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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
15 November 2007 (15.11.2007)

PCT

(10) International Publication Number
WO 2007/130415 A2

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

(21) International Application Number:
PCT/US2007/010559

(22) International Filing Date: 1 May 2007 (01.05.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/416,057 2 May 2006 (02.05.2006) US

(71) Applicant *(for all designated States except US)*: **HARRIS CORPORATION** [US/US]; 1025 W. Nasa Blvd., MS A-11i, Melbourne, FL 32919 (US).

(72) Inventors: **SMITH, Donald, L.**; 584 Hawksbill Island Drive, Satellite Beach, FL 32937 (US). **GALLUSCIO, Anthony, P.**; 549 Hummingbird Drive, Indialantic, FL 32903 (US). **KNAZIK, Robert, J.**; 112 St. Croix Avenue, Cocoa Beach, FL 32931 (US).

(74) Agents: **YATSKO, Michael, S.** et al.; Harris Corporation, 1025 W. Nasa Blvd., MS A-11I, Melbourne, FL 32919 (US).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEMS AND METHODS FOR PROTOCOL FILTERING FOR QUALITY OF SERVICE

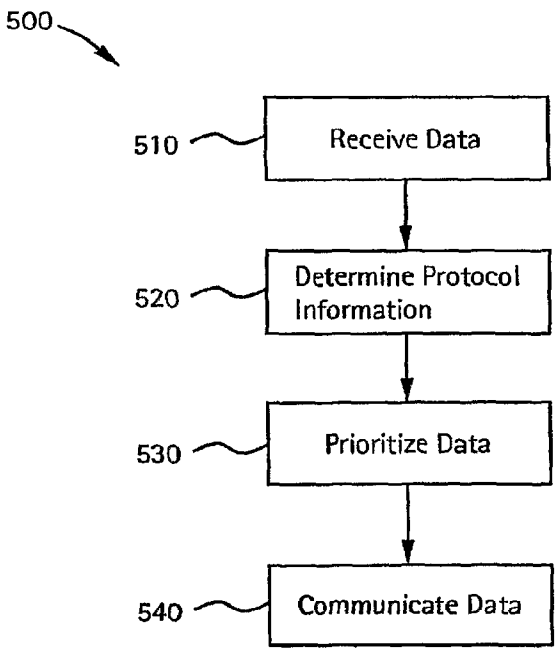

(57) Abstract: Certain embodiments of the present invention provide for a method for data communication including determining protocol information for a block of data, prioritizing the block of data, and communicating the block of data. The protocol information may include information used by a protocol to communicate the block of data. The prioritization may be based at least in part on the protocol information. The communication may be based at least in part on the prioritization of the block of data.

WO 2007/130415 A2

SYSTEMS AND METHODS FOR PROTOCOL FILTERING FOR QUALITY OF SERVICE

The presently described technology generally relates to communications networks. More particularly, the presently described technology relates to systems and methods for protocol filtering for Quality of Service.

Communications networks are utilized in a variety of environments. Communications networks typically include two or more nodes connected by one or more links. Generally, a communications network is used to support communication between two or more participant nodes over the links and intermediate nodes in the communications network. There may be many kinds of nodes in the network. For example, a network may include nodes such as clients, servers, workstations, switches, and/or routers. Links may be, for example, modem connections over phone lines, wires, Ethernet links, Asynchronous Transfer Mode (ATM) circuits, satellite links, and/or fiber optic cables.

A communications network may actually be composed of one or more smaller communications networks. For example, the Internet is often described as network of interconnected computer networks. Each network may utilize a different architecture and/or topology. For example, one network may be a switched Ethernet network with a star topology and another network may be a Fiber-Distributed Data Interface (FDDI) ring.

Communications networks may carry a wide variety of data. For example, a network may carry bulk file transfers alongside data for interactive real-time conversations. The data sent on a network is often sent in packets, cells, or frames. Alternatively, data may be sent as a stream. In some instances, a stream or flow of data may actually be a sequence of packets. Networks such as the Internet provide general

purpose data paths between a range of nodes and carrying a vast array of data with different requirements.

Communication over a network typically involves multiple levels of communication protocols. A protocol stack, also referred to as a networking stack or protocol suite, refers to a collection of protocols used for communication. Each protocol may be focused on a particular type of capability or form of communication. For example, one protocol may be concerned with the electrical signals needed to communicate with devices connected by a copper wire. Other protocols may address ordering and reliable transmission between two nodes separated by many intermediate nodes, for example.

Protocols in a protocol stack typically exist in a hierarchy. Often, protocols are classified into layers. One reference model for protocol layers is the Open Systems Interconnection (OSI) model. The OSI reference model includes seven layers: a physical layer, data link layer, network layer, transport layer, session layer, presentation layer, and application layer. The physical layer is the "lowest" layer, while the application layer is the "highest" layer. Two well-known transport layer protocols are the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). A well known network layer protocol is the Internet Protocol (IP).

At the transmitting node, data to be transmitted is passed down the layers of the protocol stack, from highest to lowest. Conversely, at the receiving node, the data is passed up the layers, from lowest to highest. At each layer, the data may be manipulated by the protocol handling communication at that layer. For example, a transport layer protocol may add a header to the data that allows for ordering of packets upon arrival at a destination node. Depending on the application, some layers may not be used, or even present, and data may just be passed through.

One kind of communications network is a tactical data network. A tactical data network may also be referred to as a tactical communications network. A tactical data network may be utilized by units within an organization such as a military (e.g., army, navy, and/or air force). Nodes within a tactical data network may include, for example, individual soldiers, aircraft, command units, satellites, and/or radios. A tactical data network may be used for communicating data such as voice, position telemetry, sensor data, and/or real-time video.

An example of how a tactical data network may be employed is as follows. A logistics convoy may be in-route to provide supplies for a combat unit in the field. Both the convoy and the combat unit may be providing position telemetry to a command post over satellite radio links. An unmanned aerial vehicle (UAV) may be patrolling along the road the convoy is taking and transmitting real-time video data to the command post over a satellite radio link also. At the command post, an analyst may be examining the video data while a controller is tasking the UAV to provide video for a specific section of road. The analyst may then spot an improvised explosive device (IED) that the convoy is approaching and send out an order over a direct radio link to the convoy for it to halt and alerting the convoy to the presence of the IED.

The various networks that may exist within a tactical data network may have many different architectures and characteristics. For example, a network in a command unit may include a gigabit Ethernet local area network (LAN) along with radio links to satellites and field units that operate with much lower throughput and higher latency. Field units may communicate both via satellite and via direct path radio frequency (RF). Data may be sent point-to-point, multicast, or broadcast, depending on the nature of the data and/or the specific physical characteristics of the network. A network

may include radios, for example, set up to relay data. In addition, a network may include a high frequency (HF) network which allows long rang communication. A microwave network may also be used, for example. Due to the diversity of the types of links and nodes, among other reasons, tactical networks often have overly complex network addressing schemes and routing tables. In addition, some networks, such as radio-based networks, may operate using bursts. That is, rather than continuously transmitting data, they send periodic bursts of data. This is useful because the radios are broadcasting on a particular channel that must be shared by all participants, and only one radio may transmit at a time.

