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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
27 December 2007 (27.12.2007)

PCT

(10) International Publication Number
WO 2007/150034 A1

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

(21) International Application Number:
PCT/US2007/071908

(22) International Filing Date: 22 June 2007 (22.06.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/816,023 22 June 2006 (22.06.2006) US

(71) Applicant *(for all designated States except US)*: **WISCONSIN ALUMNI RESEARCH FOUNDATION** [US/US]; 614 Waknut Street, 13th Floor, Madison, WI 53726 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **DONG, Qunfeng** [CN/US]; 1210 W. Dayton Street, Madison, WI 53706 (US). **BANERJEE, Suman** [IN/US]; 1210 W. Dayton Street, Madison, WI 53706 (US).

(74) Agent: **PIENKOS, John, T.**; Whyte Hirschboeck Dudek S.c., 555 East Wells Street, Suite 1900, Milwaukee, WI 53202 (US).

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

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

Published:
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

[Continued on next page]

(54) **Title:** METHOD OF DEVELOPING IMPROVED PACKET CLASSIFICATION SYSTEM

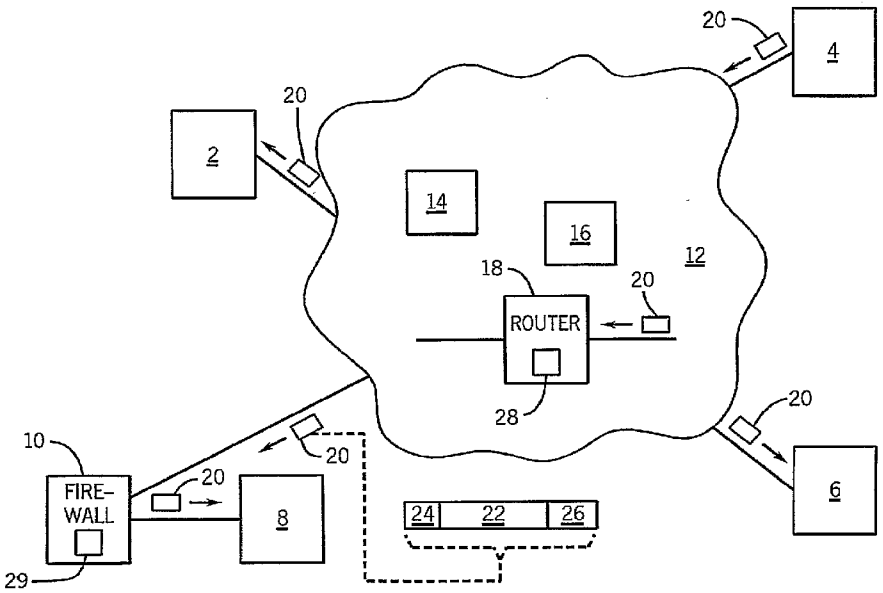

(57) **Abstract:** An method of constructing a packet classification device such as router or firewall, as well as a packet classification device constructed by way of such a method, are disclosed, In at least some embodiments, the method involves performing one or more techniques that convert an original rule set that could be implemented on a packet classification device into a modified, and simplified or shortened, rule set that instead is implemented on the packet classification device. The one or more techniques can include, for example, trimming, expansion, addition of one or more rules, and merging of one or more rules.

WO 2007/150034 A1

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.

METHOD OF DEVELOPING IMPROVED PACKET CLASSIFICATION SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of U.S. provisional patent application no. 60/816,023 filed June 22, 2006 and entitled "METHOD OF DEVELOPING IMPROVED PACKET CLASSIFICATION SYSTEM", which is hereby incorporated by reference herein.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

- -

FIELD OF THE INVENTION

[0002] The present invention relates to packet classification devices and, more particularly, relates to methods of constructing such devices so that the devices perform desired packet classification operations.

BACKGROUND OF THE INVENTION

[0003] The internet is one of the most important media for communicating many different types of information, and its importance continues to grow. Although information communicated by way of the internet can take a variety of forms and follow a variety of protocols, it often is communicated via the TCP/IP protocols (transmission control protocol/internet protocol). Further, information communicated via the internet typically is broken down into packets at the information source prior to transmission. The packets, which are sent onto the internet individually, are then reassembled upon being received at a destination.

[0004] Information packets transmitted via the internet typically pass through a variety of hubs or intermediate nodes in between the sources from which the packets are emanating and the destinations to which the packets are directed. Intermediate devices at these nodes such as routers are capable of identifying, based upon information contained in the packets (e.g., within headers and/or trailers of the packets), the sources from which the packets have come from and the destinations to which the packets are supposed to be transmitted. The routers, upon

identifying the sources and/or destinations, are capable of classifying the packets based upon such information and taking certain actions based upon how the packets are classified, for example, determining how to forward packets to their appropriate destinations.

[0005] The classification of packets conducted by routers as described above is only one type of packet classification that can be performed in relation to packets being transmitted via, provided onto, or received off of the internet. Source data and destination data are only two of a variety of types of data that can be contained in information packets based upon which classifications can be made. Additionally, while routers are one of the most common devices employed in conjunction with the internet that perform packet classification operations, a variety of other devices also are capable of making classifications such, as for example, firewalls or intrusion detection devices.

[0006] In general, the process of packet classification involves applying one or more rules to information contained in a packet, and then taking one or more actions (and/or refraining from one or more actions) based upon one or more determinations made as a result of applying the rule(s). Given the huge amount of information that is continually sent over the internet, packet classification desirably is performed in a highly-accurate, very rapid manner in order to be effective. Although conventional devices for packet classification such as ternary content addressable memories (TCAMs) are capable of operating in an accurate, rapid manner, these devices have some limitations and disadvantages.

[0007] In particular with respect to TCAMs, these devices employ hardwired circuitry that is specifically designed to implement a particular set of rules that are appropriate for a desired packet classification process. Such TCAMs are typically very fast insofar as the hardwired circuitry is capable of performing numerous determinations simultaneously in parallel fashion. Yet constructing TCAMs so that the devices are capable of handling the application of many (e.g., tens of thousands of) rules can be expensive, indeed, so expensive that the cost of a TCAM often will be 20% to 30% of the cost of a line card on which the TCAM is implemented.

[0008] Further, while packet classification can also be performed by devices other than TCAMs that are less expensive than TCAMs, these other devices tend to be slow. More particularly, while packet classification can be performed by way of software programming in combination with relatively inexpensive hardware components such as static random access memories (SRAMs) or dynamic random access memories (DRAMs), packet classification devices of this

type tend to operate fairly slowly due to the fact that the application of the many patent classification rules tends to be performed in a sequential manner.

[0009] Notwithstanding the limitations associated with the conventional use of TCAMs and other packet classification devices, the continual increase in the volume of information being transmitted via the internet continues to drive a need for patent classification devices that are capable of achieving higher levels of performance and yet are not excessively expensive.

BRIEF SUMMARY OF THE INVENTION

[0010] The present inventors have recognized that limitations affecting conventional packet classification devices, particularly limitations relating to their cost or speed of operation, are the result of the high number and/or complexity of rules or instructions implemented in those packet classification devices. The present inventors have further recognized that many given rules/instructions or sets of rules/instructions, for example, rules that employ ranges, can be expressed in a reduced or simplified form. The present inventors have also discovered multiple methodologies that are particularly applicable to expressing such rules/instructions in a reduced or simplified form including, for example, trimming, expansion, addition and/or merger operations. The present inventors have additionally recognized that a given packet classification device could be enhanced in its design and operation if it employed such a reduced or simplified set of rules/instructions instead of an original set of rules/instructions that had not been reduced/simplified.

[0011] More particularly, in at least some embodiments, the present invention relates to a method of constructing a packet classification device. The method includes providing at least one initial rule defining a first relationship between at least one packet condition and at least one operation, where the at least one initial rule is capable of being expressed as a plurality of initial instructions. The method additionally includes performing a conversion in which at least one modified rule is developed as a substitute for the at least one initial rule, where the at least one modified rule is capable of being expressed as at least one modified instruction, where the at least one modified instruction is of reduced complexity in comparison with the plurality of initial instructions, and where the at least one modified rule defines a second relationship between the at least one packet condition and the at least one operation, the second relationship being

substantially the same as the first relationship. The method additionally includes implementing the at least one modified instruction onto the packet classification device.

[0012] Additionally, in at least some embodiments, the present invention relates to a packet classification device that is constructed by way of the above-described method.

