



(51) International Patent Classification:

H04L 9/32 (2006.01) H04L 1/00 (2006.01)
H04L 9/06 (2006.01) H04L 29/08 (2006.01)
H04L 12/801 (2013.01) H04L 29/06 (2006.01)

(21) International Application Number:

PCT/IB2015/059895

(22) International Filing Date:

22 December 2015 (22.12.2015)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/579,987 22 December 2014 (22.12.2014) US

(71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (PUBL) [SE/SE]; SE-164 83 Stockholm (SE).

(72) Inventors: KOLTE, Ritesh; 350 Serra Mall, Packard 201, Stanford, California 94305 (US). ASNANI, Himanshu; 250 Holger Way, San Jose, California 95134 (US). ARISOYLU, Mustafa; 760 Overture Ct., San Jose, California 95134 (US).

(74) Agents: HENDERSON, Mymy et al.; 6300 Legacy, MS EVR 1-C-11, Plano, Texas 75024 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available):

AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LI, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: METHOD AND SYSTEM FOR PACKET REDUNDANCY REMOVAL

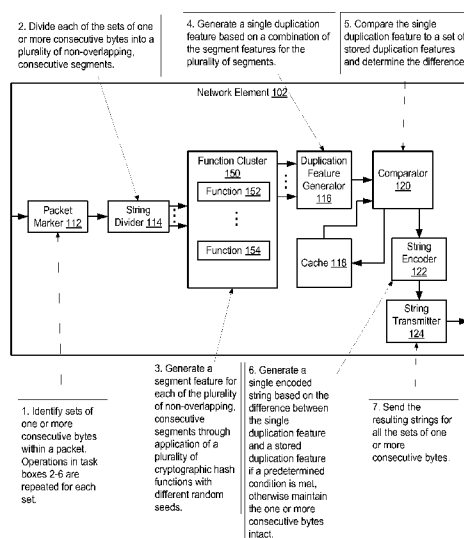



FIG. 1

(57) Abstract: A method/system for removing redundancy in packets is disclosed. In one embodiment, for each of the sets of one or more consecutive bytes within the packet, the method divides the one or more consecutive bytes into a plurality of non-overlapping, consecutive segments; generates a segment feature for each of the plurality of non-overlapping, consecutive segments through application of a plurality of cryptographic hash functions with different random seeds; generates a single duplication feature based on a combination of the segment features for the plurality of non-overlapping, consecutive segments; and generates a single compressed string when a predetermined condition is met, based on a comparison of the single duplication feature and a set of stored duplication features. The method continues with sending the resulting strings in place of the packet toward the packet's destination.



METHOD AND SYSTEM FOR PACKET REDUNDANCY REMOVAL

FIELD OF INVENTION

The embodiments of the invention are related to the field of networking. More specifically, the embodiments of the invention relate to removing redundancy in
5 packets for efficient transmission.

BACKGROUND

In networking, bandwidth and processing resources at each network are utilized in packet transmission. The amount of transmission correlates to the amount of bandwidth and processing resource consumed, and it is desirable to minimize
10 transmission of and/or remove redundancy in packets prior to transmission. When packets are updated to remove redundancy, the updated packets are restored at the destination by adding back the removed redundancy so the original packets can be utilized without the end user of the packets realizing the changes of the packets during transmission.

15 One way to remove redundancy in packets is through caching. Caching may be generally categorized into two types. One type is object caching, which removes transmission redundancy at the object level. For example, for access hypertext transfer protocol (HTTP) content, a client sends a request for an object (e.g., file, document, image, etc.) to a server. Instead of directly retrieving the object from the server, the
20 request is intercepted by a proxy server, which checks if the proxy server has the object in its local cache. If it has the object, the proxy server responds by sending the cached object without the server retrieving the object. Object caching is most useful in the following scenarios:

The transmitted content does not change much;

25 The transmitted content can be pre-populated prior to users trying to access it; and/or

The transmitted content needs to be accessed by multiple users.

Object caching is typically limited to specific protocols, and if a byte changes, the whole object has to be retrieved again. Thus, if the transmitted content is not static, another type of caching, byte caching, may be utilized. Byte caching is a protocol-independent bidirectional caching technique that functions by looking for a common
5 sequence of data. If any duplicate is found, the duplicate data is removed from the byte sequence, and in its place, a token is added. Byte caching is useful in accelerating traffic, and it may be used along with object caching to further accelerate specific protocols. Byte caching looks for repetition of exact sequences of data, and if an exact match is not found, byte caching does not compress the data to be transmitted.

10 SUMMARY

A method for removing redundancy in packets is disclosed. The method includes identifying sets of one or more consecutive bytes within a packet. For each of the sets of one or more consecutive bytes within the packet, the method divides the one or more consecutive bytes into a plurality of non-overlapping, consecutive segments;
15 generates a segment feature for each of the plurality of non-overlapping, consecutive segments through application of a plurality of cryptographic hash functions with different random seeds; generates a single duplication feature based on a combination of the segment features for the plurality of non-overlapping, consecutive segments; and generates a single compressed string when a predetermined condition is met, based on a
20 comparison of the single duplication feature and a set of stored duplication features. The method continues with sending the resulting strings in place of the packet toward the packet's destination.

A non-transitory machine-readable medium for removing redundancy in packets is disclosed. The non-transitory machine-readable medium has instructions stored
25 therein, which when executed by a processor, cause the processor to perform operations in a network device, the operations including identifying sets of one or more consecutive bytes within a packet. For each of the sets of one or more consecutive bytes within the packet, the operations include dividing the one or more consecutive bytes into a plurality of non-overlapping, consecutive segments; generating a segment feature
30 for each of the plurality of non-overlapping, consecutive segments through application of a plurality of cryptographic hash functions with different random seeds; generating a

single duplication feature based on a combination of the segment features for the plurality of non-overlapping, consecutive segments; and generating a single compressed string when a predetermined condition is met, based on a comparison of the single duplication feature and a set of stored duplication features. The operations
5 also include sending the resulting strings in place of the packet toward the packet's destination.

Embodiments of the disclosed techniques provide ways for compressing a packet even when the packet contains segments that are not identical to segments in the past; thus, an approximate matching is sufficient for packet compression.

10 BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like references indicate similar elements.

Figure 1 is a block diagram illustrating operations of packet redundancy
15 removal through approximate matching according to one embodiment of the invention.

Figure 2A is a block diagram illustrating the generation of a single duplication feature based on a combination of segment features for a plurality of segments according to one embodiment of the invention.

Figure 2B is a block diagram illustrating the generation of an exemplary single
20 duplication feature based on a combination of segment features for a plurality of segments according to one embodiment of the invention.

Figure 3A is a block diagram illustrating the generation of a single encoded string based on a single duplication feature according to one embodiment of the invention.

Figure 3B is a block diagram illustrating the generation of an exemplary single
25 encoded string based on a single duplication feature according to one embodiment of the invention.

Figure 4 is a flow diagram illustrating operations of packet redundancy removal according to one embodiment of the invention.

Figure 5 is a flow diagram illustrating the generation of a single compressed string according to one embodiment of the invention.

Figure 6A illustrates connectivity between network devices (NDs) within an exemplary network, as well as three exemplary implementations of the NDs, according to some embodiments of the invention.

Figure 6B illustrates an exemplary way to implement a special-purpose network device according to some embodiments of the invention.

Figure 6C illustrates various exemplary ways in which virtual network elements (VNEs) may be coupled according to some embodiments of the invention.

Figure 6D illustrates a network with a single network element (NE) on each of the NDs, and within this straight forward approach contrasts a traditional distributed approach (commonly used by traditional routers) with a centralized approach for maintaining reachability and forwarding information (also called network control), according to some embodiments of the invention.

Figure 6E illustrates the simple case of where each of the NDs implements a single NE, but a centralized control plane has abstracted multiple of the NEs in different NDs into (to represent) a single NE in one of the virtual network(s), according to some embodiments of the invention.

Figure 6F illustrates a case where multiple VNEs are implemented on different NDs and are coupled to each other, and where a centralized control plane has abstracted these multiple VNEs such that they appear as a single VNE within one of the virtual networks, according to some embodiments of the invention

DETAILED DESCRIPTION

In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description. It will be appreciated, however, by one skilled in the art that the invention may be practiced without such specific details. Those of ordinary skill in the art, with the

included descriptions, will be able to implement appropriate functionality without undue experimentation.

References in the specification to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

Bracketed text and blocks with dashed borders (e.g., large dashes, small dashes, dot-dash, and dots) may be used herein to illustrate optional operations that add additional features to embodiments of the invention. However, such notation should not be taken to mean that these are the only options or optional operations, and/or that blocks with solid borders are not optional in certain embodiments of the invention.

In the following description and claims, the terms “coupled” and “connected,” along with their derivatives, may be used. It should be understood that these terms are not intended as synonyms for each other. “Coupled” is used to indicate that two or more elements, which may or may not be in direct physical or electrical contact with each other, co-operate or interact with each other. “Connected” is used to indicate the establishment of communication between two or more elements that are coupled with each other. A “set,” as used herein refers to any positive whole number of items including one item.

An electronic device stores and transmits (internally and/or with other electronic devices over a network) code (which is composed of software instructions and which is sometimes referred to as computer program code or a computer program) and/or data using machine-readable media (also called computer-readable media), such as machine-readable storage media (e.g., magnetic disks, optical disks, read only memory (ROM), flash memory devices, phase change memory) and machine-readable transmission media (also called a carrier) (e.g., electrical, optical, radio, acoustical or other form of

propagated signals – such as carrier waves, infrared signals). Thus, an electronic device (e.g., a computer) includes hardware and software, such as a set of one or more processors coupled to one or more machine-readable storage media to store code for execution on the set of processors and/or to store data. For instance, an electronic

5 device may include non-volatile memory containing the code since the non-volatile memory can persist code/data even when the electronic device is turned off (when power is removed), and while the electronic device is turned on that part of the code that is to be executed by the processor(s) of that electronic device is typically copied from the slower non-volatile memory into volatile memory (e.g., dynamic random

10 access memory (DRAM), static random access memory (SRAM)) of that electronic device. Typical electronic devices also include a set or one or more physical network interface(s) to establish network connections (to transmit and/or receive code and/or data using propagating signals) with other electronic devices. A network device is an electronic device. A network element, as explained in more details herein below, is

15 implemented in one or more network devices.

Removing Redundancy Through Exact Matching And Its Drawbacks

For removing traffic redundancy in a network, one solution examines packet redundancy. It relies on finding exact matches between substrings in the current packet and substrings from recently processed packets that are stored in a cache. Instead of

20 sending the current substring, a token pointing to the matched substring from the recently processed packets is sent. The cache at a transmitting network element is synchronized with the cache at a receiving network element, and the token is used to restore the replaced substring based on the substrings from the recently processed packets at the receiving network element.

25 Since matching of substrings directly can be computationally demanding due to their long lengths, fingerprints of the substrings may be used for matching instead of the substrings themselves. A fingerprint of a substring is formed by mapping the substring to a much shorter bit string through a fingerprinting algorithm. The fingerprint aims at efficiently and uniquely identifying the original substring for all

30 practical purposes.

One way to remove traffic redundancy is to remove traffic redundancy sent to network links. The so called link redundancy elimination (sometimes referred to as the link RE) method can be run for each packet. A cache is used to hold the most recent packets, and an incoming packet is checked against the cache for redundancy. For every incoming packet, the method first generates a representative set of fingerprints and each fingerprint in the set is checked against the fingerprints stored in the cache. If an exact match is found, then a packet in the cache has the same content as the incoming packet in the packet region corresponding to the fingerprint. The matching region may then be expanded to find the largest matching region between the cached packet and the incoming packet, and the total repeated content is the union of all the extended matching regions corresponding to all the fingerprints. Instead of sending a matching region, a corresponding token is sent, and the token points to the region of the incoming packet. Then the cache may be updated by inserting the newly processed packet and the corresponding fingerprints. Through sending the tokens over the network links instead of the packets themselves, the traffic redundancy over the network links is removed.

Another way to remove traffic redundancy is to remove redundancy over the network as a whole. The so called network redundancy elimination (sometimes referred to as the network RE) method is designed for network wide redundancy removal. A packet can potentially be reconstructed or decoded several hops downstream from the location where it is compressed or encoded, which is a significant departure from the link RE where each compressed packet is reconstructed at the immediate downstream network element. A network RE may use a network-wide coordinated approach for intelligently allocating encoding and decoding responsibilities across network elements. A network RE may contain ingress network elements, which search for redundant content in incoming packets and encode them using previously seen packets. The network RE may also contain interior network elements, which store subsets of packets and perform decoding but not encoding. With the reduced capability, the interior network elements save processing resources.

These approaches look for substrings that have occurred exactly in the recent past. In case that an exact match is not found, no compression is performed for a

substring. However, redundancy removal is possible even in the case of an approximate match. In the case of an approximate match, the difference between an incoming substring and a cached substring can be encoded and the result may be smaller than the original substring, thus the process increases the transmission efficiency. For example, 5 say a packet contains a substring of A73FDFFF05275C342C and the “closest” substring in the cache is A73FDFFF4B275C342C, thus the substring in the cache is different from the substring in the packet only in the underlined fifth byte. Thus, the latter substring can be used to describe the former substring in a compressed manner using a representation including, for example, the following: (1) a pointer to the closest 10 substring in the cache (e.g., pointing to the substring or fingerprint of the substring of A73FDFFF4B275C342C); (2) the location of the difference (e.g., the fifth byte); and the value of the difference (e.g., value 05).