Tactical data networks are generally bandwidth-constrained. That is, there is typically more data to be communicated than bandwidth available at any given point in time. These constraints may be due to either the demand for bandwidth exceeding the supply, and/or the available communications technology not supplying enough bandwidth to meet the user's needs, for example. For example, between some nodes, bandwidth may be on the order of kilobits/sec. In bandwidth-constrained tactical data networks, less important data can clog the network, preventing more important data from getting through in a timely fashion, or even arriving at a receiving node at all. In addition, portions of the networks may include internal buffering to compensate for unreliable links. This may cause additional delays. Further, when the buffers get full, data may be dropped.

In many instances the bandwidth available to a network cannot be increased. For example, the bandwidth available over a satellite communications link may be fixed and cannot effectively be increased without deploying another satellite. In these situations, bandwidth must be managed rather than simply expanded to handle demand. In large systems, network bandwidth is a critical resource. It is

desirable for applications to utilize bandwidth as efficiently as possible. In addition, it is desirable that applications avoid "clogging the pipe," that is, overwhelming links with data, when bandwidth is limited. When bandwidth allocation changes, applications should preferably react. Bandwidth can change dynamically due to, for example, quality of service, jamming, signal obstruction, priority reallocation, and line-of-sight. Networks can be highly volatile and available bandwidth can change dramatically and without notice.

In addition to bandwidth constraints, tactical data networks may experience high latency. For example, a network involving communication over a satellite link may incur latency on the order of half a second or more. For some communications this may not be a problem, but for others, such as real-time, interactive communication (e.g., voice communications), it is highly desirable to minimize latency as much as possible.

Another characteristic common to many tactical data networks is data loss. Data may be lost due to a variety of reasons. For example, a node with data to send may be damaged or destroyed. As another example, a destination node may temporarily drop off of the network. This may occur because, for example, the node has moved out of range, the communication's link is obstructed, and/or the node is being jammed. Data may be lost because the destination node is not able to receive it and intermediate nodes lack sufficient capacity to buffer the data until the destination node becomes available. Additionally, intermediate nodes may not buffer the data at all, instead leaving it to the sending node to determine if the data ever actually arrived at the destination.

Often, applications in a tactical data network are unaware of and/or do not account for the particular characteristics of the network. For example, an application

may simply assume it has as much bandwidth available to it as it needs. As another example, an application may assume that data will not be lost in the network. Applications which do not take into consideration the specific characteristics of the underlying communications network may behave in ways that actually exacerbate problems. For example, an application may continuously send a stream of data that could just as effectively be sent less frequently in larger bundles. The continuous stream may incur much greater overhead in, for example, a broadcast radio network that effectively starves other nodes from communicating, whereas less frequent bursts would allow the shared bandwidth to be used more effectively.

Certain protocols do not work well over tactical data networks. For example, a protocol such as TCP may not function well over a radio-based tactical network because of the high loss rates and latency such a network may encounter. TCP requires several forms of handshaking and acknowledgments to occur in order to send data. High latency and loss may result in TCP hitting time outs and not being able to send much, if any, meaningful data over such a network.

Information communicated with a tactical data network often has various levels of priority with respect to other data in the network. For example, threat warning receivers in an aircraft may have higher priority than position telemetry information for troops on the ground miles away. As another example, orders from headquarters regarding engagement may have higher priority than logistical communications behind friendly lines. The priority level may depend on the particular situation of the sender and/or receiver. For example, position telemetry data may be of much higher priority when a unit is actively engaged in combat as compared to when the unit is merely following a standard patrol route. Similarly, real-time video data from an UAV may

have higher priority when it is over the target area as opposed to when it is merely in-route.

There are several approaches to delivering data over a network. One approach, used by many communications networks, is a "best effort" approach. That is, data being communicated will be handled as well as the network can, given other demands, with regard to capacity, latency, reliability, ordering, and errors. Thus, the network provides no guarantees that any given piece of data will reach its destination in a timely manner, or at all. Additionally, no guarantees are made that data will arrive in the order sent or even without transmission errors changing one or more bits in the data.

Another approach is Quality of Service (QoS). QoS refers to one or more capabilities of a network to provide various forms of guarantees with regard to data that is carried. For example, a network supporting QoS may guarantee a certain amount of bandwidth to a data stream. As another example, a network may guarantee that packets between two particular nodes have some maximum latency. Such a guarantee may be useful in the case of a voice communication where the two nodes are two people having a conversation over the network. Delays in data delivery in such a case may result in irritating gaps in communication and/or dead silence, for example.

QoS may be viewed as the capability of a network to provide better service to selected network traffic. The primary goal of QoS is to provide priority including dedicated bandwidth, controlled jitter and latency (required by some real-time and interactive traffic), and improved loss characteristics. Another important goal is making sure that providing priority for one flow does not make other flows fail. That is, guarantees made for subsequent flows must not break the guarantees made to existing flows.

Current approaches to QoS often require every node in a network to support QoS, or, at the very least, for every node in the network involved in a particular communication to support QoS. For example, in current systems, in order to provide a latency guarantee between two nodes, every node carrying the traffic between those two nodes must be aware of and agree to honor, and be capable of honoring, the guarantee.

There are several approaches to providing QoS. One approach is Integrated Services, or "IntServ." IntServ provides a QoS system wherein every node in the network supports the services and those services are reserved when a connection is set up. IntServ does not scale well because of the large amount of state information that must be maintained at every node and the overhead associated with setting up such connections.

Another approach to providing QoS is Differentiated Services, or "DiffServ." DiffServ is a class of service model that enhances the best-effort services of a network such as the Internet. DiffServ differentiates traffic by user, service requirements, and other criteria. Then, DiffServ marks packets so that network nodes can provide different levels of service via priority queuing or bandwidth allocation, or by choosing dedicated routes for specific traffic flows. Typically, a node has a variety of queues for each class of service. The node then selects the next packet to send from those queues based on the class categories.

Existing QoS solutions are often network specific and each network type or architecture may require a different QoS configuration. Due to the mechanisms existing QoS solutions utilize, messages that look the same to current QoS systems may actually have different priorities based on message content. However, data consumers may require access to high-priority data without being flooded by lower-priority

data. Existing QoS systems cannot provide QoS based on message content at the transport layer.

As mentioned, existing QoS solutions require at least the nodes involved in a particular communication to support QoS. However, the nodes at the "edge" of network may be adapted to provide some improvement in QoS, even if they are incapable of making total guarantees. Nodes are considered to be at the edge of the network if they are the participating nodes in a communication (i.e., the transmitting and/or receiving nodes) and/or if they are located at chokepoints in the network. A chokepoint is a section of the network where all traffic must pass to another portion. For example, a router or gateway from a LAN to a satellite link would be a choke point, since all traffic from the LAN to any nodes not on the LAN must pass through the gateway to the satellite link.

Thus, there is a need for systems and methods providing QoS in a tactical data network. There is a need for systems and methods for providing QoS on the edge of a tactical data network. Additionally, there is a need for adaptive, configurable QoS systems and methods in a tactical data network.

Embodiments of the present invention provide systems and methods for facilitating communication of data. A method includes determining protocol information for a block of data, prioritizing the block of data, and communicating the block of data. The protocol information may include information used by a protocol to communicate the block of data. The prioritization may be based at least in part on the protocol information. The communication may be based at least in part on the prioritization of the block of data.

Certain embodiments provide a system for data communication including a prioritization component and a communication component. The prioritization component may be

adapted to determine a priority for a block of data. The block of data may include protocol information. The priority may be determined based at least in part on the protocol information. The communication component may be adapted to communicate the block of data based at least in part on the priority of the block of data.