[0013] Further, in at least some embodiments, the present invention relates to a method of converting an initial rule set into a modified rule set, the modified rule set being at least one of shorter than and less complex than the initial rule set. The method includes performing a process that includes at least one of a trimming operation, an expansion operation, a rule addition operation, and a rule merging operation

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] FIG. 1 is a schematic diagram of the internet and devices that in at least one embodiment are coupled thereto and interact therewith including packet classification devices, for example, a router and a firewall;

[0015] FIGS. 2-5 are flow charts showing exemplary steps of operation of four procedures that can be performed to simplify rules that are for implementation in a packet classification device, in accordance with at least some embodiments of the present invention;

[0016] FIG. 6 is a flow chart showing how, in at least some embodiments of the present invention, the four procedures of FIGS. 2-5 would be ordered relative to one another to achieve enhanced simplification of rules;

[0017] FIG. 7 is a flow chart showing an exemplary process for constructing a packet classification device through the use of one or more rule-simplification procedures such as those shown in FIGS. 2-6; and

[0018] FIG. 8 is a graph showing exemplary data indicative of how the implementation of the procedures shown in FIGS. 2-7 can achieve simplification of rules.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0019] Referring to FIG. 1, a schematic is provided illustrating exemplary devices 2, 4, 6 and 8 that are all respectively in communication with the internet 12 and potentially thereby in communication with one another by way of the internet. The internet 12 in the present embodiment is intended to be representative of the interconnected system of networks that

connects computers and other devices around the world via the TCP/IP protocols. However, the present invention is also intended to be applicable to other networks and/or systems of networks including, for example, other forms of the internet that use other types of communication protocols, intranet networks as are commonly employed within various institutions and businesses, and a variety of other networks and systems of networks.

[0020] The devices 2, 4, 6 and 8 are generally intended to be representative of a wide variety of computerized devices, communication devices (including wireless communication devices) and other devices that are capable of communicating with the internet 12 and potentially with one another as well by way of the internet 12 (or by way of such other networks as are employed in other alternate embodiments). In particular, the devices 2, 4, 6 and 8 can be personal computers or mainframe computers, work stations, personal digital assistants (PDAs), or a variety of other devices (e.g., cellular telephones) that are in communication with the internet. The devices 2, 4, 6 and 8 could also be servers or clients, including web servers or web clients capable of interaction in accordance with the worldwide web (WWW). The device 8, which is coupled to the internet by way of an additional device 10 described in further detail below, is intended to indicate that the devices 2, 4, 6 and 8 can be indirectly or directly coupled to the internet 12.

[0021] Further as shown, within the internet 12 are typically situated one or more intermediate devices 14, 16 and 18. The intermediate devices 14, 16 and 18 serve to process and/or route/direct information being communicated by way of the internet 12 including information transmitted onto or received from the internet by the devices 2, 4, 6 and 8. Typically, such information is transmitted in the form of packets 20 as shown, which are subportions of larger portions of information that are intended for transmission over the internet 12.

[0022] For example, when a given one of the devices 2-8 is about to send information via the internet 12 to another of the devices, the one device breaks up that information into subportions of information. As illustrated in FIG. 1, the device adds various information to each subportion 22 of information before it is transmitted as one of the packets 20, including both header information in a header 24 that precedes the subportion 22 as well as trailer information in a trailer 26 that follows the subportion 22. Upon being transmitted, the packets 20 are received off of the internet 12 by a receiving device and are subsequently recombined, that is, the header and trailer information 24, 26 of those packets is stripped from the packets and the subportions 22 of

information within those packets are recombined to arrive at the original portion of information that was to be sent from the source device.

[0023] The header and trailer 24 and 26 can include a variety of different types of data. For example, referring to Table 1, the header or trailer 24, 26 can include information specifying a source internet protocol (IP) address to which a packet is intended to be directed, or a destination IP address from which the packet has originated. The information can also include a source port or device, destination port or device, or protocol(s) that are applicable to the packet. A variety of other types of information not shown in Table 1 can also be included, for example, information allowing for the detection of errors that might occur during transmission of the information via the internet 12. The header or trailer 24, 26 of a given packet can also contain other information including, for example, information associated with any of the layers of the ISO OSI 7 Layer model. It should be understood that the present invention is intended to apply generally to a variety of embodiments in which packets or other portions of data are classified, regardless of the particular format that data takes in those packets/portions, and is not limited to packets having headers or trailers, data relating to the IP protocol, or data relating to any particular layer of any particular protocol or model.

[0024] Table 1 also shows exemplary rules and actions that could be applied and taken, respectively, by an exemplary packet classification device. As shown, the exemplary rules include four different rules that classify information packets based upon any of the source IP address, the destination IP address, the source port, the destination port and the protocol that is used, all of which (as discussed above) is information potentially contained within the headers/trailers of packets. If a given one of the rules is satisfied, that is, a given packet meets the various criteria of the rule, then an action associated with that rule occurs. The actions shown in Table 1 include a deny action and a permit action, although the actions that are taken (or other determinations that are made) in response to any given rule can vary from the exemplary rules shown in Table 1. The asterisks in Table 1 are wildcard indicators signifying that any particular value associated with the particular field or position of that asterisk would satisfy the rule.

Rule #	Source IP	Destination IP	Source Port	Destination Port	Protocol	Action
1	*	10.112.*.*	5000-65535	*	UDP	deny
2	32.75.226.*	*	*	1001-2000	ICMP	deny
3	199.36.184	*	49152-65535	*	ICMP	deny
4	*	*	*	*	*	permit

Table 1: Example rules in a packet classifier.

[0025] Depending upon the embodiment and circumstance, a variety of devices can potentially be configured to perform packet classification activities, e.g., configured to process and/or take actions (or refrain from taking actions) based upon the information in the information packets 20, particularly the data contained in the headers and/or footers 24, 26 of the packets as they are being transmitted from a given source to a given destination. Such devices can include routers, such as a router 18 shown in FIG. 1. As shown the router 18 includes a processing device 28 that, upon receiving one of the packets 20, analyzes the header information 24 (and/or the trailer information 26) to make one or more determinations including, for example, a determination as to where the packet should next be transmitted within the internet, or whether the packet should be transmitted to another destination at all. These determinations can be made by applying rules/instructions to or upon relevant information contained within the header 24 and/or trailer 26 such as the types of information described above.

[0026] Routers such as the router 18 are not the only type of device that can perform packet classification operations. For example, in some circumstances, packets being received off of the internet 12 are first processed by a device such as a firewall 10 before being further directed onto their eventual destination (e.g., the device 8). The firewall 10 depending upon the embodiment can be a standalone device as shown, complete with an independent processing unit 29, or can simply be implemented by way of software running on a primary receiving device such as the device 8. In any event, the firewall 10 also can analyze packet information such as that contained in the header 24 and/or trailer 26 of a given packet 20 to make any of a variety of determinations. In addition to firewalls, packet classification devices can encompass various other devices including, for example, intrusion detection devices.

[0027] As with conventional packet classification devices, in at least some embodiments of the present invention the packet classification devices such as the router 18 and firewall 10 discussed above can employ one or more TCAMs to perform packet classification operations. The

TCAMs are preferred in at least some embodiments insofar as they are hardware-based (e.g., hardwired) devices that are capable of rapidly applying rules in a parallel manner. The TCAMs can be considered to be the processing devices (or to be encompassed within the processing devices) of the packet classification devices, such as the processing devices 28, 29 mentioned above. In still other embodiments, the packet classification devices can employ one or more memory devices such as SRAMs or DRAMs in combination with a processor (e.g., a microprocessor) and software programming as the processing devices 28, 29 in order to perform packet classification. Although the use of such software-based devices can be less expensive than the use of TCAMs, the use of such software-based devices also will tend to be slower in operation than embodiments involving TCAMs.

[0028] In contrast with conventional embodiments of devices that perform packet classification, in accordance with embodiments of the present invention, devices such as the router 18 and/or the firewall 10 that perform packet classification employ one or more processing device(s) (e.g., TCAM(s)) that are configured to perform packet classification in an improved manner such that the processing device(s) can be one or more of less costly to construct, physically smaller and/or faster in performing packet classification. More particularly, to the extent that the processing devices that are employed are TCAMs, implementation in accordance with embodiments of the present invention can reduce the size and cost of the hardware that is required. Alternatively, if the processing devices that are employed are software-based devices as described above, implementation in accordance with embodiments of the present invention can be achieved through the use of software programs that can be executed at a more rapid pace than in conventional software-based devices.

[0029] In order to achieve such less expensive, smaller and/or faster processing devices for performing packet classification, in accordance with embodiments of the present invention, one or more techniques are employed to reduce the number of (or otherwise simplify) the rules that are applied in order to achieve desired packet classification, such that the amount of hardware and/or number of software programming steps that are required in order to implement the packet classification process are reduced. Among these techniques are four techniques shown in FIGS. 2-5 and exemplified by way of Tables 2-5, namely, trimming, as shown in FIG. 2, expanding, as shown in FIG. 3, addition, as shown in FIG. 4, and merging, as shown in FIG. 5.

[0030] In general, each of the techniques shown in FIGS. 2-5 reduces the overall complexity of implementing a given rule or set of rules on a processing device such as a TCAM by reducing a total number of subrules or instructions that are executed by a processing device/TCAM to implement that given rule or set of rules. In the case of a TCAM that executes binary instructions, the techniques serve to reduce the number of binary instructions or TCAM entries that need to be executed in order to achieve a given application of rule(s).

[0031] In at least some embodiments, such as the particular embodiments exemplified by Tables 2-5, reductions in the number of instructions can be achieved when the particular numerical ranges that are considered in determining whether a given rule is met happen to include numerical values, or fail to include numerical values, that are easily represented in binary code (or by some other numerical representation format). Because in most cases a packet classification device applies multiple rules in a prioritized or hierarchal manner, in many cases the ranges of one or more of the rules can be adjusted such that the rule(s) are more easily represented in binary (or other numerical) format and yet the overall prioritized set of rules continues to operate in the same manner as it would have operated if the rule adjustment had not been made.

