

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
9 October 2008 (09.10.2008)

PCT

(10) International Publication Number
WO 2008/121690 A2

- (51) International Patent Classification:
H04L 12/24 (2006.01) H04L 12/28 (2006.01)
- (21) International Application Number:
PCT/US2008/058390
- (22) International Filing Date: 27 March 2008 (27.03.2008)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/694,703 30 March 2007 (30.03.2007) US
11/766,634 21 June 2007 (21.06.2007) US
11/771,822 29 June 2007 (29.06.2007) US
11/861,049 25 September 2007 (25.09.2007) US
11/961,895 20 December 2007 (20.12.2007) US
- (71) Applicant (for all designated States except US): **PACKE-TEER, INC.** [US/US]; 10201 N. De Anza Blvd., Cupertino, CA 95014 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **OKHOLM, Jon**

Eric [US/US]; 17660 Vista Ave., Monte Sereno, CA 95030 (US). **RIDDLE, Guy** [US/US]; 18243 Knuth Road, Los Gatos, CA 95033 (US).

(74) Agent: **SPOLYAR, Mark James**; Law Office of Mark J. Spolyar, 38 Fountain Street, San Francisco, California 94114 (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, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,

[Continued on next page]

(54) Title: DATA AND CONTROL PLANE ARCHITECTURE FOR NETWORK APPLICATION TRAFFIC MANAGEMENT DEVICE

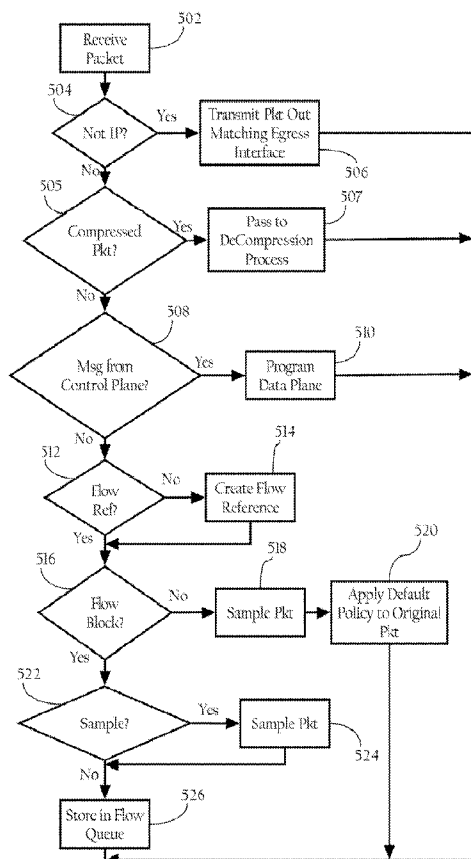


FIG. 5A

(57) Abstract: A data and control plane architecture for network devices. An example system architecture includes a network processing unit implementing one or more data plane operations, and a network device operably coupled to the network processing unit that implements a control plane. In a particular implementation, the network processing unit is configured to process network traffic according to a data plane configuration, and sample selected packets to the network device. The network device processes the sampled packets and adjusts the data plane configuration responsive to the sampled packets.

WO 2008/121690 A2



ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- *without international search report and to be republished upon receipt of that report*

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

DATA AND CONTROL PLANE ARCHITECTURE
FOR NETWORK APPLICATION TRAFFIC MANAGEMENT DEVICE

DATA AND CONTROL PLANE ARCHITECTURE
FOR NETWORK APPLICATION TRAFFIC MANAGEMENT DEVICE

TECHNICAL FIELD

[0001] The present disclosure relates to data processing in a network device including a control plane and data plane system architecture.

BACKGROUND

[0002] Enterprises have become increasingly dependent on computer network infrastructures to provide services and accomplish mission-critical tasks. Indeed, the performance, security, and efficiency of these network infrastructures have become critical as enterprises increase their reliance on distributed computing environments and wide area computer networks. To that end, a variety of network devices have been created to provide data gathering, reporting, and/or operational functions, such as firewalls, gateways, packet capture devices, bandwidth management devices, application traffic monitoring devices, and the like. For example, the TCP/IP protocol suite, which is widely implemented throughout the world-wide data communications network environment called the Internet and many wide and local area networks, omits any explicit supervisory function over the rate of data transport over the various devices that comprise the network. While there are certain perceived advantages, this characteristic has the consequence of juxtaposing very high-speed packets and very low-speed packets in potential conflict and produces certain inefficiencies. Certain loading conditions degrade performance of networked applications and can even cause instabilities which could lead to overloads that could stop data transfer temporarily.

[0003] In response, certain data flow rate control mechanisms have been developed to provide a means to control and optimize efficiency of data transfer as well as allocate available bandwidth among a variety of business enterprise functionalities. For example, U.S. Patent No. 6,038,216 discloses a method for explicit data rate

control in a packet-based network environment without data rate supervision. Data rate control directly moderates the rate of data transmission from a sending host, resulting in just-in-time data transmission to control inbound traffic and reduce the inefficiencies associated with dropped packets. Bandwidth management devices allow for explicit data rate control for flows associated with a particular traffic classification. For example, U.S. Patent No. 6,412,000, above, discloses automatic classification of network traffic for use in connection with bandwidth allocation mechanisms. U.S. Patent No. 6,046,980 discloses systems and methods allowing for application layer control of bandwidth utilization in packet-based computer networks. For example, bandwidth management devices allow network administrators to specify policies operative to control and/or prioritize the bandwidth allocated to individual data flows according to traffic classifications. In addition, network security is another concern, such as the detection of computer viruses, as well as prevention of Denial-of-Service (DoS) attacks on, or unauthorized access to, enterprise networks. Accordingly, firewalls and other network devices are deployed at the edge of such networks to filter packets and perform various operations in response to a security threat. In addition, packet capture and other network data gathering devices are often deployed at the edge of, as well as at other strategic points in, a network to allow network administrators to monitor network conditions.

[0004] Enterprise network topologies can span a vast array of designs and connection schemes depending on the enterprise's resource requirements, the number of locations or offices to connect, desired service levels, costs and the like. A given enterprise often must support multiple LAN or WAN segments that support headquarters, branch offices and other operational and office facilities. Indeed, enterprise network design topologies often include multiple, interconnected LAN and WAN segments in the enterprise's intranet, and multiple paths to extranets and the Internet. Enterprises that cannot afford the expense of private leased-lines to develop their own WANs, often employ frame relay, or other packet switched networks, together with Virtual Private Networking (VPN) technologies to connect private enterprise sites via a service provider's public network or the

Internet. Some enterprises also use VPN technology to create extranets with customers, suppliers, and vendors. These network topologies often require the deployment of a variety of network devices at each remote facility. In addition, some network systems are end-to-end solutions, such as application traffic optimizers using compression tunnels, requiring network devices at each end of a communications path between, for example, a main office and a remote facility.

[0005] Many of the network devices discussed above are typically deployed at strategic locations in the network topology such that all or nearly all network traffic flows through them. For example, firewall and intrusion detection systems are typically deployed at the edges of a network domain to filter incoming and outgoing traffic. Similarly, bandwidth management systems are typically deployed between a network and an access link to allow for more direct control of access link utilization. Given that these network devices may process large amounts of network traffic (especially during peak load conditions), they must possess sufficient computing resources to provide for sufficient performance and throughput. If the network device becomes a bottleneck, latency increases and degrades network application performance. Still further, the processes and functions performed by these network devices are becoming more complex and, thus, require higher processing power than previous generation systems. Indeed, bandwidth management systems, for example, have evolved to include complex packet inspection, classification and control mechanisms.

[0006] In some previous approaches to increasing the performance of network devices, vendors have simply relied on more powerful processors, frequently turning to customized hardware solutions. This approach, however, is inherently limited to the capability of the custom hardware. Custom hardware solutions also require increased development costs and long lead times, as well as limited flexibility for correcting bugs and adapting to changing customer requirements. In addition, while some network device manufactures have turned to systems with multiple processors, they have not addressed the challenges posed by Quality of Service (QoS) and other devices that employ stateful or flow-aware inspection, classification and control mechanisms.

SUMMARY

[0007] The present invention provides methods, apparatuses and systems directed to a network device system architecture that increases throughput of devices that process network traffic. In a particular implementation, an example system architecture includes a network device implementing a control plane, that is operably coupled to a network processing unit implementing one or more data plane operations. In a particular implementation, the network processing unit is configured to process network traffic according to a data plane configuration, and sample selected packets of a data flow to the network device. The control plane of the network device processes the sampled packets and adjusts the data plane configuration responsive to the sampled packets of the data flow. In particular implementations, the control plane and data plane implement a server-side triggered policy caching mechanism that allows for previous classification policy decisions made for previous data flows to be applied to subsequent new flows.

DESCRIPTION OF THE DRAWINGS

[0008] Figure 1 is a schematic diagram illustrating a computer network environment in which implementations of the invention may operate.

[0009] Figures 2A and 2B are schematic diagrams illustrating the interconnections between a network application traffic management device and a network processing unit according to particular implementations of the invention.

[0010] Figure 3 is a functional block diagram that shows the components and system architecture of a network application traffic management device and a network processing unit according to one particular implementation of the invention.

[0011] Figure 4A is a schematic diagram illustrating logical processing modules of an application traffic management device according to one particular implementation of the invention.

[0012] Figure 4B is a functional block diagram illustrating a process flow, according to one implementation of the present invention, among various hardware and

software modules of a network application traffic management unit.

[0013] Figure 5A is a flow chart setting forth a method, according to one possible implementation of the invention, directed to processing flows at a network processing unit.

[0014] Figure 5B is a flow chart illustrating another example method directed to processing flows at a network processing unit.

[0015] Figure 5C is a flow chart providing an example method directed to processing control messages at a network processing unit.

[0016] Figure 6 is a state diagram, according to a particular implementation of the invention, directed to controlling the sampling of packets, and the rate at which packets are sampled, to a network application traffic management device.

[0017] Figure 7 is a flow chart illustrating a method directed to sampling packets.

[0018] Figure 8 is a flow chart setting forth an example method directed to sampling packets to a control plane.

[0019] Figure 9 is a flow chart illustrating an example method directed to processing received packets sampled by a data plane.

[0020] Figure 10 is a diagram showing a flowchart of a process for passing a packet from a child partition to a parent partition or port, which process might be used with some embodiments of the present invention.

[0021] Figure 11 is a flow chart illustrating an example method directed to updating a token count for a partition.

[0022] Figure 12 is a flow chart illustrating an example method directed to scheduling a packet of a partition for forwarding.

[0023] Figure 13 is a schematic diagram graphically illustrating a hierarchical token bucket scheduling mechanism.

[0024] Figure 14 is a functional block diagram illustrating an example relationship between various hardware components of a data plane and the components of a beltway data structure.

[0025] Figure 15 is a flow chart illustrating an example process for adding command entries to a beltway data structure.

[0026] Figure 16 is a flow chart illustrating an example process for reading and

processing command entries from a beltway data structure.

[0027] Figure 17 is a flow chart illustrating an example process implemented by control plane compression logic.

[0028] Figure 18 is a flow chart illustrating an example method executed by a data plane compression process.

[0029] Figure 19 is a flow chart illustrating an example decisional logic associated with writing commands to a compression unit.

DESCRIPTION OF EXAMPLE EMBODIMENT(S)

[0030] The following example embodiments are described and illustrated in conjunction with apparatuses, methods, and systems which are meant to be examples and illustrative, not limiting in scope.

[0031] The present invention provides methods, apparatuses and systems directed to a network device system architecture that increases throughput of devices that process network traffic. In a particular implementation, an example system architecture includes a network device implementing a control plane, that is operably coupled to a network processing unit implementing one or more data plane operations. In a particular implementation, the network processing unit is configured to process network traffic according to a data plane configuration, and sample selected packets to the network device. The network device processes the sampled packets and adjusts the data plane configuration responsive to the sampled packets. In particular implementations, the present invention is directed to methods, apparatuses and systems that use fast network processors to accelerate the operation of existing slower network device hardware platforms. As described herein, the architecture allows the bulk of network traffic processing to be offloaded to the fast network processor instead of the network device.

[0032] In a particular implementation, the present invention provides a cost effective solution to increasing the throughput of existing hardware with little to no modification to the existing hardware and minimal changes to software or firmware with the use of an external appliance or device that implements a data plane can be used to increase the throughput of existing hardware with little to no modification to the existing hardware and minimal changes to software or firmware to implement control plane operations.

A. Network Environment

[0033] Figure 1 illustrates, for didactic purposes, a network 50, such as wide area network, interconnecting a first network 40a, supporting a central operating or headquarters facility (for example), and a second network 40b, supporting a branch

office facility (for example). In one implementation, network 50 may include a MPLS VPN network core interconnecting networks 40a and 40b. Network 50 may also be operably connected to other networks associated with the same administrative domain as networks 40a, 40b, or a different administrative domain. Furthermore, network 50 may allow access to a variety of hosts over the Internet, as well. As Figure 1 shows, the first network 40a interconnects several hosts or end systems 42, including personal computers and servers, and provides access to resources operably connected to computer network 50 via router 22 and access link 21. Access link 21 is a physical and/or logical connection between two networks, such as computer network 50 and network 40a. The computer network environment, including network 40a and network 50 is a packet-based communications environment, employing TCP/IP protocols (for example), and/or other suitable protocols, and has a plurality of interconnected digital packet transmission stations or routing nodes. First network 40a, and network 40b, can each be a local area network, a wide area network, combinations thereof, or any other suitable network.

[0034] As Figure 1 illustrates, network devices 30, in one implementation, are deployed at the respective edges of networks 40a and 40b. In a particular implementation, network devices 30 are network application traffic management devices operative to manage network application traffic. As discussed below, network application traffic management devices 30 may include a number of different functional modules, such as compression modules, tunneling modules, rate control modules, gateway modules, protocol acceleration modules, and the like. In addition, network application traffic management devices 30 may include functions, such as compression and/or tunneling, where cooperation with a remote device (such as another network application traffic management device) is required, while also performing other functions that can be performed independently. However, the control and data plane system architecture according to the present invention can be integrated into a variety of network devices, such as proxies, firewalls, intrusion detection systems, packet capture or network monitoring equipment, VPN servers, web services network gateways or brokers, and the like.

A.1. Example System Architecture

[0035] Figure 2A illustrates an example system architecture of network application traffic management device 30. In the implementation shown, network application traffic management device 30 comprises an application traffic management unit 200 and a network processing unit 300. Network application traffic management unit 200 may be implemented using existing hardware device platforms and functionality with small software modifications to interact with network processing unit 300. Network processing unit 300 interconnects with application traffic management unit as shown, and is further operably connected to network 50 and LAN 40a. In the implementation shown, network processing unit 300 is a discrete processing unit that does not share memory with network application traffic management unit 200, instead being interconnected via one or more network interfaces. In other implementations, the interconnections between network processing unit 300 and network application traffic management unit 200 can be implemented using other suitable interface technologies, such as other network interface technologies, and bus interfaces (e.g., Peripheral Component Interconnect (PCI) interfaces, and Industry Standard Architecture (ISA) interfaces). Furthermore, network application traffic management unit 200 and network processing unit 300 may be directly interconnected to each other with only network cabling extending between the packet interfaces. In another implementation, the packet interfaces of the network application traffic management unit 200 and network processing unit 300 may be connected to an Ethernet switch or other network fabric. The devices could use layer 2 or 3 tunneling protocols to transmit sampled packets to each other.

[0036] As Figure 2A illustrates network application traffic management unit 200 comprises packet interfaces 216a and 216b, while network processing unit 300 comprises packet interfaces 320a-d. In one implementation, packet interfaces comprise Ethernet interfaces including MAC layer functionality. However, other network interfaces can be used depending on the network environment. Generally, packets received from network 50 on packet interface 320a are transmitted to

network 40a from packet interface 320b, and vice versa. Network processing unit 300, generally speaking, may be configured to implement one or more data plane operations on the network traffic transmitted between network 50 and network 40a according to a data plane configuration. As discussed in more detail below, network processing unit 300 is configured to receive packets from network 50 or 40a and selectively sample received packets to application traffic management unit 200 for processing. In one implementation, packets received on packet interface 320a, when sampled, are transmitted from packet interface 320c to packet interface 216a of network application traffic management unit 200. Network processing unit 300, in one implementation, then processes the packet without waiting for a response from network application traffic management unit 200. In a particular implementation, the network processing unit 300, given that it processes packets according to its current data plane configuration, can continue to process packets even when the control plane crashes and/or reboots. In addition, network application traffic management unit 200 may also be configured to emit probe messages and other messages directed to device discovery, network management, and the like, directly as well. Application traffic management unit 200 performs various control plane operations on sampled packets, such as packet classification, policy identification, and the like. Responsive to one or more received sample packets, network application traffic management unit 200 may transmit one or more control messages to network processing unit 300 operative to cause changes to the data plane configuration of the network processing unit 300. For example, network application traffic management unit 200 may receive sampled packets of a data flow, and classify them to identify one or more policies or controls. Network application traffic management unit 200 may then transmit a control message identifying one or more traffic classification or service types and one or more policies to be applied to packets of the data flow. Network processing unit 300 may then apply the one or more policies to subsequent packets of the same data flow. [0037] Other implementations are possible. For example, network application traffic management unit 200 and network processing unit 300 could be connected using a single pair of packet interfaces. In other implementations, network

application traffic management unit 200 and network processing unit 300 could be connected with additional packet interfaces than that shown in Figure 2A. For example, a packet interface pair between network application traffic management unit 200 and network processing unit 300 could be configured for each packet interface of network processing unit 300 that is connected to a network, such as an internal LAN or a WAN. As Figure 2B shows, in another implementation, network processing unit 300 could be configured with multiple packet interfaces 320a, 320b, 320e, and 320f for communication with respective networks 50, 40a, 51 & 41a, and multiplex packets transmitted to a network application traffic management unit 200 over a smaller number of packet interfaces 320c and 320d.

[0038] In yet another implementation, a single network application traffic management unit 200 can be connected to multiple network processing units 300 disposed at various points in a network environment. For example, two network processing units 300 could be deployed on separate access links, and communicably coupled to a single network application traffic management unit 200. Conversely, a single network processing unit 300 could be operably coupled to multiple application traffic management units 200. In one such implementation, the network processing unit can be configured to ensure that packets of the same flow are transmitted to the same network application traffic management 200.

A.1.a. Network Application Traffic Management Unit

[0039] While network application traffic management unit 200 may be implemented in a number of different hardware architectures, some or all of the elements or operations thereof may be implemented using a computing system having a general purpose hardware architecture such as the one in Fig. 3. In one implementation, network application traffic management unit 200 comprises a processor 202, a cache memory 204, and one or more software applications and drivers directed to the functions described herein. In one implementation, network application traffic management unit 200 includes a high performance input/output (I/O) bus 206 and a standard I/O bus 208. A host bridge 210 couples processor 202 to high performance I/O bus 206, whereas I/O bus bridge 212 couples the two buses 206 and 208 to each

other. A system memory 214 and one or more network/communication interfaces 216 couple to bus 206. Mass storage 218 and I/O ports 220 couple to bus 208. Collectively, these elements are intended to represent a broad category of computer hardware systems, including but not limited to general purpose computer systems based on the Pentium® processor manufactured by Intel Corporation of Santa Clara, Calif., as well as any other suitable processor.

[0040] Network interface 216c provides communication between network application traffic management unit 200 and a network through which a user may access management or reporting functions. Mass storage 218 provides permanent storage for the data and programming instructions to perform the above described functions implemented in the system controller, whereas system memory 214 (e.g., DRAM) provides temporary storage for the data and programming instructions when executed by processor 202. I/O ports 220 are one or more serial and/or parallel communication ports that provide communication between additional peripheral devices, which may be coupled to network application traffic management unit 200.

[0041] Network application traffic management unit 200 may include a variety of system architectures; and various components of network application traffic management unit 200 may be rearranged. For example, cache 204 may be on-chip with processor 202. Alternatively, cache 204 and processor 202 may be packed together as a “processor module,” with processor 202 being referred to as the “processor core.” Furthermore, certain implementations of the present invention may not require nor include all of the above components. For example, the peripheral devices shown coupled to standard I/O bus 208 may couple to high performance I/O bus 206. In addition, in some implementations only a single bus may exist, with the components of network application traffic management unit 200 being coupled to the single bus. Furthermore, network application traffic management unit 200 may include additional components, such as additional processors, storage devices, or memories.

[0042] The operations of the network application traffic management unit 200 described herein are implemented as a series of software routines (see Figures 4A

and 4b) hosted by network application traffic management unit 200. These software routines comprise a plurality or series of instructions to be executed by a processor in a hardware system, such as processor 202. Initially, the series of instructions are stored on a storage device, such as mass storage 218. However, the series of instructions can be stored on any suitable storage medium, such as a diskette, CD-ROM, ROM, EEPROM, etc. Furthermore, the series of instructions need not be stored locally, and could be received from a remote storage device, such as a server on a network, via network/communication interface 216c. The instructions are copied from the storage device, such as mass storage 218, into memory 214 and then accessed and executed by processor 202.

[0043] An operating system manages and controls the operation of network application traffic management unit 200, including the input and output of data to and from software applications (not shown). The operating system provides an interface between the software applications being executed on the system and the hardware components of the system. According to one embodiment of the present invention, the operating system is a real-time operating system, such as PSOS, or LINUX. In other implementations, the operating system may be the Windows® 95/98/NT/XP/Vista operating system, available from Microsoft Corporation of Redmond, Wash. However, the present invention may be used with other suitable operating systems, such as the Apple Macintosh Operating System, available from Apple Computer Inc. of Cupertino, Calif., UNIX operating systems, and the like.

[0044] Figure 4B provides an illustration of the components and functional modules, and data structures, relevant to how packets are processed by network application traffic management unit 200. As Figure 4B illustrates, network application traffic management unit 200 includes inside NIC receive ring 81a, outside NIC receive ring 81b, inside NIC receive queue 84a, outside NIC receive queue 84b, NIC driver 83, and packet buffer 82. Packet buffer 82 is operative to store packets received at packet interfaces 216a, 216b. To summarize the operations associated with receiving and ultimately processing packets, network interface 216a, for example, receives and stores a packet in packet buffer 82. Network interface 216a also maintains a pointer to the packet in inside NIC receive

ring 81a. As discussed more fully below, NIC driver 83 determines whether to queue or otherwise retain the packet, or to discard it. In one implementation, NIC driver 83, operating at periodic interrupts, writes pointers out of inside NIC receive ring 81a and into inside NIC receive queue. Network application traffic management unit 200 operates substantially identically for packets received at outside packet interface 216b. Network device application 75, in one implementation, operates on packets stored in packet buffer 82 by accessing the memory address spaces (pointers) to the packets in inside NIC receive queue 84a and outside NIC receive queue 84b. In one implementation, a sample packet received at inside packet interface 216a is dropped, after processing by network device application, as opposed to being transmitted from outside packet interface 216b. The rings and other data structures supporting the transmission of packets from network interfaces 216a, 216b are not shown.

[0045] In one implementation, packet buffer 82 comprises a series of fixed-size memory spaces for each packet (e.g., 50,000 spaces). In other implementations, packet buffer 82 includes mechanisms allowing for variable sized memory spaces depending on the size of the packet. Inside NIC receive ring 81a is a circular queue or ring of memory addresses (pointers) corresponding to packets stored in packet buffer 82. In one implementation, inside NIC receive ring 81a includes 256 entries; however, the number of entries is a matter of engineering and design choice. In one implementation, each entry of inside NIC receive ring 81a includes a field for a memory address, as well as other fields for status flags and the like. For example, one status flag indicates whether the memory address space is empty or filled with a packet. Inside NIC receive ring 81a also maintains head and tail memory addresses, as described below. In one implementation, packet interface 216a also maintains the head and tail memory address spaces in its registers. The head memory address space corresponds to the next available memory space in packet buffer 82 to which the next packet is to be stored. Accordingly, when packet interface 216a receives a packet, it checks the head address register to determine where in the system memory reserved for packet buffer 82 to store the packet. After the packet is stored, the status flag in the ring entry is changed to "filled". In

addition, the system memory returns a memory address for storing the next received packet, which is stored in the next entry in inside NIC receive ring 81a, in addition, the head address register is advanced to this next memory address. The tail memory address space corresponds to the earliest received packet which has not been processed by NIC driver 83. In one implementation, packet interface 216a also maintains a copy of inside NIC receive ring 81a in a memory unit residing on the network interface hardware itself. In one implementation, packet interface 216a discards packets when inside NIC receive ring 81a is full—i.e., when the tail and head memory addresses are the same.

