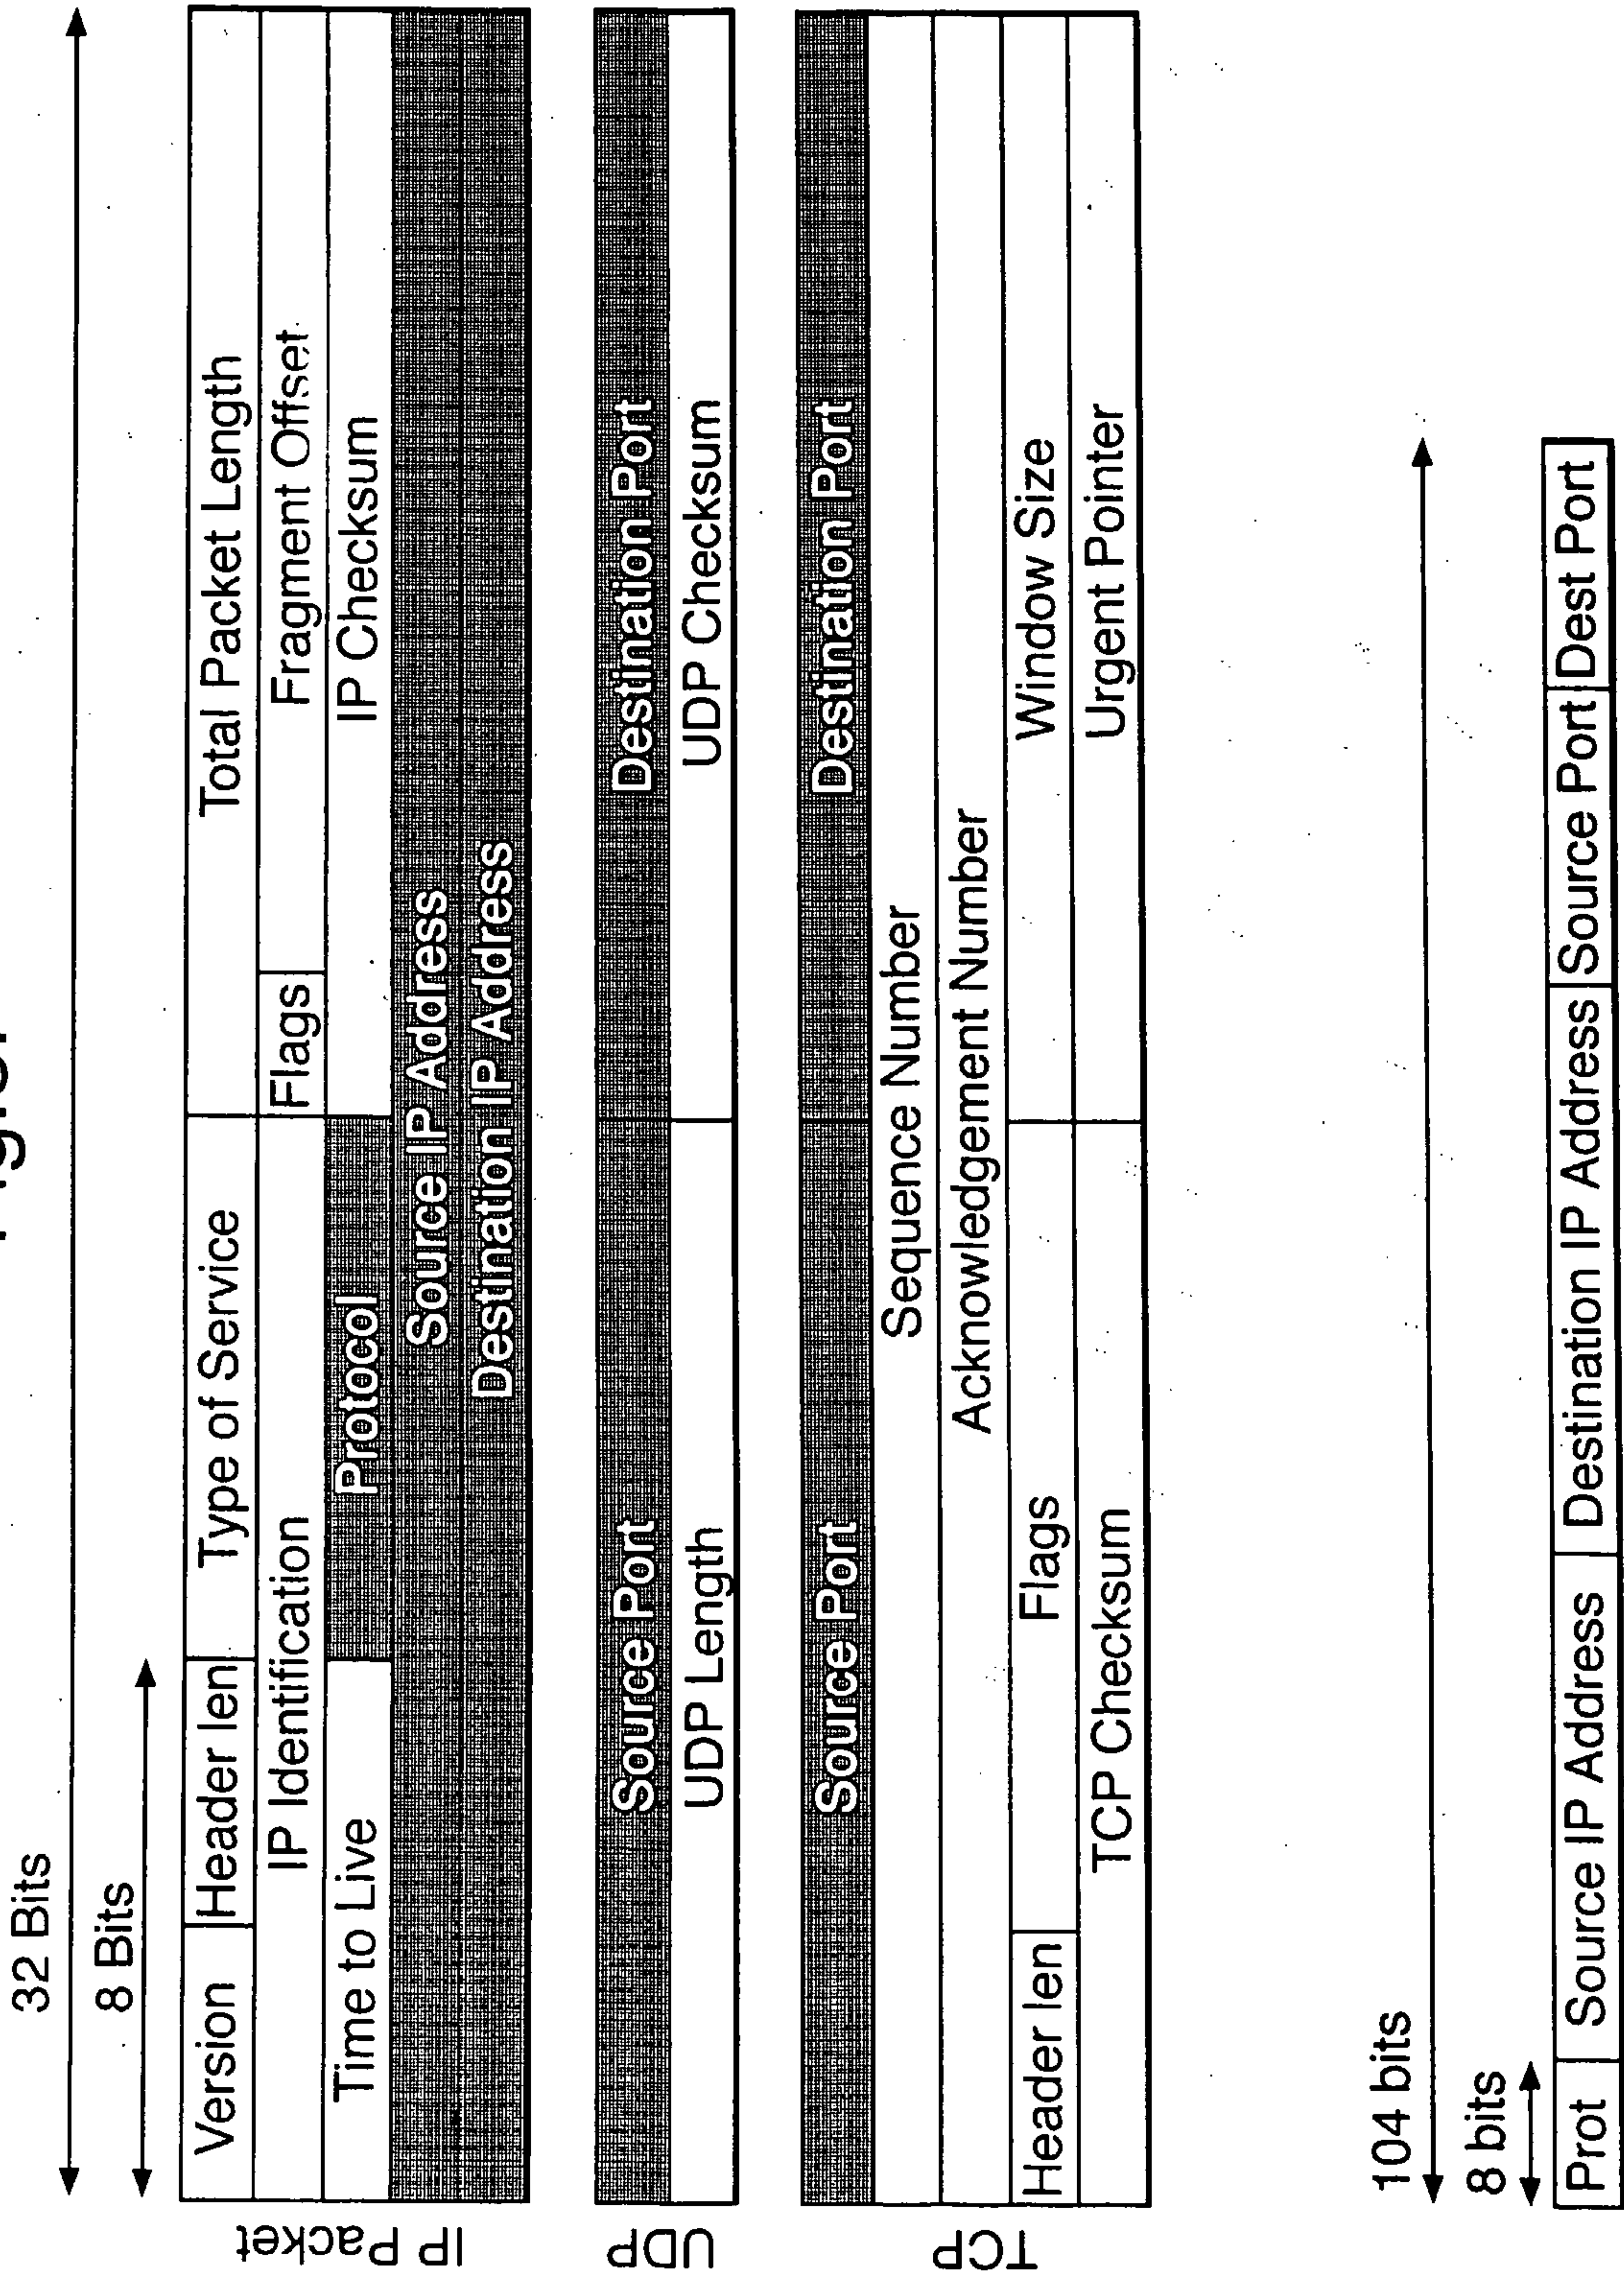


