

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
12 February 2004 (12.02.2004)

PCT

(10) International Publication Number  
**WO 2004/013720 A2**

- (51) International Patent Classification<sup>7</sup>: **G06F**
- (21) International Application Number:  
PCT/US2003/021583
- (22) International Filing Date: 10 July 2003 (10.07.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
10/211,434 2 August 2002 (02.08.2002) US
- (71) Applicant: ASTUTE NETWORKS [US/US]; 16868 Via del Campo Ct., Suite 200, San Diego, CA 92127 (US).
- (72) Inventors: **OSMAN, Fazil, Ismet**; 2471 Canyon Creek Road, Escondido, CA 92025 (US). **KNEE, Simon, John**; 12885 Caminito De Las Olas, Del Mar, CA 92014 (US).
- (74) Agent: **GALLIANI, William, S.**; COOLEY GODWARD LLP, 3000 El Camino Real, Five Palo Square, Palo Alto, CA 94306-2155 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**  
— *without international search report and to be republished upon receipt of that report*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*



WO 2004/013720 A2

(54) Title: HIGH DATA RATE STATEFUL PROTOCOL PROCESSING

(57) Abstract: A method, system, and apparatus for processing messages according to one or more stateful protocols in a data communication system. Messages belonging to various communication flows, such as TCP transport layer "connections", are processed according to a selected stateful protocol. Messages belonging to a single unified flow are distributed among different protocol processing cores (PPCs) at different times. Assignment of flows to PPCs is typically modifiable at the level of specific individual flows, permitting flexible load balancing between PPCs. The architecture of the system facilitates the method by providing a mix of pipelined and parallel processing structures.

## HIGH DATA RATE STATEFUL PROTOCOL PROCESSING

### BACKGROUND OF THE INVENTION

#### Field of the Invention

This invention relates to data transfer processing systems.

#### Description of Related Art

5 Data transfer systems typically convey data through a variety of layers, each performing different types of processing. The number of different layers, and their attributes, vary according to the conceptual model followed by a given communication system. Examples include a model having seven layers that is defined by the International Standards Organization (ISO) for Open Systems Interconnection (OSI), and a five-layer model defined by the American National Standards  
10 Institute (ANSI) that may be referred to as the "Fibre Channel" model. Many other models have been proposed that have varying numbers of layers, which perform somewhat different functions. In most data communication systems, layers range from a physical layer, via which signals containing data are transmitted and received, to an application layer, via which high-level programs and processes share information. In most of the conceptual layer models, a Transport Layer exists  
15 between these extremes. Within such transport layers, functions are performed that are needed to coordinate the transfer of data, which may have been sent over diverse physical links, for distribution to higher-level processes.

Within the transport layer, a communication system coordinates numerous messages (such as packets) that each belong to a particular "flow" or grouping of such messages. Each message  
20 may be identified by its association with a particular flow identification key (flow key), which in turn is typically defined by information about the endpoints of the communication. Transport layer processing is generally performed by processing modules which will be referred to as transport layer terminations (TLTs), which manage data received from remote TLTs (or being transmitted to the remote TLTs) according to a set of rules defined by the transport layer protocol (TLP) selected  
25 for each particular flow. A TLT examines each message that it processes for information relevant to a flowstate that defines the status of the flow to which the message belongs, updates the flowstate accordingly, and reconstitutes raw received data on the basis of the flowstate into proper form for the message destination, which is typically either a remote TLT or a local host. Flows are typically bidirectional communications, so a TLT receiving messages belonging to a particular flow from a  
30 remote TLT will generally also send messages belonging to the same flow to the remote TLT. Management of entire flows according to selected TLPs by maintaining corresponding flowstates

distinguishes transport layer processing from link level processing, which is generally concerned only with individual messages.

There are many well-known TLPs, such as Fibre Channel, SCTP, UDP and TCP, and more will likely be developed in the future. TLPs typically function to ensure comprehensible and accurate communication of information to a target, such as by detecting and requesting retransmission of lost or damaged messages, reorganizing various messages of a flow into an intended order, and/or providing pertinent facts about the communication to the target. Transmission Control Protocol (TCP) is probably the best-known example of a TLP, and is extensively used in networks such as the Internet and Ethernet applications. TCP is a connection-oriented protocol, and information about the state of the connection must be maintained at the connection endpoints (terminations) while the connection is active. The connection state information includes, for example, congestion control information, timers to determine whether packets should be resent, acknowledgement information, and connection identification information including source and destination identification and open/closed status. Each active TCP connection thus has a unique connection ID and a connection state. A TCP "connection" is an example of the more general stateful protocol processing system ("SPPS") concept that is termed "flow" herein, while TCP "connection ID" and "connection state" are examples of the more general SPPS concepts referred to herein as "flow key" and "flowstate," respectively. The flow key in a TLP may be uniquely specified by a combination of the remote link (destination) address (typically an Internet Protocol or "IP" address), the remote (destination) TCP port number, the local link (source) address (also typically an IP address), the local (source) TCP port number, and in some cases a receiver interface ID. It may also be useful to include a protocol indication as part of the general flow key, in order to distinguish flows that have otherwise identical addressing but use different TLPs.

Data communications can also occur in many layers other than the classic transport layer. For example, iSCSI communications occur at layers above the transport layer, yet the communications include stateful messages belonging to a flow and are thus analogous, in some ways, to transport layer communications.

There is a constant demand for higher data rates for data communications systems, as computers are increasingly linked both locally (*e.g.*, over local area networks) and over wide areas (*e.g.*, over the Internet). In order to achieve higher data rates, commensurately faster processing is needed for stateful protocols in transport layers and elsewhere. Faster hardware, of course, may be able to proportionally increase processing speeds. However, hardware speed increases alone will not cost-effectively increase protocol processing speeds as quickly as desired, and thus there is a need for protocol processing systems that enable faster processing, for a given hardware speed, by virtue of their architecture and methods.

## SUMMARY OF THE INVENTION

A method, a system, and an apparatus for stateful protocol processing is described herein. Stateful protocol processing entails maintaining a "state" to track the condition of a "flow" of data. The state of a particular flow is often updated to reflect the processing of individual data "messages" that belong to the flow, and the flow itself is identified by a flow key that is associated, explicitly or implicitly, with each message belonging to the flow. The protocol that is indicated for each message defines the processing steps to be performed on each message, in view of the present state of the flow.

In one aspect, the present invention relates to a method of processing data in a stateful protocol processing system ("SPPS") configured to process a multiplicity of flows of messages. Each flow is associated with a uniquely corresponding flow key that is conveyed by messages belonging to such flow. The method includes receiving a plurality of messages belonging to a particular flow. Various SPPS events associated with the particular flow are then derived from the received messages. The method further includes specifically assigning a first protocol processing core ("PPC") to process one or more events of the particular flow in accordance with a stateful protocol (SP) of the particular flow. In addition, a different second PPC is specifically assigned to process one or more other events of the particular flow in accordance with the SP of the particular flow.

The SPPS described herein may be implemented in many ways, and details of some exemplary embodiments are described below. One embodiment is a method that includes assigning processing of events (*i.e.*, information derived from messages) of a specific flow, as opposed to events from a more general class of flows, to a PPC. Another embodiment is a method that includes assigning processing of events of flows to PPCs irrespective of a preliminary queue in which the event is located, and then transferring the events to a local queue of the assigned PPC. Yet another embodiment is a method that includes steps to receive messages, steps to define events based on the received messages, steps to assign a first event of a flow to a first PPC, and steps to assign a second event of the flow to a second PPC.

A further embodiment is a system for terminating a data communication transport layer, and includes a message receiver module configured to receive messages, each having a flow ID, and to derive events therefrom (the events could, for example, simply be the message as received, but typically would include some changes to the content or form of the received message). The embodiment also includes some PPC modules, and a dispatcher module configured to receive events, determine if a PPC is already assigned to process events of the same flow, and if not then to select a PPC that has been configured to be compatible with the event to process the event without regard to the flow ID associated with the event.

Yet a further embodiment is an apparatus for processing messages, and includes means for receiving messages, more than one means for stateful protocol (SP) processing of flow information, and means for selecting a PPC for processing information of a particular flow, irrespective of the flow ID. Another embodiment is an apparatus that includes more than one PPC microprocessor,  
5 each configured to perform SP processing for messages belonging to an assigned flow, and having local memory to hold the flowstate of such assigned flow. This embodiment also has a dispatcher that receives message information, and forwards at least the flowstate-relevant part of one message to a first PPC microprocessor, and forwards at least the flowstate-relevant part of another message to a second PPC microprocessor.

10 Novel subsystems are also described within the context of a SPPS in which they may be employed. One example is a data dispatcher subsystem, and another is an event-derivation system. Each of these subsystems may in turn have further subsystems; for example, the data dispatcher subsystem may include a lookup subsystem for tracking flow information of flows, particularly those that are not presently being processed by a PPC, and/or a core activity manager subsystem for  
15 keeping track of information that is presently being processed by a PPC. The data dispatcher subsystem may direct data events that belong to a single flow to be distributed to a plurality of parallel protocol processor cores on the basis of PPC loading, and may maintain flowstate for a particular flow independently when no protocol processor core is processing data events of the particular flow.

20 Embodiment details and some alternatives are set forth in the accompanying drawings and the description below. Because all embodiments of the invention cannot reasonably be described herein, the described embodiments must be understood as illustrating, rather than limiting, the invention.

### BRIEF DESCRIPTION OF THE DRAWINGS

25 FIGURE 1A is a block diagram showing general interface connections to a stateful protocol processing system.

FIGURE 1B is a block diagram of a transport layer termination system within a typical computer system.

30 FIGURE 2 is a more detailed block diagram of a stateful protocol processing system such as that of FIGURE 1.

FIGURE 3 is a block diagram showing further details of some of the features of the stateful protocol processing system of FIGURE 2.

FIGURE 4 is a flowchart of acts used in varying the protocol core that is selected to process a flow.

FIGURE 5 is a flowchart of acts performed by a dispatcher module in response to receiving an event belonging to a flow.

5 FIGURE 6 is a flowchart of acts performed by a dispatcher module in response to receiving a "done statement" from a protocol core.

Like reference numbers and designations in the various drawings indicate like elements throughout.

### DETAILED DESCRIPTION

10 Throughout this description, embodiments and variations are described for the purpose of illustrating uses and implementations of the invention. The illustrative description should be understood as presenting examples of the invention, rather than as limiting the scope of the invention.

Stateful protocol processing entails processing data that arrives in identifiable and distinguishable units that will be referred to herein as "messages." A multiplicity of messages will belong to a "flow," which is a group of messages that are each associated with a "flow key" that uniquely identifies the flow. The methods and apparatus described herein for stateful protocol processing are most useful when a multiplicity of different flows is concurrently active. A flow is "active" whether or not a message of the flow is presently being processed, as long as further messages are expected, and becomes inactive when no further processing of messages belonging to the particular flow are expected.

20 A "stateful protocol" defines a protocol for treating messages belonging to a flow in accordance with a "state" that is maintained to reflect the condition of the flow. At least some (and typically many) of the messages belonging to a flow will affect the state of the flow, and stateful protocol processing therefore includes checking incoming messages for their effect on the flow to which they belong, updating the state of the flow (or "flowstate") accordingly, and processing the messages as dictated by the applicable protocol in view of the current state of the flow to which the messages belong.

30 Processing data communications in accordance with TCP (Transmission Control Protocol) is one example of stateful protocol processing. A TCP flow is typically called a "connection," while messages are packets. The flow key associated with each packet consists primarily of endpoint addresses (*e.g.*, source and destination "socket addresses"). A flowstate is maintained for

each active connection (or flow) that is updated to reflect each packet of the flow that is processed. The actual treatment of the data is performed in accordance with the flowstate and the TCP processing rules.

TCP is a protocol that is commonly used in TLT (transport layer termination) systems. A typical TLT accepts messages in packets, and identifies a flow to which the message belongs, and a protocol by which the message is to be processed, from information contained within the header of the packet. However, the information that is required to associate a message with a flow to which it belongs and a protocol by which it is to be processed may be provided in other ways, such as indirectly or by implication from another message with which it is associated, or by a particular source from which it is derived (for example, if a particular host is known to have only one flow active at a time, then by implication each message from that host belongs to the flow that is active with respect to that host).

Moreover, stateful protocol processing as described herein may be utilized in places other than TLT systems, in which case the information about flow and protocol may well be provided elsewhere than in an incoming packet header. For example, an incoming TCP packet may encapsulate data that is to be processed according to an entirely different protocol, in a different "layer" of processing. Accordingly, the stateful protocol processing effected within the context of a TLT system described herein provides a specific example of a general stateful protocol processing system ("SPPS"). Messages belonging to one stateful protocol flow may, for example, be encapsulated within messages belonging to a distinct stateful protocol. The well-known communication protocol referred to as "SCSI" provides examples of data communication at layers other than a transport layer. A common use of SCSI is between a host and a peripheral device such as a disk drive. SCSI communications may take place over a special purpose connection dedicated to SCSI communications, or they may be encapsulated and conveyed via a different layer. SCSI may be encapsulated within messages of some transport layer protocols, such as Fibre Channel and TCP. "FCP" is a protocol by which SCSI messages are encapsulated in Fibre Channel protocol messages, while "iSCSI" is a protocol by which SCSI messages are encapsulated in TCP messages. FCP and iSCSI are each stateful protocols.