Certain embodiments provide a computer-readable medium including a set of instructions for execution on a computer, the set of instructions including a protocol information determination routine, a prioritization routine, and a communication routine. The protocol information determination routine may be configured to determine protocol information for a block of data. The prioritization routine may be configured to determine a priority for the block of data based at least in part on the protocol information. The communication routine may be configured to communicate the block of data.

Fig. 1 illustrates a tactical communications network environment operating with an embodiment of the present invention.

Fig. 2 shows the positioning of the data communications system in the seven layer OSI network model in accordance with an embodiment of the present invention.

Fig. 3 depicts an example of multiple networks facilitated using the data communications system in accordance with an embodiment of the present invention.

Fig. 4 illustrates a data communication environment operating with an embodiment of the present invention.

Fig. 5 illustrates a flow diagram for a method for communicating data in accordance with an embodiment of the present invention.

The foregoing summary, as well as the following detailed description of certain embodiments of the present invention, will be better understood when read in conjunction

with the appended drawings. For the purpose of illustrating the invention, certain embodiments are shown in the drawings. It should be understood, however, that the present invention is not limited to the arrangements and instrumentality shown in the attached drawings.

Fig. 1 illustrates a tactical communications network environment 100 operating with an embodiment of the present invention. The network environment 100 includes a plurality of communication nodes 110, one or more networks 120, one or more links 130 connecting the nodes and network(s), and one or more communication systems 150 facilitating communication over the components of the network environment 100. The following discussion assumes a network environment 100 including more than one network 120 and more than one link 130, but it should be understood that other environments are possible and anticipated.

Communication nodes 110 may be and/or include radios, transmitters, satellites, receivers, workstations, servers, and/or other computing or processing devices, for example.

Network(s) 120 may be hardware and/or software for transmitting data between nodes 110, for example. Network(s) 120 may include one or more nodes 110, for example.

Link(s) 130 may be wired and/or wireless connections to allow transmissions between nodes 110 and/or network(s) 120.

The communications system 150 may include software, firmware, and/or hardware used to facilitate data transmission among the nodes 110, networks 120, and links 130, for example. As illustrated in Fig. 1, communications system 150 may be implemented with respect to the nodes 110, network(s) 120, and/or links 130. In certain embodiments, every node 110 includes a communications system 150. In certain embodiments, one or more nodes 110 include a communications system 150. In

certain embodiments, one or more nodes 110 may not include a communications system 150.

The communication system 150 provides dynamic management of data to help assure communications on a tactical communications network, such as the network environment 100. As shown in Fig. 2, in certain embodiments, the system 150 operates as part of and/or at the top of the transport layer in the OSI seven layer protocol model. The system 150 may give precedence to higher priority data in the tactical network passed to the transport layer, for example. The system 150 may be used to facilitate communications in a single network, such as a local area network (LAN) or wide area network (WAN), or across multiple networks. An example of a multiple network system is shown in Fig. 3. The system 150 may be used to manage available bandwidth rather than add additional bandwidth to the network, for example.

In certain embodiments, the system 150 is a software system, although the system 150 may include both hardware and software components in various embodiments. The system 150 may be network hardware independent, for example. That is, the system 150 may be adapted to function on a variety of hardware and software platforms. In certain embodiments, the system 150 operates on the edge of the network rather than on nodes in the interior of the network. However, the system 150 may operate in the interior of the network as well, such as at "choke points" in the network.

The system 150 may use rules and modes or profiles to perform throughput management functions such as optimizing available bandwidth, setting information priority, and managing data links in the network. By "optimizing" bandwidth, it is meant that the presently described technology can be employed to increase an efficiency of bandwidth use to communicate data in one or more networks. Optimizing bandwidth usage may include removing functionally redundant

messages, message stream management or sequencing, and message compression, for example. Setting information priority may include differentiating message types at a finer granularity than Internet Protocol (IP) based techniques and sequencing messages onto a data stream via a selected rule-based sequencing algorithm, for example. Data link management may include rule-based analysis of network measurements to affect changes in rules, modes, and/or data transports, for example. A mode or profile may include a set of rules related to the operational needs for a particular network state of health or condition. The system 150 provides dynamic, "on-the-fly" reconfiguration of modes, including defining and switching to new modes on the fly.

The communication system 150 may be configured to accommodate changing priorities and grades of service, for example, in a volatile, bandwidth-limited network. The system 150 may be configured to manage information for improved data flow to help increase response capabilities in the network and reduce communications latency. Additionally, the system 150 may provide interoperability via a flexible architecture that is upgradeable and scalable to improve availability, survivability, and reliability of communications. The system 150 supports a data communications architecture that may be autonomously adaptable to dynamically changing environments while using predefined and predictable system resources and bandwidth, for example.

In certain embodiments, the system 150 provides throughput management to bandwidth-constrained tactical communications networks while remaining transparent to applications using the network. The system 150 provides throughput management across multiple users and environments at reduced complexity to the network. As mentioned above, in certain embodiments, the system 150 runs on a host node in and/or at the top of layer four (the transport layer) of the

OSI seven layer model and does not require specialized network hardware. The system 150 may operate transparently to the layer four interface. That is, an application may utilize a standard interface for the transport layer and be unaware of the operation of the system 150. For example, when an application opens a socket, the system 150 may filter data at this point in the protocol stack. The system 150 achieves transparency by allowing applications to use, for example, the TCP/IP socket interface that is provided by an operating system at a communication device on the network rather than an interface specific to the system 150. System 150 rules may be written in extensible markup language (XML) and/or provided via custom dynamic link libraries (DLLs), for example.

In certain embodiments, the system 150 provides quality of service (QoS) on the edge of the network. The system's QoS capability offers content-based, rule-based data prioritization on the edge of the network, for example. Prioritization may include differentiation and/or sequencing, for example. The system 150 may differentiate messages into queues based on user-configurable differentiation rules, for example. The messages are sequenced into a data stream in an order dictated by the user-configured sequencing rule (e.g., starvation, round robin, relative frequency, etc.). Using QoS on the edge, data messages that are indistinguishable by traditional QoS approaches may be differentiated based on message content, for example. Rules may be implemented in XML, for example. In certain embodiments, to accommodate capabilities beyond XML and/or to support extremely low latency requirements, the system 150 allows dynamic link libraries to be provided with custom code, for example.

Inbound and/or outbound data on the network may be customized via the system 150. Prioritization protects client applications from high-volume, low-priority data, for example.

The system 150 helps to ensure that applications receive data to support a particular operational scenario or constraint.

In certain embodiments, when a host is connected to a LAN that includes a router as an interface to a bandwidth-constrained tactical network, the system may operate in a configuration known as QoS by proxy. In this configuration, packets that are bound for the local LAN bypass the system and immediately go to the LAN. The system applies QoS on the edge of the network to packets bound for the bandwidth-constrained tactical link.