[0032] Referring to FIG. 2 in particular, a flowchart 30 shows exemplary steps for performing a first technique that can result in a reduced set of instructions, namely, a trimming operation. As shown, upon starting the trimming operation, a next lower priority (or, initially, the highest priority) rule is identified. Then, at a step 34, a core region of the rule is computed. The core region is the part of a rule's range that is not covered by higher rules of the same "color" or lower rules of the same "color", where rules of the same color are rules that, if applicable, result in the same action (or other determination).

[0033] Next, at a step 36, the rule is then trimmed to be the minimum hypercube that encloses its core region, in order to preserve the semantics of the rule set. The "minimum hypercube" is a geometric term intended to indicate the minimum-sized multi-dimensional volume that would encompass the core region of a rule, where the core region could be defined along multiple dimensions or axes. For rules having only a one-dimensional range, the minimum hypercube merely refers to a simple range of values that are within the core region. Fourth, at a step 38, if a range clause originally specifies a prefix, it must be expanded to be the minimum prefix. This is performed in order to avoid unnecessary increases in the number of instructions (e.g., TCAM

entries) that are needed. Finally at a step 39, it is determined whether the last rule (e.g., the lowest priority rule) of the overall rule set has been reviewed or not. If not, then the next rule is obtained at step 32. However, if the last rule has been evaluated, then the trimming operation is complete.

[0034] Turning to Table 2, an exemplary trimming operation is shown. In particular as shown, prior to the trimming operation, a desired packet classification operation might include first and second rules having two ranges, 96-127 and 100-255, respectively. In order to implement these rules, however, a packet classification device would require five different instructions/TCAM entries, namely, a first entry corresponding to the first rule and four entries corresponding to the second rule. Trimming in this circumstance is possible with respect to the second rule because a portion of the range covered by the second, lower priority, rule is not encompassed within the range of the first, higher priority, rule.

[0035] More particularly, while the range 100-127 of the second rule is encompassed both by the lower priority rule and the higher priority rule, the range of 128 to 255 is a "core region" of the second rule. Further, the range 128 to 255 is more easily representable in binary than the range of 100 to 255, with a minimum prefix of only "1" followed by seven wild card indicators (e.g., "1xxxxxxx"). Thus, by trimming the range of the second, lower priority rule from the original 100-255 to only 128-255, the number of instructions/TCAM entries performed by the processing device can be reduced from the original five instructions to only two instructions.

	Before trimming		After trimming	
	Rule	TCAM entries	Rule	TCAM entries
r_1:	x ∈ [96, 127] → *deny*	011xxxxx→*deny*	*x* ∈ [96, 127] → *deny*	011xxxxx→*deny*
r_2:	x ∈ [100, 255] → *permit*	011001xx→*permit*		1xxxxxxx-*permit*
		01101xxx→*permit*	*x* ∈ [128, 255] → *permit*	
		0111xxxx→*permit*		
		1xxxxxxx→*permit*		

Table 2: Packet classifiers and their TCAM representations before/after trimming.

[0036] Turning to FIG. 3, a flowchart 40 shows exemplary steps for performing a second technique that can result in a reduced set of instructions, namely, an expansion operation. In particular, the expansion process begins by picking a range clause that can be expanded, at a step 42. Proper performance of the expansion process does not depend upon the rules of any prioritized set of rules being considered in the order of priority, albeit proper performance should

involve consideration of all of the rules of such a set. Next, at a step 44, it is determined whether an expansion of a range of the rule under consideration is allowed in view of the higher priority rule or rules that exist in relation to the rule under consideration. Typically, an expansion is allowed if the range of a higher priority rule overlaps (or at least begins at an endpoint of) the rule under consideration and further extends beyond the range of the rule under consideration.

[0037] If the expansion is not allowed at the step 44, then the procedures proceeds to a step 48, at which it is determined whether any other range of any other rule can be expanded. If another range can be expanded, then the process returns to step 42 and, if not, then the expansion process is completed. If at the step 44 it is determined that the expansion is allowed, then a minimum expansion of the range of the rule under consideration is performed, at a step 46. A minimum expansion of the range (as opposed to something more) typically will result in the largest decrease in the number of instructions/TCAM entries that are needed to implement the desired rule. Subsequent to the expansion at step 46, the process again proceeds to step 48.

[0038] Table 3 shows an exemplary expansion that could be performed with respect to a pair of rules. In this example, the rule set includes a first rule having a range of 32-79, and a second rule having a range of 72-255. As shown, to perform these first and second rules, six instructions/TCAM entries are required. However, by expanding the range of the second, lower priority, rule, the number of TCAM entries can be reduced. In particular, the lower boundary of the range 72-255 can be lowered from the number 72 to a lower value since the range of the first, higher priority rule encompasses not only the number 72 but all of the values from 32 through 72. Further as shown, in particular the lower range of the second rule can be reduced from 72 to 64. The range 64-255 is much easier to represent in terms of the TCAM entries required in comparison with the original range 72-255, since only two TCAM entries (as opposed to four entries) are required to represent the second rule after its range has been expanded from 72-255 to 64-255.

	Before expanding		After expanding	
	Rule	TCAM entries	Rule	TCAM entries
R_1:	x ∈ [32, 79] → *deny*	001xxxxx→*deny* 0100xxxx→*deny*	x ∈ [32, 79] → *deny*	001xxxxx→*deny* 0100xxxx→*deny*
R_2:	x ∈ [72, 255] → *permit*	01001xxx→*permit* 0101xxxx→*permit* 011xxxxx→*permit* 1xxxxxxx→*permit*	x ∈ [64, 255] → *permit*	01xxxxxx→*permit* 1xxxxxxx→*permit*

Table 3: Packet classifiers and their TCAM representations before/after expanding.

[0039] Turning to FIG. 4, a third technique for improving the simplicity of rules implemented on a packet classification device involves adding one or more new rules to an existing set of rules. This technique can, and typically is, performed in conjunction with the expansion technique described with reference to FIG. 3. Thus, a flowchart 50 of FIG. 4 includes each of the steps 42-48 shown in FIG. 3 and also includes several additional steps associated with the possible addition of one or more rules. However, in contrast to the flowchart 40 (in which when expansion is not allowed at step 44 the process proceeds to step 48), in the flowchart 50 if expansion is not allowed at the step 44, then the process proceeds to a step 52. At step 52, a new rule is added (or possibly rules are added) before the original rule, and the original rule is expanded in terms of its range. Next, at a step 54, it is determined whether the semantics of the rule set are preserved as a result of the operation in step 52. The "semantics" of a given rule set are preserved by a modified rule set if application of the modified rule set upon all possible different input packet information (e.g., input values) would result in the same actions being taken as if the given rule set had been applied.

[0040] If in step 54 it is determined that the semantics are not preserved, then the operation performed in step 52 is rolled back or reversed, at a step 56, after which the process returns to the step 48. However, if it is determined at step 54 that the semantics of the rule set are preserved, then the process proceeds to a step 58, at which it is determined whether the overall number of instructions (e.g., TCAM entries) needed to execute the original rule have been reduced by adding the additional rule (or rules). That is, it is determined whether the overall number of instructions needed to execute the modified original rule (with expanded range) plus the new rule (or rules) is less than the number of instructions needed to execute the original rule. If not, then the process again rolls back/reverses the modification to the rule set at the step 56. However, if the number of instructions is reduced, then the step 58 proceeds immediately to the step 48 and the added, expanded rule set is utilized.

[0041] Table 4 exemplifies an addition operation that could be performed in an effort to reduce the complexity of rules in a packet classification device. As shown, an initial set of rules could include a first rule having a range 64-119 and a second, lower priority rule having a range of 0-255. Implementation of these rules in a packet classification device would require four instructions/TCAM entries as shown. Further as shown, however, representing the first rule having the range of 64-119 requires in particular three instructions. The representation of this

rule could actually be simplified by expanding the range of that rule, from 64-119 to 64-127, and then further adding an additional, higher priority rule having a range of 120-127. The addition of this new, highest priority rule as shown actually decreases the overall number of instructions that are necessary to represent the overall rule set by one instruction. It should be noted that the newly-added rule specifies an action ("permit") that is opposite that of the original rule as well as the modified version of the original rule with the expanded range ("deny"), so that the semantics of the overall rule set are preserved.

	Before adding r_0		After adding r_0	
	Rule	TCAM entries	Rule	TCAM entries
r_0:			$x \in [120, 127] \rightarrow permit$	01111xxx→*permit*
R_1:	$x \in [64, 119] \rightarrow deny$	010xxxxx→*deny* 0110xxxx→*deny* 01110xxx→*deny*	$x \in [64, 127] \rightarrow deny$	01xxxxxx→*deny*
R_2:	$x \in [0, 255] \rightarrow permit$	xxxxxxxx→*permit*	$x \in [0, 255] \rightarrow permit$	xxxxxxxx→*permit*

Table 4: Packet classifiers and their TCAM representations before/after adding r_0.