One example of such encapsulation involves information belonging to a first stateful flow, such as an iSCSI flow, that is communicated over a local network within messages belonging to a distinct second stateful flow, such as a TCP connection. A first SPPS may keep track of the state of the encapsulating TCP connection (flow). The same SPPS, or a different second one, may determine that some of the messages conveyed by the encapsulating flow form higher-level messages that belong to an encapsulated iSCSI flow. The flow key of the encapsulated iSCSI flow may be contained within each encapsulated message, or it may be determined by implication from

the flow key of the encapsulating TCP/IP packets that are conveying the information. Given knowledge of the flow key of the encapsulated flow, and of the protocol (iSCSI) by which the encapsulated flow is to be processed, the SPPS may maintain a state for the iSCSI flow, and may identify and process the messages associated with the flow in accordance with the specified protocol (iSCSI, in this example).

Thus, a transport layer termination system may provide a good example of a SPPS (stateful protocol processing system). Indeed, a TLT is likely to include at least some stateful processing, thus qualifying as a SPPS. However, a SPPS can be utilized for other data communication layers, and for other types of processing, as long as the processing includes updating the flowstate of a flow to which a multiplicity of messages belong, in accordance with a stateful protocol that is defined for the messages. Therefore, although the invention is illustrated primarily with respect to a TLT system, care should be taken not to improperly infer that the invention is limited to TLT systems.

FIGURE 1A illustrates interface connections to a SPPS 100. A SPPS packet input processing block 102 may accept data in packets from any number of sources. The sources typically include a host connection, such as "Host 1" 104, and a network connection, such as "Network 1" 106, but any number of other host connections and/or network connections may be used with a single system, as represented by "Host N" 108 and "Network M" 110. A protocol processing block 112 processes incoming data in accordance with the appropriate rules for each flow of data (*i.e.*, stateful protocol rules such as are defined by the well-known TCP, for stateful messages specified for processing according to such stateful protocol). Flows generally involve bidirectional communications, so data is typically conveyed both to and from each host connection and/or network connection. Consequently, a packet output processing block 114 delivers data to typically the same set of connections ("Host 1" 104 to "Host N" 108 and "Network 1" 106 to "Network M" 110) from which the packet input processing block 102 receives data.

FIGURE 1B provides an overview of connections to a TLTS 150 that provides an example of a simple SPPS as implemented within a computing system 152. A single host system 154 is connected to the TLTS 150 via a connection 156 that uses a well-known SPI-4 protocol. The host 154 behaves as any of the hosts 104-108 shown in FIGURE 1A, sending messages to, and receiving messages from, the TLTS 150. The TLTS 150 is connected to a Media Access Control ("MAC") device 158 via another SPI-4 protocol connection 160. The MAC 158 is connected to a network 162 via a suitable connection 164. The MAC converts between data for the TLTS (here, in SPI-4 format), and the physical signal used by the connection 164 for the network 162. The network 162 may have internal connections and branches, and communicates data to and from remote communications sources and/or targets, exemplified by as "source/target system 1" 170,



"source/target system 2" 180, and "source/target system 3" 190. Any number of communication source/targets may be accessed through a particular network. Source/target systems may be similar to the computing system 152. More complicated source/target systems may have a plurality of host and network connections, such as is illustrated in FIGURE 1A. Thus, some source/target systems may effectively connect together a variety of different networks.

FIGURE 2 is a block diagram showing modules of an exemplary SPPS 200. In one embodiment, two SPI-4 Rx interface units 202 and 204 receive data over standard SPI-4 16-bit buses that accord with "System Packet Interface Level 4 (SPI-4) Phase 2: OC-192 System Interface for Physical and Link Layer Devices. Implementation Agreement OIF-SPI4-02.0," Optical Internetworking Forum, Fremont, CA, January 2001 (or latest version). The number of connections is important only insofar as it affects the overall processing capability needed for the system, and from one to a large number of interfaces may be connected. Each individual interface may process communications to any number of network and/or host sources; separate physical host and network connections are not necessary, but may be conceptually and physically convenient. Moreover, while SPI-4 is used for convenience in one embodiment, any other techniques for interface to a physical layer (e.g., PCI-X) may be used alternatively or additionally (with processing in the corresponding input blocks, e.g., 202, 204, conformed) in other embodiments.

### Message Splitting

Still referring to FIGURE 2, data received by the interfaces 202 and 204 is conveyed for processing to message splitter modules 206 and 208, respectively. The transfer typically takes place on a bus of size "B." "B" is used throughout this document to indicate a bus size that may be selected for engineering convenience to satisfy speed and layout constraints, and does not represent a single value but typically ranges from 16 to 128 bits. The message splitter modules 206 and 208 may perform a combination of services. For example, they may reorganize incoming messages (typically packets) that are received piece-wise in bursts, and may identify a type of the packet from its source and content and add some data to the message to simplify type identification for later stages of processing. They may also split incoming messages into "payload" data and "protocol event" (hereafter simply "event") data.

As the data arrives from the SPI-4 interface, a message splitter module such as 206 or 208 may move all of the data into known locations in a scratchpad memory 210 via a bus of convenient width B. Alternatively, it may send only payload data to the scratchpad, or other subset of the entire message. The scratchpad 210 may be configured in various ways; for example, it may function as a well-known first-in, first-out (FIFO) buffer. In a more elaborate example, the scratchpad 210 may be organized into a limited but useful number of pages. Each page may have a

relatively short scratchpad reference ID by which a payload (or message) that is stored in the scratchpad beginning on such page can be located. When the payload overruns a page, an indication may be provided at the end of the page such that the next page is recognized as concatenated, and in this manner any length of payload (or message) may be accommodated in a block of one or more pages that can be identified by the reference ID of the first page. A payload length is normally part of the received header information of a message. The scratchpad reference ID may provide a base address, and the payload may be disposed in memory referenced to the base address in a predetermined manner. The payload terminates implicitly at the end of the payload length, and it may be useful to track the number of bytes received by the scratchpad independently, in order to compare to the payload length that is indicated in the header for validation. If the scratchpad also receives the header of a message, that header may be similarly accessed by reference to the scratchpad reference ID. Of course, in this case the payload length validation may be readily performed within the scratchpad memory module 210, but such validation can in general be performed many other places, such as within the source message splitter (206, 208), within the dispatcher 212, or within a PPC 216-222, as may be convenient from a data processing standpoint.

#### Event Derivation

A typical function of a message splitter 206, 208 is to derive, from the incoming messages, the information that is most relevant to stateful processing of the messages, and to format and place such information in an "event" that is related to the same flow as the message from which it is derived. For example, according to many transport layer protocols, "state-relevant" data including flow identification, handshaking, length, packet order, and protocol identification, is disposed in known locations within a packet header. Each stateful protocol message will have information that is relevant to the state of the flow to which it belongs, and such state-relevant information will be positioned where it can be identified. (Note that systems that perform stateful protocol processing may also process stateless messages. TLPs, for example, typically also process packets, such as Address Request Protocol or ARP packets, which are not associated with an established flow and thus do not affect a flowstate. Such "stateless" packets may be processed by any technique that is compatible with the presently described embodiments. However, these techniques are not discussed further herein because the focus is on the processing of stateful messages that do affect a flowstate for a message flow.)

The event that is derived from an incoming message by a message splitter module such as 206 or 208 may take a wide range of forms. In the simplest example, in some embodiments it may be the entire message. More typically, the event may exclude some information that is not necessary to make the decisions needed for SP processing. For example, the payload may often be excluded, and handled separately, and the event may then be simply the header of a message, as

received. However, in some embodiments information may be added or removed from the header, and the result may be reformatted, to produce a derived event that is convenient for processing in the SPPS.

### Event Typing

5           Received messages may, for example, be examined to some extent by the interface (202, 204) or message splitter (206, 208) modules, and the results of such examination may be used to derive a "type" for the event. For example, if a packet has no error-checking irregularities according to the protocol called for in the flow to which the packet belongs, then the event derived from such package may be identified with an event "type" field that reflects the protocol and  
10           apparent validity of the message. Each different protocol that is processed by the SPPS may thus have a particular "type," and this information may be included in the event to simplify decisions about subsequent processing. Another type may be defined that is a message fragment; such fragments must generally be held without processing until the remainder of the message arrives. Message fragments may have subtypes according to the protocol of the event, but need not. A  
15           further event type may be defined as a message having an error. Since the "type" of the event may be useful to direct the subsequent processing of the event, messages having errors that should be handled differently may be identified as a subtype of a general error. As one example, error type events may be identified with a subtype that reflects a SP of the event.

          Any feature of a message (or of a derived event) that will affect the subsequent processing  
20           may be a candidate for event typing. Thus, event typing may be very simple, or may be complex, as suits the SPPS embodiment from an engineering perspective. Event typing is one example of augmentation that may be made to received message information in deriving an event. Other augmentation may include revising or adding a checksum, or providing an indication of success or failure of various checks made upon the validity of the received message. Relevant locations may  
25           also be added, such as a scratchpad location indicating where the message information may be found within the scratchpad memory 210. Note that if a message source that uses the SPPS, such as a host, is designed to provide some or all of such "augmenting" information within the message (e.g., the header) that it conveys to the SPPS, then the message splitter may not need to actually add the information in order to obtain an "augmented" event.

30           In addition to augmenting message information, event derivation may include reformatting the event information to permit more convenient manipulation of the event by the SPPS. For example, processing may be optimized for certain types of events (such as TCP events, in some systems), and deriving events of other types may include reformatting to accommodate such optimized processing. In general, then, events may be derived by doing nothing to a received

message, or by augmenting and/or reformatting information of the message, particularly state-relevant information, to aid later processing steps. For TCP, for example, the resulting event may consist primarily of the first 256 bytes of the packet, with unnecessary information removed and information added to reflect a scratchpad location in which it is copied, the results of error checking, and the event typing. If a host is configured to prepare data in a form that is convenient, a resulting host event issued from the message splitter may be the first bytes of the message (*e.g.*, the first 256 bytes), with few or no changes.

It may be convenient to implement the message splitter functions using an embedded processor running microcode, which lends itself to reprogramming without a need to change the device design. However, the message splitter function may alternatively be implemented via software executed in a general-purpose processor, or in an application specific integrated circuit (ASIC), or in any other appropriate manner.

Many alternatives are possible for the particular set of processing steps performed by message splitter modules such as 206 and 208. For example, a "local proxy" of the flow ID (*i.e.*, a number representing the flow ID of the message that suffices to identify the flow within the SPPS and is more useful for local processing) could be determined and added to the event at the message splitter – a step that is performed during a later processing block in the illustrated embodiments. Also, it is not necessary that incoming messages be split at all. Instead, incoming messages may be kept together: for example, they may be stored intact in the scratchpad memory so as to be available to many parts of the system, or they may be forwarded in their entirety directly to the event dispatcher 212 and thence to the protocol processing cores (PPCs) 216-222 that are described below in more detail. If incoming messages are not split, then these modules 206, 208 might, for example, be renamed "packet preprocessors" to reduce confusion. The skilled person will understand that, in many cases, design convenience primarily determines which module performs any particular acts within a complex system.

### Event Dispatcher

As shown in FIGURE 2, the events prepared by the message splitters 206, 208 are forwarded to an event dispatcher module 212, where they may be entered into a queue. The event dispatcher module 212 (or simply Dispatcher) may begin processing the incoming event by initiating a search for a local flow ID proxy, based on the flow identification "key" that arrives with the message.

Local Flow ID Proxy

The flow identification key (or simply "flow key") uniquely identifies the flow to which the message belongs in accordance with the SP (for example, a TLP) used by the flow. The flow key can be very large (typically 116-bits for TCP) and as such it may not be in a format that is convenient for locating information maintained by the SPPS that relates to the particular flow. A local flow ID proxy may be used instead for this purpose. A local flow ID proxy (or simply "local proxy ID," "local flow ID," or "proxy ID") generally includes enough information to uniquely identify the particular flow within the SPPS, and may be made more useful for locating information within the SPPS that relates to the particular flow. For example, a local flow ID proxy may be selected to serve as an index into a flowstate memory 214 to locate information about a particular flow (such as a flowstate) that is maintained within the SPPS. Not only may a local flow ID proxy be a more convenient representative of the flow for purposes of the SPPS, it will typically be smaller as well.

A local flow ID proxy may be determined within the Dispatcher module or elsewhere, such as within the message splitter modules 206, 208 as described previously. Given the very large number of local flow ID proxies that must be maintained, for example, in large TLTSs (transport layer termination systems), determining the proxy ID may be a nontrivial task. If so, it may be convenient from an engineering perspective to make such determination by means of a separate "lookup" module, as described below. In some embodiments, such a lookup module may be a submodule of the message splitter modules 206, 208, or it may be a submodule of the Dispatcher module, or it may be best designed as independent and accessible to various other modules.

