

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 October 2007 (25.10.2007)

PCT

(10) International Publication Number
WO 2007/120165 A2

- (51) International Patent Classification: **Not classified**
- (21) International Application Number:
PCT/US2006/026069
- (22) International Filing Date: 30 June 2006 (30.06.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/174,373 30 June 2005 (30.06.2005) US
- (71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **BALAKRISHNAN, Santosh** [IN/—]; 1215 E. Carla Vista Dr., Gilbert, AZ 85296 (US). **RAMAKRISHNA, Venkatraman** [IN/US]; 3777 Mentone Ave., #304, Los Angeles, CA 90034 (US).
- (74) Agents: **SHELLER, James, C.** et al.; Blakely Sokoloff Taylor & Zafman LLP, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025 (US).
- (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, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, ME, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, 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, 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:**
— without international search report and to be republished upon receipt of that 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: STATEFUL PACKET CONTENT MATCHING MECHANISMS

(57) Abstract: Methods and apparatus for performing content inspection across packet boundaries using stateful packet content matching mechanisms. Pattern data structures are employed to facilitate string search algorithm machines. Incoming packets are classified to flows, and stored in corresponding flow queues. Flow table entries are used to identify the pattern data structure for a given flow. During content inspection, the algorithm machine updates state information to identify partial or full matches. When a partial match is detected towards the end of an earlier packet state information in the flow table is updated to reflect the state such that content inspection of a following packet begins at that state. As a result, content inspection across packet boundaries is supported. The techniques may further be employed in a packet processing context pipeline stage.



WO 2007/120165 A2

STATEFUL PACKET CONTENT MATCHING MECHANISMS

FIELD OF THE INVENTION

[0001] The field of invention relates generally to network equipment and, more specifically but not exclusively relates to techniques for performing packet content
5 inspection across packet boundaries.

BACKGROUND INFORMATION

[0002] Enterprise networks are threatened with a variety of security issues on a daily basis. These threats can be either internal or external. Internal threats range
10 from break in attempts by disgruntled employees, virus or worms let loose by employees or just employee network access behavior deemed unacceptable per enterprise security policies. External threats range from dedicated attacks on the enterprise network in an attempt to steal intellectual property, denial of service attacks, unauthorized intrusions, viruses and worms *etc.* In all of these cases, a
15 packet firewall is a network administrator's primary form of defense.

[0003] Packet firewalls sit inline with network traffic, intercept incoming packets, and verify each packet against a set of firewall rules to accept, reject and optionally log the packet. In addition to packet filtering, network administrators sometimes also use packet filters to enforce traffic management policies. Such
20 policies are useful in limiting or controlling offensive behavior. Due to the fact that packet firewalls sit inline with and inspect all network traffic, it is important that the firewall should be able to provide sufficient network throughput to keep up with network traffic demands. In today's enterprise networks, firewall rules are typically limited to 2500 Cisco ACL (access control list) rules. These rules are fairly specific,
25 and are designed to allow or reject specific activities or hosts. The ACL rules are limited to this number for performance and manageability reasons. The number of rules directly affects router performance, hence these rules are maintained at a low number. Moreover, large number of rules also makes rules management more error-

prone and difficult to verify or modify. In a large organization, the rules are distributed across firewalls in various sites, and adapted as necessary, increasing chances of an error. For these reasons, the number of rules must be maintained at a manageable level.

5 [0004] Under a typical advanced firewall implementation, two levels of filtering are employed. At the first level, filtering is performed based on applicable ACL rules. In this instance, a highest-priority rule corresponding to the ACL database is identified based on the packet header information. For example, the rule may be identified based on a five-tuple input corresponding to values for the source and
10 destination addresses, source and destination ports, and protocol using well-known classification algorithms. Since many attacks (particularly denial of service attacks) will originate from a known source address using a particular port, packets corresponding to these attacks can be readily identified, and appropriate rules (*e.g.*, drop packet) may be employed to effect a desired firewall policy. This first level of
15 filtering can be implemented at line-rate speeds using modern networking equipment. Under some implementations, dedicated components or separate computers are employed for performing these filtering operations.

[0005] The second level of filtering relates to packet inspection. In this case, the actual packet payload is searched for a particular string or set of strings. For
20 example, the firewall applications may need to search for certain strings indicative of a virus or Internet worm that is present in the packet. In addition, other non-security applications may likewise need to peek into the packet payload, such as for load balancing or billing purposes. These operations, known as “content inspection” or “(deep) packet inspection,” involve inspecting the packet payload for
25 candidate patterns and taking actions based on the presence or absence of these patterns.

[0006] Under some firewall implementations, packet/content inspection is off-loaded to a separate application or sub-system that does not support line-rate speeds. For example, these operations may be performed by a separate computer host or embedded general-purpose processor coupled to or provided by a network device.

5 Since the operations are not performed at line-rate (and thus not restricted to corresponding processing latencies), they can employ larger but slower, less-expensive memory (*e.g.*, DRAM (dynamic random access memory)), and employ conventional string search techniques.

[0007] Network processors (also referred to as network processor units (NPUs))
10 are increasingly being used in a variety of networking equipment due to their cost effectiveness, processing speed, flexibility, and upgradeability. In constructing next-generation networking platforms, it is desirable that robust firewall functionality be added without requiring the addition of specialized firewall components, instead utilizing network processor technology and adding firewall
15 functionality to NPU code in a reusable, scaleable fashion.

[0008] One of the key challenges in content/packet inspection concerns scanning for patterns that span multiple packets. Under a conventional approach using a dedicated firewall resource (*e.g.*, a content inspection application running on a separate component or system), this problem is solved by simply buffering packet
20 streams. The payload content of multiple packets are assembled into a single buffer (thus eliminating the packet boundaries) to form a copy of the data stream being transported, and an appropriate pattern-matching algorithm is employed to scan the data in the buffer for matches. While this technique requires no modifications to conventional pattern-matching algorithms, it imposes significant storage
25 requirements. Furthermore, while this mechanism is relatively simple to implement in systems with small number of flows, it does not scale well for systems that need to support large number of flows simultaneously. Moreover, this technique is not

amenable to use with network processors, and cannot be performed at line-rates when large numbers of simultaneous flows are to be supported.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same becomes better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein like reference numerals refer to like parts throughout the various views unless otherwise specified:

[0010] Figure 1 is a schematic diagram illustrating a technique for processing multiple functions via multiple compute engines using a context pipeline;

[0011] Figure 2 is a schematic diagram illustrating a software block architecture for implementing aspects of the embodiments of the invention described herein;

[0012] Figure 3 is a flowchart illustrating operations and logic performed by one embodiment to determine whether content inspection is to be performed on a packet;

[0013] Figure 4a is shows various parts of a packet that are employed for flow classification;

[0014] Figure 4b shows a set of classification rules and associated flows;

[0015] Figure 5 is illustrates one embodiment of a hash table lookup scheme employing primary and secondary hash tables;

[0016] Figure 6 is a flowchart illustrating operations and logic performed during pre-processing operations to generate pattern data structures;

[0017] Figure 7 is a schematic diagram illustrating data flows and functional blocks employed by a network line card including an NPU having multiple compute engines used to execute instruction threads to perform packet processing operations including content inspection;

[0018] Figure 8a is a flowchart illustrating operations and logic performed during content inspection of a first packet in a flow;

- [0019] Figure 8b is a flowchart illustrating operations and logic performed during content inspection of packets other than a first packet in a flow;
- [0020] Figure 9 is a flowchart illustrating content inspection operations that employ offset and depth search information;
- 5 [0021] Figure 10 shows a goto function, and failure function, and an output function corresponding to an exemplary AC state machine data structure;
- [0022] Figure 11 shows one embodiment of a pattern data structure to be stored in memory;
- [0023] Figure 12 shows one embodiment of an AC machine node memory data
10 structure;
- [0024] Figure 13 shows one embodiment of an AC machine index memory data structure;
- [0025] Figure 14 shows one embodiment of a DAWG data structure to be stored in memory;
- 15 [0026] Figure 15 shows one embodiment of a suffix tree memory data structure;
- [0027] Figure 16 shows one embodiment of a SHIFT table memory data structure;
- [0028] Figure 17 shows one embodiment of a HASH table memory data structure;
- 20 [0029] Figure 18 shows one embodiment of a PREFIX table memory data structure;
- [0030] Figure 19 shows one embodiment of a root node memory data structure;
- [0031] Figure 20 shows one embodiment of a child node memory data structure;
and
- 25 [0032] Figure 21 shows one embodiment of a condition-action rule memory data structure.

DETAILED DESCRIPTION

[0033] Embodiments of methods and apparatus for performing string searches across packet boundaries are described herein. In the following description, numerous specific details are set forth to provide a thorough understanding of
5 embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, *etc.* In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

10 [0034] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily
15 all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0035] In accordance with aspects of the embodiments described herein, mechanisms are disclosed that facilitate content inspection across packet boundaries
20 while maintaining line-rate speeds and providing scalable solutions. These mechanisms employ state-based content inspection algorithms that allow content inspection to stop and restart at packet boundaries, while maintaining information about the results of partial string searches. At the same time, the mechanisms can be implemented with minimal buffering, while supporting line-rate speeds using
25 NPUs. In order to better appreciate and understand the implementation of the mechanisms on modern network processors, a description of typical pipelined-based NPU processing operations is first described. Following this, details of the various embodiments are discussed.

- [0036] Modern network processors, such as Intel's® IXP2xxx family of network processors, employ multiple multi-threaded processing elements (*e.g.*, microengines (MEs)) to facilitate line-rate packet processing operations. In order to process a packet, the network processor (and/or network equipment employing the network processor) needs to extract data from the packet header indicating the destination of the packet, class of service, *etc.*, store the payload data in memory, perform packet classification and queuing operations, determine the next hop for the packet, select an appropriate network port via which to forward the packet, *etc.* These operations are generally referred to as "packet processing" operations.
- 10 [0037] Some of the operations on packets are well-defined, with minimal interface to other functions or strict order implementation. Examples include update-of-packet-state information, such as the current address of packet data in a DRAM buffer for sequential segments of a packet, updating linked-list pointers while enqueueing/dequeueing for transmit, and policing or marking packets of a connection flow. In these cases, the operations can be performed within the predefined cycle-stage budget. In contrast, difficulties may arise in keeping operations on successive packets in strict order and at the same time achieving cycle budget across many stages. A block of code performing this type of functionality is called a context pipe stage.
- 15 [0038] In a context pipeline, different functions are performed on different microengines (MEs) as time progresses, and the packet context is passed between the functions or MEs, as shown in Figure 1. Under the illustrated configuration, z MEs 100_{0-z} are used for packet processing operations, with each ME running n threads. Each ME constitutes a context pipe stage corresponding to a respective function executed by that ME. Cascading two or more context pipe stages constitutes a context pipeline. The name context pipeline is derived from the observation that it is the context that moves through the pipeline.
- 20
- 25

[0039] Under a context pipeline, each thread in an ME is assigned a packet, and each thread performs the same function but on different packets. As packets arrive, they are assigned to the ME threads in strict order. For example, there are eight threads typically assigned in an Intel IXP2800® ME context pipe stage. Each of the
5 eight packets assigned to the eight threads must complete its first pipe stage within the arrival rate of all eight packets. Under the nomenclature illustrated in Figure 1, $ME_{i,j}$, i corresponds to the i th ME number, while j corresponds to the j th thread running on the i th ME.

[0040] A more advanced context pipelining technique employs interleaved
10 phased piping. This technique interleaves multiple packets on the same thread, spaced eight packets apart. An example would be $ME_{0,1}$ completing pipe-stage 0 work on packet 1, while starting pipe-stage 0 work on packet 9. Similarly, $ME_{0,2}$ would be working on packet 2 and 10. In effect, 16 packets would be processed in a pipe stage at one time. Pipe-stage 0 must still advance every 8-packet arrival rates.
15 The advantage of interleaving is that memory latency is covered by a complete 8-packet arrival rate.