[0042] Referring to FIG. 5, an additional, fourth technique for improving the simplicity of rules implemented on a packet classification device involves a merging operation. As shown in FIG. 5, this merging operation can be implemented in conjunction with the expansion and addition operations described with reference to FIGS. 3 and 4. In particular, FIG. 5 shows a flowchart 60 having the steps 42, 44, 46 and 48 of FIGS. 3 and 4 relating to expansion, as well as steps 52, 54, 56 and 58 of FIG. 4 relating to the addition of rules. In contrast to the flowchart 50 of FIG. 4, the flowchart 60 does not proceed directly to the rollback/reversal step 56 if at the step 58 it is determined that the number of instructions/TCAM entries of the rule is not reduced.

[0043] Rather, if this is determined, then a redundancy is removed at a step 62. This entails combining two rules that are complementary in terms of their ranges (e.g., their ranges if combined result in a larger range) into a single rule having a combination range extending to encompass the largest range determined by the outer bounds of the individual rules' ranges. Subsequently, it is determined whether the number of instructions/TCAM entries of the overall rule set have been reduced due to the operation performed in step 62, at a step 64. If it is determined at step 64 that the number of instructions/TCAM entries of the rule set has been successfully reduced, then the rule modification is made and the process returns to the step 48. However, if the number of instructions/TCAM entries of the rule set has not been successfully

reduced, then the procedure again returns to the step 56, in which case the redundancy operation of step 62 is rolled back/reversed and then the procedure again returns to the step 48.

[0044] Referring additionally to Table 5, an exemplary merging operation is shown in more detail. As shown, prior to a merging operation, an overall rule set could include four rules for implementation, namely, a first rule having a range of 96 to 111, a second rule having a range of 64 through 95, a third rule having a range of 100-127, and a fourth rule having a range of 0-255, where the first and fourth rules specified an action of "permit" and the second and third rules specified an action of "deny". As shown, implementation of the first through third rules would require five instructions/TCAM entries.

[0045] However, upon further analysis, it is apparent that the ranges of the second and third rules can be combined to arrive at an overall range of 64-127. This is possible because the first rule extends to encompass a gap existing between the upper extent of the second rule (95) through the lower extent of the third rule (100), and because application of the first rule results in an action that is opposite to the action required by each of the second and third rules. That is, even though the second and third rules are merged to apply to a larger range of values 64-127 than the original second and third rules, the inclusion of the intermediate range 96-99 does not impact the semantics of the packet classification device since the action taken with respect to those particular values is fully determined by the highest priority first rule.

	Before merging		After merging	
	Rule	TCAM entries	Rule	TCAM entries
R_1:	x ∈ [96 111] →*permit*	0110xxxx→*permit*	*x* ∈ [96, 111] →*permit*	0110xxxx→*permit*
r_2 (r'_2):	x ∈ [64, 95] →*deny*	010xxxxx→*deny*	*x* ∈ [64, 127] →*deny*	01xxxxxx→*deny*
R_3:	x ∈ [100, 127] →*deny*	011001xx→*deny*		
		01101xxx→*deny*		
		0111xxxx→*deny*		
R_4:	x ∈ [0, 255] →*permit*	xxxxxxxx→*permit*	*x* ∈ [0, 255] →*permit*	xxxxxxxx→*permit*

Table 5: Packet classifiers and their TCAM representations before/after merging.

[0046] It should also be understood that, with respect to the above-described procedure involving the removal of a redundancy in rules (or reduction or merger of rules), as well as each of the previously-described procedures involving trimming, expansion or addition, it is necessary that the semantics of the overall rule set (and of the overall set of instructions/TCAM entries used to implement that rule set) be preserved notwithstanding the change(s) to the rules and

instructions/TCAM entries. Thus, in making changes to the rules in the manner described above (for example, in removing a redundancy in the rules), it is presumed that such changes will only be made if the changes do not alter the overall operation of the overall collective rule set in terms of taking actions (or making determinations) in response to received packet information.

[0047] Turning to FIG. 6, the four techniques described above with reference to FIGS. 2-5 and Tables 2-5 can be, in at least some embodiments, combined into an overall procedure represented by a flowchart 70. The flowchart 70 in particular is intended to be representative of a preferred (albeit not the only) order of applying the various techniques described above. This order is preferred insofar as, by performing the techniques in this order, a maximum potential reduction in the number of instruction/TCAM entries that are employed to implement a particular packet classification rule set is realized. As shown in FIG. 6, the trimming procedure represented by the flowchart 30 is the first step of the overall process represented by the flow chart 70. Upon completion of the trimming operation at the step 30, then a multi-step routine 66 is performed that encompasses each of the remaining expansion, addition and merging techniques. Once the routine 66 is performed, then the overall process concludes at a step 68, in which any remaining redundancies in the rule set are eliminated.

[0048] Performance of the routine 66 begins by obtaining a particular rule, which is either the first rule of a set of rules or the next rule subsequent to the processing of an earlier rule, at a step 72. Upon obtaining a particular rule, then it is determined at a step 74 whether expansion of the rule set as described with reference to FIG. 3 will be of use. If so, then the routine proceeds to step 76, at which an expansion procedure is performed, after which time the process returns back to step 74. If, however, it is determined at step 74 that expansion of a rule will not help, then it is further determined at a step 78 whether adding a rule will help reduce the overall size of the rule set. If so, then the subroutine proceeds to a step 80, in which a rule addition procedure such as that described with reference to FIG. 4 is performed. After the additional procedure is performed, then the subroutine returns to the step 74.

[0049] If it is determined at step 78 that adding a rule will not help, then the routine 66 proceeds to a step 82, at which it is determined if a merger with another rule will help to simplify the overall rule set. If a merger of rules will help, then the subroutine procedure 66 proceeds to a step 84, at which a merger is performed, and then subsequently the subroutine returns to the step 74. However, if it is determined at step 82 that a merger of rules will not help, then the routine

proceeds to a step 86, at which it is determined whether the present rule that is the subject of consideration is the last rule that could be analyzed. If it is not the last rule, then the routine 66 returns to the step 72, at which another rule is obtained. However if it is the last rule, then the subroutine proceeds from the step 86 to the step 68.

[0050] Although FIG. 6 shows one manner of performing all of the techniques shown in FIGS. 2-5 in a comprehensive manner, it will be understood that these various techniques can be performed independently and also can be performed in conjunction with one another (and possibly with other techniques) in a variety of other manners. It should be also noted that FIG. 6 is intended to be indicative, on a general level, of a preferred order of performing the four techniques described above with respect to FIGS. 2-5. For that reason, the particular steps shown in FIG. 6 do not necessarily clearly correspond on a one-to-one basis with the particular steps shown in FIGS. 2-5.

[0051] Turning to FIG. 7, processes for reducing/simplifying rule sets such as those described above with respect to FIGS. 2-5 can be implemented in developing, programming, making and/or manufacturing a variety of devices that are intended to perform packet classification including, for example, routers, firewalls, line cards, intrusion detection systems, TCAMs, software-based packet classification devices, etc. As shown in FIG. 7, in order to make/manufacture a given packet classification device, an initial set of rules for implementation in that packet classification device is developed at a step 92. Then, at a step 94, the rules are reduced in terms of their number and/or complexity using one or more of the above described techniques, or similar techniques. Once the reduced set of rules is determined, then at a step 96 the packet classification device is built, configured, programmed, and/or otherwise made/manufactured. Whether the device is assembled as opposed to programmed can depend upon the type of device, for example, whether the device is a TCAM or a software-based device as described above. Finally, once the device is made, then it can be implemented within a network or network application, as shown by a step 98.

[0052] The implementation of techniques such as those described above can result in modified rule sets that are significantly reduced in terms of their overall size/complexity relative to the initial rule sets that might have been proposed. For example, as shown in FIG. 8, the fraction of rule sets that can be reduced by as much as 25% to 50% of their original size through the use of the preferred procedure shown in FIG. 6 is approximately 44.1% of all rule sets, and an

additional 49.7% of rule sets can be compressed by as much as 50% to 75% of their original size through the use of that preferred procedure.

[0053] Because the implementation of techniques such as those described above can result in a significant reduction in the size and/or complexity of rule sets, packet classification devices within which such reduced-size/complexity rule sets are implemented can achieve dramatic reductions in size, complexity and/or cost, as well as possibly increases in operational speed. More particularly, where the rule sets are implemented on TCAMs, which often is preferred insofar as TCAMs can typically perform these rule sets in a more rapid than other types of embodiments, the overall size, complexity and cost of the hardware can be dramatically reduced. Additionally, in software-based devices for performing packet classification, the use of rule sets of reduced-size and/or complexity developed by virtue of techniques such as those described above can result in significantly reduced rule sets, and correspondingly reduced-length software programs, which can be executed more rapidly than would otherwise be the case and possibly be implemented at lower cost.

[0054] It is specifically intended that the present invention not be limited to the embodiments and illustrations contained herein, but include modified forms of those embodiments including portions of the embodiments and combinations of elements of different embodiments as come within the scope of the following claims.

