

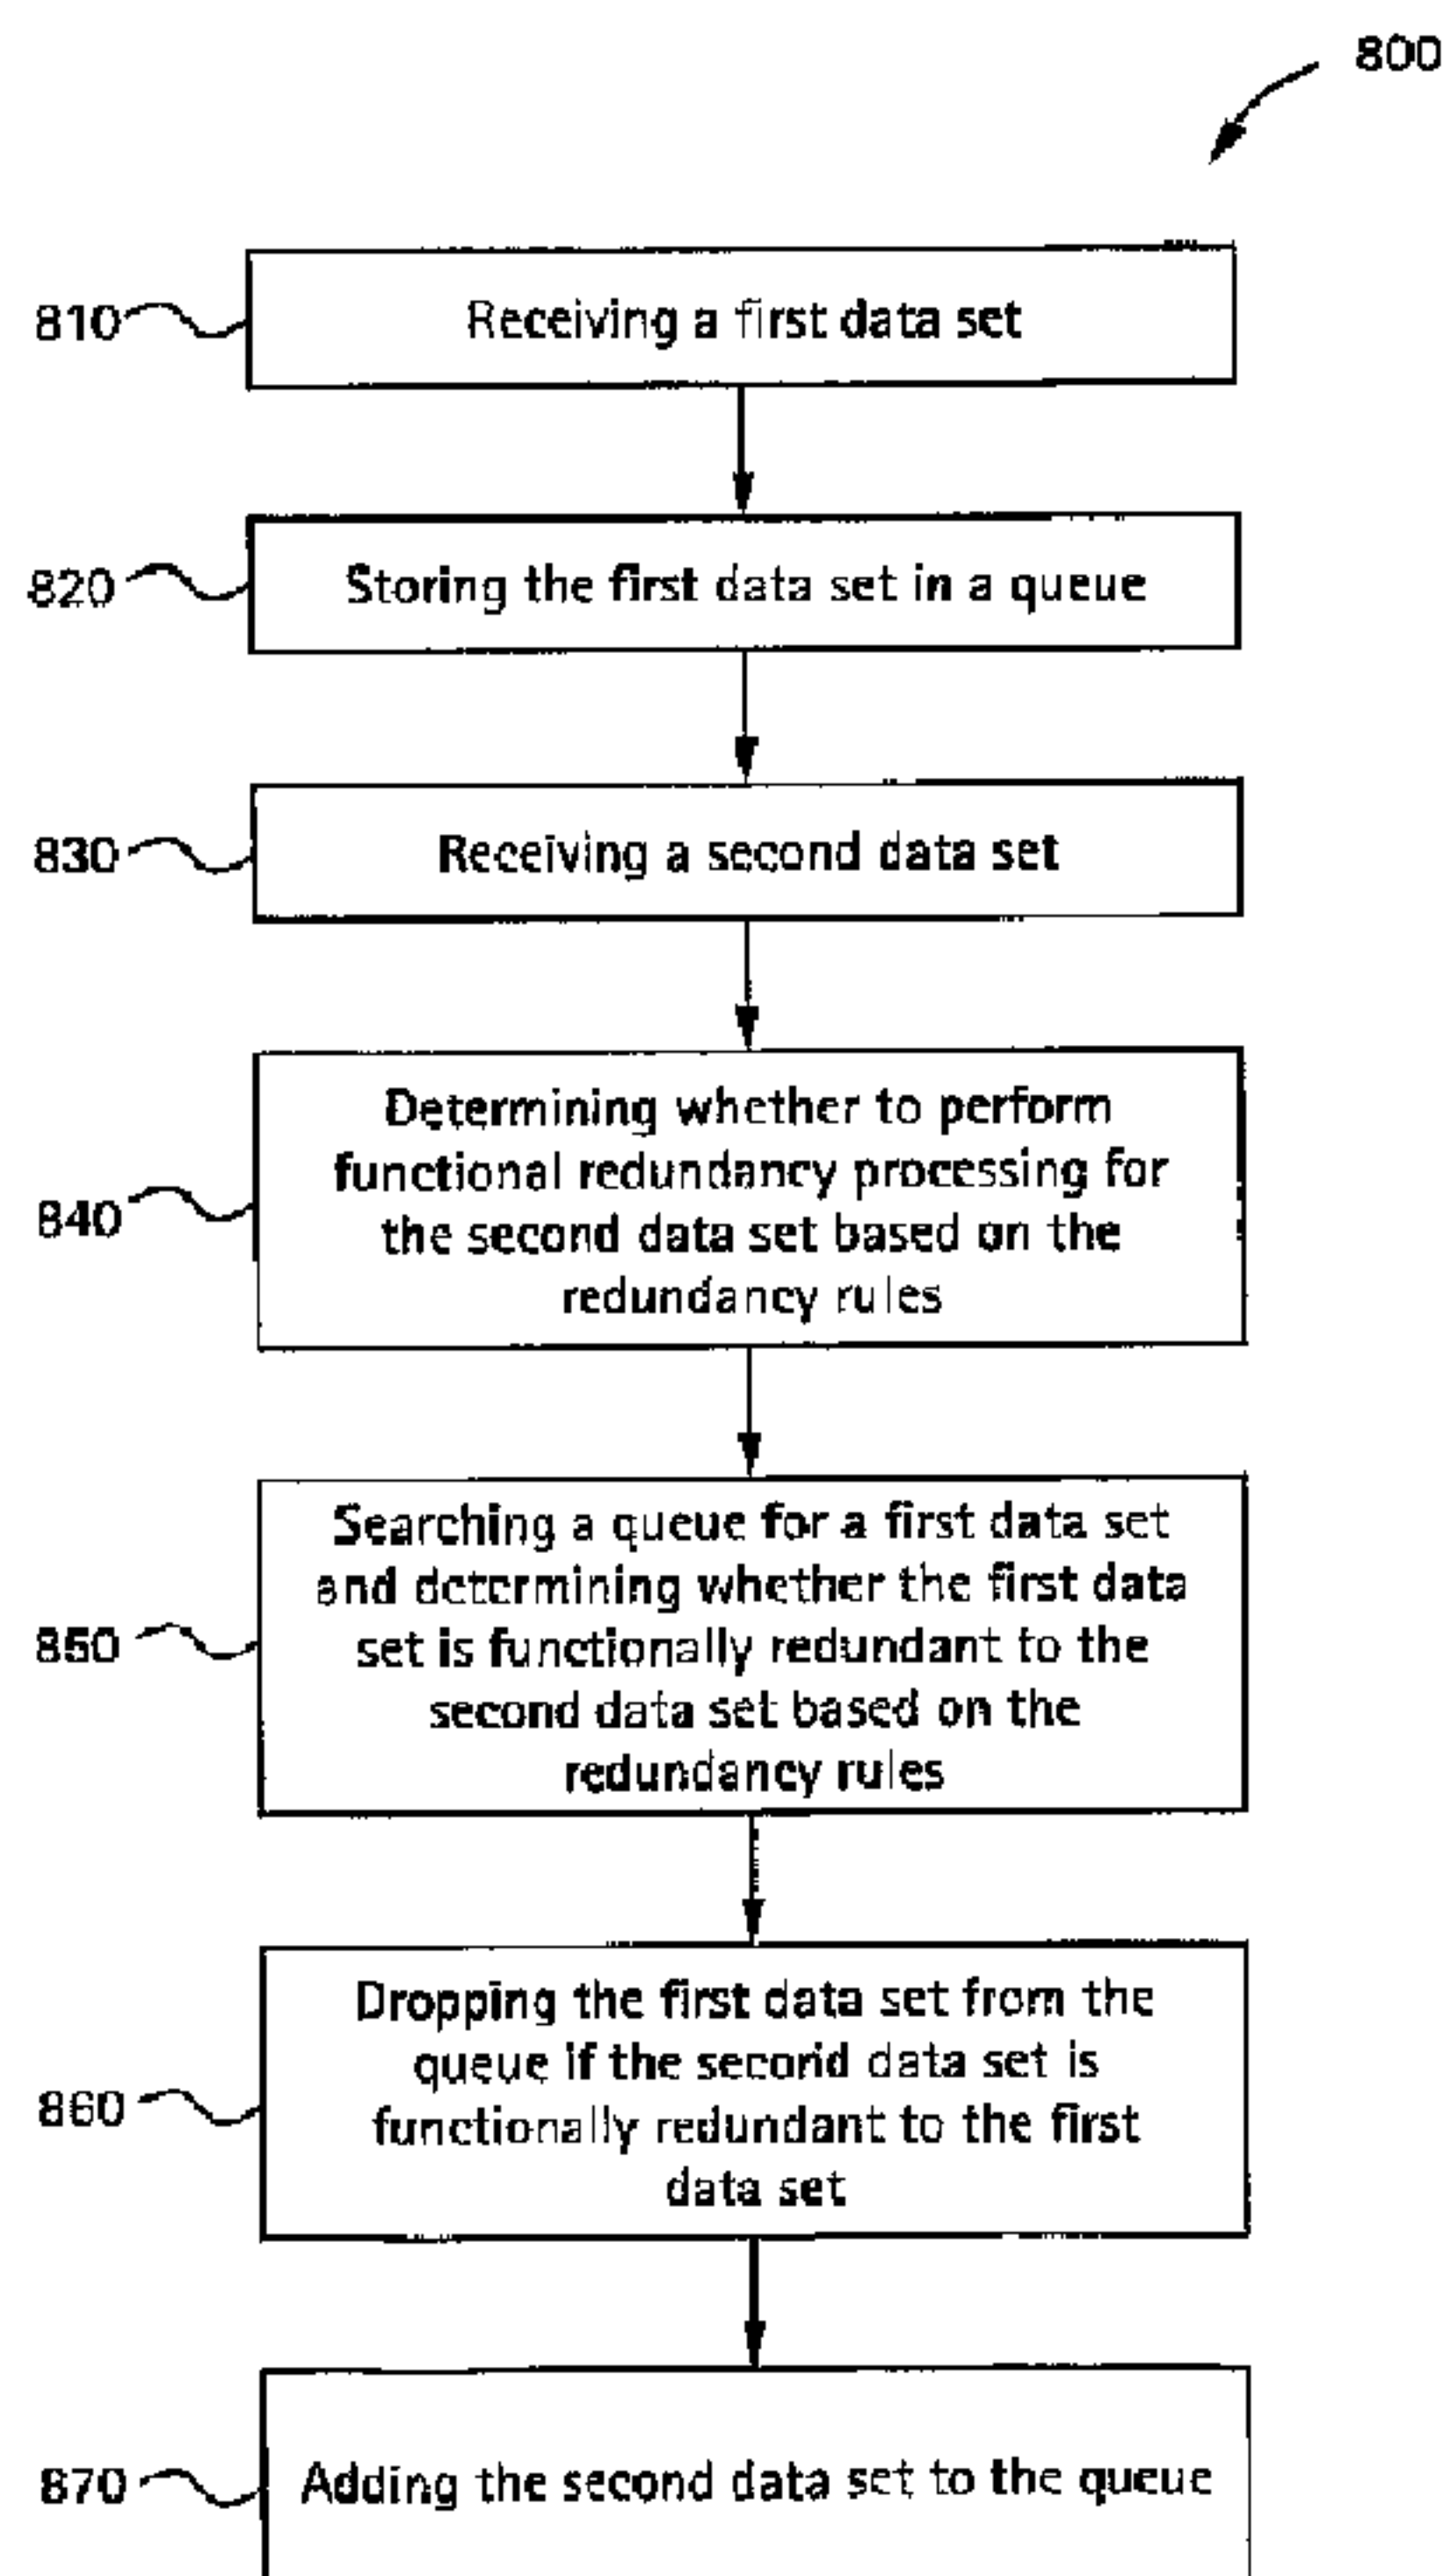


(86) Date de dépôt PCT/PCT Filing Date: 2007/05/15
 (87) Date publication PCT/PCT Publication Date: 2007/11/29
 (45) Date de délivrance/Issue Date: 2012/11/27
 (85) Entrée phase nationale/National Entry: 2008/11/18
 (86) N° demande PCT/PCT Application No.: US 2007/011649
 (87) N° publication PCT/PCT Publication No.: 2007/136625
 (30) Priorité/Priority: 2006/05/18 (US11/436,912)

(51) Cl.Int./Int.Cl. *H04L 29/02* (2006.01)
 (72) Inventeurs/Inventors:
 SMITH, DONALD L., US;
 GALLUSCIO, ANTHONY P., US;
 KNAZIK, ROBERT J., US
 (73) Propriétaire/Owner:
 HARRIS CORPORATION, US
 (74) Agent: GOUDREAU GAGE DUBUC

(54) Titre : PROCÉDE ET SYSTÈME POUR LA QUALITÉ DE SERVICE BASÉE SUR LA REDONDANCE FONCTIONNELLE

(54) Title: METHOD AND SYSTEM FOR FUNCTIONAL REDUNDANCY BASED QUALITY OF SERVICE



(57) Abrégé/Abstract:

Certain embodiments of the present invention provide for a system (100) and method (800) (900) for preserving bandwidth in data networks. The method (800) (900) includes determining whether to perform functional redundancy processing for a current data set (840) (940). Determining whether to perform functional redundancy processing for a current data set may be conducted according to redundancy rules. In performing functional redundancy processing, the method includes searching at least one queue for a data set that is functionally redundant to the current data set (850) (950). The searching may be conducted according to redundancy rules. If a queued data set is found to be functionally redundant to the current data set, the queued data set may be dropped from the queue and the current data set added to the queue (860) (960). In such a manner, functionally redundant messages are dropped prior to transmission and bandwidth is optimized.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
29 November 2007 (29.11.2007)