[0041] According to aspects of the embodiments now described, novel content inspection schemes and mechanisms are implemented using NPUs that employ multiple multi-threaded processing elements. The schemes are implemented via a
20 context pipelining technique under which content inspection of respective packets assigned to a common flow are performed using different instruction threads executing on one or more compute engines (*e.g.*, microengines). To facilitate the schemes, content inspection state information is maintained in data structures that are accessible to the instruction threads and compute engines. For example, a first
25 thread executing on a first compute engine is used to perform a string search on a first of multiple packets that collectively contain a string that matches a candidate pattern, yet is not contained entirely within a single packet (*i.e.*, respective portions

of the string are contained in respective packets). The first thread maintains state information pertaining to the string search, with corresponding information identifying a partial match when the end of the first packet is encountered. Meanwhile, a second thread executing on either the same or a different compute engine is employed to perform a string search on a second packet that contains the remaining portion of the candidate pattern at its beginning. The second thread employs the state information generated by the first thread to continue the search at the point the first thread left off. Upon identifying a match, appropriate action may be taken based on corresponding predefined rules or actions (*a.k.a.* “action targets”).

10 The entire content inspection process is performed in a manner that is parallel with conventional NPU packet-processing operations in a manner that supports line-rate speeds. Furthermore, the mechanism is easily scaled to support large numbers of simultaneous flows.

[0042] Figure 2 shows a high-level view of various components, data structures and process blocks of a software architecture 200 for implementing aspects of packet-processing operations described herein. The components include a Deep Packet Inspector (DPI) module 202, a flow manager module 204, a rules engine module 206, a logger module 208, and a target handler module 210. The data structures include a flow table 212 and rules tables 214. The process blocks include

20 an input verification block 216, a flow validation block 218, a rule match block 220, a log exception block 222, an exception generation block 224, a policy performance block 226, and a response generation block 228.

[0043] The Deep Packet Inspection module 202 is used to perform various string search operations based on applicable search patterns, as described below in further detail. The flow manager 204 provides an API (Application Program Interface) to access and manage flow table 212. The rules engine 206 provides an API to the rules table 214, implementing a level of translation. For example, a user

25

specified rule will undergo multiple levels of translations before it gets to a form usable by the fast path (*e.g.*, compute engine) processing code. Rather than embodied as a single module (as shown herein for simplicity) DPI module 202 should be considered as a collection of various protocol-specific functions that are implemented via corresponding threads run on NPU compute engines.

[0044] The logger module 208 manages the logging function. This module may either handle logging directly or use services from another component in a different processor (*e.g.*, a general-purpose processor in an NPU or a processor in a separate computer host) to perform the logging function. Typically, the logging function in a firewall might range from simple text-based logging to dumping data into relational databases for extensive report generation. Complexity of this module will vary depending on the functionality provided.

[0045] As a packet enters the system, it is first verified for protocol conformance, checksum *etc.*, by input verification block 216. Once this sanity check is performed, the flow validation block 218 verifies whether the packet belongs to a known (*e.g.*, existing) flow or requires a new flow. If the packet does not belong to any known flow, a new entry is added to flow table 212, and the packet metadata (*e.g.*, header data) is passed to rules matching block 220. This block will iterate through a set of match structures, evaluating conditions for every match. The target actions are encoded in the packet metadata as they are matched. Thus a packet can have more than one target action associated with it. For example, it might be required to log the packet as well as forward it. Once the packet is matched against the rules, it is passed to the target handler block 210 for further processing. This block performs the actions associated with the packet metadata.

[0046] If the incoming packet belongs to a known flow, it is not matched against the access rules again. Instead, it is passed directly to target handler module 210, which performs any necessary processing for the applicable action(s). Actions that

need to be performed on the packets of the flow are stored in or otherwise pointed to by the corresponding flow table entry that exists for the flow.

[0047] With reference to the flowchart of Figure 3, the following operations are performed during processing of each input packet. The process begins at a block 300, wherein flow classification is performed. In general, flow classification may be performed by one of several schemes that are commonly employed for this purpose. Typically, flows will be classified based on an ACL rule database; however, other classification mechanisms may also be employed. The purpose of flow classification is to assign the packet to a corresponding flow.

10 [0048] Traditionally, the rules for classifying a message (*e.g.*, packet) are called *filters* (or *rules* in firewall terminology), and the packet or flow classification problem is to determine the lowest cost matching filter or rule for each incoming message at the network element. The relevant information is contained in K distinct *fields* in each message (packet). For instance, the relevant fields for an IPv4 packet could comprise the Destination Address (32 bits), the Source Address (32 bits), the Protocol Field (8 bits), the Destination Port (16 bits), the Source Port (16 bits), and, optionally, the TCP flags (8 bits). Since the number of flags is limited, the protocol and flags may be combined into one field in some implementations.

[0049] Figure 4a shows an exemplary packet structure used illustrating a 5-tuple flow classification scheme. In general, this scheme can be extended for any K -tuple classification. Each packet will include various fields in its header, with the number and specific fields being dependent on its transport mechanisms (*e.g.*, TCP/IP packets will have different fields than UDP packets). The packet payload data, also commonly referred to packet data or payload follows the packet header data. An optional packet footer may be included at the end of the packet.

[0050] In the illustrated example of Figure 4a, five fields are selected for perform classification on, as depicted by FLDs 1-5. Each packet field contains a

corresponding field value ($V_1 - V_N$). FLDs 1-5 represent header information, such as source and destination addresses, source and destination ports, protocol information, error checking information, and the like. Payload field 400 defines that portion of the packet containing payload data. Information in the footer, as depicted
5 by FLD N, may also be employed for flow classification in some cases.

[0051] Under most classification schemes, only header information is used to perform the classification. However, depending on the particular ACL rules, it may be necessary to include additional information contained in the packet payload and/or footer. For brevity, the following discusses packet classification based on
10 only header information.

[0052] When a packet is received, it is parsed to extract selected field values from its header based on the classification scheme. For the illustrated example, these will be values V_1 - V_5 from FLDs 1-5. The extracted field values are then used to search a rule database to determine the set of classification rules that the packet
15 matches, and hence the associated set of resulting flows.

[0053] As illustrated in Figure 4b, the filter (ACL) database consists of a finite set of filters, $filt_1, filt_2 \dots filt_N$, with each filter associated with a corresponding classification rule and flow. Each filter is a combination of K values, one for each field (or partial field) applicable to the classification scheme. Each field in a filter is
20 allowed three kinds of matches: *exact match*, *prefix match*, or *range match*. In an exact match, the header field of the packet should exactly match the filter field. In a prefix match, the filter field should be a prefix of the header field. In a range match, the header values should lie in the range specified by the filter. Each filter $filt_i$ has an associated directive (*i.e.*, action), which specifies how to handle a packet
25 matching the filter.

[0054] In general, classification rule matching may be performed using one of many known techniques. For example, under one embodiment an aggregated bit

vector scheme is employed, in which each field value is associated with a corresponding bit vector identifying matching rules in the ACL table. (In practice, the bit vectors are stored using an aggregated or hierarchical scheme to reduce the storage size of each bit vector, since many if not most of the bit values will be 0's, indicating no match.) The bit vectors are then logically ANDed, with the result identifying the matching rules. The highest-priority matching rule may then be easily identified. Other schemes, such as but not limited to Recursive Flow Classification (RFC) (Pankaj Gupta and Nick McKeown, *Packet Classification on Multiple Fields*, ACM SIGCOMM 1999) may also be employed

10 [0055] Returning to Figure 3, once the packet has been classified to a flow, a determination is made in a decision block 302 to whether a corresponding flow data structure (*i.e.*, flow table entry) exists in flow table 212. If a flow table entry already exists, the logic proceeds to a block 304 in which the flow structure and corresponding state information is retrieved. If a flow table entry does not already exist, a new flow structure is allocated in the flow table, and its state is initialized, as shown in a block 306. Further details of the flow data structure and state information are discussed below.

[0056] In one embodiment, the flow table entries are maintained in a hash table that uses a double hashing scheme. As shown in Figure 5, the hash table is organized as two tables, a primary table maintained in SRAM (synchronous random access memory) and a secondary table maintained in the DRAM (dynamic random access memory). The primary table is divided into a set of slots, while each slot entry corresponds to a flow entry in the secondary table. The primary table is indexed by two indices: the first value identifies the slot word (a slot is an 8 bit value – 7 bits of signature and 1 flag bit), while the second value (the signature) is stored in the slot and is used to identify the slot. The flow entry in the secondary table is indexed by the slot word index and the slot number with a matching

20
25

signature value. The flow table is maintained as two sets of hash table; if a search for a flow entry in the first table fails, a search of that entry in the second table is performed.

[0057] Depending on the associated flow, content inspection may or may not be required. For example, packets that are sent from a “trusted” source address or sent to a particular destination port may be deemed by the rule database as trustworthy and assigned to a corresponding trusted flow having an associated policy (*i.e.*, action) of “forward.” In contrast, packets sent for an untrusted source address might be assigned to an untrusted flow that has an associated policy of “drop.” In other cases, the flow may have an associated policy of “inspect content,” which requires content inspection to be performed. As depicted by a decision block 308, if content inspection is required, content inspection operations are initiated with the retrieved state information in a block 312. For packets for which content inspection is not required, continuation of packet processing is performed, as depicted by a continuation block 310.

[0058] Overall, the implementation of the search mechanism involves two phases: a build or pre-process phase and a run-time phase. During the build phase, various data structures and tables are built and saved to be later stored in memory during associated initialization operations. During the run-time phase, which is generally ongoing, searching is performed on received packets using the data structures and tables.

[0059] Details of one embodiment of the build phase are shown in Figure 6. The process begins in a block 600, wherein the search criteria are defined. This can be stated in a very general way as the problem of searching for the occurrence of a pattern (string) within another string (packet data in this instance). This problem is identical to the classic string-matching problem, which is stated in formal terms as:

Given a text string T of 'n' characters and a pattern string P of 'm' characters, find the positions of all occurrences of P in T.

[0060] In the context of firewall use, the above problem can be mapped to the problem of finding the presence of a virus or worm signature, protocol command or the name of a known executables in packet data. It will usually be sufficient to just search for the presence of P in T, rather than to find all occurrences (which does not really mitigate the problem, given that finding the first occurrence is usually as “hard” as finding multiple occurrences, in theoretical terms). In most practical firewall and IDS applications, the signature string to search for may not be known beforehand; the best that can be done is to narrow the field down to a set of patterns that need to be search for, in the context of a particular flow. Therefore, a more general problem is to search for occurrences of any, or all, of a set of patterns {P} within a given text T.

[0061] Sometimes, the exact string, or strings, to be matched will not be known beforehand. Signatures can also occur in morphed forms, where the morphing is usually a small change to the original pattern. In such cases, there is a need to handle approximate matching of strings also. Sometimes there may also be a need to match regular expressions instead of known strings. Techniques for all of these are currently being employed in data mining and search applications. In general, these algorithms are somewhat more complicated than the simple string searching one, and have higher memory and cycle requirements.

[0062] In the context of the aforementioned policies (conditions), different pattern sets {P} may apply to different policies. For example, certain flows may only need to be searched for a particular pattern set, while other flows may include searches for broader or entirely different pattern sets. Accordingly, an appropriate pattern set is specified for each policy. In general, and depending on the search

algorithm that is employed, the pattern sets may include exact matches and/or approximate matches (*e.g.*, matches using wildcard characters).

[0063] Next, in a block 602, data structures and related tables (collectively referred to as pattern data structures (DS) are built based on the selected search
5 algorithm(s) to be employed. This aspect is covered in further detail below.

[0064] As depicted by a decision block 604, offset and or depth information may also be associated with the search criteria. Oftentimes, more information is known about a particular pattern than just the string of bytes it is comprised of. Such information may help narrow down the region in the packet (or associated bit
10 stream) where the pattern could occur. As used herein, this information includes the Offset and Depth. The *Offset* value indicates the number of bytes that can be skipped from the beginning of a packet (or bit stream) before starting the search. The *Depth* value indicates the maximum number of bytes within the packet (or bit stream) that need to be searched to have a chance of finding the pattern. The *Offset*
15 and *Depth* values are also referred to as “Tagged” values.