[0046] As discussed above, NIC driver 83 is operative to read packet pointers from inside NIC receive ring 81a to inside NIC receive queue 84a. In one implementation, NIC driver 83 operates on inside NIC receive ring 81a by accessing the tail memory address to identify the earliest received packet. To write the packet in the inside NIC receive queue 84a, NIC driver 83 copies the memory address into inside NIC receive queue, sets the status flag in the entry in inside NIC receive ring 81a corresponding to the tail memory address to empty, and advances the tail memory address to the next entry in the ring. NIC driver 83 can discard a packet by simply dropping it from inside NIC receive ring 81a, and not writing it into inside NIC receive queue 84a. As discussed more fully below, this discard operation may be performed in connection with random early drop mechanisms, or the load shedding mechanisms, according to those described in commonly-owned U.S. Application No. 10/843,185, entitled “Packet Load Shedding”, filed on May 11, 2004, whose disclosure is incorporated herein by reference for all purposes. Still further, NIC driver 83, in one implementation, is a software module that operates at periodic interrupts to process packets from inside NIC receive ring 81a to inside NIC receive queue 84a. At each interrupt, NIC driver 83 can process all packets in receive ring 81a or, as discussed more fully below, process a limited number of packets. Furthermore, as discussed more fully below, a fairness algorithm controls which of inside NIC receive ring 81a and outside NIC receive ring 81b to process first at each interrupt.

[0047] In one implementation, inside NIC receive ring 81a, outside NIC receive ring

81b, inside NIC receive queue 84a, outside NIC receive queue 84b, and packet buffer 82 are maintained in reserved spaces of the system memory of network application traffic management unit 200. As discussed above, network device application 75, operating at a higher level, processes packets in packet buffer 82 popping packet pointers from receive queues 84a, 84b. The system memory implemented in network application traffic management unit 200, in one embodiment, includes one or more DRAM chips and a memory controller providing the interface, and handling the input-output operations, associated with storing data in the DRAM chip(s). In one implementation, the hardware in network application traffic management unit 200 includes functionality allowing first and second network interfaces 216a, 216b to directly access memory 82 to store inbound packets received at the interfaces in packet buffer. For example, in one implementation, the system chip set associated with network application traffic management unit 200 can include a Direct Memory Access (DMA) controller, which is a circuit that allows for transfer of a block of data from the buffer memory of a network interface, for example, directly to memory 82 without CPU involvement. A variety of direct memory access technologies and protocols can be used, such as standard DMA, first-party DMA (bus mastering), and programmed I/O (PIO). In one implementation, each network interface 216a and 216b is allocated a DMA channel to the memory 82 to store packets received at the corresponding interfaces. [0048] Figure 4A is a block diagram illustrating functionality, according to one embodiment of the present invention, included in network application traffic management unit 200. In one embodiment, network application 75 of network application traffic management unit 200 comprises packet processor 131, data plane interface module 132, measurement engine 140, traffic classification engine 137, management information base (MIB) 138, control plane compression management module 139, and administrator interface 150. The co-pending and commonly owned patents and patent applications identified above describe various functions and operations that can be incorporated into network application traffic management unit 200. Packet processor 131 is operative to detect new data flows and construct data structures including attributes characterizing the data flow.

Data plane interface module 132 is operative to generate control messages and transmit them to network processing unit 300, as well as receive packets (control messages and sampled packets) from network processing unit 300 and selectively forward sampled packets to packet processor 131. Traffic classification engine 137 is operative to analyze data flow attributes and identify traffic classes corresponding to the data flows. In one embodiment, traffic classification engine 137 stores traffic classes, in association with pointers to bandwidth utilization controls or pointers to data structures defining such bandwidth utilization controls. Management information base 138 is a database of standard and extended network objects related to the operation of network application traffic management unit 200. Control plane compression management module 139 handles management tasks such as remote compression device discovery, and the like. Measurement engine 140 maintains measurement data relating to operation of network application traffic management unit 200 to allow for monitoring of bandwidth utilization across access link 21 with respect to a plurality of bandwidth utilization and other network statistics on an aggregate and/or per-traffic-class level. Network application traffic management unit 200, in one embodiment, further includes a persistent data store (not shown), such as a hard disk drive, for non-volatile storage of data.

[0049] Administrator interface 150 facilitates the configuration of network application traffic management unit 200 to adjust or change operational and configuration parameters associated with the device. For example, administrator interface 150 allows administrators to select identified traffic classes and associate them with bandwidth utilization controls (e.g., a partition, a policy, etc.).

Administrator interface 150, in one implementation, also displays various views associated with a traffic classification scheme and allows administrators to configure or revise the traffic classification scheme. Administrator interface 150 can be a command line interface or a graphical user interface accessible, for example, through a conventional browser on client device 42. In addition, since in one implementation, network processing unit 300 may not be a network addressable device and only responds to control messages transmitted from

network application traffic management unit 200, administrator interface 150 provides a unified user interface for network application traffic management unit 200 and network processing unit 300 in the aggregate.

[0050] As disclosed in U.S. Application No. 10/843,185, the number of packets in the inside or outside NIC receive queues 84a, 84b can be monitored to signal a possible overload condition. That is when the number of packets in one of the queues exceeds a threshold parameter, network application traffic management unit 200 may perform one or more actions. In one implementation, network application traffic management unit 200 may transmit a message to network processing unit 300 signaling that it is at or near an overload state. As described in more detail below, network processing unit 300 responsive to such a message may stop sampling packets to network application traffic management unit 200 or reduce the rate at which packets are sampled. Still further, as described in U.S. Application No. 10/843,185, network application traffic management unit 200 may access host database 134 to compare certain observed parameters corresponding to the source hosts identified in received packets, and compare them against corresponding threshold values to determine whether to discard received packets. For example, a host identified as being part of a Denial-of-Service attack may be deemed a "bad host." In one implementation, network application traffic management unit 200 may transmit control messages to network processing unit 300 directing it to drop packets from an identified bad host.

[0051] When network application traffic management unit 200 operates without network processing unit 300 it generally operates to receive packets at a first interface (e.g., packet interface 216a), process the packets, and emit the packets at a second interface (e.g., packet interface 216a), or vice versa. When configured to operate in connection with network processing unit 300, however, network application traffic management unit 200 is configured to receive and process the packets sampled to it, but to drop the packets instead of emitting them. As part of this process, network application traffic management unit 200, in one implementation, receives a sampled packet, processes the packet, and may transmit one or more control messages to network processing unit 300 indicating

how subsequent packets of a data flow should be handled.

A.1.b. Network Processing Unit

[0052] Figure 3 also illustrates an example system architecture for a network processing unit 300 according to a particular implementation of the invention. In one implementation, network processing unit 300 comprises a multi-core processor such as the Octeon Plus CN58XX 4 to 16-Core MIPS64-Based SoCs offered by Cavium Networks of Mountain View, California. Of course, any other suitable single or multi-core processor can be used. In the implementation shown in Figure 3, network processing unit 300 comprises processor 302 (comprising one to a plurality of processing cores), cache 304 (e.g., L2 cache shared among the processing cores), memory controller 306 and random access memory (RAM) 308. Network processing unit 300, in one implementation, further comprises packet input module 310, I/O bridge 312, packet output module 314, TCP unit 316, scheduler or work order module 318, compression/decompression unit 324 (hereinafter, "compression unit 324"), and packet interfaces 320a-d. Network processing unit 300 may also include other functional modules such as a regular expression unit for string matching, and an encryption/decryption unit. Still further, network processing unit 300, in some implementations, may include multiple multi-core processors.

[0053] Although not illustrated, in one implementation, network processing unit 300 may also include a power supply, RJ-45 or other physical connectors, and a chassis separate from network application traffic management unit 200. For example, as discussed above, network processing unit 300 may be a separate physical unit in the form factor of a 1U or 2U appliance. The network processing unit 300 may be used to accelerate and enhance the throughput of an existing network application traffic management device, such as network application traffic management unit 200. In one implementation, without network processing unit 300, application traffic management unit 200 would be directly connected to the network path segment between network 50 and network 40a. For example, packet interface 216a would be operably connected to network 50, while packet interface

216b would be operably connected to network 40a. To increase throughput, however, network processing unit 300 may be interconnected as shown in Figures 2A, 2B and 3.

[0054] Network processing unit 300 can be configured to perform various data plane operations, and to selectively forward packets to application traffic management unit 200. In one possible configuration, application traffic management unit 200 performs higher-level processing of packets of respective data flows to classify the data flows and identify one or more policies to be applied to the data flows.

Throughput can be increased, due to the capabilities of the network processing unit 300 to perform data plane operations on packets at wireline or near wireline speeds, and that network application traffic management unit 200 sees only a subset of all packets traversing networks 40a and 50. Accordingly, particular implementations of the invention provide for an inexpensive and easily deployable solution that accelerates the performance of existing hardware and allows an end-user to preserve investments in existing hardware platforms. In one such implementation, the network processing unit 300 can be distributed with a computer-readable media, such as optically or magnetically recorded disks or tapes, that include one or more software modules that, when installed, modify the operation of the network application traffic management unit 200 to interact with the network processing unit 300 as described herein. The computer-readable media may also include a copy of firmware for the network processing unit 300. In one implementation, network application traffic management unit 200 can store the firmware and provide it to network processing unit 300 during a configuration session. In other implementations, however, the network processing unit 300 and network application traffic management unit 200 may be incorporated into the same chassis.

B. Control Messages

[0055] As described herein, network application traffic management unit 200 (Control Plane) and network processing unit 300 (Data Plane) implement a two-way message path by which network application traffic management unit 200

directs network processing unit 300 which policies should be applied to the data flows traversing it. In a particular implementation, network processing unit 300 also returns network statistics, such as Measurement Samples, to be integrated into the measurement and reporting functionality of measurement engine 140 of network application traffic management unit 200. The Control Plane makes the flow decision after completing classification of the flow, including peeking at the data packets as necessary and consulting the policies stored in the traffic classification engine 137. Example network traffic classification mechanisms are described in U.S. Application Ser. Nos. 11/019,501, as well as other patents and patent applications identified above. Still further, as discussed in more detail below, the Control Plane may direct the Data Plane to cache the traffic classification and policy decisions made in connection with a given data flow, and apply the same policies to subsequent data flows that include a set of matching attributes, such as destination network address, port and protocol identifiers.

[0056] Control messages between the network application traffic management unit 200 and network processing unit 300, in one implementation, use a specific VLAN to facilitate identification of control messages and other communications between them. In some implementations, VLAN tagging is not employed. Flow Information Messages have the same IP and TCP/UDP protocol headers as the flow they refer to in order to get the same tuple hash from the network processor hardware. Alternatively, flow information messages can be encapsulated in IP-in-IP or Generic Routing Encapsulation (GRE) or other tunneling protocols. Other control messages use specific addresses for the network application traffic management unit 200 and network processing unit 300. These are local to the two units (in one implementation, chosen from the 127 class A address range) and need no configuration.

[0057] In a particular implementation, there are 5 types of control messages from the Control Plane to the Data Plane, and 3 types of control messages in the reverse direction. The first message sent to the Control Plane is the SizingData message describing one or more attributes of various operational data structures, such as the sizes of tables, buffer spaces reserved for various functions, etc. PartitionInfo

messages are sent to describe the configuration of partitions, and any subsequent changes. A FlowInfo message is sent when network application traffic management unit 200 decides on the partition and policy to apply to a flow. Two message types, the OverloadStatus and the BadHostInfo inform the Data Plane when the network application traffic management unit 200 enters or leaves an overloaded condition and of any hosts the Load Shedding feature decides are behaving badly.

[0058] The three types of messages sent from the Data Plane to the Control Plane are the ReTransmitRequest to recover from possible lost messages or to resynchronize, the MeasurementSample message to transmit measurement samples for the configured traffic classes and partitions, and the LittleNote to transmit status messages to be logged.

[0059] Other message types may also be implemented for different functions. For example, one or more message types may be configured for compression functions, such as a message for setting up Layer 3 tunnels with remote nodes (in implementations that would use such tunnels), and specifying the compression algorithm to be used. Other message types may include encryption message types as well. In yet other embodiments, network application traffic management unit 200 may store a firmware image for network processing unit 300 and interact (typically during initialization) to determine the firmware image stored on network processing unit 300. Network application traffic management unit 200, if it determines that a firmware update is required, may transmit the firmware image to network processing unit 300 in one to a plurality of control messages.

B.1. SizingData Message

[0060] Network application traffic management unit 200 transmits a SizingData message to provide an initial configuration to the network processing unit 300. In a particular implementation, network processing unit 300 simply forwards received packets along the network path to their destination without processing, until it receives a configuration from the network application traffic management unit 200. The SizingData message indicates the capacities of the Control Plane. In a

particular implementation, the Data Plane allocates its memory to be aligned with these capacities, such as the number of partitions, the number of supported traffic classes, the number of supported flow blocks. The following illustrates an example format of a SizingData message according to one particular implementation of the invention. In a particular implementation, objects, such as data flows, partitions, and classes are referenced relative to an index and an instance identifier.

```

typedef struct _ObjectReference {
    uint16_t      index;
    uint16_t      instance;
} ObjectReference;

typedef struct _SizingData {
    uint16_t      type;
#define          kSizingDataType          787
    uint16_t      seq;                    // sequence number

    uint32_t      ipaddr;                 // main address of NATM
    uint32_t      ptncount;               // # of partitions
    uint32_t      classcount;             // # of traffic classes
    uint32_t      flowcount;              // # of flow blocks
    uint16_t      servicecount;           // # of services
    uint16_t      triggerexpiration;      // in minutes
    uint32_t      triggerlines;           //number of DTP hash groups
} SizingData, *SizingDataPtr;

```

The triggerexpiration parameter indicates the length of time that a cached policy decision is valid. In one implementation, a null or zero value indicates that the cached policy decision does not time out. The triggerlines parameter indicates the number of hash groups the Data Plane should allocate in memory.

B.2. PartitionInfo Message

[0061] Network application traffic management unit 200 sends PartitionInfo messages when a partition is created, deleted, moved, or resized. A PartitionInfo message can also be transmitted in response to a ReTransmitRequest message sent by the Data Plane (see below).

```

typedef struct _PartitionInfo {
    uint16_t      type;
#define          kPartitionInfoType      789
    uint16_t      seq;

```

```

        ObjectReference    partition;    // this partition
        uint8_t            direction;
#define kInboundDirection    0
#define kOutboundDirection  1
        uint8_t            isroot;
        uint8_t            action;
#define kPartitionActionNew    1
#define kPartitionActionResize 2    // linkages unchanged
#define kPartitionActionDetach 3    // detach, don't
                                        delete, uses old parent
#define kPartitionActionReattach 4    // is detached, parent
                                        is new parent
#define kPartitionActionDelete 5    // should be leaf,
                                        parent is old parent
#define kPartitionActionRetransmit 6
        uint8_t            isdefault;
        ObjectReference    parent;
        uint32_t            minbw;
        uint32_t            maxbw;
    } PartitionInfo, *PartitionInfoPtr;

```

[0062] Some partition attributes in the PartitionInfo message include the minimum (minbw) and maximum (maxbw) bandwidth allocated to the partition, the identity of the parent of the partition, the direction of traffic flow (direction) to which the partition corresponds, and whether the partition is the default partition (isdefault) or root (isroot) for that direction.

B.3. FlowInfo Message

[0063] A major aspect of the control functions performed by the Control Plane is embodied in the FlowInfo message sent by the Control Plane when it has decided what policy or policies should be applied to a new data flow. In one implementation, the Control Plane is operative to create a data structure for the flow, and transmit a FlowInfo message to the Data Plane. The FlowInfo message causes the Data Plane to create a flow block, which is a data structure or object for storing various attributes of the data flow. The flow block is identified by a FlowIndex and an instance value. Attributes of the flow block may include one or more of the attributes defined in the FlowInfo message set forth below. The following illustrates attributes that may be included in a FlowInfo message according to one particular implementation of the invention.

```

typedef struct _FlowInfo {
    uint16_t      type;
#define          kFlowInfoType      788
    uint16_t      seq;

    uint32_t      flowindex;          // identifies flows
    uint16_t      flowinstance;
    uint16_t      service;           //service identifier for flow
    uint8_t       serverside;

#define          kInSide              0          //server located inside
#define          kOutSide             1          //server located outside
    uint8_t       sendmore;         // keep sending packets
    uint8_t       policies[2];

#define          kPolicyPriority       0x01
#define          kPolicyRate          0x86      // any of these bits
#define          kPolicyPassThru     0x08
#define          kPolicyDiscard      0x10
#define          kPolicyNever        0x20
#define          kPolicyCompressible 0x40

    ObjectReference classes[2];      // by direction
    ObjectReference partitions[2];
    uint8_t       priorities[2];
    uint8_t       trigger;

#define          kTriggerDont         0
#define          kTriggerRemember    1
#define          kTriggerRecycled     2
#define          kTriggerServiceOnly 3
    uint8_t       direction;         // to hash flow
                                        attribute tuple
} FlowInfo, *FlowInfoPtr;

```

[0064] Each data flow is identified by its FlowIndex, a number uniquely determined by which flow block (TCB or UCB type) was allocated to it by the Control Plane. The FlowInfo message, in a particular implementation, contains the determined policy or policies for the data flow (for example, one of Priority, Rate, PassThru, Discard, or Never). In addition, the kPolicyCompressible variable, in one implementation, is a one-bit variable or flag that indicates to the Data Plane that the data flow is compressible (see below). Additional compression parameters or policies can also be included if the variable width (in bits) is expanded. Still further, the FlowInfo messages may also include a service parameter which is a value that maps to a network application type (such as Oracle® database, FTP, Citrix®, HTTP, and other network applications). The serverside parameter indicates whether the location of the server of the data flow relative to the Data and Control plane. A server is typically the host that received the initial packet of the data flow (such as a TCP SYN) from a client host. The inside or outside server

determination is based on detection of the initial packets of a data flow and their direction. With reference to Figure 1, an "inside" server relative to network application traffic management device 30 associated with network 40a, is a host connected to network 40a, while an outside server host is a host located across network 50. In one particular implementation, there are policies for each direction ("inbound" and "outbound") or "half-flow" of the traffic flow. There are also two traffic class indices, partition numbers, and priorities in the FlowInfo message. The partition parameter indicates the partition to which the data flow should be associated for each traffic direction, while the priority parameter indicates the priority of the data flow.

[0065] The FlowInfo message may also contain control variables related to interaction between the Control Plane and Data Plane relative to the data flow. For example, the Control Plane may set the sendmore variable to false to indicate that the Data Plane should completely take over handling packets of the data flow. For example, as described in more detail below, the Data Plane will continue to sample packets of a data flow to the Control Plane until it receives a FlowInfo message for that data flow, where the sendmore variable is set to "false." If the sendmore variable is set to true, the Data Plane will continue to sample packets to the Control Plane until the Control Plane transmits another FlowInfo message with sendmore set to false. In a particular implementation, when packet sampling stops for a given data flow is defined by the Control Plane, which can use this mechanism to implement one or more value added features, such as packet capture. For example, if a data flow hits a traffic class with packet capture enabled, the Control Plane can set sendmore to true and never clear it for the life of the data flow. Anything that required the Control Plane to handle all the packets of a flow could be handled in this manner.

[0066] FlowInfo messages may also contain a trigger parameter indicating whether the Control Plane should cache the service identification (service), traffic classification (classes) and policy (partitions, priorities) decisions contained in the FlowInfo message. For example, if the trigger parameter is 0, the Data Plane does not cache this information. If the trigger parameter is 1, the Data Plane caches

this information for use in connection with subsequent data flows that match the server-side attributes of the current data flow. Furthermore, the Control Plane may set the trigger parameter to 2, indicating that the Control Plane acknowledges the matching of the data flow to cached information and should not be cached another time. The Control Plane may also set the trigger parameter to 3 to indicate that the flow information (such as network application type) should be cached in the triggering cache, but that default policies (instead of cached policies in the triggering) should be applied. Processing of FlowInfo messages is described in more detail below.

[0067] In a particular implementation, FlowInfo messages have the same IP and TCP/UDP protocol headers as the data flow to which they refer. In such a configuration, the network processing unit 300 computes the same hash value for the 5-tuple (see below) of header attributes that are used to identify data flows. Network processing unit 300 has functionalities that allow for the packets of the same data flow to be processed by a common processor core. Addressing the FlowInfo messages in this manner allows the control messages for a flow to be processed by the same processor core handling data packets of the flow. Alternatively, the attributes of the 5-tuple for the data flow can also be included in the FlowInfo message, and the addresses in the headers can correspond to the addresses of the Data Plane and Control Plane.

B.4. OverloadStatus and BadHostInfo Messages

[0068] The Control Plane uses the OverloadStatus and BadHostInfo messages to control the flow of sampled packets from the Data Plane. The following defines the formats of the OverloadStatus and BadHostInfo messages according to an implementation of the invention.

```

typedef struct _OverloadStatus {
    uint16_t          type;
#define              kOverloadStatusType          791
    uint16_t          seq;
    uint8_t           overloaded;
} OverloadStatus, *OverloadStatusPtr;

```

```

typedef struct _BadHostInfo {
    uint16_t          type;
#define kBadHostType 792
    uint16_t          seq;

    uint32_t          ipaddr;
    uint8_t           client;
    uint8_t           direction;
} BadHostInfo, *BadHostInfoPtr;

```

[0069] In one implementation, the Data Plane is not configured with a “maximum rate” the Control Plane is capable of handling. Rather, the Control Plane learns this from the OverloadStatus messages sent from the Control Plane when it senses an overload condition, such as a threshold number of packets in one or more receive queues. This signaling scheme allows the Data Plane to automatically adjust to interfacing with other models of a network application traffic management unit 200 or recognizing that different network traffic mixes may place different loads on the classification mechanisms of the Control Plane.

[0070] In a particular implementation, the Control Plane also indicates to the Data Plane when hosts are behaving badly. For example, the Control Plane may send a BadHostInfo message to inform the Data Plane of any hosts the Load Shedding feature decides are behaving badly. The Data Plane can reduce or block traffic for a period of time in response to the BadHostInfo messages. In one implementation, the Data Plane can grow the packet rate sampled to the Control Plane (relative to a given host) until it receives a subsequent BadHostInfo message from the Control Plane.

B.5. ReTransmitRequest Message

[0071] As discussed above, the Data Plane may also transmit messages to the Control Plane. For example, the Data Plane may send a ReTransmitRequest message that lets the Data Plane ask for a replay of certain of the downward control messages. In a particular implementation, the Data Plane may transmit a ReTransmitRequest message each time it sees an object referenced in a control message for which it has no information. For example, the Data Plane may request a replay of the SizingData message, which may get lost while the Control Plane is booting up, or the OverloadStatus message, which might get lost in an overload

condition, and the PartitionInfo message, which is helpful for resynchronization when the Control Plane comes up after the Data Plane. ReTransmitRequest messages also facilitate resynchronization between the Control Plane and the Data Plane in the event of a fault or crash of either the Data Plane or the Control Plane. The following illustrates the format of a ReTransmitRequest message according to one particular implementation of the invention.

```

typedef struct _ReTransmitRequest {
    uint16_t      type;
#define          kRetransmitType          775
    uint16_t      seq;

    uint16_t      what;          // message type to resend
    uint16_t      index;        // partition index (kPartitionInfoType)
} RetransmitRequest, *RetransmitRequestPtr;

```

B.6. MeasurementSample Message

[0072] In one implementation, the Control and Data Planes implement a measurement data signaling scheme to allow measurement engine 140 to maintain network statistics relative to data flows, partitions and traffic classes. In a particular implementation, the Data Plane transmits MeasurementSample messages to the Control Plane such that it can update the values of various statistics it maintains. The following illustrates the format of a MeasurementSample message according to one possible implementation of the invention.

```

typedef struct _MeasurementSample {
    ObjectReference obj;
    uint32_t        packets;
    uint32_t        bytes;
} MeasurementSample, *MeasurementSamplePtr;

typedef struct _MeasurementData {
    uint16_t      type;
#define          kMeasurementDataType    777
    uint16_t      seq;

    uint16_t      count;
    uint8_t       flavor;
#define          kMeasurementFlavorPartition 1
#define          kMeasurementFlavorClass    2
#define          kMeasurementFlavorService  3
    uint8_t       pad;

```

```

    MeasurementSample    samples[1];                //count samples here
} MeasurementData, *MeasurementDataPtr;

```

In one implementation, the Data Plane maintains byte and packet counts per traffic class and per partition (excluding the “sampled” packets which the Control Plane has already counted). On a periodic basis, a background task will bundle up samples for active classes and partitions, and forward the data back to the Control Plane for recording in MeasurementSample messages. In a particular implementation, the MeasurementSample messages may also include a compressed byte count (the byte count after compression or before decompression) and an uncompressed byte count (byte count before compression or after decompression).

B.7. LittleNote Message

[0073] For diagnostic purposes, the Control Plane may send log data (such as “printf” output) LittleNote messages. At the Control Plane, the events may be logged into the Control Plane “system event” log as well as copied to any configured **syslog** servers. The following illustrates the format of a LittleNote message according to one possible implementation of the invention.

```

typedef struct _LittleNote {
    uint16_t                type;
#define kLittleNoteType    779
    uint16_t                seq;

    uint8_t                level;           // SYSLOG_LEVEL
#define kNoteEmerg        0
#define kNoteAlert       1
#define kNoteCrit        2
#define kNoteErr         3
#define kNoteWarning     4
#define kNoteNotice     5
#define kNoteInfo        6
#define kNoteDebug       7

    uint8_t                pad;
    uint16_t               reserved;
    char                   note[1];        // null terminated
} LittleNote, *LittleNotePtr;

```

C. Example Process Flows

[0074] Figure 5A illustrates an example process flow, according to one possible implementation of the invention, executed by the network processing unit 300.