A search for the local flow ID proxy may be simplified, or even eliminated, for events received from a host that is configured to include the local flow ID proxy rather than (or in addition to) the usual SP flow key that will accompany flow messages on a network. Such a host configuration can reduce the workload of whatever module would otherwise determine the local flow ID proxy, *e.g.*, the Dispatcher. Another way to reduce the local flow ID proxy lookup effort may be to maintain a "quick list" of the most recently used flow IDs, and their associated proxies, and to check this list first for each arriving message or event.

If a message arrives that belongs to a flow for which no local flow ID proxy or flowstate is known, the Dispatcher 212 may create a new local flow proxy ID. In many cases the Dispatcher (or a lookup submodule) may then initialize a flowstate for such new flow. It may be useful to select such proxy ID as a value that will serve as a table entry into memory that may be used to store a flowstate for such new flow in a convenient memory, such as flowstate memory 214. Such memory may be quite large in large systems, requiring special management.

Memories

Each distinct "memory" described herein, such as the scratchpad memory 210 and the flowstate memory 214, typically includes not only raw memory but also appropriate memory controller facilities. However, the function of the memory controller is generally not central to the present description, which merely requires that the memory either store or return specified blocks of data in response to requests. Because SPPSs as described herein may be made capable of concurrently processing millions of active flows (or may be limited to processing a few thousand, or even fewer, active flows), and because a typical flowstate may be approximately 512 bytes, multiple GB of memory may be needed to implement the SPPS of FIGURE 2. Techniques for implementing such large memories are known and constantly evolving, and any such known or subsequently developed technique may be used with any type of memory to form the SPPS of FIGURE 2, so long as adequate performance is achieved with such memory. Memories are distinguished from each other as distinct memories if they function in a substantially independent manner. For example, distinct memories may be independently addressable, such that addressing a data item stored in one memory does not preclude simultaneously addressing an unrelated item in a distinct memory. Distinct memories may also be independently accessible, such that accessing an item in one memory does not preclude simultaneously accessing an unrelated item in a distinct memory. Due to such independence, distinct memories may in some cases avoid data access bottlenecks that may plague common (or shared) memories.

Lookup Submodule

The Dispatcher module 212 illustrated in FIGURE 2 may include submodules that perform particular subsets of the Dispatcher tasks. For example, it may be useful to incorporate a separate "lookup" module to perform the function of looking up a local flow ID proxy based on the flow key that is included in the arriving event. Another function of the Dispatcher 212 may be to establish and maintain flow timers for active flows, as may be required by the particular SP associated with each flow. When it is convenient to maintain such flow timers in memory that is indexed by the local flow ID proxy, the lookup module may also conveniently perform the function of monitoring the flow timers. Also, the Dispatcher 212 may provide the flowstate to a PPC when assigning it to process events of a flow. If the flowstate is similarly maintained in memory at a location indexed by the local flow ID proxy, then this may be another function that may conveniently be performed by the lookup module. Such a lookup module may be independent, or it may be essentially a submodule of the Dispatcher. The lookup module could also be associated primarily with other sections of the system. For example, it could be primarily associated with (or even a submodule of)

a message splitter module 206, 208, if that is where the lookup tasks are performed, or it could be primarily associated with the PPCs 216-222 if the lookup tasks were performed primarily there.

The lookup process may require extensive processing, such as a hash lookup procedure, in order to select or determine a local flow ID proxy based on raw flow identification or "flow key." As such, a lookup module (or submodule) may be implemented with its own microprocessor system and supporting hardware. When flow ID proxy determination is performed by a lookup module (or submodule), the Dispatcher may assign and transfer an event to a PPC without waiting for the determination to be completed, and the lookup module can later transfer flow information (obtained by use of the local flow ID proxy) to the assigned PPC without further interaction with the Dispatcher.

Once a "lookup" or other submodule is established as a distinct entity, it may as a matter of design convenience be configured to perform any of the tasks attributed to the Dispatcher (or other module in which it is located or with which it is associated, and indeed in many cases may perform tasks that are attributed, in the present description, to other modules, such as the message splitter modules. The ability to move functionality between different functional modules is a common feature of complex processing systems, and the skilled person will understand that moving functionality between modules does not, in general, make a system significantly different.

Many other functions may be performed by the Dispatcher 212, or by its submodules. For example, the Dispatcher may request a checksum from the scratchpad memory 210 reflecting the payload of the message, combine it with a checksum included with the event that covers that portion of the message converted into the event, and incorporate the combined checksum into the event. A bus of modest size is shown between the Dispatcher 212 and the other processing blocks that is sufficient for this purpose. As with many Dispatcher functions, this function could be performed elsewhere, such as in the message splitter blocks 206, 208, or during later processing.

#### Director Submodule

Another module, or Dispatcher submodule, may be created to perform some or all of the decision making for the Dispatcher. Such submodule, which will be referred to as a "Director," may perform the steps involved in selecting a particular PPC to handle a particular event of a flow, and keeping track, for the overall SPPS (stateful protocol processing system), of the status of active flow processing in the various PPCs.

The "flow processing status" maintained by the Director submodule may indicate, for example, that other events of the flow are presently being processed by a PPC, or that a new flowstate generated after PPC processing of a previous event (of the flow) is presently being written to the flow state memory. It may also indicate if the flow is being torn down, or that a timer event is pending for that flow. Such flow processing status information may be used, for example, to cause

the Director submodule to delay the forwarding of an event to a PPC when appropriate, such as to avoid overwriting a flowstate while the flow processing status of a flow says that its flowstate is being written from a PPC to the flowstate memory. Once the update of the flowstate memory is complete, as reflected by the flow processing status, the new event may be forwarded to a PPC.

5           The Director submodule's flow processing status information may also be used, for example, to prevent timer expirations from being improperly issued while a flow is being processed by a PPC. Such timer events should not be issued if the very act of processing an event may cause such a timer expiration to be cancelled. The Director submodule may refer to the flow processing status information before allowing timer events to be issued to PPCs, so that such timer events are  
10           issued only when no other events are active for that flow. As with the lookup submodule, organization of the Director as a distinct module may permit the Dispatcher to simply hand off an incoming event to the Director.

### Protocol Processing Cores and Buses - Structural Introduction

          Having established a local flow ID proxy for a message, the Dispatcher 212 determines  
15           where the message event (or entire message, if messages and events are not split) should be processed in accordance with the SP associated with the flow. In some embodiments, the bulk of such SP processing is performed by a Protocol Processing Core ("PPC"). A cluster having a number of PPCs is represented by the PPCs 216 through 218, while PPCs 220 and 222 represent another cluster of PPCs. Two PPC clusters are shown, but any number of such PPC clusters may be  
20           used. For example, one TLTS embodiment may comprise only a single cluster of PPCs, while a complex SPPS embodiment may include hundreds of clusters. Two of the PPCs in a cluster are shown in FIGURE 2, but two or more PPCs may be used in any given cluster, with five PPCs per cluster being typical. Though it may be convenient for design symmetry, the number of PPCs in each cluster need not be identical. The particular organization of PPCs into clusters is selected, in  
25           part, to facilitate the transfer of data by reducing bus congestion. Each cluster may utilize an intracluster intercore bus 224 (or 226) interconnecting PPCs of each cluster, and each cluster will typically be connected to a bus network and control block 228 by a bus 230 or 232. Data between the Dispatcher 212 and the PPCs may be organized by a bus network and control block 228. The bus network and control block 228 functions primarily as a "crossbar" switch that facilitates  
30           communication between a variety of modules, as described in more detail below.

          PPCs (*e.g.*, 216-222) typically include a processor core and microcode (*i.e.*, some form of sequential instructions for the processor core) that enables the PPC to process events that are submitted to it. They also typically include local memory, which the PPC can access without interfering with other PPCs, sufficient to hold the relevant flowstate data of a flow that the PPC is



processing. It will typically be convenient to maintain much or all of the flowstate of a particular flow in the local memory of the PPC processing a message event for that flow. The PPC local memory may be organized into a number of blocks or "workspaces" that are each capable of holding a flowstate. PPCs will typically have a queue for incoming events, and workspaces for several  
5 different flows having events in the queue that are concurrently being processed by the PPC.

The buses represented herein are described as being bidirectional in nature. However, if convenient, the buses may be implemented as two one-way buses that in some cases will not be of equal bit-width in both directions. Thus, a bus indicated as having a width B bits represents a bus width that that may be selected for convenience in a particular implementation, and may be  
10 directionally asymmetrical. The typical considerations for bus size apply, including space and driver constraints of the physical layout, and the required traffic needed to achieve a particular performance target. The buses are not shown exhaustively; for example, a message bus may usefully be connected (for example by daisy-chaining) between all of the physical pieces of the TPTS, even though such a bus is not explicitly shown in FIGURE 2. Moreover, if the SPPS is  
15 implemented as program modules in software or firmware running on a general processing system, rather than in a typical implementation that employs ASICs having embedded microprocessors, the buses represented in FIGURE 2 may represent data transfer between software modules, rather than hardware signals.

#### Assigning Events to a PPC

In some embodiments of the present invention, the Dispatcher 212 selects a particular PPC  
20 to process events associated with a particular flow. There are a number of considerations for such assignment. First, the PPC must be one of the PPCs that are compatible with, or configured to process, events of the type in question. Such compatibility may be determined in the Dispatcher, or in a flow processing status subsystem of the Dispatcher, by means of a table of PPCs that indicates  
25 the event types or protocols the PPC is compatible with, which may in turn be compared with the protocol or event type requirements of the incoming event. In some embodiments the event is marked with an indication of its "type" at another stage of processing, for example in the message splitter module. The Dispatcher then needs only select a PPC that is compatible based on the predetermined "type" of the event. Typically, the event types will be so defined that all messages  
30 having state-relevant information for a particular flow will also have the same event type, and can be processed by the same PPC. Thus, a PPC will be selected from the constellation of PPCs that can process the indicated event type.

A PPC is selected from this constellation of compatible PPCs according to an algorithm that may, for example, compare PPC loading to find a least-loaded PPC, or it may select a PPC in a

round-robin manner, or it may select PPCs randomly. Typically, events of each flow are specifically directed to a PPC, rather than being directed to a PPC as a member of a class of flows. Such individualized processing of each flow permits load balancing irrespective of the attributes of a class of flows. When flows are assigned as members of a class, such as one that shares certain features of a flow ID (or flow key), it may happen that a large number of such a class needs to be processed concurrently, overwhelming the capacity of a PPC, while another PPC is unloaded. This effect may be accentuated when flows are assigned to PPC in classes that have a large number of members. While many embodiments assign each flow uniquely (in a class size of one), it may be effective in some embodiments to assign flows in classes, particularly small classes or classes whose membership can be changed to balance loading.

Similar effects for load balancing may be achieved, even if flows have been assigned in a large class, if a mechanism is provided for releasing specific flows from assignment to particular PPCs. In many embodiments, both assignment and release of flows to PPCs is done for individual or specific flows. Finally, even if both assignment of flows to a PPC, and release of flows from a PPC, is performed for classes of flows, an equivalent effect may be achieved by making the classes flexibly reassignable to balance loading. That is, if the class that is assigned to a PPC can be changed at the level of specific flows, then loading can be balanced with great flexibility. In each case it is possible to change the flows assigned to a PPC in singular units, such that a flow is ultimately assigned to a PPC essentially irrespective of any fixed class attributes, such as characteristics that hash a flow ID to a particular value, and similarly irrespective of other flows that may be assigned to that PPC (or to another PPC).

After selecting a PPC, the Dispatcher 212 forwards the event to the PPC together with instructions regarding a flowstate "workspace." As mentioned above, the decisions for selecting a PPC may be performed in the Director submodule of the Dispatcher. In a typical embodiment, the Dispatcher 212 first determines if an incoming event belongs to a flow that already has events assigned to a particular PPC. A submodule, such as a Core Activity Manager that tracks the activity of PPCs, may perform this determination in some embodiments, while in others embodiments the Director submodule may perform these functions. In the case that a PPC is already assigned for events of the flow of the incoming event, the incoming event is typically forwarded to the same PPC, which may already have the flowstate present within its local memory.

However, if no PPC is presently assigned to the flow, then the Dispatcher selects a particular PPC, for example the PPC 216, to process the incoming event (or assigns the flow to the particular PPC). Selection may be based upon information of the Core Activity Manager, which maintains activity status that can be used to balance loading on the various (compatible) PPCs. The Director submodule may perform the actual assignment and balancing decisions, and in some

embodiments the Director and the Core Activity Manager are substantially a single submodule having a dedicated processor and program code to perform these tasks. The assignment may be simply "round robin" to the compatible PPC that has least recently received an event, or on the basis of PPC queue fullness, or otherwise.