[0065] Although such information is very useful where a single pattern is concerned, it is a little more difficult to extend it to the multi-pattern case where each pattern might have a different offset and depth. The naïve method of using this information would be simply to check these tagged values whenever we reach an
20 accepting state in the applicable state machine, but this would be both inefficient and wrong; if the offset or depth does not match when we reach an accepting state, we could have potentially missed other (valid) patterns by going off on a different path (in the machine). Therefore, we need to incorporate offset and depth checks at every step of the algorithm, so that the correct state transitions are made. One way
25 this is achieved is by incorporating (*e.g.*, adding) appropriate tagged value information in the applicable Pattern data structures, as depicted in a block 606 and explained in further detail below. After the pattern data structures are fully built,

they are saved to a file or the like in a block 608. During subsequent initialization operations, the pattern data structures will be loaded into memory and/or code stores, and corresponding pattern data lookup data will be loaded into memory, as depicted by respective blocks 610 and 612.

5 [0066] To illustrate how and where this data might be stored, attention is directed to the execution environment illustrated in Figure 7. The execution environment pertains to a network line card 700 including an NPU 702 coupled to an SRAM store (SRAM) 704 via an SRAM interface (I/F) 705, and coupled to a DRAM store (DRAM) 706 via a DRAM interface 707. Selected modules are also
10 depicted for NPU 702, including a deep packet inspection module 202, a flow manager module 204, a queue manager module 708, a receive engine 710, and a buffer manager 712. In the manner described above, the operations associated with each of these modules are facilitates by corresponding instruction threads executing on MEs 714. In one embodiment, the instruction threads are initially stored in an
15 instruction store 715 comprising a non-volatile storage device, such as flash memory or a mass storage device or the like.

[0067] As illustrated in Figure 7, various data structures and tables are stored in SRAM 704. These include a flow table 212, a policy data structure table 716, and a pattern data structure table 718. Also, packet metadata 720 is typically stored in
20 SRAM as well. As described above with reference to Figure 5, respective portions of a flow table may be split between SRAM 704 and DRAM 706; for simplicity, all of the flow table 212 data is depicted in SRAM 704 in Figure 7.

[0068] Typically, information that is frequently accessed for packet processing (e.g., metadata 720) will be stored in SRAM, while bulk packet data (either entire
25 packets or packet payloads) will be stored in DRAM, with the latter having higher access latencies but costing significantly less. Accordingly, under a typical

implementation, the memory space available in the DRAM store is much larger than that provided by the SRAM store.

[0069] As shown in the lower left-hand corner of Figure 7, each ME 714 includes a local memory 722, local registers 724, separate SRAM and DRAM read and write buffers 726 (depicted as a single block for convenience), a code store 728, and a compute core (*e.g.*, Arithmetic Logic Unit (ALU)) 730. In general, information may be passed to and from an ME via the SRAM and DRAM write and read buffers, respectively. In addition, in one embodiment a next neighbor buffer (not shown) is provided that enables data to be efficiently passed between ME's that are configured in a chain or cluster.

[0070] As describe below, each pattern data structure will provide information for effectuating a corresponding state machine or table-based machine. The various pattern data structures will typically be stored in pattern data structure table 718, as illustrated in Figure 7. However, there may be instances in which selected pattern data structures are stored in selected code stores that are configured to store both instruction code and data.

[0071] As described above for Figure 6, in addition to storing the pattern data structures, associated pattern data structure lookup data is likewise stored in block 612. In the embodiment illustrated in Figure 7, the lookup data is stored as pointers associated with a corresponding policy in the policy data structure table 716. Other schemes may also be employed. The pattern data structure lookup data is used, in part, to build flow table entries in the manner described below.

[0072] An overview of operations performed during the run-time content inspection phase are illustrated in Figures 8a and 8b, wherein Figure 8a depicts operations performed during processing of a first packet in a flow, and Figure 8b depicts operations performed during processing of any subsequent packet in the flow (as depicted by the i^{th} packet). As shown in a block 800 in each of Figures 8a

and 8b, in response to receiving an input packet initial packet processing operations are performed via execution of one or more threads on one or more MEs 714. To distinguish these threads from other threads used for content inspection, the threads are labeled pre-search threads in Figures 8a and 8b.

5 [0073] In further detail and with reference to Figure 7, as input packets 732 are received at line card 700, they are processed by receive engine 710, which temporarily stores them in receive (Rx) buffers 734 in association with ongoing context pipeline packet processing operations. The packet headers are extracted, and corresponding packet metadata 720 is stored in SRAM 704. The packets are
10 also classified in the manner discussed above, and assigned to existing or new flows. Each flow has an associated queue that is stored in DRAM 706, as depicted by a flow queue 736. The queues are generally managed by queue manager 708 and/or flow manager 204. Upon creation of a flow queue, associated information is added to metadata 720, such that the first (head) and last (tail) packet in the queue
15 can be readily identified for subsequent operations, such as dequeuing.

[0074] Packet streams are received from various network nodes in an asynchronous manner, based on flow policies and other criteria, as well as less predictable network operations. As a result, on a sequential basis packets from different flows may be received in an intermixed manner, as illustrated in Figure 7.
20 For example, each of input packets 732 is labeled with F#-#, wherein the F# identifies the flow, and the -# identifies the sequential packet for a given flow. As will be understood, packets do not contain information specifically identifying the flow to which they are designed, but rather such information is determined during flow classification. However, the packet sequence data is provided in applicable
25 packet headers, such as TCP headers (*e.g.*, TCP packet sequence #). In Figure 7, flow queue 736 contains the first five packets in a Flow #1. As before, packets are

added to associated flow queues after they have been classified during enqueueing operations performed in block 800.

[0075] Following the operations of block 800, operations analogous to those shown in Figure 3 will be performed, resulting in a content search being required.

5 As discussed above, if the packet is assigned to a new flow, a corresponding flow table entry will be created in block 306. In the embodiment illustrated in Figure 7, each flow table entry includes a flow ID, a buffer pointer, a policy pointer, a pattern data structure pointer, state information, and statistics information. (A flow table entry may also include other information that is not shown.) The flow ID is used for
10 flow table entry lookup, and in one embodiment comprises a hash index in accordance with the flow table index scheme shown in Figure 5. In one embodiment, the buffer pointer points to applicable metadata 720, which in turn includes pointers to the head and tail packets in the flow queue associated with the flow table entry. The policy pointer points to a policy data structure in policy data
15 structure table 716. As discussed above, each flow is associated with a policy. During flow classification, the policy is identified, and a corresponding policy pointer is provided to flow manager 204.

[0076] As discussed above, each policy that requires content inspection is also associated with a pattern data structure. Thus, in one embodiment, the pattern data
20 structure pointer may be extracted from the policy data structure.

[0077] The state information is used to store the state of a corresponding algorithm state used by deep packet inspection module 202, and is described in further detail below. The statistics information is used for various purposes, including identifying the current sequence number of the last packet in the flow for
25 which content inspection was completed. In this manner, the next packet to be inspected for a given flow may be identified based on its relative location to the head packet in an associated flow queue.

[0078] As shown in a block 802A of Figure 8a, if the packet is the first packet in the flow, a first thread will be launched on an ME 714 to perform a pattern search of the packet to identify any matches in view of the applicable pattern set {P}. Accordingly, data corresponding to the new flow table entry for this packet is read
5 by the thread, including the state information, which will indicate the state machine is in the initial state. In a block 804, the location of the packet in the associated flow queue is identified using the buffer pointer and metadata, and the packet is copied to local memory 722 of the ME. In a block 806B, the applicable pattern data structure (as identified by the pattern data structure pointer) is retrieved from pattern data
10 structure table 718 and copied into local memory 722. The existing state 738 of the corresponding search algorithm is initialized to the initial state, which indicates no partial matches currently exist.

[0079] The pattern search operations begin at the start of the packet (*e.g.*, start of the packet payload data) and is performed in an iterative manner until the end of
15 the packet is reached. As depicted by start and end loop blocks 808 and 816, the applicable algorithm machine employs a search increment that may comprise one or more characters, depending on the particular search algorithm. In response to each increment, the existing state 738 of the machine is updated, as applicable. In a block 812, a determination is made to whether a full match has been identified by
20 the machine. If so, the associated action for the search condition is performed, as depicted by a continuation block 814, and the thread is terminated.

[0080] In general, the associated action may be performed by the same or another (more likely) thread on the same or another ME. Optionally, the associated action may be performed by the NPU's general-purpose processor (not shown in
25 Figure for clarity). In the case that the matching string identifies the existence of a malicious worm or the like, the associated action will typically be to drop all remaining packets in the flow, and to possibly update the ACL database action for

the source address. Dropping the remaining packets may be accomplished by updating the policy pointer in the associated flow table entry such that the policy for the flow is effectively switched from “inspect content” to “drop.”

[0081] Once the pattern search algorithm reaches the end of the packet (without
5 a match), the state and statistics information in the associated flow table entry is updated in a block 818. At this point, execution of the first thread is terminated.

[0082] Turning to Figure 8b, upon receipt of the subsequent packets classified to the same flow, similar operations are performed, as indicated by like-numbered blocks in Figure 8a and 8b. At a block 802B, an i th thread is launched on the same
10 or another ME 714, wherein i represents the packet sequence number (normalized from the actual sequence number in the header). After processing the first packet, the flow table entry that is read will be an existing entry that includes an existing state that was updated by the previous packet. As before, the flow table entry information, including the statistics, is used to locate the packet in the flow queue,
15 whereupon it is copied to the ME’s local memory 722 in block 804.

[0083] At block 806B, the applicable pattern data structure is retrieved by the ME in a manner similar to block 806A of Figure 8a. However, as a result of processing the previous packet (*e.g.*, the first packet), the state of the machine will be initialized to the existing state updated by that previous packet. If a partial match
20 was present at the end of the previous packet, the existing state will be somewhere in the machine other than the initial state (*e.g.*, 0). If no partial match was present, the existing state will be the initial state.

[0084] The pattern search operations of blocks 808, 810, 812, 814, and 816 are preformed in a similar manner as applied to the first packet, with the only difference
25 being the initial state of the machine. Accordingly, if a partial match existed at the end of the previous packet, the pattern search is began at the point the state machine left off, looking for the second portion (the remaining matching portion) one or

more strings in the pattern set {P} (depending on the existing state). As before, if a full match is identified, the associated action is performed in continuation block 814.

[0085] As discussed above, the search criteria may specify offset and depth values for respective pattern strings. In this case the pattern search processing is augmented to account for this additional search information. As shown in Figure 9, we begin at the start of a packet, with the prior operations of Figure 8a or 8b (as applicable) having already been performed. At a block 900, a portion of the packet is skipped to reach the offset, at which point the search will begin. In further details discussed below, this operation may be implemented on an individual pattern search basis, such that some patterns are offset at different points than others when a multi-pattern search is employed. If the packet is a first packet in a flow, the number of bytes corresponding to the offset will be skipped from the beginning of the packet.

[0086] It is theoretically possible that the offset for one or more search patterns will be greater than the entire length of the first packet. Accordingly, in this case, the entire packet is skipped for those search patterns, and the length of the first packet will be deducted from the offset value, such that the portion of the offset to be skipped for the following packet will consider how much of the bit stream will have been already skipped. For example, consider a case in which the offset for a particular search pattern is greater than the maximum size of the packet payload for a give protocol. For this particular search pattern, the entire first packet will be skipped, and the search will commence at an offset in the second packet that equals the difference between the original offset and the size of the first packet. In one embodiment, this multi-packet offset scheme may be facilitated by updating appropriate statistics information in the associated flow table entry, as discussed below with reference to block 920.

[0087] As depicted by start and end loop blocks 902 and 914, the pattern search involves an iteration over the packet payload data, with the location within the data being advanced by the applicable search increment being employed by the search algorithm. The starting point for each search pattern in pattern set {P} will be the effective offset for that search pattern. In response to moving each search increment, the machine state is updated in a block 904. As indicated by a decision block 906, if a search string match is indicated by the machine, an associated action is performed in a continuation block 908. If no match exists, the logic proceeds to a decision block 910 to determine whether the depth for the applicable search string has been reached. If so, searching for that string is complete, as depicted by an end block 912. If not, the process loops back to start loop block 902 to begin the next iteration.