PCT

(10) International Publication Number
WO 2007/136625 A3(51) International Patent Classification:
H04L 12/66 (2006.01)(74) Agents: YATSKO, Michael, S. et al.; Harris Corporation,
1025 W. NASA Blvd., MS A-11I, Melbourne, FL 32919
(US).(21) International Application Number:
PCT/US2007/011649(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, ME, MG, MK, MN, MW, MX,
MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO,
RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date: 15 May 2007 (15.05.2007)

(25) Filing Language: English

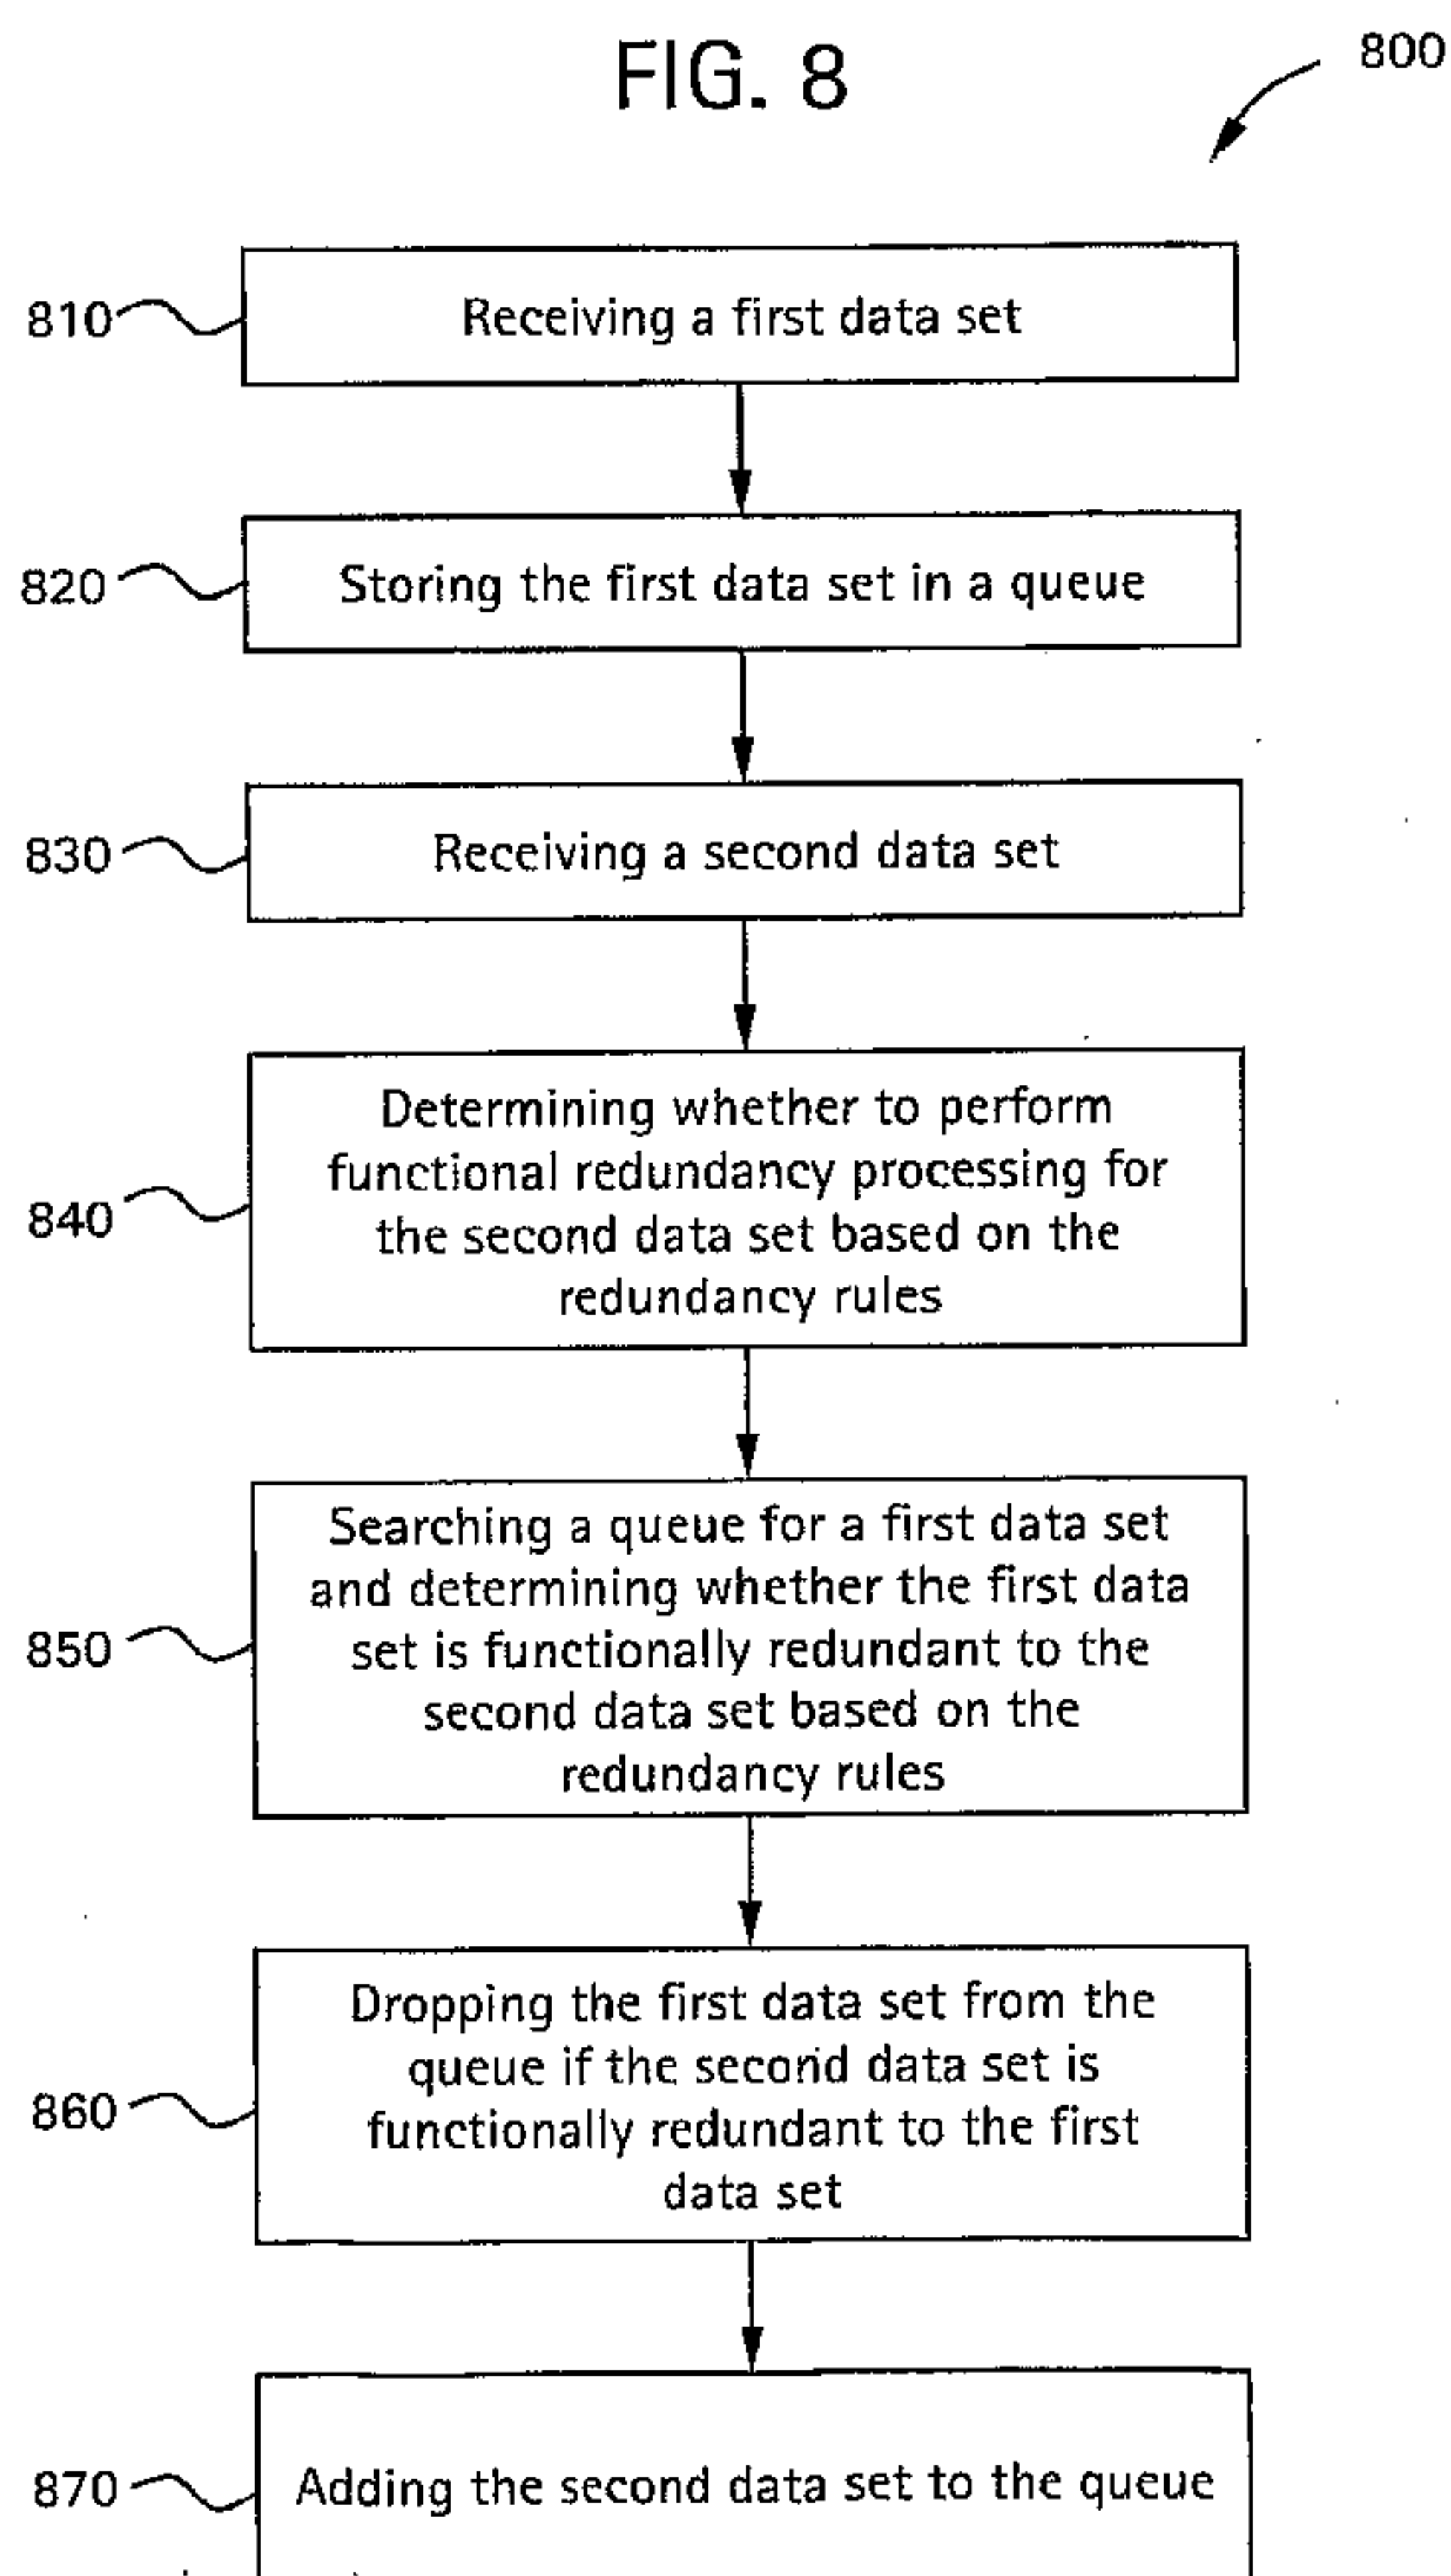
(26) Publication Language: English

(30) Priority Data:
11/436,912 18 May 2006 (18.05.2006) US(71) Applicant (*for all designated States except US*): HARRIS
CORPORATION [US/US]; 1025 W. NASA Blvd., MS
A-11I, Melbourne, FL 32919 (US).(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).(72) Inventors: SMITH, Donald, L.; 584 Hawksbill Island
Drive, Satellite Beach, Florida 32937 (US). GALLUS-
CIO, Anthony, P.; 549 Hummingbird Drive, Indialantic,
Florida 32903 (US). KNAZIK, Robert, J.; 112 St. Croix
Avenue, Cocoa Beach, FL 32931 (US).

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR FUNCTIONAL REDUNDANCY BASED QUALITY OF SERVICE

FIG. 8



(57) Abstract: Certain embodiments of the present invention provide for a system (100) and method (800) (900) for preserving bandwidth in data networks. The method (800) (900) includes determining whether to perform functional redundancy processing for a current data set (840) (940). Determining whether to perform functional redundancy processing for a current data set may be conducted according to redundancy rules. In performing functional redundancy processing, the method includes searching at least one queue for a data set that is functionally redundant to the current data set (850) (950). The searching may be conducted according to redundancy rules. If a queued data set is found to be functionally redundant to the current data set, the queued data set may be dropped from the queue and the current data set added to the queue (860) (960). In such a manner, functionally redundant messages are dropped prior to transmission and bandwidth is optimized.