Removing Redundancy Through Approximate Matching

Figure 1 is a block diagram illustrating operations of packet redundancy 15 removal through approximate matching according to one embodiment of the invention. Network element 102 is a network element that receives a packet and compresses the packet prior to sending the packet towards its destination. Task boxes 1-7 illustrate the order in which operations of approximate matching are performed according to one embodiment of the invention.

20 At task box 1, a packet marker 112 identifies sets of one or more consecutive bytes within a packet. Each of the sets, say M sets of one or more consecutive bytes, is a candidate for compression. Packet marker 112 identifies the positions of the bytes to divide the packet. In one embodiment, packet marker 112 divides the whole packet into the M sets, thus the whole packet is considered for compression. In another 25 embodiment, packet marker 112 does not mark bytes in the header of the packet, and only the packet payload is considered for compression. The rationale is that the packet header may be needed for packet forwarding at a downstream network element, and compression of the packet header would disrupt the capability of the downstream network element to decode the packet. On the other hand, in some networks (e.g., 30 point-to-point connection with no routing needed), the information in the packet header is not needed for downstream network elements to decode the packet, in which case,

the whole packet may be a candidate for compression. Each or some of the M sets may have a different byte length in one embodiment. In an alternative embodiment, all the M sets have the same byte length, in which case, the byte length is denoted here to be N bytes. A byte (also referred to as an octet) contains eight bits, and two entities having
5 the same byte length means that the two entities have the same bit length too. It is to be noted that the operations in task boxes 2-6 below are repeated for each of the sets of one or more consecutive bytes within the packet.

At task box 2, a string divider 114 divides each of the sets of one or more consecutive bytes into a plurality of non-overlapping, consecutive segments. Each of
10 the segments may also be referred to as a substring of the packet. Each byte in the set of one or more consecutive bytes belongs to one and only one of the plurality of non-overlapping, consecutive segments. The segments may have different byte lengths in one embodiment. In an alternative embodiment, all segments of a set have the same byte length, in which case, the byte length is denoted here to be k bytes.

At task box 3, a function cluster 150 generates a segment feature for each of the plurality of non-overlapping, consecutive segments of a set of one or more consecutive bytes. The generation of the segment features is through application of a plurality of cryptographic hash functions with different random seeds. Functions 152-154 are the plurality of cryptographic hash functions with different random seeds. The
15 cryptographic hash functions are designed to take a segment and produce a fixed length hash value. The cryptographic hash functions are designed to behave as much as possible like a random function while still being deterministic and efficiently computable. Thus, a pseudo random number generator can be used, and different random seeds can be applied for these cryptographic hash functions. In one
20 embodiment, a cryptographic hash function with a different random seed is applied to each segment, so that no segment in a set shares the same random seed of a cryptographic hash function with another segment in the set.
25

One advantage of applying a plurality of cryptographic hash functions with different random seeds is that if a cached string is similar to the set of one or more
30 consecutive bytes, but the byte order of the cached string is different from the order of the segments of the set, the cached string and the set are not considered similar as the

orders of the segments cannot be shuffled in packet transmission. Using an identical hash function with the same random seed may make the hash values of an “out-of-order” cached string (with regard to segments of a set of one or more consecutive bytes) and the set appear similar, and the “false positive” is not desirable for packet
5 redundancy removal.

At task box 4, the resulting plurality of segment features from function cluster 150 are sent to a duplication feature generator 116, which generates a single duplication feature based on the combination of the segment features for the plurality of segments. The single duplication feature is used to determine the redundancy of the set of one or
10 more consecutive bytes. The determination is made at a comparator 120.

At task box 5, the comparator 120 compares the single duplication feature to a set of stored duplication features and determines the difference. The result of the comparison is sent to a string encoder 122. Based on the comparison of the single duplication feature and the set of stored duplication features, string encoder 122
15 generates a single compressed string when a predetermined condition is met at task box 6. In one embodiment, the predetermined condition is that a number of bit value differences of corresponding bits between the single duplication feature and a stored duplication feature is the smallest within the stored duplication features and no more than a predetermined number. The stored duplication feature may then be deemed to be
20 close enough to represent the single duplication feature of the set of one or more consecutive bytes. In one embodiment, the set of stored duplication features are fingerprints of a set of stored strings. The corresponding set of stored strings may be stored in a cache 118.

If the predetermined condition cannot be met by the comparison of the single
25 duplication feature with all of the set of stored duplication features, the set of one or more consecutive bytes (which is a byte string) is not compressed. Thus, for each of the set of one or more consecutive bytes, either the string encoder 122 does not perform compression (thus, the set is intact), or the string encoder 122 generates a single compressed string. The compressed string includes a representation of the selected
30 stored duplication feature (i.e., the stored duplication feature with the smallest difference with respect to the single duplication feature). If the selected stored

duplication feature is not exactly the same as the single duplication feature (i.e., not an exact match), the compressed string further includes the difference between the single duplication feature and the selected stored duplication feature. Thus, through task boxes 2-6, the packet is compressed to a set of strings – strings of uncompressed one or more
5 consecutive bytes and compressed strings. Then at task box 7, instead of the packet itself, the resulting strings for all the sets of one or more consecutive bytes, compressed or not, are sent toward to the destination of the packet. At the destination, since the cache of the sending network element (network element 102 in this example) is synchronized with the cache of the destination network element, the destination
10 network element may decode the compressed strings to its original form, and restore the packet.

Figure 2A is a block diagram illustrating the generation of a single duplication feature based on a combination of segment features for a plurality of segments according to one embodiment of the invention. The operations in **Figure 2A** are an
15 embodiment of implementation of task boxes 1-4 of **Figure 1**.

Referring to **Figure 2A**, packet 201 contains a header 202 and a payload 204. The packet is marked for compression, and the marking identifies sets of one or more consecutive bytes within the packet. Each of the sets of one or more consecutive bytes is a set of N consecutive bytes, where N is an integer. The sets of N consecutive bytes
20 are in the payload 204 in this example, while an alternative embodiment may have the sets of N consecutive bytes including both the header 202 and the payload 204.

Each of the sets of consecutive N bytes is divided into a plurality of non-overlapping, consecutive segments. As illustrated at reference 220, the N bytes at reference 212 are divided into a plurality of k byte segments, and the segments are non-
25 overlapping (e.g., the last byte of the first segment, X_k , is followed by X_{k+1} , the first byte of the second segment). In this example, these segments are byte strings of equal bit length. Each segment has applied to it a function; thus, functions $function_l()$ to $function_{N/k}()$ are applied to respective segments. The functions are different functions such as cryptographic hash functions with different random seeds as illustrated at
30 reference 222.

The application of the functions results in a segment feature for each segment. For example, the application of hash functions results in hash values of a fixed number of bits. The fixed number of bits is F bits in this example, and the bit values are represented by Y , which is either one or zero. With all the segment features being of
5 equal length in this example, the bit positions of all the segment features may be aligned respectively, and the combination of the values of each bit position may result in a value of the single duplication feature. In one embodiment, each of the bit positions of the single duplication feature is set to be one if a majority of the values at corresponding bit positions at the segment features are one; otherwise the bit positions
10 are set to be zero. D_1 to D_F form the resulting single duplication feature as illustrated at reference 224.

Figure 2B is a block diagram illustrating the generation of an exemplary single duplication feature based on a combination of segment features for a plurality of segments according to one embodiment of the invention. The segments in **Figure 2B**
15 are embodiments of segments of the N consecutive byte string of equal length of **Figure 2A**.

The exemplary segments 252 are non-overlapping, consecutive segments forming a complete string of N consecutive bytes. In this example, each segment of the N consecutive bytes contains the same number of bytes, and each byte contains eight
20 bits expressed in a hexadecimal format. Through hashing using a plurality of cryptographic hash functions with different random seeds, each segment results in a fixed-length hash value (in bits) referred to as a segment feature. The segment features are aligned according to the relative bit positions. Thus, all the first bits of the segment features are in the first column in reference 254, all the second bits of the segment
25 features are in the second column, and so on. Then all the bits in the same column are combined to generate a value of the bit position of the single duplication feature. The generation of the value is based on the majority of the bit values in the column in this example. If the majority of the bits are ones, as illustrated in the first two columns at reference 254, the bit position value is one. Otherwise, the bit position value is zero. As
30 illustrated at reference 254, $D_1 = 1$, $D_2 = 1$, $D_3 = 0$, and $D_F = 0$.

Figure 3A is a block diagram illustrating the generation of a single encoded string based on a single duplication feature according to one embodiment of the invention. The operations in **Figure 3A** are an embodiment of implementation of task boxes 5-6 of **Figure 1**. The single duplication feature **D** at reference 224 is compared to
5 a set of stored duplication features 302. The stored duplication features include SD(1) and SD(2), each having a bit length equal to that of the single duplication feature **D**. The comparison is to find a “closest” stored duplication feature, where the number of bit value differences of corresponding bits between the single duplication feature and the stored duplication feature is the smallest, and the number of bit value difference is
10 no more than a predetermined number at reference 312.

In one embodiment, the bit value difference is determined based on the Hamming distance between the single duplication feature and each of the set of stored duplication features. A Hamming distance between two strings of equal length is the number of positions at which the corresponding bit values are different. Finding the
15 number of bit value differences is one way to find the closest approximation of the single duplication feature, and other approaches are possible. Embodiments of the invention are agnostic to the way to find a stored duplication feature that is the least different from the single duplication feature.

In this example, the bit value difference between the single duplication feature
20 **D** and the stored duplication feature SD(2) is the smallest, and the difference is below a predetermined number of bits; thus, SD(2) is selected. SD(2) is used to generate a single compressed string that includes the identifier of SD(2), and the difference between **D** and SD(2) at reference 314. It is to be noted that the difference may include only the bit position without specifying the value difference, as the value difference is a
25 bit value, which can only be one or zero. Indicating a bit position being different is sufficient for the receiving network element to know the difference (e.g., if SD(2) at the bit position is one, **D** would be zero at a bit position that is different). It is also to be noted that if there is an exact match between the single duplication feature **D** and a stored duplication feature, then the stored duplication feature, without the difference, is
30 sufficient to represent the single duplication feature **D** and the corresponding one or

more consecutive bytes; thus, in this case, the single compressed string for the set of one or more consecutive bytes is the identifier of the stored duplication feature itself.

The single compressed string may be transmitted toward the destination network element in the place of the corresponding N consecutive bytes. At the destination network element, it identifies $SD(2)$ based on the identifier of $SD(2)$. As the destination network element is synchronized with the transmitting network element with regard to the stored set of duplication features, the destination network element knows the value of $SD(2)$. Based on the value of $SD(2)$ and the transmitted difference between D and $SD(2)$, the destination network element can determine the value of D . From the value of D , the destination network element can restore the compressed N consecutive bytes.

In one embodiment, the set of stored duplication features and their corresponding set of stored strings are stored in a cache at both the source and destination network elements. In an alternative embodiment, only the set of stored strings, fingerprints of which are the set of stored duplication features, are stored in the cache at both the source and destination network elements.

As the single duplication feature D is generated, the transmitting network element may consider that the corresponding N consecutive bytes are likely to be transmitted again. Thus, the transmitting network element may store the single duplication feature D , replacing a least used duplication feature in the set of stored duplication features at reference 316. In one embodiment, the corresponding stored string of the least used duplication feature is replaced by the corresponding N consecutive bytes in the cache of the transmitting network element for transmission.

Figure 3B is a block diagram illustrating the generation of an exemplary single encoded string based on a single duplication feature according to one embodiment of the invention. The duplication feature D in **Figure 3B** is an embodiment of the duplication feature of **Figure 3A**.

Referring to **Figure 3B**, the Hamming distances between the duplication feature D and $SD(1)$, D and $SD(2)$, D and $SD(n)$ are computed at references 342 and 344. As illustrated, the Hamming distance between D and $SD(1)$ is equal to two, where the bit

positions that are different are underlined in **Figure 3B**. The Hamming distance between D and SD(2) is equal to one, and the bit position, bit 0, is also underlined in **Figure 3B**.

If there is a stored duplication feature whose Hamming distance to D is zero, the stored duplication feature SD matches D exactly, in which case the SD will be the single encoded string, and the single encoded string may include only the identifier of the exactly matching SD without the difference field (i.e., as there is no difference). In the example of **Figure 3B**, it is assumed that there is no exact match of the duplication feature D in the set of stored duplication features. Since the Hamming distance of one for SD(2) is no more than that of the Hamming distance of other stored duplication features (since no exact match is assumed, the Hamming distances for the stored duplication features are larger than zero), and assuming the predetermined condition is that the Hamming distance cannot be more than three, for example, which SD(2) meets, SD(2) is selected to generate a single encoded string representing the duplication feature D at reference 364. The single encoded string includes the identifier of SD(2) and bit zero; the former of the two-tuple indicates the approximation of the duplication feature and the latter of the two-tuple indicates the bit position of the difference between the approximation and the duplication feature D.

At reference 366, the duplication feature D is stored as a new member of the set of stored duplication features SD, replacing a member of the set of existing stored duplication features SD, which is the least utilized in compression. The corresponding set of stored strings of the set of stored duplication features SD are updated too, similarly, to the stored duplication features SD.

Flow Diagrams

Figure 4 is a flow diagram illustrating operations of packet redundancy removal according to one embodiment of the invention. Method 400 may be implemented in a network device implementing a network element such as network element 102 of **Figure 1**.