[0088] Upon reaching the end of the packet, two actions may result, depending on the implementation, as depicted by a decision block 916. For embodiments that only search the first packet in a flow, the search process for the flow is complete, with no match being found. As a result, the policy pointer in the flow table entry is updated in a block 918 to point to a policy that does not perform a content search, such that a search is not performed on the remaining packets in the flow. If a multi-packet search (with regard to the offset aspect) is supported, the statistics information is updated in the flow table entry to account for the size of the skipped packet in the offset to be employed for the following packet, as depicted in block 920.

[0089] In one embodiment, the effective depth employed for a given search pattern can be adjusted to support multi-packet searches. For example, consider a situation in which the offset is less than the size of a first packet, while the depth is greater than the size of the first packet. This situation may be handled by reducing the depth value that is applied to the second packet by the size of the first packet.

This operation may also be performed in block 920. It is further noted that under this scenario, the state information would also be updated in the flow table entry to reflect the state of the state machine at the end of the first packet in the same manner discussed above.

5 [0090] In one embodiment the mechanism for effecting offset and depth operations is provided via data contained in the applicable pattern data structure. Accordingly, a block 922 is included in Figure 9 to illustrate that the pattern data structure is adjusted to account for the offset and/or depth under the multi-packet implementation. This operation will be performed by another thread (launched to
10 process a next packet in the flow) using the statistics information that was updated in block 920.

Multi-pattern Search Algorithms

[0091] The simplest way of searching for the occurrence of multiple patterns is by iterative application of any of a number of well-known single pattern algorithms
15 for each pattern. A number of multi-pattern algorithms have been proposed since the late 1970s. The FSM (Finite State Machine) algorithms have traditionally been the focus of research, but there are other efficient algorithms that don't use FSMs, like the ones proposed by Wu and Manber.

[0092] An FSM algorithm is based on the following considerations. Given a
20 finite alphabet and a pattern, it is possible to construct a finite state automaton in an efficient manner that can process text strings to return occurrences of the pattern. Every state (denoted by a number) in the machine corresponds to the longest prefix of the pattern matched so far. This has the property of needing to scan every character in the text exactly once. The preprocessing mainly involves computing
25 the machine transition function. A run of the algorithm involves feeding the search text into the FSM as input. An occurrence of P is detected whenever an accepting state is reached.

[0093] In general, any type of FSM algorithm may be implemented for the pattern searches employed by the embodiments described herein. In order to better understand how an FSM algorithm operates, a brief description of the well-known Aho-Corasick (AC) algorithm (A. V. Aho, M. Corasick, "Efficient string matching: an aid to bibliographic search", *Communications of the ACM*, vol. 18 (1975) 333—
 5 340) will now be discussed. The following description is taken from the foregoing document.

[0094] Let $K = \{y_1, y_2 \dots, y_k\}$ be a finite set of strings called *keywords* and let x be an arbitrary string which we shall call the *text string*. Our problem is to locate
 10 and identify all substrings of x which are keywords in K . Substrings may overlap with one another. A pattern matching machine for K is a program which takes as input the text string x and produces as output the locations in x at which keywords of K appear as substrings. The pattern matching machine consists of a set of *states*. Each state is represented by a number. The machine processes the text string x by
 15 successively reading the symbols in x , making state transitions and occasionally emitting output. The behavior of the pattern matching machine is dictated by three functions: a goto function g , a failure function f , and an output function *output*.

[0095] Figure 10 shows the functions used by a pattern matching machine for the set of keywords {he, she, his, hers}, wherein the states are 0, 1 . . . , 9. The goto
 20 function g maps a pair consisting of a state and an input symbol into a state or the message *fail*. The directed graph represents the goto function. For example, the edge labeled h from 0 to 1 indicates that $g(0, h) = 1$. The absence of an arrow indicates *fail*. Thus, $g(1, \sigma) = \text{fail}$ for all input symbols σ that are not e or i. All the pattern matching machines have the property that $g(0, \sigma) \neq \text{fail}$ for all input
 25 symbols σ . This property of the goto function on state 0 ensures that one input symbol will be processed by the machine in every machine cycle. The failure function f maps a state into a state. The failure function is consulted whenever the

goto function reports *fail*. Certain states are designated as output states that indicate that a set of keywords has been found. The output function formalizes this concept by associating a set of keywords (possibly empty) with every state.

[0096] An *operating cycle* of a pattern matching machine is defined as follows.

5 Let s be the current state of the machine and a the current symbol of the input string x .

1. If $g(s, a) = s'$, the machine makes a *goto transition*. It enters state s' and the next symbol of x becomes the current input symbol. In addition, if $output(s') \neq empty$, then the machine emits the set $output(s')$ along with the
10 position of the current input symbol. The operating cycle is now complete.

2. If $g(s, a) = fail$, the machine consults the failure function f and is said to make a *failure transition*. If $f(s) = s''$ the machine repeats the cycle with s' as the current state and a as the current input symbol.

[0097] Initially, the current state of the machine is the start state and the first
15 symbol of the text string is the current input symbol. The machine then processes the text string by making one operating cycle on each symbol of the text string.

DAWG-Match Implementation

[0098] In one embodiment, an FSM-based algorithm known as the DAWG-Match algorithm (M. Crochemore, A. Czumaj, L. Gasieniec, T. Lecroq, W.
20 Plandowski, and W. Rytter, "Fast practical multi-pattern matching", *Inf. Process. Lett.*, 71((3-4)):107-113, 1999) is employed. This algorithm combines the ideas of Aho-Corasick with Boyer-Moore to get a better average running time. It requires computation of both an AC state machine as well as a DAWG (Directed Acyclic Word Graph) for the reverse patterns of the given pattern set. While computing the
25 transition functions for the AC state machine, a *shift* table is also computed, which stores the maximal shift at a particular state that guarantees no occurrence of a pattern in the skipped area (the state corresponds to the longest prefix of a pattern

matched to a suffix of the text scanned so far); this function corresponds to the Apostolico-Giancarlo *skip* function. In a step of the algorithm, the DAWG is used to first scan from right to left until a mismatch, upon which the AC machine is used to scan from left to right to recognize prefixes, and decide on the shift value. If the set of character matched using the DAWG overlap with the prefix matched by AC in the previous phase, a match is reported.

[0099] One embodiment of the pattern data structure for the DAWG-Match algorithm is presented below, with a corresponding memory data structure shown in Figure 11.

```

10 struct Pattern {
    char *string;          /* Pointer to string of bytes */
    int length;            /* Number of bytes in string */
    int offset;            /* string MUST NOT occur within first offset bytes of
15 text */
    int depth;             /* string MUST occur within first depth bytes of text */
    char fcase;            /* Flag (0=case independent; 1=case dependent) */
    struct Pattern *prev;  /* Pointer to previous pattern (for conditional search;
20 search
                        for string only if prev has been found */
    struct Pattern *next;  /* Pointer to next pattern (next must follow string in the
text
                        for condition to be satisfied (or modification of a
                        RegEx) */
25 int prev_pat_num;      /* Pattern number of prev in a pre-ordered list */
    int next_pat_num;     /* Pattern number of next in a pre-ordered list */
    int distance;         /* Same meaning as offset; start position of text is first
byte
                        after end of prev (if found) */
30 int within;           /* Same meaning as depth; start position of text is first
byte
                        after end of prev (if found) */
    int exact;            /* Exact start position of string in text */
    int relative;         /* Same meaning as exact; start position of text is first
35 byte
                        after end of prev (if found) */
};

```

[00100] The *Prev* and *Next* values will be NULL if not defined. Default values of *offset* and *distance* will be 0, unless explicitly defined. Default values of *depth*, *within*, *exact* and *relative* will be -1, unless explicitly defined. For pre-processing, pointer values *prev* and *next* are used. For the algorithm implementation in

microcode, *prev_pat_num* and *next_pat_num* are employed as pattern index values are used more ubiquitously (actual pattern address can be computed from these values, as pattern structure is of constant size).

[00101] In the memory map data structure of Figure 11, addresses are aligned to word (4 bytes) boundaries. The actual pattern string is a sequence of bytes starting at the byte address indicated in the first field of the data structure. In practice, pattern strings are stored contiguously and immediately after the end of the last pattern data structure. This enables pre-allocation of memory, computation of relative addresses during pre-processing and requires only one pointer (address of first pattern structure) to be maintained.

[00102] The AC machine employs a deterministic finite state automaton (Q, δ, f, s_0, T) , where Q is the set of states, δ is the transition function, f is the failure function, s_0 is the start state and T is the set of accepting states. The transition function determines which state to move to from a particular state on a given input. The failure function indicates *exception* transitions, for $(state, input)$ pairs that are not defined by the transition function.

[00103] The AC machine is structured as a *trie* of nodes, and can be specified just by the pointer to the start state. In one embodiment, the follow data structure is employed, while a corresponding memory data structure is shown in Figure 12.

```

20  struct transition {
        char input;                /* Byte value that triggers state transition */
        struct AC_node *destination; /* Pointer to destination AC-machine state
(node) */
25  };

    struct AC_node {
        int node_id;                /* Unique integer as state id */
        int shift;                  /* Maximum safe shift value (for DAWG-
30  Match):

                                shift is the minimal distance that can be
                                skipped
                                with a guarantee that no pattern can occur in
                                the
35  skipped area */

```

```

        int offset;                /* Prefix indicated by this state must end on or
after                               first offset bytes of text */
        int depth;                /* Prefix indicated by this state must end
5  within depth                  bytes of text */

        struct AC_node *fail;      /* Pointer to fail state */
        int num_transitions;       /* Number of transitions from this state */
        struct transition *next;   /* List of transitions from this state, ordered by
10                             ASCII value of input byte */

        char faccepting;          /* Flag (1 = accepting; 0 = not accepting) */
        int num_outputs;          /* Number of patterns recognized, if state is
                                accepting */
        struct Pattern **output;   /* Pattern(s) recognized, if state is
15  accepting */
        int *pattern_numbers;      /* List of patterns recognized, as index values
in a                               pre-ordered list of patterns */

    };

```

20 **[00104]** The default values of *offset*, *faccepting*, *num_transitions* and *num_outputs* are 0, unless explicitly defined or calculated. The default value of *depth* is -1, unless explicitly defined or calculated. Pointers are used for *destination*, *fail* and *output* only during pre-processing (computation of the machine). The microcode implementation uses *pattern_numbers* and *node_id* for

25 states (actual memory addresses are computed from these).

[00105] In one embodiment, an AC machine index structure is employed, as shown in Figure 13. Although the index and each of the AC machine nodes are designed to be independently relocatable in memory, it is advantageous to store the state information in contiguous fashion ordered by state ID, and to store the

30 structure information immediately after the index. This allows different addresses (relative) to be computed during pre-processing, reduces the number of pointers to be maintained to one (base address of index structure), and enables pre-allocation of memory.

[00106] A directed acyclic word graph, or a *suffix automaton*, recognizes all

35 suffixes of patterns in a given set (the DAWG-match algorithm takes reversed patterns as input). A DAWG does not have a failure function, and states don't need to have output patterns associated with them. One embodiment of a DAWG data

structure is presented below, while a corresponding memory data structure is shown in Figure 14:

```

5      struct suf_transition {
          char input;                /* Byte value that triggers state transition */
          char flag;                 /* Flag (0=primary edge; 1=secondary edge) */
          struct DAWG_node *destination; /* Pointer to destination DAWG state (node) */
      };

10     struct DAWG_node {
          int node_id;               /* Unique integer as state id */
          int num_transitions;       /* Number of transitions from this state */
          struct suf_transition *next; /* List of transitions from this state, ordered
15         struct DAWG_node *suffix_pointer; /* Pointer to parent in a virtual tree
that is                                     used for computation of the DAWG */
      };

```

[00107] The Default values of *flag* and *num_transitions* are 0, unless explicitly specified or calculated. The *flag* is used only for construction of the DAWG, and has no role to play in actual string matching. Pointers are used for *destination* only during pre-processing (computation of the machine). The microcode implementation uses *node_id* for states (actual memory addresses are computed from these). The DAWG machine index structure is identical to the AC machine index structure. The DAWG indices and structures are stored contiguously in a manner similar to that discussed above for the AC machine.