WO 2007/136625 A3

WO 2007/136625 A3



Published:

— *with international search report*

(88) Date of publication of the international search report:
14 August 2008

**METHOD AND SYSTEM FOR FUNCTIONAL REDUNDANCY BASED
QUALITY OF SERVICE**

The presently described technology generally relates
5 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
10 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
15 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.

20 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
25 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
30 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
5 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-
10 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
15 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
20 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.

25 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
30 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 range 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 is shared by participants, and 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
5 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.

10 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
15 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
20 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
25 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
30 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
5 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
10 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
15 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
20 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
25 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
30 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
5 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
10 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
15 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,
20 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
25 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
30 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.

Certain embodiments of the present invention provide for a method for data communication. The method includes receiving a first data set and storing the first data set in a queue. The method also includes receiving a second data set. Next, the method includes determining whether to perform functional redundancy processing for the second data set based on redundancy rules. The redundancy rules may be controlled by a selected mode. Next, the method includes searching the queue for the first data set and determining whether the first data set is functionally redundant to the second data set based on the redundancy rules. If the first data set is functionally redundant to the second data set, the first data

set is dropped from the queue and the second data set is added
to the queue. In an embodiment, the spot in the queue is
preserved as the first data set is replaced by the second data
set. In another embodiment, the second data set is added to
5 the end of the queue.

Certain embodiments of the present invention provide
for a computer-readable medium having a set of instructions
for execution on a processing device. The set of instructions
includes a receiving routine for receiving a first data set
10 and a second data set. The set of instructions also includes
a storing routine for storing the first data set in a queue.
The set of instructions also includes a determining routine
for determining whether to perform functional redundancy
processing for the second data set based on redundancy rules.
15 The redundancy rules are controlled by a selected mode. The
set of instructions also includes a searching routine for
searching the queue for the first data set and determining
whether the first data set is functionally redundant to the
second data set based on the redundancy rules. If the first
20 data set is functionally redundant to the second data set,
dropping the first data set from the queue and adding the
second data set to the queue.

Certain embodiments of the present invention include
a method for data communication, the method comprising
25 receiving a first data set and receiving a second data set.
The method also includes reviewing a selected mode to
determine whether to perform functional redundancy processing
on the second data set. The selected mode having a set of
redundancy rules for determining whether to perform functional
30 redundancy processing. The method also includes performing
functional redundancy processing by determining whether the
second data set is functionally redundant to the first data
set. The second data set being functionally redundant to the
first data set according to the set of redundancy rules. If

the first data set is functionally redundant to the second data set, dropping the first data set from a queue and adding the second data set to the queue.

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 data communication environment operating with an embodiment of the present invention.

Fig. 6 illustrates a flow diagram in accordance with an embodiment of the present invention.

Fig. 7 illustrates a flow diagram in accordance with an embodiment of the present invention.

Fig. 8 illustrates a method in accordance with an embodiment of the present invention.

Fig. 9 illustrates a method 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.

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

15 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.

20 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.

30 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
5 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.

10 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
15 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.

20 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. Optimizing bandwidth usage may include removing functionally redundant messages,
25 message stream management or sequencing, and message compression, for example. 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. Setting information priority may
30 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 identifies 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 may 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 from 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 from 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
5 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.

10 Fig. 4 illustrates a data communication environment 400 operating with an embodiment of the present invention. The environment 400 includes one or more source nodes 420, a data communication system 410, and one or more destination nodes 430. The data communication system 410 is in
15 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 wire, radio, satellite, network links, and/or through inter-process communication, for
20 example. In certain embodiments, the data communication system 410 may communicate with one or more source nodes 420 and/or destination nodes 430 over one or more tactical data networks. The components of the system 400 may be single units, separate units, may be integrated in various forms, and may be
25 implemented in hardware and/or in software.

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.
30 In certain embodiments, the data communication system 410 may include a memory unit and/or data base for storing computer instructions and rules. The data communication system 410 may also include a processor for processing data, rules, and instructions. 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
5 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.

In certain embodiments, the data communication system 410 is transparent to other applications. For example,
10 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
15 the data communication system 410, may be unaware of the prioritization of data performed by the data communication system 410.

The components, elements, and/or functionality of the data communication system 410 may be implemented alone or
20 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
25 other processing device.

The source node 420 may include a sensor or measurement device to collect data or telemetry information. For example, the source node 420 may be a Global Positional System sensor to indicate positional data for a mobile
30 vehicle, such as a tank, humvee unit, personal transporter, or individual soldier. In another example, the source node 420 may be a photography unit, such as a video or still picture camera that acquires video or images. In another example, the source node may be a communication module, such as a radio or

microphone. The destination node 430 may be any device or system interested in the data acquired by the source node 420. For example, the destination node 430 may be a receiver, a central computer system, and/or computers utilized by a
5 command post or reconnaissance unit.

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
10 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
15 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
20 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
25 (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
30 Transport Protocol (RTP). In certain embodiments, the data may also include time stamp information. The time stamp information may indicate, for example, the time of data acquisition by the source node 420.

In certain embodiments, the data includes a header and a payload. The header may include some or all of the protocol information and the time stamp information, for example. In certain embodiments, some or all of the protocol information and the time stamp 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, and the time stamp information is stored in yet another buffer.

In an embodiment, the source node 420 and the data communication system 410 may be part of the same mobile unit. A mobile unit may be a tank, humvee unit, personal transporter, individual soldier, unmanned aerial vehicle (UAV), or other mobile unit. A tank may have a GPS sensor to indicate positional data as a source unit 420. The positional data may be communicated to the data communication system 410. The data communication system 410 may be located on the tank. The data communication system 410 may prepare the data for communication to the destination node 430. As part of the preparation for communication to the destination node, the data communication system 410 may execute some form of network access protocol. The network access protocol may include requesting network access from a control unit, sensing carrier availability, or other form of access control.

In an example, the network for which the data communication system 410 is attempting to acquire access may be bandwidth-constrained. In addition, one or more links may be unreliable and/or intermittently disconnected. Accordingly, the data communication system 410 may temporarily

queue data received from the source 420 until the data communication system 410 has been able to access the network to communicate the data to the destination 430. For example, the source 420 may acquire a first data set. The source 420
5 may communicate the first data set to the data communication system 410. The data communication system 410 may not currently have network access to transmit the first data set to the destination 430. The first data set may then be temporarily queued in a queue until the data communication
10 system 410 has network access. The source, in the mean time, may acquire a second data set. The second data set may be communicated to the data communication system 410. The data communication system 410 may not yet have network access to transmit the first data set or the second data set to the
15 destination 430. If the data is the type in which the timing of the data is relevant, for example the most recent data is relevant to the destination 430, such as for example positional data, the first data set is no longer relevant. In other words, the first data set has become functionally
20 redundant in view of the second data set. Accordingly, transmission of the first data set to the destination 430 may unnecessarily consume network bandwidth.

In another example, a source may acquire a first data set and communicate the first data set to the data
25 communication system 410. The data communication system 410 may or may not currently have network access to transmit the first data set to the destination 430. If the data communication system 410 has access to the network, the data communication system 410 may transmit the first data set to
30 the destination 430. The source 420 may acquire a second data set. The data communication system 410 may or may not currently have network access to transmit the second data set to the destination 430. If the data communication system 410 has access to the network, the data communication system 410

may transmit the second data set to the destination 430. If the data is the type in which the time of acquisition causes the data to change a relatively small amount between consecutive acquisitions, such as for example pictures taken every second from an unmanned aerial vehicle (UAV) from 5000 ft. above ground, the consecutive data sets may be functionally redundant, in other words capturing primarily the same image. Accordingly, transmission of the consecutive data sets to the destination 430 may unnecessarily consume network bandwidth.

In another example, a source may acquire a first data set and communicate the first data set to the data communication system 410. The data communication system 410 may or may not currently have network access to transmit the first data set to the destination 430. If the data communication system 410 has access to the network, the data communication system 410 may transmit the first data set to the destination 430. The source 420 may acquire a second data set. The data communication system 410 may or may not currently have network access to transmit the second data set to the destination 430. If the data communication system 410 has access to the network, the data communication system 410 may transmit the second data set to the destination 430. If the content of the first data set is the same, and/or similar to the content of the second data set, such as a lack of an audio component in data sets that are expected to have an audio component, the consecutive data sets may be functionally redundant. Accordingly, transmission of the consecutive data sets to the destination 430 may unnecessarily consume network bandwidth.

Fig. 5 illustrates a data communication environment 500 operating with an embodiment of the present invention. The environment 500 includes the data communication system 410, the source 420, and the destination 430 as shown in

Figure 4. The data communication system 410 is expanded and detailed to illustrate certain embodiments of the invention.

The data communication system 410, as shown in Fig. 5, includes a redundancy rules data base 510, a receiver 520, an operational processor 530, one or more queues 540, and a transmitter 550. The redundancy rules data base 510 is in communication with operational processor 530. The receiver 520 is in communication with the operational processor 530 and source 420. The operational processor 530 is in communication with the queue 540. The data communication system 410, as shown in Fig. 5, may communicate with the source node(s) 420 and/or destination node(s) 430 over links, such as wire, radio, satellite, network links, and/or through inter-process communication, for example. In certain embodiments, the data communication system 410 may communicate with one or more source nodes 420 and/or destination nodes 430 over one or more tactical data networks. The components of the system 500, the connections in the system 500, and the system 500, may be single units, separate units, may be integrated in various forms, and may be implemented in hardware and/or in software.

In an embodiment, the redundancy rules data base 510 may be a data base for storing and communicating various rules and/or protocol to an operational processor 530. The redundancy rules data base 510 and the connection shown between the redundancy rules data base 510 and the operational processor 530 is representative of various components and/or software for performing operations and storing rules and instructions. The receiver 520 is representative of various components and/or software for receiving data from the source 420. The operational processor 530 is representative of a processor for processing data and performing operations based on rules communicated by the redundancy rules data base 510. The queue 540 is representative of various components and/or software for queuing data. The transmitter 550 is

representative of various components and/or software for transmitting data to destination 430.

As discussed above, 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 receiver 520 is shown having a block of data. A data set may include a single block of data, or a plurality of blocks of data.

10 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 may also include time stamp information. The time stamp information may indicate, for example, the time of data acquisition by the source node 420.

30 In certain embodiments, the data includes a header and a payload. The header may include some or all of the protocol information and the time stamp information, for example. In certain embodiments, some or all of the protocol information and the time stamp 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
5 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, and the time stamp information is stored in yet another buffer.

10 In operation, a data set may be provided and/or generated by one or more data sources 420. The data set is received at the receiver 520. The data set may be received over one or more links, for example. For example, the data set may be received at the data communication system 410 from
15 a radio over a tactical data network. As another example, a data set 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, a data set may include a single block of data, or a plurality of
20 blocks of data.

In certain embodiments, the receiver 520 may communicate the data set to an operational processor 530. The operational processor 530 may receive the data set and determine whether to perform functional redundancy processing
25 on the data set. The operational processor 530 may base the determination whether to perform functionally redundant processing on the data set on redundancy rules from the redundancy rules data base 510.

30 In an embodiment, the redundancy rules may be rules that control whether to perform functional redundancy processing on a particular data set, and if functional redundancy processing is to be performed, how to perform the functional redundancy processing. The mode, as selected by a user or selected by computer software based on various

factors, may define "sets" of redundancy rules applicable to the data set.

For example, the redundancy rules may be set as either "on" or "off" based the "mode" selected by a user. As
5 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 affect changes in rules, modes, and/or
10 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" or
15 selection of a mode by a user, for example.

If the selected mode utilizes a set of redundancy rules, then the operational processor 530 may analyze the data set to determine whether to perform functional redundancy processing based on the redundancy rules. In an embodiment,
20 the available modes may have different redundancy rules. For example, mode A may have a first set of redundancy rules and mode B may have a second set of redundancy rules. A set of redundancy rules may belong to a single mode, or a plurality of modes. A mode may include more than one set of redundancy
25 rules.

In determining whether to perform functional redundancy processing, the operational processor 530 may read information from the data block, such as the protocol information, payload, and/or time stamp of the data block. As
30 discussed above, the protocol information may include, for example, a source address, a destination address, a source port, a destination port, a protocol type and/or a time stamp. If, for example, the redundancy rules in the redundancy rules data base 510, for the selected mode, indicate that functional

redundancy should be performed for data from a particular source, then the operational processor 530 performs functional redundancy processing for data from the particular source. For example, the redundancy rules of the selected mode may indicate to perform functional redundancy processing for data from source node A, but not for data from source node B. Accordingly, functional redundancy processing would not be performed for data from source node B in this example.

If the operational processor 530 determines that functional redundancy processing should be performed on the data set, the operational processor 530 may perform functional redundancy processing on the data set according to the redundancy rules. The redundancy rules may be determined by the selected mode. For example, the redundancy rules may dictate that the operational processor 530 search the queue 540 for a data set that originates from the same source as the current data set. The redundancy rules may dictate that if an older data set is found in the queue from a particular source, the older data set is functionally redundant to the current, newer, data set. Alternatively, the redundancy rules may dictate that the operational processor 530 search the queue 540 for a data set that originates from the same source as the current data set. The redundancy rules may instruct the operational processor to review the time stamp information of the queued data set and the current data set to determine if the queued data set and the current data set contain functionally redundant information. In another alternative, the redundancy rules may instruct the operational process to review the payload information to determine if the queued data set and the current data set have similar and/or the same functional contents and thus functionally redundant.