Fig.3.



4/6

Fig.4.

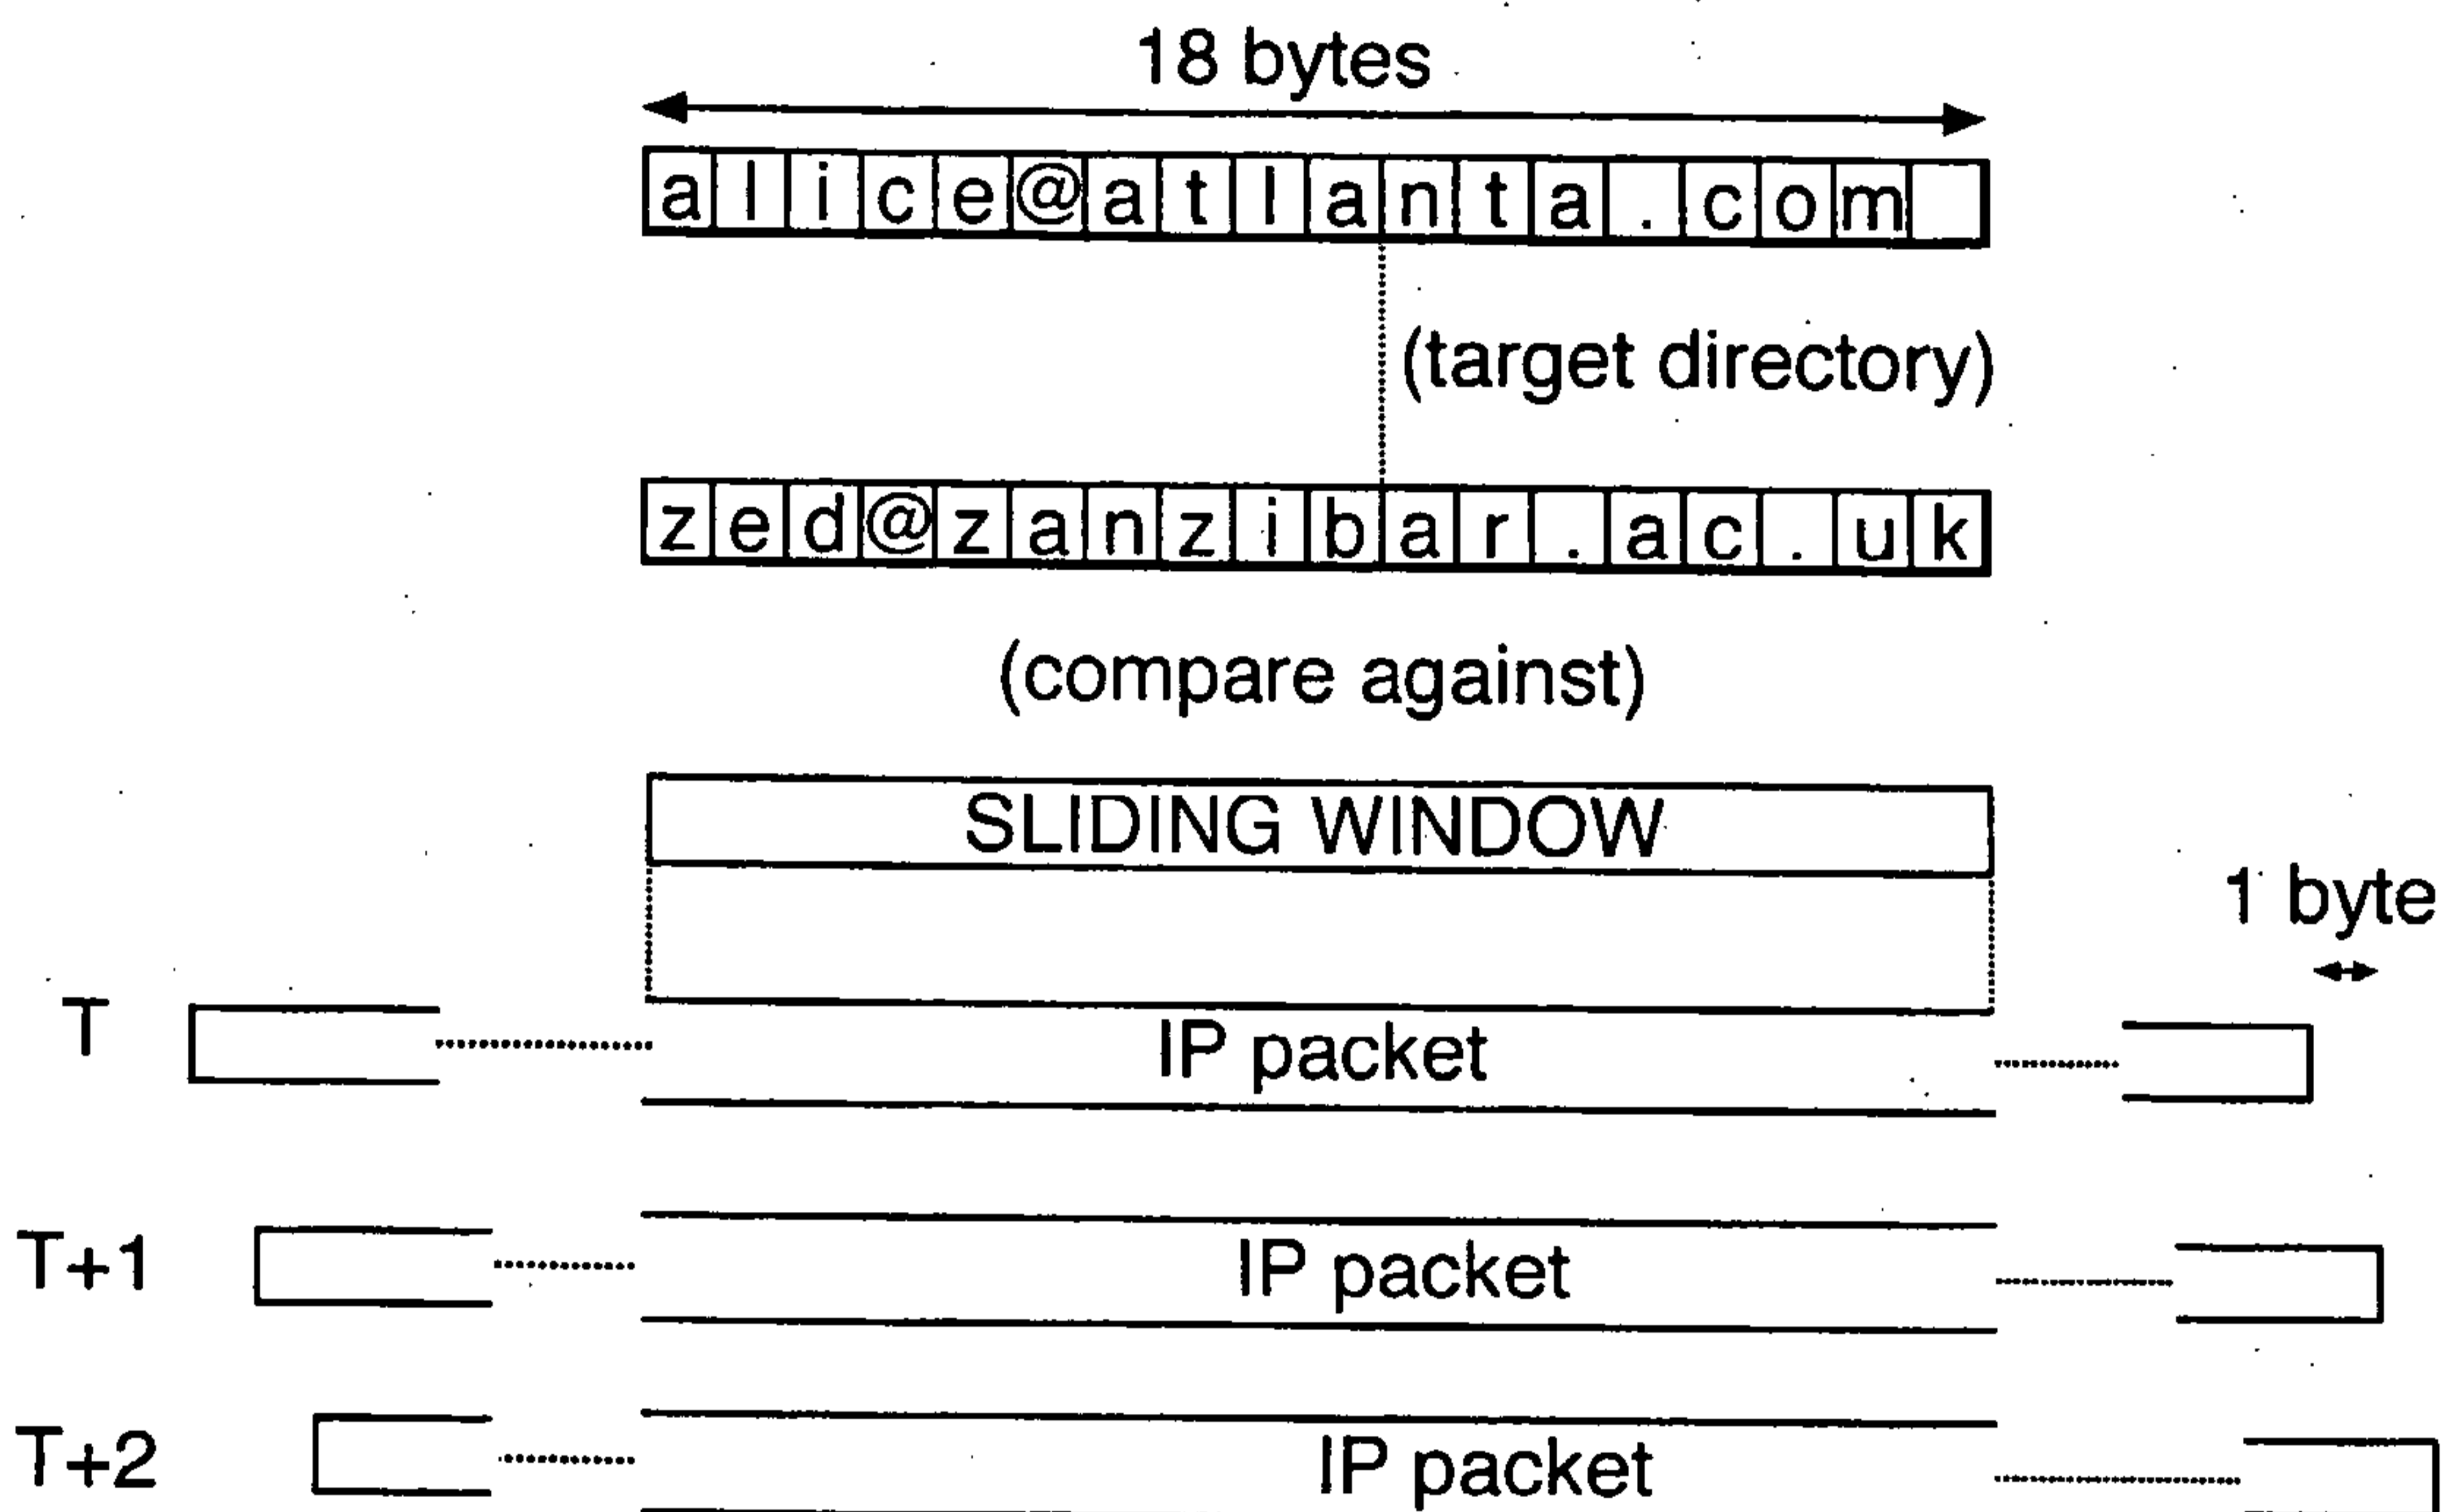


Fig.5.