In certain embodiments, the system 150 offers dynamic support for multiple operational scenarios and/or network environments via commanded profile switching. A profile may include a name or other identifier that allows the user or system to change to the named profile. A profile may also include one or more identifiers, such as a functional redundancy rule identifier, a differentiation rule identifier, an archival interface identifier, a sequencing rule identifier, a pre-transmit interface identifier, a post-transmit interface identifier, a transport identifier, and/or other identifier, for example. A functional redundancy rule identifier specifies a rule that detects functional redundancy, such as from stale data or substantially similar data, for example. A differentiation rule identifier specifies a rule that differentiates messages into queues for processing, for example. An archival interface identifier specifies an interface to an archival system, for example. A sequencing rule identifier identifies a sequencing algorithm that controls samples of queue fronts and, therefore, the sequencing of the data on the data stream. A pre-transmit interface identifier specifies the interface for pre-transmit processing, which provides for special processing such as encryption and compression, for example. A post-transmit interface identifier identifies an interface for post-transmit

processing, which provides for processing such as de-encryption and decompression, for example. A transport identifier specifies a network interface for the selected transport.

A profile may also include other information, such as queue sizing information, for example. Queue sizing information identifiers a number of queues and amount of memory and secondary storage dedicated to each queue, for example.

In certain embodiments, the system 150 provides a rules-based approach for optimizing bandwidth. For example, the system 150 may employ queue selection rules to differentiate messages into message queues so that messages may be assigned a priority and an appropriate relative frequency on the data stream. The system 150 may use functional redundancy rules to manage functionally redundant messages. A message is functionally redundant if it is not different enough (as defined by the rule) from a previous message that has not yet been sent on the network, for example. That is, if a new message is provided that is not sufficiently different from an older message that has already been scheduled to be sent, but has not yet been sent, the newer message may be dropped, since the older message will carry functionally equivalent information and is further ahead in the queue. In addition, functional redundancy many include actual duplicate messages and newer messages that arrive before an older message has been sent. For example, a node may receive identical copies of a particular message due to characteristics of the underlying network, such as a message that was sent by two different paths for fault tolerance reasons. As another example, a new message may contain data that supersedes an older message that has not yet been sent. In this situation, the system 150 may drop the older message and send only the new message. The system 150 may also

include priority sequencing rules to determine a priority-based message sequence of the data stream. Additionally, the system 150 may include transmission processing rules to provide pre-transmission and post-transmission special processing, such as compression and/or encryption.

In certain embodiments, the system 150 provides fault tolerance capability to help protect data integrity and reliability. For example, the system 150 may use user-defined queue selection rules to differentiate messages into queues. The queues are sized according to a user-defined configuration, for example. The configuration specifies a maximum amount of memory a queue may consume, for example. Additionally, the configuration may allow the user to specify a location and amount of secondary storage that may be used for queue overflow. After the memory in the queues is filled, messages may be queued in secondary storage. When the secondary storage is also full, the system 150 may remove the oldest message in the queue, logs an error message, and queues the newest message. If archiving is enabled for the operational mode, then the de-queued message may be archived with an indicator that the message was not sent on the network.

Memory and secondary storage for queues in the system 150 may be configured on a per-link basis for a specific application, for example. A longer time between periods of network availability may correspond to more memory and secondary storage to support network outages. The system 150 may be integrated with network modeling and simulation applications, for example, to help identify sizing to help ensure that queues are sized appropriately and time between outages is sufficient to help achieve steady-state and help avoid eventual queue overflow.

Furthermore, in certain embodiments, the system 150 offers the capability to meter inbound ("shaping") and

outbound ("policing") data. Policing and shaping capabilities help address mismatches in timing in the network. Shaping helps to prevent network buffers form flooding with high-priority data queued up behind lower-priority data. Policing helps to prevent application data consumers from being overrun by low-priority data. Policing and shaping are governed by two parameters: effective link speed and link proportion. The system 150 may form a data stream that is no more than the effective link speed multiplied by the link proportion, for example. The parameters may be modified dynamically as the network changes. The system may also provide access to detected link speed to support application level decisions on data metering. Information provided by the system 150 may be combined with other network operations information to help decide what link speed is appropriate for a given network scenario.

Fig. 4 illustrates a data communication environment 400 operating with an embodiment of the present invention. The environment 400 includes a data communication system 410, one or more source nodes 420, and one or more destination nodes 430. The data communication system 410 is in communication with the source node(s) 420 and the destination node(s) 430. The data communication system 410 may communicate with the source node(s) 420 and/or destination node(s) 430 over links, such as radio, satellite, network links, and/or through inter-process communication, for example. In certain embodiments, the data communication system 410 may communication with one or more source nodes 420 and/or destination nodes 430 over one or more tactical data networks.

The data communication system 410 may be similar to the communication system 150, described above, for example. In certain embodiments, the data communication system 410 is adapted to receive data from the one or more source nodes 420.

In certain embodiments, the data communication system 410 may include one or more queues for storing, organizing, and/or prioritizing the data. Alternatively, other data structures may be used for storing, organizing, and/or prioritizing the data. For example, a table, tree, or linked list may be used. In certain embodiments, the data communication system 410 is adapted to communicate data to the one or more destination nodes 430.

The data received, stored, prioritized, processed, communicated, and/or transmitted by data communication system 410 may include a block of data. The block of data may be, for example, a packet, cell, frame, and/or stream. For example, the data communication system 410 may receive packets of data from a source node 420. As another example, the data communication system 410 may process a stream of data from a source node 420.

In certain embodiments, the data includes protocol information. The protocol information may be used by one or more protocols to communicate the data, for example. The protocol information may include, for example, a source address, a destination address, a source port, a destination port, and/or a protocol type. The source and/or destination address may be an IP address, for example, of a source node 420 and/or a destination node 430. The protocol type may include the kind of protocol used for one or more layers of communication of the data. For example, the protocol type may be a transport protocol such as Transmission Control Protocol (TCP), User Datagram Protocol (UDP), or Stream Control Transmission Protocol (SCTP). As another example, the protocol type may include Internet Protocol (IP), Internetwork Packet Exchange (IPX), Ethernet, Asynchronous Transfer Mode (ATM), File Transfer Protocol (FTP), and/or Real-time Transport Protocol (RTP).

In certain embodiments, the data includes a header and a payload. The header may include some or all of the protocol information, for example. In certain embodiments, some or all of the protocol information is included in the payload. For example, protocol information may include information regarding a higher-level protocol stored in the payload portion of a block of data. In certain embodiments, the data is not contiguous in memory. That is, one or more portions of the data may be located in different regions of memory. For example, protocol information may be stored in one region of memory while the payload is stored in another buffer.

Source node(s) 420 provide and/or generate, at least in part, data handled by the data communication system 410. A source node 420 may include, for example, an application, radio, satellite, or network. The source node 420 may communicate with the data communication system 410 over a link, as discussed above. Source node(s) 420 may generate a continuous stream of data or may burst data, for example. In certain embodiments, the source node 420 and the data communication system 410 are part of the same system. For example, the source node 420 may be an application running on the same computer system as the data communication system 410.

Destination node(s) 430 receive data handled by the data communication system 410. A destination node 430 may include, for example, an application, radio, satellite, or network. The destination node 430 may communicate with the data communication system 410 over a link, as discussed above. In certain embodiments, the destination node 430 and the data communication system 410 are part of the same system. For example, the destination node 430 may be an application running on the same computer system as the data communication system 410.

The data communication system 410 may communicate with one or more source nodes 420 and/or destination nodes 430 over links, as discussed above. In certain embodiments, the one or more links may be part of a tactical data network. In certain embodiments, one or more links may be bandwidth constrained. In certain embodiments, one or more links may be unreliable and/or intermittently disconnected.