[00108] A suffix tree, like a DAWG, recognizes all suffixes of patterns in a given set (the AC-Suffix-Tree algorithm takes reversed patterns as input). This structure is essentially a DAWG unraveled into a tree, thus having a unique path from source to any other state. Details of one embodiment of a suffix tree data structure are provided below, while a corresponding memory data structure is shown in Figure 15:

```

35     struct tree_transition {
          char input;                /* Byte value that triggers state
transition */
          struct suffix_node *destination; /* Pointer to destination suffix tree
state (node) */
      };

```

```

    struct suffix_node {
        int node_id;                /* Unique integer as state id */
        int num_transitions;        /* Number of transitions from this state */
5      struct tree_transition *next; /* List of transitions from this state, ordered
                                   by ASCII value of input byte */
        struct AC_node *ACnode;    /* AC-machine state where scan will end,
starting
                                   from the source and the reverse of the
10     pattern
                                   recognized by this suffix tree state as input
        */
    };

```

15 [00109] The Default value of *num_transitions* is 0, unless explicitly specified or calculated. The *ACnode* is what differentiates a DAWG and a Suffix Tree. In the latter, each state is reached by a unique path of bytes (substring of a pattern) from the start state. Therefore, each state will have a unique AC machine node representing that traversed string. For example, let state *s* in the suffix tree be

20 reachable from the start state by a string $c_1c_2c_3\dots c_n$. Let this string be provided as input to the AC machine. If the AC machine scan stops at a state *d*, the value of *s*->*ACnode* will be set to *d*. Pointers are used for *destination* and *ACnode* only during pre-processing (computation of the machine). The microcode implementation uses *node_id* for states (actual memory addresses are computed from these). The suffix

25 tree machine employs an index structure that is identical to the AC machine index structure, while the suffix tree indices and structures are stored contiguously in a manner similar to that discussed above for the AC machine.

[00110] During pre-processing operations used to build the pattern data structures, the pattern list is scanned for case-independent patterns. In order to keep

30 the state machines deterministic and to be able to recognize such patterns, whatever combinations of case they occur in (within the packet payload), all combinations of byte strings that the particular pattern could occur as are generated in one embodiment. All these new combinations are added to the pattern list, but the *prev*

pointer is used to point to the original pattern, so that memory space is not wasted in storing these exponential numbers of new patterns.

[00111] In practice, manual or programmatic selection operations may be employed to generate and/or remove case-independent patterns. For example, while it might be reasonable to store combinations of the basic same character string using all small, and caps, and first letter only capitalization forms, it would typically be far less beneficial to generate forms of the character string that include random capitalization.

[00112] The foregoing implementation provides several additions to the basic AC machine construction algorithm as proposed by Aho and Corasick and referenced above, including:

Shift computation: This is necessary for the DAWG-Match algorithm as proposed by Crochemore *et al.*;

Offset and Depth computation: If patterns have associated offset and depth that are not equal to the default start and end of packet, we need to compute values for each state. If we did not compute these values for each state, we would be making unnecessary state transitions; only when an accepting state was found would we know whether or not the output pattern was valid, and in the process, we might miss some valid patterns; and

Incorporating case independent patterns: By computing all possible combinations of strings for case independent patterns, we ensure that a case independent pattern will always be recognized. In one embodiment, the output at any accepting state as the original pattern is marked, rather than having to store unnecessary combinations of patterns.

[00113] In view of the embodiments illustrated in Figures 6, 7 8a, 8b, and 9 and discussed above, the DAWG and suffix tree algorithms are implemented via execution of microcode threads on selected ME's. In addition to the foregoing

description, the operations are summarized as follows. First, an attempt is made to try to match the longest prefix possible of any pattern in the pattern set {P} to a suffix of T. For the DAWG and suffix tree algorithms described above, the prefix value can be represented by an equivalent state of the AC machine that recognizes
5 patterns in the set {P}.

[00114] For the first packet in the flow, the prefix of a pattern matched in the previous pattern is the empty string, so the *prev_state* input will be the start state of the AC machine. If, for any packet, this value is not equal to the start state, the AC machine is used to scan T from its first character in a forward direction, starting
10 from *prev_state*, till a suitable shift value is obtained and the DAWG can start its scan.

[00115] Towards the end of the packet scan, during a loop iteration, a substring of the packet content to be scanned by the DAWG will cross the end of the packet (*i.e.* *pos* >= size of packet). When this happens, we need to break form the loop and
15 resume scanning of the packet where the AC machine left off in the last loop iteration (*i.e.* at the same state and the same character position in the packet). The AC machine scan continues until the end of the packet is reached, at which point, the AC machine state is recorded as the *prev_state* in the associated flow table entry, which is employed to initialize the state for processing the next packet in the
20 sequence.

[00116] This mechanism for detecting patterns spanning multiple packets will work only if packets of the flow arrive in sequence. The *prev_state* value makes sense only if the current packet has a sequence number that is one higher than the packet that arrived previously in the flow. If packets don't arrive in order (this
25 applies to TCP and other connection-oriented protocols), such pattern detection cannot be performed just by storing one *prev_state* value in the flow entry. This problem may be addressed for some situations by storing multiple instances of such

state values, along with associated packet sequence numbers. However, this solution does not work if a packet arrives whose previous packet in sequence has not been scanned (this is because the nature of our machine data structures makes it difficult to detect pattern suffixes at the beginning of packets; detecting pattern
5 prefixes at the end of packets is easy in comparison). Another approach would be to use a combination of buffering and storing state values. In this implementation, if a packet arrives out of sequence, it could be buffered until its previous packet in sequence arrives and is scanned.

[00117] As discussed above, offset and depth tagged values may be provided to
10 reduce the portion of a packet or packets that need to be searched. The idea is that at during every step in the scan of the AC machine, we should be able to check offset and depth values against the character position to be scanned, and decide whether or not we are proceeding on the correct path. If the values associated with the current state don't match up to the current text position, a *fail* transition will be
15 made. An AC-machine state represents a prefix of one or more of the patterns we are searching for. Therefore, to ensure that we don't miss any pattern, we need to set the most conservative offset and depth value for that state. The formula we use for this is the same that is used to compute *shift* values, as described in Crochemore *et al.* as well as in the pre-processing algorithms described above. Whenever a
20 pattern is processed, the offset and depth for a state are calculated based on distance from start state and distance until the end of the pattern, respectively. When the state is reached again during the processing of a subsequent pattern, the offset is calculated as the minimum of the current computed value and the previously set value; the depth is calculated as the maximum of the current computed value and the
25 previously set value (an unspecified depth will be *infinity*, a value higher than all specified values). Offset and depth need not be changed during fail state

computation; if we come across a state with an invalid offset or depth during packet scan, a jump is automatically made to the fail state.

[00118] The foregoing scheme works for cases in which a single packet is searched. As discussed above, provisions are made in one embodiment to handle cases in which an offset and/or depth applies across packets. In order to implement this scheme, it is necessary to modify the application pattern data structure once it is copied into local memory to adjust the offset and/or depth values. At the same time, corresponding data needs to be stored in the associated flow table entry (*e.g.*, in the statistics information) to reflect: 1) that the pattern data structure needs to be adjusted; and 2) what those adjustments need to be.

[00119] In most instances, it will be unnecessary to perform offset and depth searches across multiple packets. In cases in which the applicable pattern set includes a large number of search strings, it will also usually be impractical to maintain offset and/or depth state information for each pattern in the statistics information. At the same time, offset and depth searches are typically employed to look for a particular signature (or only a few signatures). As a result, statistics information corresponding to the(se) signature(s) could be maintained in the flow table entry without consuming an excessive amount of memory. Whether or not offset and depth searches across packets is to be supported is left to the particular implementation.

Wu-Manber Algorithm Implementations

[00120] In some embodiments, search schemes based on the Wu-Manber algorithms ([WM92] S. Wu, U. Manber, "Fast Text Searching Allowing Errors", *Communications of the ACM* 35(10): 83-91(1992); and [WM94] S. Wu and U. Manber. "A fast algorithm for multi-pattern searching". *Technical Report TR-94-17*, University of Arizona, 1994) are employed. Wu and Manber take a very different approach from the state-machine based algorithms described above. These

algorithms also differ from other approaches by matching blocks of characters at a time, rather than single characters. If the total length of all patterns is M , the recommended size of the block (B) is $\log_{|\Sigma|} 2M$. The algorithm pre-computes three lookup tables called SHIFT, HASH and PREFIX. The SHIFT table is indexed by an integer that is a hashed mapping of a block of characters; this table is equivalent to the *bad character* table in the well-known Boyer-Moore algorithm (Boyer and J. Moore. "A fast string searching algorithm". *Communications of the ACM*, 20:762-772, 1977.). Entries for substrings common to multiple patterns contain the minimum possible shift value, to be more conservative. For larger values of B , tables could be compressed by mapping several strings to the same location, setting the minimum of the shift values in each case. The integer used to index to SHIFT is also used to index to HASH, which stores pointers to strings that map to that integer. The PREFIX table contains hash values for prefixes of each pattern, the prefix size being implementation-dependent. In each algorithm step, a set of B characters (at the right end of the text substring being scanned) are hashed to an integer, which is used to lookup the SHIFT value. If this value is greater than 0, the text is shifted. Otherwise, the hash of the prefix of this substring (of size m_s , where m_s is the length of the shortest pattern) is computed, and matched to the prefix of each string in the corresponding HASH entry (using the PREFIX table). If a prefix matches, actual text match takes place. This algorithm depends on the pattern sizes being mostly uniform, since only the shortest pattern length is considered during matching.

[00121] The pattern data structure as well as its corresponding memory data structure are identical to that used for the DAWG-Match algorithms. Standardizing the pattern data structure makes it easy to extend the content inspection framework and implement to different search algorithms. The shift table is structured as a 2-dimensional array of integers $S[n][4]$. As shown in the exemplary shift table

memory data structure of Figure 16 and below, each table entry $S[k]$ comprises a pointer to an integer array, which has the following structure:

```

5      {
          S[k][0]: int;      /* Primary shift value */
          S[k][1]: int;      /* Secondary shift value */
          S[k][2]: int;      /* Offset */
          S[k][3]: int;      /* Depth */
10     }

```

[00122] 'k' is the value (usually a long integer) used to index to the shift table. If the size of the alphabet is A and the size of the substring (chunk of string) that we will use as hash value is B, the maximum size of the table is A^B . 'k' is a (unique) hash value obtained from a block of B characters (the simplest hash function is to consider the block to be a base-A number of B digits). The *primary shift* refers to the distance by which the scan position must be shifted to the right. This value is 0 only if the scanned substring matches a substring at a suitable position in one of the patterns. The *secondary shift* refers to the distance by which the scan position must be shifted to the right only when the *primary* value is zero. This shift value replaces the default shift of 1, after processing a text position for *primary_shift*=0. This value will always be higher than 0 (and therefore, higher than *primary_shift* when it is equal to 0).

[00123] The *offset* is the least distance from the beginning of the packet content where a valid pattern ending with the current substring under consideration (given that all patterns are truncated to the least pattern size) may occur. The *depth* is the maximum distance from the beginning of the packet content within which a valid pattern ending with the current substring under consideration (given that all patterns are truncated to the least pattern size) must end. The *offset*, *depth* and *secondary_shift* values make sense, and are used, only if the *primary shift* value is 0, i.e., when a substring in the text (packet content) matches the end of one of the patterns under consideration.

[00124] The HASH table is structured as a single dimensional array of pointers $H[n]$. As shown in the exemplary HASH table memory data structure of Figure 17, each table entry $H[k]$ comprises a pointer to a pattern data structure. The index into the table is identical to that for the Shift table. The patterns (truncated to the
5 minimum pattern size) are sorted by the hash values of the chunk of B characters they end with. The value of $H[k]$ is the first pattern in this sorted list that has a hash value greater than or equal to 'k'.