For example, the redundancy rules may identify a first, earlier data set as functionally redundant to a second, later data set if the first data set and second data set are

acquired from the same source. In another example, the
redundancy rules may identify a first, earlier data set as
functionally redundant to a second, later data set if the
first data set and second data set were acquired with less
5 than a predetermined time period between the first data set
and second data set. In another example, the redundancy rules
may identify a second, later data set as functionally
redundant to a first, earlier data set if the first data set
and second data set have a functionally similar and/or the
10 same content, such as a lack of an audio component in a data
set expected to have an audio component, for example.

An example of data that may be classified as
functionally redundant may be positional data. A source 420,
for example a GPS indicator, may generate a first data set for
15 reporting the position of a humvee at a particular time. Due
to network constraints, the first data set may be stored in
the queue 540. As the humvee may be moving, the source 420
may generate a second data set reporting a different position
from the first data set. When the operational processor 530
20 receives the second data set, the operational processor 530
may determine that the selected mode utilizes functional
redundancy for positional data from the source 420.

The operational processor 530 may then search the
queue 540 to determine if a data set from the source 420 is
25 stored in the queue 540. If a data set from the source 420 is
located, the operational processor 530 may determine that the
first data set of positional data is functionally redundant to
the second data set of positional data. The operational
processor 530 may drop the earlier first data set of
30 positional data from the queue 540, as the first data set of
positional data, acquired earlier in time, is now less
accurate than the second data set, acquired later in time, of
positional data. The operational processor 530 may then add
the second data set of positional data to the queue 540.

The operational processor 530 may add the second data set to the queue 540 such that the order of transmission of the queue 540 is unchanged. For example, the second data set may replace the first data set in the location of the queue. In an embodiment, the spot in the queue is preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. The operational processor 530 may add the second data set to the queue 540 in a first-in-first-out protocol. In such a manner, the most recent positional data is sent to the destination 430 without burdening the network with positional data that is "old" and no longer accurate nor relevant.

Another example of data that may be classified as functionally redundant may be consecutive data sets acquired over a short period of time. For example, an unmanned aerial vehicle (UAV) may be flying at 5000 feet and snapping one picture per second. The UAV may be flying relatively slowly, so the pictures taken by the UAV are not substantially different from one another other. As consecutive pictures taken over a short period of time may not illustrate new information, the consecutive pictures may be functionally redundant. Transmitting the functionally redundant pictures may unnecessarily consume valuable bandwidth.

In an example provided for better understanding, a UAV may capture a first picture (first data set in this example) from a camera (the source 420 in this example), and communicate the first picture to the receiver 520. The receiver 520 may communicate the first picture to the operational processor 530. The operational processor 530 may communicate the first picture to the queue 540. The first picture may wait in the queue 540 until the transmitter 550 has network access. In certain embodiments, the transmitter 550 may not have network access, or may not request network

access, before a second picture (second data set in this example) from the camera (the source 420 in this example) is communicated to the operational processor 530.

5 The operational processor 530, upon receiving the second data set, determines whether to perform functional redundancy processing on the second data set. As discussed above, the operational processor 530 may read information from the data block, such as the protocol information, payload, and/or time stamp of the data block. The protocol information
10 may include, for example, a source address, a destination address, a source port, a destination port, a protocol type and/or a time stamp. When the operational processor 530 receives the second data set, the operational processor 530 may determine that the selected mode utilizes functional
15 redundancy for data from the source 420 (camera on the UAV in this example) according to redundancy rules. For example, the redundancy rules may specify that a time threshold separate data sets transmitted to the destination 430 from the source 420.

20 The operational processor 530 may search the queue 540 to determine if a data set from the source 420 is stored in the queue 540. If a data set from the source 420 is located, the operational processor 530 may review the time stamp of the first data set and the time stamp of the second
25 data set. The operational processor 530 may determine that the difference in acquisition time, as identified in the time stamps, between the first data set and the second data set is less than a particular time threshold. If the difference between the acquisition time of the first data set and the
30 second data set is less than a particular time threshold, the operational processor 540 may determine that the first data set is functionally redundant to the second data set. The time threshold value may be determined by the selected mode.

In an embodiment, the operational processor 530 may drop the earlier first data set (first picture) from the queue 540, as the first data set, acquired earlier in time, is now less relevant than the second data set (second picture),
5 acquired later in time. The operational processor 530 may then add the second data set (second picture) to the queue 540. The operational processor 530 may add the second data set to the queue 540 such that the order of transmission of the queue 540 is unchanged. For example, the second data set may
10 replace the first data set in the location of the queue. In an embodiment, the spot in the queue is preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. The operational processor 530 may add the second data
15 set to the queue 540 in a first-in-first-out protocol.

Another example of data that may be classified as functionally redundant may be data sets having the same or similar content. If the content of the first data set is the same, and/or similar to the content of the second data set,
20 such as a lack of an audio component in data sets that are expected to have an audio component, the consecutive data sets may be functionally redundant. For example, data sets expected to contain audio data that do not contain audio data may indicate that the data set is transmitting radio silence.
25 If radio silence is not useful information to the destination 430, transmitting data sets containing radio silence may unnecessarily consume bandwidth. Accordingly, consecutive data sets having little or no audio component, for data sets where an audio component is expected, may be functionally
30 redundant.

In an example provided for better understanding, a radio or microphone (the source 420 in this example) may acquire a first data set and communicate the first data set to the receiver 520. The receiver 520 may communicate the first

data set to the operational processor 530. The operational processor 530 may communicate the first data set to the queue 540. The first data set may wait in the queue 540 until the transmitter 550 has network access. In certain embodiments, 5 the transmitter 550 may not have network access, or may not request network access, before a second data set is communicated to the operational processor 530 from the source 420.

10 The operational processor 530, upon receiving the second data set, determines whether to perform functional redundancy processing on the second data set. As discussed above, the operational processor 530 may read information from the data block, such as the protocol information, payload, and/or time stamp of the data block. The protocol information 15 may include, for example, a source address, a destination address, a source port, a destination port, a protocol type and/or a time stamp. When the operational processor 530 receives the second data set, the operational processor 530 may determine that the selected mode utilizes functional 20 redundancy for data from the source 420 (a radio microphone in this example) according to the redundancy rules. For example, the redundancy rules may specify that an audio component be present in data payload from the source 420.

25 The operational processor 530 may search the queue 540 to determine if a data set from the source 420 is stored in the queue 540. If a data set from the source 420 is located, the operational processor 530 may review the payload of the first data set at the expected location of the audio component. The operational processor 530 may also review the 30 payload of the second data set at the expected location of the audio component. The operational processor 530 may determine that the first data set does not contain an audio component and the second data set does not contain an audio component. If the first data set and the second data set both do not have

an audio component, and both data sets are expected to have an audio component, transmitting both the first and second data sets does not communicate useful information to the destination 430. Accordingly, the operational processor 540
5 may determine that the first data set is functionally redundant to the second data set or that the second data set is functionally redundant to the first data set.

In an embodiment, the operational processor 530 may drop the earlier first data set from the queue 540. The
10 operational processor 530 may then add the second data set to the queue 540. The operational processor 530 may add the second data set to the queue 540 such that the order of transmission of the queue 540 is unchanged. For example, the second data set may replace the first data set in the location
15 of the queue. In an embodiment, the spot in the queue is preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. Alternatively, the operational processor 530 may add the second data set to the queue 540 in
20 a first-in-first-out protocol.

In another embodiment, the operational processor 530 may drop the later, second data set and leave the first data set in the queue 540. In an embodiment, it may be more efficient for the operational processor 530 to drop the second
25 data set rather than add the second data set to the queue and drop the first data set. In the embodiment where the first data set and the second data set generally contain similar and/or the same content, the system 500 may be generally indifferent on selecting either the first data set or the
30 second data set for transmission. Other factors being substantially equal, the operational processor 530 may operate more efficiently to transmit the first data set and delete the second data set for certain data.

The above examples of functional redundancy are only examples. The redundancy rules may be crafted to define redundant data based on source, time, payload, or other factors. The redundancy rules may be applicable to a single mode or multiple modes. A mode may utilize a single set of redundancy rules or multiple sets of redundancy rules.

Fig. 6 illustrates a flow diagram 600 for the operational processor 530 in accordance with an embodiment of the present invention. At step 610, a first data set has been received and stored in the queue 540. A second data set has been received and communicated to the operational processor 530. The data communication system 410 is in a mode that utilizes functional redundancy for data from the source 420. The second data set is operated upon by the operational processor 530. The operations by the operational processor 530 may include determining whether to perform functional redundancy processing for the second data set based on the redundancy rules from the redundancy rules data base 510, as governed by the selected mode. For example, the operational processor 530 may read information from the second data set, such as the protocol information, payload, and/or time stamp of the data set. As discussed above, the protocol information may include, for example, a source address, a destination address, a source port, a destination port, a protocol type and/or a time stamp. If, for example, the redundancy rules in the redundancy rules data base 510, for the selected mode, indicate that functional redundancy should be performed for the current data set, then the operational processor performs functional redundancy processing.

If the redundancy rules dictate that functional redundancy processing is to be performed, the flow diagram indicating the functionality of the operational processor 530, moves on to step 620. If the functional redundancy rules dictate that functional redundancy processing is not to be

performed, then the operational processor 530 (flow diagram) moves on to step 640 and adds the second data set to the queue 540. In an embodiment, the spot in the queue is preserved as the first data set is replaced by the second data set. In
5 another embodiment, the second data set is added to the end of the queue. Once the second data set has been added to the queue 540, the second data set awaits transmission to the destination 430.

At step 620, the operational processor 530 searches
10 the queue 540 according to the redundancy rules. In an embodiment, the operational processor 530 searches the queue 540 for a first data set that is functionally redundant to the second data set. As explained above, the operational
processor determines whether a first data set is functionally
15 redundant to the second data set based on the rules in the redundancy rules data base 510, as dictated by the selected mode. As an example, the operational processor 530 may decide that the first data set is functionally redundant to the second data set if the first and second data sets originate
20 from the same source. As another example, the operational processor 530 may decide that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source and the difference between the time stamps is not greater than a
25 predetermined threshold. As yet another example, the operational processor 530 may decide that the first data set is functionally redundant to the second data set if the first and second data sets have a common element that may not be useful to transmit to the destination 430 multiple times.

30 If, at step 620, the operational processor 530 locates functionally redundant data in the queue 540, the operational processor 530 moves on to step 630. If, at step 630, the operational processor does not locate functionally redundant data in the queue 540, the operational processor 530

moves on to step 640 and adds the second data set to the queue 540 to await transmission.

At step 630, the operational processor 530 may drop the earlier first data set of from the queue 540. The
5 operational processor 530 may then add the later second data set to the queue 540. In an embodiment, the operational processor 530 may add the second data set to the queue 540 such that the order of transmission of the queue 540 is unchanged. In an embodiment, the spot in the queue is
10 preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. The operational processor 530 may add the second data set to the queue 540 in a first-in-first-out protocol. The steps of the flow diagram 600 may execute
15 repeatedly in order to utilize bandwidth efficiently.

Fig. 7 illustrates a flow diagram 700 for the operational processor 530 in accordance with an embodiment of the present invention. At step 710, a first data set has been received and stored in the queue 540. A second data set has
20 been received and communicated to the operational processor 530. The data communication system 410 is in a mode that utilizes functional redundancy for data from the source 420. The second data set is operated upon by the operational processor 530. The operations by the operational processor
25 530 may include determining whether to perform functional redundancy processing for the second data set based on the redundancy rules from the redundancy rules data base 510, as governed by the selected mode. For example, the operational processor 530 may read information from the second data set,
30 such as the protocol information, payload, and/or time stamp of the data set. As discussed above, the protocol information may include, for example, a source address, a destination address, a source port, a destination port, a protocol type and/or a time stamp. If, for example, the redundancy rules in

the redundancy rules data base 510, for the selected mode, indicate that functional redundancy should be performed for the current data set, then the operational processor performs functional redundancy processing.

5 If the redundancy rules dictate that functional redundancy processing is to be performed, the flow diagram indicating the functionality of the operational processor 530, moves on to step 720. If the functional redundancy rules dictate that functional redundancy processing is not to be
10 performed, then the operational processor 530 (flow diagram) moves on to step 740 and adds the second data set to the queue 540. Once the second data set has been added to the queue 540, the second data set awaits transmission to the destination 430.

15 At step 720, the operational processor 530 searches the queue 540 according to the redundancy rules. In an embodiment, the operational processor 530 searches the queue 540 for a first data set that is functionally redundant to the second data set. As explained above, the operational
20 processor determines whether a first data set is functionally redundant to the second data set based on the rules in the redundancy rules data base 510, as dictated by the selected mode. As an example, the operational processor 530 may decide
25 that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source. As another example, the operational processor 530 may decide that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source and the
30 difference between the time stamps is not greater than a predetermined threshold. As yet another example, the operational processor 530 may decide that the first data set is functionally redundant to the second data set if the first

and second data sets have a common element that may not be useful to transmit to the destination 430 multiple times.

If, at step 720, the operational processor 530 locates functionally redundant data in the queue 540, the operational processor 530 moves on to step 730. If, at step 730, the operational processor does not locate functionally redundant data in the queue 540, the operational processor 530 moves on to step 740 and adds the second data set to the queue 540 to await transmission.

At step 730, the operational processor 530 may drop the later, second data set and leave the first data set in the queue 540. In an embodiment, it may be more efficient to for the operational processor 530 to drop the second data set rather than add the second data set to the queue and drop the first data set. In the embodiment where the first data set and the second data set generally contain similar and/or the same content, the system 500 may be generally indifferent on selecting either the first data set or the second data set for transmission. Other factors being substantially equal, the operational processor 530 may operate more efficiently to transmit the first data set and delete the second data set for certain data. The steps of the flow diagram 700 may execute repeatedly in order to utilize bandwidth efficiently.