In operation, data is provided and/or generated by one or more data sources 420. The data is received at the data communication system 410. The data may be received over one or more links, for example. For example, data may be received at the data communication system 410 from a radio over a tactical data network. As another example, data may be provided to the data communication system 410 by an application running on the same system by an inter-process communication mechanism. As discussed above, the data may be a block of data, for example.

Data is received by the data communication system 410. In certain embodiments, the data communication system 410 may not receive all of the data. For example, some of the data may be stored in a buffer and the data communication system 410 may receive only header information and a pointer to the buffer. For example, the data communication system 410 may be hooked into the protocol stack of an operating system and when an application passes data to the operating system through a transport layer interface (e.g., sockets), the operating system may then provide access to the data to the data communication system 410.

In certain embodiments, the data communication system 410 may organize and/or prioritize the data. In certain embodiments, the data communication system 410 may determine a priority for a block of data. For example, when a block of data is received by the data communication system 410, a prioritization component of the data communication

system 410 may determine a priority for that block of data. As another example, a block of data may be stored in a queue in the data communication system 410 and a prioritization component may extract the block of data from the queue based on a priority determined for the block of data and/or for the queue.

The priority of the block of data may be based at least in part on protocol information associated and/or included in the block of data. In certain embodiments, the data communication system 410 may determine protocol information for the data. The protocol information may be similar to the protocol information described above, for example. For example, the data communication system 410 may determine a priority for a block of data based on the source address of the block of data. As another example, the data communication system 410 may determine a priority for a block of data based on the transport protocol used to communicate the block of data.

The prioritization of the data by the data communication system 410 may be used to provide QoS, for example. For example, the data communication system 410 may determine a priority for a data received over a tactical data network. The priority may be based on the source address of the data, for example. For example, a source IP address for the data from a radio of a member of the same platoon as the platoon the data communication system 410 belongs to may be given a higher priority than data originating from a unit in a different division in a different area of operations. The priority may be used to determine which of a plurality of queues the data should be placed into for subsequent communication by the data communication system 410. For example, higher priority data may be placed in a queue intended to hold higher priority data, and in turn, the data

communication system 410, in determining what data to next communicate may look first to the higher priority queue.

The data may be prioritized based at least in part on one or more rules. As discussed above, the rules may be user defined. In certain embodiments, rules may be written in XML and/or provided via custom DLLs, for example. A rule may specify, for example, that data received using one protocol be favored over data utilizing another protocol. For example, command data may utilize a particular protocol that is given priority, via a rule, over position telemetry data sent using another protocol. As another example, a rule may specify that position telemetry data coming from a first range of addresses may be given priority over position telemetry data coming from a second range of addresses. The first range of addresses may represent IP addresses of other aircraft in the same squadron as the aircraft with the data communication system 410 running on it, for example. The second range of addresses may then represent, for example, IP addresses for other aircraft that are in a different area of operations, and therefore of less interest to the aircraft on which the data communication system 410 is running.

In certain embodiments, the data communication system 410 does not drop data. That is, although data may be low priority, it is not dropped by the data communication system 410. Rather, the data may be delayed for a period of time, potentially dependent on the amount of higher priority data that is received.

In certain embodiments, the data communication system 410 includes a mode or profile indicator. The mode indicator may represent the current mode or state of the data communication system 410, for example. As discussed above, the data communications system 410 may use rules and modes or profiles to perform throughput management functions such as optimizing available bandwidth, setting information priority,

and managing data links in the network. The different modes may affecting changes in rules, modes, and/or data transports, for example. A mode or profile may include a set of rules related to the operational needs for a particular network state of health or condition. The data communication system 410 may provide dynamic reconfiguration of modes, including defining and switching to new modes "on-the-fly," for example.

In certain embodiments, the data communication system 410 is transparent to other applications. For example, the processing, organizing, and/or prioritization performed by the data communication system 410 may be transparent to one or more source nodes 420 or other applications or data sources. For example, an application running on the same system as data communication system 410, or on a source node 420 connected to the data communication system 410, may be unaware of the prioritization of data performed by the data communication system 410.

Data is communicated from the data communication system 410. In certain embodiments, a communication component is used to communicate the data. The data may be communicated to one or more destination nodes 430, for example. The data may be communicated over one or more links, for example. For example, the data may be communicated by the data communication system 410 over a tactical data network to a radio. As another example, data may be provided by the data communication system 410 to an application running on the same system by an inter-process communication mechanism.

In one embodiment, for example, a bandwidth-constrained network, such as a tactical data network, includes one or more source nodes and one or more destination nodes. The nodes may be aircraft radios, satellites, and/or software applications, for example. The data from the source node(s) is communicated to the data communication system. The data communication system may be on the same node as a source node,

a destination node, or on an intermediate node. For example, the data communication system may be on a fighter aircraft, with source nodes such as an application on the aircraft, other aircraft in the squadron, a headquarters unit, and a ground unit. The data may be communicated over a link such as a satellite link, a radio link, and/or inter-process communication. The data from the source node(s) includes protocol information such as source address, destination address, source port, and destination port. The data communication system determines a priority for the data received from the source node(s) based on the priority information. For example, data from one source node may receive a higher priority than data from another source node because the first source node is in a headquarters unit and the second node is in a ground unit hundreds of miles away. The data communication system then communicates the data to the destination node(s) based on the priority. For example, higher priority data, such as orders from the headquarters unit, may be communicated to an application on the aircraft so it can be displayed to the pilot ahead of lower priority data, such as position telemetry from a ground unit hundreds of miles away.

As discussed above, the components, elements, and/or functionality of the data communication system 410 may be implemented alone or in combination in various forms in hardware, firmware, and/or as a set of instructions in software, for example. Certain embodiments may be provided as a set of instructions residing on a computer-readable medium, such as a memory, hard disk, DVD, or CD, for execution on a general purpose computer or other processing device.

Fig. 5 illustrates a flow diagram for a method 500 for communicating data in accordance with an embodiment of the present invention. The method 500 includes the following steps, which will be described below in more detail. At step

510, data is received. At step 520, protocol information is determined. At step 530, data is prioritized. At step 540, data is communicated. The method 500 is described with reference to elements of systems described above, but it should be understood that other implementations are possible.

At step 510, data is received. Data may be received at the data communication system 410, for example. The data may be received over one or more links, for example. The data may be provided and/or generated by one or more data sources 420, for example. For example, data may be received at the data communication system 410 from a radio over a tactical data network. As another example, data may be provided to the data communication system 410 by an application running on the same system by an inter-process communication mechanism. As discussed above, the data may be a block of data, for example.

In certain embodiments, the data communication system 410 may not receive all of the data. For example, some of the data may be stored in a buffer and the data communication system 410 may receive only header information and a pointer to the buffer. For example, the data communication system 410 may be hooked into the protocol stack of an operating system, and, when an application passes data to the operating system through a transport layer interface (e.g., sockets), the operating system may then provide access to the data to the data communication system 410.

At step 520, protocol information is determined. The protocol information may be determined based at least in part on data that has been received. For example, the protocol information may be determined at least in part based on the data receive at step 510. The protocol information may be used by one or more protocols to communicate the data, for example. The protocol information may include, for example, a source address, a destination address, a source port, a destination port, and/or a protocol type. The source and/or

destination address may be an IP address, for example, of a source node 420 and/or a destination node 430. The protocol type may include the kind of protocol used for one or more layers of communication of the data.