When network processing unit 300 receives a packet (502), the hardware of the network processing unit 300 includes a packet parsing logic circuit that parses a received packet and computes a hash of an attribute tuple of the received packet. In one implementation, the tuple comprises the source IP address, destination IP address, source port number, destination port number, and a protocol identifier (such as TCP, UDP, etc.). Network processing unit 300 uses this 5-tuple to identify different data flows between hosts. In one implementation, the values of these fields are arranged in an order, depending on the direction the packet is traveling, to ensure that the resulting hash is the same for data flows of the packet transmitted in either direction. In one implementation, the hardware-computed hash is a 16-bit hash. A secondary longer (e.g., 32-bit) hash, or a hash using a different algorithm, of the same 5-tuple is also computed to identify the data flow. In one implementation, the 16-bit hash computed by hardware may map to one or more secondary hashes. To identify a data flow, the 16-bit hash essentially narrows the search space to a subset of the secondary hashes that are mapped to the 16-bit hash. In some implementations, the first 16-bit hash, for packets received that are compressed packets, can be based on a tuple including only the source and destination IP addresses. This hash can also be used as a tag for a work queue entry that is passed to the compression unit 324, and to otherwise identify the data flow prior to decompression. The secondary hash can be computed after decompression. In such an implementation, the 2-tuple can be used to narrow the search space, while the 5-tuple is used to actually identify the flow after compression.

[0075] As Figure 5A illustrates, if the received packet is not an IP packet (504), network processing unit 300 forwards the packet along toward its destination from an egress interface that corresponds to the packet interface on which the packet was received (506). In other implementations, network processing unit 300 and the control plane can be configured to process non-IP packets as well. If the packet is a compressed packet received from a remote device (such as another network application traffic management device) (505), network processing unit 300 passes the packet to a de-compression process that schedules the packet for decompression

by compression unit 324 (507). De-compressed packets, in one implementation, are re-injected into the packet processing path where the operations of Figure 5A are repeated. Furthermore, in connection with processing of the compressed packets, network processing unit 300 may create a flow reference (see below) and associate the flow reference of the compressed packets with a flow reference corresponding to the de-compressed packets that have been re-injected into the processing path. [In an alternative embodiment, the compressed packets, in one implementation, include IPComp headers that identify the original source and destination hosts. In one implementation, a flow reference can be created based on a tuple of the source and destination IP addresses, and run through the compression unit 324, where the uncompressed output is associated with the original flow reference. In one implementation, the uncompressed output can be hashed by software to yield a more granular flow tuple (including port numbers and protocol identifiers) that is used instead of the originally computed hash.]

[0076] Furthermore, if the received packet is a control message (see above) from the Control Plane (508), network processing unit 300 passes the control message to a process that programs the data plane by changing one or more attributes of the data plane configuration (510). For example, network processing unit 300 may create a flow block in response to a FlowInfo message. The FlowInfo message may identify a partition and a priority for the data flow. In one implementation, flow blocks are identified by the FlowIndex values of FlowInfo messages. See also Figure 5C, discussed below. Otherwise, network processing unit 300 determines whether it has a flow reference that matches the hash it previously computed for the packet (512). If not, network processing unit 300 creates a flow reference in response to the data packet (514). A flow reference includes a key (typically a hash of the 5-tuple attribute values, see above), a FlowIndex value (indexing into an array or table (or other data structure) of flow blocks, and a flow instance identifier. When initially created, the flow reference includes a null FlowIndex value, which may subsequently be modified to index to a flow block when created. As discussed in connection with Figure 5B, however, the FlowIndex value may be set to an entry of cached flow information in a triggering cache.

[0077] As Figure 5A illustrates, network processing unit 300 may be configured to sample packets to network application traffic management unit 200 (518), if there is no flow block created for the data flow (516). If there is no flow block for the data flow, network processing unit 300 may apply one or more default policies to the packet (520). For example, network processing unit 300 may assign the packet to a default partition having a maximum bandwidth parameter enforced by scheduler 318. If a flow block exists (516), network processing unit 300 determines whether to sample the packet to network application traffic management unit 200 (522, 524). For example, the sendmore attribute of the flow block may be set to true. In addition, the received packet may be a data flow or connection-terminating packet (such as a TCP FIN or RST). In a particular implementation, network processing unit 300 is configured to transmit connection-initiating (e.g., handshake or TCP SYNs and SYN/ACKs), and connection-terminating packets (e.g., TCP FINs, RSTs, etc.) to allow network application traffic management unit 200 to set up and tear down data structures as required. Lastly, as Figure 5A illustrates, network processing unit 300 then stores the packet in a flow queue corresponding to the flow reference (526). In one implementation, if there are no packets currently stored in the flow queue, the process generates a message to the partition identified in the flow object (or the default partition), which causes the network processing unit hardware to generate a work entry identifying the partition. As discussed in more detail below, a policy enforcement mechanism schedules packets from the flow queues for output using a hierarchical partition configuration.

[0078] In one implementation, the internal processes of network application traffic management unit 200 assume that a data flow has terminated if a packet associated with the data flow has not been encountered in a threshold period of time. Termination of a data flow may cause the network application traffic management unit 200 to tear down various data structures for the data flow (to allow the memory space to be used for other data flows). In such implementations, the network processing unit 300 may be configured to periodically sample packets to network application traffic management unit 200 (even after sendmore has been set to false) to ensure that the network application traffic management unit 200

does not deem the flow terminated. The rate at which these packets are sampled will depend on the configuration of the network application traffic management unit 200 and the threshold values it uses to deem flows terminated. In such an implementation, the decisional logic represented in 522 of Figure 5A can be augmented to include a check that compares the last sample time to the current time and to conditionally sample the packet if the time difference is greater than a threshold.

C.1. Server-Side Triggered Policies

[0079] According to the process flow illustrated in Figure 5A, the Data Plane applies one or more default policies to new data flows until it receives a FlowInfo message from the Control Plane identifying one or more policies to be applied. Figures 5B and 5C illustrate an alternative implementation where the Data Plane may possibly apply cached policies, as opposed to default policies, to new data flows that match one or more criterion, such as server-side related attributes of IP address and port number.

[0080] As discussed above, the Data Plane may maintain a set of flow objects in a reserved memory space. The flow objects include flow attribute information and one or more applicable policies (see above). The Data Plane may also reserve memory (a triggering cache) for the caching of flow attribute information for possible re-use for subsequent data flows. As discussed above, the Data Plane may store in a triggering cache certain flow information responsive to the value of the trigger parameter in FlowInfo messages. Figure 5C illustrates an example process that a Data Plane may apply to a received FlowInfo message. As Figure 5C shows, when the Data Plane receives a FlowInfo message, it maps the FlowIndex value in the FlowInfo message to a flow object entry in the reserved memory space (560), and saves the flow configuration data in the Flow Block entry (562). If the trigger parameter is set to "remember" (564), the Data Plane then accesses the triggering cache to identify whether a matching entry exists (566). A cache entry may comprise a key value, a time stamp, and flow attribute and configuration information (such as service identifiers, policies and the like). The key comprises a

tuple of an IP address, port number and a protocol identifier. In one implementation, the IP address and port number used for the key is determined with reference to the serverside parameter in the FlowInfo message. That is, the Data Plane uses the serverside parameter to identify the IP address and port number of the server of the data flow and uses this information and the protocol identifier to generate the key.

[0081] In one implementation, the cache is organized as a set of hash groups, where each hash group includes a set of N cache entries. In addition, each cache entry is also addressed relative to an index value that maps to the memory space reserved for the cache. To identify a possible matching entry, the Data Plane may compute a hash of the key to select a hash group, and then search (possibly computing a secondary hash) the entries for a key that matches. If a matching entry in the cache exists, the Data Plane updates or replaces the flow configuration information in the cache entry and updates the time stamp (568). If no matching entry exists, the Data Plane selects a cache entry within the hash group to store the flow configuration data (570), setting a time stamp as well. If there is no empty slot in the hash group, the Data Plane may use a least-recently used (LRU) algorithm to select an entry based on the value of the time stamps.

[0082] The cached flow information may be used for subsequent data flows, where the server-side attributes of the flow match the cached information. In other words, use of the cached flow information may be triggered by the destination (server-side) IP address, destination (server-side) port number and protocol identifier of the initial packet of the data flow. How the Control Plane sets the trigger value of the FlowInfo message may also be specified by one or more policies. The configuration of server-side triggered caching policies can be based on a variety of factors, and can be fully automated, partially automated based on a rule set, and/or manually performed by a network administrator. For example, a network administrator may have configured a traffic classification based on a network application where the server port is static. The Control Plane may set the triggering policy to "remember" (1) by default. In addition, a network administrator may have configured a traffic class with a matching rule based on a host list. In one implementation, the Control

Plane might, as a default operation, want the Data Plane to cache the partition and policies for a server found on a user-created host list. The traffic classification database of the Control Plane may be configured to associate server-side trigger policies to various network application types. Application of the server-side triggered caching policies can be based on a variety of factors associated with the behavior of the network application or service type and/or other factors, such as the configuration of the network application traffic management device 30. For example, the Control Plane may implement a default rule where server-side triggered policies are implemented for all traffic classifications, except where the traffic classifications are based on network applications that have certain specified behaviors. For example, network applications where data flows use dynamic port number assignments, as FTP data flows, may not be suitable candidates for caching server-side triggered policy information. However, information relevant to FTP control flows can be cached since the server-side port for the control flow is generally static. In addition, a traffic classification database may store a traffic class based on known network applications, such as YouTube, MySpace, among others. Recognizing that the IP address and port number of the sites hosting such network applications is not likely to be dynamic, a server-side trigger policy may be used for the traffic class that causes the trigger value to be set to "remember" when data flows hit that traffic class. Accordingly, packets of subsequent data flows sourced from a client host initiating a connection to a video sharing site at the IP address and port number will hit the cache, causing the cached policies to be applied to the data flows. In addition, server-side triggered caching policies can also be based on whether the network administrator has specified a default policy for a given traffic class, and an exception list (based on client IP address or other attribute). Use of an exception list may indicate that flow information caching may not be suitable. Still further, traffic classification can also be based on attributes that are orthogonal to server-side attributes, such as diffserv or DSCP markings. The Control Plane may be configured to set the trigger policy to "don't" (0) or "service-only" (3) for such traffic classes.

[0083] Figure 5B illustrates a process flow that may be implemented by the Data Plane in connection with server-side triggered policies. The illustrated process is quite similar to that described in connection with Figure 5A. However, in Figure 5B, after the Data Plane creates a flow reference (514), it accesses the triggering cache to determine if a matching cache entry exists (550). To find a match, the Data Plane generates a key comprising the server-side IP address, server-side port number and the protocol identifier contained in the received packet. As discussed above, the Data Plane may hash this key value to identify a hash group, and then search the entries of the hash group to find a matching cache entry. Additionally, a cache hit may also be conditioned on the value of the time stamp in the cache entry relative to a time out value (triggerexpiration) transmitted by the Control Plane in a SizingData control message (see above). If no matching entry is found, the Data Plane operates as discussed above relative to the data flow. If a matching entry is found, however, the Data Plane sets the Flow Index value in the flow reference to the index value (K) of the matching cache entry, and sets a flag associated with the packet that causes the data plane to insert a header when sampling the packet to the control plane (see Figure 8, and accompanying description) (552). Accordingly, when the process proceeds to step 516, the flow reference identifies a flow index value, causing the cached policies (partitions, compression settings, and/or priorities) in the cache entry to be applied (526), as opposed to the default policy (520). However, if the trigger parameter has been set to 3 (ServiceOnly), as discussed above, then the Data Plane applies the default policy. This configuration allows the Data Plane to apply desired or more appropriate policies to the data flow from the first packet, as opposed to later in the data flow. The Data Plane will also continue to sample the packets until a FlowInfo message for the flow is received. Other implementations are possible. For example, after setting the packet flag and flow index value (552), the process flow may proceed directly to sampling the packet (524).

C.2. Packet Sampling

[0084] Figures 6, 7 and 8 set forth process flows, according to one particular implementation of the invention, directed to sampling packets to network application traffic management unit 200. Figure 6 is a state diagram showing when, and the aggregate rate at which, packets across all data flows are sampled to network application traffic management unit 200. Figure 7 illustrates a process implemented by the data plane when sampling packets to network application traffic management unit 200.

[0085] As Figure 6 provides, network processing unit 300 is configured to sample packets to network application traffic management unit 200 at a packet sampling rate (602). If the number of packets to be sampled would exceed the current packet sampling rate, network processing unit 300 samples a subset of the packets. For example, if network processing unit 300 receives during a given interval 100 packets to be sampled per second, and the current sampling rate is 50 packets per second, network processing unit 300 would not sample 50 of the packets in a one-second interval. Selection of which packets to sample or not can be based on one or more sampling policies, the details of which are beyond the scope of this disclosure. As Figure 6 illustrates, as long as the network processing unit 300 does not receive an OverloadStatus message indicating that network application traffic management unit 200 is overloaded (604), it continues sampling packets at the packet sampling rate, periodically growing the packet sampling rate (606, 608). However, if network processing unit 300 receives an OverloadStatus message indicating that network application traffic management unit 200 is overloaded (604), it stops sampling packets to network application traffic management unit 200 and decreases the packet sampling rate (in one implementation, by half of the current sampling rate) (610). As Figure 6 shows, network processing unit 300 discontinues packet sampling until it receives an OverloadStatus message indicating that network application traffic management unit 200 is ready to receive packets (612). In the implementation shown, network processing unit 300 may also transmit a ReTransmitRequest (616) asking for an OverloadStatus message at periodic intervals (614).

[0086] As illustrated in Figure 7, network processing unit 300 may be configured to conditionally sample packets to network application traffic management unit 200. For example, if a host identified in the packet is identified as a bad host in a BadHostInfo message (702), network processing unit 300 may drop the packet entirely (703), or pass the packet to a process that decides whether to drop the packet or pass the packet through. Furthermore, if network application traffic management unit 200 is currently overloaded (704), network processing unit 300 does not sample the packet. Otherwise, network processing unit 300 samples the packet to network application traffic management unit 200 (706), and resets a last sample time value of the flow block. The last sample time value can be used to control sample of packets, such as the leaking of packets to network application traffic management unit 200, see above. As discussed above, in one implementation, network processing unit 300 samples the packet by transmitting it to network application traffic management unit 200. Network processing unit 300 then resumes processing of the packet, such as applying one or more policies.

[0087] Figure 8 illustrates an example process directed to sampling packets to the Control Plane, where the Data Plane constructs and inserts service identifying headers into the first packet of a flow. Referring to Figures 7 and 8, when the Data Plane samples a packet to the control plane (706), it determines whether the packet of the data flow matches a cache entry (802), and whether the packet flag has been set (803). If so, the Data Plane constructs a header including a service identifier contained in the matching cache entry, and inserts the header into the sampled packet (804). In a particular implementation, the header is a 4-byte header comprising a 16-bit operation code identifying a service hint type, and a 16-bit value of the service identifier. In one implementation, the header is inserted after the MAC layer header, where VLAN tags are normally inserted. The Data Plane then sends the sampled packet to the Control Plane (806), resetting a last sample packet time. When the network interface returns an indication that the sampled packet has been sent (808), the Data Plane, rather than deleting the packet from the buffer, passes the packet to a policy enforcement mechanism (812), removing the header if one has been added (809, 810).

[0088] Figure 9 shows a process that a Control Plane may apply to sampled packets. As Figure 9 illustrates, when the Control Plane receives a sampled packet (852), it scans for a header including the service hint operation code (854). If a header exists, the Control Plane extracts the service identifier from the header and adds it to a data structure, such as a flow object, for the data flow (856), and passes the sampled packet to the classification mechanism (858). The existence of a service identifier in the flow object indicates to the classification mechanism that service type identification for the data flow has been accomplished. This saves computing resources for other tasks, since the Control Plane need not process the data flow to identify a service type and can proceed to classifying the data flow.

[0089] In the implementation described above, a large portion of the network processing is offloaded to the network processing unit 300, which with its dedicated hardware-level processing features allows for faster processing of network traffic. In the implementation described above, the network processing unit 300 handles network traffic using pre-existing programming. If it does not have a record of a flow and its class, policy, or partition (or matching cached information), it applies defaults to the traffic, and samples the traffic to the network application traffic management unit 200. In this manner, the performance requirements on the network traffic management unit 200 are significantly reduced since it sees only a limited subset of the traffic (typically, the initial packets, one or more leaked packets to prevent flow termination processes of the Control Plane, and possibly terminating packets of a data flow). The network application traffic management unit 200 can classify the traffic fully and report back the class, partition, and other policies (e.g., compression, etc.) of the data flow when it is done. In the meantime, the network processing unit 300 continues to use defaults or server-side triggered cached policies until it receives programming for the specific flow. Once programming is received, it handles the traffic using the policies specified by the network application traffic management unit 200. Further, the use of server-side triggered cache policies provides additional advantages. The major work done by the Control Plane is classifying new flows. Another way to offload more cycles from it is to recognize that once the Control Plane has carefully examined a new flow and

decided it is a given traffic class, it is likely that the next new flow to the same server IP address and port is also likely to be the same traffic class. Because the Data Plane caches this possibly expensive determination and applies it to subsequent new flows (within a reasonable time) made to the same port on the same server host, much duplicate work is eliminated from the Control Plane, allowing it to handle more new flows per minute. Furthermore, server-side triggered policies also facilitate the application of policies and other operations, such as compression, tunneling and protocol acceleration, where it is desirable to begin with the first packet of the data flow.

D. Data Plane Flow and Partition Processing

[0090] Network processing unit 300 also implements bandwidth enforcement functionality to partition bandwidth among data flows. In particular implementations, the bandwidth enforcement and partition scheduling functionality described herein operates separately for each traffic direction. In one implementation, packets are compressed prior to partition scheduling to allow partition scheduling decisions to be based on the size of the compressed packets actually transmitted. As noted above, network processing unit 300 includes a scheduler or work order module 318, in some embodiments. In embodiments that use the Cavium NPU, this work order module 318 might comprise the work order module described in U.S. Patent Application Publication No. 2006/0056406, entitled "Packet Queuing, Scheduling, and Ordering", which includes a hardware data structure for the ordered and synchronized processing of atomic entries in a FIFO work-queue in parallel. In particular embodiments, the processes described herein might make use of this hardware data structure in program modules that execute on the NPU's multiple processors. In particular implementations, a work order module 318 schedules work entries for multiple processor cores in a manner where a data structure labeled atomic is operated on by one processor at any given time. [0091] The work order module 318 implements hardware work queuing, hardware work scheduling and tag-based synchronization and ordering to queue and schedule work for the processors. A work queue entry may be added to a work queue by the

packet input unit 214 for each packet arrival, and by other processes, such as a processor core executing a partition task. The work queue entry is the primary descriptor that describes work to be performed by the cores. Work is defined to be any task to be performed by a core that is identified by an entry on a work queue. The task can include packet processing operations, for example, packet processing operations (such as identifying flow objects, and partition scheduling operations) to be performed on a received packet identified by a work queue entry on a work queue. Each separate packet processing operation is a piece of the work to be performed by a core on the received packet stored in memory. The work order module selects (i.e. schedules) work for a processor and returns a pointer to the work queue entry that describes the work to the processor. Each work entry may have a tag. The tag can be used by the processor to order and synchronize the scheduled work. The tag includes a tag value and a tag type. The tag type indicates how the work is synchronized and ordered. There are three different tag types: ordered (ordering is guaranteed), atomic (ordering and atomicity are guaranteed) and null (no ordering is guaranteed.) A processor core accesses a data structure (or has a tag) atomically when no other processor cores can simultaneously access the same data structure. Thus, atomicity guarantees that at most one processor core has a given tag. The initial tag associated with the work is dependent on the originator of the work.

[0092] The work order module 318 may also include a timer unit used to schedule work for the processors. The timer unit, in one implementation, manages rings having a programmable number of buckets stored in a level 2 cache memory or DRAM. Each bucket stores a list of timer entries. Each timer entry stores a work queue entry pointer. The time period between buckets in each ring is programmable. At each periodic bucket time expiration, the next bucket in the ring is processed and the work defined by the work queue entry pointers stored in the timer entries in the bucket is scheduled.

[0093] A processor requests work from work order module 318. Typically, the processor polls the work order module 318 to find work. However, in some cases the work order module 318 can be selected to interrupt a processor core when it has

work for it. The work order module 318, in one implementation, selects or schedules the work to the processor core based on the groups from which the processor core accepts work.

[0094] Some embodiments might associate a packet with a work-queue entry and use the entry's 32-bit tag (which can be labeled "atomic") as an index into an array of partitions (e.g., where a partition might be a struct or class), each of which includes the tag/index of its parent partition, if any (e.g., as a field in a struct or a data member of a data object). Then, in order to transfer a packet from a partition to its parent (or to transfer a newly-arrived packet to its initial partition), a process might receive a work-queue entry corresponding to a packet from the work queue, perform one or more operations on the partition data structure associated with the atomic tag, identify the atomic tag associated with the parent partition, and add the work-queue entry, modified to include the atomic tag of the parent partition, back onto the work queue maintained by the work order module when it, for example, passes the packet from a child partition to a parent partition. Here it will be appreciated that if two or more work-queue entries have the same tag, only one work entry having that tag can be processed by the NPU's processors at one point in time, according to a particular embodiments of the ordered and synchronized processing described above. In particular implementations, when a processor forwards a packet to a parent partition, it accesses a routine that returns the tag of the parent partition. A work entry including the identified tag is then modified to include this tag and submitted to the work order module.

[0095] As discussed above, packets received at the data plane are associated with corresponding data flows based on source and destination attributes of the flow. In a particular implementation, the network processing unit hardware, when it receives a packet on a network interface, reserves memory in a packet buffer space, allocates a work queue entry, and adds a pointer to the received packet stored in the buffer to the work queue entry. The data plane hardware may map a tuple of packet attributes to a tag, which is also added to the work queue entry. As discussed above, this tuple of packet attributes may map to a flow data structure stored in the memory. This flow data structure may identify one or more

bandwidth policies, such as a partition and/or a priority, as well as compression policies. When this work entry is provided to a processor, the processor may map the packet to a partition; the processor may also map the packet to a priority. Still further, for new flows, the data plane may sample one or more packets of a given data flow to the Control Plane, which classifies the data flow and assigns one or more policies (such as a partition and a priority). For such new flows, a default or cached partition may be identified for the data flow. Still further, after a data flow for a given packet has been identified, the network processor stores the work queue entry for the packet in a flow queue. If the queue for the data flow was empty, the processor changes the tag of the work queue entry to the tag corresponding to the partition and re-submits the work entry to the work order module 318. As discussed below, the work order module distributes work entries to processors to perform processes related to the partitions, such as re-configuring the partition or passing a packet to a parent partition.

[0096] The partition initially assigned to packet of a data flow may be a leaf partition or a non-leaf partition. The priority indicates the packets priority relative to other packets. In one implementation, each partition comprises N partition queues, where N is the number of priority levels that can be assigned. In one implementation, packets placed on a higher priority queue are selected for processing, relative to a given partition, before lower priority queues. In other implementations, weighted fair queuing algorithms can be used to arbitrate among the priority queues. In some implementations, to promote fairness among data flows, only one packet of a given data flow is processed in the hierarchical partition configuration at a given time. In such embodiments, a work queue-entry (or token packet) might include a pointer to a FIFO flow queue which is accessed to process the next packet in the data flow after the work-queue entry corresponding to a data flow has passed through the root partition.

[0097] Some embodiments of that system utilize a partition tree implemented as a hierarchy of independent token buckets, where each partition is a bucket which releases packets at its assigned bandwidth rate and the lowest bucket in the hierarchy is the root partition (e.g., inbound or outbound). In these embodiments,

one might think of each bucket as having a “control knob” and as having a “high water mark” and a “low water mark”, which represent configured minimum and maximum settings. In these embodiments, a bucket regularly suggests a target rate to its children. Each of the children will use the suggested target rate provided it is between the child’s configured minimum and maximum settings. Since all children receive the same target rate, bandwidth can be fairly shared between them. Figure 13 is a schematic diagram showing such a hierarchy of buckets, which hierarchy might be used with some embodiments of the present invention.

[0098] Further, in some embodiments, the bucket hierarchy uses tokens, with each partition starting with enough tokens to forward N (such as 2) full size packets. Whenever a partition has a packet that it cannot forward immediately due to insufficient tokens, the partition sets a timer for a small period of time (e.g., using a hardware timer). The hardware has a mechanism that, when the timer expires, returns the work entry to a processor for execution. Upon receiving the work entry, the processor implementing the partition determines how many tokens it can add to its bucket and resumes forwarding packets if the added tokens are sufficient to do so. However, other embodiments might employ a leaky bucket algorithm, rather than a token bucket algorithm, to control target rate in conjunction with a hardware timer.