Fig. 8 illustrates a method 800 in accordance with an embodiment of the present invention. At step 810, a first data set may be received. At step 820, the first data set may be stored in a queue. At step 830, a second data set may be received. At step 840, it may be determined whether to perform functional redundancy processing for the second data set. The determination whether to perform functional redundancy processing for the second data set may depend on the selected mode, and the redundancy rules associated with the selected mode. The mode may be selected either manually, or automatically, based on for example, network conditions.

At step 850, the queue may be searched for a data set that may be functionally redundant to the second data set. Whether a data set located in a queue is functionally redundant to a current data set may be determined by the
5 redundancy rules. The redundancy rules may be determined by the selected mode. As an example, the redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source. As another example, the
10 redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source and the difference between the time stamps is not greater than a predetermined threshold. As yet another example, the
15 redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets have the same and/or similar contents.

At step 860, if the second data set is determined to be functionally redundant to the first data set, the, earlier,
20 first data set may be dropped from the queue. At step 870, the, later, second data set may then be added to the queue. In an embodiment, the second data set may be added to the queue such that the order of transmission of the queue is unchanged. In an embodiment, the spot in the queue is
25 preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. Alternatively, the second data set may be added to the queue in a first-in-first-out protocol. The steps of the flow diagram 800 may execute repeatedly in order
30 to utilize bandwidth efficiently.

One or more of the steps of the method 800 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.

5 Certain embodiments of the present invention may omit one or more of the method 800 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, 10 including simultaneously, than listed above.

The system and method 800 described above may be carried out as part of a computer - readable storage medium including a set of instructions for a computer. The set of instructions may include a receiving routine for receiving a 15 first data set and a second data set. The set of instructions may also include a storing routine for storing the first data set in a queue. The set of instructions may also include a determining routine for determining whether to perform functional redundancy processing for the second data set based 20 on redundancy rules.

The redundancy rules may be controlled by a selected mode routine. The selected mode routine may be selected by a selection routine. The selection routine may be selected by a user or may be selected dynamically based on network 25 conditions. The redundancy rules may include a first redundancy routine for determining whether the second data set and the first data set originate from the same source node. The redundancy rules may also include a second redundancy routine for determining the time between the time stamp of the second data set and the time stamp of the first data set. The 30 redundancy rules may also include a third redundancy routine for determining whether a common element exists between the first data set and the second data set.

The set of instructions may also include a searching routine for searching the queue for the first data set. If the first data set is functionally redundant to the second data set, the first data set may be dropped from the queue and the second data set added to the queue. In an embodiment, the spot in the queue is preserved as the first data set is replaced by the second data set. In another embodiment, the second data set is added to the end of the queue. The step of adding the second data set to the queue includes a first-in-first-out routine for adding the second data set to the queue in a first-in-first-out protocol.

Fig. 9 illustrates a method 900 in accordance with an embodiment of the present invention. At step 910, a first data set may be received. At step 920, the first data set may be stored in a queue. At step 930, a second data set may be received. At step 940, it may be determined whether to perform functional redundancy processing for the second data set. The determination whether to perform functional redundancy processing for the second data set may depend on the selected mode, and the redundancy rules associated with the selected mode. The mode may be selected either manually, or automatically, based on for example, network conditions.

At step 950, the queue may be searched for a data set that may be functionally redundant to the second data set. Whether a data set located in a queue is functionally redundant to a current data set may be determined by the redundancy rules. The redundancy rules may be determined by the selected mode. As an example, the redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source. As another example, the redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets originate from the same source and the

difference between the time stamps is not greater than a predetermined threshold. As yet another example, the redundancy rules may dictate that the first data set is functionally redundant to the second data set if the first and second data sets have a common element that may not be useful to transmit multiple times.

At step 960, if the second data set is determined to be functionally redundant to the first data set, the second data set may be dropped and the first data set may remain in the queue. In an embodiment, it may be more efficient to drop the second data set rather than add the second data set to the queue and drop the first data set. In the embodiment where the first data set and the second data set generally contain similar information, either the first data set or the second data set may be selected for transmission. Other factors being substantially equal, it is more efficient to transmit the first data set and drop the second data set for certain data. The steps of the flow diagram 900 may execute repeatedly in order to utilize bandwidth efficiently.

One or more of the steps of the method 900 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 the method 900 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.

The system and method 900 described above may be carried out as part of a computer - readable storage medium including a set of instructions for a computer. The set of instructions may include a receiving routine for receiving a first data set and a second data set. The set of instructions may also include a storing routine for storing the first data set in a queue. The set of instructions may also include a determining routine for determining whether to perform functional redundancy processing for the second data set based on redundancy rules.

The redundancy rules may be controlled by a selected mode routine. The selected mode routine may be selected by a selection routine. The selection routine may be selected by a user or may be selected dynamically based on network conditions. The redundancy rules may include a first redundancy routine for determining whether the second data set and the first data set originate from the same source node. The redundancy rules may also include a second redundancy routine for determining the time between the time stamp of the second data set and the time stamp of the first data set. The redundancy rules may also include a third redundancy routine for determining whether the content of the first data set is functionally redundant to the content of said second data set.

The set of instructions may also include a searching routine for searching the queue for the first data set. The set of instructions may also include a transmitting routine for transmitting the first data set and dropping the second data set if the contents of the first data set are functionally redundant to the contents of the second data set.

CLAIMS

1. A method for data communication, the method comprising:
receiving a first data set;
5 storing said first data set in a queue;
receiving a second data set;
reviewing a selected processing mode to determine whether to
perform functional redundancy processing for said second data set
based on redundancy rules, said selected mode having a set of
10 redundancy rules for determining whether to perform functional
redundancy processing;
based on the performed determination of whether to perform
functional redundancy processing, searching said queue for said
first data set and determining whether said first data set is
15 functionally redundant to said second data set based on said
redundancy rules; and,
if said first data set is functionally redundant to said second
data set, dropping said first data set from said queue and adding
said second data set to said queue, wherein said first data set is
20 functionally redundant to said second data set based at least in
part on a comparison of said first and second data set resulting in
a degree of difference between said first and second data sets
being below a threshold, said threshold set to include non-
duplicate data.
25
2. The method of claim 1, wherein said selected mode is selected
by a user.
3. The method of claim 1, wherein said redundancy rules
30 determine whether to perform functional redundancy processing based
on information contained in the header of data in the second data
set.
4. The method of claim 1, wherein said redundancy rules
35 determine that the first data set is functionally redundant to said

second data set if the time between the time stamp of said second data set and the time stamp of said first data set is below a predetermined time period.

5 5. The method of claim 1, wherein said step of adding said second data set to said queue includes adding said second data set to said queue in a first-in-first-out protocol.

6. A computer-readable medium having a set of instructions for
10 execution on a processing device, said set of instructions comprising:

a receiving routine for receiving a first data set and a second data set;

a storing routine for storing said first data set in a queue;

15 a determining routine for determining whether to perform functional redundancy processing for said second data set based on redundancy rules, said redundancy rules being controlled by a selected mode; and,

20 a searching routine for searching said queue for said first data set and determining whether said first data set is functionally redundant to said second data set based on said redundancy rules,

25 wherein if said first data set is functionally redundant to said second data set, dropping said first data set from said queue and adding said second data set to said queue, wherein said first data set is functionally redundant to said second data set based at least in part on a comparison of said first and second data sets resulting in a degree of difference between said first and second data sets being below a threshold, said threshold set to include
30 non-duplicate data.

7. The computer-readable medium of claim 6, wherein said selected mode is selected by a selected mode routine, said selected mode routine being controlled by a user.

35

8. A method for data communication, the method comprising:
receiving a first data set;
receiving a second data set;
reviewing a selected mode to determine whether to perform
5 functional redundancy processing on said second data set, said
selected mode having a set of redundancy rules for
determining whether to perform functional redundancy
processing;
performing functional redundancy processing by determining
10 whether said second data set is functionally redundant to said first
data set, said second data set being functionally redundant to said
first data set according to said set of redundancy rules; and,
if said first data set is functionally redundant to said
second data set, dropping said first data set from a queue and
15 adding said second data set to said queue, wherein said first data
set is functionally redundant to said second data set based at
least in part on a comparison of said first and second data sets
being below a threshold, said threshold set to include non-
duplicate data.
20
9. The method of claim 8, wherein said redundancy rules determine
whether to perform functional redundancy processing based on the
source node of the second data set.
- 25 10. The method of claim 8, wherein said redundancy rules determine
that the first data set is functionally redundant to said second
data set if said second data set and said first data set
originate from the same source node.