Method 400 starts at reference 402, where sets of one or more consecutive bytes within a packet are identified. The sets of one or more consecutive bytes may include

all the bytes of the packet in one embodiment. In an alternative embodiment, the sets of one or more consecutive bytes may include only the packet payload. In one embodiment, the sets of one or more consecutive bytes are equal in bit length.

At reference 404, a set of one or more consecutive bytes not being processed yet
5 (“unprocessed” one or more consecutive bytes) is selected. The set of one or more consecutive bytes are divided into a plurality of non-overlapping, consecutive segments at reference 406. The plurality of non-overlapping, consecutive segments are equal in bit length in one embodiment.

Then, at reference 408, for each of the plurality of non-overlapping, consecutive
10 segments, a segment feature is generated through application of a plurality of cryptographic hash functions with different random seeds. In one embodiment, a cryptographic hash function with a different random seed is applied to each of the plurality of non-overlapping, consecutive segments.

At reference 410, a single duplication feature is generated based on a
15 combination of the segment features for the plurality of non-overlapping, consecutive features. In one embodiment, the combination of the segment features for the plurality of non-overlapping, consecutive segments comprises setting a value of each bit position of the single duplication feature based on values at corresponding bit positions of each of the segment features. The value of each bit position of the single duplication feature
20 may be set to be one upon determining that a majority of the values at corresponding bit positions of the segment features are one, and the value of each bit position may be set to be zero otherwise in one embodiment.

At reference 414, a single compressed string is generated based on a
comparison of the single duplication feature and a set of stored duplication features,
25 when a predetermined condition is met. The operation is to find a closest approximation of the single duplication feature within the set of stored duplication features, and when one is found, it is used to represent the set of one or more consecutive bytes. In one embodiment, each single compressed string is surrounded by a character sequence (e.g., an escape character at each end of the string).

At reference 416, it is determined whether or not all the sets of one or more consecutive bytes are processed, and if they are, the resulting strings are sent toward the destination of the packet in place of the packet at reference 418. Otherwise, if it is determined that not all the sets of one or more consecutive bytes are processed, the
5 method 400 returns to reference 404 to select the next set of unprocessed one or more consecutive bytes.

Figure 5 is a flow diagram illustrating the generation of a single compressed string according to one embodiment of the invention. The single compressed string is generated from a set of one or more consecutive bytes, and it is generated based on a
10 generated single duplication feature of the set of one or more consecutive bytes. Method 500 may be an implementation of reference 414 in one embodiment.

At reference 502, the single duplication feature is compared to a set of stored duplication features. Each of the set of stored duplication features corresponds to one stored string of one or more consecutive bytes in a cache of the network element. In
15 one embodiment, the set of stored duplication features are stored in the cache too.

At reference 504, it is determined whether a predetermined condition is met. In one embodiment, the predetermined condition is that the difference between the single duplication feature and a stored duplication feature is no more than that of the single duplication feature and any other stored duplication feature of the set, and the
20 difference is below a predetermined number of bits. If the predetermined condition is not met, the flow goes to reference 510, and the corresponding set of one or more consecutive bytes is maintained intact (i.e., not compressed).

If the predetermined condition is met at reference 504, a single compressed string is generated based on the difference between the single duplication feature and
25 the stored duplication feature with the smallest difference. In one embodiment, the single compressed string includes a representation of the stored duplication feature and the difference. The representation of the stored duplication feature is an identifier of the stored duplication feature in one embodiment.

At reference 508, the single duplication feature is stored to replace a stored
30 duplication feature whose corresponding stored string of one or more consecutive bytes

is the least utilized in the cache. Thus, the set of stored duplication feature is updated, reflecting that the set of one or more consecutive bytes within the packet, corresponding to the single duplication feature, is likely to be used again in the near future.

5 Through methods 400 and/or 500, not only a packet with sets of one or more consecutive bytes that have been transmitted in the recent past may be compressed, the packet with sets of one or more consecutive bytes that are close enough to sets of bytes transmitted in the recent past (thus the sets of bytes are stored in the cache) can also be compressed. Thus, the redundancy removal of approximate matching can be used more
10 widely than the redundancy removal based on exact matching.

SDN and NFV Environment Utilizing Embodiments of the Invention

Embodiments of the invention may be utilized in traditional communication networks, where the data plane and control plane are integrated in a network element implemented by a network device. Embodiments of the invention may also be utilized
15 in a SDN and NFV network containing network devices. A network device (ND) is an electronic device that communicatively interconnects other electronic devices on the network (e.g., other network devices, end-user devices). Some network devices are “multiple services network devices” that provide support for multiple networking functions (e.g., routing, bridging, switching, Layer 2 aggregation, session border
20 control, Quality of Service, and/or subscriber management), and/or provide support for multiple application services (e.g., data, voice, and video).

Figure 6A illustrates connectivity between network devices (NDs) within an exemplary network, as well as three exemplary implementations of the NDs, according to some embodiments of the invention. **Figure 6A** shows NDs 600A-H, and their
25 connectivity by way of lines between A-B, B-C, C-D, D-E, E-F, F-G, and A-G, as well as between H and each of A, C, D, and G. These NDs are physical devices, and the connectivity between these NDs can be wireless or wired (often referred to as a link). An additional line extending from NDs 600A, E, and F illustrates that these NDs act as ingress and egress points for the network (and thus, these NDs are sometimes referred
30 to as edge NDs; while the other NDs may be called core NDs).

Two of the exemplary ND implementations in **Figure 6A** are: 1) a special-purpose network device 602 that uses custom application-specific integrated-circuits (ASICs) and a proprietary operating system (OS); and 2) a general purpose network device 604 that uses common off-the-shelf (COTS) processors and a standard OS.

5 The special-purpose network device 602 includes networking hardware 610 comprising compute resource(s) 612 (which typically include a set of one or more processors), forwarding resource(s) 614 (which typically include one or more ASICs and/or network processors), and physical network interfaces (NIs) 616 (sometimes called physical ports), as well as non-transitory machine readable storage media 618
10 having stored therein networking software 620, which contains compression software 627 including instructions for the approximate matching based redundancy removal as discussed herein above. A physical NI is hardware in a ND through which a network connection (e.g., wirelessly through a wireless network interface controller (WNIC) or through plugging in a cable to a physical port connected to a network interface
15 controller (NIC)) is made, such as those shown by the connectivity between NDs 600A-H. During operation, the compression software 627 may be executed by the networking hardware 610 to instantiate a set of one or more compression software instance 621A-R. Each of the compression software instances 621A-R, and that part of the networking hardware 610 that executes that compression software instance (be it
20 hardware dedicated to that load balancer instance and/or time slices of hardware temporally shared by that compression software instance with others of the networking software instance(s) 622), form a separate virtual network element 630A-R. Each of the virtual network element(s) (VNEs) 630A-R includes a control communication and configuration module 632A-R (sometimes referred to as a local control module or
25 control communication module) and forwarding table(s) 634A-R, such that a given virtual network element (e.g., 630A) includes the control communication and configuration module (e.g., 632A), a set of one or more forwarding table(s) (e.g., 634A), and that portion of the networking hardware 610 that executes the virtual network element (e.g., 630A).

30 The special-purpose network device 602 is often physically and/or logically considered to include: 1) a ND control plane 624 (sometimes referred to as a control

plane) comprising the compute resource(s) 612 that execute the control communication and configuration module(s) 632A-R; and 2) a ND forwarding plane 626 (sometimes referred to as a forwarding plane, a data plane, or a media plane) comprising the forwarding resource(s) 614 that utilize the forwarding table(s) 634A-R and the physical NIs 616. By way of example, where the ND is a router (or is implementing routing functionality), the ND control plane 624 (the compute resource(s) 612 executing the control communication and configuration module(s) 632A-R) is typically responsible for participating in controlling how data (e.g., packets) is to be routed (e.g., the next hop for the data and the outgoing physical NI for that data) and storing that routing information in the forwarding table(s) 634A-R, and the ND forwarding plane 626 is responsible for receiving that data on the physical NIs 616 and forwarding that data out the appropriate ones of the physical NIs 616 based on the forwarding table(s) 634A-R.

Figure 6B illustrates an exemplary way to implement the special-purpose network device 602 according to some embodiments of the invention. **Figure 6B** shows a special-purpose network device including cards 638 (typically hot pluggable). While in some embodiments the cards 638 are of two types (one or more that operate as the ND forwarding plane 626 (sometimes called line cards), and one or more that operate to implement the ND control plane 624 (sometimes called control cards)), alternative embodiments may combine functionality onto a single card and/or include additional card types (e.g., one additional type of card is called a service card, resource card, or multi-application card). A service card can provide specialized processing (e.g., Layer 4 to Layer 7 services (e.g., firewall, Internet Protocol Security (IPsec) (RFC 4301 and 4309), Secure Sockets Layer (SSL) / Transport Layer Security (TLS), Intrusion Detection System (IDS), peer-to-peer (P2P), Voice over IP (VoIP) Session Border Controller, Mobile Wireless Gateways (Gateway General Packet Radio Service (GPRS) Support Node (GGSN), Evolved Packet Core (EPC) Gateway)). By way of example, a service card may be used to terminate IPsec tunnels and execute the attendant authentication and encryption algorithms. These cards are coupled together through one or more interconnect mechanisms illustrated as backplane 636 (e.g., a first full mesh coupling the line cards and a second full mesh coupling all of the cards).

Returning to **Figure 6A**, the general purpose network device 604 includes hardware 640 comprising a set of one or more processor(s) 642 (which are often COTS processors) and network interface controller(s) 644 (NICs; also known as network interface cards) (which include physical NIs 646), as well as non-transitory machine readable storage media 648 having stored therein software 650, which contains compression software 627. During operation, the processor(s) 642 execute the software 650 to instantiate one or more sets of one or more applications 664A-R, including instances of compression software 627. While one embodiment does not implement virtualization, alternative embodiments may use different forms of virtualization – represented by a virtualization layer 654 and software containers 662A-R. For example, one such alternative embodiment implements operating system-level virtualization, in which case the virtualization layer 654 represents the kernel of an operating system (or a shim executing on a base operating system) that allows for the creation of multiple software containers 662A-R that may each be used to execute one of the sets of applications 664A-R. In this embodiment, the multiple software containers 662A-R (also called virtualization engines, virtual private servers, or jails) are each a user space instance (typically a virtual memory space); these user space instances are separate from each other and separate from the kernel space in which the operating system is run; the set of applications running in a given user space, unless explicitly allowed, cannot access the memory of the other processes. Another such alternative embodiment implements full virtualization, in which case: 1) the virtualization layer 654 represents a hypervisor (sometimes referred to as a virtual machine monitor (VMM)) or a hypervisor executing on top of a host operating system; and 2) the software containers 662A-R each represent a tightly isolated form of software container called a virtual machine that is run by the hypervisor and may include a guest operating system. A virtual machine is a software implementation of a physical machine that runs programs as if they were executing on a physical, non-virtualized machine; and applications generally do not know they are running on a virtual machine as opposed to running on a “bare metal” host electronic device, though some systems provide para-virtualization which allows an operating system or application to be aware of the presence of virtualization for optimization purposes.

The instantiation of the one or more sets of one or more applications 664A-R, as well as the virtualization layer 654 and software containers 662A-R if implemented, are collectively referred to as software instance(s) 652. Each set of applications 664A-R, corresponding software container 662A-R if implemented, and that part of the
5 hardware 640 that executes them (be it hardware dedicated to that execution and/or time slices of hardware temporally shared by software containers 662A-R), forms a separate virtual network element(s) 660A-R.

The virtual network element(s) 660A-R perform similar functionality to the virtual network element(s) 630A-R - e.g., similar to the control communication and
10 configuration module(s) 632A and forwarding table(s) 634A (this virtualization of the hardware 640 is sometimes referred to as network function virtualization (NFV)). Thus, NFV may be used to consolidate many network equipment types onto industry standard high volume server hardware, physical switches, and physical storage, which could be located in Data centers, NDs, and customer premise equipment (CPE).
15 However, different embodiments of the invention may implement one or more of the software container(s) 662A-R differently. For example, while embodiments of the invention are illustrated with each software container 662A-R corresponding to one VNE 660A-R, alternative embodiments may implement this correspondence at a finer level granularity (e.g., line card virtual machines virtualize line cards, control card
20 virtual machine virtualize control cards, etc.); it should be understood that the techniques described herein with reference to a correspondence of software containers 662A-R to VNEs also apply to embodiments where such a finer level of granularity is used.

In certain embodiments, the virtualization layer 654 includes a virtual switch
25 that provides similar forwarding services as a physical Ethernet switch. Specifically, this virtual switch forwards traffic between software containers 662A-R and the NIC(s) 644, as well as optionally between the software containers 662A-R; in addition, this virtual switch may enforce network isolation between the VNEs 660A-R that by policy are not permitted to communicate with each other (e.g., by honoring virtual local area
30 networks (VLANs)).

The third exemplary ND implementation in Figure 6A is a hybrid network device 606, which includes both custom ASICs/proprietary OS and COTS processors/standard OS in a single ND or a single card within an ND. In certain embodiments of such a hybrid network device, a platform VM (i.e., a VM that that
5 implements the functionality of the special-purpose network device 602) could provide for para-virtualization to the networking hardware present in the hybrid network device 606.