[0099] It will be appreciated that such a bucket hierarchy and the partition-centric processing described herein is advantageous for several reasons. A non-root partition is independent insofar as it can autonomously reset its “rate knob” to its parent’s current suggested child target rate. This avoids the need to walk-through of the partition tree, resetting all rate knobs at the same time. Moreover, when a partition is configured, it is given an index for its parent. The partition can quickly verify that the index remains valid, before “dripping” a packet down to its parent. If the index is not valid, the partition might use the default root partition and request reprogramming from the control plane. Moreover, the bucket hierarchy prevents CPU cycles from being wasted on update processing for partitions that are not currently receiving traffic.

[00100] Figure 10 is a diagram showing a flowchart of a process for passing a

packet from a first partition to a parent partition or a port, which process might be used with some embodiments of the present invention. It will be appreciated that the process shown in Figure 10 might run on one or more processors of a multiprocessor network processor unit in the data plane. The process flow illustrated in Figure 10 is executed in connection with a work entry transmitted to a processor core of the multi-core network processing unit. Processor cores perform the process on a partition identified in the work entry and request more work entries from the work order module. In addition to work entries that cause packets to traverse partitions, other work entries may include jobs that cause a processor core to change the configuration of a partition (such as minimum and maximum bandwidth parameters, parent child relationships). In some implementations, these work entries include tags that correspond to the respective partitions on which they operate. It will be appreciated that using the same work order synchronization mechanism for scheduling and configuration of a partition allows altering the partition tree programming without locking the entire or substantial portions of the tree or using other software synchronization mechanisms that could potentially stall packet processing.

[00101] In the process's first step 1002, the partition receives a packet from a child partition in the form of a work entry removed from a work queue such as the work queue provided by the work order module of the network processor unit 300. Alternatively, the partition might be the initial partition, such as a leaf partition, associated with a data flow, in which case it will receive the packet (i.e., in the form of a work entry) as a packet from a process that associates packets to data flows and adds packets to flow queues (see above). The partition process then goes to step 1004, where a conditional operation is performed. Specifically, the partition process periodically adjusts its target rate and also computes a new suggested target rate for child partitions. In one implementation, at initialization of the network processing unit 300, a task is created that is called at a periodic time interval (such as 1 second). At each periodic time interval, the task is called to increment a global heart beat counter. The partition processes, responsive to detected changes in the heart beat counter, compute new suggested target rates for their respective child

partitions (1010) and, for non-root partitions, adopt the child target rate suggested by respective parent partitions within minimum and maximum bandwidth constraints (1006). As discussed below, a partition maintains a count of its active child partitions by incrementing a counter each time a work entry for a packet is flagged with a reserved bit. Child partition processes are configured to set this flag for only one packet after the global counter increments. In this manner, the parent partition estimates the number of active child partitions by accessing the child counter at the global counter interval, which essentially is a count of the child partitions active over the previous interval. In one implementation, the child counter is reset at each global counter interval (1010).

[00102] To compute a suggested child target rate, the partition process divides its target rate by the number of active child partitions (as indicated by the child counter value). In addition, for non-root partitions, the partition process checks the target rate suggested by its parent partition and conditionally adopts it, setting the partition rate to the minimum rate if the suggested is lower than the configured range, or the maximum rate if the suggested rate is higher than the configured range. The partition process then adds tokens to its bucket based on the current target rate for that partition (1012), see below. If the packet is flagged as the first packet from that child partition during the predefined time interval (1014), the partition process increments an active child partition counter (1016) and then schedules the packet for transmission (1018) (see also Figure 12).

[00103] In another implementation, the computation of a suggested target rate for child partitions further accounts for the amount of data stored in a queue for the partition. A partition's "size" is the maximum number of bytes it is allowed to queue. This should normally be adjusted for the partition's configured maximum rate, and it is based on a system's resource allocation scheme. In one implementation, each partition is allowed to store up to 1 second of traffic (`gMaxPartitionDelay`) at its maximum rate. So, a partition with a 1 mbps maximum rate (`ptn->pMaxRate`) would be allowed to store 125 Kbytes (`ptn->pMaxBytesQueued`) (1 million bits / 8 bits-per-byte * 1 second). A low water mark (`ptn->pLowWaterBytes`) for the partition can be defined as 1/3 of the size of the

partition (`ptn->pMaxBytesQueued`), while a high water mark (`ptn->pHighWaterBytes`) can be set to 2/3 of the size of the partition. Other high and low water levels can be used. When a packet is received, the partition process drops the packet if the amount of queued data (`ptn->pBytesQueued`) resulting from adding the packet would exceed the size of the partition. Otherwise, the partition process adds the number of bytes in the packet (`packet->length`) to `ptn->pBytesQueued`. When a packet is forwarded, the partition process subtracts `packet->length` from `ptn->pBytesQueued`. The following pseudo-code illustrates how a partition may determine the suggested child target rate based on the amount of queued data for the partition. This determination, as above, is computed at each heartbeat or rate adjustment interval.

```

oldSuggestedTargetRate = ptn->pChildTargetRate;
newSuggestedTargetRate = ptn->pTargetRate / ptn->pNumChildren;
newSuggestedRate = (oldSuggestedTargetRate + newSuggestedTargetRate) / 2;

if (partition is full) {
    newSuggestedRate = newSuggestedRate / 2;
}
else if (partition is above high water mark) {
    if (partition is filling up (more full than last interval)) {
        newSuggestedRate = newSuggestedRate * 0.8;
    } else if (partition is emptying (less full than last interval)) {
        newSuggestedRate = newSuggestedRate * 0.95;
    }
}
else if (partition is between high and low water mark) {
    if (partition is filling up) {
        newSuggestedRate = newSuggestedRate * 0.99;
    } else if (partition is emptying) {
        newSuggestedRate = newSuggestedRate * 1.01;
    }
}
else if (partition is below low water mark) {
    if (partition is filling up) {
        newSuggestedRate = newSuggestedRate * 1.05;
    } else if (partition is emptying) {
        newSuggestedRate = newSuggestedRate * 1.2;
    }
}
else {
    // partition is empty
    newSuggestedRate = newSuggestedRate * 1.25;
}
ptn->pChildTargetRate = newSuggestedRate;

```

By using the foregoing algorithm, the partition can manage its own children and its queued backlog while still forwarding traffic at the rate its parent partition has advised it to use. It will tend to keep the hierarchical partition configuration working efficiently by avoiding partitions either becoming empty or full too often. Of course, the values used to adjust the child target rate can be tuned for the individual application based on a variety of engineering and design considerations.

[00104] In a particular implementation, tokens are implemented in each partition as a counter. The counter contains the number of bytes that the partition is allowed to send without pausing. The counter does not go below zero and, in some implementations, is capped at a maximum number. When a partition is created, it is given a configured amount of “tokens”, and its initial update timestamp is set.

```

    ptn->pLastUpdated    = eonGetClockCount();
    ptn->pBytesAllowed    = PTN_STARTUP_BYTES_ALLOWED;
    ptn->pMaxBytesAllowed = 2 * (pip->maxbw / 8);

```

[00105] This notes that the partition was last updated at the current clock count. It is given a starting number of bytes it can send

(PTN_STARTUP_BYTES_ALLOWED, in one implementation, 6000 bytes). It also has a maximum token credit limit based on its maximum configured bandwidth setting. So, for a partition with a maximum rate setting (pip->maxbw) of 100 mbps, for example, the maximum bytes the partition accumulates would be 25,000,000 bytes, which would drain in 2 seconds at the maximum rate.

[00106] Figure 11 illustrates a method according to one implementation for updating the tokens for a partition. Every time the partition is activated, it updates its own token counter, based on its current target rate. In a particular implementation, it determines how many clock cycles have passed since it was last updated, then it adds in the correct number of new tokens, and it finally checks to make sure it isn't exceeding the allowed limit. As Figure 11 illustrates, the partition process first determines the number of clock cycles since the last update, where clockNow = eonGetClockCount() and clockDiff = clockNow - ptn->pLastUpdated (1030). The partition process then computes the amount of bytes that could be transmitted in the clockDiff time interval based on the current target

rate (ptn->pTargetRate) (1032). The partition process then updates the update the time stamp of the partition (ptn->pLastUpdated) (1034) and sets the number of tokens for the partition (ptn->pBytesAllowed) based on the computed number of bytes in step 1032, or the maximum bytes allowed for the partition (ptn->pMaxBytesAllowed) (1036, 1038, 1040).

[00107] Figure 12 sets forth a method directed to scheduling packets of a partition for transmission to either a parent partition or an interface port. When a partition gets a new packet, it queues the packet in a partition queue (1051) and checks to see if it can send it. In one implementation, a priority determines to which partition queue the packet is added. The first check is to determine if a previous packet has already been delayed (1052). If there are no other previous packets stored on the partition queue (1052), the partition process determines whether sufficient tokens exist to forward the packet (1056). If not, the partition process computes a delay interval and sets a timer based on the computed delay interval (1058). The hardware of the network processor unit, as discussed above, has a mechanism that, when the timer expires, returns the work entry to a processor for execution. In one implementation, the delay interval can be computed based on the difference between the size of the packet (packet->length) and the current number of tokens [bytesNeeded = packet->length - ptn->pBytesAllowed], and the current target rate. In one implementation, the delay interval is computed by dividing the number of bytes (tokens) needed to transmit by the current target rate. In another implementation, however, the delay interval can be tuned for the speed of the networks involved. For example, rather than setting up many very short timers that are triggered at short intervals, a longer delay that causes timers to be triggered less frequently can be used. In a multi-gigabit system, such delays are not usually noticeable, and are more efficient for the system to implement. In a particular implementation, if the partition needs (bytesNeeded) less than 2000 bytes, the delay is set, based on the current target rate (ptn->pTargetRate), to allow 2000 bytes (tokens) to be accumulated. In a large system with many partitions, this can be more efficient.

[00108] As Figure 12 illustrates, however, if the partition has enough tokens to

send the packet (1056), the partition process deducts tokens from the partition based on the size of the packet (packet->length) (1060), conditionally setting a packet flag if the packet is the first packet forwarded since the last rate update interval (1062, 1064), and forwards the packet to the next partition or port (if the current partition is a root partition) (1066). In one implementation, the partition process can attempt to empty its partition queue (1068) if sufficient tokens exist (1056).

[00109] To forward a packet to a parent partition (1026), the partition process accesses a routine that returns the tag corresponding to the parent partition, modifies the work entry for the packet to add the parent partition tag, and resubmits the work entry to the work order module 318. If the partition process currently implements a root partition process, it forwards the packet to a process, by transmitting the work entry, that causes the packet to be transmitted from a network interface. For root partitions, the partition process schedules the packets for transmission from a network interface. When the packet is transmitted, the work entry for the packet is transmitted back to the flow queuing process that deletes the packet from the buffer, and checks whether another packet work entry is stored in the flow queue. If so, the flow queuing process, as discussed above, changes the tag to the initial partition tag, causing the packet to be processed through the hierarchical partition configuration described above and ultimately scheduled for transmission. The embodiments described above can be configured to achieve a number of advantages. For example, since a partition process can be implemented to forward packets without accessing data structures of other partitions (other than to check for suggested rates of parent partitions), the processing described above can be implemented in a multi-processor environment in an efficient manner, such as without software locking or other delay-inducing methods.

[00110] In some implementations, the root partition may be implemented by a set of partition processes with different tags. It will be appreciated that a root partition can become a bottleneck since every packet must pass through it. So to prevent such a bottleneck, a particular embodiment employs a technique called

“root spreading”. In this technique, a root partition is implemented as a set of k (for example, 4) “ghost partitions”, where each ghost partition comprises $1/k$ of the total or maximum bandwidth in a given network traffic direction, inbound or outbound, of the root partition. Using this technique, up to $2*k$ processors in an NPU might be able to simultaneously work on root partition tasks, instead of 2 (i.e., one processor for the inbound root partition and one for the outbound root partition), for a given full-duplex access link. When a child partition passes a packet to a parent partition (1066), it looks up the atomic tag of the parent, accessing a routine that normally returns a tag. For root partitions, however, each processor, in reserved memory, maintains counters and other information directed to looking up a tag to be used to access a given ghost partition corresponding to the root. In one implementation, each time a processor looks up a tag for a root partition, the processor accesses a counter, c , and increments it. The processor then takes c modulo k which, in connection with the direction (inbound or outbound) indexes to one of the $2*k$ ghost root partitions to use. The processor then modifies the work entry for the packet to include the tag corresponding to the ghost root partition which is resubmitted to the work order module 318.

E. Beltway Mechanism to Reduce Overhead Associated with Atomic Scheduling [00111] As discussed above, the partition scheduling process, viewed as a whole, is implemented by a plurality of processors operating on a sequence of work queue entries, and relies on the work order module 318 to provide for atomic access to shared data structures when necessary. There are a number of other processes or operations that also require atomic access to shared data structures. Such operations are typically much smaller operations, such as writing a packet to an output port of a network interface, passing a packet to a compression module, or setting a timer. Indeed, the partition scheduling process discussed above utilizes these operations. For example, execution of the root partition process results in a packet being written onto a hardware queue of a network interface or output port. The output port hardware operates to separately transmit the packet. Furthermore, the partition scheduling processes also rely in some instances on

hardware-based timers, as well.

[00112] Access to hardware resources (such as output ports, timers, compression modules, etc.) are somewhat constrained in multiple processor environments. More specifically, hardware resources are often accessed via a defined set of registers (e.g., configuration and status registers (CSRs)) and command buffers, which may be accessed by only one processor at a time. For example, a network port is accessible via a set of registers, to which a processor is provided atomic access. To write a packet onto a given output port, a processor typically accesses a command buffer to write one or more commands and packet pointers, and writes commands (e.g., operation codes or opcodes) to configuration and status registers. Atomic access to these registers and command buffers is required because another processor should not be able to write commands to these registers or command buffers during execution of a current operation. Other examples are hardware-based timers or compression units. In a multiple processor environment, the processors, and possible other processing units (e.g., regular expression engines, etc.) essentially compete for access to the hardware registers and command buffers used for various functions, such as writing packets to output ports, setting the timers, or compressing packets.

[00113] Using spin locks to control access to shared data structures on such hardware facilities is inefficient and can lead to significant throughput degradations in systems incorporating a large number of processors operating in parallel. The work order module 318 does have facilities that may be utilized to provide for atomic access to such data structures, as discussed above. Use of such facilities, however, comes with a cost. Specifically, there is overhead associated with obtaining work queue entry buffers, storing data in a work queue entry, submitting the work queue entry to the work order module 318, as well as receiving new work queue entries from the work order module 318. While this overhead is acceptable for processes that perform a number of operations in connection with a single work queue entry, this is inefficient for smaller operations (such as setting timers, passing packets to compression logic, etc.).

[00114] The following describes a so-called beltway mechanism that takes

advantage of atomic locking mechanisms supported by certain classes of hardware processors, such as the work order module 318 discussed above, to handle the tasks that require atomic access to data structures without wasting processor cycles spinning while also reducing the overhead associated with these atomic locking mechanisms. The beltway mechanisms described herein can be used to control access to software and hardware facilities in an efficient manner.

[00115] As an overview, the beltway mechanism can be implemented around a given task, such as writing packets to a specific output port or specific group of output ports, setting a hardware-based timer, using hardware-based compression logic, etc. An instance of a beltway data structure is created for each such task. Relative to each task, the beltway data structure includes a set of circular buffers or rings, and global variables that point to consumer processes or contain state information of the particular beltway instance. In the implementation described herein, the number of circular buffers or rings of a given beltway instance equals the number of processors of the data plane hardware. Of course, the number of processors and corresponding circular buffers may be limited by configuration parameters. In addition, so-called producer functions add entries to the beltway data structure, while a so-called consumer process reads entry from the beltway data structure and calls a consumer function associated with the beltway data structure for processing of the entry.

[00116] Figure 14 is a functional block diagram illustrating the relationship between a beltway data structure 1408 and the number of processors (N) 1406. In the data plane 300, discussed above, each processor 1406 includes L1 or on-chip cache memory. The data plane 300 further includes main memory 1402 and L2 cache memory 1404 shared by processors 1406. The beltway data structure 1408 includes circular buffers or rings 1410 for corresponding ones of the processors. As discussed in more detail below, when a particular processor 1406 executes a process that results in adding an entry to a given beltway data structure, the entry is added to the circular buffer that corresponds to that processor. Accordingly, entries added by processor 1 would be stored on circular buffer 1. Each beltway data structure 1408 is maintained in main memory 1402 and accessible by each of the processors

1406. Of course, at run time, portions or all of a given beltway data structure may also be stored in L2 cache memory 1404, as well as the L1 caches of each processor 1406.

E.1. Beltway Data Structures

[00117] As discussed above, a beltway data structure 1408 includes N circular buffers (rings) for corresponding ones of the N processors, as well as global variables, the types of which may vary depending on the type of task to which the beltway is applied (writing to output port, timers, compression, etc.). The following object definition provides a beltway data structure according to one possible implementation:

```
typedef struct _BeltType {
    BeltConsumerFunction    consumer;
    int32_t                 max;
    int32_t                 size;
    int32_t                 limit;
    bool                    scheduled;
    int64_t                 did;
    int64_t                 ran;
    BeltQueue               queues[kMaxProcessors];
} BeltType, *BeltTypePtr;
```

BeltConsumerFunction points to the function that consumes entries in the circular buffers of the beltway data structure, and provides access to a resource, such as a hardware output port or group of hardware output ports. In a particular implementation, each consumer function has exclusive access to a given resource or group of resources, such that there is a 1 to 1 relationship between a given beltway data structure and a corresponding resource or group of resources. The max variable defines the maximum size of each circular buffer of the beltway data structure, while the size variable defines the size of each entry in the circular buffers. The limit variable defines the maximum number of entries a consumer process can consume in a given circular buffer before moving on to another circular buffer in the same beltway data structure. The scheduled variable is a Boolean indicating whether the beltway data structure has been scheduled for processing

(e.g., whether a work queue entry for the beltway module has been added to the atomic scheduling process). The variable `queues[]` is an array of the circular buffers of the beltway module in main memory 1402.

[00118] Additionally, the circular buffers can be defined by the following object definition:

```
typedef struct _BeltQueue {
    MinimalWorkQueueEntry    wq;
    VoidPtr                  ring;
    int32_t                   first;
    int32_t                   last;
} BeltQueue, *BeltQueuePtr;
```

In this implementation, the `ring` variable contains the starting memory address (e.g., a 64-bit pointer) reserved for the circular buffer in main memory. The `first` variable contains the offset (relative to the `ring` variable) corresponding to the first entry in the circular buffer to be read, while the `last` variable is the offset (relative to the `ring` variable) points to the entry following the last entry to be read in the circular buffer. As discussed below, the size variable defining the size of each entry is used to increment the `first` and `last` variables, as needed, when entries are read from or written to a given circular buffer. In a particular implementation, the `first` and `last` variables are defined relative to the `ring` variable as the origin. For example, the first memory space of the circular buffer is zero. According to this convention, if the last entry were the third memory space relative to the `ring`, the `last` variable would be set to the product of two times the size variable. The `wq` variable stores a work queue entry that a given processor, when executing a producer function, may use when submitting a work queue entry to the work order module 318 in order to schedule execution of a consumer process.

[00119] Each entry in a circular buffer can comprise a set of commands specific to the type of task to which a given beltway data structure is associated. For example, the following object definition shows the structure of an entry for passing packets to an output port:

```
typedef struct _OutputBeltEntry {
    CycleTime                time;
    int                      queue;
```

```

OutputCommandWord        command;
WorkQueueEntryPtr        work;
BufferHandle              packet;
} OutputBeltEntry, *OutputBeltEntryPtr;

```

In the foregoing code section, the time variable indicates the current clock counter maintained by the data plane hardware. The queue variable identifies to which port queue supported by the output port the packet should be written. The command variable identifies a command or operation code associated with transmission of the packet. The work variable is a work queue entry returned to the calling function when the transmit operation is complete. The packet variable identifies the memory address of the buffer space storing the packet. In some implementations, a single consumer function can be configured to access a group of output ports, such as 4 output ports.

[00120] Other entry data structures can be defined for other tasks. For example, the following objects define an entry structure for setting a hardware timer:

```

typedef union _TimerWorkEntry {
    WorkQueueEntry        w;

    struct _t {
        uint64_t          wqe[4];
        CycleTime         deadline;
        uint8_t           cancelled;
        uint8_t           istimer;
        uint64_t          location :50;
        uint8_t           packet_data[80];
    } t;
} TimerWorkEntry, *TimerWorkEntryPtr;

typedef struct _TimerBeltEntry {
    CycleTime              when;
    TimerWorkEntryPtr      work;
} TimerBeltEntry, *TimerBeltEntryPtr;

```

In the foregoing, TimerBeltEntry is the entry stored in the circular buffer and includes a when variable, and a work variable that points to a work queue entry

stored in a buffer in main memory. The when variable indicates the clock cycle or time value at which point the work queue entry identified in the work variable should be returned by the timer hardware.

[00121] Still further, an entry structure for a compression task may be defined as follows:

```
typedef struct _CompressBeltEntry {
    WorkQueueEntryPtr      work;
    uint8_t                compress;
    uint16_t               dataOffset;
    uint16_t               dataLength;
} CompressBeltEntry, *CompressBeltEntryPtr;
```

The work variable, as above, identifies a pointer to commands and data used by the hardware compression logic to compress the data. The compress variable contains a command, such as compress, or decompress, used to instruct the compression unit. The work variable identifies a work queue pointer that is returned to the calling function when the packet compression operation is completed. In one implementation, it further points to the input buffer that stores the packet to be compressed or decompressed. The dataOffset variable identifies the starting point within the packet where the compression or decompression operation should be performed, while the dataLength variable indicates (in bytes) the length of the data block to be compressed or decompressed. In another implementation, the CompressBeltEntry can simply contain a pointer to a work queue entry. The compression unit 324 can figure out whether to compress or decompress a packet based on one or more attributes of the packet. For example, packets with IPComp headers are decompressed, while all other packets are compressed.

[00122] As one skilled in the art will recognize, the foregoing are examples purely for didactic purposes. The type, length and structure of each entry can vary considerably and will depend on the type of task, and the application programming or other interfaces used to access the particular functionality.

[00123] In a particular implementation, an initialization function or process can create instances of beltway data structures, allocating memory resources as necessary, in response to commands issued by an application, such as during a

system start up or recovery process. These commands can identify a beltway type (e.g., output port, compression, timer), and a hardware resource, such as an output port, or a group of similar hardware resources, such as a group of output ports to be associated with a given instance of a beltway data structure. In other implementations, the hardware resource or group of hardware resources can be associated with a function (consumer) that is operative to access such resource(s).

E.2. Producer and Consumer Processes

[00124] As mentioned above, a producer function, when called, adds entries to a beltway data structure, while a consumer process reads entries from a beltway data structure in order to process the entries. The consumer process and producer functions, in one implementation, are embodied as processor readable instructions stored in memory that the processors may execute. A producer function may be called in connection with execution of another process. For example, the root partition process described above in Section D may call a producer function to pass a packet to an output port for transmission. Additionally, the partition processes described above may call a producer function to set a hardware-based timer. In addition, a compression process may call a producer function to add a packet to a beltway data structure associated with compression unit 324. Producer functions executed by multiple processors may add entries to the same beltway data structure. The facilities of the data plane hardware handle memory coherence tasks associated with writing entries to memory.

[00125] Figure 15 illustrates a process flow associated with adding an entry to a beltway data structure according to one possible implementation of the invention. The beltway add function receives as input parameters a pointer to a beltway data structure that corresponds to the task or resource to be accessed, and a pointer to command data stored in a stack or other memory location. The beltway add function identifies a circular buffer, and corresponding memory address of the circular buffer (bq), based on the processor currently executing the beltway add function (GetProcessorNumber) (1502). For example, assuming that processor 1 calls the beltway add function in connection with writing a packet from a root

partition to an output port, the circular buffer, and the memory address of the circular buffer, corresponding to processor 1 is identified. The beltway add function obtains the memory address (wqe) for a work queue entry (bq->wq) required to schedule a consumer process with the work order module 318 (1504). The beltway add function then identifies the memory address (be) of the last entry in the circular buffer based on the starting memory address of the circular buffer (bq->ring) and the last variable (bq->last) (1506). The beltway add function computes a next variable by incrementing bq->last by the size of each entry (bt->size) (1508). If the result is greater than the maximum size of the buffer, next is set to zero. The beltway add function then determines whether the circular buffer is full by comparing the next variable to the first variable (bq->first) corresponding to the circular buffer. If the circular buffer is full (1510), the beltway add function returns "false" as an indicator to the calling process to allow it to perform error handling operations or a retry. If the circular buffer is not full, the beltway add function copies the contents of the beltway command indicated at the location of the command entry pointer to the next memory space of the circular buffer, and sets the last variable (bq->last) to next (1512). The beltway add function then accesses the global state variables (bt->scheduled & wqe->scheduled) of the beltway data structure to determine whether a consumer process has been scheduled (1514). If a consumer process for the beltway data structure has not been scheduled, the beltway add function sets both bt->scheduled & wqe->scheduled variables to true, and submits a work queue entry to the work order module 318 to start a consumer process (1516). In a particular implementation, a tag identifying the work queue entry as atomic is added. Eventually, a processor will pick up the work queue entry from the work order module 318 and execute a consumer process, such as the consumer process described below.