[00125] Since the patterns data structures have a fixed size, they can be accessed in a random manner from memory. Therefore, the address of a pattern data
10 structure can be calculated using the sequence number of the pattern in the array. Accordingly, the values stored in the HASH table in memory comprise pattern index numbers (ranging from 0 to *number_of_patterns* - 1), rather than actual addresses, in order to make the patterns relocatable.

[00126] The PREFIX table is structured as a single dimensional array of integers
15 (usually long integers). The size of the table is equal to the number of patterns in the pattern set, and the pattern (or pattern index number) is used as index to any entry of the Prefix table. The value $Prefix[k]$ is a hash of the first B' characters in pattern number 'k' ('k' ranging from 0 to *number_of_patterns* -1). The hash is typically computed in an identical way to that computed for the Shift table (by
20 considering the prefix to be a base- A number consisting of B' digits). In the memory representation (one embodiment of which is shown in Figure 18), each table entry is the unmodified prefix value, since the alphabet to be employed comprises the entire ASCII set.

[00127] Both the SHIFT and HASH tables contain a large amount of redundant
25 information. The size of these tables is exponential, while the set of patterns is likely to be few in number. Therefore, we could expect most of the SHIFT table values to be equal to the default value. Similarly, a large number of HASH table entries

would contain useless information that would never be looked up. On the other hand, this significant waste of memory is compensated by the performance benefits of looking up information in constant time. For larger values of B (3 or more), the memory usage becomes prohibitive, so we use a B+ tree as a compressed data structure with a relatively small number of lookups required to fetch any piece of information.

[00128] This B+ tree is structured as a tree of nodes, where each leaf has constant distance from the root. Though the number of levels could be arbitrary, we have chosen the depth to be two, which will suffice for most practical purposes. Each child node is an array of the following structure:

```

struct hash_record {
    long key_value;           /* Hash value (index to SHIFT and HASH
tables) */
    int primary_shift;        /* Primary shift (as in the SHIFT table) */
    int secondary_shift;      /* Secondary shift (as in the SHIFT table) */
    int offset;               /* Offset (as in the SHIFT table) */
    int depth;                /* Depth (as in the SHIFT table) */
    struct Pattern *patlist_start; /* Pointer to start of pattern list (HASH table
entry) */
    struct Pattern *patlist_end; /* Pointer to end of pattern list (next HASH
table entry) */
    int pstart;               /* Pattern number in sorted list of patlist_start
*/
    int pend;                 /* Pattern number in sorted list of patlist_end
*/
};

```

[00129] The root node is special and is a constant-size array, each element being structured as follows:

```

struct bplus_tree_node {
    struct hash_record *list; /* Pointer to child node */
    int num_elements;         /* Number of elements in list */
};

struct bplus_tree_node root[A];

```

An exemplary memory structure of a root node is shown in Figure 19. The root has 'A' elements of type *struct bplus_tree_node*, where A is the alphabet size (256 in the case of the ASCII set).

[00130] Consider the SHIFT table. If the index value is a hash of a chunk of B characters $c_1c_2..c_B$, each element of the root can be accessed randomly by computing the hash value of the *most significant byte* of that chunk c_1 . The remainder of that chunk ($c_2..c_B$) is mapped to the *key_value* of an elements of a child node. Each child node (*list*) need not contain entries for all possible hash values for strings of size B-1; it only contains entries for those elements that don't have associated default primary shift value (which will be equal to *Minimum_pattern_size-B+1*). To access the information located in *SHIFT_table*[$c_1c_2..c_B$], we look up *root*[c_1].*list*(k) (where *list*(k) refers to the element in *root*[c_1].*list* having key value = hash value of $c_2..c_B$).

5

10 HASH table values are looked up in the same way.

[00131] The meaning of each entry in the foregoing structure is generally self-evident. The *patlist_end* and *pend* parameters do not refer to the last pattern in the list of patterns that are obtained from a particular HASH table entry; but rather refer to the pattern pointed to by the next HASH entry, or the pattern following this list of patterns in the global sorted pattern list.

15

[00132] Figure 20 shows one embodiment of a child node memory data structure. Although the root nodes and each of the child nodes are designed to be independently relocatable in memory, the embodiment of Figure 20 is implemented such that the child nodes immediately follow the root structure. Also, child nodes are written into memory in contiguous fashion, ordered by the index value of the parent element. Storing this data contiguously makes memory pre-allocation easy, and only one pointer (the address of the first word of the root) need be specified explicitly.

20

[00133] The Wu-Manber algorithm assumes that all patterns are of equal length; its correct operation relies on this assumption. Since that is not a practical assumption, we need to truncate all strings to the size of the shortest pattern. During an algorithm run, the complete pattern is matched when a candidate position is

25

discovered. The patterns that are passed to all pre-processing functions are truncated.

[00134] Pseudo-code examples used build the SHIFT, HASH, PREFIX, and B+ trees are provided in the appendix at the end of this specification. Additionally,
5 pseudo-code for implementing variations of the Wu-Manber algorithms is likewise provided in the appendix.

[00135] The Wu-Manber algorithm in the appendix significantly varies from that described in [WM94]. This algorithm supports extra features of patterns, such as offset, depth and case independence. There is also one more difference in the basic
10 algorithm, which is the *secondary* shift value. In the original algorithm proposed by Wu and Manber [WM94], if the primary shift value turned out to be 0, the skip value was always forced to be 1. This certainly ensures correct operation, but a skip of 1 is too conservative. The secondary shift indicates the least distance that needs to be skipped in order to match the substring (chunk of size B) under consideration.
15 The largest safe value for this skip will be the second lowest shift calculated during the computation of the SHIFT table, which is exactly our definition of the *secondary* shift.

[00136] As is evident from the Wu-Manber algorithm description, the largest prefix of any pattern that is matched at the end (suffix) of a previous packet is
20 provided as input to the algorithm when it is executed for the next packet. This information is provided in two parts: *prefix_address*, the address of the first byte of the pattern string whose prefix was matched, and *prefix_length*, the size of the prefix (or the number of bytes of the pattern that were matched successfully). For the first packet in the flow, the prefix of a pattern matched in the previous pattern is
25 the empty string, so the *prefix_length* parameter value will be 0, and the *prefix_address* uninitialized.

- [00137] Towards the end of the packet scan, during a loop iteration, the chunk to be used as index to the table entries will cross the end of the packet (*i.e.*, *position* $\geq$ size of packet). When this happens, we have to revert partially to the naïve string-matching algorithm. While there are still some inspected characters left at the end of the packet, those are matched to the prefix (of the same size) of every pattern in our pattern set. As soon as a match is obtained, the length and the address of the match are recorded in *prefix_length* and *prefix_address* respectively. When the next packet in sequence is to be scanned, this prefix matched is appended to the packet content and the scan proceeds in normal fashion.
- 10 [00138] This technique for matching patterns that span multiple packets fails when packets arrive out of sequence. Matching arbitrary suffixes at the beginning of a packet, hoping that it is part of a pattern that began in a previous packet will not work. As mentioned for the DAWG-Match case, a combination of buffering and state maintenance could provide a comprehensive solution.
- 15 [00139] In some embodiments, the foregoing Wu-Manber algorithms may also be employed for offset/depth searching. While attempting to match patterns, we also need to ensure that the offset and depth constraints are met, and we don't waste cycles in trying to match a pattern that might be present but turn out to have invalid limit values. Therefore, we associate valid offset and depth values with each SHIFT
- 20 table entry. The offset and depth values are set in the same way as the shift values are, by selecting the safest possible value. If a chunk of bytes *S* is used to index into the SHIFT table, the corresponding offset is set to be the minimum offset of any pattern that ends with this chunk (the pattern is truncated first to the shortest pattern length). Similarly the depth is set to be the maximum value of (*depth-length+minimum_pattern_size*), as this indicates where the truncated pattern must
- 25 end within the text.

[00140] During a packet scan, when a primary shift value of 0 is observed, we first check the offset and depth (looked up from the table) and match them to the current text position. Only if the constraints are met do we proceed to cycle through the list of patterns indicated by the HASH table. During this procedure, after a
5 successful PREFIX table match and before the entire pattern match is attempted byte-by-byte, the offset and depth associated with the patterns themselves are matched with the current text position. If the current position is within limits, the complete pattern match is attempted. If any of the offset and depth checks described above fail, the position is incremented by a value equal to the secondary
10 shift and the search continued.

[00141] Case independence, fortunately, can be handled with much greater memory efficiency by the Wu-Manber algorithm as compared to the DAWG-Match algorithm. The extra work that needs to be done is mostly shifted to the pre-processing phase. While updating a SHIFT table, if a case independent pattern is
15 processed, all case-dependent combinations of the B-length chunk under consideration are computed and the SHIFT table entries corresponding to each of those combinations adjusted. For the HASH and the PREFIX tables, index and prefix values are computed from lower case equivalents of case-independent patterns.

20 [00142] During a packet scan, a SHIFT table lookup is straightforward and no addition is required to the original algorithm proposed by Wu and Manber, as valid SHIFT entries exist for all potential case combinations. For the HASH table lookup though, we need to compute two separate index values:

Normal index, using the characters of the chunk as they are; and
25 Lower case index, computed by converting all upper case characters in the B-chunk to lower case.

[00143] The second index is necessary as the HASH entry for a case independent pattern is computed from lower case characters only. The entire cycle of checking prefixes, followed by matching complete patterns is done twice, one for the normal index value and another for the lower case index value. Likewise, both normal and lower case versions of the pattern prefixes are computed. If the pattern under consideration is case independent, the lower case prefix is matched, otherwise the normal one is matched. Once the algorithm gets down to matching patterns byte-by-byte with the packet content, case-independent matching becomes straightforward.

10 [00144] In one embodiment, the content inspection framework is designed to be a module that could be plugged into a previously-implemented stateful firewall framework. Under the Intel IXP processing model, the content inspection framework can be plugged in as a single microblock into any packet-processing pipeline. In one embodiment, the firewall uses CA (Condition-Action) pairs for rule processing. For example, certain CA rules could pertain to content inspection, like checking for a particular offending pattern with packet content (condition) and dropping the packet if the pattern is detected (action).

[00145] The memory map for a typical content inspection CA rule is shown in Figure 21. The first two words (or the first long word) are reserved, and the actual information begins at an offset of 8 bytes. The meanings of each of the fields is as follows:

1. **Algorithm code:** This indicates the type of search algorithm to use.
 - 1 = DAWG-Match algorithm
 - 25 2 = Suffix Tree-Match algorithm
 - 3 = Wu-Manber algorithm
 - 4 = Wu-Manber algorithm using B+ tree structure in place of SHIFT and HASH tables
 - 5 = Match pattern given exact start position (*i.e.* it is assumed that the pattern starts at the specified position or doesn't occur at all)
 - 30

2. **SRAM address:** This indicates the address (or pointer to index of addresses) of data structures used for pattern searching. For the first 4 algorithms, this address points to the first word of the *master index*; this contains information about the addresses of the various data structures that are used for inspection.

3. **‘Spanning pattern check’ policy:** For the first 4 algorithm codes, this is just a flag that has a value 0 (which indicates that checks for patterns that could span multiple packets should not be made) or 1 (which indicates that such checks must be made). For the exact position match, this field indicates the position in the packet content (as an offset from the first byte) where the pattern match should be performed.

[00146] In order to check for patterns spanning multiple packets, we need to store the partial match information (AC state for DAWG match, pattern address and prefix length for Wu-Manber) with the entry corresponding to the current flow in the flow table. We also need to store some information regarding the validity of the stored information when the next packet arrives to be processed. For TCP flows, therefore, we store the sequence number of the next expected offset along with the state information. Similar kinds of sequence information could be stored for any other connection-oriented protocol.

[00147] In one embodiment of the flow table entry, this information is stored in the following manner: the first word is the expected sequence number of the first byte in the next packet, the second word is the AC machine state from the end of the previous scan (for algorithm codes 1 and 2) or the prefix length (algorithm codes 3 and 4) and the third word is the partially matched pattern address (used only for algorithm codes 3 and 4).