Regardless of the above exemplary implementations of an ND, when a single one of multiple VNEs implemented by an ND is being considered (e.g., only one of the
10 VNEs is part of a given virtual network) or where only a single VNE is currently being implemented by an ND, the shortened term network element (NE) is sometimes used to refer to that VNE. Also in all of the above exemplary implementations, each of the VNEs (e.g., VNE(s) 630A-R, VNEs 660A-R, and those in the hybrid network device 606) receives data on the physical NIs (e.g., 616, 646) and forwards that data out the
15 appropriate ones of the physical NIs (e.g., 616, 646). For example, a VNE implementing IP router functionality forwards IP packets on the basis of some of the IP header information in the IP packet; where IP header information includes source IP address, destination IP address, source port, destination port (where “source port” and “destination port” refer herein to protocol ports, as opposed to physical ports of a ND),
20 transport protocol (e.g., user datagram protocol (UDP) (RFC 768, 2460, 2675, 4113, and 5405), Transmission Control Protocol (TCP) (RFC 793 and 1180), and differentiated services (DSCP) values (RFC 2474, 2475, 2597, 2983, 3086, 3140, 3246, 3247, 3260, 4594, 5865, 3289, 3290, and 3317).

Figure 6C illustrates various exemplary ways in which VNEs may be coupled
25 according to some embodiments of the invention. Figure 6C shows VNEs 670A.1-670A.P (and optionally VNEs 670A.Q-670A.R) implemented in ND 600A and VNE 670H.1 in ND 600H. In Figure 6C, VNEs 670A.1-P are separate from each other in the sense that they can receive packets from outside ND 600A and forward packets outside of ND 600A; VNE 670A.1 is coupled with VNE 670H.1, and thus they communicate
30 packets between their respective NDs; VNE 670A.2-670A.3 may optionally forward packets between themselves without forwarding them outside of the ND 600A; and

VNE 670A.P may optionally be the first in a chain of VNEs that includes VNE 670A.Q followed by VNE 670A.R (this is sometimes referred to as dynamic service chaining, where each of the VNEs in the series of VNEs provides a different service – e.g., one or more layer 4-7 network services). While Figure 6C illustrates various exemplary
5 relationships between the VNEs, alternative embodiments may support other relationships (e.g., more/fewer VNEs, more/fewer dynamic service chains, multiple different dynamic service chains with some common VNEs and some different VNEs).

The NDs of **Figure 6A**, for example, may form part of the Internet or a private network; and other electronic devices (not shown; such as end user devices including
10 workstations, laptops, netbooks, tablets, palm tops, mobile phones, smartphones, phablets, multimedia phones, Voice Over Internet Protocol (VOIP) phones, terminals, portable media players, GPS units, wearable devices, gaming systems, set-top boxes, Internet enabled household appliances) may be coupled to the network (directly or through other networks such as access networks) to communicate over the network
15 (e.g., the Internet or virtual private networks (VPNs) overlaid on (e.g., tunneled through) the Internet) with each other (directly or through servers) and/or access content and/or services. Such content and/or services are typically provided by one or more servers (not shown) belonging to a service/content provider or one or more end user devices (not shown) participating in a peer-to-peer (P2P) service, and may include,
20 for example, public webpages (e.g., free content, store fronts, search services), private webpages (e.g., username/password accessed webpages providing email services), and/or corporate networks over VPNs. For instance, end user devices may be coupled (e.g., through customer premise equipment coupled to an access network (wired or wirelessly)) to edge NDs, which are coupled (e.g., through one or more core NDs) to
25 other edge NDs, which are coupled to electronic devices acting as servers. However, through compute and storage virtualization, one or more of the electronic devices operating as the NDs in Figure 6A may also host one or more such servers (e.g., in the case of the general purpose network device 604, one or more of the software containers 662A-R may operate as servers; the same would be true for the hybrid network device
30 606; in the case of the special-purpose network device 602, one or more such servers could also be run on a virtualization layer executed by the compute resource(s) 612); in which case the servers are said to be co-located with the VNEs of that ND.

A virtual network is a logical abstraction of a physical network (such as that in Figure 6A) that provides network services (e.g., L2 and/or L3 services). A virtual network can be implemented as an overlay network (sometimes referred to as a network virtualization overlay) that provides network services (e.g., layer 2 (L2, data link layer) and/or layer 3 (L3, network layer) services) over an underlay network (e.g., an L3 network, such as an Internet Protocol (IP) network that uses tunnels (e.g., generic routing encapsulation (GRE), layer 2 tunneling protocol (L2TP), IPSec) to create the overlay network).

A network virtualization edge (NVE) sits at the edge of the underlay network and participates in implementing the network virtualization; the network-facing side of the NVE uses the underlay network to tunnel frames to and from other NVEs; the outward-facing side of the NVE sends and receives data to and from systems outside the network. A virtual network instance (VNI) is a specific instance of a virtual network on a NVE (e.g., a NE/VNE on an ND, a part of a NE/VNE on a ND where that NE/VNE is divided into multiple VNEs through emulation); one or more VNIs can be instantiated on an NVE (e.g., as different VNEs on an ND). A virtual access point (VAP) is a logical connection point on the NVE for connecting external systems to a virtual network; a VAP can be physical or virtual ports identified through logical interface identifiers (e.g., a VLAN ID).

Examples of network services include: 1) an Ethernet LAN emulation service (an Ethernet-based multipoint service similar to an Internet Engineering Task Force (IETF) Multiprotocol Label Switching (MPLS) or Ethernet VPN (EVPN) service) in which external systems are interconnected across the network by a LAN environment over the underlay network (e.g., an NVE provides separate L2 VNIs (virtual switching instances) for different such virtual networks, and L3 (e.g., IP/MPLS) tunneling encapsulation across the underlay network); and 2) a virtualized IP forwarding service (similar to IETF IP VPN (e.g., Border Gateway Protocol (BGP)/MPLS IPVPN RFC 4364) from a service definition perspective) in which external systems are interconnected across the network by an L3 environment over the underlay network (e.g., an NVE provides separate L3 VNIs (forwarding and routing instances) for different such virtual networks, and L3 (e.g., IP/MPLS) tunneling encapsulation across

the underlay network)). Network services may also include quality of service capabilities (e.g., traffic classification marking, traffic conditioning and scheduling), security capabilities (e.g., filters to protect customer premises from network – originated attacks, to avoid malformed route announcements), and management
5 capabilities (e.g., full detection and processing).

Fig. 6D illustrates a network with a single network element on each of the NDs of **Figure 6A**, and within this straight forward approach contrasts a traditional distributed approach (commonly used by traditional routers) with a centralized approach for maintaining reachability and forwarding information (also called network
10 control), according to some embodiments of the invention. Specifically, **Figure 6D** illustrates network elements (NEs) 670A-H with the same connectivity as the NDs 600A-H of **Figure 6A**.

Fig. 6D illustrates a network with a single network element on each of the NDs of **Figure 6A**, and within this straight forward approach contrasts a traditional
15 distributed approach (commonly used by traditional routers) with a centralized approach for maintaining reachability and forwarding information (also called network control), according to some embodiments of the invention. Specifically, **Figure 6D** illustrates network elements (NEs) 670A-H with the same connectivity as the NDs 600A-H of **Figure 6A**.

Figure 6D illustrates a network with a single network element on each of the NDs of **Figure 6A**, and within this straight forward approach contrasts a traditional distributed approach (commonly used by traditional routers) with a centralized approach for maintaining reachability and forwarding information (also called network
20 control), according to some embodiments of the invention. Specifically, **Figure 6D** illustrates network elements (NEs) 670A-H with the same connectivity as the NDs 600A-H of **Figure 6A**.

For example, where the special-purpose network device 602 is used, the control communication and configuration module(s) 632A-R of the ND control plane 624 typically include a reachability and forwarding information module to implement one
30 or more routing protocols (e.g., an exterior gateway protocol such as Border Gateway Protocol (BGP) (RFC 4271), Interior Gateway Protocol(s) (IGP) (e.g., Open Shortest

Path First (OSPF) (RFC 2328 and 5340), Intermediate System to Intermediate System (IS-IS) (RFC 1142), Routing Information Protocol (RIP) (version 1 RFC 1058, version 2 RFC 2453, and next generation RFC 2080)), Label Distribution Protocol (LDP) (RFC 5036), Resource Reservation Protocol (RSVP) (RFC 2205, 2210, 2211, 2212, as well
5 as RSVP-Traffic Engineering (TE): Extensions to RSVP for LSP Tunnels RFC 3209, Generalized Multi-Protocol Label Switching (GMPLS) Signaling RSVP-TE RFC 3473, RFC 3936, 4495, and 4558)) that communicate with other NEs to exchange routes, and then selects those routes based on one or more routing metrics. Thus, the NEs 670A-H (e.g., the compute resource(s) 612 executing the control communication and
10 configuration module(s) 632A-R) perform their responsibility for participating in controlling how data (e.g., packets) is to be routed (e.g., the next hop for the data and the outgoing physical NI for that data) by distributively determining the reachability within the network and calculating their respective forwarding information. Routes and adjacencies are stored in one or more routing structures (e.g., Routing Information Base
15 (RIB), Label Information Base (LIB), one or more adjacency structures) on the ND control plane 624. The ND control plane 624 programs the ND forwarding plane 626 with information (e.g., adjacency and route information) based on the routing structure(s). For example, the ND control plane 624 programs the adjacency and route information into one or more forwarding table(s) 634A-R (e.g., Forwarding
20 Information Base (FIB), Label Forwarding Information Base (LFIB), and one or more adjacency structures) on the ND forwarding plane 626. For layer 2 forwarding, the ND can store one or more bridging tables that are used to forward data based on the layer 2 information in that data. While the above example uses the special-purpose network device 602, the same distributed approach 672 can be implemented on the general
25 purpose network device 604 and the hybrid network device 606. In the distributed approach, the caches of the network elements are synchronized through communications between themselves or communicating with a network management system (NMS).

Figure 6D illustrates that a centralized approach 674 (also known as software
30 defined networking (SDN)) that decouples the system that makes decisions about where traffic is sent from the underlying systems that forwards traffic to the selected destination. The illustrated centralized approach 674 has the responsibility for the

generation of reachability and forwarding information in a centralized control plane 676 (sometimes referred to as a SDN control module, controller, network controller, OpenFlow controller, SDN controller, control plane node, network virtualization authority, or management control entity), and thus the process of neighbor discovery and topology discovery is centralized. The centralized control plane 676 has a south bound interface 682 with a data plane 680 (sometime referred to the infrastructure layer, network forwarding plane, or forwarding plane (which should not be confused with a ND forwarding plane)) that includes the NEs 670A-H (sometimes referred to as switches, forwarding elements, data plane elements, or nodes). The centralized control plane 676 includes a network controller 678, which includes a centralized reachability and forwarding information module 679 that determines the reachability within the network and distributes the forwarding information to the NEs 670A-H of the data plane 680 over the south bound interface 682 (which may use the OpenFlow protocol). Thus, the network intelligence is centralized in the centralized control plane 676 executing on electronic devices that are typically separate from the NDs. The centralized reachability and forwarding information module 679 contains a cache coordination module 675. Cache coordination module 675 coordinates cache synchronization of network elements within the network. The cache content and the identification/content of the duplication features need to be synchronized, and cache coordination module 675 communicates network elements within the network and make it happen.

Where the special-purpose network device 602 is used in the data plane 680, each of the control communication and configuration module(s) 632A-R of the ND control plane 624 typically include a control agent that provides the VNE side of the south bound interface 682. In this case, the ND control plane 624 (the compute resource(s) 612 executing the control communication and configuration module(s) 632A-R) performs its responsibility for participating in controlling how data (e.g., packets) is to be routed (e.g., the next hop for the data and the outgoing physical NI for that data) through the control agent communicating with the centralized control plane 676 to receive the forwarding information (and in some cases, the reachability information) from the centralized reachability and forwarding information module 679 (it should be understood that in some embodiments of the invention, the control

communication and configuration module(s) 632A-R, in addition to communicating with the centralized control plane 676, may also play some role in determining reachability and/or calculating forwarding information – albeit less so than in the case of a distributed approach; such embodiments are generally considered to fall under the centralized approach 674, but may also be considered a hybrid approach).

While the above example uses the special-purpose network device 602, the same centralized approach 674 can be implemented with the general purpose network device 604 (e.g., each of the VNE 660A-R performs its responsibility for controlling how data (e.g., packets) is to be routed (e.g., the next hop for the data and the outgoing physical NI for that data) by communicating with the centralized control plane 676 to receive the forwarding information (and in some cases, the reachability information) from the centralized reachability and forwarding information module 679; it should be understood that in some embodiments of the invention, the VNEs 660A-R, in addition to communicating with the centralized control plane 676, may also play some role in determining reachability and/or calculating forwarding information – albeit less so than in the case of a distributed approach) and the hybrid network device 606. In fact, the use of SDN techniques can enhance the NFV techniques typically used in the general purpose network device 604 or hybrid network device 606 implementations as NFV is able to support SDN by providing an infrastructure upon which the SDN software can be run, and NFV and SDN both aim to make use of commodity server hardware and physical switches.

Figure 6D also shows that the centralized control plane 676 has a north bound interface 684 to an application layer 686, in which resides application(s) 688. The centralized control plane 676 has the ability to form virtual networks 692 (sometimes referred to as a logical forwarding plane, network services, or overlay networks (with the NEs 670A-H of the data plane 680 being the underlay network)) for the application(s) 688. Thus, the centralized control plane 676 maintains a global view of all NDs and configured NEs/VNEs, and it maps the virtual networks to the underlying NDs efficiently (including maintaining these mappings as the physical network changes either through hardware (ND, link, or ND component) failure, addition, or removal).