[00126] Figure 16 illustrates a consumer process according to one implementation of the invention. In the implementation shown, a processor, when it executes a consumer process, accesses a given beltway and, within certain optional constraints, attempts to process entries across the circular buffers of the beltway data structure. As discussed above, the consumer process, relative to a

given beltway data structure, is labeled as atomic to ensure that no other processor picks up a consumer process implicating the same beltway data structure. In the implementation shown, the beltway consumer process receives a work queue entry (wqe) from the work order module 318. The beltway consumer process identifies the beltway data structure based on the atomic tag of the work queue entry (wqe->tag) (in one implementation, by looking at the low order byte of the tag (ANDing tag with 0xFF) (1602). The beltway consumer process also resets certain state variables (bt->scheduled & wqe->scheduled) to false, and synchronizes these state variables across the N processors. Setting these variables to false causes subsequently called producer functions to submit a work queue entry to the work order module 318, as discussed above. The beltway consumer process then processes the entries, if any, contained in each circular buffer (1604). As to each circular buffer (bq), the beltway consumer process sets the memory address (be) of the circular buffer to be read (bq->ring + bq->first), and fetches the entry at the defined memory address (1608). The beltway consumer process also computes the next variable by incrementing the first variable (bq->first) by the size of the buffer entry (bt->size), setting next to zero if a wrap around occurs (1610). The beltway consumer process then calls or accesses the consumer function (bt->consumer) associated with the beltway data structure, passing the memory address of the entry (1612). In some implementations, this access may be an attempted access to a set of hardware registers, such as configuration and status registers of a hardware output port or a timer mechanism. If the entry is not consumed (1614), the beltway consumer process leaves the entry on the circular buffer and proceeds to the next circular buffer. If the entry is consumed, the beltway consumer process sets the first entry (bq->first) to the next variable computed above (1610, 1616). As Figure 16 illustrates, the beltway consumer process may process additional entries in the same circular buffer until the buffer is empty (bq->first == bq->last) or a configured limit (bt->limit) on the number of entries that can be processed is exceeded (1606). In this manner, a given processor may sequentially process a plurality of command entries in a given beltway data structure in response to a single work queue entry provided by the work order module 318. If command

entries remain in the beltway data structure after a first pass through the circular buffers, the beltway consumer process may either make additional passes through the circular buffers, or resubmit a work queue entry to the work order module 318.

[00127] Other implementations are possible. For example, each beltway entry may include a time stamp, such as a clock counter value. The beltway consumer process may be configured to select entries based on the time stamps, rather than the round-robin arbitration described above. For example, the beltway consumer process may select a circular buffer with an entry having the lowest time stamp, process the entry, and scan for the next entry across the beltway data structure with the lowest time stamp.

[00128] Given that a single processor has exclusive access to the portions of the beltway data structure relevant to the consumer process and that a single beltway data structure is exclusively associated with a hardware resource or group of hardware resources, atomic access to these resources is also ensured to the command entries while reducing overhead associated with the work order module 318. That is, during higher demand situations, the beltway data structure for a given task may store a number of command operations that are executed by a consumer process in response to a single work queue entry provided by the work order module 318 in step 1514 illustrated in Figure 15. Indeed, as demand increases, more command operations are performed for each work queue entry maintained by the work order module 318. Furthermore, processes, such as partition processes, access resources indirectly via a corresponding beltway data structure, as opposed to directly accessing the resource directly. For example, to set a timer, a process would call the beltway add function, as opposed to accessing the timer registers directly. Given that beltway add function operates only to add entries, one or more of the processors may continue to add entries, while another processor executes a consumer process or while the work order module 318 continues to buffer the work queue entry for the consumer process.

F. Compression Across Control and Data Plane Architecture

[00129] Compression functions are executed in a distributed manner between

the control plane and the data plane. In a particular implementation, the control plane handles compression management functions, such as discovery of remote compression units, setting compression policies for individual flows, and the like. In a particular implementation, compression unit 324 of network processing unit 300 is utilized to actually compress and de-compress network traffic. Packets may be passed to the compression unit 324 at various points in the packet processing path of the data plane. For example, compressed packets may be passed to compression unit 324, which re-injects the uncompressed packets into the packet processing path as discussed above. In addition, uncompressed packets can be passed to the compression unit 324 for compression prior to partition processing and other packet processing operations. The specific point at which packets are compressed depends on the exact configuration of the data plane and the packet processing functions that it has been configured to perform. Still further, the beltway producer and consumer functions discussed above can be utilized to facilitate writing commands to the compression unit 324.

F.1. Control Plane Compression Functions

[00130] In a particular implementation, the control plane decides which data flows can or should be compressed based on a set of policies associated with the traffic classifications or network application types. These policies are applied to sampled packets of a data flow (see above). The decision of the control plane is reflected in the kPolicyCompressible bit, where 0 indicates that the data flow should not be compressed and a 1 indicates that the data flow can be compressed. It will be appreciated that the compression policy decisions of the control plane are not mandatory in that a particular packet of a data flow may not be compressed for other reasons, such as the packet size being below a threshold, or insufficient hardware resources to compress the packet.

[00131] Figure 17 illustrates decisional logic that can be implemented by the control plane when setting the kPolicyCompressible bit. As discussed herein, the control plane may classify data flows based on one or more packets sampled by the control plane (see Figure 9, above). These classification operations may yield a

policy that indicates whether the data flow should be compressed. These policies may be default policies or policies configured by a network administrator. For example, encrypted traffic or the data of VoIP traffic is usually compressed and may not be a good candidate for compression (in fact, subsequent compression operations may even expand the data). Still further, a network administrator may decide to implement a policy configuration where only a subset of the network application traffic is to be compressed. As Figure 17 illustrates, if the policy indicates that the data flow should be compressed (1702), the control plane then determines whether there is a compatible de-compression unit in the path to the destination host (1704). If both conditions are met, the control plane sets the kPolicyCompressible bit in the FlowInfo message that is transmitted to the data plane (see above). As discussed above, the control plane may also set the server-side trigger bit to cause the compression policy decision to be implemented for subsequently encountered data flows.

[00132] The control plane may be manually configured with knowledge of the destination IP addresses and/or subnets that include one or more remote nodes with the capability to de-compress the compressed packet. For example, a network administrator that utilizes an MPLS VPN can configure the control plane with the IP addresses and subnets where a compatible de-compression unit is available. In addition, the control plane may also probe paths to various destination hosts or subnets to discover one or more remote nodes in the path with de-compression capabilities. U.S. Patent No. 7,013,342, which is incorporated by reference herein, discloses methods for dynamically discovering remote devices with compatible compression capabilities. To discover such remote nodes, the control plane may inject probe packets in the path to a destination host and monitor for responses, which are passed to it by the data plane for processing. In yet another implementation, the control plane can initially assume that there is a remote de-compression unit in the path to a destination host and set the compression bit. The control plane then learns which paths have remote compression units as error messages (such as Internet Control Message Protocol (ICMP) packets) are received from intermediate or destination hosts.

F.2. Data Plane Compression Functions and Operation

[00133] Compression unit 324 is a hardware-based compression circuit that can be accessed via a defined set of registers (e.g., configuration and status registers (CSRs)) and command buffers, which may be accessed by only one processor at a time. To compress a packet, a processor typically accesses a command buffer to write one or more commands and packet pointers, and writes commands (e.g., operation codes or opcodes) to configuration and status registers. When compression unit 324 completes the command, the command buffer space is free for allocation to another command. In a particular implementation, the network processing unit 300 configures a command buffer space for the compression unit 324, having a fixed number of entries. As discussed below, the state of these command buffers may influence whether a packet is ultimately compressed. A compression command typically identifies a compression operation (compress or decompress), one or more compression policies, a pointer to a packet buffer space storing the input packet, and a pointer to a packet buffer space allocated to the output. When compression unit 324 completes processing, it returns a pointer to the output. For compression operations, if the output is actually larger than the input, a process may cause the original packet to be output.

[00134] Compression unit 324, in a particular implementation, employs a stateless, packet-to-packet compression algorithm. For example, the compression unit 324 may include functionality operative to perform IP Payload Compressing Using DEFLATE, as set forth in RFC 2394, R Pereira, "IP Payload Compression Using DEFLATE" (December 1998). Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", RFC 1951, May 1996, discloses aspects of the DEFLATE compression algorithm, which is a lossless compressed data format that compresses data using a combination of the LZ77 algorithm and Huffman coding. Of course, other stateless and stateful compression algorithms can be used. The selection of a compression algorithm is considered an engineering or design choice depending on a number of different factors, such as throughput requirements, network data traffic types, and the like.

[00135] As discussed above, the compression unit 324 may also use IPComp headers, as disclosed in Shacham et al., "IP Payload Compression Protocol (IPComp)", RFC 2393, December 1998. In one mode employed by compression unit 324, the IPComp header includes the original source and destination address of the packet. IPComp header includes fields set to identify the compression algorithm and optionally a dictionary. In one implementation, the identifiers used in the IPComp header may be set to values outside the well-known identifiers specified by standards or RFCs such that the data plane does not attempt to de-compress packets not intended for it. In other implementations, a unique signature can be used instead. Furthermore, other modes can be employed, such as the use of tunneling/encapsulation protocols.

[00136] In one implementation, the compression unit 324 also supports different compression modes, each with its own processing overhead or resource requirements. Generally, the compression modes can range from a best mode that may be slow, and a fast mode that does not compress as well. For example, compression unit 324, in a first mode, may employ fixed Huffman encoding, which consumes a first number of processing cycles. In a second mode, the compression unit 324 may employ dynamic Huffman coding, which generally consumes a second, larger number of processing cycles. In a third compression mode, the compression unit 324 selects between fixed and dynamic Huffman coding, the processing cycles consumed in this third mode include resources consumed to make the selection plus the processing cycles required to perform the actual compression. In a particular implementation, the compression process may instruct the compression unit 324 as to which mode to employ by configuring appropriate command words or setting bits in the command buffer entry for the packet compression operation. As discussed below, a compression operation may select from one of these modes based on observed load, such as available space in the command buffers reserved from compression unit 324. Compression unit 324 may also support the use of custom or different dictionaries, as well.

[00137] Still further, the compression unit 324 can also be configured to accept information that sets a policy as to which compression mode (fixed v. dynamic) to

select. For example, a policy can be specified that favors speed/throughput over data size reduction, or vice versa. This policy information can also be included in the commands passed to the compression unit 324 via the command buffers. As discussed herein, a beltway producer and/or consumer process may set these policies based on current loading conditions, such as the number of available compression command buffer entries. In addition, the compression policy can also be set based on packet size or network application service type.

[00138] In the data plane packet processing path, a process may access compression unit 324 via a corresponding beltway data structure. Figure 18 illustrates a compression operation that utilizes or calls a beltway producer function (above) to conditionally add a packet to the beltway. In one implementation, the control plane invokes the compression process when the kPolicyCompressible bit is set for the data flow to which the packet corresponds. The packets added to the beltway may be uncompressed packets or compressed packets. The logic associated with adding packets assesses first whether the packet is to be compressed or decompressed by the compression unit 324 (1802). If a packet is to be compressed, the compression process accesses the size of the packet to determine whether it is large enough such that compression is worthwhile (1804). In a particular implementation, the compression process compares the packet size to a threshold value, which can be configured or changed depending on the type of compression being used. The compression process also determines whether there is sufficient command buffer space available for the compression command (1806). In one implementation, this check is performed relative to a threshold percentage of the command buffer space reserved for the compression unit 324. In one implementation, if less than 20 percent of the command buffer space reserved for the compression unit 324 is available, the packet is not compressed. This threshold has the effect of reserving at least the remaining 20 percent of the command buffer space for de-compression commands responsive to received packets. This reservation is implemented to pre-allocate or reserve the availability of compression resources for received packets that have been compressed by a remote unit. Accordingly, the threshold value can be configured

based on a variety of performance objectives. The compression process also determines whether there is sufficient memory space for the compressed or decompressed output packet (1808). If sufficient buffer space is available, the compression process reserves packet buffer space for the compressed or decompressed output, and calls the beltway add function (see Figure 15) to add a work queue entry including the compression command to the compression beltway data structure (1810). As discussed above, the beltway add function may schedule the beltway consumer process if it has not already been scheduled. As Figure 18 also shows, if there are insufficient resources for the compression operation, the compression process returns the packet to the calling function or process (1812). For compressed packets to be decompressed, the calling function or process may cause the packets to be buffered to wait for sufficient resources or dropped. For packets to be compressed, the calling function or process can decide to simply bypass compression, or buffer the packets to wait for sufficient resources, depending on the error code provided by the compression process. The logic of Figure 18 can also be extended to include other checks. For example, if compression operations performed on prior packets of the data flow did not actually compress (within some threshold) the packets, the compression process may not pass the packet to the compression unit 324.

[00139] As discussed above, the beltway consumer process, when executed, writes commands to the registers and command buffers allocated to the compression unit 324. As discussed above, the beltway for the compression unit 324 stores work queue entries and commands for both compression and decompression operations. In a particular implementation, the beltway consumer process logic can also be extended to modify compression commands for compression operations depending on observed loading conditions. De-compression operations, on the other hand, will be largely responsive to the modes and parameters used by the remote device in compressing the packet.

[00140] Figure 19 illustrates example decision logic that can be invoked when a compression command operation is written to the compression unit 324 by the beltway consumer process in step 1612 of Figure 16. When the available command

buffer space for compression unit 324 is below a first threshold (such as 40 percent) (1902), the beltway consumer process may select the fastest compression mode supported by the compression unit 324 (in one implementation, by setting the appropriate command words or bits in the command buffer to force fixed compression) (1904). Otherwise, if the available command buffer space is less than a second threshold (such as 60 percent) (1906), the beltway consumer process may select a fast compression mode (similar to setting a Z_BEST_SPEED bit in zlib compression engines) (1908). If more buffer space is available, the beltway consumer process may select a compression mode that allows compression unit 324 to select the fixed or dynamic compression mode without a selection bias toward speed.

[00141] Other embodiments are also possible. For example, the beltway consumer process or beltway add function could set the compression mode and policies based on the size of the packet and/or the network application type associated with the packet. In other implementations, selection of compression mode can be based on observed processing latency, such as the amount of time previous packet had been buffered in the compression unit 324 prior to output. In another alternative, selection of compression mode (or even the decision to compress) can be based on link utilization observed by the control plane. For example, the control plane can set a compression probability value (within a normalized range, for example, between 0 and 99) in the FlowInfo messages based on link utilization along the path for the specific flow. In one implementation, the probability value can be based on the percentage of the observed link utilization. In another implementation, the probability value can be set to the greater of 1) zero or 2) $(\text{the observed percentage link utilization} - 75) * 4$. In such an implementation, if the available command buffers are below a threshold value (e.g., 30 percent), the compression process can compute a random number (e.g., between 0 and 99). If the random number is less than the compression probability value associated with the flow to which the packet is a part, the compression process passes the packet to the compression unit 324. Otherwise, if the available command buffers exceed the threshold, the packet is passed to the compression unit 324 via the command

buffers without regard to probability values.

[00142] Particular embodiments of the above-described process might be comprised of instructions that are stored on storage media. The instructions might be retrieved and executed by a processing system. The instructions are operational when executed by the processing system to direct the processing system to operate in accord with the present invention. Some examples of instructions are software, program code, firmware, and microcode. Some examples of storage media are memory devices, tape, disks, integrated circuits, and servers. The term "processing system" refers to a single processing device or a group of inter-operational processing devices. Some examples of processing devices are integrated circuits and logic circuitry. Those skilled in the art are familiar with instructions, storage media, and processing systems.

[00143] Those skilled in the art will appreciate variations of the above-described embodiments that fall within the scope of the invention. In this regard, it will be appreciated that there are many possible orderings of the steps in the process described above and many possible modularizations of those orderings. Further, in embodiments where processing speed is not determinative, the process might run in the control plane rather than the data plane. As a result, the invention is not limited to the specific examples and illustrations discussed above, but only by the following claims and their equivalents.

CLAIMS

What is claimed is:

1. An apparatus comprising
a control plane; and
a network processing unit operative to:
 apply one or more policies to received packets of respective data flows according to a data plane configuration;
 selectively sample the received packets of the respective data flows to the control plane;
wherein the control plane is operative to
 process packets sampled by the network processing unit; and
 transmit control messages to the network processing unit to change the data plane configuration.
2. The apparatus of claim 1 wherein the data plane configuration comprises a default policy, and wherein one or more control messages identify one or more policies for respective data flows, and wherein the network processing unit is operative to apply a current data plane configuration to all received packets.
3. The apparatus of claim 1 wherein the network processing unit is operative to sample connection-initiating packets of respective data flows to the control plane.
4. An apparatus comprising
 N processors, where N is greater than 1;
 a work order module operative to provide work queue entries to the N processors, wherein the work queue entries correspond to respective tasks, and to provide, to a given processor of the N processors, exclusive access to one or more shared data structures for work queue entries including atomic tags;

a compression unit operatively coupled to a command buffer space, wherein the compression unit is operative to compress or decompress packets responsive to commands stored in the command buffer space;

a memory operative to store a beltway data structure comprising N circular buffers for corresponding ones of the N processors, wherein the N circular buffers each comprise a plurality of entries; the memory further operative to store processor-readable instructions for a beltway producer function, a compression module, and a beltway consumer process;

wherein the beltway producer function includes instructions operative to cause the N processors, when executing the beltway producer function, to

add entries to corresponding ones of the circular buffers of the beltway data structure;

submit work queue entries to the work order module including an atomic tag associated with the beltway data structure; and

wherein the compression module includes instructions operative to cause the N processors, when executing the compression module, to

responsive to an indication of a packet to be compressed, call the beltway producer function to add a command identifying the packet to an entry of a select circular buffer of the beltway data structure;

wherein the beltway consumer process includes instructions operative to cause a processor of the N processors, responsive to a work queue entry identifying the beltway consumer process and an atomic tag corresponding to the beltway data structure, to:

read entries from the circular buffers of the beltway data structure;

and

write commands stored in the entries to the command buffer space coupled to the compression unit.

5. The apparatus of claim 4 wherein the compression module further includes instructions operative to cause the N processors, when executing the compression module, to

conditionally call the beltway producer function based on the size of the packet.

6. The apparatus of claim 4 wherein the compression module further includes instructions operative to cause the N processors, when executing the compression module, to

conditionally call the beltway producer function based on an assessment of load associated with the compression unit or resources consumed by the compression unit.

7. The apparatus of claim 4 wherein the compression module further includes instructions operative to cause the N processors, when executing the compression module, to

conditionally call the beltway producer function based on utilization of command buffer space coupled to the compression unit.

8. An apparatus comprising
a control plane; and

a network processing unit operative to:

selectively sample received packets of the respective data flows to the control plane;

maintain a flow database of flow information entries, each comprising one or more policies and indexed by one or more data flow attributes;

maintain a triggering cache of flow information entries, each comprising one or more cached policies and indexed by one or more server-side attributes;

match packets of respective data flows to flow information entries in the flow database; else match packets of new respective data flows to flow information entries in the triggering cache based on the one or more server-side attributes;

apply the one or more policies to received packets of respective data flows that match corresponding flow information entries in the triggering cache or the flow database, otherwise apply one or more default policies to received packets of respective data flows according to a current data plane configuration;

wherein the control plane is operative to

process packets sampled by the network processing unit; and

transmit control messages to the network processing unit identifying one or more policies for the respective data flows;

wherein one or more of the control messages include indications operable to cause the data plane to add flow information entries in the flow database, and store the one or more policies in the triggering cache in association with the one or more server-side attributes.

9. The apparatus of claim 8 wherein the one or more server-side attributes comprise an IP address and a port number.

10. An apparatus comprising

N processors, where N is greater than 1;

a work order module operative to provide work queue entries to the N processors, wherein the work queue entries correspond to respective tasks, and to provide, to a given processor of the N processors, exclusive access to one or more shared data structures for work queue entries including atomic tags;

a memory operative to store one or more instances of a beltway data structure, each instance of a beltway data structure comprising N circular buffers for corresponding ones of the N processors, wherein the N circular buffers each comprise a plurality of entries; the memory further operative to store processor-readable instructions for a beltway producer function and a beltway consumer process;

wherein the beltway producer function includes instructions operative to cause the N processors, when executing the beltway producer function, to

add entries to corresponding ones of the circular buffers of an identified instance of a beltway data structure;

conditionally submit work queue entries to the work order module including an atomic tag associated with the identified instance of the beltway data structure; and

wherein the beltway consumer process includes instructions operative to cause a processor of the N processors, responsive to a work queue entry identifying the beltway consumer process and an atomic tag corresponding to the instance of a beltway data structure, to:

identify the instance of the beltway data structure based on one or more attributes of the work queue entry;

read entries from the circular buffers of the identified instance of the beltway data structure; and

call a consumer function associated with the identified instance of the beltway data structure to process the read entries.

11. The apparatus of claim 10 wherein the work order module is operative to maintain a queue of work entries, one or more of the work entries including an atomic tag; and

schedule work entries in the queue for the N processors such that only a single processor of the N processors is provided a work entry labeled with a given atomic tag.

12. An apparatus comprising

one or more network interfaces;

a plurality of processors;

a work order module operative to

maintain a queue of work entries, one or more of the work entries including a task description, a packet pointer, and a tag; and

schedule work entries in the queue for the plurality of processors such that only a single processor of the plurality of processors is provided a work entry labeled with a given tag;

a memory operative to define a hierarchical partition configuration, the hierarchical partition configuration comprising a plurality of partitions, the memory further operative to buffer packets received at the one or more network interfaces;

wherein the plurality of processors, to schedule received packets for transmission according to the hierarchical partition configuration, are each operative to:

receive, from the work order module, an indication of a first work entry, wherein the first work entry is associated with a packet to be processed and includes a tag identifying a partition of the plurality of partitions,

forward the packet corresponding to the first work entry to a parent partition of the identified partition by modifying the first work entry to include a tag of the parent partition, and

resubmit the first modified work entry to the work order module.