Initialization and Code Load

[00148] Returning to Figure 7, during initialization of a network line card or network device, coded instructions (*e.g.*, microcode) to facilitate various packet-processing and content inspection functions and operations described herein are loaded into the control stores 728 of the selected ME's. In general, the instructions
5 may be loaded from an instruction store 715 comprising a non-volatile storage device, such as a flash memory device. Other examples of non-volatile stores include read-only memories (ROMs), programmable ROMs (PROMs), and electronically erasable PROMs (EEPROMs).

[00149] In addition to loading the instructions from a local (to line card 700)
10 storage device, instructions may be loaded from an external source. For example, in one embodiment, the instructions are stored on a disk drive hosted by another line card (not shown) or otherwise provided by the network element in which line card 700 is installed. In yet another embodiment, the instructions are downloaded from a remote server or the like via a network as a carrier wave.

[00150] In general, each of the functional blocks illustrated in the figures and discussed herein may be implemented using programmable logic in hardware, programmable logic in firmware/software via execution of corresponding threads on one or more compute engines, or a combination of the two. For example, in one embodiment the operations depicted for each of deep packet inspection module 202,
15 flow manager 204, queue manager 708, and receive engine 710 are facilitated by executing corresponding instruction threads on one or more compute engines in an NPU in a manner integrated into a context pipeline used for packet processing. Optionally, all or a portion of the various operations provided by the foregoing functional blocks may be implemented by appropriately-programmed logic
20 embedded on an NPU or the like.

[00151] Thus, embodiments of this invention may be used as or to support firmware/software modules and/or programs executed upon some form of

processing core or otherwise implemented or realized upon or within a machine-readable medium. A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (*e.g.*, a computer). For example, a machine-readable medium can include such as a read only memory (ROM); a random access memory (RAM); a magnetic disk storage media; an optical storage media; and a flash memory device, *etc.* In addition, a machine-readable medium can include propagated signals such as electrical, optical, acoustical or other form of propagated signals (*e.g.*, carrier waves, infrared signals, digital signals, *etc.*).

10 [00152] The above description of illustrated embodiments of the invention, including what is described in the Abstract, is not intended to be exhaustive or to limit the invention to the precise forms disclosed. While specific embodiments of, and examples for, the invention are described herein for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as
15 those skilled in the relevant art will recognize.

[00153] These modifications can be made to the invention in light of the above detailed description. The terms used in the following claims should not be construed to limit the invention to the specific embodiments disclosed in the specification and the drawings. Rather, the scope of the invention is to be
20 determined entirely by the following claims, which are to be construed in accordance with established doctrines of claim interpretation.

APPENDIX

1. The algorithm for the construction of the **SHIFT table** is given below in pseudo-C code.

5 *Input:* Array of patterns, Chunk size used for hashing into table (B), Alphabet size (A)

Output: Pointer to SHIFT table

```

int **build_SHIFT_table(struct Pattern *P, int num_pat, int B, int A)
{
10     /* Initialization */
    table_size = AB;
    min_pattern_size = minimum(Pi.length);
    default_offset = minimum(Pi.offset);
15     default_depth = minimum(Pi.depth);

    for(each entry of the SHIFT_table, i: 0 to table_size-1)
    {
        /* Initialize each entry to the default value before pattern
processing */
20         SHIFT_table[0][0] = min_pattern_size-B+1;
        SHIFT_table[0][1] = min_pattern_size-B+1;
        SHIFT_table[0][2] = default_offset;
        SHIFT_table[0][3] = default_depth;
    }
25     /* Process patterns and update SHIFT_table values */
    for(each pattern Pi, i:0 to num_pat-1)
    {
        for(each substring S = Pi.string[j...j+B-1], j: 0 to min_pattern_size-
30         B)
        {
            /* Compute hash value to index into SHIFT_table */
            Sval = HASH_value(S);
            /* Compute primary shift as least offset from start position of
35         substring
                to end of pattern */
            SHIFT_table[Sval][0] = minimum(SHIFT_table[Sval][0], min_pattern_size-
j-B);
40         /* Compute secondary shift as second least shift value for the
            particular
                substring */
            SHIFT_table[Sval][1] = least value of (min_size-j-B) that is strictly
greater than SHIFT_table[Sval][0];
45         /* Compute offset as least offset among all patterns containing this
            substring */
            SHIFT_table[Sval][2] = minimum(SHIFT_table[Sval][2], Pi.offset);
            /* Compute depth as maximum depth among all patterns containing this
            substring */
50         SHIFT_table[Sval][3] = maximum(SHIFT_table[Sval][3], Pi.depth -
Pi.length +
            min_pattern_size);

            /* Update SHIFT_table for case independent patterns */
            if(Pi.fcase == 0)
55             {
                for(each possible case combination of S)
                {
                    val = HASH_value(instance of S);
                    /* Update all four fields: both shifts, offset and
60                 depth in
                        the same way as above */
                    update(SHIFT_table[val]);
                }
            }
65         }
    }

    return SHIFT_table;
}
70

```

2. The algorithm for the construction of the **HASH table** is given below in pseudo-C code.

Input: Array of patterns, Chunk size used for hashing into table (B), Alphabet size (A)

Output: Pointer to HASH table

```

5  struct Pattern **build_HASH_table(struct Pattern *P, int num_pat, int B,
   int A)
   {
10     /* Initialization */
       table_size = AB;
       min_pattern_size = minimum(Pi.length);
       sort(P, in ascending order based on hash value computed from each
       Pi.string[min_pattern_size-B...min_pattern_size-1]);

15     /* Initialize index value into HASH_table */
       index = 0;
       for(each pattern Pi, i:0 to num_pat-1)
       {
20         /* Obtain substring of B characters that each pattern ends with */
           S = Pi.string[min_pattern_size-B...min_pattern_size-1];
           /* If pattern is case independent, map it to its lower case
           equivalent */
           if(Pi.fcase == 0)
25             S = convert_to_lower_case(S);

           Sval = HASH_value(S);
           while(Sval < count)
           {
30             /* Keep updating table values and pointers until hash value
               gets
               large enough */
               HASH_table[index] = Pi;
               HASH_table_pattern_number[index] = i;
               index ++;
35           }
       }
       /* For higher index values, fill up with non-existent pointer entry; this
       does not
       matter as such index values will never be looked up */
40     for(each i:index to table_size-1)
       {
           HASH_table[index] = Pnum_pat;
           HASH_table_pattern_number[index] = num_pat;
45       }

       return HASH_table;
   }

```

3. The algorithm for the construction of the **PREFIX table** is given below in pseudo-C code.

Input: Array of patterns, Size of prefix (B'), Alphabet size (A)

Output: Pointer to PREFIX table

```

55 long int *build_PREFIX_table(struct Pattern *P, int num_pat, int B', int A)
   {
       /* Initialization */
       table_size = num_pat;

60     for(each pattern Pi, i:0 to num_pat-1)
       {
           /* Compute hash of prefix value; if pattern is case independent,
           convert
           prefix into lower case and compute hash */
65         S = Pi.string[0...B'-1];
           if(Pi.fcase == 0)
               S = convert_to_lower_case(S);

           PREFIX_table[i] = HASH_value(S);
70     }
   }

```

```

    return PREFIX_table;
}

```

5

4. The algorithm for the construction of the **B+ tree** is given below in pseudo-C code.

Input: Array of patterns sorted by hash value (hash computed in the same way as for the SHIFT and HASH tables), Size of prefix (B'), Alphabet size (A), SHIFT table, HASH table

10 *Output:* Pointer to B+ tree root

```

struct bplus_tree_node *build_Bplus_tree(int **SHIFT_table, struct Pattern
**HASH_table, int *HASH_table_pattern_number, struct Pattern *P, int
num_pat, int  $B'$ , int  $A$ )
15 {
    /* Compute default primary shift; these are the entries we will eliminate */
    min_pattern_size = minimum( $P_i$ .length);
    default_primary_shift = min_pattern_size- $B+1$ ;

20    start = 0;    /* Initialize SHIFT table entry position to scan from */

    /* Generate a child node per iteration */
    for( $i:0$  to  $A-1$ )
    {
25        /* Process  $A^{B-1}$  SHIFT and HASH table entries during each loop iteration
and
        generate hash records for each entry with a non-default primary shift
        */
        counter = 0;
        for( $j:start$  to  $start + A^{B-1}$ )
        {
            if(SHIFT_table[j][0] != default_primary_shift)
            {
35                /* Add all SHIFT and HASH table information and set
key */
                root[i].list[counter].primary_shift =
SHIFT_table[j][0];
                root[i].list[counter].secondary_shift =
SHIFT_table[j][1];
40                root[i].list[counter].offset = SHIFT_table[j][2];
                root[i].list[counter].depth = SHIFT_table[j][3];
                root[i].list[counter].patlist_start = HASH_table[j];
                root[i].list[counter].patlist_end = HASH_table[j+1];
                root[i].list[counter].pstart =
45                    HASH_table_pattern_numbe
                    r[j];
                root[i].list[counter].pend =
                    HASH_table_pattern_number[j+1];
50                root[i].list[counter].key_value = j-start;
                counter++;
            }
            root[i].num_elements = count;
            start = start +  $A^{B-1}$ ;
55        }
    }

    return root;
}
60

```

The field *key_value* is actually the index of the *hash_record* structure list, if all records (inclusive of the ones with default shift values) were added). Therefore, the above algorithm ensures that *list* (for each tree child node) is sorted in ascending order of *key_value*.

65

Wu-Manber algorithm

Input: Packet content (string of bytes) T[0...n-1], SHIFT table, HASH table, PREFIX table, Length of shortest pattern, Number of patterns, Chunk size (B), Prefix size (B'), Flag (policy) indicating whether or not we need to check for patterns spanning packet boundaries. Length of a pattern prefix matched at the end of the previous packet, Address of pattern string whose prefix was matched at the end of the previous packet.

Output: Identity of pattern detected (or -1 if none detected), Position of start of pattern as offset from the start of the string T.

10 The algorithm is described below in pseudo C-code.

```

15 Wu_Manber(int **SHIFT_table, struct Pattern **HASH_table, long int
    *PREFIX_table, char *T, int span_policy, int *prefix_length, char
    *prefix_address, int MIN_PATTERN_SIZE, int NUM_PATTERNS, int B, int B')
{
    /* Initialize position to earliest possible end position of a truncated
    pattern */
    position = MIN_PATTERN_SIZE-1;

20    /* If pattern overlap check needs to be made, append prefix matched in
    previous packet at the beginning of the packet content */
    if(span_policy == CHECK_SPANNING_PATTERN)
    {
25        S = Prefix(pattern at address prefix_address, length = prefix_len)
        T = append(S to the beginning of T);
    }

    while(position < PACKET_SIZE)
    {
30        /* Get chunk ending at 'position' and compute table index value */
        chunk = T[position-B....position];
        chunk_val = HASH_value(chunk);

35        /* Convert to lower case and compute index; this will be used later
        */
        chunk_lower_case = convert_to_lower_case(chunk);
        chunk_val_lower_case = HASH_value(chunk_lower_case);

40        if(SHIFT_table[chunk_val][0] > 0)
        {
            /* No pattern can be found at this position, so shift and try
            to
            match at the next position */
45            position = position + SHIFT_table[chunk_val][0];
            continue;
        }

        /* Primary shift value is 0, so we have a potential match starting at
        'position-MIN_PATTERN_SIZE+1' */

        /* Check offset for candidate position */
        if(position-MIN_PATTERN_SIZE+1 < SHIFT_table[0][2])
        {
55            /* No pattern ending with chunk satisfies offset limit at
            current position, so skip using secondary shift value and try to
            match at the next position */
60            position = position + SHIFT_table[chunk_val][1];
            continue;
        }

        /* Check depth for candidate position */
65        if(position >= SHIFT_table[0][3])
        {
            /* No pattern ending with chunk satisfies depth limit at
            current position, so skip using secondary shift value and try to
70            match