1/9

FIG. 1




FIG. 2

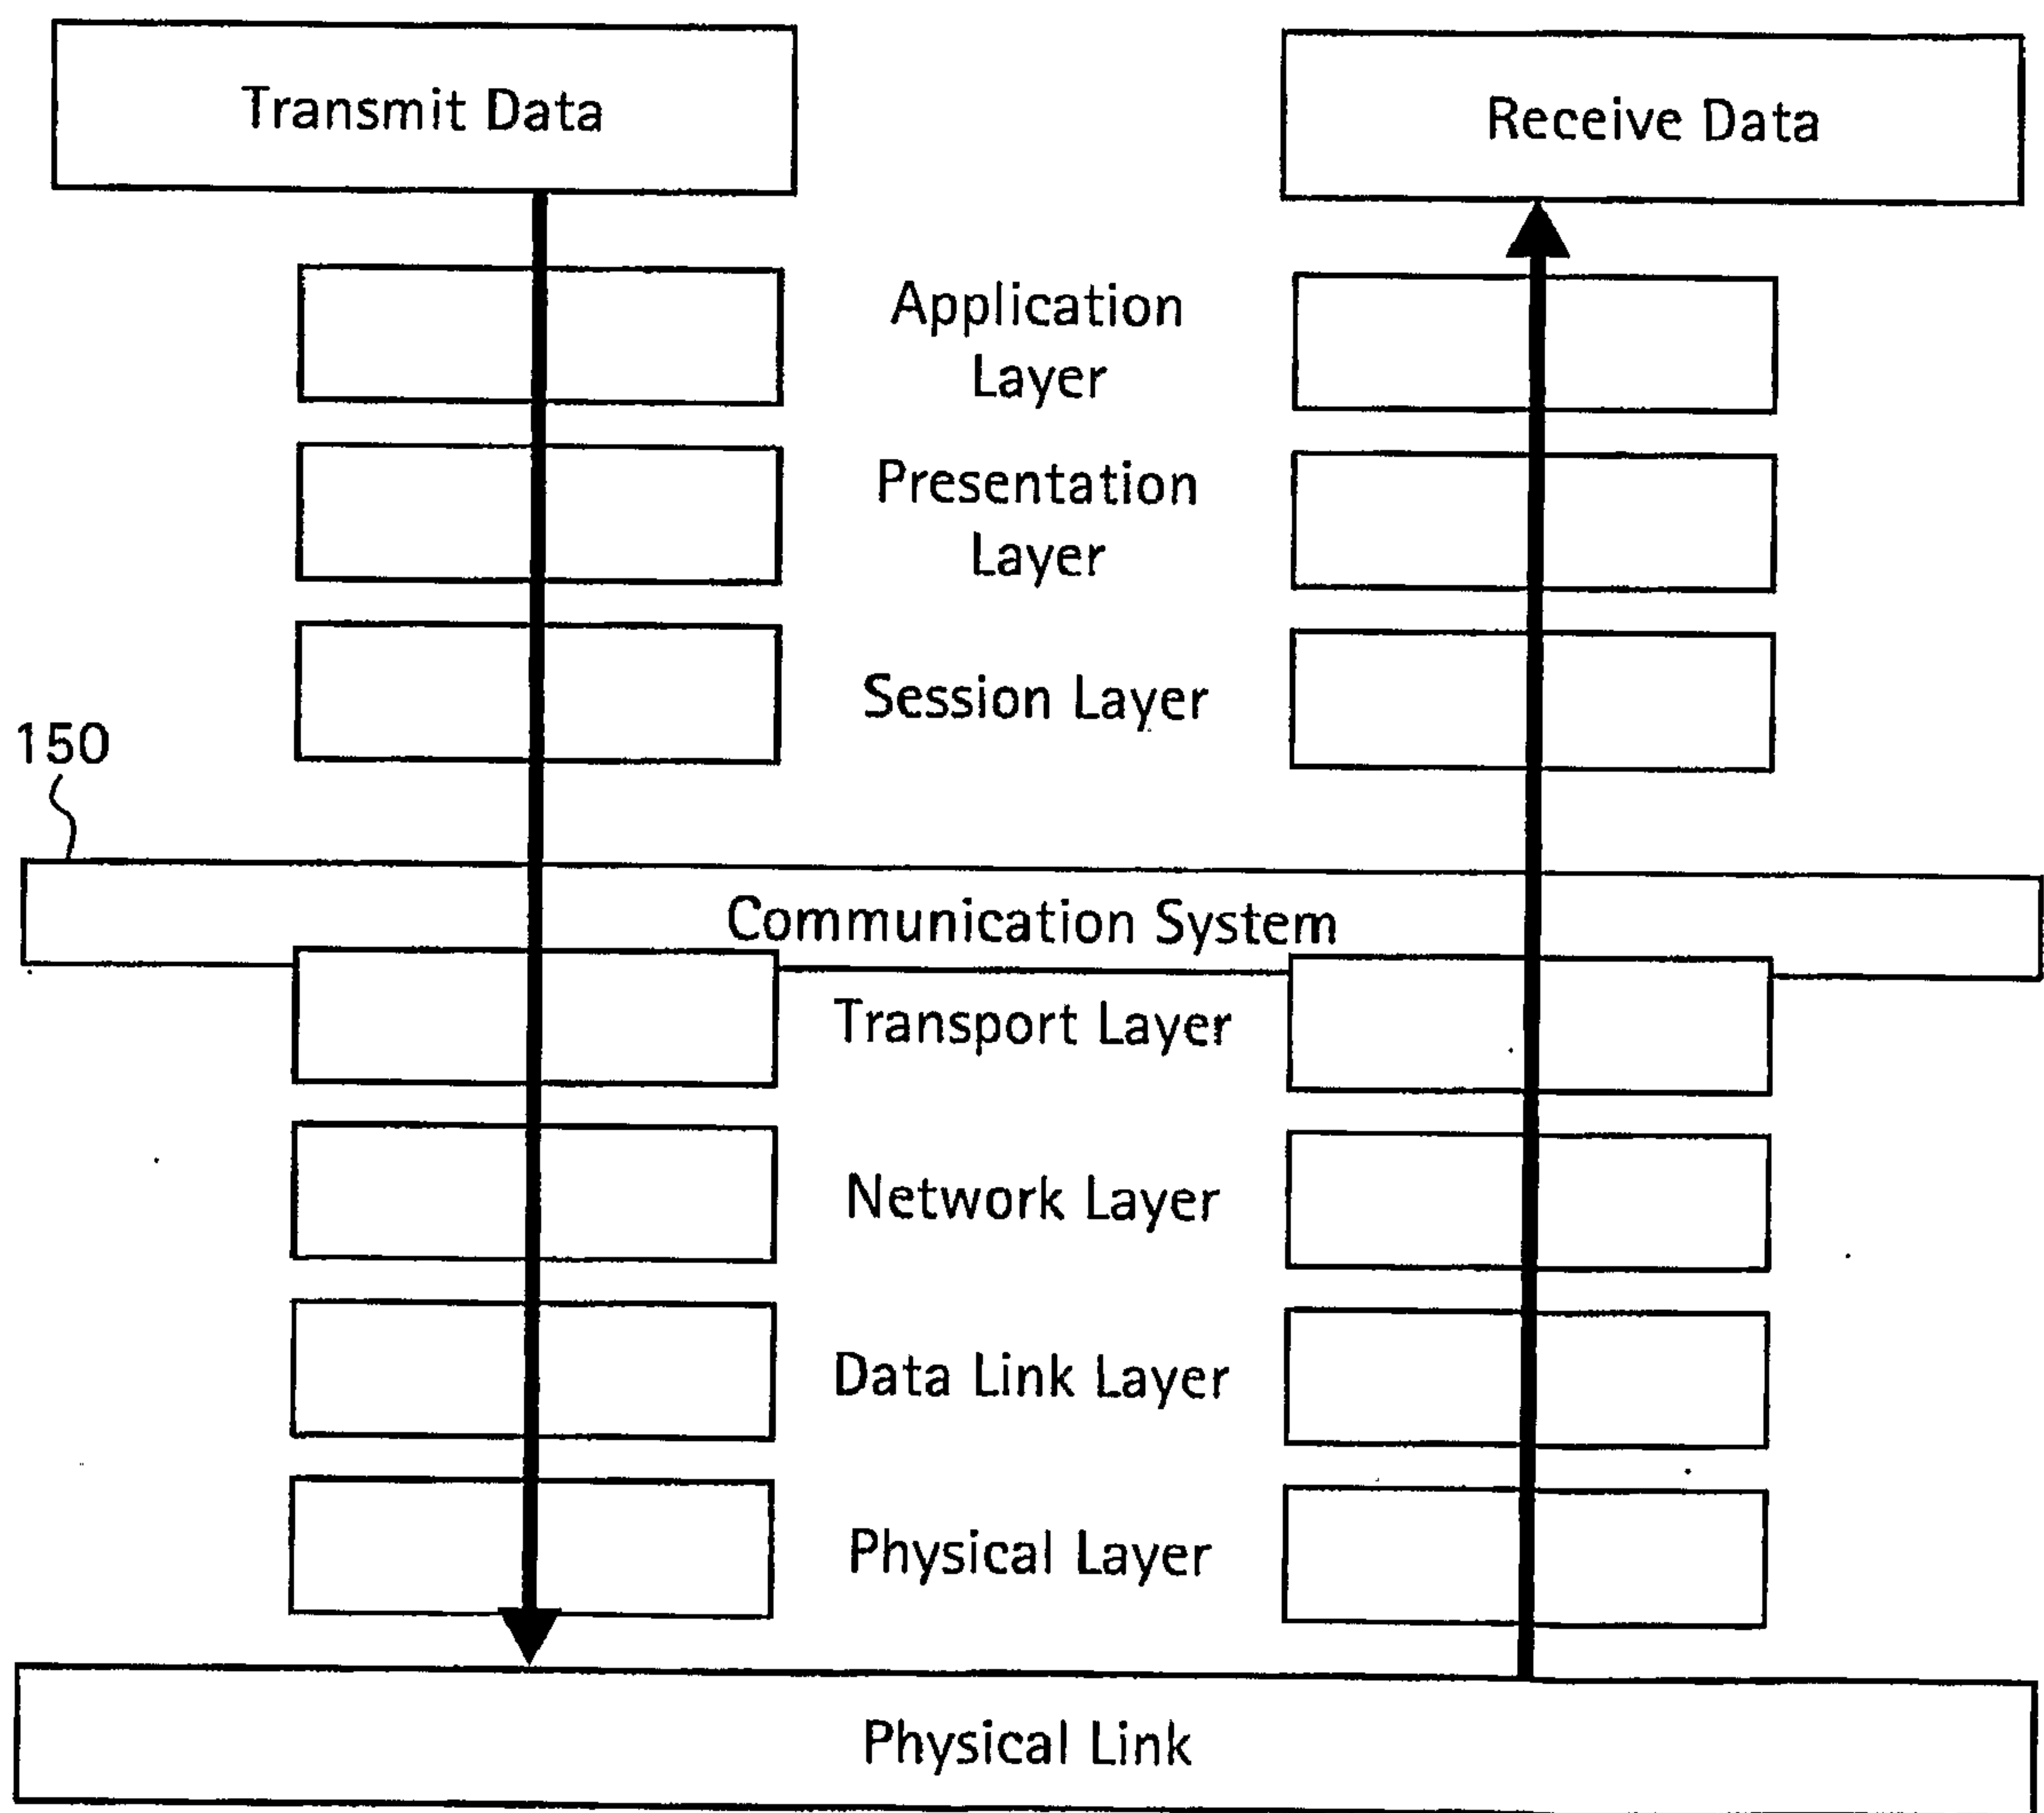


FIG. 3




FIG. 4