FIG. 1





FIG. 2A


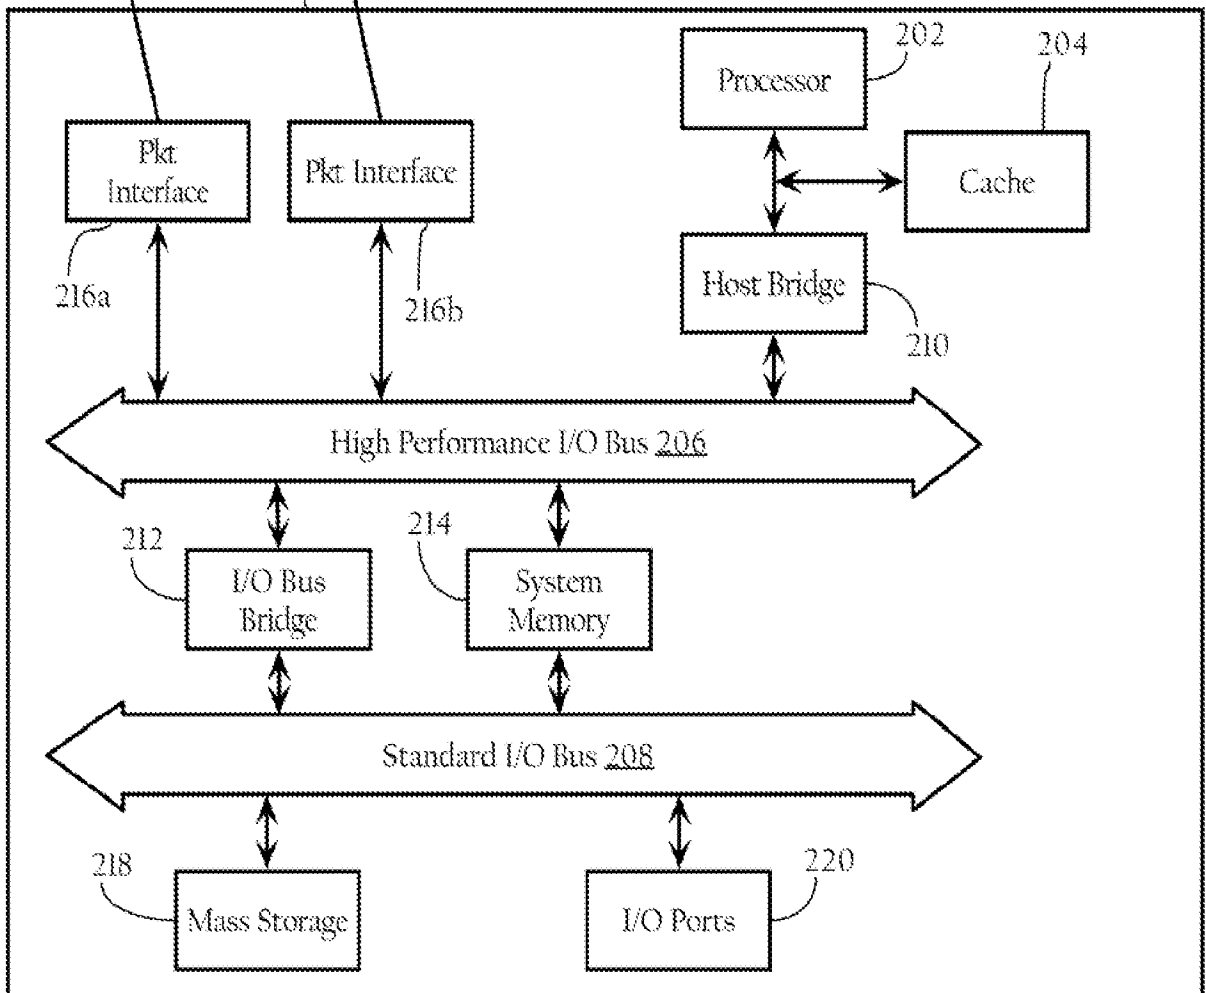

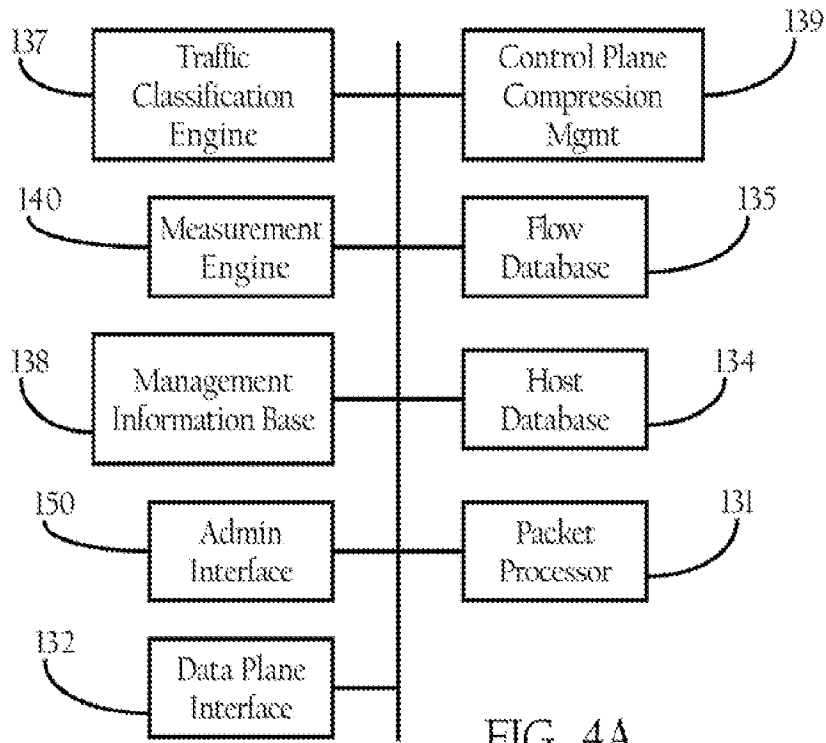


FIG. 3





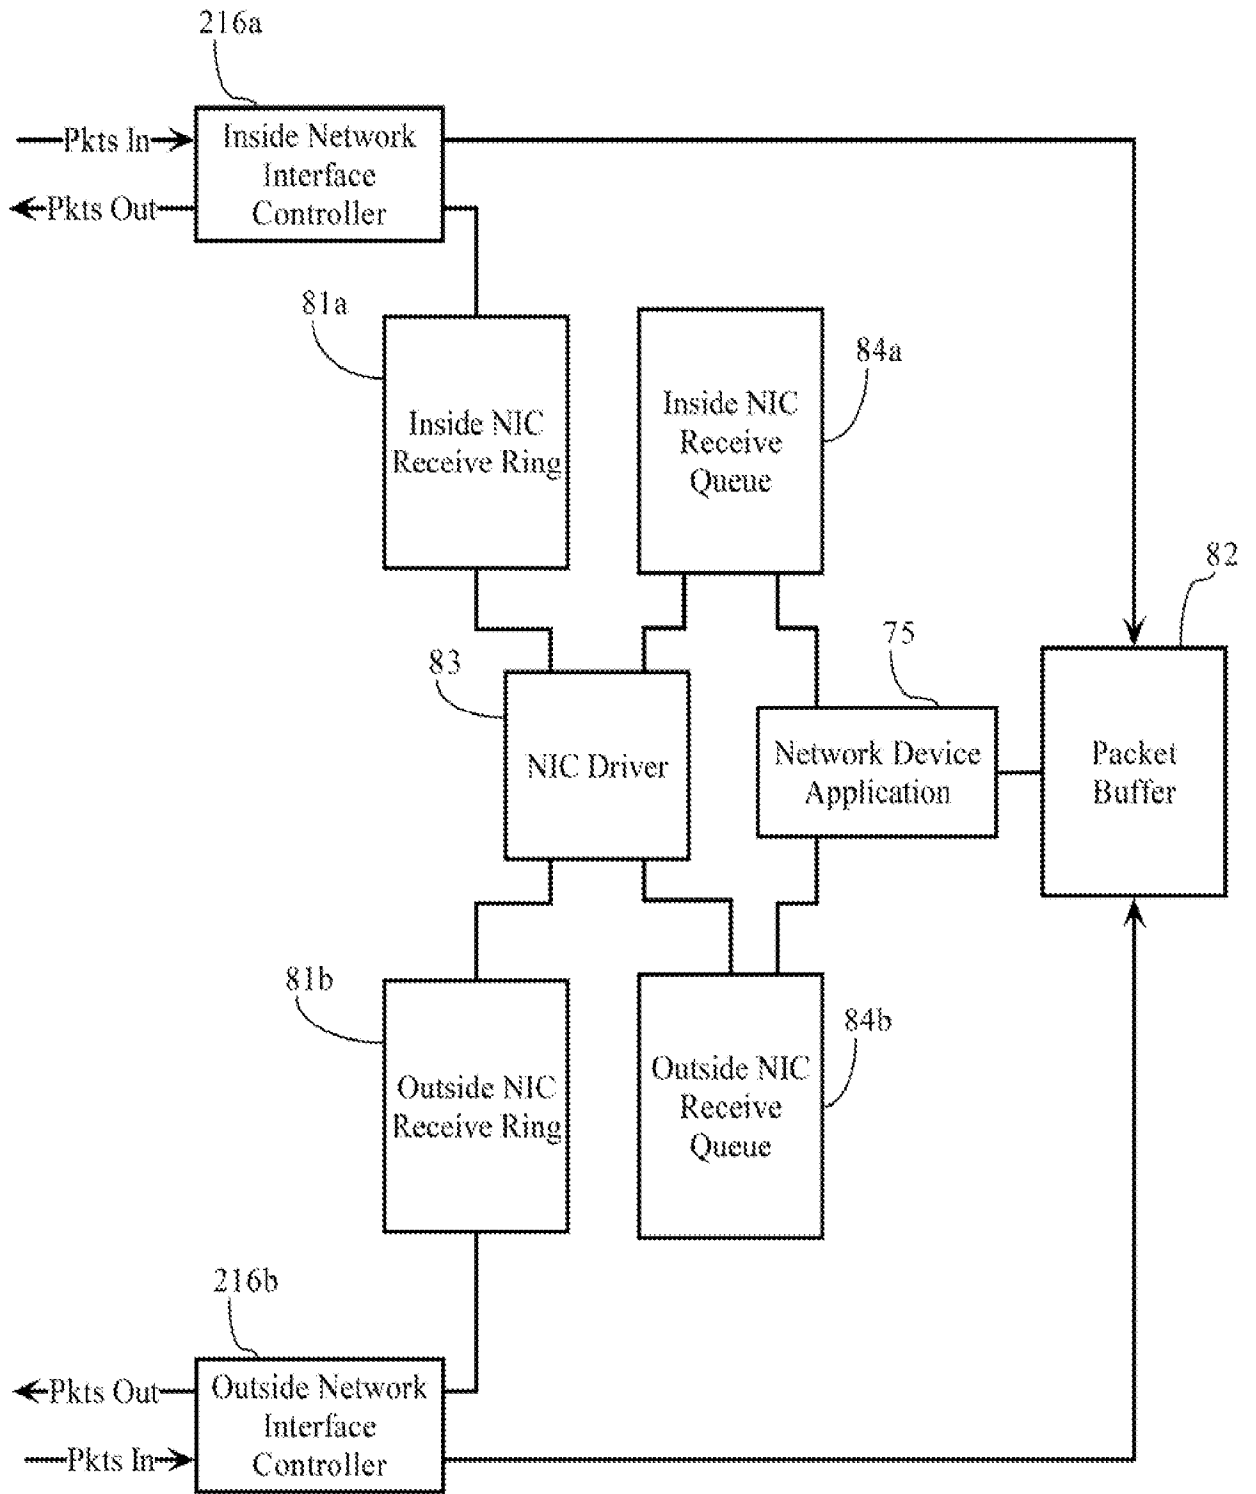


FIG. 4B

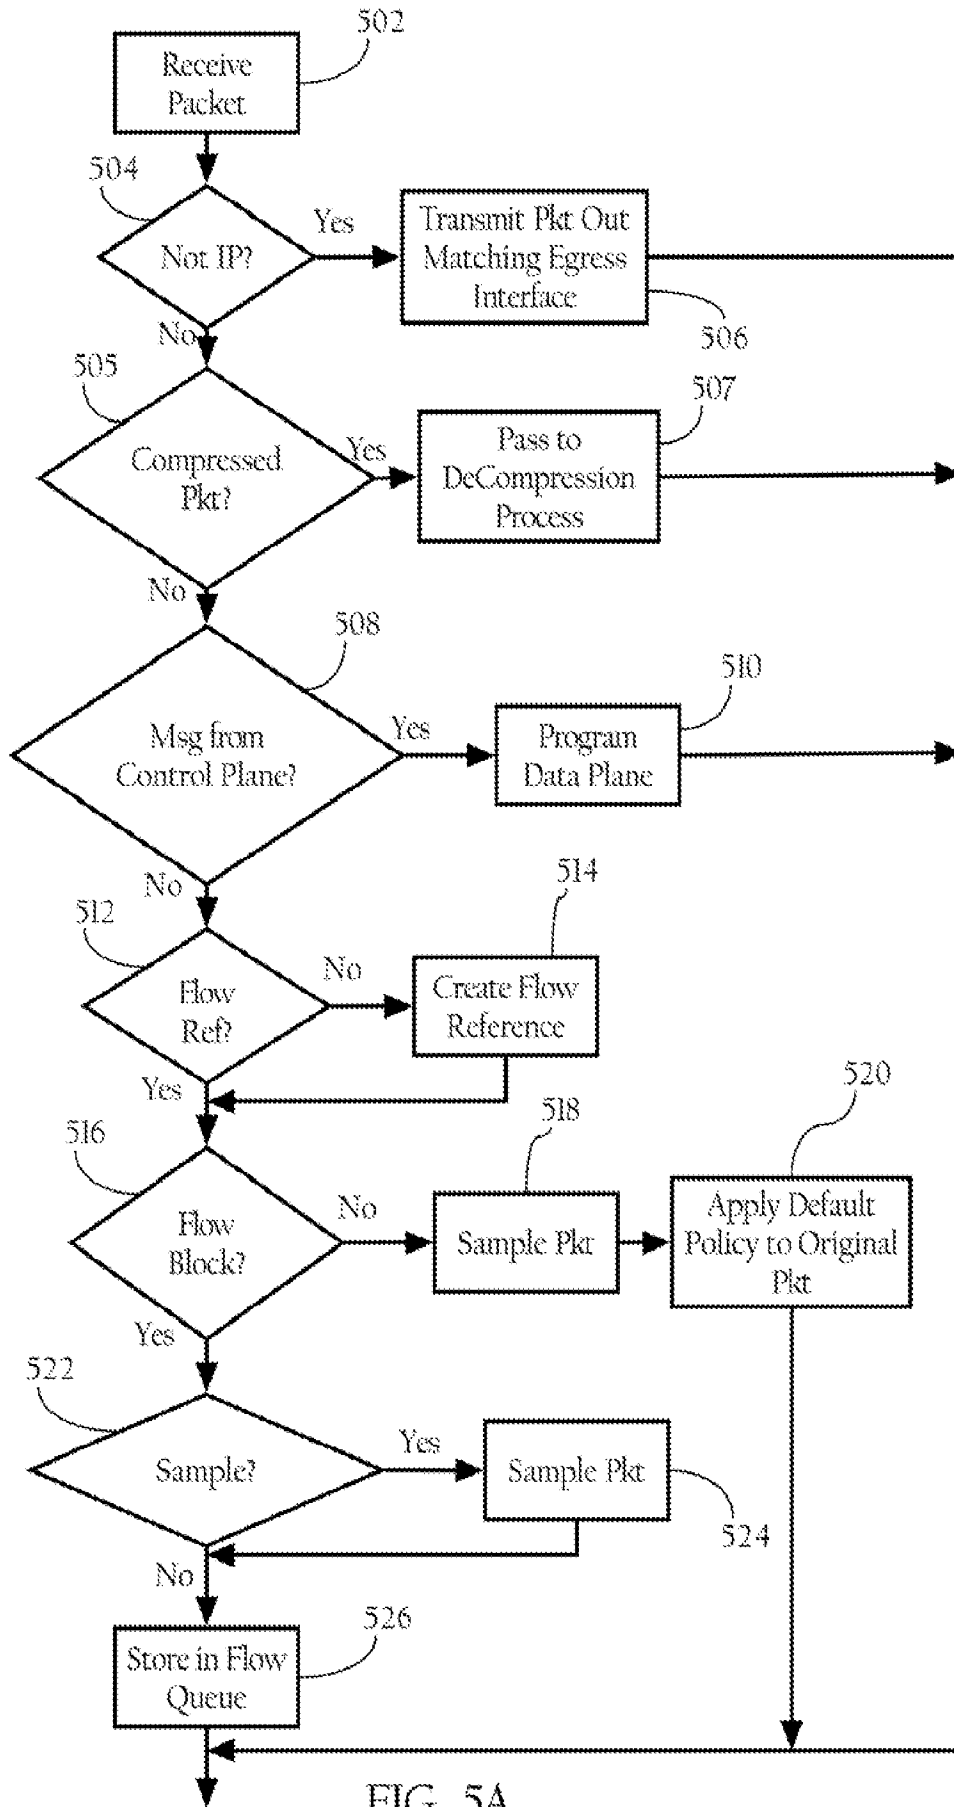


FIG. 5A

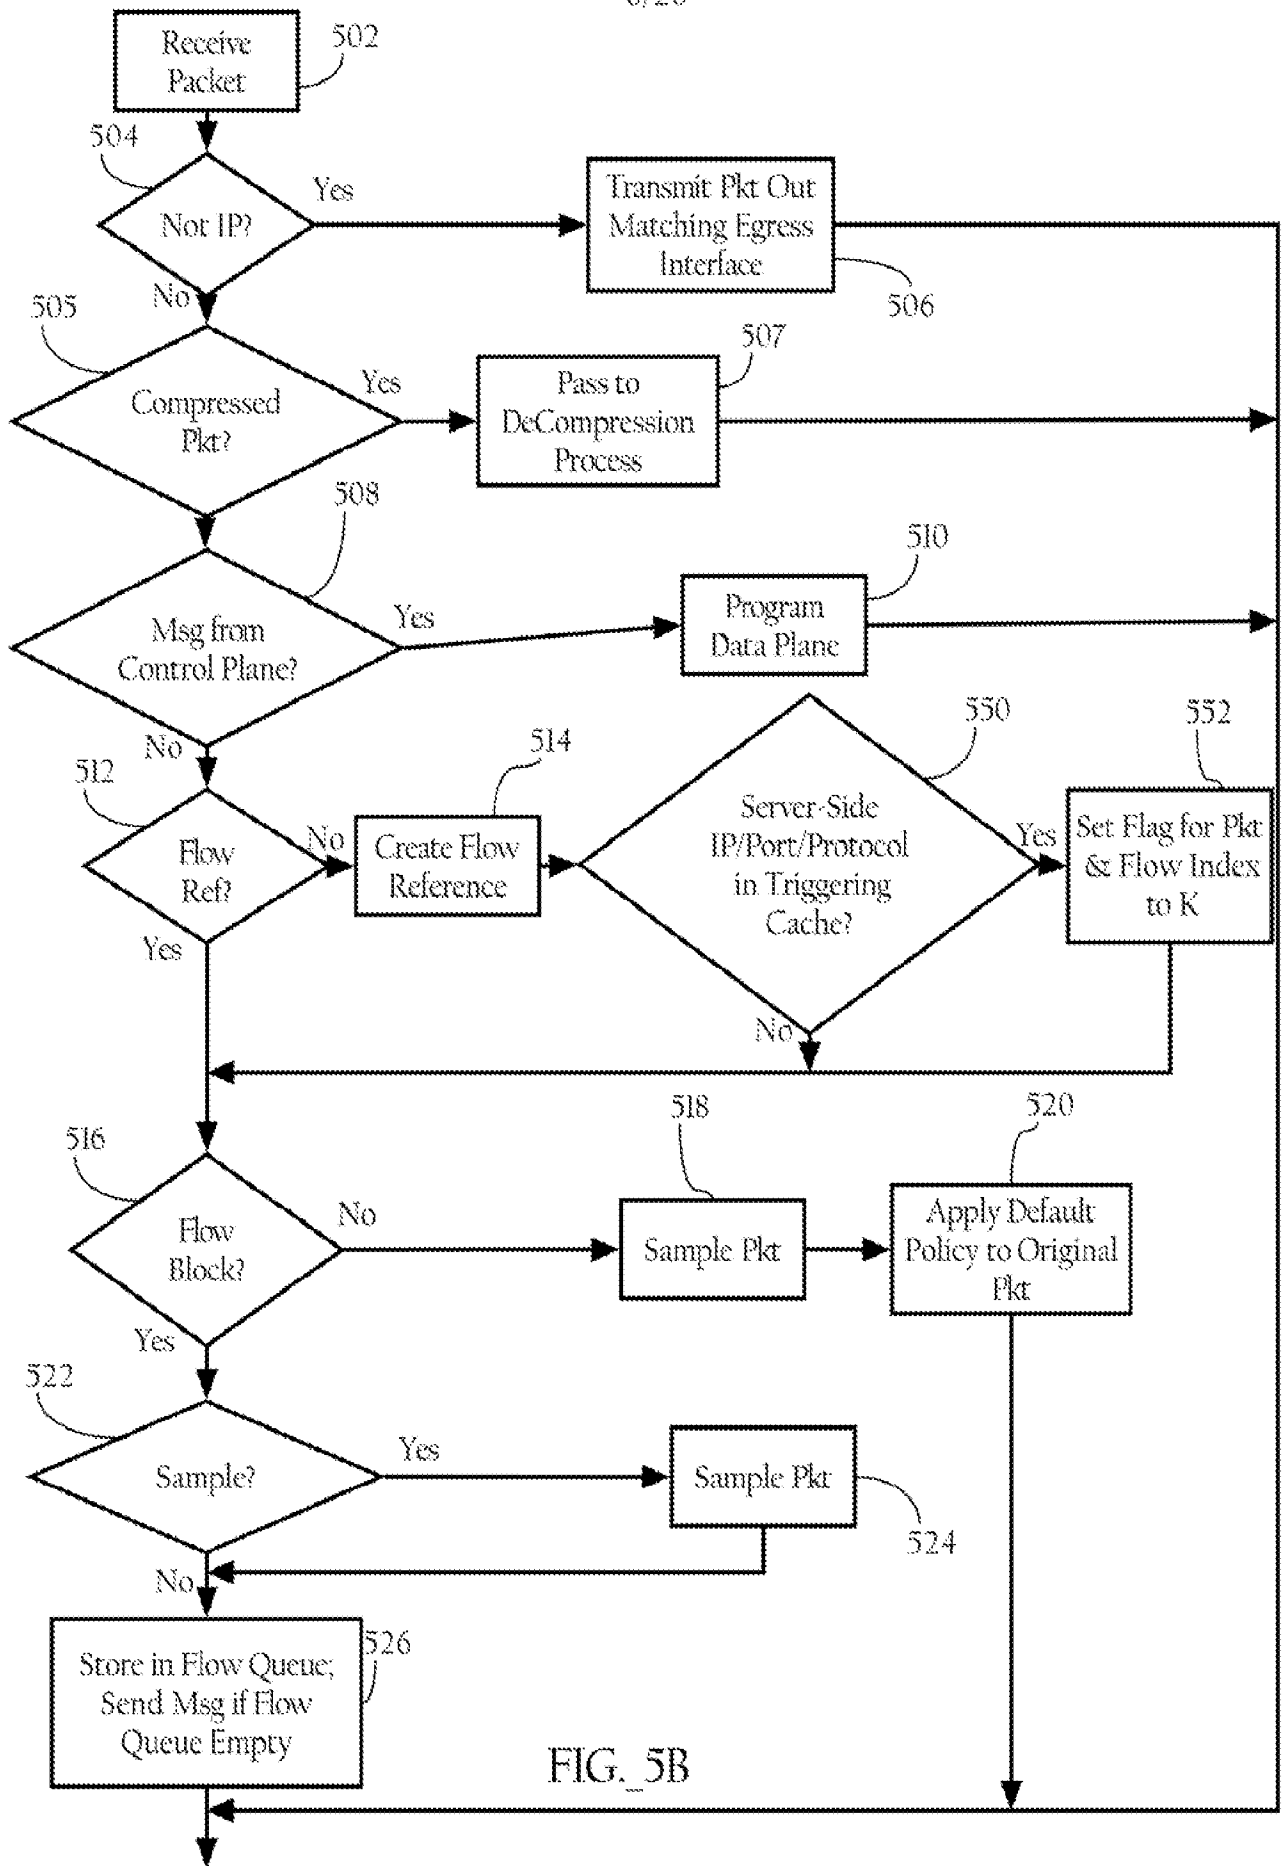