5 After a PPC 216 is assigned to process the incoming event, a workspace is selected in the local memory of the PPC 216 and the current flowstate of the flow of the incoming event is established in the selected workspace. Selection of the workspace may be done by the Dispatcher module (for example, by its Director submodule), or otherwise, such as by the PPC on a next-available basis. The flowstate may be established in the selected workspace in any convenient  
10 manner. For example, the Dispatcher may send the flowstate to the PPC via the Dispatcher (e.g., as an action of the lookup submodule), or the PPC itself may request the flowstate from a memory (e.g., the flowstate memory 214). The event is typically delivered from the Dispatcher 212 to an input queue of the PPC 216, and is associated with the selected workspace. Also, separately or as  
15 part of the event, the size and location of the data payload in scratchpad memory (if any) is typically conveyed to the PPC 216. Having this information, the PPC 216 will be able to process the event when it is reached in the queue, as described subsequently in more detail. When the PPC 216 finishes processing a particular event, it will, in some embodiments, transmit a "done" message to the Dispatcher 212, so that the Dispatcher can track the activity of the PPC. A submodule such as the Core Activity Module or the Director may, of course, perform such tracking.

#### 20 Counting Events to Track Active Flow Processing

Having transmitted an event to a selected PPC (216), the Dispatcher 212 increments an event counter in a location associated with the flow (and thus with the PPC 216). The event counter may be maintained in a local memory block, associated with the local flow ID proxy, that is reserved for such information about current PPC processing (e.g., in the core activity manager  
25 within the Dispatcher), or in another convenient location. The event counter is incremented each time an event is sent to the PPC, and is decremented each time the PPC returns a "done" message for that flow. As long as the event counter is non-zero, a PPC is currently processing an event for the associated flow. When the event counter reaches zero for a particular flow, the PPC (216) no longer has an event to process for the particular flow, and those of its resources that were allocated  
30 for processing the particular flow may be released to process other flows. Note that the PPC 216 may be processing events of other flows, and that its release from processing the particular flow may be made irrespective of such other flows.

If the event counter associated with the flow of an event arriving at the Dispatcher 212 is not zero, then it may be preferable to assign and transfer the arriving event to the same PPC. In  
35 some embodiments, if a PPC is already processing an event, the global (i.e., flowstate memory 214)

version of the flowstate is no longer valid. Rather, only the flowstate in the PPC workspace is valid. In such embodiments, the valid flowstate in the present PPC workspace should be made available to a subsequently selected PPC, which in turn should be done only after the present PPC is finished processing the event. Accordingly, at least in such embodiments, it will generally be more convenient to assign the same PPC to process arriving events belonging to a selected flow until that PPC completes all pending events for the selected flow.

An event arriving at the Dispatcher 212 for a specified flow that is already assigned to a PPC may sometimes need to be transferred, or assigned to a different PPC. In such a case it may be convenient to retain the event in the Dispatcher 212 until the current PPC completes processing all events it has been assigned. Holding the event in the Dispatcher 212 avoids the need to coordinate two PPCs that are simultaneously updating a flowstate for the particular flow. Before such handover occurs, it may also be convenient to allow the PPC to "check-in" its workspace (memory reflecting the present flowstate) to the Flow Memory before assigning the new PPC. Alternatively, the workspace may be transferred from the current PPC directly to the new PPC after all events of the current PPC queue have been processed.

If an event arrives at the Dispatcher for a flow that is active, but the related event counter is zero when an event arrives at the Dispatcher 212 (indicating that no PPC is presently assigned to the flow), then the Dispatcher (or its Director submodule) will select a PPC that is available to process that event type. The selection is typically independent of previous processing, and may be based on various factors such as load sharing and event-type processing capability. As such, the PPC selected next will likely differ from the PPC that previously processed events belonging to the flow. However, in some embodiments consideration may be given to previous processing of a particular flow by a particular PPC, such as when the PPC in fact retains useful state information. Once the PPC selection is made, processing continues as described previously, with the event conveyed to the new PPC, and the flowstate disposed in a local workplace selected within the PPC. The Dispatcher 212 either transfers the current flowstate to the new PPC or indicates where in the flowstate memory 214 the present flowstate is to be found.

An event counter is just one means that may be used to determine whether a particular PPC is presently processing a previous event of the same flow. Alternatively, for example, the PPC presently processing an event of a flow might flag the Dispatcher 212 when it finds no events in its input queue associated with an active workspace. Any other appropriate procedure may also be used to determine whether a PPC is presently assigned to processing a particular flow.

#### Updating Flowstate and Releasing a PPC

A PPC may be released from responsibility for processing events of a particular flow after the associated event counter reaches zero. Such a release means that the PPC may be assigned to

process events of a different flow, since it will generally therefore have a workspace free. In general, the PPC may be processing other flows at the same time, and the release does not affect the responsibilities of the PPC for such other flows. In the typical circumstance that the event counter (or other indication) shows that events of a particular flow may be reassigned to another PPC for processing, the SPPS is enabled to balance PPC processing loads by shifting specific individual flows between different PPCs (of those able to handle the event types of the flow) independently of other flows that may be handled by the PPCs. As compared with techniques that cause PPCs to handle events for a class of flows (such as a class of flows whose flow keys have certain characteristics), such independent flow assignment may reduce the statistical probability that one or more PPCs are idle while another PPC is processing events continuously.

Before a PPC is released, the flow memory that has been updated by the PPC (216) is stored where it will be available to a different PPC that may be selected at a later time to process the same flow. This may be accomplished in any appropriate manner, for example by transferring the contents of the relevant PPC workspace to the Dispatcher 212 and thence to the flowstate memory 214. Alternatively, the PPC (216) may convey the flowstate information to a known location in the flowstate memory 214 in cooperation with the Dispatcher 212, so that the Dispatcher is aware that the flowstate has been updated and is ready for future access. The flowstate may be conveyed more directly from the PPC (216) to the flowstate memory 214, such as via a bus 234 from the bus network and control block 228. The bus 234 may be used for either "checkout" of a flowstate from the flowstate memory 214 to a PPC, or for "check-in" of an updated flowstate from a PPC to the flowstate memory 214. When the event counter reaches zero, and the flowstate has been checked-in to the flowstate memory 214, the present PPC may be released and the flow will revert to a condition reflecting that no PPC is currently assigned to it. Within the PPC, the flowstate workspace may be indicated as free.

An alternative to storing flowstates in the flowstate memory 214 may be used in some embodiments. For a SPPS that is provided with sufficient memory local to the PPCs, the flowstate may be maintained in a workspace of the last PPC that processed it until such time as it is needed elsewhere, such as in another PPC. In such embodiments, the flowstate may be transferred to the appropriate workspace in the new PPC via an intra-cluster bus such as 224 or 226. This is more likely to be a practical alternative for small SPPSs that handle a limited number of concurrent flows.

### Socket Memory and Output Processing

In TLP applications that guarantee message delivery, for example TCP, one requirement is the confirmation that a sent message was correctly received. In these TLPs, if the message is not correctly received, the message should be retransmitted. Because it may be some time before a

request for retransmission arrives, transmitted messages need to be maintained in memory (e.g., in a "send buffer") for some period of time. Send buffering may be required even before first transmission, for example when the output target (e.g., Host1 104 or Network 1 106 in FIGURE 2) is not ready to accept data. Similarly, a "receive buffer" is frequently required. For example, messages may be received out of order, or as fragments, and these must be saved for a period of time to comply with TCP rules that require completing the messages and putting them in the correct order. While messages could simply be stored in the scratchpad memory 210, for large systems entailing large send and receive buffers, it may be more convenient to establish a separate "socket memory" 236 to store large quantities of data for somewhat extended times. Such a socket memory 236 may interface with the scratchpad memory 210 via a bus 238 as shown in FIGURE 2, and with the bus network and PPC cluster control 228 via another bus 240. (Due to substantial traffic, in some embodiments, the bus 240 may actually comprise several individual bus structures.)

The socket memory 236 may provide data intended for output to an output processor and SPI-4 Interfaces 242, 244 via buses 246 and 248. However, when data to be output is still present in the scratchpad memory 210, in some instances it may be quicker to provide the data to the output processors 242, 244 directly via buses 250, 252. The output processing may include tasks such as the preparation of message headers, primarily from the event data, calculation of checksums, and assembly of completed output messages ("reassembly"). The event typically retains some type of SP or event type identification, and the output processors may use this information to determine the proper format for headers, cyclic redundancy checks (CRCs) and other SP bookkeeping information. After a message is reassembled by the output processor, the SPI-4 portion of the output units 242 and 244 formats the message according to the SPI-4 (or other selected) interface protocol, so that the data may be output to the same connections (for example "Host 1" 104 and "Network 1" 106), from which data is received at the input to the SPPS.

### Protocol Processor Core Functions

Once a PPC has received an event of an appropriate type, and has information reflecting the size and location of any payload, it may direct treatment of the entire message in accordance with the SP being used. The PPC may direct actions regarding the flow to which the event belongs, e.g. requesting retransmission, resending previously transmitted messages, and so on, and may update the flowstate for the flow as is appropriate. In some embodiments, traffic congestion can be reduced if the PPCs do not physically transfer messages directly to the output processors (242, 244), but instead simply direct other circuits to transfer the messages for reassembly at the output processors 242, 244.

Some outgoing messages contain very little information (e.g., little or nothing more than a header), such as acknowledgements or requests for retransmission. In these cases, the PPC that is

processing the event (*e.g.*, PPC 216) may form a header based upon the event information and pass it to the socket memory 236. The socket memory 236 may, in turn, do little or nothing to the header information before passing it on to one of the output processors 242, 244. Other outgoing messages will include a substantial payload, which may, for example, have been received with an incoming message and stored in the scratchpad memory 210. The PPC may direct such payloads to be moved from the scratchpad memory 210 to the socket memory 236, and may separately direct one of such payloads to be concatenated, for example in one of the output processors 242, 244, with an appropriate header formed by the PPC. The skilled person in the computer architecture arts will recognize that the PPC can control the output message and flowstate information in many ways.

PPCs may be implemented in any manner consistent with their function. For example, a microprogrammable processor provides one level of flexibility in processing varying communication needs. Some or all PPCs could alternatively be implemented as fixed state machines, in hardware, possibly reducing circuit size and/or increasing processing speed. Yet again, some or all PPCs may comprise embedded microprocessors that are operable out of program memory that can be modified "on the fly," even while the SPPS is active. Such an implementation permits adjusting the number of PPCs able to process particular types of events, adding further load-balancing flexibility. PPCs may be configured to process some stateful protocols, and not others, and the configuration may be fixed or alterable. For example, in a PPC based on a microprogrammable processor, the microprogram (or software) typically determines which event types, or protocols, the PPC is configured to process. A PPC is "compatible" with a particular event type, or protocol, when it is configured to process such event types, or to process messages (events) according to such a protocol.

#### Bus Network and PPC Cluster Control

FIGURE 3 illustrates an exemplary architecture for the bus network and PPC cluster controller 228 of FIGURE 2. In this embodiment, the cluster of PPCs (from 216 – 218) is controlled in part via a cluster bus interface 302. Through the cluster bus interface 302, instructions are available for all of the PPCs (216 – 218) in the cluster from an instruction memory 304, typically implemented using RAM. The cluster bus interface 302 may also provide access to a routing control table 306 for all of the PPCs in the cluster. A cluster DMA controller 308 ("C DMA") may be provided, and may have an egress bus that delivers data from a FIFO of the DMA controller 308 to the cluster bus interface 302, as well as to one side of a dual port memory (*e.g.*, the DPMEM 310, 312) of each of the PPCs 216 – 218 of the cluster. The DPMEM 310, 312 is accessible on the other side from the DMA controller to the corresponding processor with which it is associated as part of a PPC 216, 218. As shown in FIGURE 3, the cluster DMA controller 308 may have a separate ingress bus by which the FIFO receives data from the dual port memory (*e.g.*,

the DPMEM 310, 312) and from the cluster bus interface 302. The DMA controller 308 may be used, for example, to transfer flowstates between the PPC local memory and the flowstate memory 214. As shown in FIGURE 3, the cluster bus controller 302 also provides bidirectional bus connections to a message bus 314, and a further bidirectional bus connection 240b to the socket memory 236. Some or substantially all of the local memory of a PPC may be DPMEM such as the DPMEM 310, but any suitable local memory may be used instead, as may be convenient for design and fabrication.

The bus 240 interconnecting the socket memory 236 and the bus network and PPC cluster control 228 is shown in FIGURE 3 as being implemented by three distinct bidirectional buses: the bus 240a interconnecting the socket memory 236 and the message bus 314; the bus 240b as mentioned above; and the bus 240c to a further cluster bus interface 316. The cluster bus interface 316 operates with respect to the cluster of PPCs 220 – 222 analogously to the cluster bus interface 302, as a crossbar switch to facilitate communication between the PPCs and the message bus 314, the socket memory 236, and to provide access to common instruction memory 318 and a routing table 320. A further cluster DMA 322 similarly manages data flow between the dual port memory of the PPCs 220 – 222 and the cluster bus interface 316. Further sets of similar modules (routing, instruction, cluster bus interface and cluster DMA) may, of course, be provided and similarly interconnected.