While **Figure 6D** illustrates the simple case where each of the NDs 600A-H implements a single NE 670A-H, it should be understood that the network control approaches described with reference to **Figure 6D** also work for networks where one or more of the NDs 600A-H implement multiple VNEs (e.g., VNEs 630A-R, VNEs 660A-R, those in the hybrid network device 606). Alternatively or in addition, the network controller 678 may also emulate the implementation of multiple VNEs in a single ND. Specifically, instead of (or in addition to) implementing multiple VNEs in a single ND, the network controller 678 may present the implementation of a VNE/NE in a single ND as multiple VNEs in the virtual networks 692 (all in the same one of the virtual network(s) 692, each in different ones of the virtual network(s) 692, or some combination). For example, the network controller 678 may cause an ND to implement a single VNE (a NE) in the underlay network, and then logically divide up the resources of that NE within the centralized control plane 676 to present different VNEs in the virtual network(s) 692 (where these different VNEs in the overlay networks are sharing the resources of the single VNE/NE implementation on the ND in the underlay network).

On the other hand, **Figures 6E and 6F** respectively illustrate exemplary abstractions of NEs and VNEs that the network controller 678 may present as part of different ones of the virtual networks 692. **Figure 6E** illustrates the simple case of where each of the NDs 600A-H implements a single NE 670A-H (see **Figure 6D**), but the centralized control plane 676 has abstracted multiple of the NEs in different NDs (the NEs 670A-C and G-H) into (to represent) a single NE 670I in one of the virtual network(s) 692 of **Figure 6D**, according to some embodiments of the invention. **Figure 6E** shows that in this virtual network, the NE 670I is coupled to NE 670D and 670F, which are both still coupled to NE 670E.

Figure 6F illustrates a case where multiple VNEs (VNE 670A.1 and VNE 670H.1) are implemented on different NDs (ND 600A and ND 600H) and are coupled to each other, and where the centralized control plane 676 has abstracted these multiple VNEs such that they appear as a single VNE 670T within one of the virtual networks 692 of **Figure 6D**, according to some embodiments of the invention. Thus, the abstraction of a NE or VNE can span multiple NDs.

While some embodiments of the invention implement the centralized control plane 676 as a single entity (e.g., a single instance of software running on a single electronic device), alternative embodiments may spread the functionality across multiple entities for redundancy and/or scalability purposes (e.g., multiple instances of software running on different electronic devices).

Standards such as OpenFlow define the protocols used for the messages, as well as a model for processing the packets. The model for processing packets includes header parsing, packet classification, and making forwarding decisions. Header parsing describes how to interpret a packet based upon a well-known set of protocols. Some protocol fields are used to build a match structure (or key) that will be used in packet classification (e.g., a first key field could be a source media access control (MAC) address, and a second key field could be a destination MAC address).

Packet classification involves executing a lookup in memory to classify the packet by determining which entry (also referred to as a forwarding table entry or flow entry) in the forwarding tables best matches the packet based upon the match structure, or key, of the forwarding table entries. It is possible that many flows represented in the forwarding table entries can correspond/match to a packet; in this case the system is typically configured to determine one forwarding table entry from the many according to a defined scheme (e.g., selecting a first forwarding table entry that is matched). Forwarding table entries include both a specific set of match criteria (a set of values or wildcards, or an indication of what portions of a packet should be compared to a particular value/values/wildcards, as defined by the matching capabilities – for specific fields in the packet header, or for some other packet content), and a set of one or more actions for the data plane to take on receiving a matching packet. For example, an action may be to push a header onto the packet, for the packet using a particular port, flood the packet, or simply drop the packet. Thus, a forwarding table entry for IPv4/IPv6 packets with a particular transmission control protocol (TCP) destination port could contain an action specifying that these packets should be dropped.

Making forwarding decisions and performing actions occurs, based upon the forwarding table entry identified during packet classification, by executing the set of actions identified in the matched forwarding table entry on the packet.

However, when an unknown packet (for example, a “missed packet” or a “match-miss” as used in OpenFlow parlance) arrives at the data plane 680, the packet (or a subset of the packet header and content) is typically forwarded to the centralized control plane 676. The centralized control plane 976 will then program forwarding table entries into the data plane 680 to accommodate packets belonging to the flow of the unknown packet. Once a specific forwarding table entry has been programmed into the data plane 680 by the centralized control plane 676, the next packet with matching credentials will match that forwarding table entry and take the set of actions associated with that matched entry.

10 A network interface (NI) may be physical or virtual; and in the context of IP, an interface address is an IP address assigned to a NI, be it a physical NI or virtual NI. A virtual NI may be associated with a physical NI, with another virtual interface, or stand on its own (e.g., a loopback interface, a point-to-point protocol interface). A NI (physical or virtual) may be numbered (a NI with an IP address) or unnumbered (a NI without an IP address). A loopback interface (and its loopback address) is a specific type of virtual NI (and IP address) of a NE/VNE (physical or virtual) often used for management purposes; where such an IP address is referred to as the nodal loopback address. The IP address(es) assigned to the NI(s) of a ND are referred to as IP addresses of that ND; at a more granular level, the IP address(es) assigned to NI(s) assigned to a NE/VNE implemented on a ND can be referred to as IP addresses of that NE/VNE.

Each VNE (e.g., a virtual router, a virtual bridge (which may act as a virtual switch instance in a Virtual Private LAN Service (VPLS) (RFC 4761 and 4762) is typically independently administrable. For example, in the case of multiple virtual routers, each of the virtual routers may share system resources but is separate from the other virtual routers regarding its management domain, AAA (authentication, authorization, and accounting) name space, IP address, and routing database(s). Multiple VNEs may be employed in an edge ND to provide direct network access and/or different classes of services for subscribers of service and/or content providers.

30 Within certain NDs, “interfaces” that are independent of physical NIs may be configured as part of the VNEs to provide higher-layer protocol and service

information (e.g., Layer 3 addressing). The subscriber records in the AAA server identify, in addition to the other subscriber configuration requirements, to which context (e.g., which of the VNEs/NEs) the corresponding subscribers should be bound within the ND. As used herein, a binding forms an association between a physical
5 entity (e.g., physical NI, channel) or a logical entity (e.g., circuit such as a subscriber circuit or logical circuit (a set of one or more subscriber circuits)) and a context's interface over which network protocols (e.g., routing protocols, bridging protocols) are configured for that context. Subscriber data flows on the physical entity when some higher-layer protocol interface is configured and associated with that physical entity.

10 The operations of the flow diagrams **Figures 4 and 5** are described with reference to the exemplary embodiment of **Figures 1-3 and 6**. However, it should be understood that the operations of flow diagrams can be performed by embodiments of the invention other than those discussed with reference to the exemplary embodiment of **Figures 1-3 and 6**, and the exemplary embodiment of **Figures 1-3 and 6** can
15 perform operations different than those discussed with reference to the flow diagrams of **Figures 4 and 5**.

While the flow diagrams in the figures herein above show a particular order of operations performed by certain embodiments of the invention, it should be understood that such order is exemplary (e.g., alternative embodiments may perform the operations
20 in a different order, combine certain operations, overlap certain operations, etc.).

Different embodiments of the invention may be implemented using different combinations of software, firmware, and/or hardware. Thus, the techniques shown in the figures can be implemented using code and data stored and executed on one or more electronic devices (e.g., an end system, a network device). Such electronic
25 devices store and communicate (internally and/or with other electronic devices over a network) code and data using computer-readable media, such as non-transitory computer-readable storage media (e.g., magnetic disks; optical disks; random access memory; read only memory; flash memory devices; phase-change memory) and transitory computer-readable transmission media (e.g., electrical, optical, acoustical or
30 other form of propagated signals – such as carrier waves, infrared signals, digital signals). In addition, such electronic devices typically include a set of one or more

processors coupled to one or more other components, such as one or more storage devices (non-transitory machine-readable storage media), user input/output devices (e.g., a keyboard, a touchscreen, and/or a display), and network connections. The coupling of the set of processors and other components is typically through one or more
5 busses and bridges (also termed as bus controllers). Thus, the storage device of a given electronic device typically stores code and/or data for execution on the set of one or more processors of that electronic device.

While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments
10 described, can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting.

CLAIMS

What is claimed is:

1. A method implemented in a network device for removing redundancy in packets, the method comprising:
 - 5 identifying (402) sets of one or more consecutive bytes within a packet;
for each of the sets of one or more consecutive bytes within the packet,
dividing (406) the one or more consecutive bytes into a plurality of non-overlapping, consecutive segments;
generating (408) a segment feature for each of the plurality of non-overlapping, consecutive segments through application of a
10 plurality of cryptographic hash functions with different random seeds;
generating (410) a single duplication feature based on a combination of the segment features for the plurality of non-overlapping, consecutive segments; and
15 generating (414) a single compressed string when a predetermined condition is met, based on a comparison of the single duplication feature and a set of stored duplication features; and
sending (418) the resulting strings in place of the packet toward the packet's
20 destination.
2. The method of claim 1, wherein the sets of one or more consecutive bytes are equal in bit length.
- 25 3. The method of claim 1, wherein the plurality of non-overlapping, consecutive segments are equal in bit length.
4. The method of claim 1, wherein the generating the segment feature for each of the plurality of non-overlapping, consecutive segments comprises applying a
30 cryptographic hash function with a different random seed to each of the plurality of non-overlapping, consecutive segments.

5. The method of claim 1, wherein the generating the single duplication feature comprises setting a value of each bit position of the single duplication feature based on values at corresponding bit positions of each of the segment features.
- 5 6. The method of claim 5, wherein the value of each bit position of the single duplication feature is set to be one upon determining that a majority of the values at corresponding bit positions of the segment features are one, and the value of each bit position is set to be zero otherwise.
- 10 7. The method of claim 1, wherein the generating the single compressed string based on the comparison of the single duplication feature and the set of stored duplication features comprises:
comparing (502) the single duplication feature to the set of stored duplication features, each corresponding to one stored string of one or more
15 consecutive bytes in a cache of the network device;
upon determining that the predetermined condition is met, wherein the predetermined condition is that difference between the single duplication feature and a stored duplication feature is no more than that of the single duplication feature and any other stored duplication feature of the set,
20 and the difference is below a predetermined number of bits, generating (506) the single compressed string based on the difference and the stored duplication feature; and
upon determining that the difference between the single duplication feature and
25 each of the set of stored duplication features fails to meet the predetermined condition, maintaining (510) the set of one or more consecutive bytes within the packet intact.
8. The method of claim 7, wherein each single compressed string includes a representation of the stored duplication feature and the difference.
- 30 9. The method of claim 7, further comprising:

storing (508) the single duplication feature to replace a stored duplication feature whose corresponding stored string of one or more consecutive bytes is least utilized in the cache.

- 5 10. The method of claim 7, wherein the set of stored duplication features is stored in the cache.
11. The method of claim 1, wherein each single compressed string is surrounded by a character sequence.
- 10 12. The method of claim 1, wherein all the sets of one or more consecutive bytes are within a payload of the packet.
13. A non-transitory machine-readable medium having instructions stored therein,
15 which when executed by a processor, cause the processor to perform operations in a network device, the operations comprising:
identifying (402) sets of one or more consecutive bytes within a packet;
for each of the sets of one or more consecutive bytes within the packet,
dividing (406) the one or more consecutive bytes into a plurality of non-
20 overlapping, consecutive segments;
generating (408) a segment feature for each of the plurality of non-overlapping, consecutive segments through application of a plurality of cryptographic hash functions with different random seeds;
25 generating (410) a single duplication feature based on a combination of the segment features for the plurality of non-overlapping, consecutive segments; and
generating (414) a single compressed string when a predetermined condition is met, based on a comparison of the single duplication
30 feature and a set of stored duplication features; and
sending (418) the resulting strings in place of the packet toward the packet's destination.

14. The non-transitory machine-readable medium of claim 13, wherein the
generating the segment feature for each of the plurality of non-overlapping,
consecutive segments comprises applying a cryptographic hash function with a
5 different random seed to each of the plurality of non-overlapping, consecutive
segments.
15. The non-transitory machine-readable medium of claim 13, wherein the
generating the single duplication feature comprises setting a value of each bit
10 position of the single duplication feature based on values at corresponding bit
positions at each of the segment features.
16. The non-transitory machine-readable medium of claim 13, wherein the
generating the single compressed string based on the comparison of the single
15 duplication feature and the set of stored duplication features comprises:
comparing (502) the single duplication feature to the set of stored duplication
features, each corresponding to one stored string of one or more
consecutive bytes in a cache of the network device;
upon determining that the predetermined condition is met, wherein the
20 predetermined condition is that difference between the single duplication
feature and a stored duplication feature is no more than that of the single
duplication feature and any other stored duplication feature of the set,
and the difference is below a predetermined number of bits, generating
(506) the single compressed string based on the difference and the stored
25 duplication feature; and
upon determining that the difference between the single duplication feature and
each of the set of stored duplication features fails to meet the
predetermined condition, maintaining (510) the set of one or more
consecutive bytes within the packet intact.
30

17. The non-transitory machine-readable medium of claim 16, wherein each single compressed string includes a representation of the stored duplication feature and the difference.
- 5 18. The non-transitory machine-readable medium of claim 16, the operations further comprising:
storing (508) the single duplication feature to replace a stored duplication feature whose corresponding stored string of one or more consecutive bytes is least utilized in the cache.
- 10
19. The non-transitory machine-readable medium of claim 13, wherein each single compressed string is surrounded by a character sequence.
20. The non-transitory machine-readable medium of claim 13, wherein all the sets
15 of one or more consecutive bytes are within a payload of the packet.