```

z e d @ z a n z i b a r . a c . u k
z e d @ z a n z i b a r . a c . u
z e d @ z a n z i b a r . a c .
z e d @ z a n z i b a r . a c

```

5/6
Fig.6.

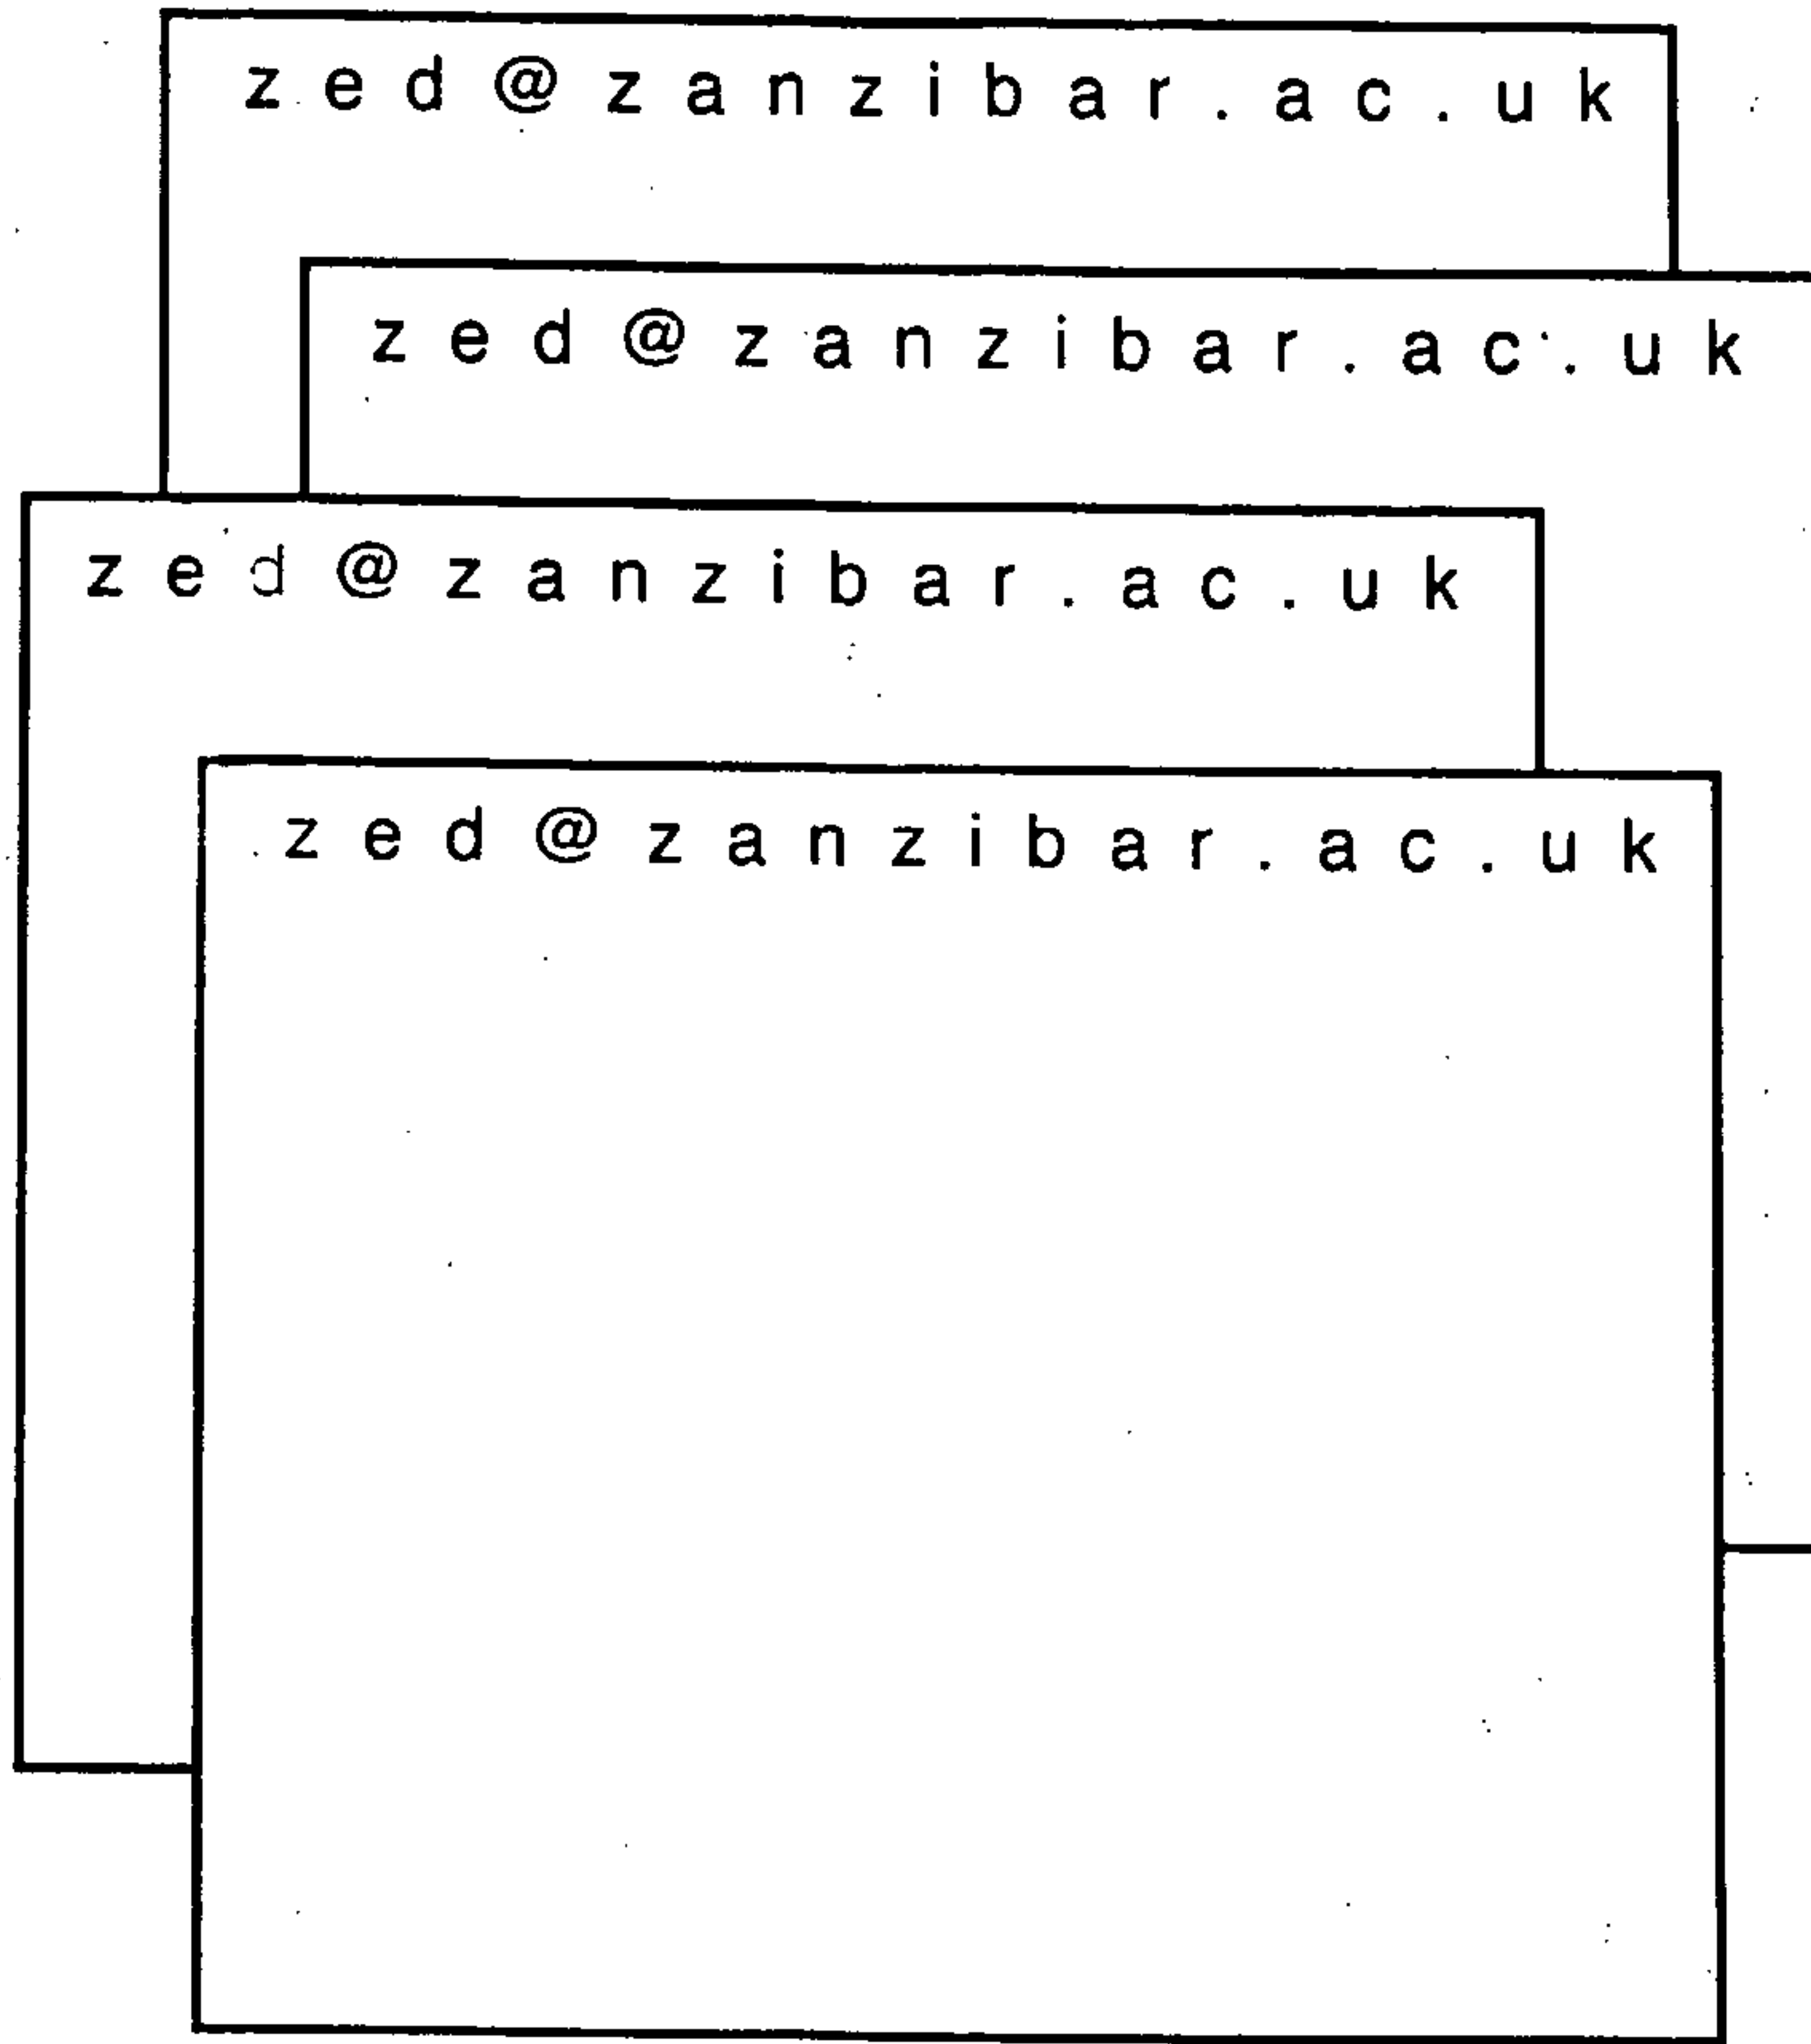