```

```

        at the next position */
        position = position + SHIFT_table[chunk_val][1];
        continue;
5      }

      /* Compute prefix values (normal and lower case) for candidate string
      */
      prefix = T[position-MIN_PATTERN_SIZE+1....position-
MIN_PATTERN_SIZE+B'];
10     prefix_val = HASH_value(prefix);
      prefix_lower_case_val = HASH_value(convert_to_lower_case(prefix));

      /* Candidate patterns: pointer to patterns are in ascending order */
      for(each pattern Pi: HASH_table[chunk_val] <= Pi <
15     HASH_table[chunk_val+1])
      {
          if(Pi.fcase == 0 && prefix_lower_case_val == PREFIX_TABLE[i])
          {
              /* Convert both strings to be matched to lower case*/
              String1 = convert_to_lower_case(Pi.string);
              String2 = convert_to_lower_case(T[position-
MIN_PATTERN_SIZE+1.....position-
MIN_PATTERN_SIZE+Pi.length]);
20             if(String1 == String2)
              {
                  Record Pi; /* Pattern match */
                  Record 'position';
                  Depending on policy, return or continue
30             scanning T;
              }
          }
          else if(Pi.fcase == 1 && prefix_val == PREFIX_TABLE[i])
          {
              if(Pi.string == T[position-
MIN_PATTERN_SIZE+1....position-
MIN_PATTERN_SIZE+Pi.length])
              {
                  Record Pi; /* Pattern match */
                  Record 'position';
                  Depending on policy, return or continue
40             scanning T;
              }
          }
      }
45     }

    /* Now replace chunk matched to index into SHIFT_table with its lower
    case equivalent */
    Repeat(above 'for' loop replacing 'chunk_val' by
50     'chunk_val_lower_case');

    /* Skip using secondary shift value and continue scanning */
    position = position+SHIFT_table[chunk_val][1];
55     }

    /* If we don't need to check if a pattern might span multiple packets, return
    */
    if(span_policy == CHECK_SPANNING_PATTERN)
        return;
60     }

    /* Calculate start position of a candidate pattern that could overflow into
    the next packet */
    position = position-MIN_PATTERN_SIZE+1;
65     while(position < PACKET_SIZE)
    {
        for(each pattern Pi, i:0 to num_pat)
        {
70             /* Length of packet content remaining that could match a
            prefix of
            one of the patterns */
            length = PACKET_SIZE-position+1;

```

```

/* Match prefix from the PREFIX table before attempting complete
match */
5  if(Pi.fcase == 0)
{
    if(length >= B')
        if(PREFIX_table[i] != T[position...position+B'-1])
            continue;
10  if(T[position...PACKET_SIZE-1] == Pi.string[0...length-1])
    {
        Record prefix_length = length;
        Record prefix_address = Pi.string;
        return;
    }
15  }
    If pattern Pi is case independent, attempt above match in a case-
    independent way;
    }
20  }
    Record prefix_length = 0;
    return;
}
25

```

Wu-Manber algorithm using B+ trees

30 **Input:** Packet content (string of bytes) T[0...n-1], B+ tree root pointer, PREFIX table, Length of shortest pattern, Number of patterns, Chunk size (B), Prefix size (B'), Flag (policy) indicating whether or not we need to check for patterns spanning packet boundaries. Length of a pattern prefix matched at the end of the previous packet, Address of pattern string whose prefix was matched at the end of the previous packet
35 **Output:** Identity of pattern detected (or -1 if none detected), Position of start of pattern as offset from the start of the string T.

40 The algorithm is identical to the earlier described Wu-Manber algorithm, except for the data structures used for shift values and pointer lookup. The primary and secondary shifts, as well as the offset and depth values are now looked up from a B+ tree child entry (from a *hash_record* structure) instead of a SHIFT table. Pointers to the candidate patterns that might occur at a given position are also obtained from the same B+ tree child node entry instead of a HASH table. But unlike a SHIFT or a HASH table, these values cannot be looked up directly (in constant time). A sequential scan would be too inefficient. Therefore, we use binary search to locate the appropriate record; the lookup
45 steps are described below:

1. Get the chunk (substring) $c_1c_2..c_B$ used to compute the hash value.
2. Use most significant byte (c_1) to locate the appropriate child node pointer and number of child nodes ($root[c_1].list$ and $root[c_1].num_elements$).
3. Use the rest of the chunk $c_2..c_B$, mapped to an integer, to locate the
50 appropriate *hash_record* structure. This integer will be the *key_value* of any matching structure. Since all structures are sorted in ascending order based on *key_value*, a simple binary search will return the required data). For an alphabet of size 256, a maximum of
55 8 lookups will be required to locate a record.

All algorithm steps apart from the data lookup are identical to those in the Wu-Manber algorithm described above.

60 Practical values for B would be 2 and 3 (having B = 1 would defeat the purpose of building a tree, though it would work). Practical values for B' would be 4 or less. Both B and B' must be smaller than the length of the shortest pattern.

CLAIMS

What is claimed is:

5

- 1 1. A method comprising:
2 performing content inspection across multiple packets to identify a search
3 string having a first portion in an earlier packet and a second portion in a following
4 packet,
5 wherein the content inspection is performed via execution of a plurality of
6 instruction threads.

- 1 2. The method of claim 1, wherein the content inspection includes concurrently
2 searching for multiple different strings.

- 1 3. The method of claim 2, wherein the content inspection is performed via a
2 computer implementation of a string search algorithm based on the DAWG-
3 (Directed Acyclic Word Graph) match algorithm.

- 1 4. The method of claim 2, wherein the content inspection is performed via a
2 computer implementation of a string search algorithm based on the Wu-Manber
3 algorithm.

- 1 5. The method of claim 1, further comprising:
2 receiving a first packet;
3 classifying the first packet to assign it to a flow;

4 identifying one or more strings to be searched for in packets assigned to the
5 flow;
6 employing a content inspection of the first packet using a first instruction
7 thread;
8 detecting a partial match of one of the one or more search strings that is
9 located toward the end of the first packet;
10 saving string search state information corresponding to the partial match;
11 receiving a second packet;
12 classifying the second packet to assign it to the flow;
13 retrieving the string search state information;
14 employing the string search state information to detect a second portion of
15 the search string, the second portion of the search string beginning toward the start
16 of the second packet.

1 6. The method of claim 5, further comprising:
2 identifying a string match condition has occurred; and
3 performing an associated action for the condition to process subsequently-
4 received packets assigned to the flow.

1 7. The method of claim 1, further comprising:
2 employing an offset mechanisms for at least one search string, the offset
3 mechanism skipping a first portion of a first packet defined as an offset for an
4 associated search string.

1 8. The method of claim 1, further comprising:

2 employing a depth mechanisms for at least one search string, the depth
3 mechanism tracking a depth of a search in a packet and stopping the search upon
4 reaching a depth defined for an associated search string.

1 9. The method of claim 1, further comprising:
2 performing a multi-packet search for a case-independent string.

1 10. The method of claim 1, wherein the content inspection comprises an
2 operation in a context pipeline of operations that are performed to process packets at
3 a network device, the context pipeline of operations being performed via execution
4 of multiple instruction threads on multiple compute engines.

1 11. A machine-readable medium to provide instructions, which if executed on a
2 multi-threaded network processor perform operations comprising:
3 performing content inspection across multiple packets to identify a search
4 string having a first portion in an earlier packet and a second portion in a following
5 packet,
6 wherein the content inspection for a given packet is performed by execution
7 of a separate instance of an instruction thread.

1 12. The machine-readable medium of claim 11, wherein execution of the
2 instructions performs further operations comprising:
3 classifying a first packet that is received to assign it to a flow;
4 identifying one or more strings to be searched for in packets assigned to the
5 flow;
6 employing a content inspection of the first packet using a first instruction
7 thread;

8 detecting a partial match of one of the one or more search strings that is
9 located toward the end of the first packet;
10 saving string search state information corresponding to the partial match;
11 classifying a second packet that is received to assign it to the flow;
12 retrieving the string search state information; and
13 employing the string search state information to detect a second portion of
14 the search string, the second portion of the search string beginning toward the start
15 of the second packet.

1 13. The machine-readable medium of claim 11, wherein execution of the
2 instructions performs further operations comprising:
3 maintaining a flow table in at least one shared memory resource, the flow
4 table including a plurality of entries associated with a respective flow, each flow
5 table entry including,
6 information identifying a pattern data structure employed for a multi-
7 pattern search associated with the respective flow; and
8 state information to identify a state of a the multi-pattern search;
9 in response to receiving a packet,
10 classifying the packet to a flow;
11 enqueueing the packet into a flow queue allocated to the flow;
12 locating the flow table entry corresponding to the flow;
13 copying the packet into a local memory of a compute engine on the network
14 processor;
15 copying the pattern data structure identified by the flow table entry into the
16 local memory of the compute engine; and
17 performing a multi-pattern string search of the packet on the compute engine
18 using the pattern data structure.

1 14. The machine-readable medium of claim 13, wherein execution of the
2 instructions performs further operations comprising:
3 classifying the packet to a new flow; and
4 generating a new entry in the flow table corresponding to the new flow.

1 15. The machine-readable medium of claim 11, wherein the content inspection
2 is performed via a computer implementation of a string search algorithm based on
3 the DAWG-(Directed Acyclic Word Graph) match algorithm.

1 16. The machine-readable medium of claim 11, wherein the content inspection
2 is performed via a computer implementation of a string search algorithm based on
3 the Wu-Manber algorithm.

1 17. A network line card, comprising:
2 a network processor unit (NPU) having a plurality of multi-threaded
3 compute engines;
4 a dynamic random access memory (DRAM) store, communicatively-coupled
5 to the NPU;
6 an static random access memory (SRAM) store, communicatively-coupled
7 to the NPU; and
8 an storage device, communicatively-coupled to the NPU, having instructions
9 stored therein to be executed on the NPU to perform operations comprising:
10 classifying received packets to associated flows; and
11 performing content inspection across multiple packets for a given
12 flow to identify a search string in a data stream transported via the multiple
13 packets having respective portions in a sequential pair of packets.

1 18. The network line card of claim 17, wherein the NPU comprises an Intel IPX-
2 based NPU.

1 19. The network line card of claim 17, wherein the content inspection is
2 performed via a computer implementation of a string search algorithm based on the
3 DAWG-(Directed Acyclic Word Graph) match algorithm.

1 20. The network line card of claim 17, wherein the content inspection is
2 performed via a computer implementation of a string search algorithm based on the
3 Wu-Manber algorithm.

1/14




*Fig. 1**Fig. 3*

Fig. 2






Fig. 4a




Fig. 4b




Fig. 5

4/14




*Fig. 6*


Fig. 7




6/14

**Fig. 8a**

7/14

**Fig. 8b**

8/14









Fig. 10

Fig. 11






AC Machine Node

Fig. 12**Fig. 13**

AC Machine Index Structure

Fig. 14

DAWG Data Structure

11/14

0	8	16	24	32

# Trans.		AC state number (23 bits)		
Input 1		Next state number 1		

.....				

Input n		Next state number n		

Suffix Tree Data Structure

Fig. 15

0	8	16	24	32

Primary Shift ($S[0][0]$)				

Secondary Shift ($S[0][1]$)				

Offset ($S[0][2]$)				

Depth ($S[0][3]$)				

.....				

.....				

.....				

Primary Shift ($S[N-1][0]$)				

Secondary Shift ($S[N-1][1]$)				

Offset ($S[N-1][2]$)				

Depth ($S[N-1][3]$)				

SHIFT Table

Fig. 16

12/14

0	8	16	24	32

# of first pattern in list with hash = 0				

.....				

# of first pattern in list with hash = N-1				

HASH Table

Fig. 17

0	8	16	24	32

Prefix of size B' of pattern # 0				

.....				

Prefix of size B' of pattern # N-1				

PREFIX Table

Fig. 18

0	8	16	24	32

Number of records in child node #1				

SRAM address of child node #1				

.....				

Number of records in child node #A-1				

SRAM address of child node #A-1				

Root Node

Fig. 19

13/14

0	8	16	24	32

Record #1: primary shift				

Record #1: secondary shift				

Record #1: offset				

Record #1: depth				

Record #1: start pattern index number				

Record #1: end pattern index number				

Record #1: key value				

.....				

Record #k: primary shift				

Record #k: secondary shift				

Record #k: offset				

Record #k: depth				

Record #k: start pattern index number				


Record #k: end pattern index number				

Record #k: key value				

Child Node

Fig. 20

14/14



Condition Action Rule

Fig. 21