FIG. 5B

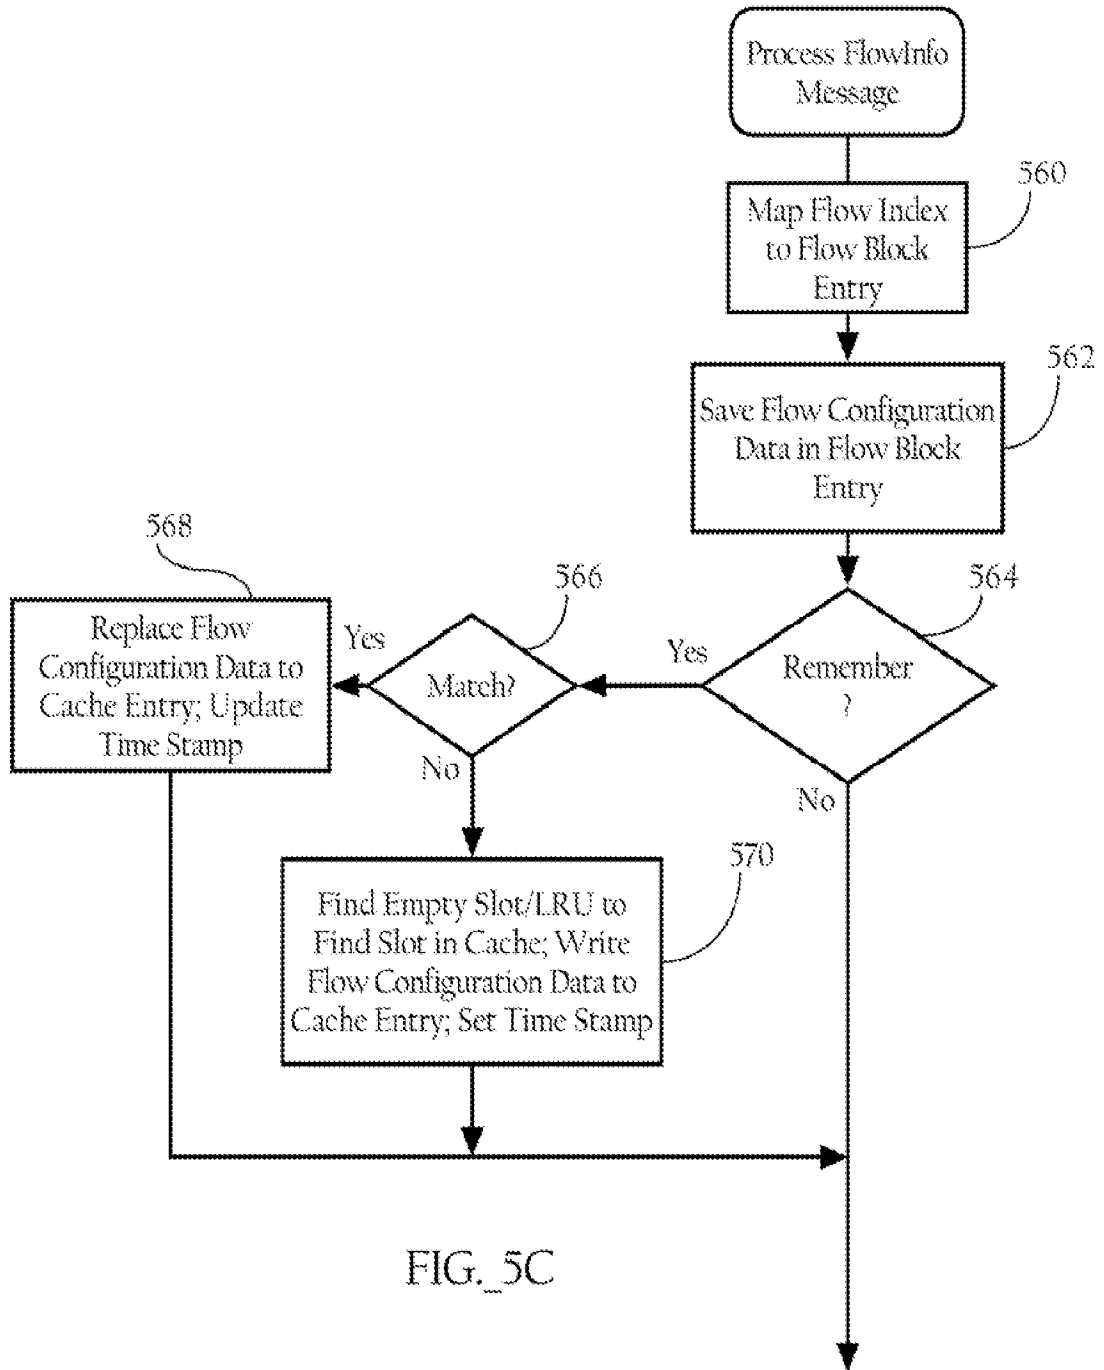


FIG. 5C




FIG._6

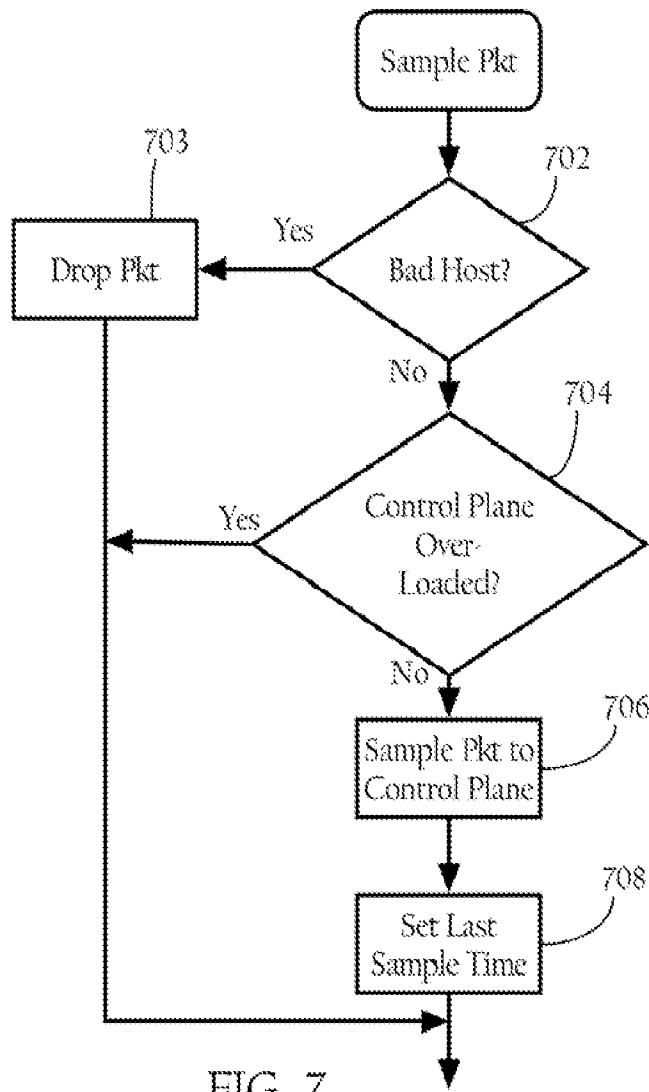
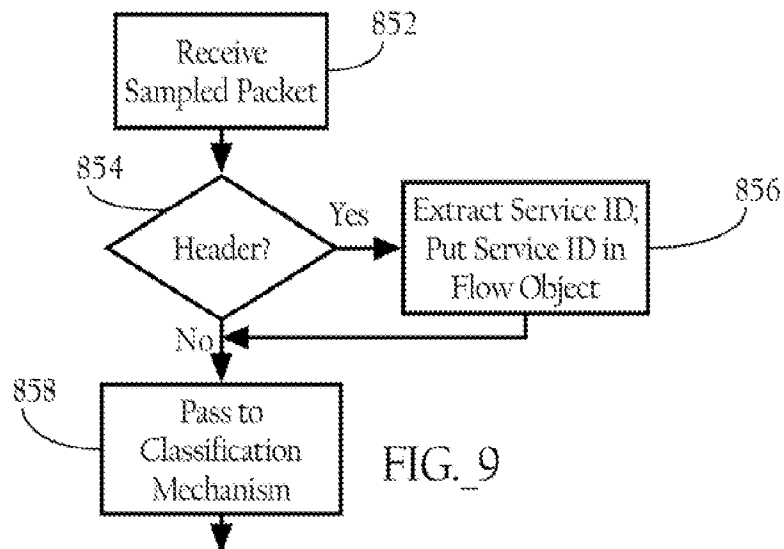
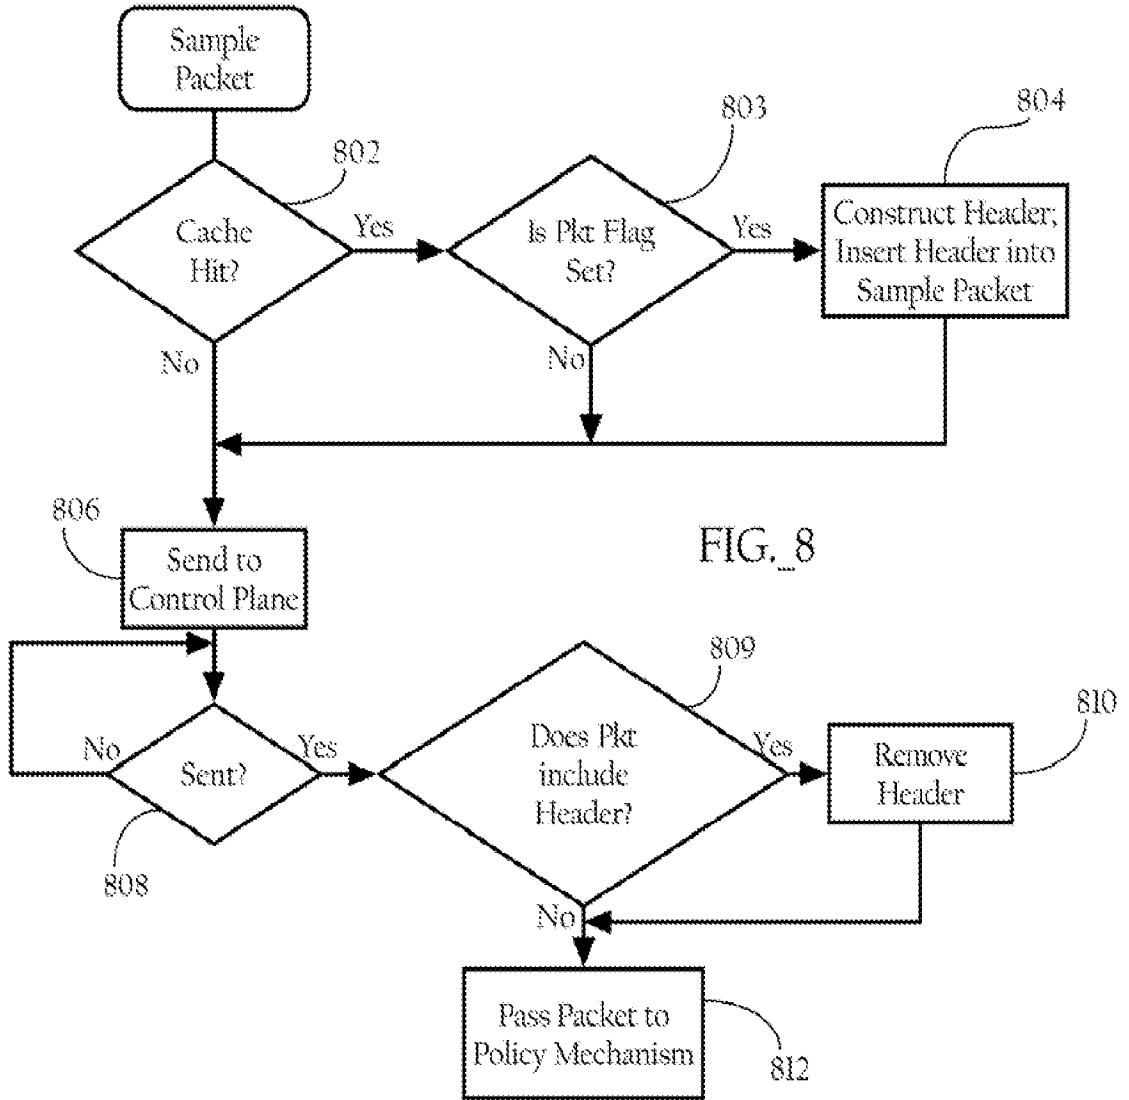
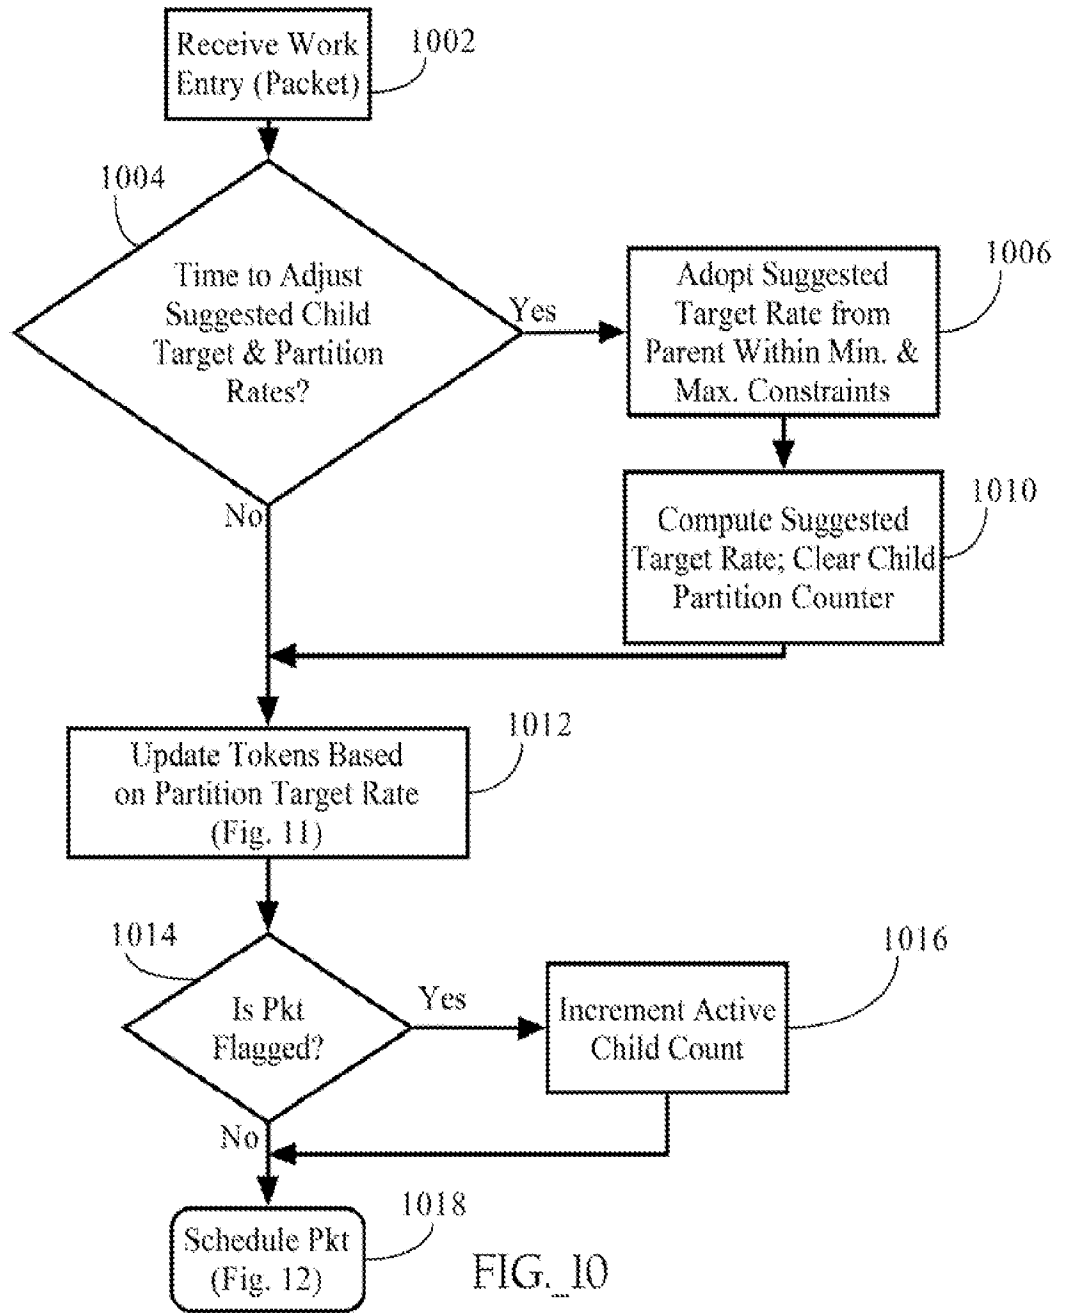