The skilled person in the computer architecture arts will appreciate that any suitable bus control may be used to implement the connections shown for the bus network and PPC cluster control 228. For example, the routing and instruction information may be maintained within individual PPCs. In addition, the PPC memory need not be dual-port, nor is a DMA controller such as 308 or 322 necessary. In somewhat less complex embodiments, the cluster bus interfaces 302, 316 may simply be part of the message bus 314, or the interfaces may be omitted entirely. Conversely, even more elaborate bus architectures may be employed to increase the speed and power of some embodiments.

#### Flow Processing with Alternate Protocol Cores

FIGURE 4 is a flowchart showing acts that may be performed by an exemplary SPPS to perform stateful protocol processing of messages belonging to a flow, generally alternating PPCs (protocol processing cores), that is, using different PPCs at different times. As shown in FIGURE 4, at a step 402 a message is received. This step may include various substeps, such as reconstructing complete messages from packet fragments, performing validity checks, and/or establishing checksums. Next, at a step 404, the payload of the message may be moved to a scratchpad memory. The step 404 is optional, insofar as it indicates splitting the message and storing part of the message in a temporary memory location that is especially available to both input



and output processing facilities. Alternatively, for example, the message may be kept together, and/or it may be moved directly to a more permanent memory location.

Proceeding to a step 406, an event portion of the message may be defined. Event definition typically includes the state-relevant portion of the message, and may entail reformatting a header of the message and adding information, such as checksums and event type indication, to facilitate further processing of the event, as discussed in more detail hereinabove. If the message is not split, the "event" may include the payload information, and may even be an incoming message substantially as received. Processing of the event proceeds at a step 408 where data contained within the event that uniquely identifies the flow (the "flow key") is examined to begin a process of determining a location of flowstate information and a local flow ID proxy. A decision step 410 checks whether a PPC is actively processing an event of the same flow. This check may be effected by searching for the flow key in a local "active flow" table. If the flow key is found in the "active flow" table, then a PPC is presently processing another event belonging to the same flow, and the process exits the decision step 410 on the "yes" branch. If the flow is not active (*e.g.*, if the flow key of the flow is not found in the "active flow" table), then processing continues at a decision step 412. Other techniques may be used in the step 410 to determine if events associated with the flow key are presently being processed by any PPC, such as searching an area of memory reserved for the status of message flows that are presently being processed by a PPC (*e.g.*, within a Dispatcher's Core Activity Management submodule). Alternatively, for example, a single flowstate location may be examined for an indication (*e.g.*, a flag) that processing is in progress at a PPC. Further techniques and criteria for determining whether a PPC is actively processing the flow are described below with reference to a decision step 428.

At the decision step 412 a check is made as to whether the flow associated with the flow key is active at the SPPS. This may be performed by checking for a valid flow location in a flow memory that maintains flowstates for active flows when no PPC is presently processing events of the flow. (Since the number of active flows can be very large, the flow memory is typically distinct, separately accessible, and much larger, than the local flow table used for flows presently being processed by a PPC.) This step typically includes a "lookup" task of determining a local flow ID proxy related to the flow key, a task which may involve processing the flow key information according to hash algorithms. Once the local flow ID proxy is determined, it can generally be used to locate an existing flowstate for the flow corresponding to the flow key. The mere existence of a valid flowstate may cause an affirmative result at the decision step 412.

If the flow is not active at all, so that no valid flowstate exists in either general flowstate memory or in a PPC actively processing a flow, then the process proceeds to an initialization step 414 to create and initialize a valid flowstate area within flowstate memory. Note that some stateless

"events" exist that do not require a flowstate, such are Address Resolution Protocol (ARP) events which do not belong to a flow, and for which no flow need be created. ARP, and other such "stateless" events, may be processed independently of the processing steps of FIGURE 4, which are primarily relevant to "stateful" events.

5           Once an active flow is established (whether located at the decision step 412, or initialized at the initialization step 414), the method may proceed to assign a PPC to process the event at an assignment step 416. This step may involve several substeps, such as determining and identifying which PPCs are compatible (*i.e.*, capable of processing events of the present type) and available (*e.g.*, have room in their queues) to process the event. A PPC may be selected from those satisfying  
10 both of these criteria in many ways, such as in a round-robin fashion, or by selecting the least full PPC local queue, or randomly, or by other load balancing algorithms. Because the PPC has just been newly assigned to process the event, the flowstate is made available to the PPC at a step 418. The flowstate may be delivered by the Dispatcher (or submodule) as described above; or, if a global flowstate memory is shared with the assigned PPC, then this step may comprise identifying the  
15 flowstate memory location to the PPC. The step 418 also typically includes identifying the location of a "workspace" where the PPC can access the flowstate during processing. Such workspace is typically maintained locally at the PPC, but in some embodiments may be maintained more globally, or split to be both local and global.

20           Once a PPC has been assigned and has a valid flowstate, which occurs after the step 418 (or after an affirmative step 410), processing proceeds at the steps 420 and 422. Step 420 tracks the activity of a PPC processing a flow. In one embodiment of the present invention, step 420 includes incrementing an event counter associated with the assignment of the PPC to process the flow, but alternatives are described below with regard to the decision step 428.

25           At a step 422 the contents of the event are provided to the assigned PPC. This may be accomplished by physically copying the event contents to a queue in the local memory of the PPC, or, as an alternative example, by identifying a location of the event data to the PPC. Such queue may contain events from different flows, for example from as many different flows as workspace storage is available for corresponding flowstates. If either event queue or flowstate workspace is not available in (or for) a compatible PPC, then the Dispatcher may temporarily withhold effecting  
30 part or all of the event/workspace transfer to the PPC.

35           Once transfer is completed, the assigned PPC has access to the flowstate of the flow, and to the event data, which typically includes information regarding the size and location of the payload associated with the event. At a step 424, the PPC may perform much of the transport layer protocol processing for the message that is associated with the event. The protocol defines the net effect that such processing must achieve, but of course the effect may be accomplished in any manner either

presently practiced or later developed for such transport layer processing. Actions by the PPC may include, as examples, updating the flowstate, creating a header for a message to be output, directing that a previously transmitted message be retransmitted, or sending a request for retransmission of a message that was received with an error. Actions by the PPC may also include directing the  
5 reassembly of a header it constructs to a received payload, and transmission to a different TLTS connected to the network at another end of the flow, or to a local host. Upon completing the event, a done statement is asserted at a step 426. In one embodiment, the done statement is returned to a global dispatcher used to track PPC activity.

### Releasing an Active PPC

10 After the PPC completes processing the present event, a determination is made at a decision step 428 whether the PPC has completed all processing for the flow to which the event belongs. In one embodiment, such determination may be made by a dispatcher module decrementing an event counter associated with a PPC in response to a "done" statement, and determining that the event counter has reached zero. However, many alternatives for establishing that a PPC is done with the  
15 flow will be appropriate in different embodiments. For example, a PPC may be considered "done" with a flow when it completes processing the last event of that flow that exists in its queue. As another example, the PPC may be considered done with a flow when the flowstate in its local memory is overwritten or invalidated by processing in another PPC. These, or other definitions of "done," may be tracked in one (or more) of various places, such as within the PPC itself, or at a  
20 more global module such as a dispatcher (*e.g.*, within a core activity manager submodule).

If, at the decision step 428, the PPC is determined to be actively processing the flow, the method may proceed to a conclusion step 430 with no further processing, since the flowstate local to the PPC has been updated and the global flowstate need not necessarily be updated. However, upon determining that the PPC is done with processing the flow, the local flowstate that has been  
25 updated at the PPC is transferred to a more global flowstate location at a step 432, so that the PPC workspace becomes available for processing events of a different flow. The global flowstate can then be subsequently accessed when further events arrive that belong to the flow. The PPC may be deemed "done" based on event processing completion for the flow as determined by the Dispatcher, by a submodule or other module, or by the PPC itself. The "done" designation may also be  
30 postponed after the processing of all events from the flow is completed, for example until the PPC has no other room for new flows and events. Once the PPC is deemed "done" at a step 434, the PPC may be released from "assignment" to processing the flow, which may, for example, include setting a flag that indicates that the flowstate memory in the PPC is no longer valid, or is available

for further storage of a different flowstate. After the step 434, the PPC will be treated as free of the event, and of the flow to which the event belongs.

A decision step 436 will typically occur at some point to determine whether the last event processed by the PPC permits the flow to be completely closed. This decision step 436 may be  
5 made even before the occurrence of the decision step 428, or before the steps 432 and/or 434, because such a decision to terminate the flow may obviate a need to write the flowstate to memory.

Such a decision may also subsume the decision that the PPC is "done" with the flow. However, for processing convenience, the termination decision may be considered as occurring in the sequence shown in FIGURE 4. The PPC itself will typically determine, as part of its SP processing duties,  
10 whether the last event completed the flow (e.g., whether the flowstate is advanced to the "connection closed" condition). However, a decision to actually close the flow may be made more globally, such as at the dispatcher (or a submodule). If it is determined at the decision step 436 not to terminate the flow, the system is generally done processing the message and proceeds to the done step 430. However, if it is determined at the step 436 to terminate the flow, the local flow ID proxy  
15 and flowstate memory location may thereupon be released for other uses. Since PPCs are generally assigned to, and released from processing events belonging to a flow at the level of a specific flow, largely irrespective of where other flows are assigned (at least within the universe of compatible PPCs), it is possible, indeed highly probable, that a PPC is assigned to process events (or messages) belonging to a flow that was previously processed by another PPC. Such flow-PPC reassignments  
20 may be rather frequent, and under some circumstances may even occur for each event of a flow.

### Dispatcher Processing

FIGURE 5 is a flowchart showing acts that may be taken by a "dispatcher" module within an exemplary SPPS to dispatch events belonging to a flow to different PPCs at different times. FIGURE 5 is focused on acts, which may be generally attributed to a dispatcher module (and  
25 submodules), to effect distribution of incoming events. Thus, FIGURE 5 steps may substantially be a subset of steps of the overall SPPS, such as are illustrated in FIGURE 4, although FIGURE 5 steps are from the dispatcher module perspective and may also include different details than are shown in FIGURE 4. The dispatcher module is conceptually separate from the PPCs to which it dispatches events, and from input processing from which it receives events, and may be connected  
30 within a SPPS like the dispatcher 212 in FIGURE 2, or may be otherwise connected. The dispatcher module may also be conceptually or even physically subdivided; for example, reference is made to a local flow ID proxy (and/or flowstate) "lookup" module, and to a Director Core Activity Manager, each of which may either conceptually or physically be a submodule of the dispatcher module, or an ancillary module associated with the dispatcher.

As shown in FIGURE 5, at a step 502 an event is received from an input source. The event will typically contain only a "state-relevant" part of a message being processed by the SPPS. That is, the event will typically contain only the information necessary for a PPC (protocol processing core) to control the maintenance of the flowstate of the flow associated with the message, and not the payload of the message. However, in some embodiments the payload, or parts of it, may be kept with the state-relevant data. The dispatcher examines a "flow key" contained within the event that uniquely identifies the flow to which the event belongs. At a step 504, the dispatcher searches for a match to the flow key in a Core Activity Manager (or "CAM"), which would indicate that a PPC was actively processing an event related to that flow. If a match is not found in the CAM (which may be physically or conceptually separate from the dispatcher), then in this exemplary embodiment it is presumed that no PPC is actively processing an event of the flow, and at a step 506 a CAM entry will be initialized to track the activity of the PPC assigned to process the event.

At a step 508, the dispatcher searches for a local flow ID proxy that corresponds to the flow key. For SPPSs that handle a large number of flows, this search may be performed by a distinct lookup module which may, for example, perform a hash lookup to locate a local flow ID proxy as quickly as possible. A decision step 510 depends on whether a local flow ID proxy matching the flow key was found. If not, then the SPPS may not yet be processing any data from the flow, and accordingly at a step 512 a flowstate ID may be selected to be associated with the flow that is uniquely identified by the flow key of the event. Thereafter (or if a local flow ID proxy was found and the decision at the step 510 was "yes"), processing may proceed at a step 514.

A PPC is selected to handle the event at the step 514. This step may include a substep of determining the type of event being processed, though in some embodiments this substep is performed by earlier processing modules (e.g., a message splitter such as 206 or 208 of FIGURE 2).

An "event-type mask" maintained in the dispatcher for each PPC may be compared to bits indicating the type of event to determine which PPCs are compatible with the event type. Another substep may include examining the relative activity levels of those PPCs that are configured to handle the event type. The least busy PPC may be selected, or the next PPC that has any room in its input queue may be selected in a round-robin fashion. As a further example, data may be maintained on recent PPC activity (e.g., in a core activity manager submodule) including assignment of local workspaces, and a PPC may be selected that has not yet overwritten its flowstate memory for the flow of the event, even though it is otherwise considered "done" with the flow. A director submodule, either in combination with or as part of a core activity manager (CAM) submodule, may perform these acts. Selection of a PPC (within the universe of compatible PPCs) is generally made for the flow of each incoming event specifically, without regard to an *a priori* class of flows to which the flow might belong (such as by virtue of characteristics of its flow

key). As a result of such individual assignment techniques, the PPC selected to handle a particular event of a flow frequently differs from a PPC that handled previous events of the same flow (unless the particular flow is presently active in a PPC, as explained elsewhere).