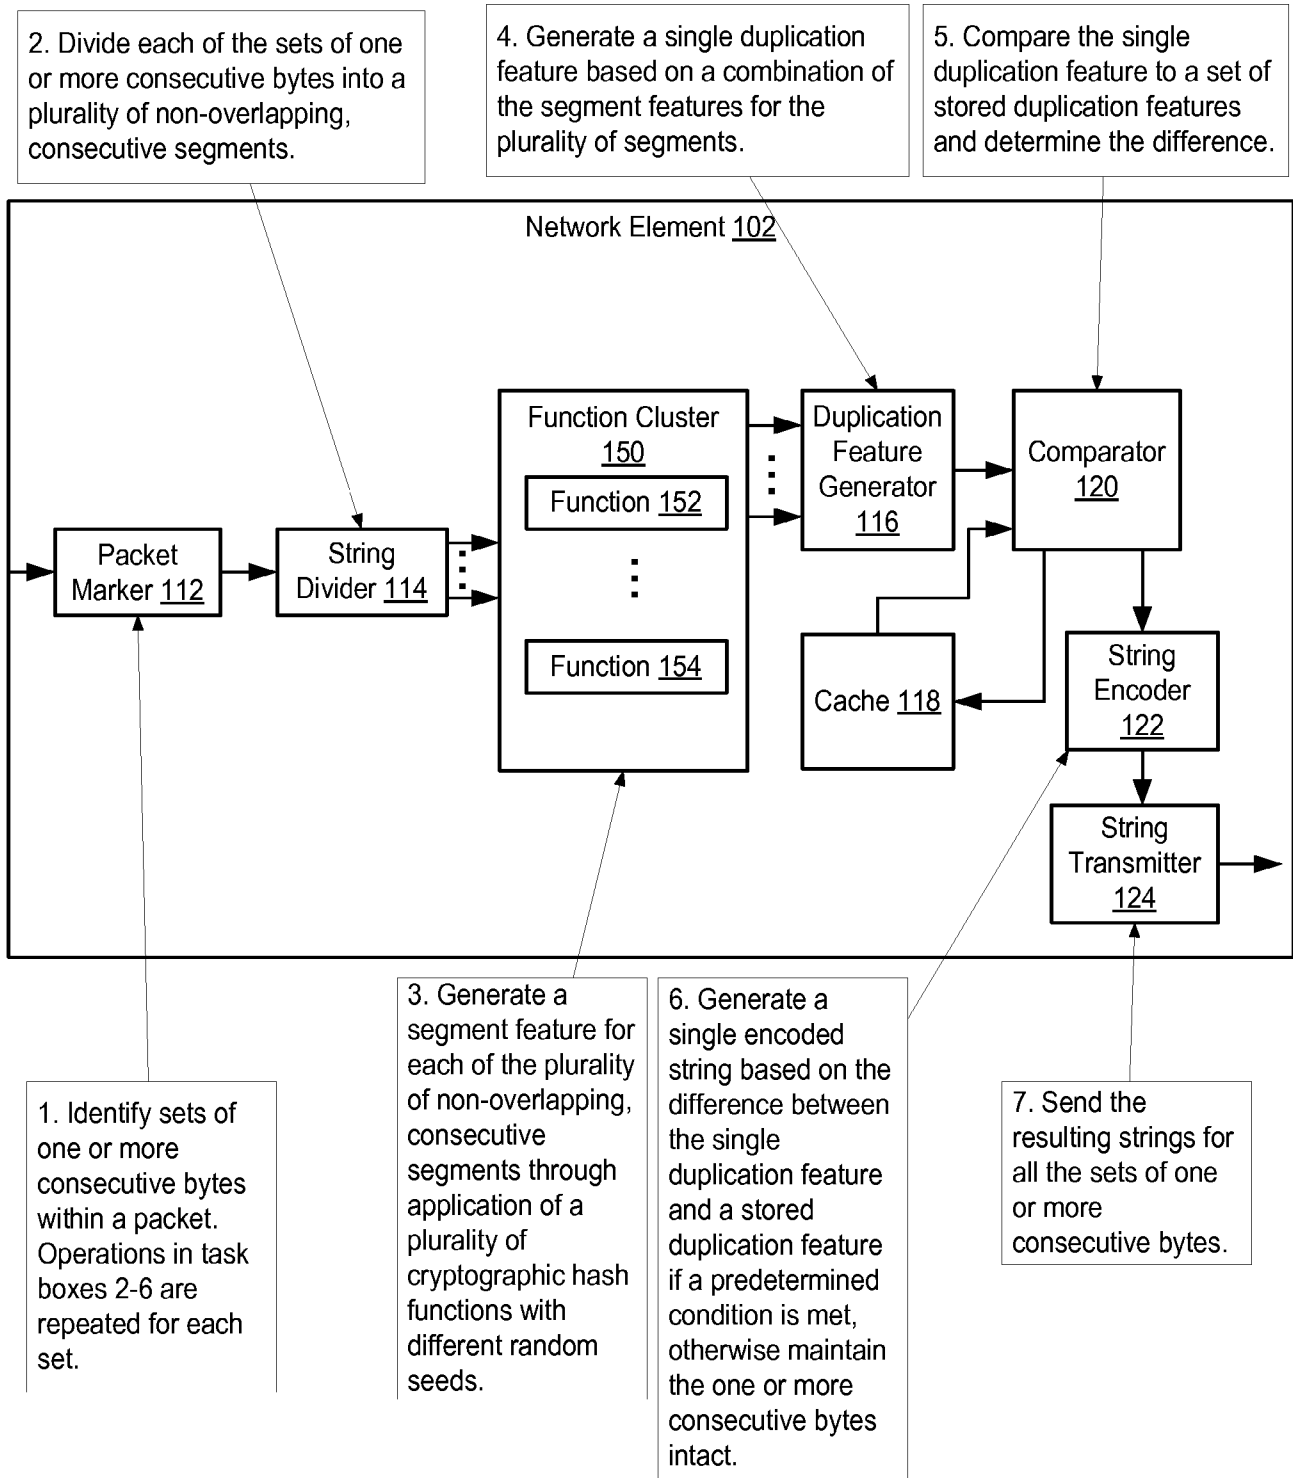


FIG. 1

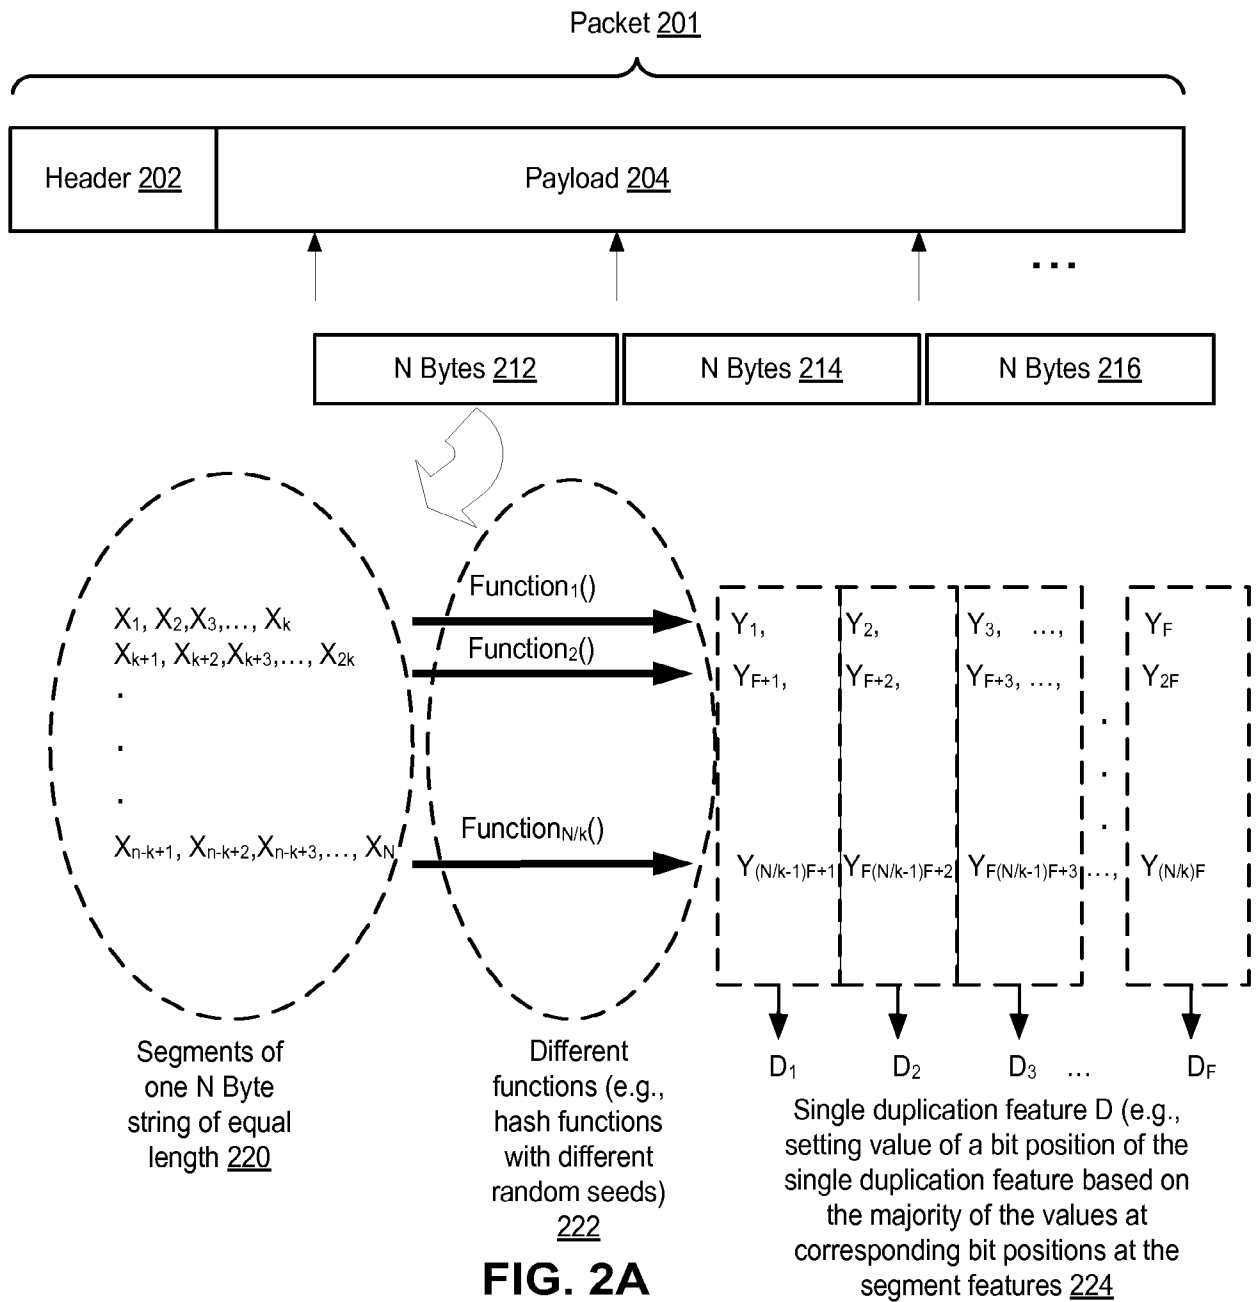


FIG. 2A




FIG. 2B

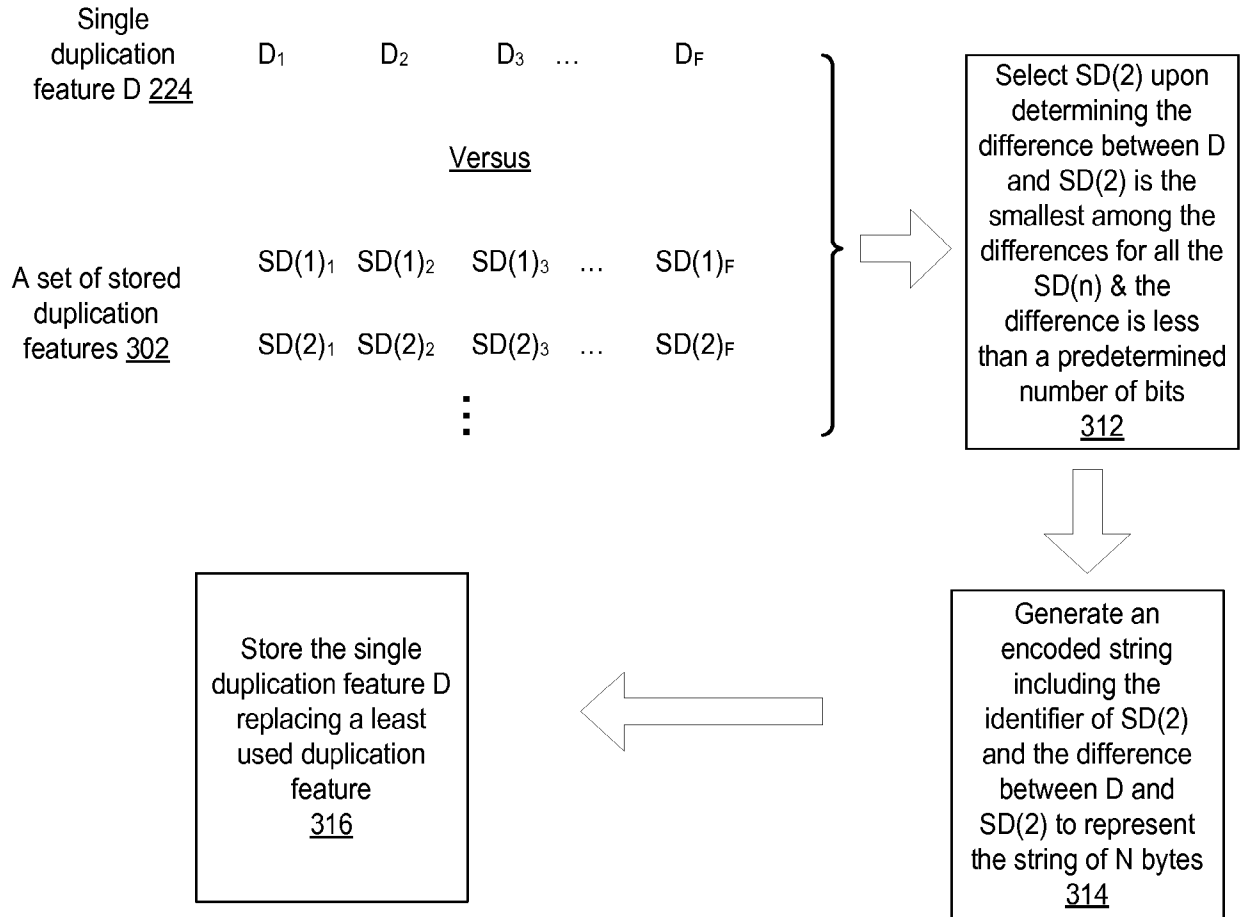


FIG. 3A

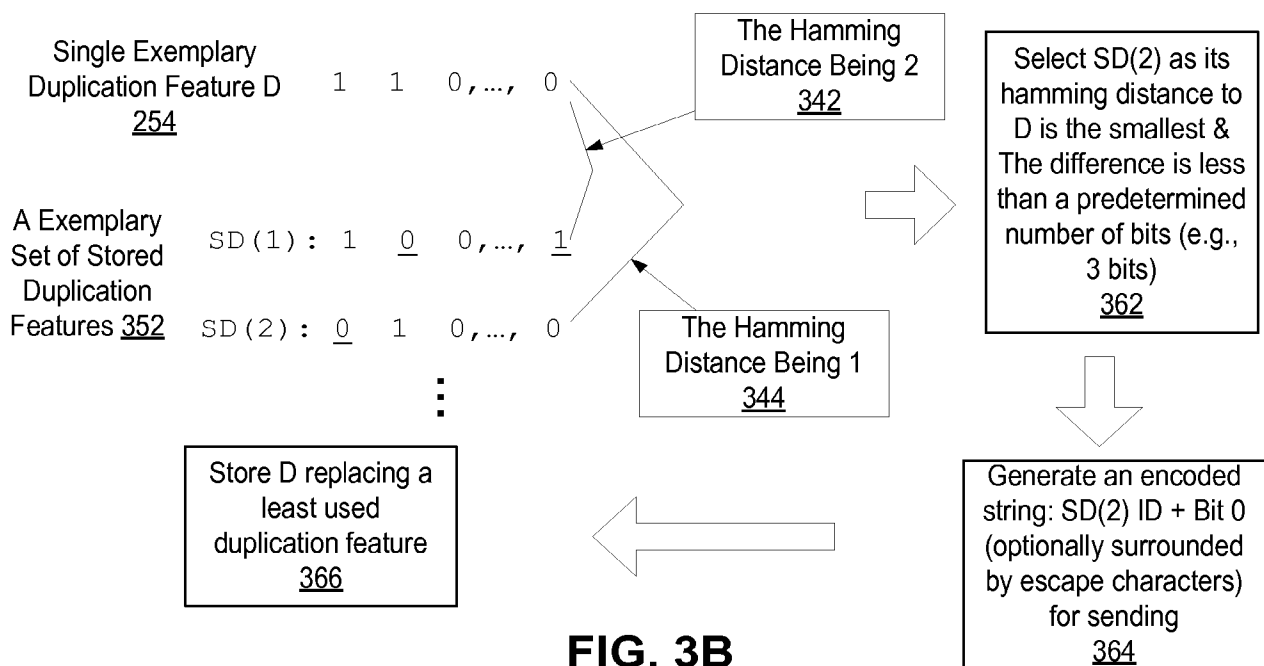


FIG. 3B

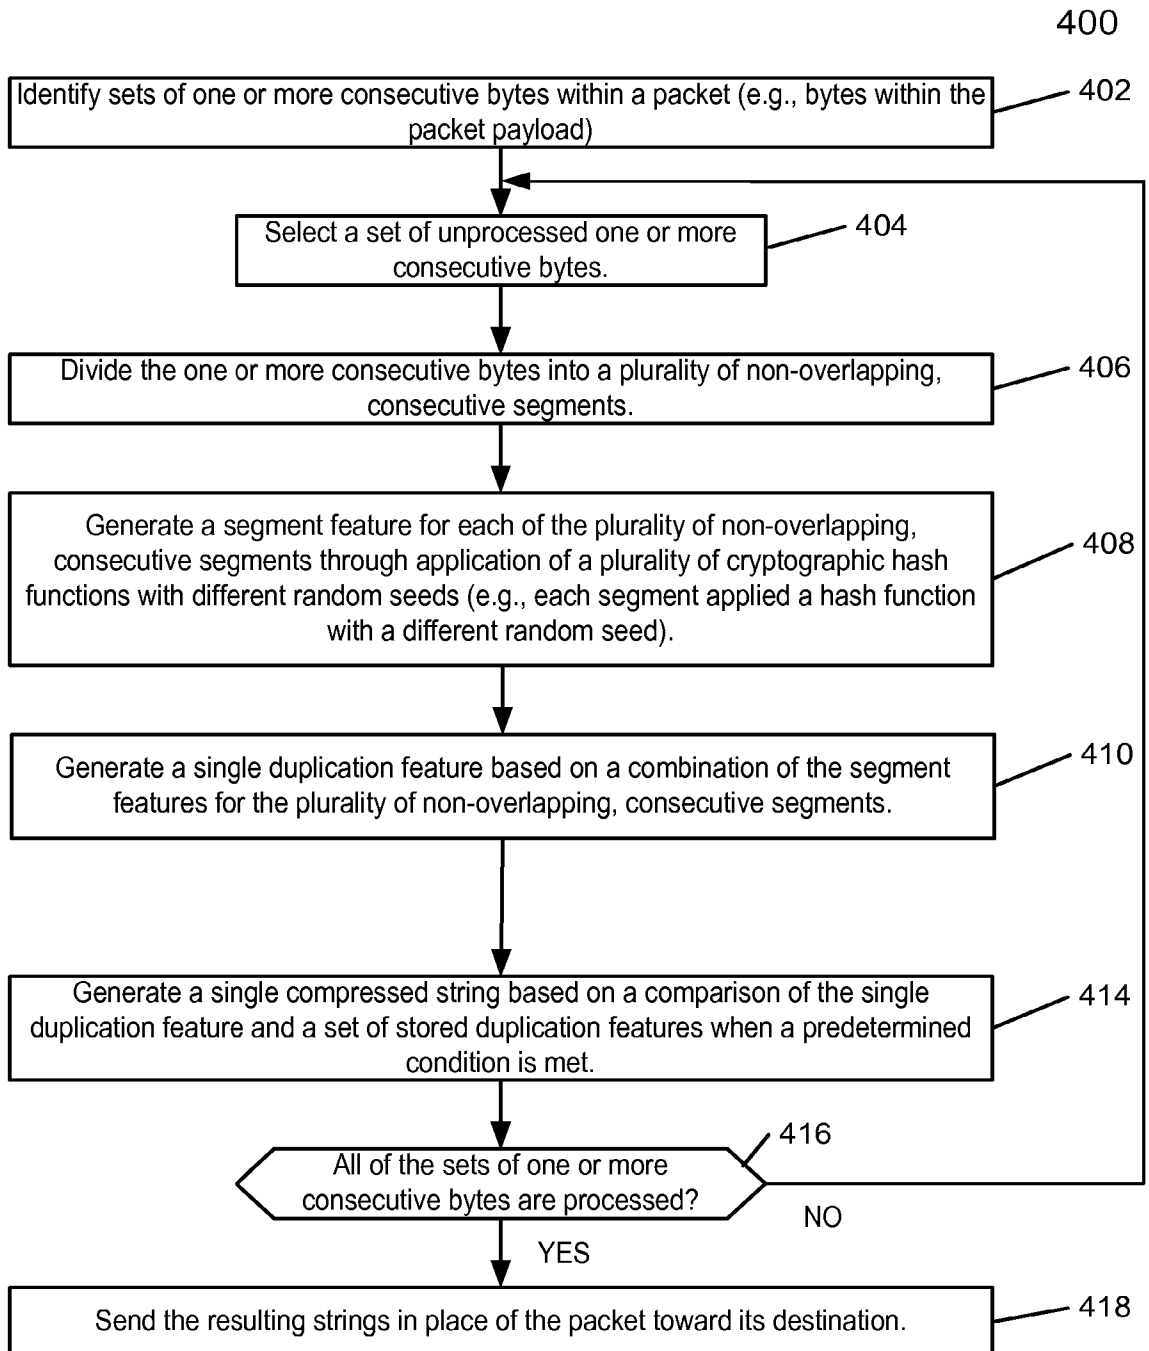


FIG. 4

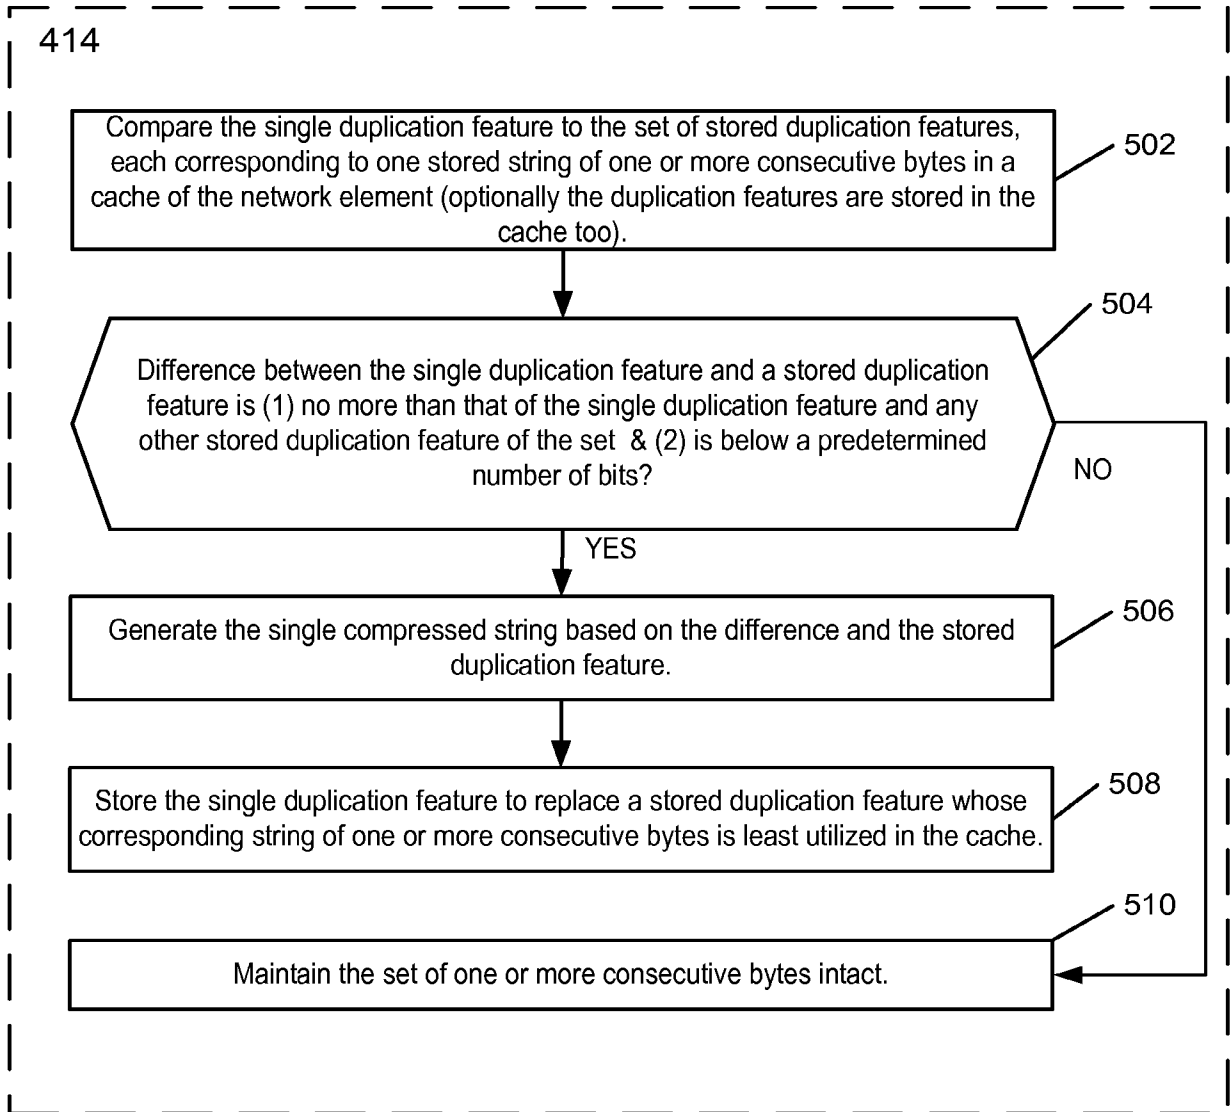


FIG. 5

Fig. 6A

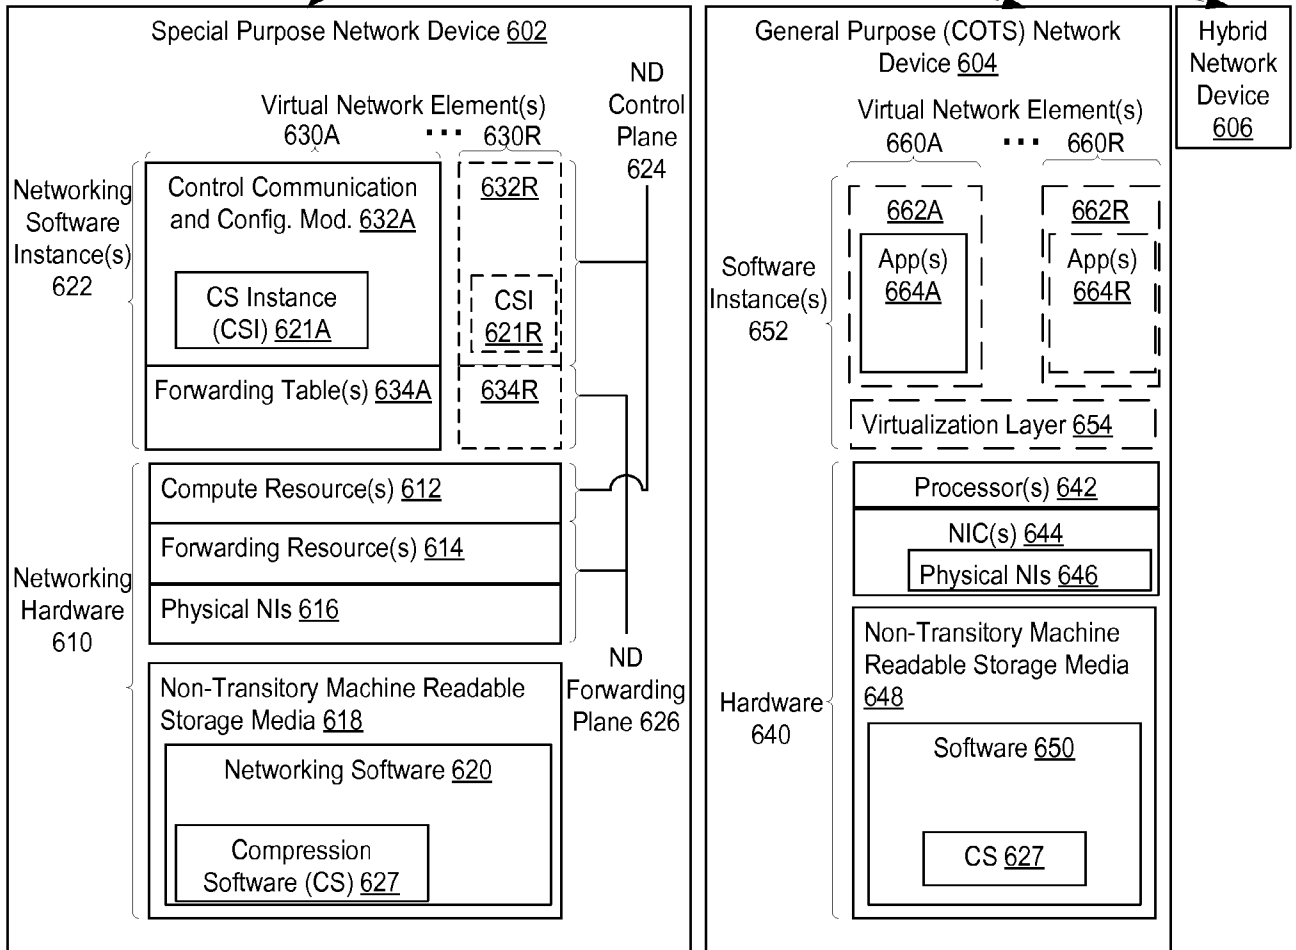
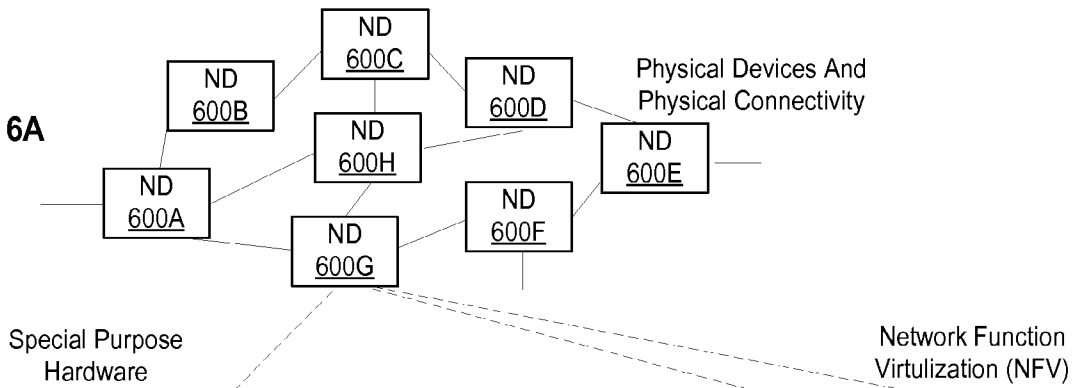



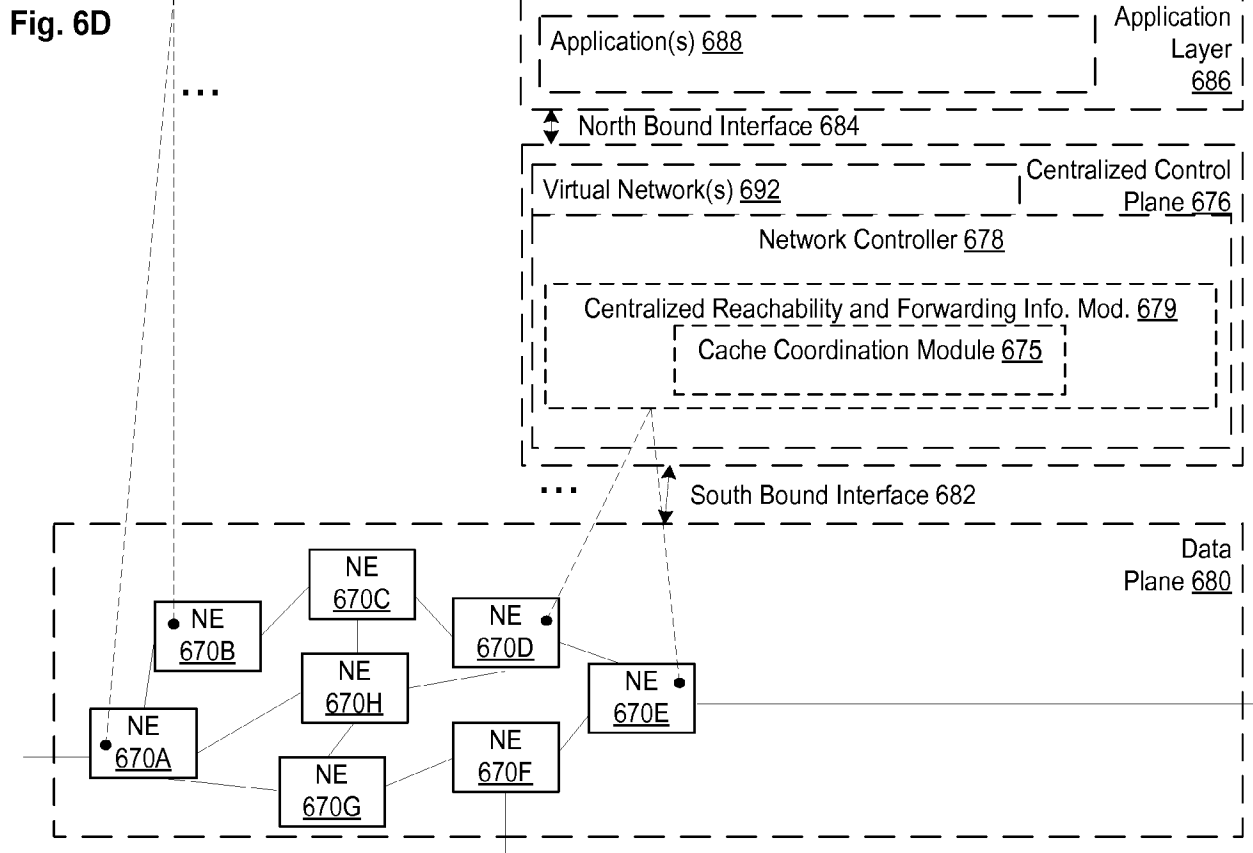
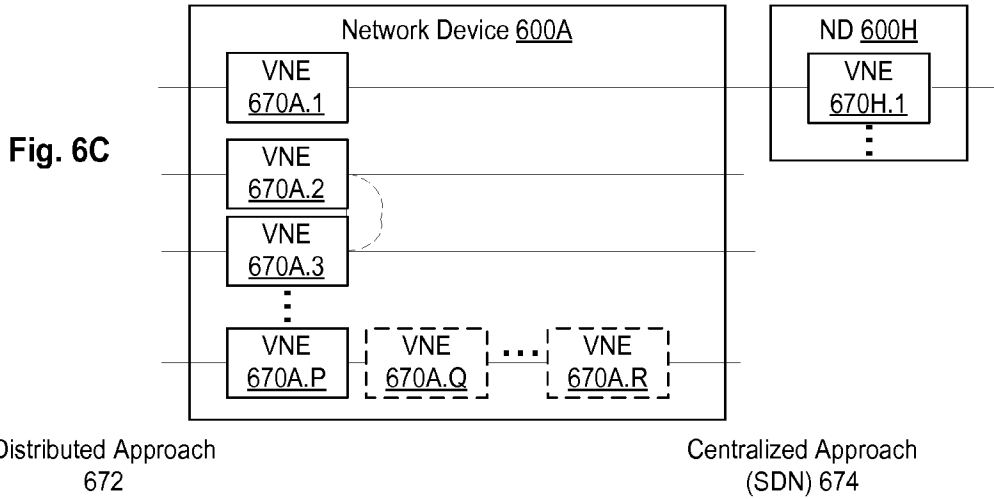


Fig. 6B





INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2015/059895

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04L9/32 H04L9/06 H04L12/801 H04L1/00 H04L29/08
 H04L29/06
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 H04L