FIG. 7





12/20

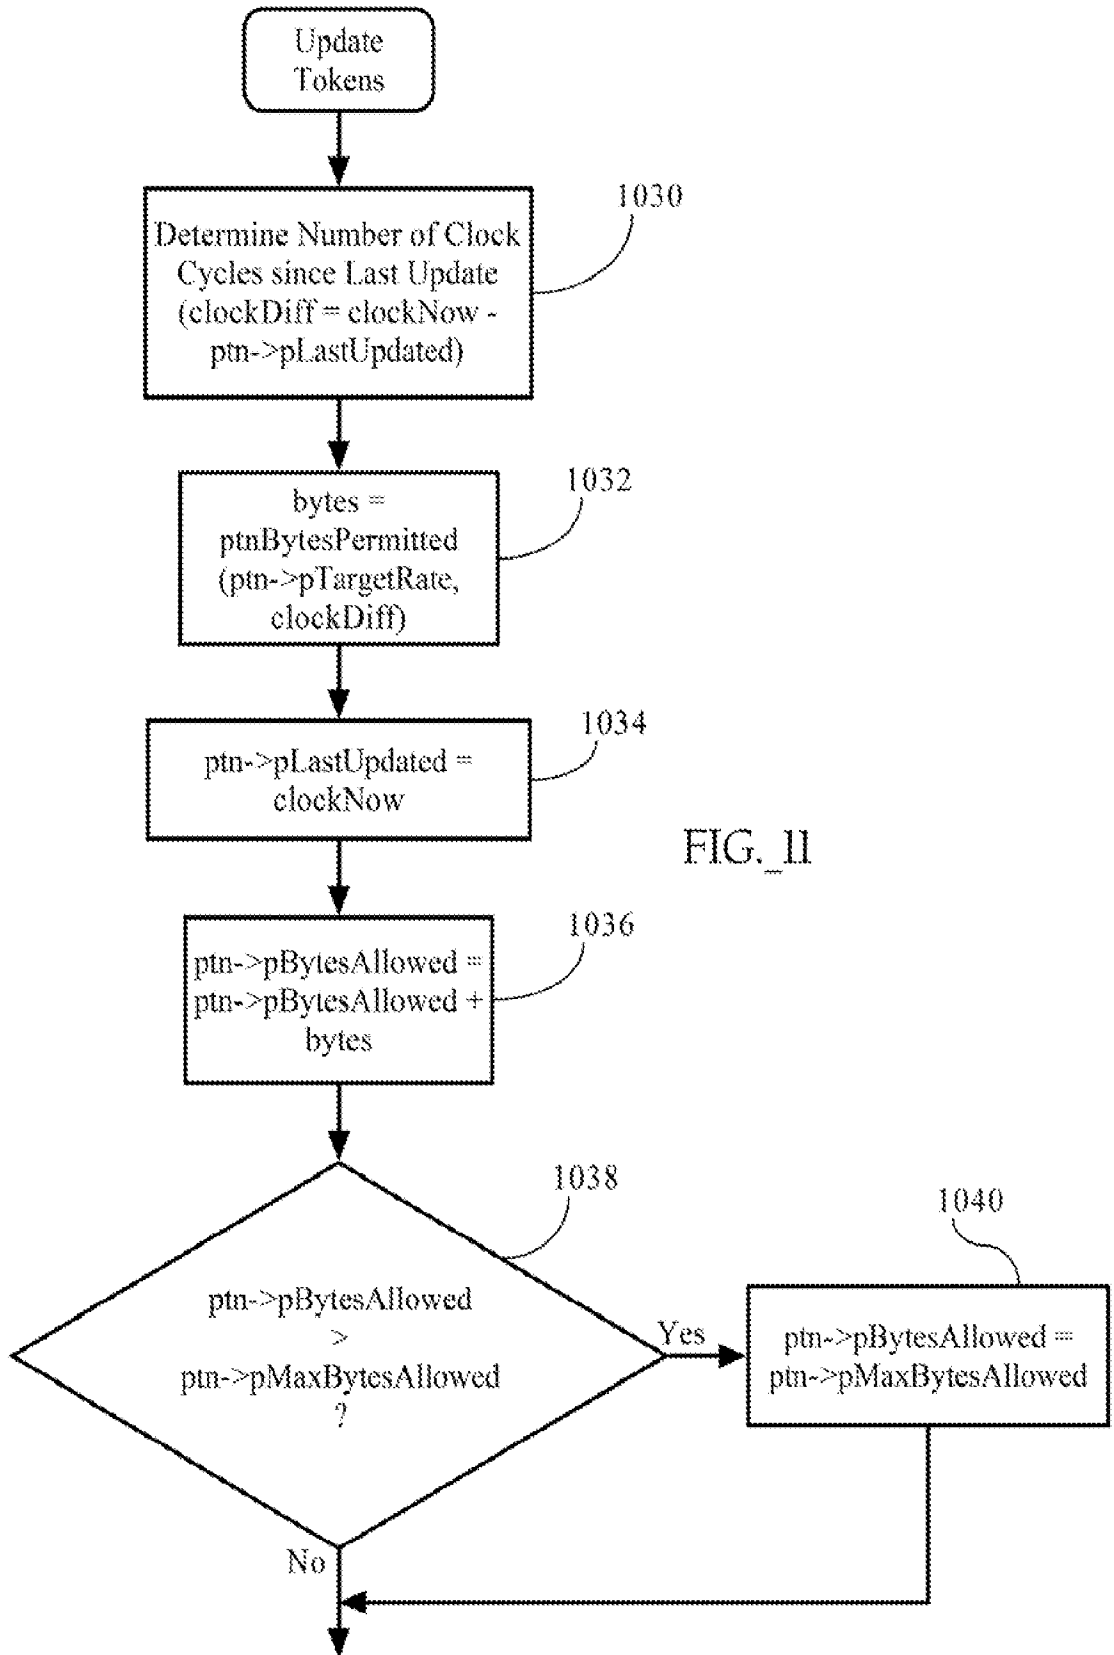


FIG. II

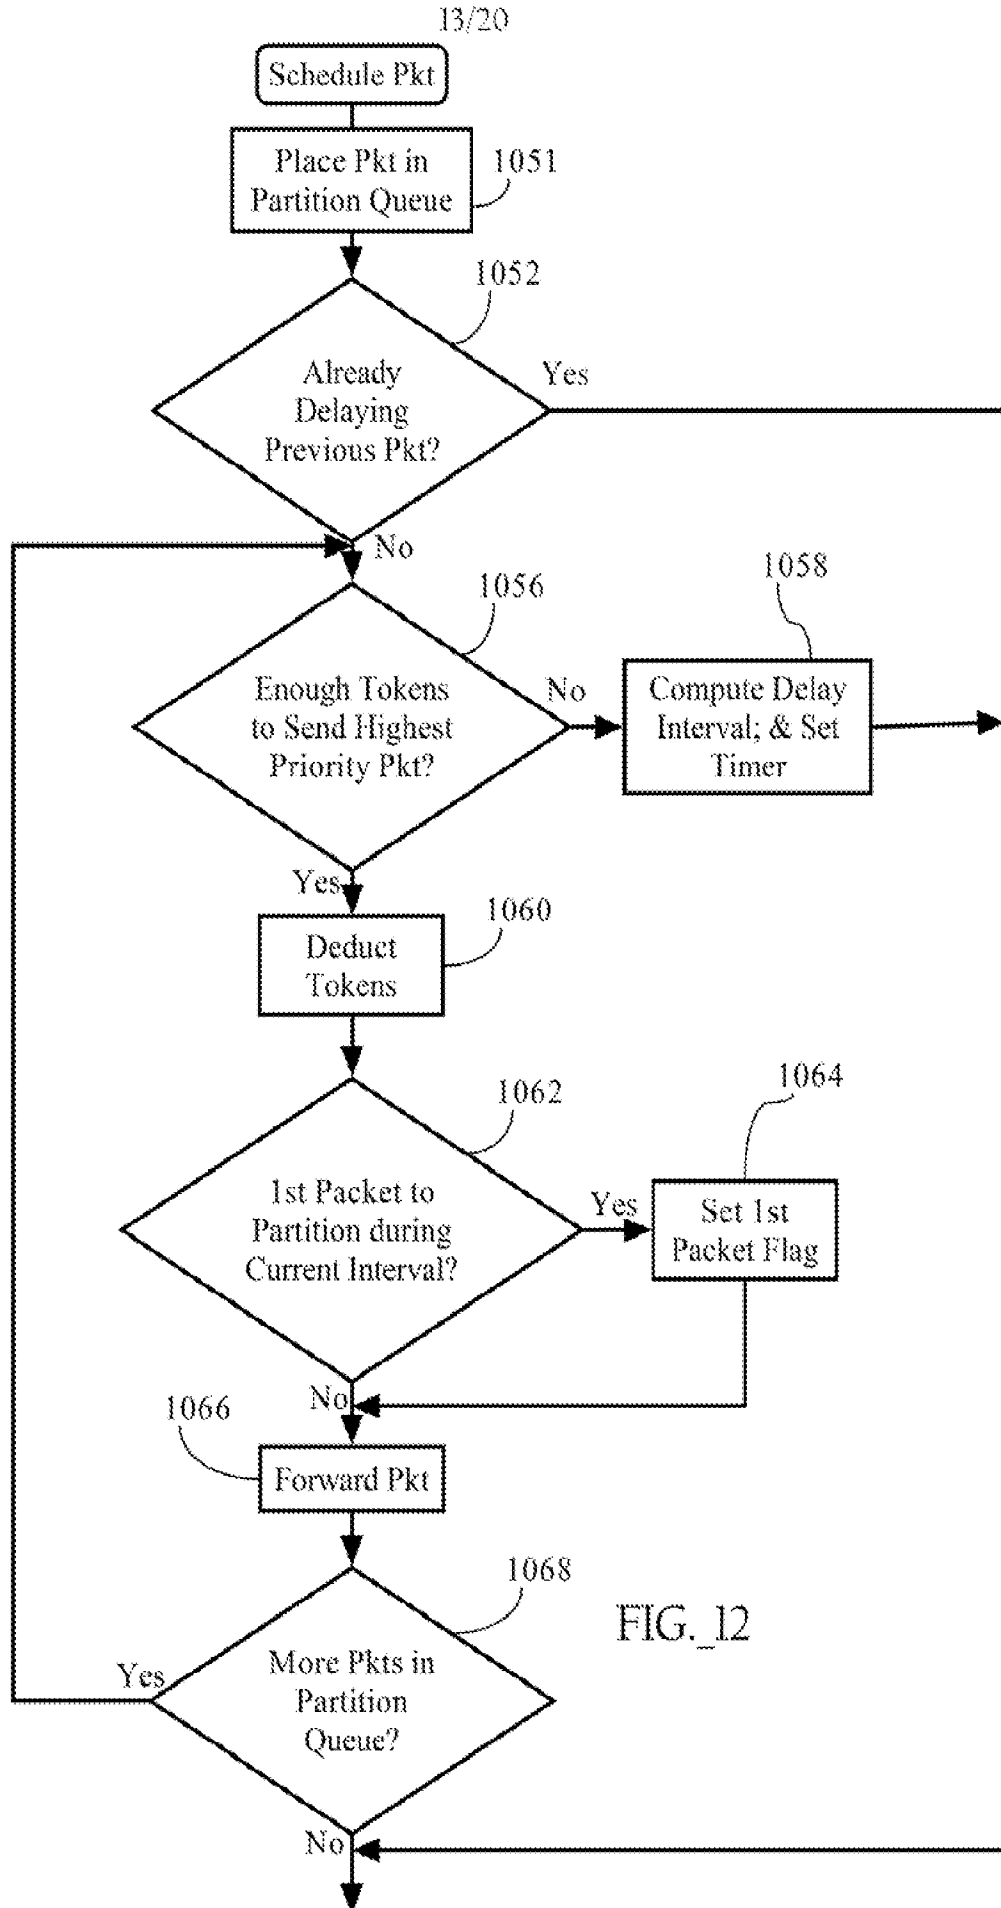



FIG. 12

FIG. 13






FIG. 14

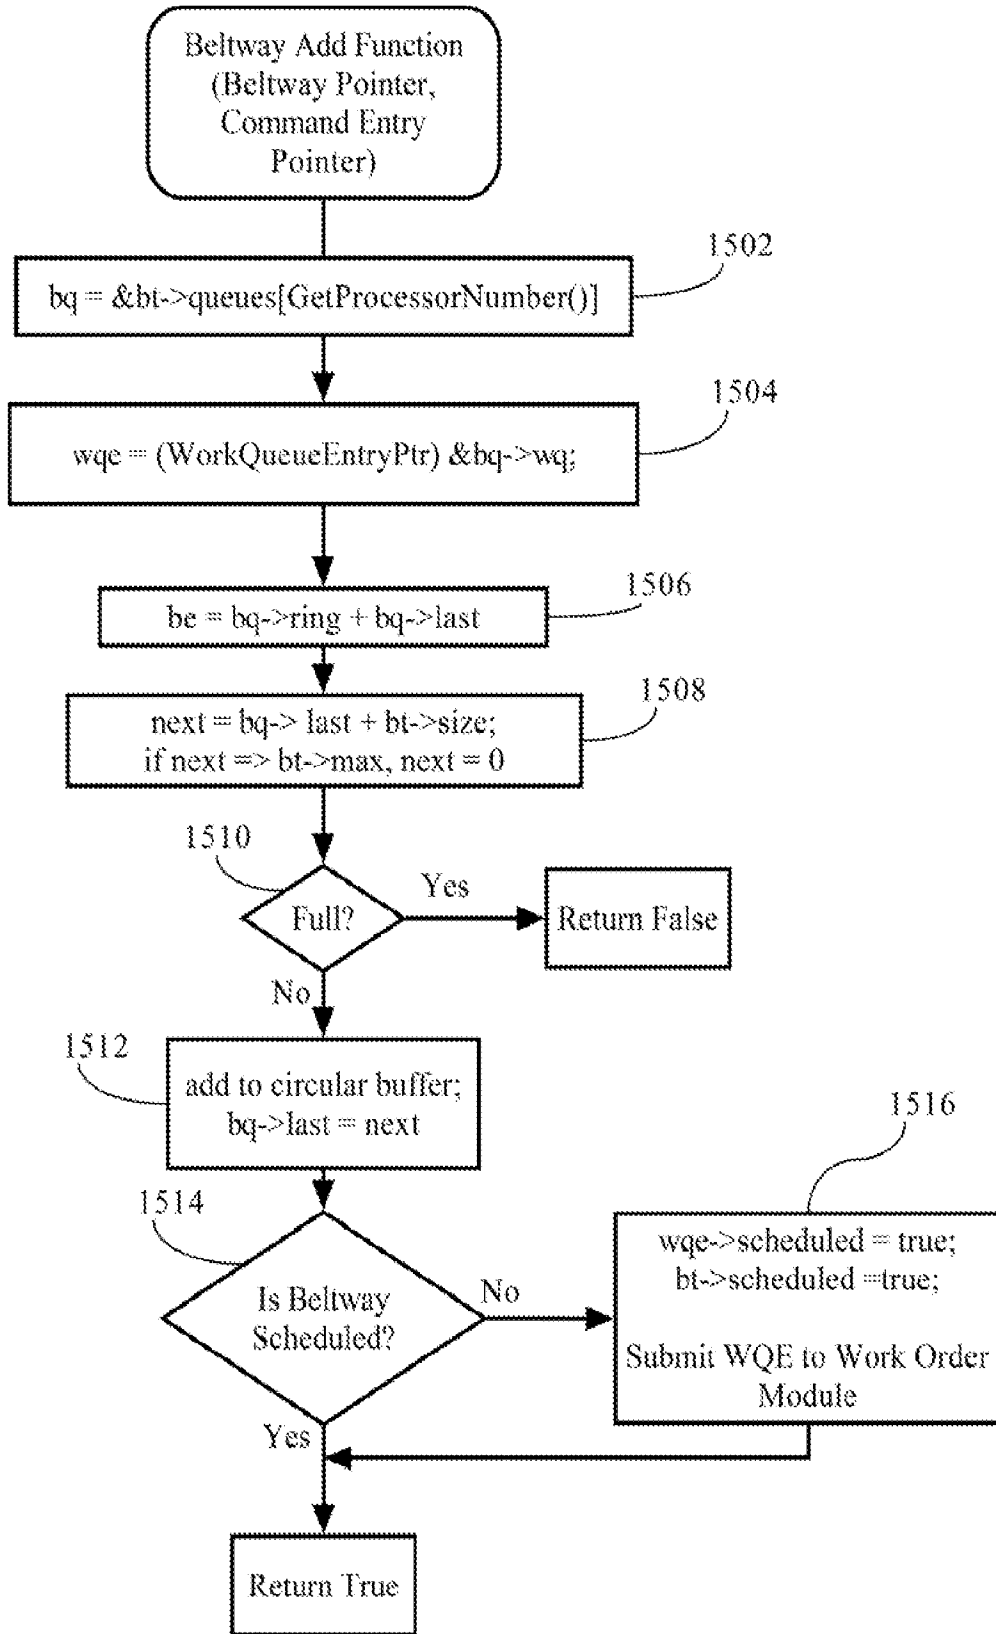


FIG. 15

17/20

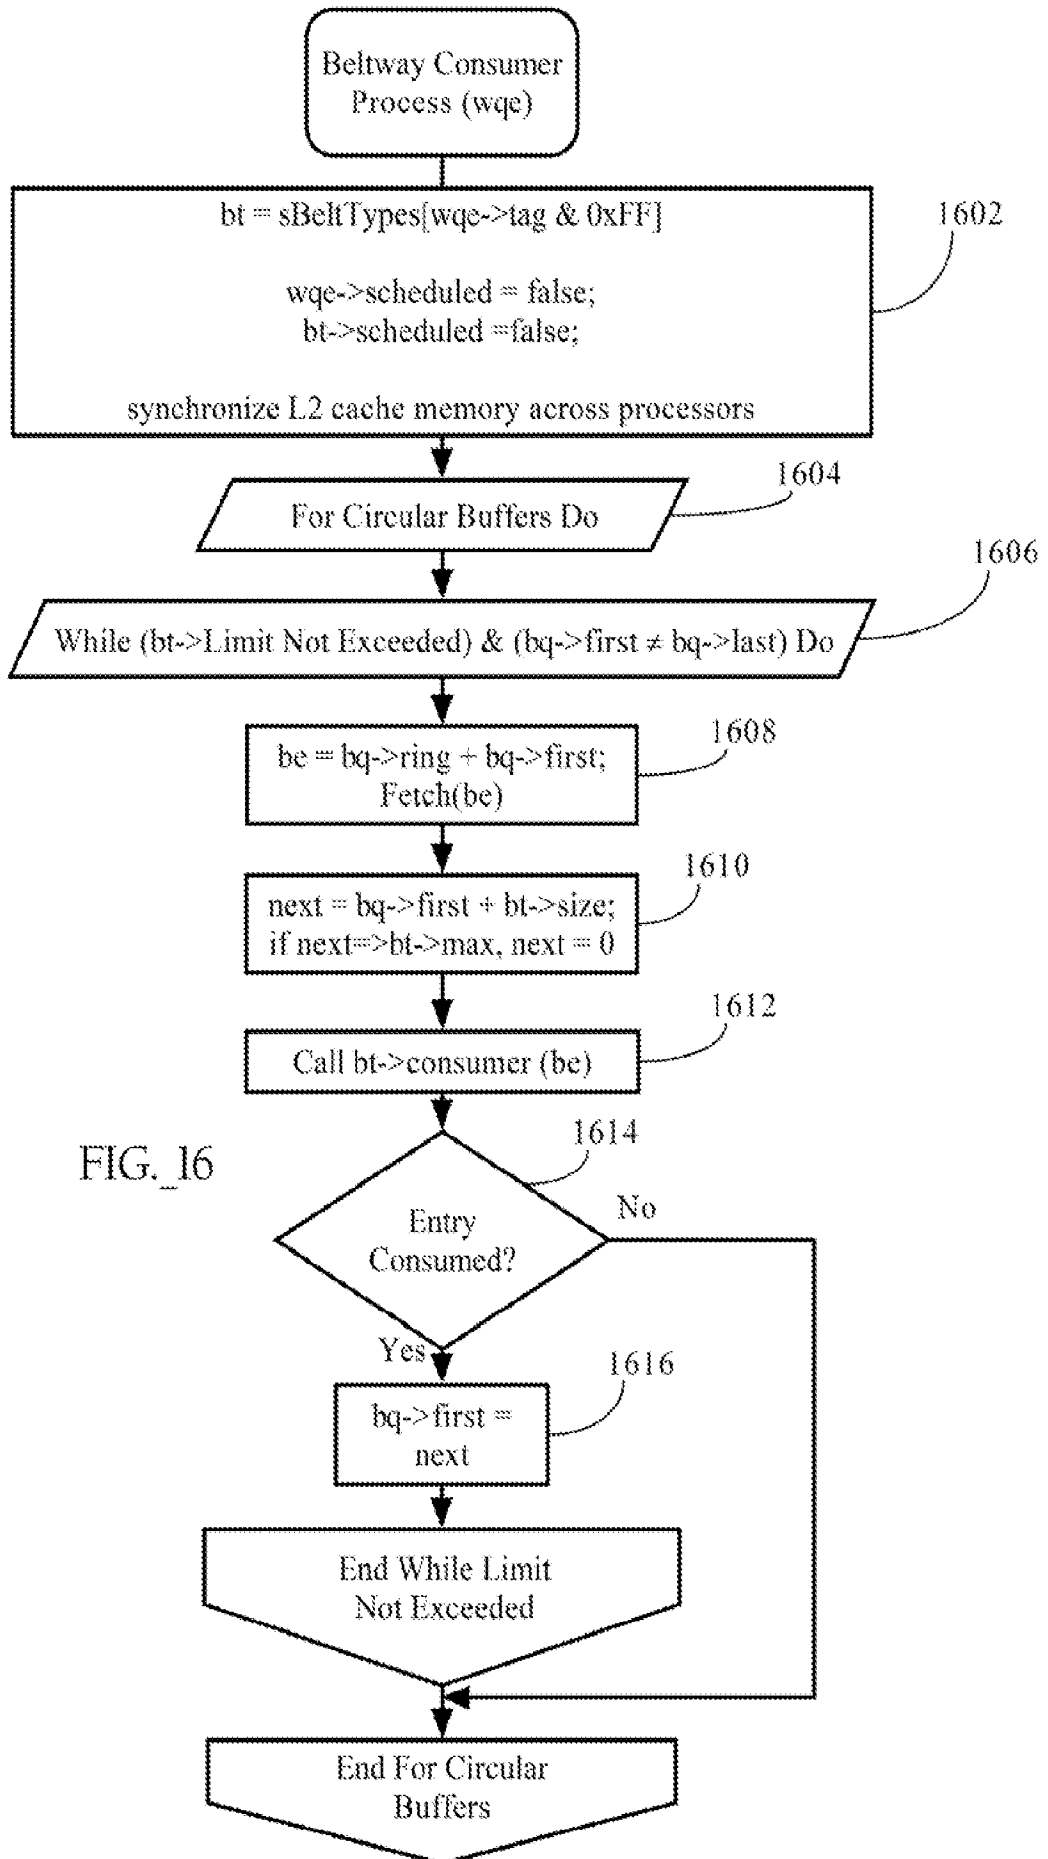


FIG. 16

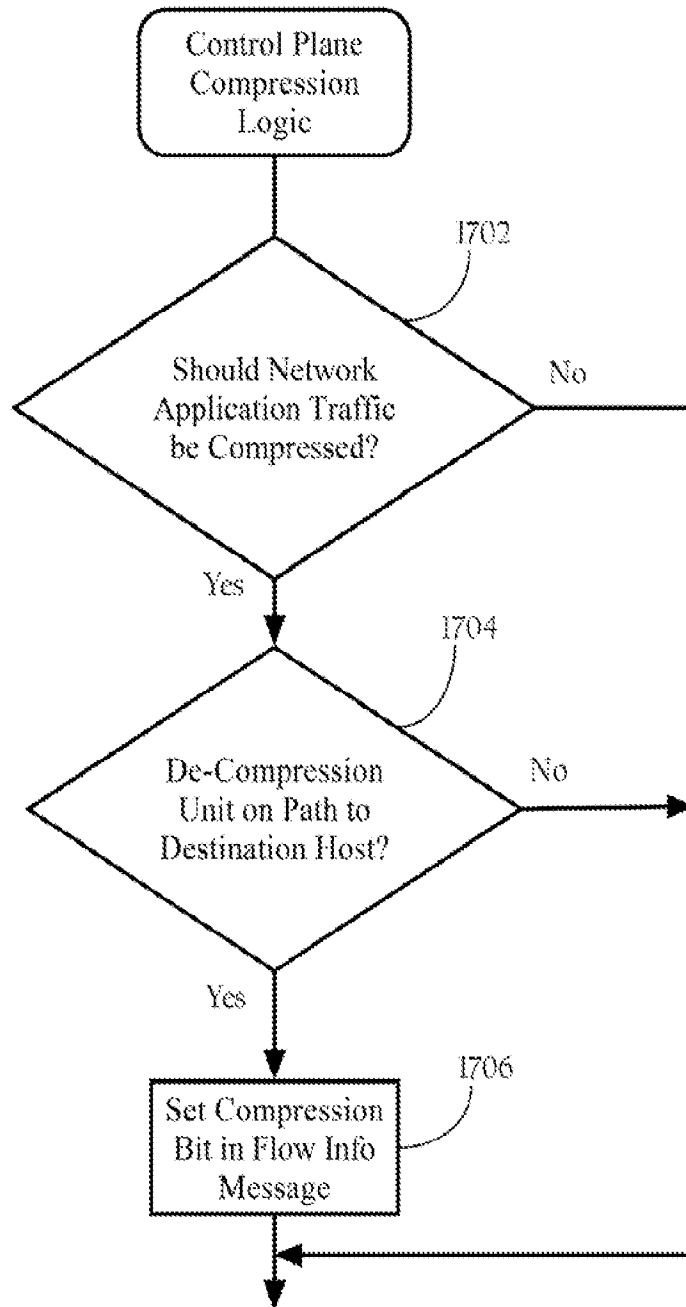


FIG. 17

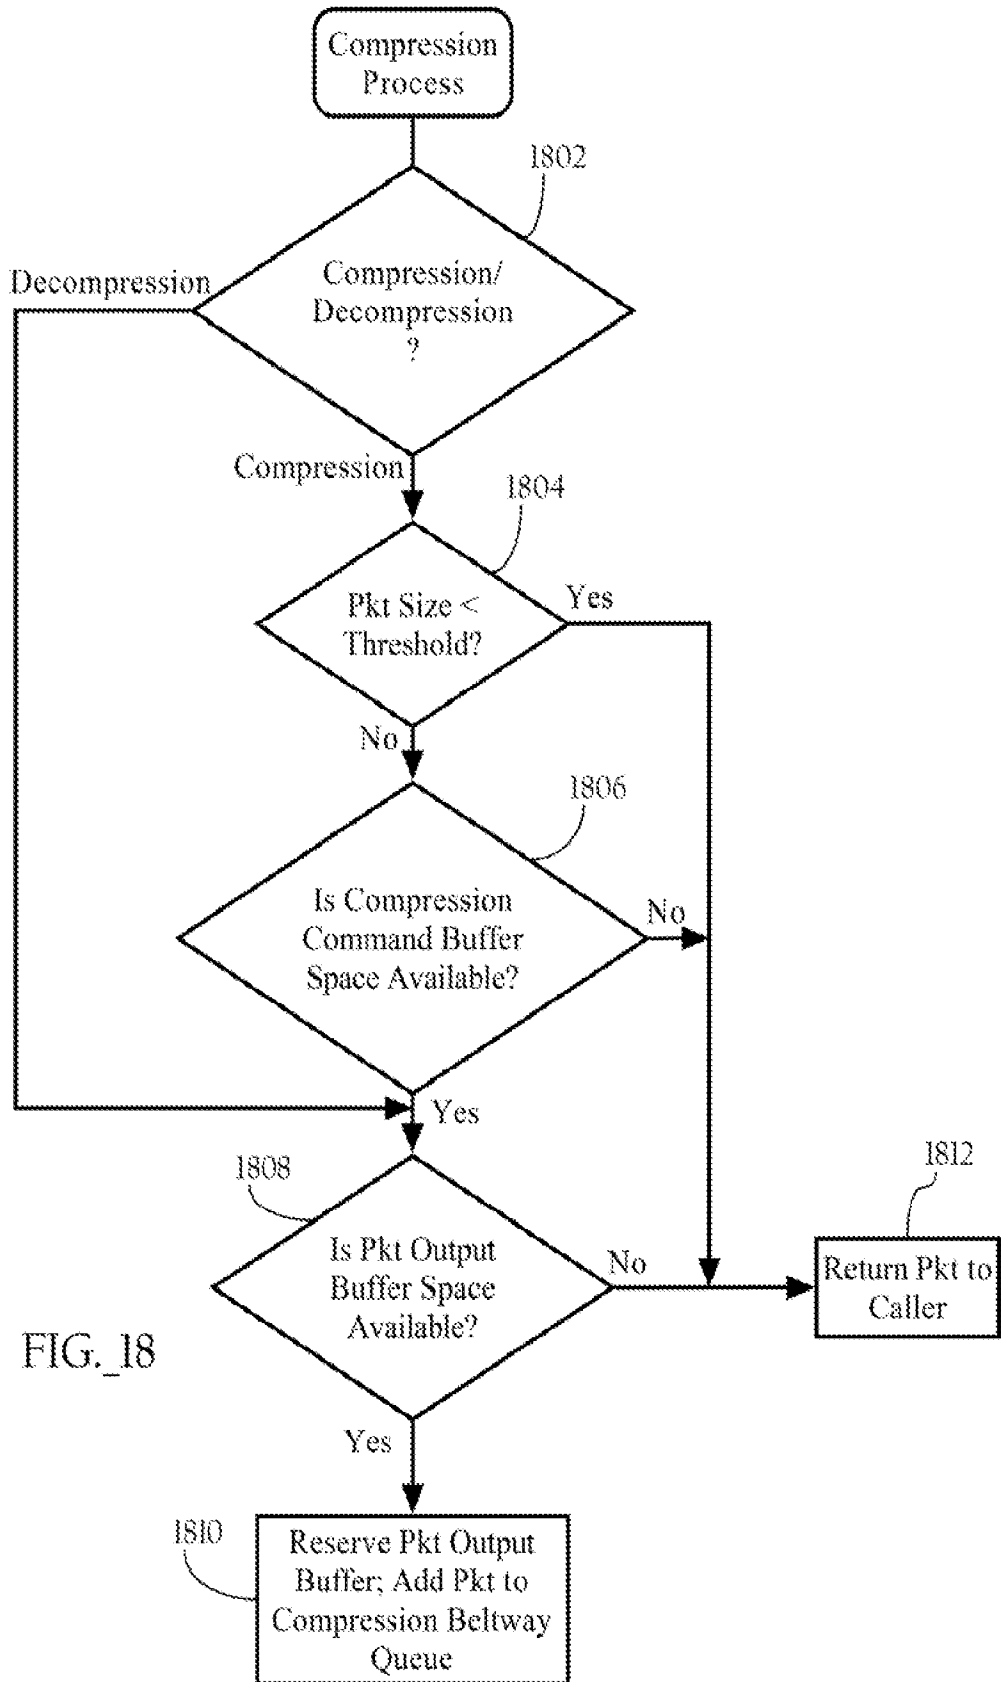


FIG. 18

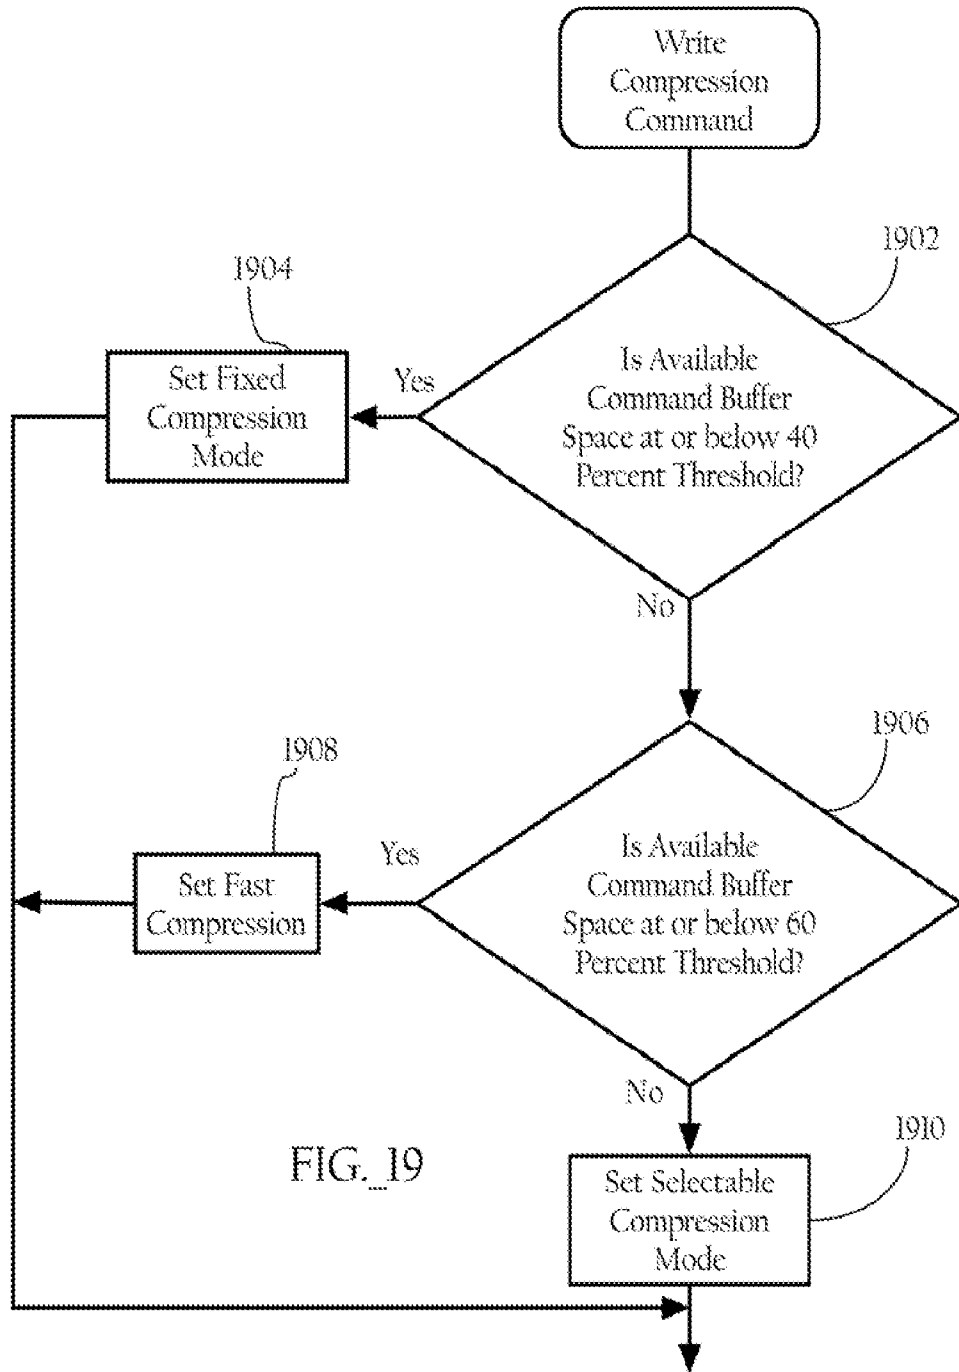


FIG. 19