Since the flowstate was initialized in the step 512, or was located in the steps 508-510, and the PPC was selected at the step 514, the flowstate may now be transferred to the PPC at a step 516.

In some embodiments such transfer may be "virtual," merely providing an indication of where the flowstate exists in memory so that the PPC can access it. Next, processing can proceed to a step 518. This same step may be reached directly from the decision step 504, since if that decision was "yes" then a PPC is already processing an earlier event belonging to the same flow. Such an active PPC will (in many embodiments) already have the most valid flowstate for the flow, and in that case will generally be selected to process the present event. Therefore, at the step 518, the event itself may be forwarded to an input area or queue of the selected PPC. Along with the step 518, an event counter may be incremented at a step 520. The event counter is one way to determine when a PPC is actively processing another event of the flow of the present event, but other ways may be used, such as waiting for the PPC to indicate that it is done processing all present events of a particular flow. This is the end of the receive processing for the dispatcher.

FIGURE 6 is a flowchart illustrating some steps that the dispatcher (or its submodules) may perform in response to feedback from the PPC. As in FIGURE 5, the steps of FIGURE 6 are largely a subset of steps taken by the overall SPPS, but they are described from the perspective of the dispatcher, and may contain more or different steps than are illustrated in FIGURE 4 for the overall SPPS.

The illustrated response acts of FIGURE 6 start at a step 602 during which the dispatcher receives a "done statement" or other indication that a PPC has completed processing an event of a particular flow. The dispatcher may then decrement the event counter for the flow (discussed with respect to the step 520 of FIGURE 5). If, at a decision step 606, the event counter is found to have reached zero, or if the completion of a "burst" of events by a PPC is otherwise indicated, then the dispatcher may cause the flowstate, as updated by the PPC, to be stored in more permanent memory to free up the memory of the PPC. (Note that this step is not needed for embodiments in which the same memory is used by the PPC during processing as when no PPC is processing the memory, a circumstance that may occur, for example, when the flowstate is always maintained in the same global location, and is merely accessed by a PPC processing the flow, as needed). A flag or other indication may then be included in the CAM, or sent to the PPC, to indicate that the flowstate stored in the PPC is no longer valid. Then at a step 612, the PPC may be released from handling the particular flow it was processing. Since no active PPC is now processing an event of a particular flow, the CAM block in which the PPC activity was maintained can also be released at a step 614.

Note that "release" may amount to merely setting a flag showing that the PPC (or the CAM memory block) is available. Such flag may indicate availability, but a PPC may be treated for all intents and purposes as if it is still actively processing events of a flow after such indication, as long as no essential data has been overwritten. In that case, the decision step 606 would return a "no" until the data blocks are actually overwritten and thus destroyed. In any case, if the decision step 606 returns a "no," then processing is done, since the steps 608 – 614 are generally not needed in that event. Otherwise, processing is done after the CAM block is released at the step 614.

### Encapsulated Stateful Flow Processing

One manner in which a SPPS (stateful protocol processing system) such as described herein may process flows of layers other than transport layers is by extracting the encapsulated messages and recirculating the extracted messages for further processing. Such further processing may be performed in accordance with the appropriate protocol for the encapsulated message, which is typically different from the protocol (typically a TLP) used for the encapsulating messages.

After an encapsulated stateful message is retrieved, reformatted and provided to a SPPS as an input (a non-transport layer input), the SPPS can process the message in accordance with the appropriate protocol as long as one or more of the PPCs are configured with the steps required by the particular protocol (e.g., iSCSI). Thus, it is straightforward to simply use a SPPS for non-transport layer processing.

There are numerous ways in which a SPPS may be notified that encapsulated data requires recirculation. Notification may be implicit, for example, if all processed data requires recirculation.

Alternatively, one or more portions of the header or payload of the encapsulating messages may contain information indicating a need for such recirculation. A SPPS may examine each payload for an indication to recirculate encapsulated information, or it may examine payloads only when an indication is provided in the header. Thus, the SPPS may receive instruction as to whether a payload is to be examined, whether it requires further processing, and by what protocol such further processing should be performed, by any combination of implicit and explicit information in the header and/or payload of the encapsulating message.

A "recirculation" protocol may first be invoked such that the payload (and/or portions of the header) of an encapsulating message is segmented and reassembled as a message for the encapsulated flow. Note that a single encapsulating message may contain all or part of a plurality of encapsulated messages, and that conversely a single encapsulated message may require a plurality of encapsulating messages to be conveyed (for example, when a large message is encapsulated in a plurality of small packets, such as ATM packets). The recirculation protocol defines appropriate reassembly of the encapsulated message, and also directs that it be returned to

the input of the SPPS for further processing. Such a recirculation protocol may format the recirculated message in a particularly efficient format, such as by specifying the local flow ID proxy, the event type, and other useful information as is known. In this manner the SPPS recirculation protocol processor(s) would function similarly to a host operating in close conjunction with the SPPS. Such a host, having knowledge of an ideal format for messages to the SPPS, may speed processing by formatting messages in such ideal format.

It should also be noted that recirculation may be effected by a modified communication path, such that the reassembly or "output processors" 242 and/or 244 transfer the reassembled encapsulated message directly back to a message splitter 206 or 208, rather than passing it through interfaces such as the SPI-4 interfaces in 242, 244, 202 and 204 which may be unnecessary for recirculation. Indeed, the recirculated message may be entirely preformatted in the manner that would otherwise be effected by the message splitters 206 or 208. The selected PPC processing the encapsulating message (or a related processor) may perform such preformatting and direct the information to be delivered directly from the reassembly processors in 242/244 to the scratchpad memory 210 and the dispatcher 212, thus bypassing the message splitters entirely.

Once recirculation has been effected, further processing of the encapsulated information may proceed just as described hereinabove, that is, in substantially the same manner that a TLP message is processed. In the case of interest, the encapsulated information is stateful and belongs to a flow, so an event may be created that reflects the state-relevant portion of the message, a local proxy of the flow key will be determined, a state for the flow will be created or located, and a PPC (protocol processing core) compatible with the protocol will be assigned to process the event derived from the (previously encapsulated, now recirculated) message. These steps may be performed not only for recirculated messages, but for messages of any flow, whether transport layer or not, that is provided to an input of the SPPS.

Processing a non-transport layer message may, of course, require that information be sent to a further subsystem. For example, data within an encapsulated message may require delivery to a host. The assigned PPC may effect such sending by directing that the information be reassembled in a manner acceptable to the target host, and then directing that the reassembled message be transferred to the target host. In an alternative, sending the encapsulated message to a network connection may require that the outgoing message be reencapsulated in a TLP message (typically, but not necessarily, the same TLP, such as TCP, that was used for the original encapsulating message). Thus, further recirculation may be required at this point to reencapsulate such message. In theory, at least, messages may be "nested" in a series of any number of encapsulations that must be stripped off before the innermost stateful message can be processed. Similarly, processing such



innermost stateful message may require symmetrical reencapsulation of a message. In practice, excessive encapsulation will be avoided in the interests of efficiency.

A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the scope of the invention. For example, the methods of the present invention can be executed in software or hardware, or a combination of hardware and software embodiments. As another example, it should be understood that the functions described as being part of one module may in general be performed equivalently in another module. As yet another example, steps or acts shown or described in a particular sequence may generally be performed in a different order, except for those embodiments described in a claim that include a specified order for the steps.

Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiments, but only by the scope of the appended claims. The description may provide examples of similar features as are recited in the claims, but it should not be assumed that such similar features are identical to those in the claims unless such identity is essential to comprehend the scope of the claim. In some instances the intended distinction between claim features and description features is underscored by using slightly different terminology.

## CLAIMS

What is claimed is:

1. A method of processing data in a stateful protocol processing system ("SPPS") that processes a multiplicity of flows of messages, each flow being associated with a uniquely corresponding flow identification ("FID") that is conveyed by messages belonging to such flow, the method comprising:
  - a) receiving a plurality of messages belonging to a particular flow;
  - b) deriving events associated with the particular flow from the plurality of messages;
  - c) specifically assigning a first protocol processing core ("PPC") to process one or more of the events in accordance with a stateful protocol (SP) of the particular flow; and
  - 10 d) specifically assigning a different second PPC to process one or more other of the events in accordance with the SP of the particular flow.
2. The method of Claim 1, wherein assigning PPCs in steps c) and d) further comprises following an algorithm to balance loading between different PPCs.
3. The method of Claim 2, wherein the algorithm to balance loading between different PPCs  
15 includes round-robin distribution among a set of PPCs.
4. The method of Claim 2, further comprising:
  - e) storing some number of events in a local queue for each PPC; and
  - f) identifying a low-load PPC within a relevant group of PPCs that has less events stored in the local queue of the low-load PPC than are stored in the queue of a different PPC within the  
20 relevant group; and wherein
  - g) the algorithm to balance loading between different PPCs includes assigning the low-load PPC to process an event of a flow that has no events presently assigned to a PPC.
5. The method of Claim 1, further comprising:
  - e) assigning events of a combination of different flows concurrently to the first PPC; and
  - 25 f) assigning subsequent events of one of the combination of different flows to the second PPC while the first PPC continues to process events of another of the combination of different flows.

6. The method of Claim 1, further comprising:
- e) assigning events of a combination of different flows concurrently to the first PPC; and
  - f) explicitly releasing the assignment of the first PPC to process events of one of the combination of different flows while the first PPC continues to be assigned to process events of another of the combination of different flows.
7. The method of Claim 1, further comprising:
- e) receiving a message of an unassigned flow for which event processing is presently assigned to no PPC, the unassigned flow having a corresponding FID;
  - f) identifying a plurality of compatible PPCs that are configured to process events of a general type of events of the unassigned flow; and
  - g) selecting a PPC to process one or more events of the unassigned flow from among the compatible PPCs without regard to the FID corresponding to the unassigned flow.
8. The method of Claim 1, further comprising:
- e) receiving a message of an unassigned flow for which event processing is presently assigned to no PPC; and
  - f) selecting a PPC to process one or more events of the unassigned flow from among a plurality of PPCs compatible with such events without regard to FIDs of flows presently queued to be processed by the plurality of compatible PPCs.
9. The method of Claim 1, further comprising explicitly determining whether a particular event belongs to a flow for which a PPC is presently assigned before selecting a PPC to process the particular event.
10. The method of Claim 9, further comprising processing all incoming events that will affect a state of a corresponding flow through a previously selected PPC if the previously selected PPC is determined to be presently assigned to the corresponding flow.
11. The method of Claim 9, further comprising releasing a previous assignment of a PPC to process events of a particular flow without regard to the FID of any flow.
12. The method of Claim 1, further comprising releasing a previous assignment of a PPC to process events of a first flow based on an indication that no events that are relevant to affect a state of the first flow are presently being processed.

13. The method of Claim 12, wherein the indication that no relevant events of the first flow are presently being processed is based on a counter that counts introduction and completion of relevant events of the first flow.

14. The method of Claim 1, further comprising determining a type for an event before selecting a PPC to process the event.

15. The method of Claim 1, further comprising:

e) determining event type prior to selecting a PPC for a number of events derived from received messages belonging to various flows, a multiplicity of distinct event types being determined including at least:

- i) good events of a flow employing a first stateful protocol ("SP");
- ii) good events of a flow employing a different second SP; and
- iii) events of a flow employing the first SP that are determined to have errors.

16. The method of Claim 15, wherein the distinct event types determined further include iv) packet fragments.

17. The method of Claim 1, further comprising recirculating information encapsulated in a first message for stateful protocol processing within the SPPS as a second stateful protocol message.

18. A method of processing data in a data communication stateful protocol processing system that processes a multiplicity of flows of data communication messages, each flow being associated with a uniquely corresponding flow identification ("FID") that is conveyed by messages belonging to such flow, the method comprising:

- a) receiving messages belonging to a particular flow and messages belonging to other flows;
- b) deriving events from the received messages that are associated with the flow indicated by the FID of the message from which they are derived, including events associated with the particular flow and events associated with the other flows;
- c) placing each event in one of a group of one or more preliminary processing queues;
- d) assigning a first protocol processor core ("PPC") to process a first event of the particular flow without regard to the preliminary processing queue in which the first event is located, and subsequently transferring the first event to a local queue of the assigned first PPC; and
- e) assigning a different second PPC to process a different second event of the particular flow without regard to the preliminary processing queue in which the second event is located, and subsequently transferring the second event to a local queue of the assigned second PPC.

19. The method of Claim 18, wherein the first and second PPCs are substantially dedicated to performing stateful protocol message processing.

20. The method of Claim 18, wherein step b) further comprises substantially excluding payload data of the received messages from the derived events.

5 21. The method of Claim 20, wherein step b) further comprises disposing an indication of event type in events that do not include an event type indication.