CLAIMS

WE CLAIM:

1. A method of constructing a packet classification device, the method comprising:

providing at least one initial rule defining a first relationship between at least one packet condition and at least one operation, wherein the at least one initial rule is capable of being expressed as a plurality of initial instructions;

performing a conversion in which at least one modified rule is developed as a substitute for the at least one initial rule, wherein the at least one modified rule is capable of being expressed as at least one modified instruction, wherein the at least one modified instruction is of reduced complexity in comparison with the plurality of initial instructions, and wherein the at least one modified rule defines a second relationship between the at least one packet condition and the at least one operation, the second relationship being substantially the same as the first relationship; and

implementing the at least one modified instruction onto the packet classification device.

2. The method of claim 1, wherein the at least one operation includes at least one of a first action and a first determination that one of the first action and a second action should not be taken.

3. The method of claim 2, wherein the first action is a granting of a permission, and the first determination is a determination that the permission should not be granted.

4. The method of claim 1, wherein the conversion includes at least one of a trimming operation, an expansion operation, an addition operation, and a merger operation.

5. The method of claim 1, wherein the conversion includes modifying a numerical range of one initial rule to arrive at one modified rule.

6. The method of claim 1, wherein the conversion includes a trimming operation in which a range of one initial rule is reduced to arrive at one modified rule.

7. The method of claim 1, wherein the conversion includes an expansion operation in which a range of one initial rule is increased to arrive at one modified rule.

8. The method of claim 1, wherein the conversion includes an addition operation and the at least one modified rule is greater in number than the at least one initial rule.

9. The method of claim 8, wherein the at least one modified rule includes a higher priority rule and a lower priority rule, wherein the lower priority rule has a first range that is greater than an initial range of one initial rule and the higher priority rule has a second range that encompasses a difference between the initial range and the first range.

10. The method of claim 1, wherein the conversion includes a merger operation and the at least one modified rule is lesser in number than the at least one initial rule.

11. The method of claim 1, wherein the at least one initial rule includes a higher priority rule and first and second lower priority rules, wherein the higher priority rule has a middle range, wherein the first lower priority rule has a first range that extends from a first value below the middle range into the middle range, wherein the second lower priority rule has a second range that extends from within the middle range to a second value above the middle range, and wherein the at least one modified rule includes both the higher priority rule and a combination rule having a combination range that extends from the first value to the second value.

12. The method of claim 1, wherein the conversion includes a trimming operation followed by at least one of an expansion operation, an addition operation, and a merger operation.

13. The method of claim 12, wherein the conversion includes each of the expansion operation, the addition operation and the merger operation, and wherein the addition operation occurs subsequent to the expansion operation and the merger operation occurs subsequent to the addition operation.

14. The method of claim 1, wherein the instructions are ternary content addressable memory (TCAM) entries.

15. The method of claim 1, wherein each of the instructions identifies a range as a binary number by including at least one wild card as at least one digit.

16. The method of claim 1, wherein at least one of the following is true:

the at least one packet condition is determined based upon first information contained in at least one of a source address field, a destination address field, a source port identifier field, a destination port identifier field, and a protocol type field; and

the at least one packet condition is determined based upon second information contained in at least one layer of a packet in accordance with the ISO OSI 7 Layer model.

17. The method of claim 1, wherein the at least one packet condition is a value existing within at least one of a header and a trailer of a packet.

18. The method of claim 1, wherein the packet classification device includes at least one of a router, a line card, a firewall, an intrusion detection device, a TCAM, and a software-based device with programming for performing packet classification.

19. The method of claim 1, wherein the at least one modified instruction is implemented by way of hardware into the packet classification device.

20. The method of claim 19, wherein the at least one modified instruction is implemented by way of software into the packet classification device.

21. The method of claim 1, wherein the packet classification device is configured for operation in conjunction with at least one of the internet and an intranet.

22. A packet classification device constructed according to the method of claim 1.

23. A method of converting an initial rule set into a modified rule set, the modified rule set being at least one of shorter than and less complex than the initial rule set, the method comprising:

performing a process including at least one of a trimming operation, an expansion operation, a rule addition operation, and a rule merging operation.

FIG. 1

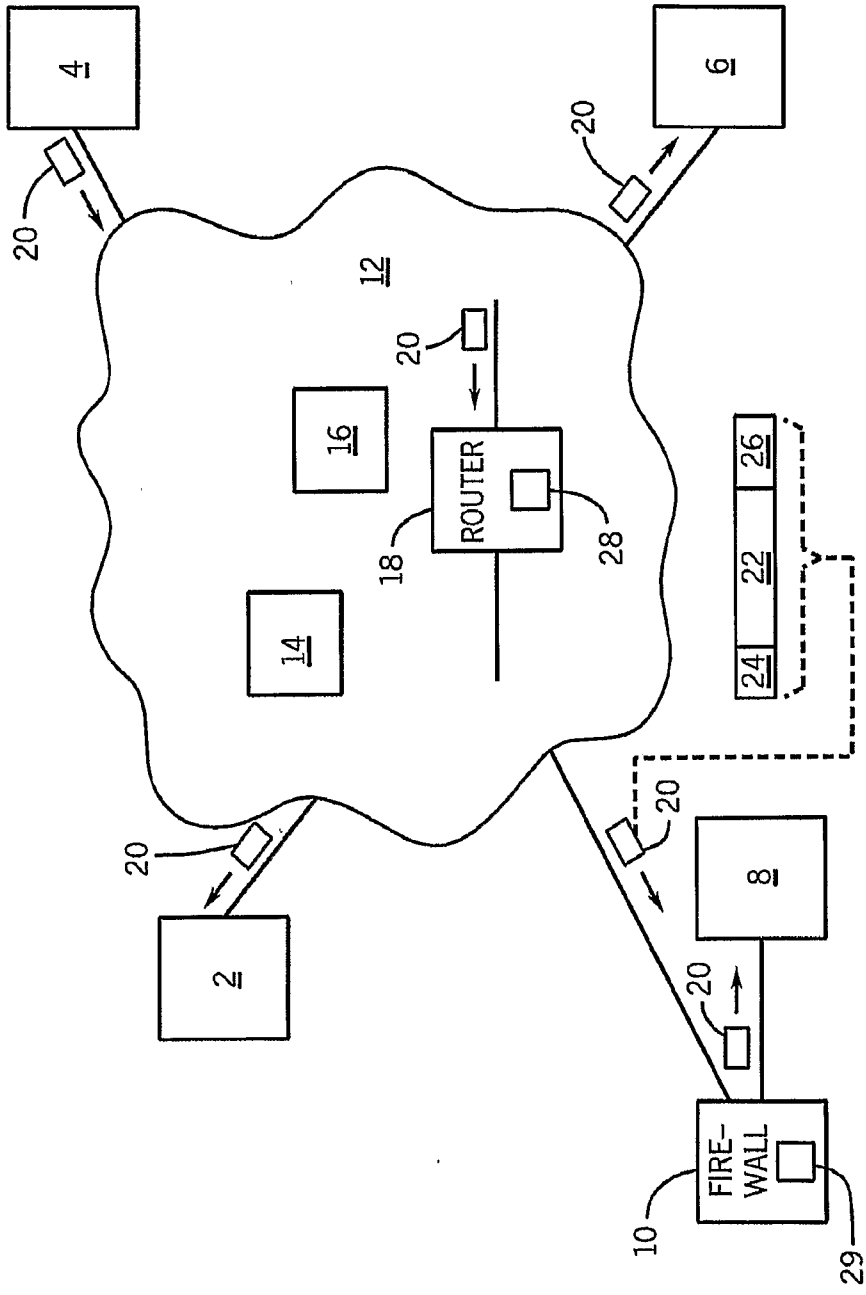
4
6
20
20
12
20
16
ROUTER
18
28
14
24
22
26
20
20
8
2
20
FIRE-
WALL
10
29