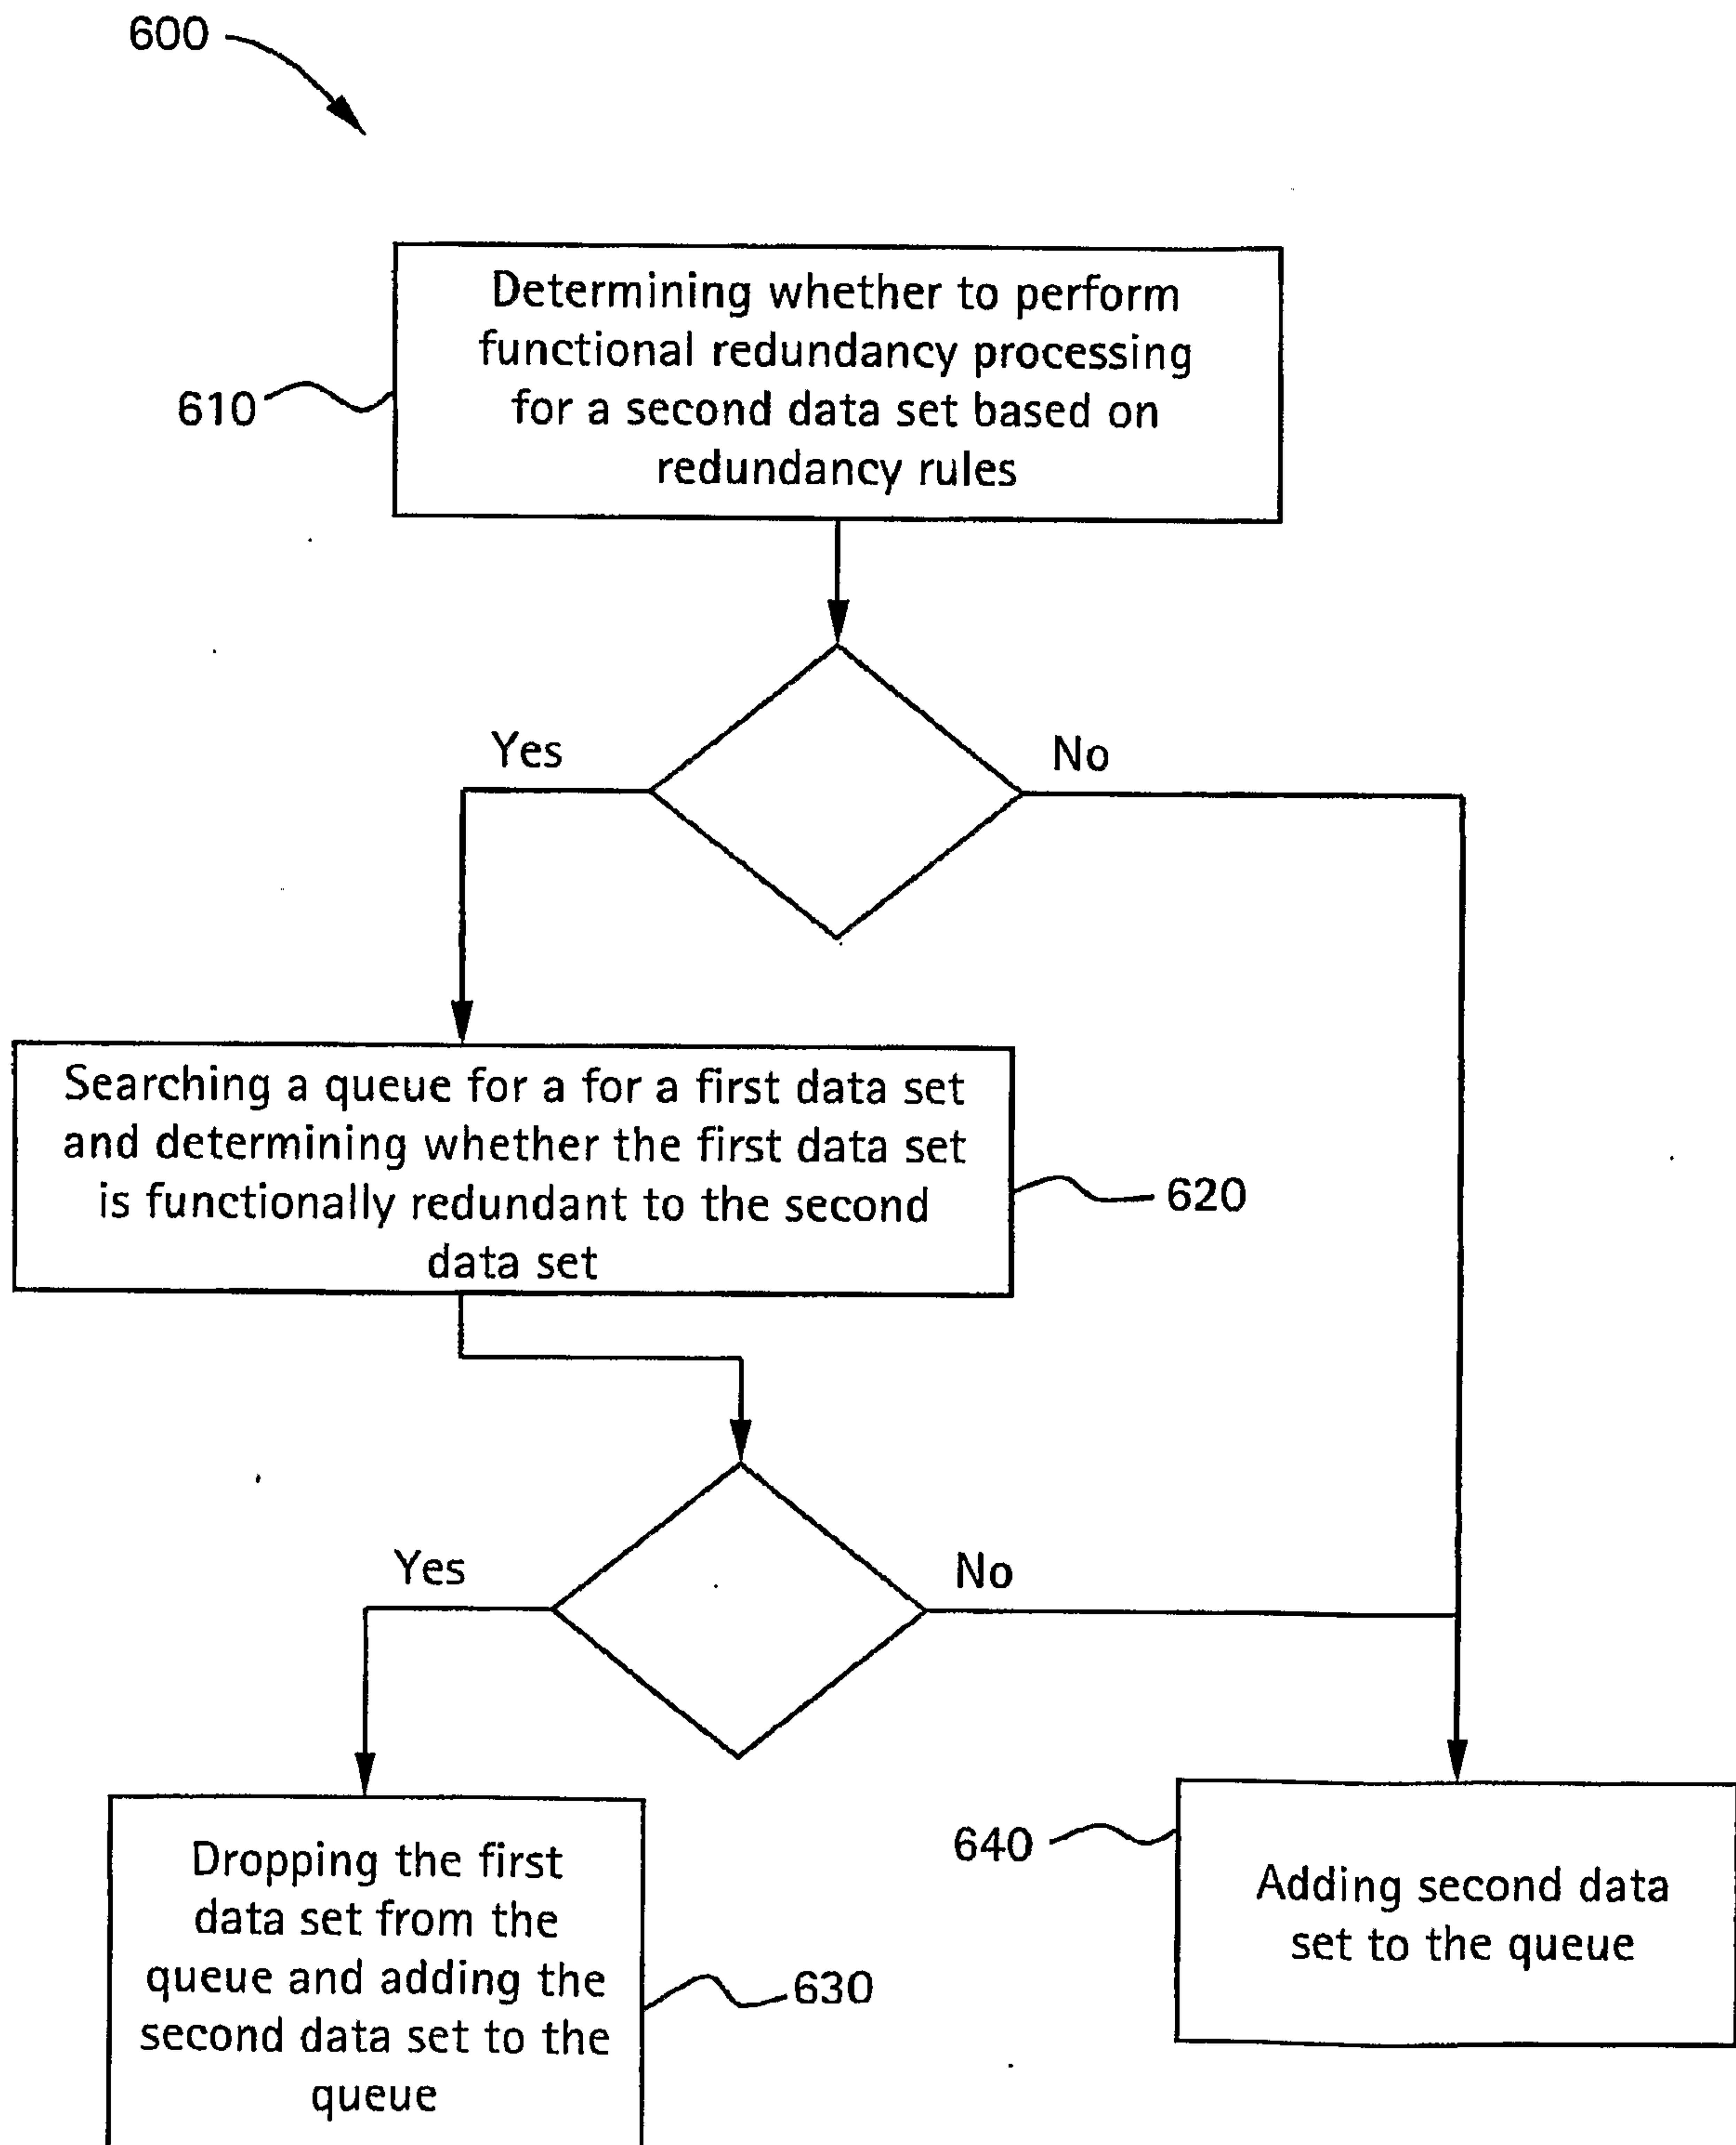


FIG. 5



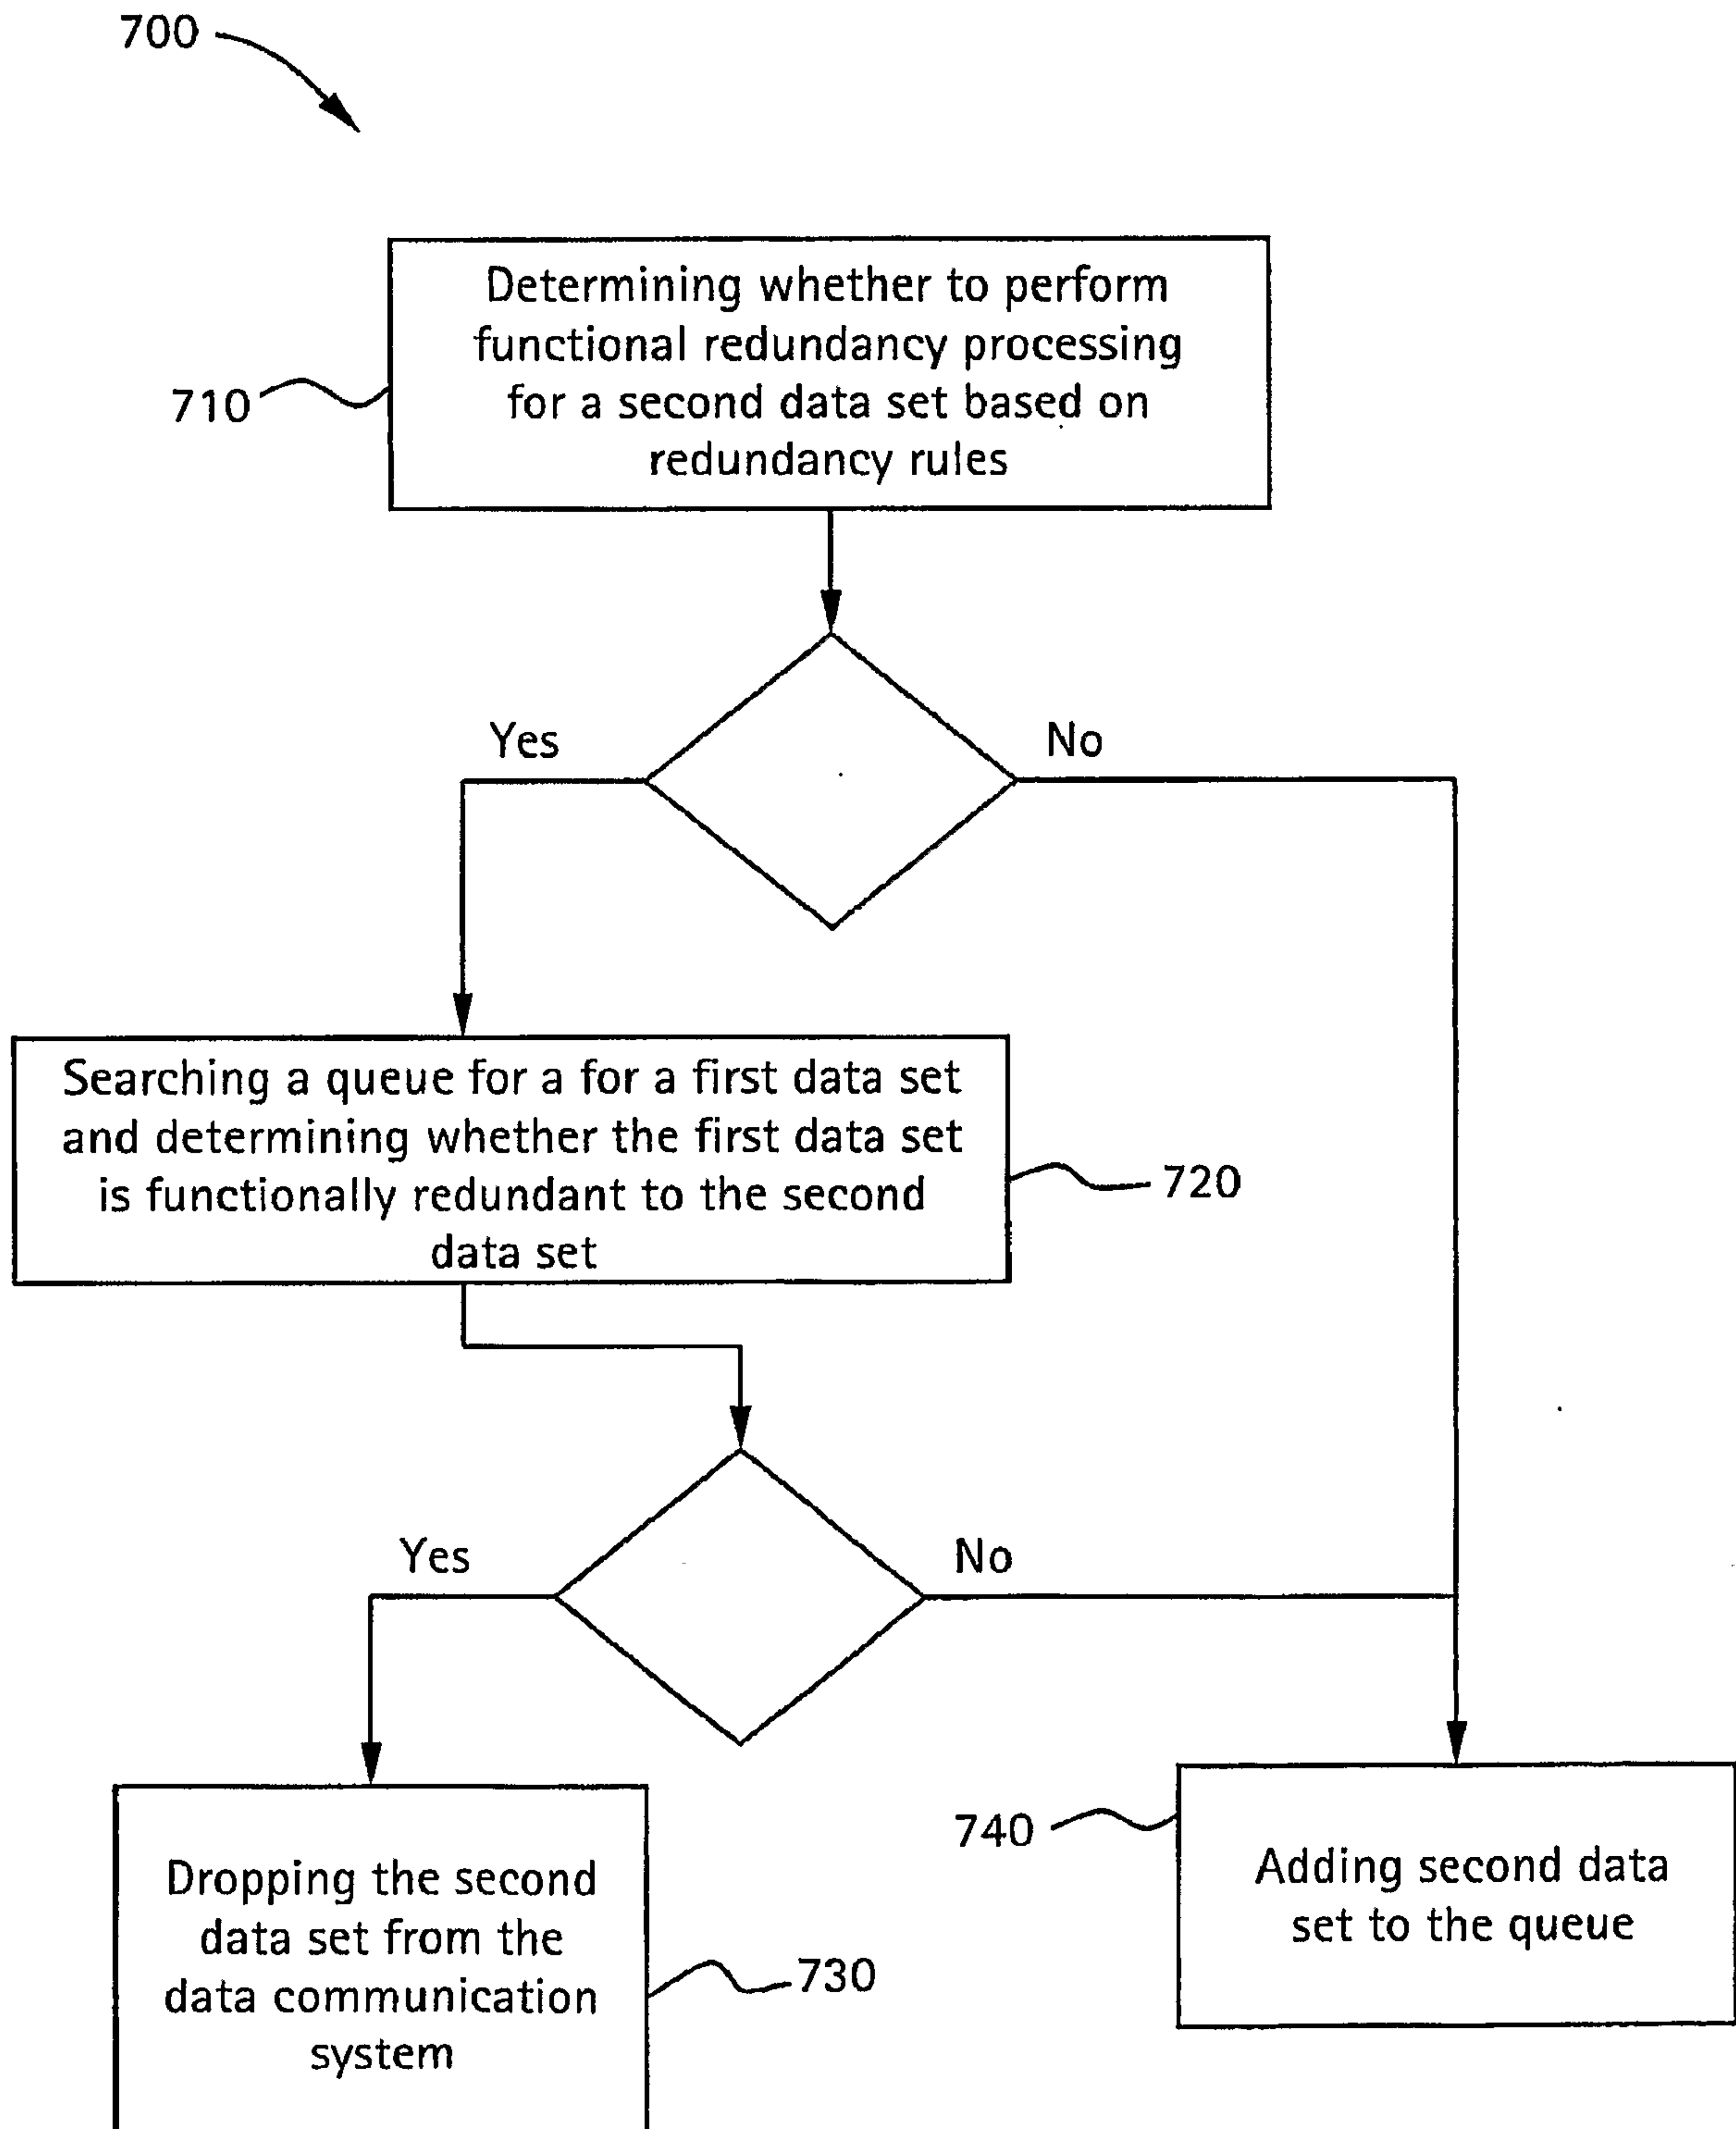
6/9

FIG. 6