22. The method of Claim 18, further comprising f) validating integrity of events prior to transferring the events to PPCs.

10 23. The method of Claim 18, further comprising performing step b) in a packet processor that is physically distinct from all stateful protocol PPCs.

24. The method of Claim 23, further comprising f) determining a local proxy ID for an event of a flow that is not presently assigned for processing by a PPC, via acts of a lookup processor that performs program steps that are not performed by the packet processor and are not performed by the PPCs.

15 25. The method of Claim 24, further comprising f) accessing a flowstate at a memory address dependent upon the local proxy ID, the flowstate having been previously stored for the flow that is not presently assigned for processing by a PPC.

20 26. The method of Claim 18, further comprising f) determining a local flow proxy ID, at least for events belonging to flows that are not presently assigned for processing by a particular PPC, via actions of dedicated lookup hardware.

27. The method of Claim 26, wherein the dedicated lookup hardware includes a microprocessor.

28. The method of Claim 18, further comprising:

- f) maintaining data of the first event and associated flowstate data in a first memory uniquely associated with the first PPC; and
- g) concurrently maintaining data of an event of a different flow and a corresponding flowstate of the different flow in a memory that is accessible by a different PPC without regard to access of the first memory by the first PPC.

29. The method of Claim 28, further comprising maintaining both event data and corresponding flowstate data in memory local to a PPC assigned to process such event until processing of such event by the assigned PPC is completed.

30. The method of Claim 18, further comprising assigning PPCs to process events based upon an algorithm for balancing loading between different PPCs.

31. A method of processing data in a stateful protocol processing system ("SPPS") that processes a multiplicity of flows of messages, each flow being associated with a uniquely corresponding flow identification ("FID") that is conveyed by messages belonging to such flow, the method comprising:

- a) steps to effect receipt of a plurality of messages belonging to a particular flow;
- b) steps to define events based upon the plurality of messages;
- c) steps to effect assignment of a first event of a particular flow to a first protocol processor core ("PPC") for stateful protocol ("SP") processing; and
- d) steps to effect assignment of a second event of the particular flow to a different second PPC for SP processing.

32. The method of Claim 31, further comprising steps to determine whether any PPC is presently processing an event of a flow associated with the FID of a new event.

33. The method of Claim 31, further comprising steps to determine when processing has been completed for all present events associated with a particular FID.

34. A system for terminating a data communication transport layer, comprising:

- a) a message receiver module configured to
  - i) receive numerous data communication messages, each such message associated with a flow uniquely identified by a flow identification ("FID") of such message, and
  - ii) derive events from the received messages, each event associated with the flow associated with the message from which it is derived;

b) a plurality of protocol processor core ("PPC") modules configured to process events in accordance with a stateful protocol of the flow of the event; and

c) an event dispatching module for dispatching events among the plurality of PPC modules, the event dispatching module configured to:

5           iv) receive events from the message receiver module;

          v) determine, for each received event, whether a PPC module is presently assigned to process events of the flow of the received event; and

          vi) select, from among a plurality of compatible PPC modules without regard to the FID of the flow of a particular received event, a PPC module to process a received event that  
10           was determined in step ii) to have no PPC module presently assigned for processing.

35. The system of Claim 34, wherein the event dispatching module further comprises a lookup submodule configured to determine local proxy identification for flows of received events determined in step ii) to have no PPC module presently assigned for processing.

36. The system of Claim 34, wherein the event dispatching module further comprises

15           d) an event-tracking submodule configured to

          i) store flow processing status information regarding events dispatched to PPC modules so as to determine whether a PPC module is presently assigned to process events of particular flows.

37. The system of Claim 36, wherein (d) the event-tracking submodule is further configured to

20           ii) retain a count of events according to the flow of the event, and to

          iii) increase the count of a particular flow when an event of the particular flow is dispatched to a PPC module assigned to the flow, and

          iv) reduce the count of the particular flow upon indication that the assigned PPC module has completed processing an event of the particular flow.

25 38. The system of Claim 34, further comprising e) a scratchpad memory configured to store payload data of messages that is separated from events derived from the messages.

39. The system of Claim 34, further comprising e) a message output processing module configured to assemble an output message according to instructions from a PPC module processing a flow of the output message.

30 40. The system of Claim 34, wherein the message receiver module further comprises a processor configured to establish an event type field within events.

41. An apparatus for processing messages according to a data communication stateful protocol, the apparatus comprising:

a) means for receiving messages from a physical layer, said message corresponding to one or more communication flows;

5 b) a plurality of means for stateful protocol (SP) processing of information of a particular communication flow of said one or more communication flows;

c) means for selecting one among the plurality of SP processing means to process information of the particular communication flow so as to balance processor loading irrespective of flow identification ("FID") data provided in messages of the particular communication flow.

10 42. The apparatus of Claim 41, wherein the

c) means for selecting comprises

i) means for assigning a specific SP processing means for processing information of the particular communication flow, and

15 ii) means for releasing the assigned SP processing means from processing information of the particular communication flow without a need to release the assigned SP processing means from processing information of any other communication flow.

43. The apparatus of Claim 42, further comprising d) means for deriving events of corresponding flows from received messages of such corresponding flows.

20 44. The apparatus of Claim 43, wherein the c) means for selecting further comprises means to determine when processing is completed for all events of a flow that have been transferred to an assigned SP processing means.

45. A stateful protocol processing apparatus, comprising:

25 a) more than one protocol processing core ("PPC") microprocessor, each PPC microprocessor including:

i) local memory sufficient to substantially store a flowstate that is accessible by the PPC without interference to other PPCs, and

30 ii) program memory configured to contain instruction code, during operation of the apparatus, sufficient to cause the microprocessor to process a received data message in accordance with an appropriate stateful protocol (SP), including updating an associated flowstate stored in local memory in response to state-relevant information in the received data message; and



b) a dispatcher including an input for receiving data messages, and processing circuitry that is configured, at least during operation of the apparatus, to

iv) forward information of a first received data message associated with a particular flow to a first PPC for SP processing, and

v) forward information of at least a different second message associated with the particular flow to a different second PPC for SP processing.

46. The apparatus of Claim 45, wherein each PPC further comprises dual-port local memory.

47. A method within a data communication system for distributing communication events belonging to corresponding communication flows among a plurality of compatible protocol processors, the method comprising:

a) receiving a first event belonging to a first flow in a queue for distribution;

b) selecting a first protocol processor to process the first event from among a plurality of protocol processors that are compatible with the first event;

c) transferring the first event from the queue to the first protocol processor;

d) transferring other events belonging to other flows from the queue to the first protocol processor;

e) receiving a second event belonging to the particular flow in the queue;

f) selecting a different second protocol processor to process the second event from among the plurality of compatible protocol processors; and

g) transferring the second event from the queue to the second protocol processor while the first protocol processor continues to process other events from the queue.

48. The method of Claim 47, further comprising releasing the first protocol processor from processing events of the first flow in response to completion by the first protocol processor of all events of the first flow that have been transferred to the first protocol processor.

49. The method of Claim 47, further comprising:

h) receiving, after step c) and before the first protocol processor has completed processing the first event, a third event of the particular flow; and

i) consequently selecting the first protocol processor to process the third event of the particular flow.