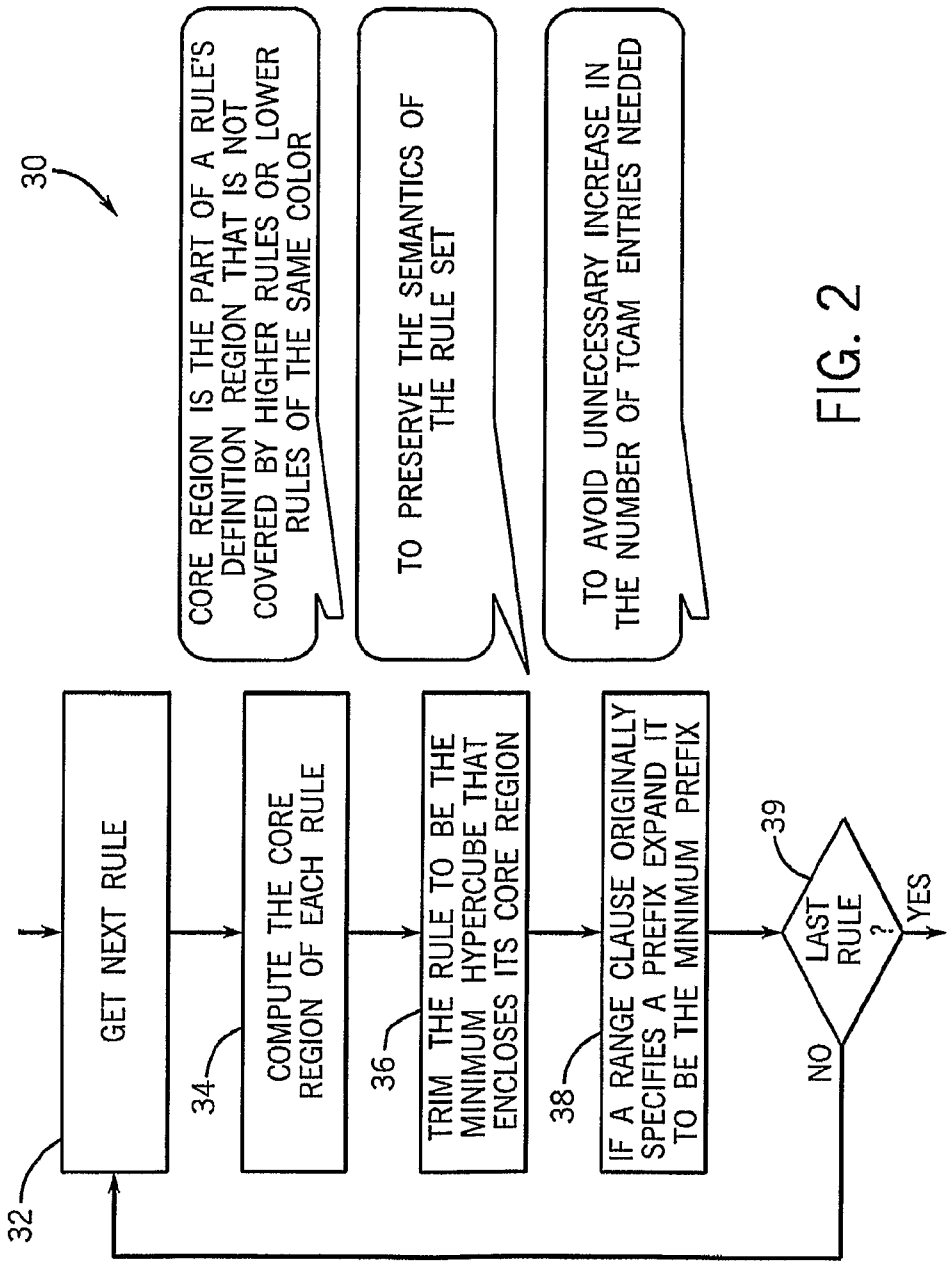
30
CORE REGION IS THE PART OF A RULE'S DEFINITION REGION THAT IS NOT COVERED BY HIGHER RULES OR LOWER RULES OF THE SAME COLOR
TO PRESERVE THE SEMANTICS OF THE RULE SET
TO AVOID UNNECESSARY INCREASE IN THE NUMBER OF TCAM ENTRIES NEEDED
GET NEXT RULE
COMPUTE THE CORE REGION OF EACH RULE
TRIM THE RULE TO BE THE MINIMUM HYPERCUBE THAT ENCLOSES ITS CORE REGION
IF A RANGE CLAUSE ORIGINALLY SPECIFIES A PREFIX EXPAND IT TO BE THE MINIMUM PREFIX
LAST RULE ?
NO
YES
32
34
36
38
39

FIG. 2

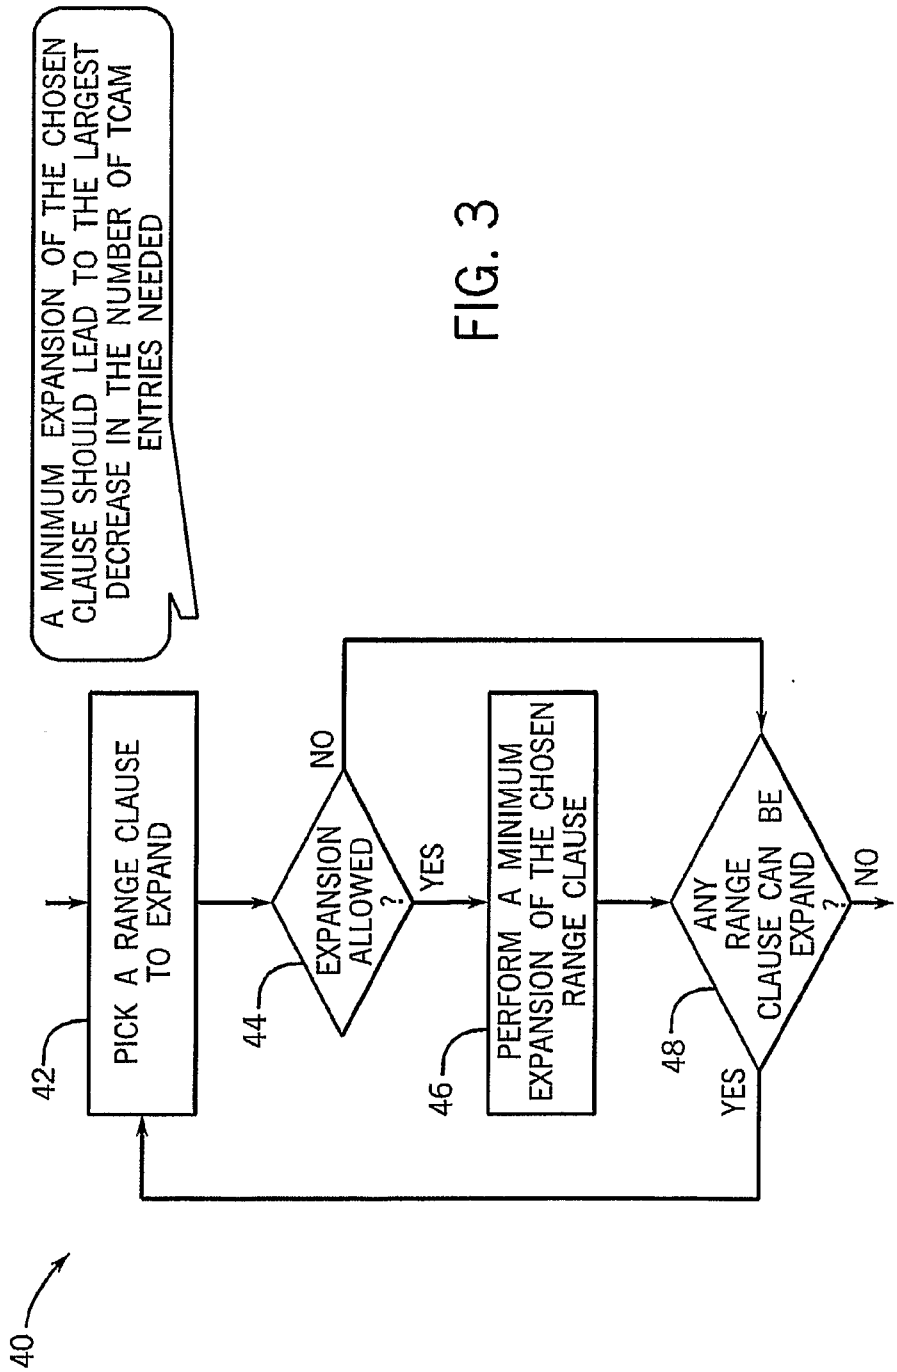
A MINIMUM EXPANSION OF THE CHOSEN CLAUSE SHOULD LEAD TO THE LARGEST DECREASE IN THE NUMBER OF TCAM ENTRIES NEEDED
PICK A RANGE CLAUSE TO EXPAND
42
EXPANSION ALLOWED ?
44
NO
YES
PERFORM A MINIMUM EXPANSION OF THE CHOSEN RANGE CLAUSE
46
ANY RANGE CLAUSE CAN BE EXPAND ?
48
YES
NO
40