At step 530, data is prioritized. The data may be prioritized and/or organized by data communication system 410, for example. The data to be prioritized may be the data that is received at step 510, for example. In certain embodiments, the data communication system 410 may determine a priority for a block of data. For example, when a block of data is received by the data communication system 410, a prioritization component of the data communication system 410 may determine a priority for that block of data. As another example, a block of data may be stored in a queue in the data communication system 410 and a prioritization component may extract the block of data from the queue based on a priority determined for the block of data and/or for the queue. The priority of the block of data may be based at least in part on protocol information associated and/or included in the block of data. The protocol information may be similar to the protocol information described above, for example. For example, the data communication system 410 may determine a priority for a block of data based on the source address of the block of data. As another example, the data communication system 410 may determine a priority for a block of data based on the transport protocol used to communicate the block of data.

The prioritization of the data may be used to provide QoS, for example. For example, the data communication system 410 may determine a priority for a data received over a tactical data network. The priority may be based on the source address of the data, for example. For example, a source IP address for the data from a radio of a member of the same platoon as the platoon the data communication system 410

belongs to may be given a higher priority than data originating from a unit in a different division in a different area of operations. The priority may be used to determine which of a plurality of queues the data should be placed into for subsequent communication by the data communication system 410. For example, higher priority data may be placed in a queue intended to hold higher priority data, and in turn, the data communication system 410, in determining what data to next communicate, may look first to the higher priority queue.

The data may be prioritized based at least in part on one or more rules. As discussed above, the rules may be user defined and/or programmed based on system and/or operational constraints, for example. In certain embodiments, rules may be written in XML and/or provided via custom DLLs, for example. A rule may specify, for example, that data received using one protocol be favored over data utilizing another protocol. For example, command data may utilize a particular protocol that is given priority, via a rule, over position telemetry data sent using another protocol As another example, a rule may specify that position telemetry data coming from a first range of addresses may be given priority over position telemetry data coming from a second range of addresses. The first range of addresses may represent IP addresses of other aircraft in the same squadron as the aircraft with the data communication system 410 running on it, for example. The second range of addresses may then represent, for example, IP addresses for other aircraft that are in a different area of operations, and therefore of less interest to the aircraft on which the data communication system 410 is running.

In certain embodiments, the data to be prioritized is not dropped. That is, although data may be low priority, it is not dropped by the data communication system 410. Rather, the data may be delayed for a period of time,

potentially dependent on the amount of higher priority data that is received.

In certain embodiments, a mode or profile indicator may represent the current mode or state of the data communication system 410, for example. As discussed above, the rules and modes or profiles may be used to perform throughput management functions such as optimizing available bandwidth, setting information priority, and managing data links in the network. The different modes may affecting changes in rules, modes, and/or data transports, for example. A mode or profile may include a set of rules related to the operational needs for a particular network state of health or condition. The data communication system 410 may provide dynamic reconfiguration of modes, including defining and switching to new modes "on-the-fly," for example.

In certain embodiments, the prioritization of data is transparent to other applications. For example, the processing, organizing, and/or prioritization performed by the data communication system 410 may be transparent to one or more source nodes 420 or other applications or data sources. For example, an application running on the same system as data communication system 410, or on a source node 420 connected to the data communication system 410, may be unaware of the prioritization of data performed by the data communication system 410.

At step 540, data is communicated. The data communicated may be the data received at step 510, for example. The data communicated may be the data prioritized at step 530, for example. Data may be communicated from the data communication system 410, for example. The data may be communicated to one or more destination nodes 430, for example. The data may be communicated over one or more links, for example. For example, the data may be communicated by the data communication system 410 over a tactical data network to

a radio. As another example, data may be provided by the data communication system 410 to an application running on the same system by an inter-process communication mechanism.

One or more of the steps of the method 500 may be implemented alone or in combination in hardware, firmware, and/or as a set of instructions in software, for example. Certain embodiments may be provided as a set of instructions residing on a computer-readable medium, such as a memory, hard disk, DVD, or CD, for execution on a general purpose computer or other processing device.

Certain embodiments of the present invention may omit one or more of these steps and/or perform the steps in a different order than the order listed. For example, some steps may not be performed in certain embodiments of the present invention. As a further example, certain steps may be performed in a different temporal order, including simultaneously, than listed above.

Thus, certain embodiments of the present invention provide systems and methods for protocol filtering for QoS. Certain embodiments provide a technical effect of protocol filtering for QoS.

CLAIMS

1. A method for data communication, the method including:

determining protocol information for at least first and second blocks of data, wherein the protocol information includes information used by a protocol to communicate the blocks of data;

prioritizing at least the first and second blocks of data by enqueuing the blocks of data in at least one queue based at least in part on the protocol information; and

communicating the blocks of data based at least in part on the prioritizations of the blocks of data, wherein the order in which the blocks of data are communicated is based on the prioritizing step.

2. The method of claim 1, wherein the prioritizing step includes prioritizing at least the first and second blocks of data based at least in part on a mode indicator associated with one or more of the blocks of data.

3. The method of claim 1, wherein the communicating step includes transmitting the blocks of data at least in part over a tactical data network.

4. The method of claim 1, wherein the order is determined based at least in part on a range of addresses.

5. The method of claim 1, wherein the order is determined based at least in part on a user defined rule.

6. The method of claim 5, wherein the rule is based at least in part on a mode.

7. The method of claim 1, wherein the prioritizing step is transparent to an application program.

8. A system for data communication, the system including:

a prioritization component adapted to determine a first priority for a first block of data and a second priority for a second block of data in at least one queue, the first and second blocks of data each including protocol information, wherein the first and second priorities are determined based at least in part on the protocol information; and

a communication component adapted to communicate the first block of data before at least the second block of data based at least in part on the priority of the first block of data.

9. The system of claim 8, wherein the blocks of data are received at least in part over a tactical data network.

10. The system of claim 8, further including a mode indicator, wherein the mode indicator indicates a current mode, wherein the prioritization component is adapted to prioritize the blocks of data based at least in part on the mode indicator.