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EMIR HALEPOVIC ET AL: "On the performance of Redundant Traffic Elimination in WLANs", COMMUNICATIONS (ICC), 2012 IEEE INTERNATIONAL CONFERENCE ON, IEEE, 10 June 2012 (2012-06-10), pages 5434-5439, XP032274418, DOI: 10.1109/ICC.2012.6364592 ISBN: 978-1-4577-2052-9 page 5434 - page 5436 ----- -/--	1-20

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 23 March 2016	Date of mailing of the international search report 01/04/2016
--	--

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

INTERNATIONAL SEARCH REPORT

 International application No
 PCT/IB2015/059895

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	BHAVISH AGGARWAL ET AL: "EndRE: An End-System Redundancy Elimination Service for Enterprises", INTERNET CITATION, 3 March 2010 (2010-03-03), page 14pp, XP007918513, Retrieved from the Internet: URL:http://pages.cs.wisc.edu/~akella/paper s/endre-nsdi10.pdf [retrieved on 2011-05-04] page 4 - page 7 -----	1-20
Y	WO 02/091668 A2 (ERICSSON TELEFON AB L M [SE]; GEHRMANN CHRISTIAN [SE]; BLOM ROLF [SE]) 14 November 2002 (2002-11-14) page 13, line 15 - page 27, line 13 -----	1-20
A	WO 2007/026287 A1 (KONINKL PHILIPS ELECTRONICS NV [NL]; NIKOV VENTZISLAV [BE]; GRUMIAUX F) 8 March 2007 (2007-03-08) page 3, line 1 - page 4, line 20 -----	1-20
A	US 2011/176556 A1 (GUO KATHERINE H [US] ET AL) 21 July 2011 (2011-07-21) paragraph [0102] - paragraph [0118] -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/IB2015/059895

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 02091668	A2	14-11-2002	NONE

WO 2007026287	A1	08-03-2007	NONE

US 2011176556	A1	21-07-2011	US 2011176556 A1 21-07-2011
			US 2013279524 A1 24-10-2013
			US 2014247836 A1 04-09-2014
			WO 2011087956 A1 21-07-2011