FIG. 3

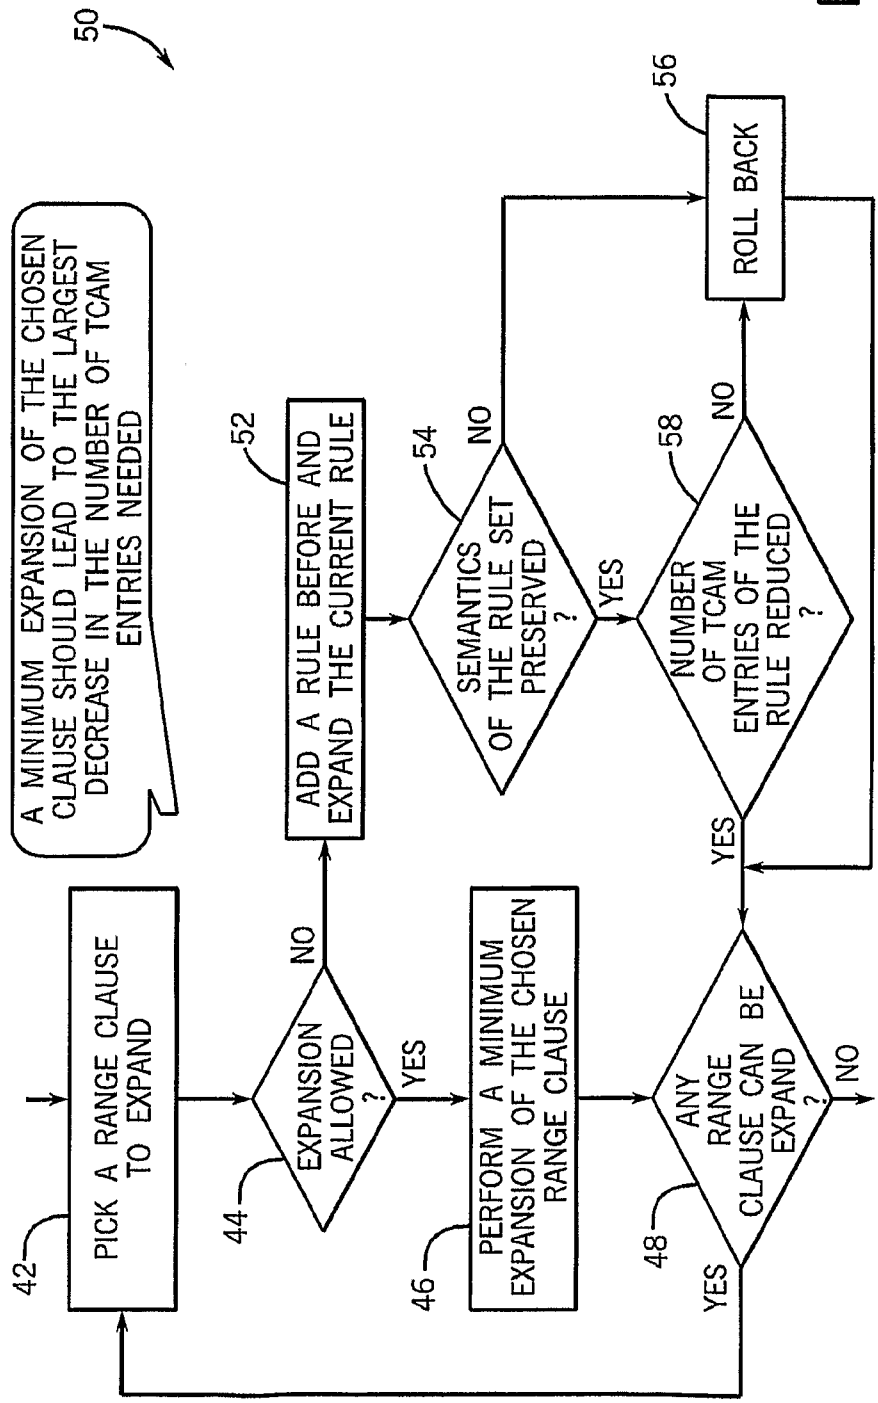
50
A MINIMUM EXPANSION OF THE CHOSEN CLAUSE SHOULD LEAD TO THE LARGEST DECREASE IN THE NUMBER OF TCAM ENTRIES NEEDED
PICK A RANGE CLAUSE TO EXPAND
42
EXPANSION ALLOWED ?
44
NO
YES
ADD A RULE BEFORE AND EXPAND THE CURRENT RULE
52
PERFORM A MINIMUM EXPANSION OF THE CHOSEN RANGE CLAUSE
46
SEMANTICS OF THE RULE SET PRESERVED ?
54
NO
YES
ROLL BACK
56
NUMBER OF TCAM ENTRIES OF THE RULE REDUCED ?
58
NO
YES
ANY RANGE CLAUSE CAN BE EXPAND ?
48
YES
NO

FIG. 4

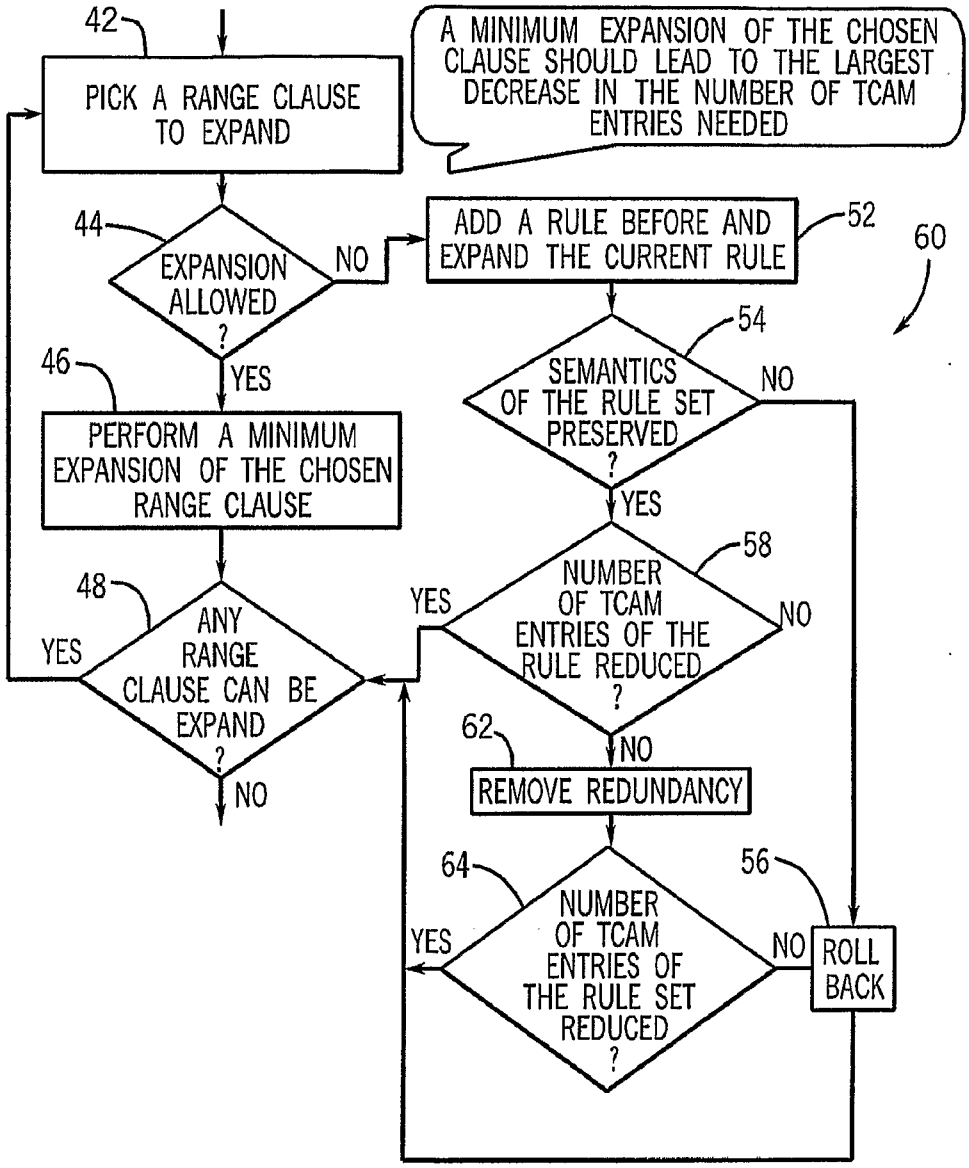
42
PICK A RANGE CLAUSE TO EXPAND
A MINIMUM EXPANSION OF THE CHOSEN CLAUSE SHOULD LEAD TO THE LARGEST DECREASE IN THE NUMBER OF TCAM ENTRIES NEEDED
44
EXPANSION ALLOWED ?
NO
ADD A RULE BEFORE AND EXPAND THE CURRENT RULE
52
60
YES
46
PERFORM A MINIMUM EXPANSION OF THE CHOSEN RANGE CLAUSE
54
SEMANTICS OF THE RULE SET PRESERVED ?
NO
YES
58
NUMBER OF TCAM ENTRIES OF THE RULE REDUCED ?
YES
NO
48
ANY RANGE CLAUSE CAN BE EXPAND ?
YES
NO
62
NO
REMOVE REDUNDANCY
64
NUMBER OF TCAM ENTRIES OF THE RULE SET REDUCED ?
56
YES
NO
ROLL BACK