FIGURE 1A



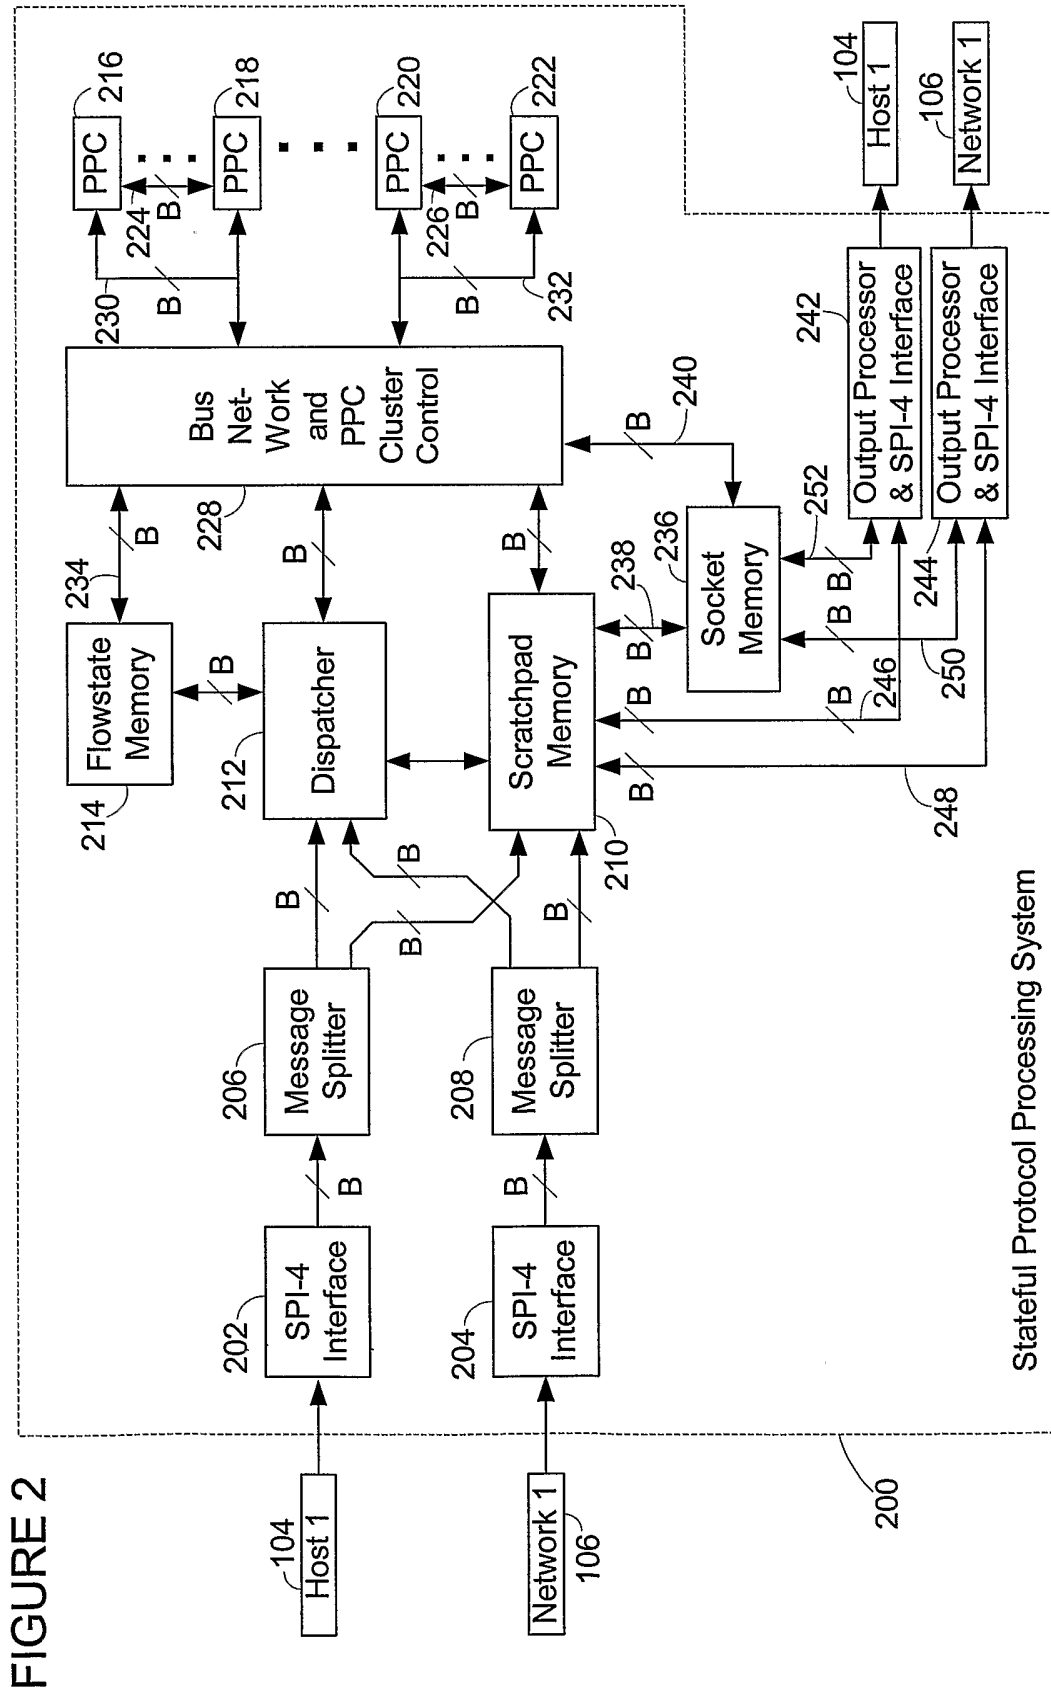


FIGURE 1B





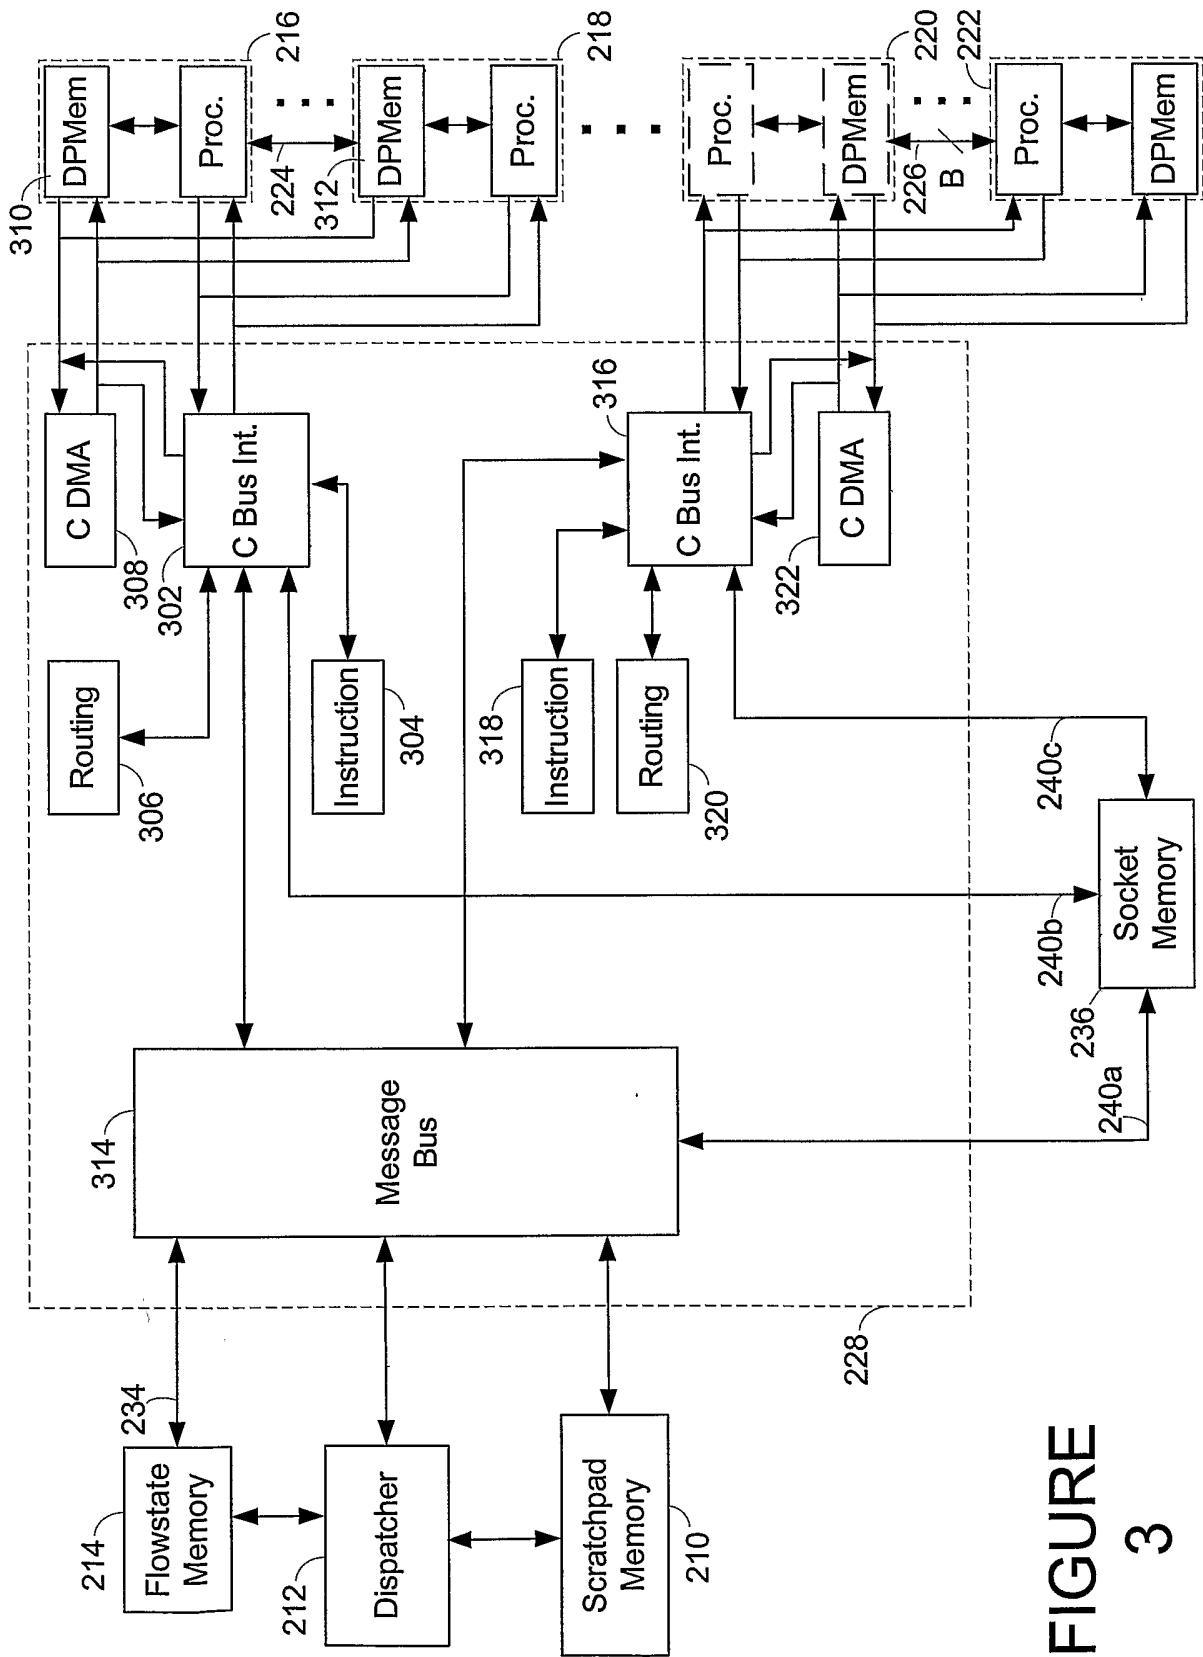


FIGURE 3

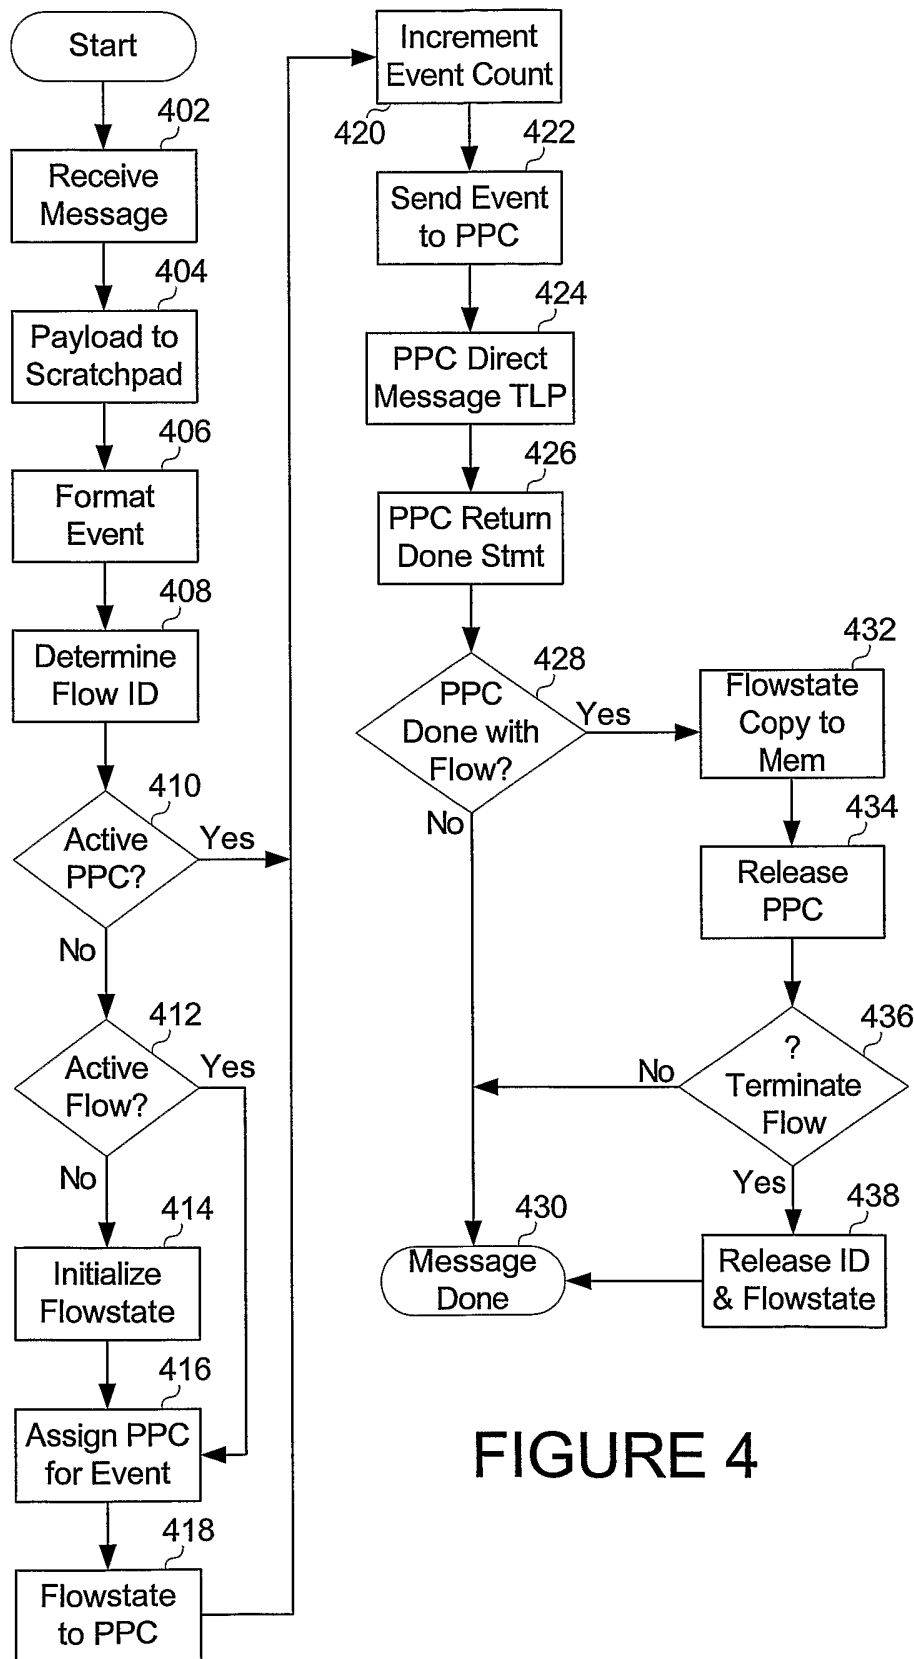


FIGURE 4

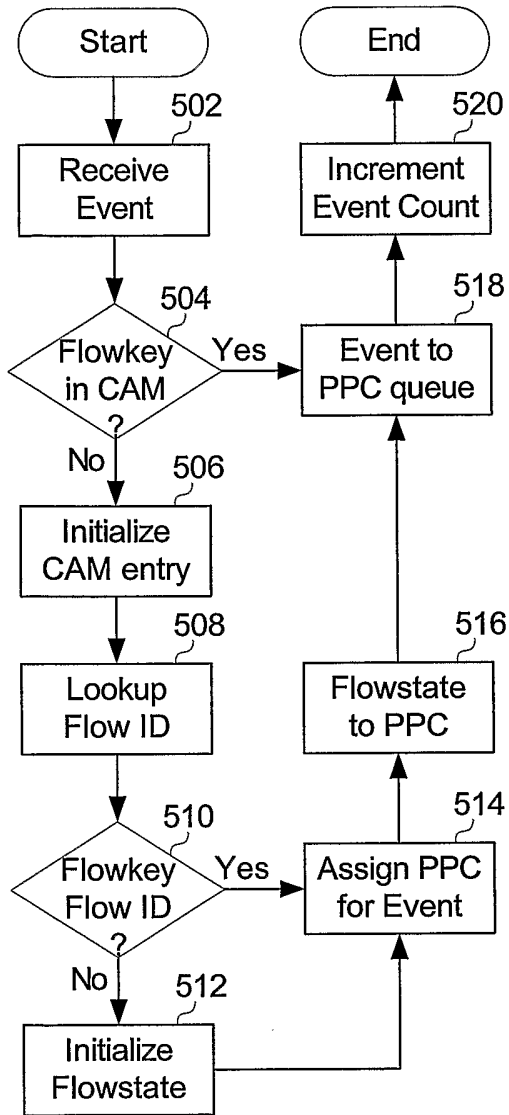


FIGURE 5

FIGURE 6