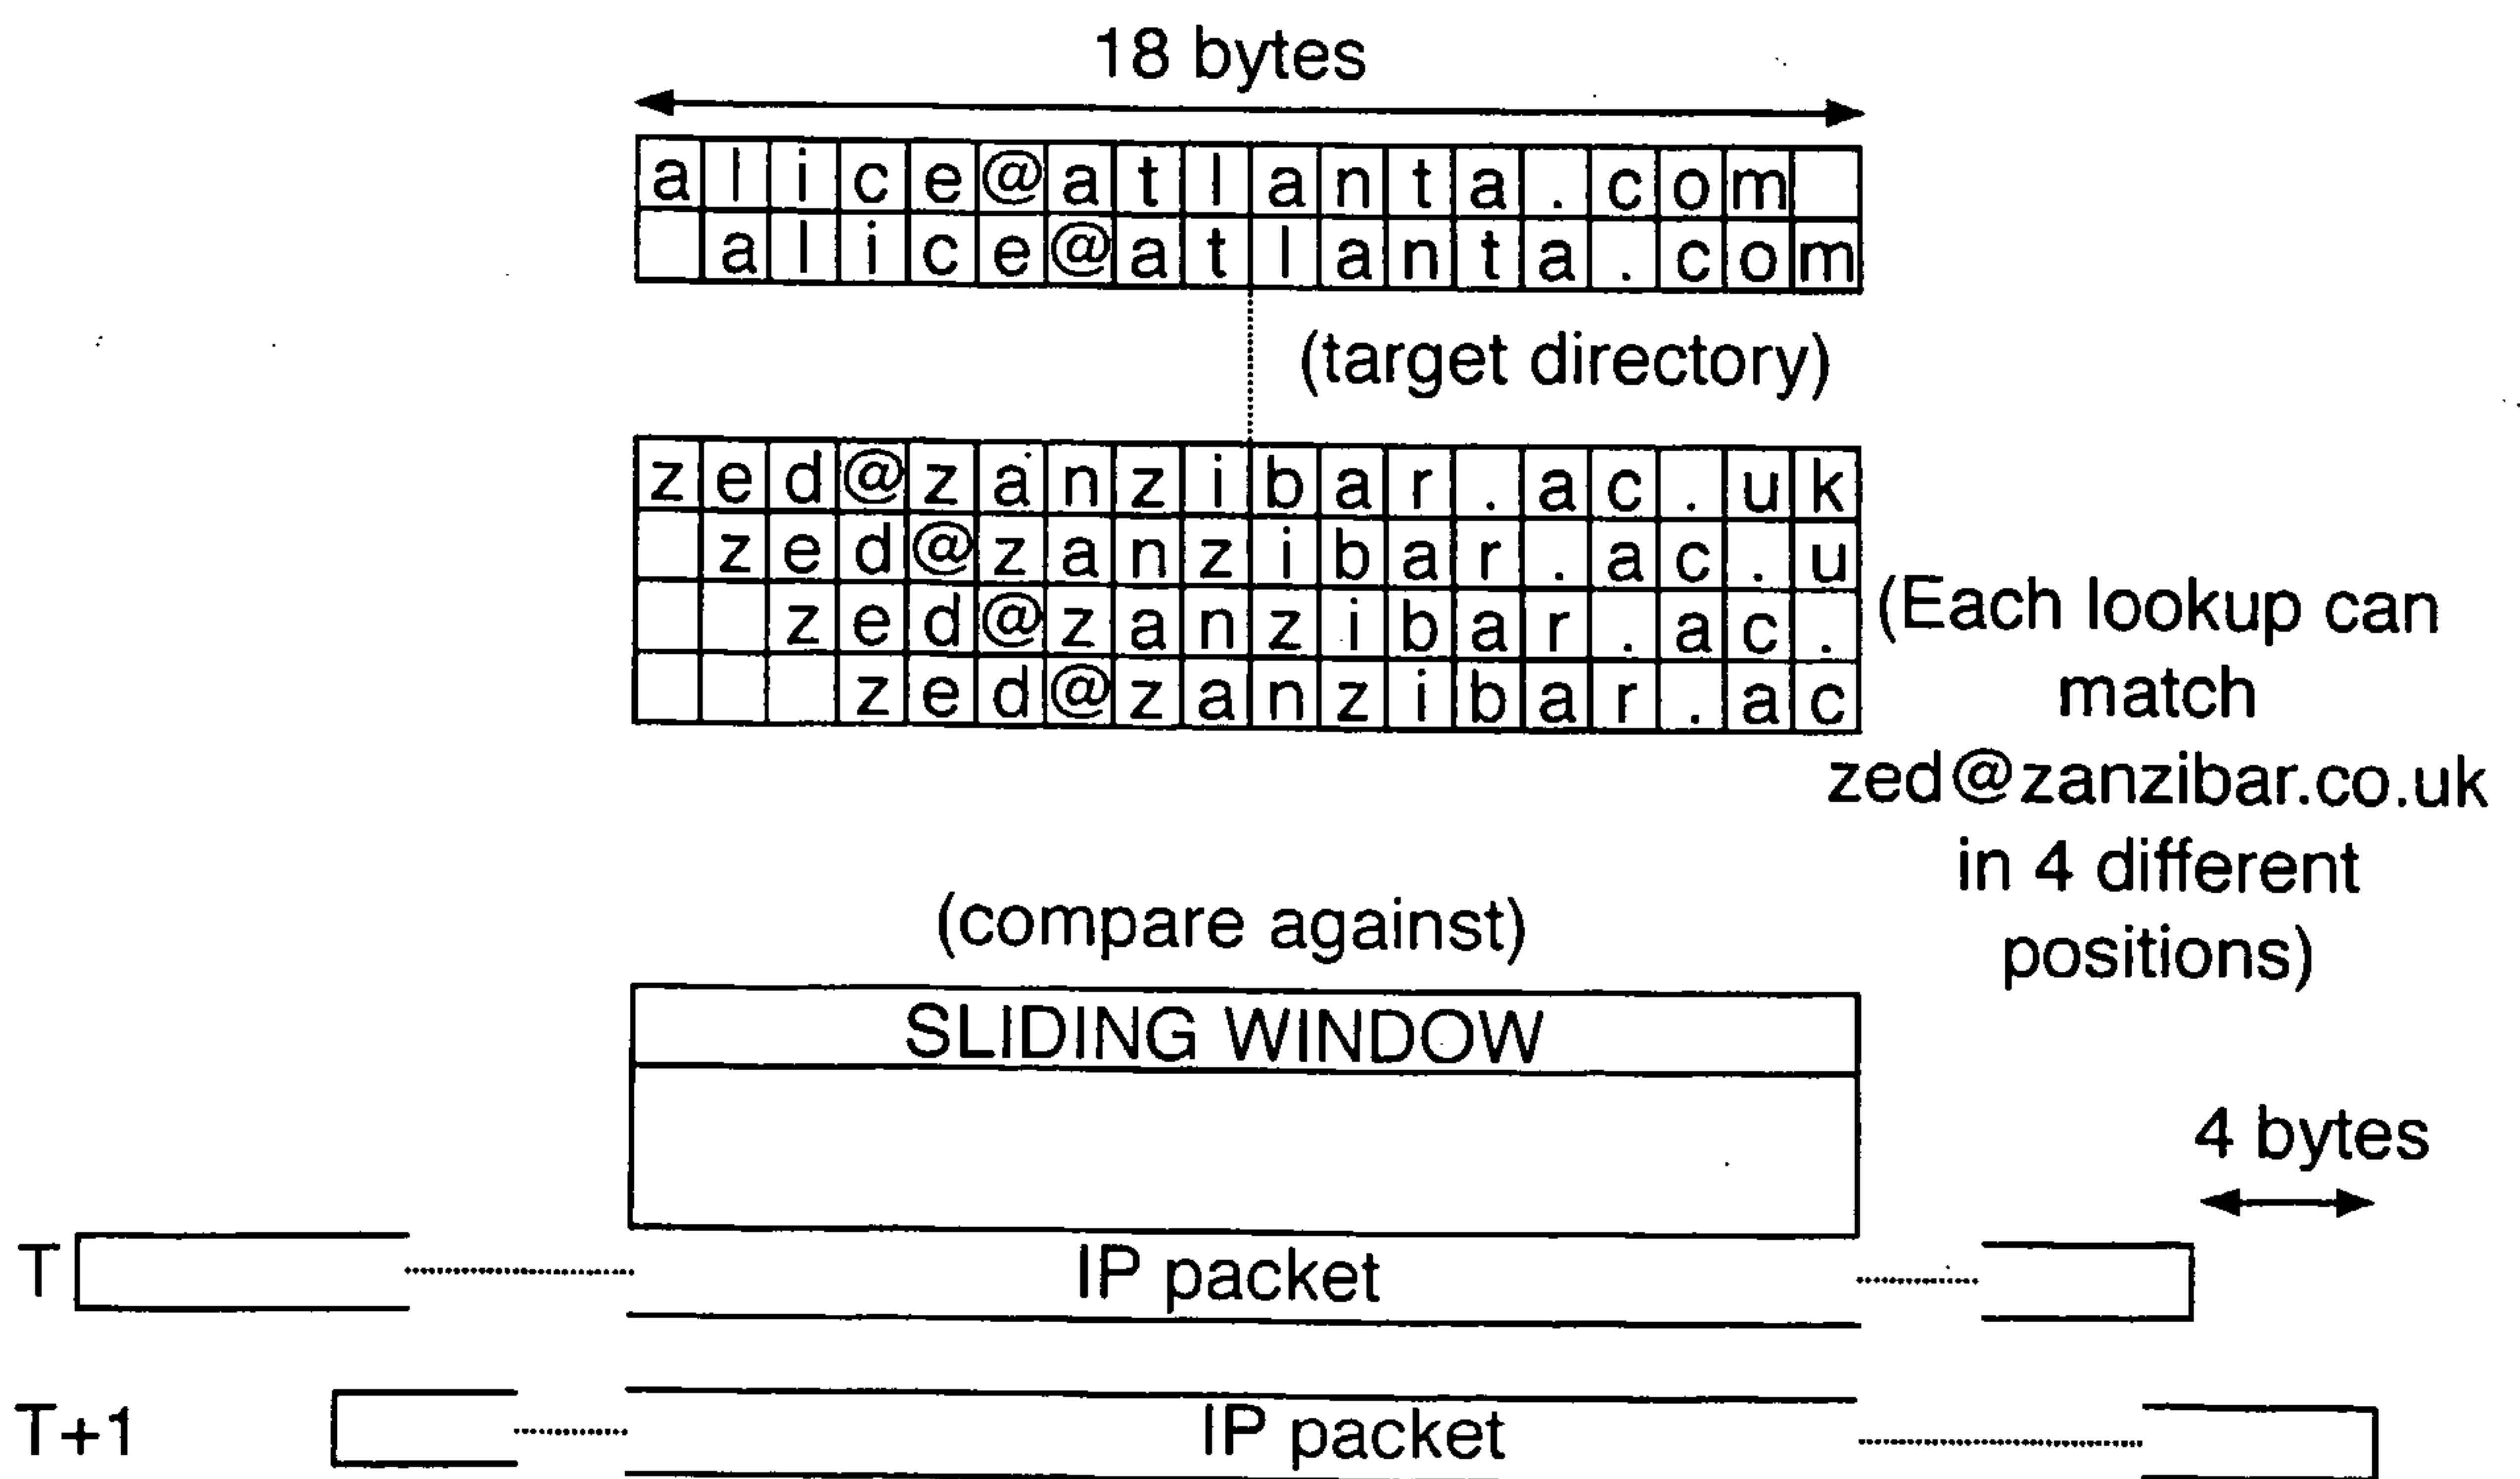


Fig.7.