FIG. 5

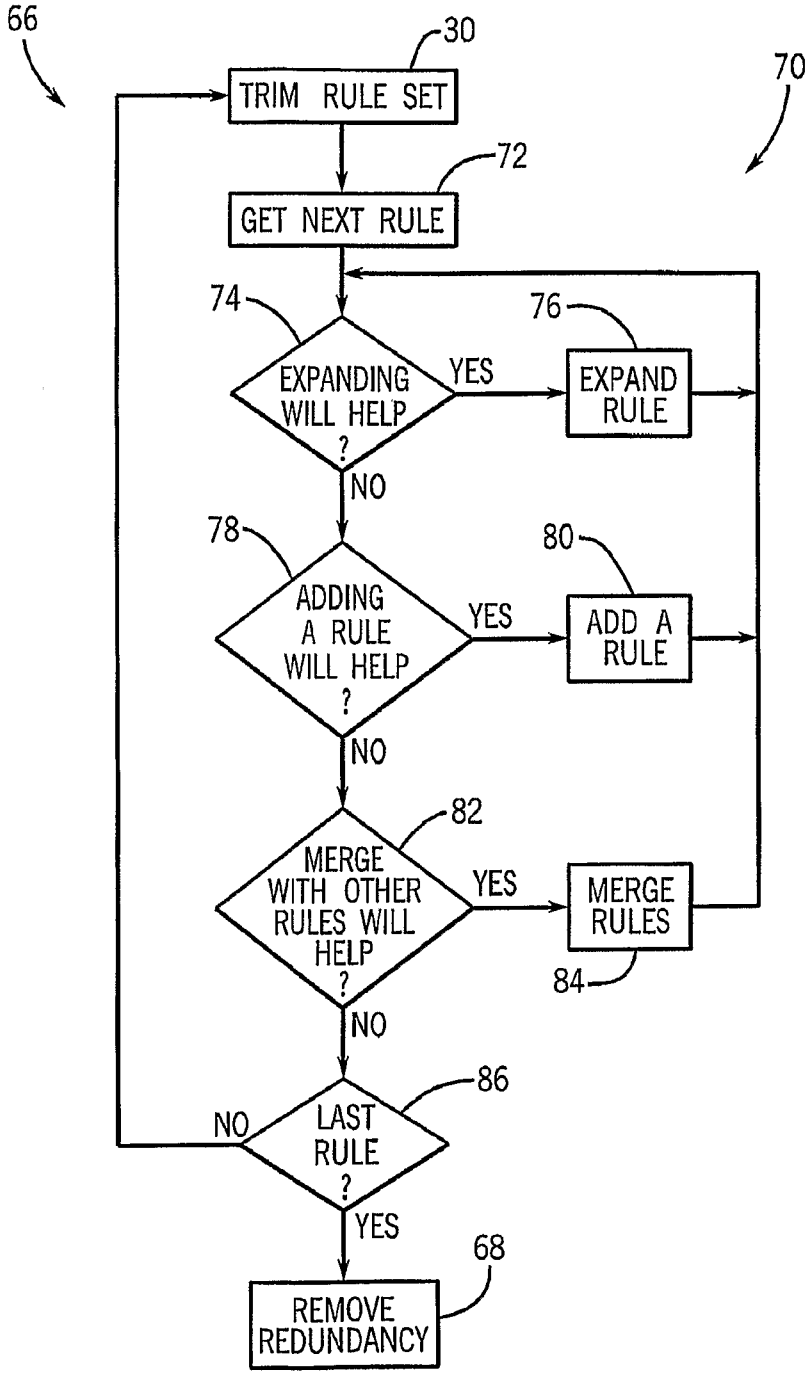
66
30
TRIM RULE SET
70
72
GET NEXT RULE
74
EXPANDING WILL HELP ?
YES
76
EXPAND RULE
NO
78
ADDING A RULE WILL HELP ?
YES
80
ADD A RULE
NO
82
MERGE WITH OTHER RULES WILL HELP ?
YES
MERGE RULES
84
NO
86
NO
LAST RULE ?
YES
68
REMOVE REDUNDANCY

FIG. 6

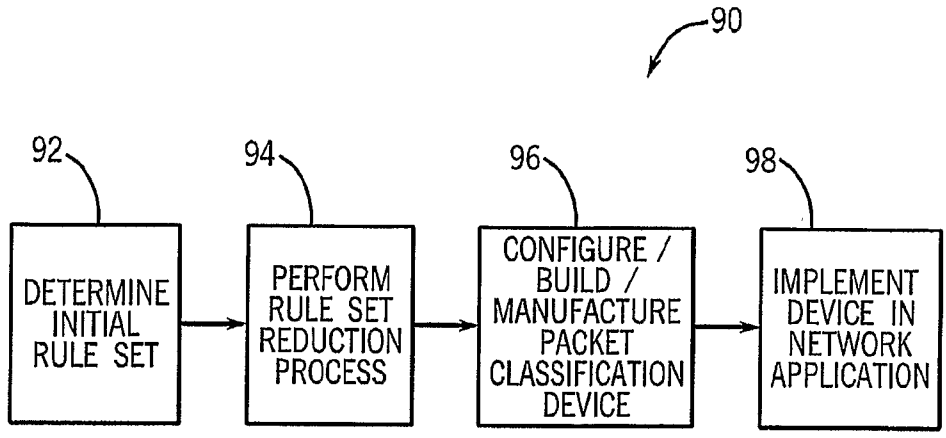
90
92
94
96
98
DETERMINE INITIAL RULE SET
PERFORM RULE SET REDUCTION PROCESS
CONFIGURE / BUILD / MANUFACTURE PACKET CLASSIFICATION DEVICE
IMPLEMENT DEVICE IN NETWORK APPLICATION

FIG. 7

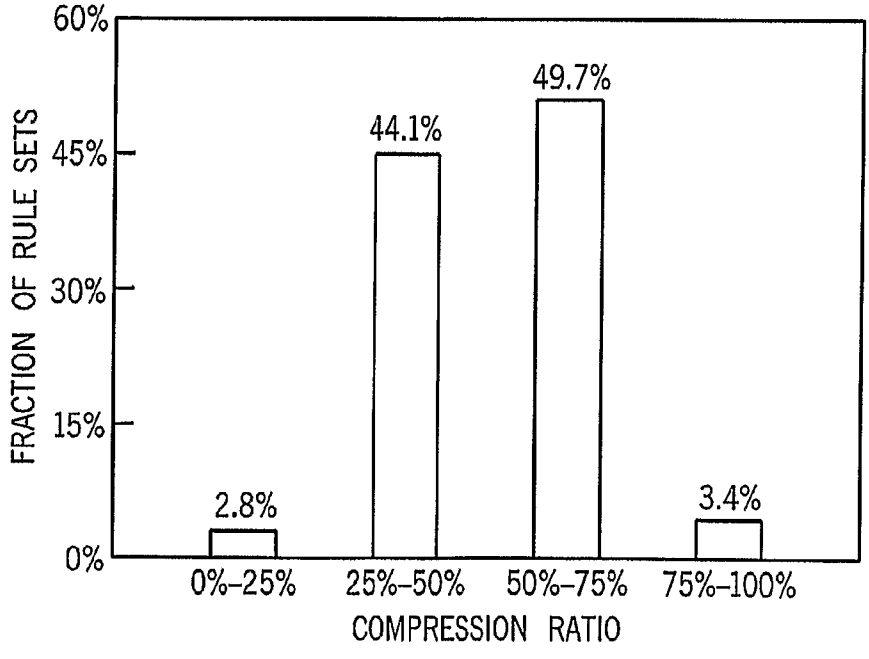
60%
45%
30%
15%
0%
FRACTION OF RULE SETS
2.8%
44.1%
49.7%
3.4%
0%–25%
25%–50%
50%–75%
75%–100%
COMPRESSION RATIO

FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2007/071908

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/56

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

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

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/213235 A1 (MARSHALL JOHN W [US] ET AL) 28 October 2004 (2004-10-28) abstract paragraphs [0008] - [0014] paragraphs [0037] - [0047]; figures 4-7 paragraphs [0053] - [0077]; claims 1-3; figures 10-14	1-23
X	US 2002/009076 A1 (ENGBERSEN TON [CH] ET AL) 24 January 2002 (2002-01-24) paragraphs [0003] - [0009] paragraphs [0023] - [0060]; figures 1-3 paragraphs [0061] - [0081]; figures 6.7A,7B,8 paragraphs [0086] - [0104] paragraphs [0129] - [0140]; claims 1-6	1-23

-/--

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

[X] See patent family annex.

* Special categories of cited documents :

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

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

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
27 November 2007	03/12/2007

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31–70) 340–2040, Tx. 31 651 epo nl, Fax: (+31–70) 340–3016	Gavriliu, Bogdan

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2007/071908

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 03/058450 A (SANCTUM INC [US]) 17 July 2003 (2003-07-17) abstract the whole document	1-23
A	US 2004/190526 A1 (KUMAR ALOK [US] ET AL) 30 September 2004 (2004-09-30) paragraphs [0025] - [0062]; figures 1-10 paragraphs [0097] - [0104]	1-23
A	US 2005/254502 A1 (CHOI LYNN [KR]) 17 November 2005 (2005-11-17) paragraphs [0008] - [0023] paragraphs [0038] - [0061]; figures 1,5A,5B,5C paragraphs [0066] - [0115]; figures 3,4,6	1-23

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2007/071908

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004213235 A1	28-10-2004	CA 2520693 A1	28-10-2004
		EP 1614264 A1	11-01-2006
		WO 2004093399 A1	28-10-2004
US 2002009076 A1	24-01-2002	AT 319249 T	15-03-2006
		CA 2330222 A1	27-07-2001
		DE 60026229 T2	14-12-2006
		JP 3485262 B2	13-01-2004
		JP 2001274837 A	05-10-2001
		KR 20010077983 A	20-08-2001
WO 03058450 A	17-07-2003	AU 2002364055 A1	24-07-2003
US 2004190526 A1	30-09-2004	CN 1534942 A	06-10-2004
		WO 2004095784 A2	04-11-2004
US 2005254502 A1	17-11-2005	NONE	