FIG. 1

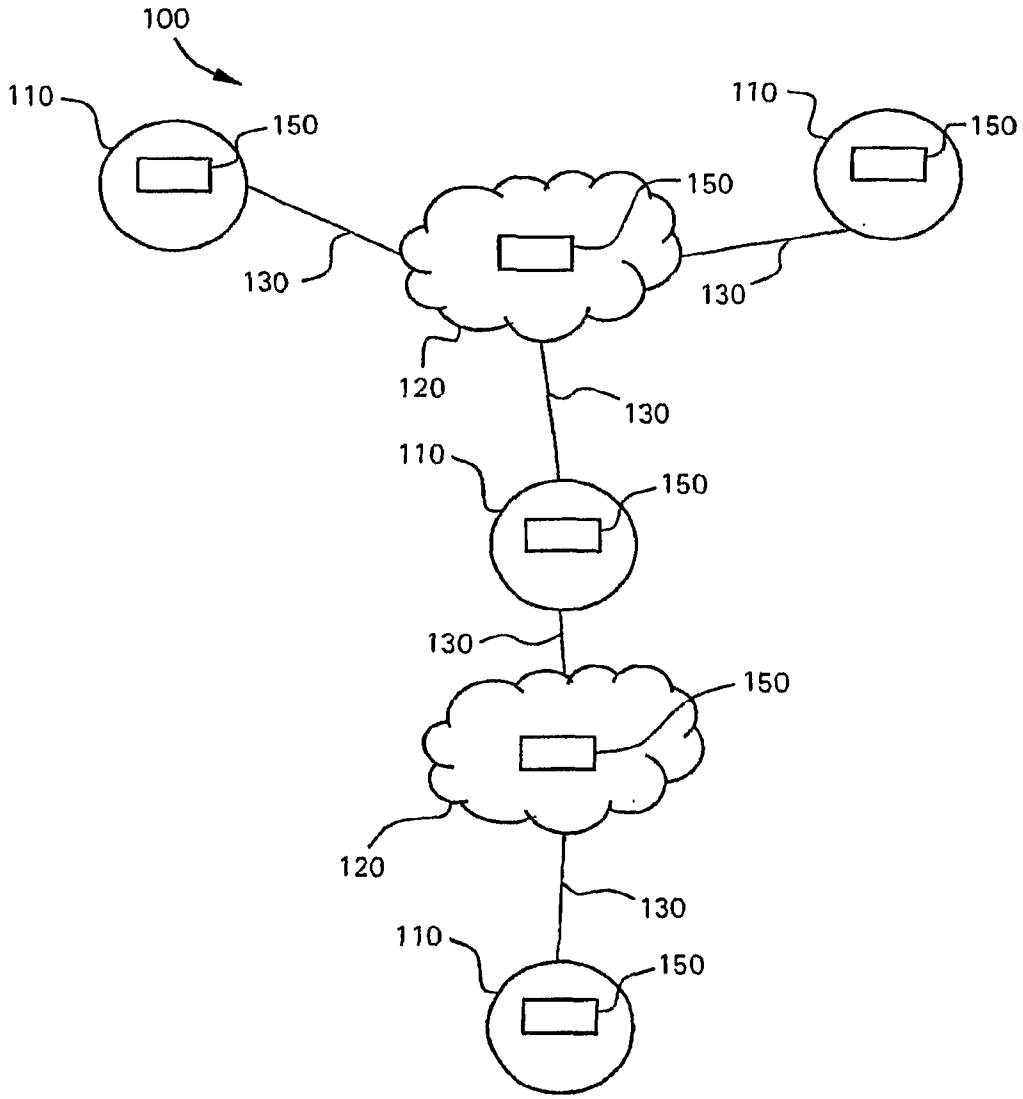
100
110
150
110
150
150
130
130
120
130
110
150
130
150
120
130
110
150

FIG. 2

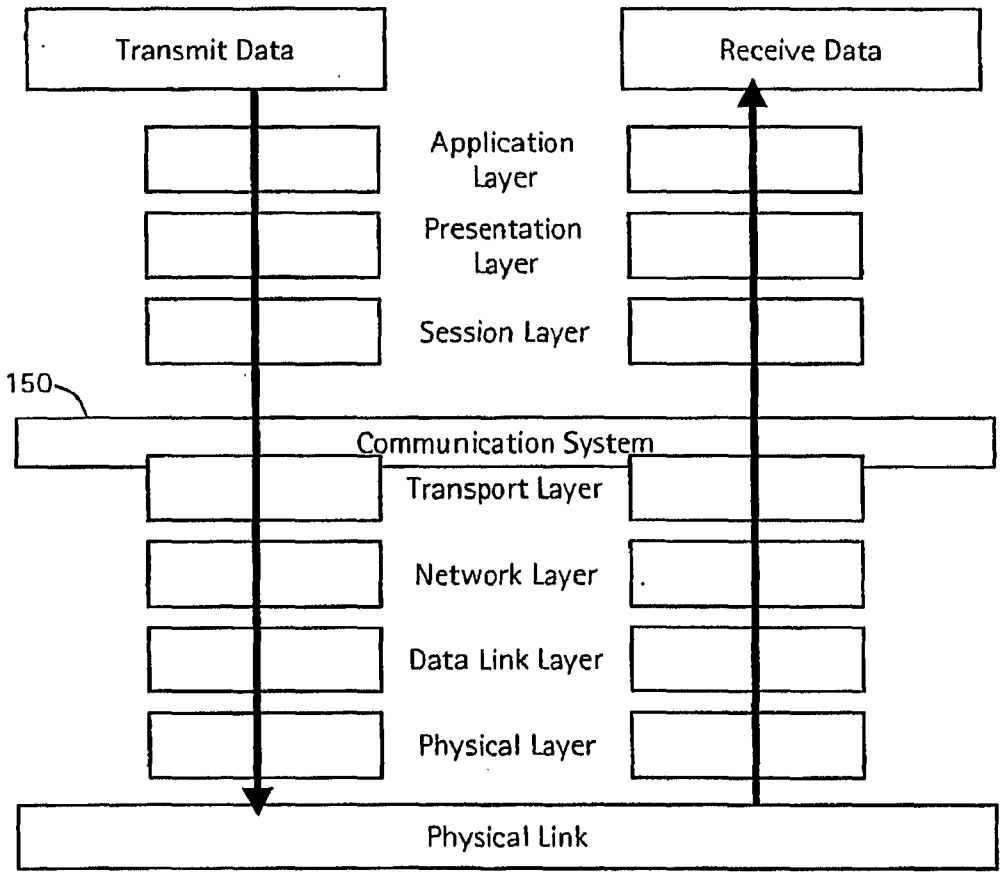
Transmit Data
Receive Data
Application Layer
Presentation Layer
Session Layer
150
Communication System
Transport Layer
Network Layer
Data Link Layer
Physical Layer
Physical Link