6/6

Fig.8.

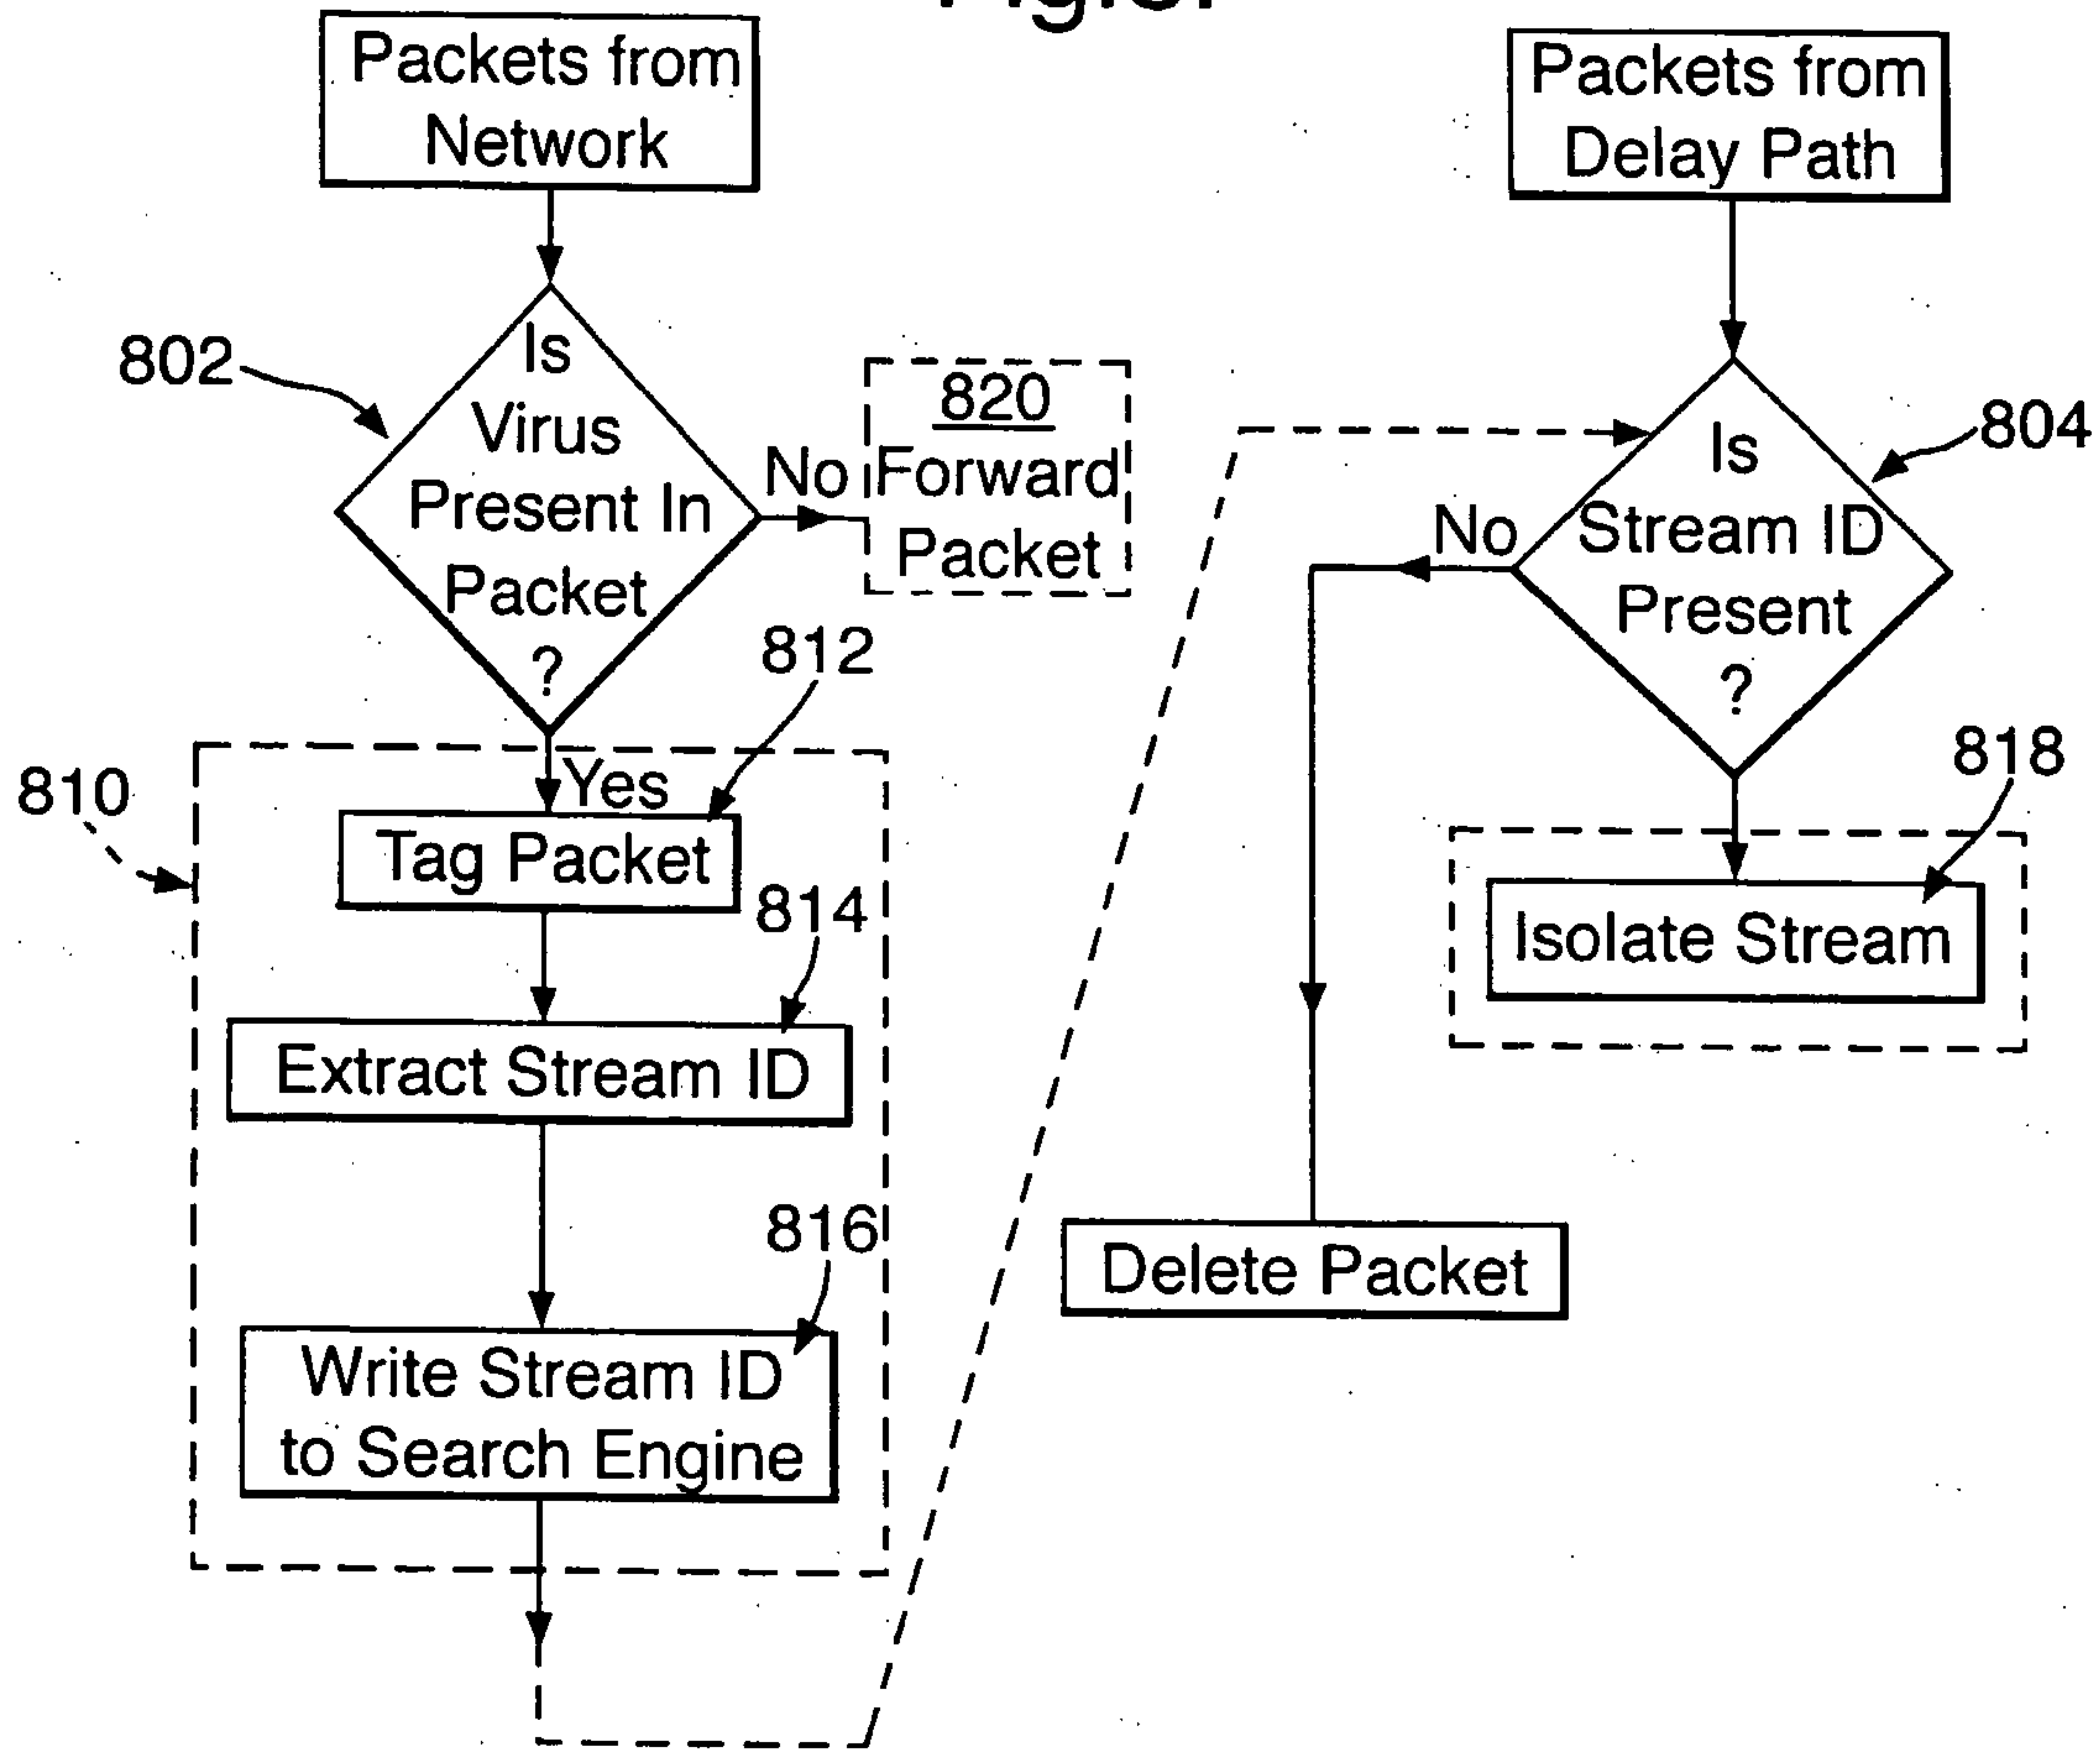


Fig.9.