FIG. 3

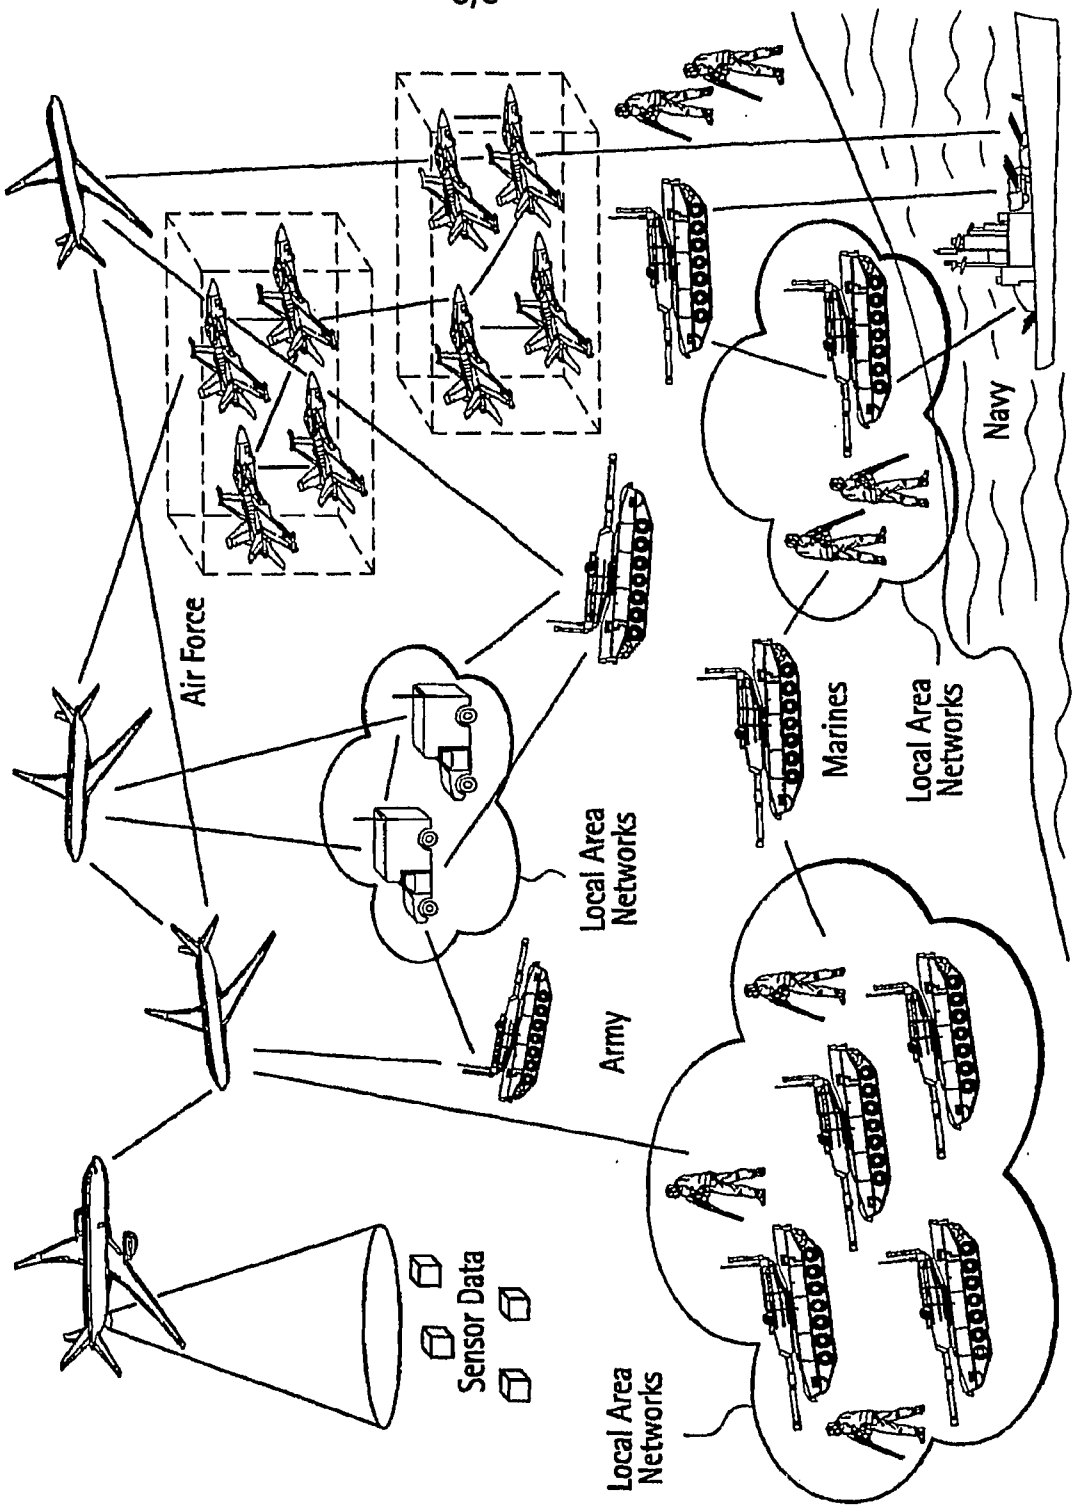
Air Force
Local Area Networks
Army
Marines
Local Area Networks
Navy
Sensor Data
Local Area Networks

FIG. 4

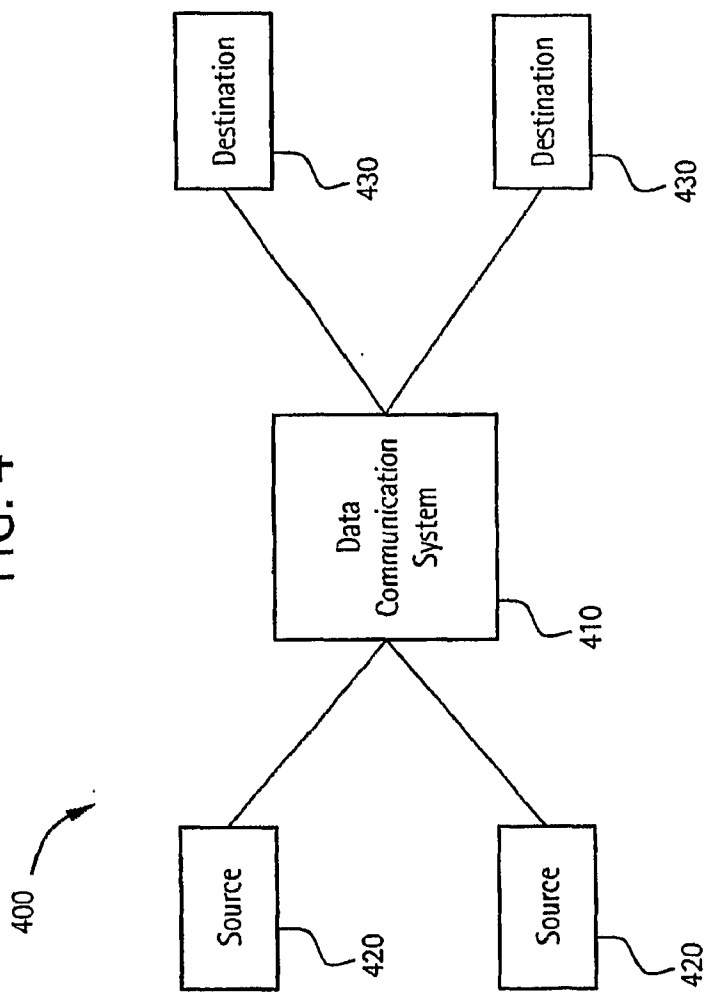
Destination
430
Destination
430
Data Communication System
410
Source
420
Source
420
400

FIG. 5

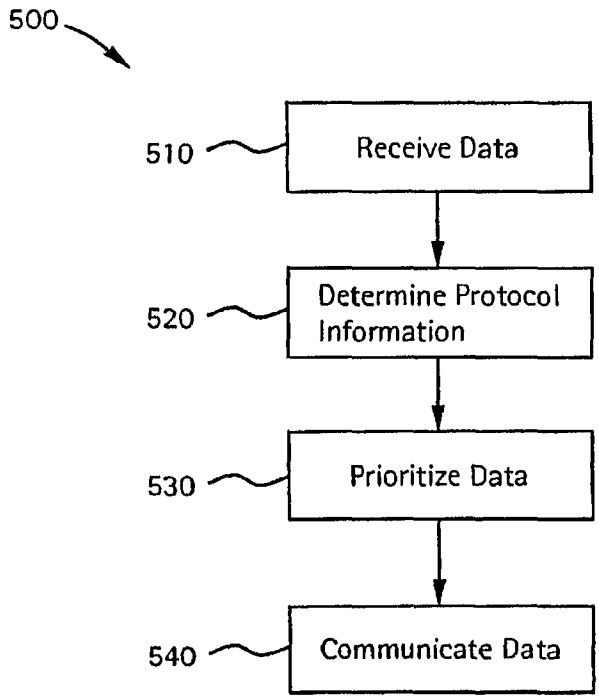
500
510
Receive Data
520
Determine Protocol Information
530
Prioritize Data
540
Communicate Data