7/9

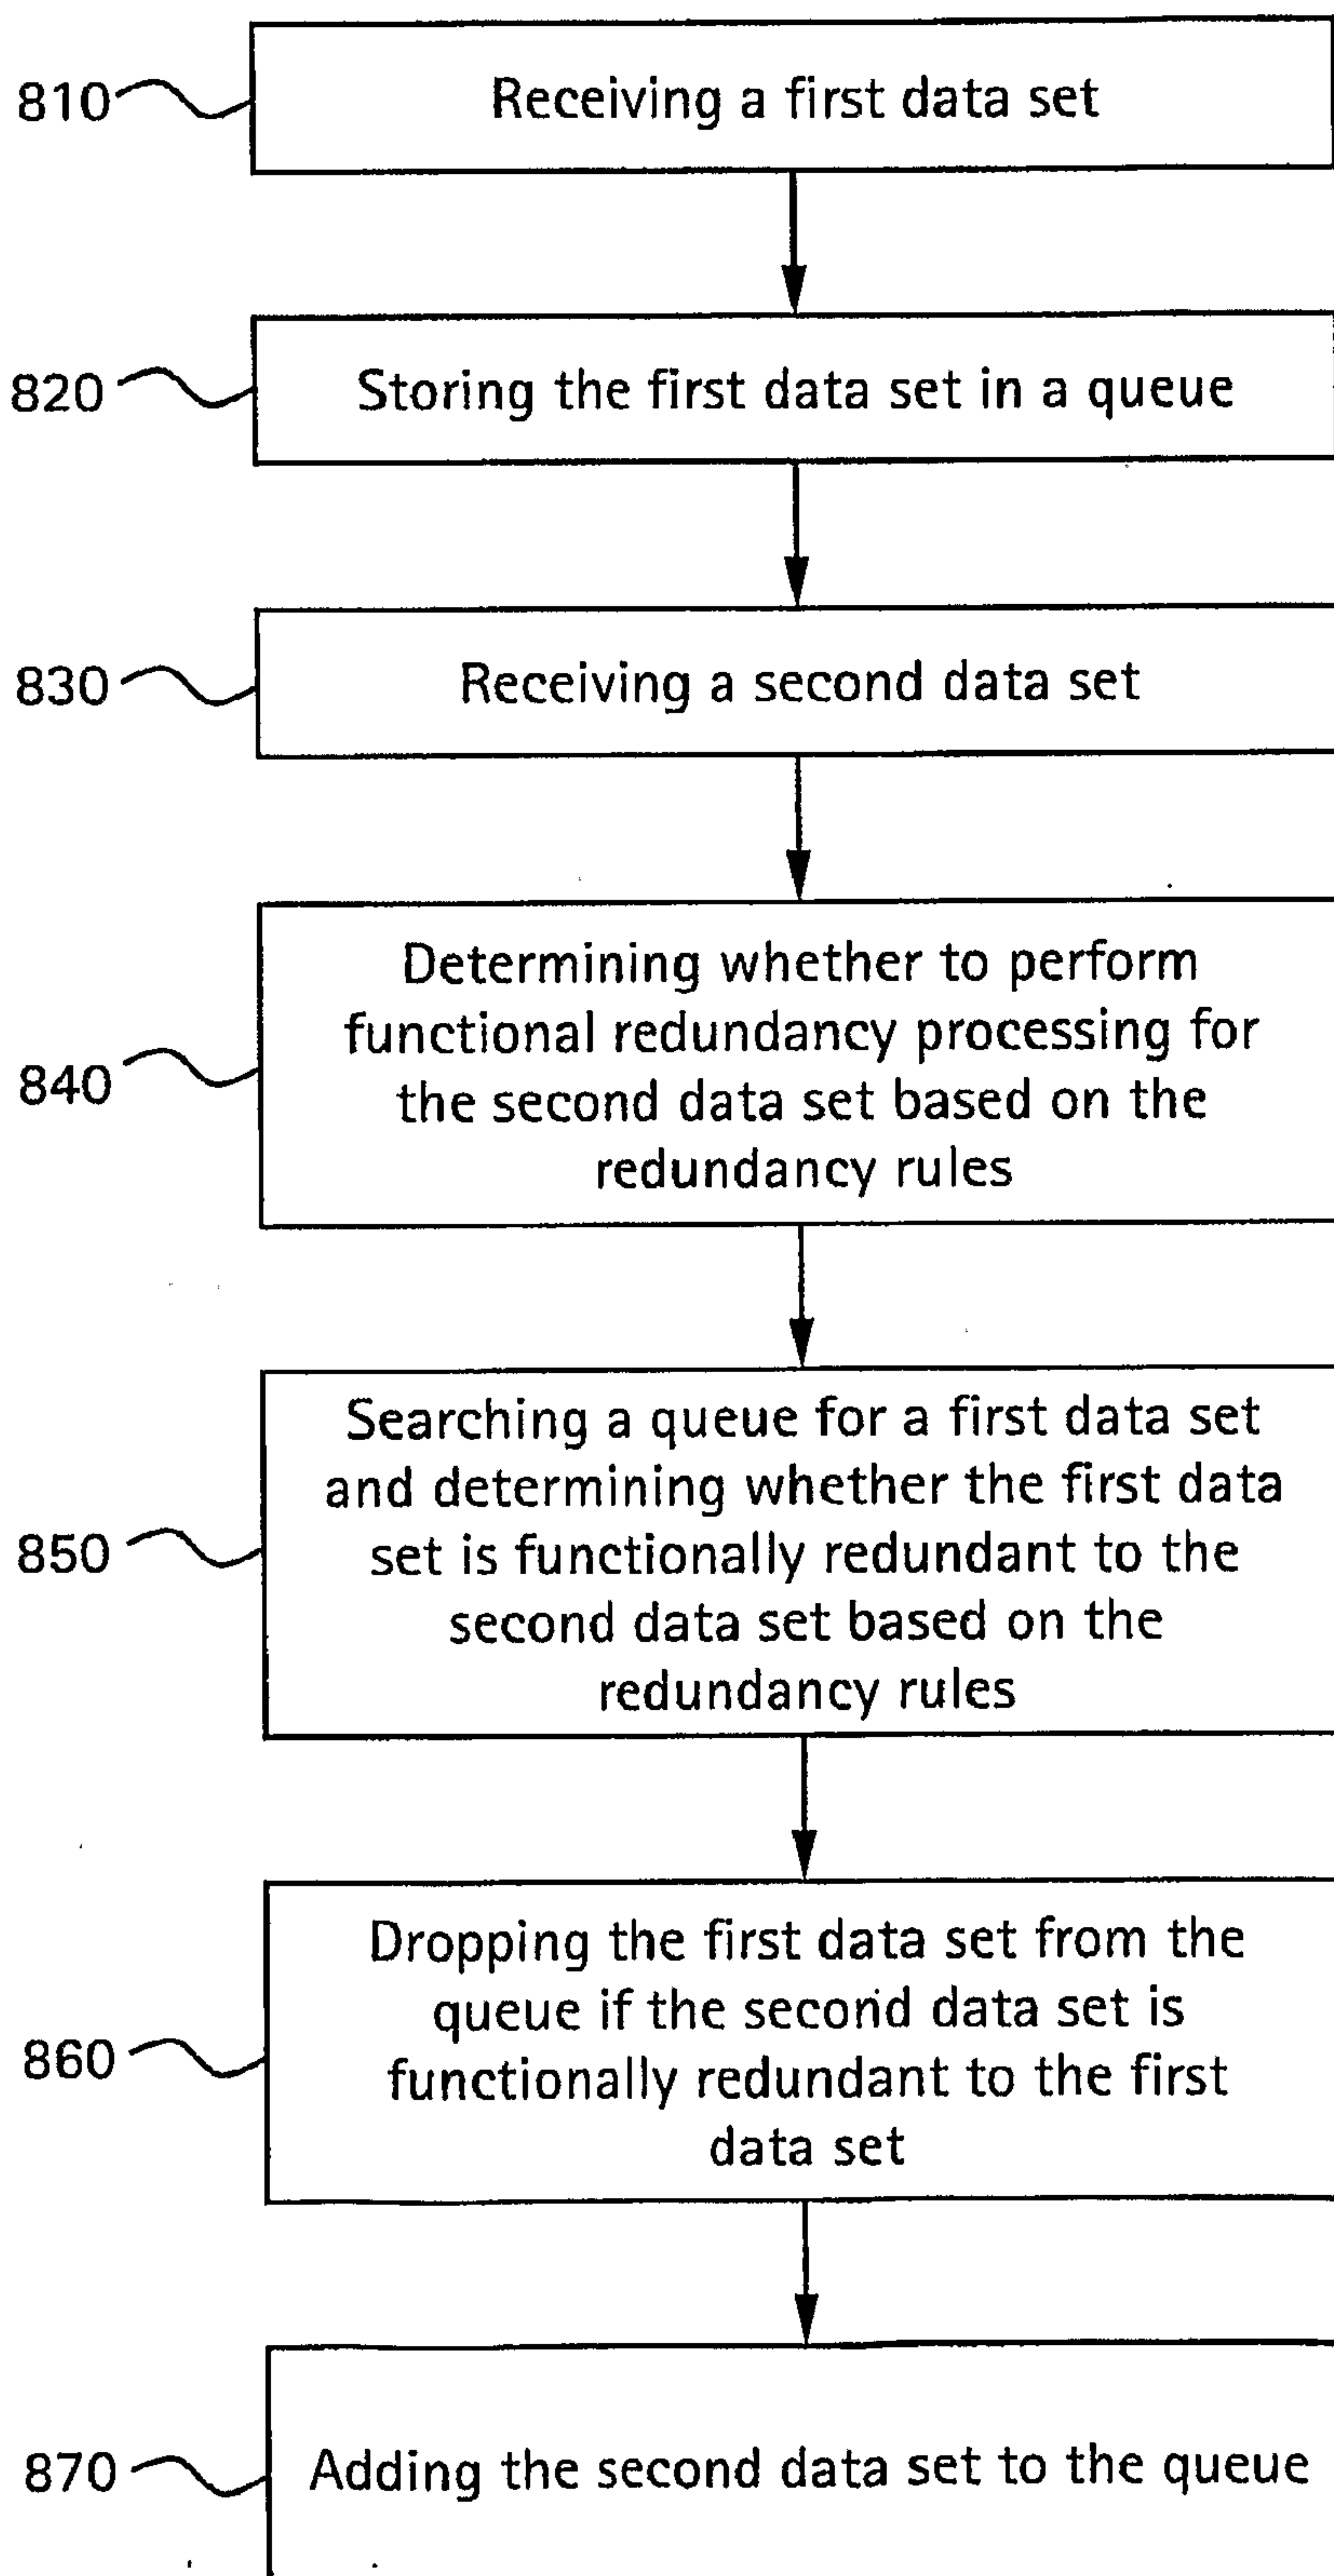
FIG. 7



8/9

FIG. 8

800



9/9

FIG. 9

900

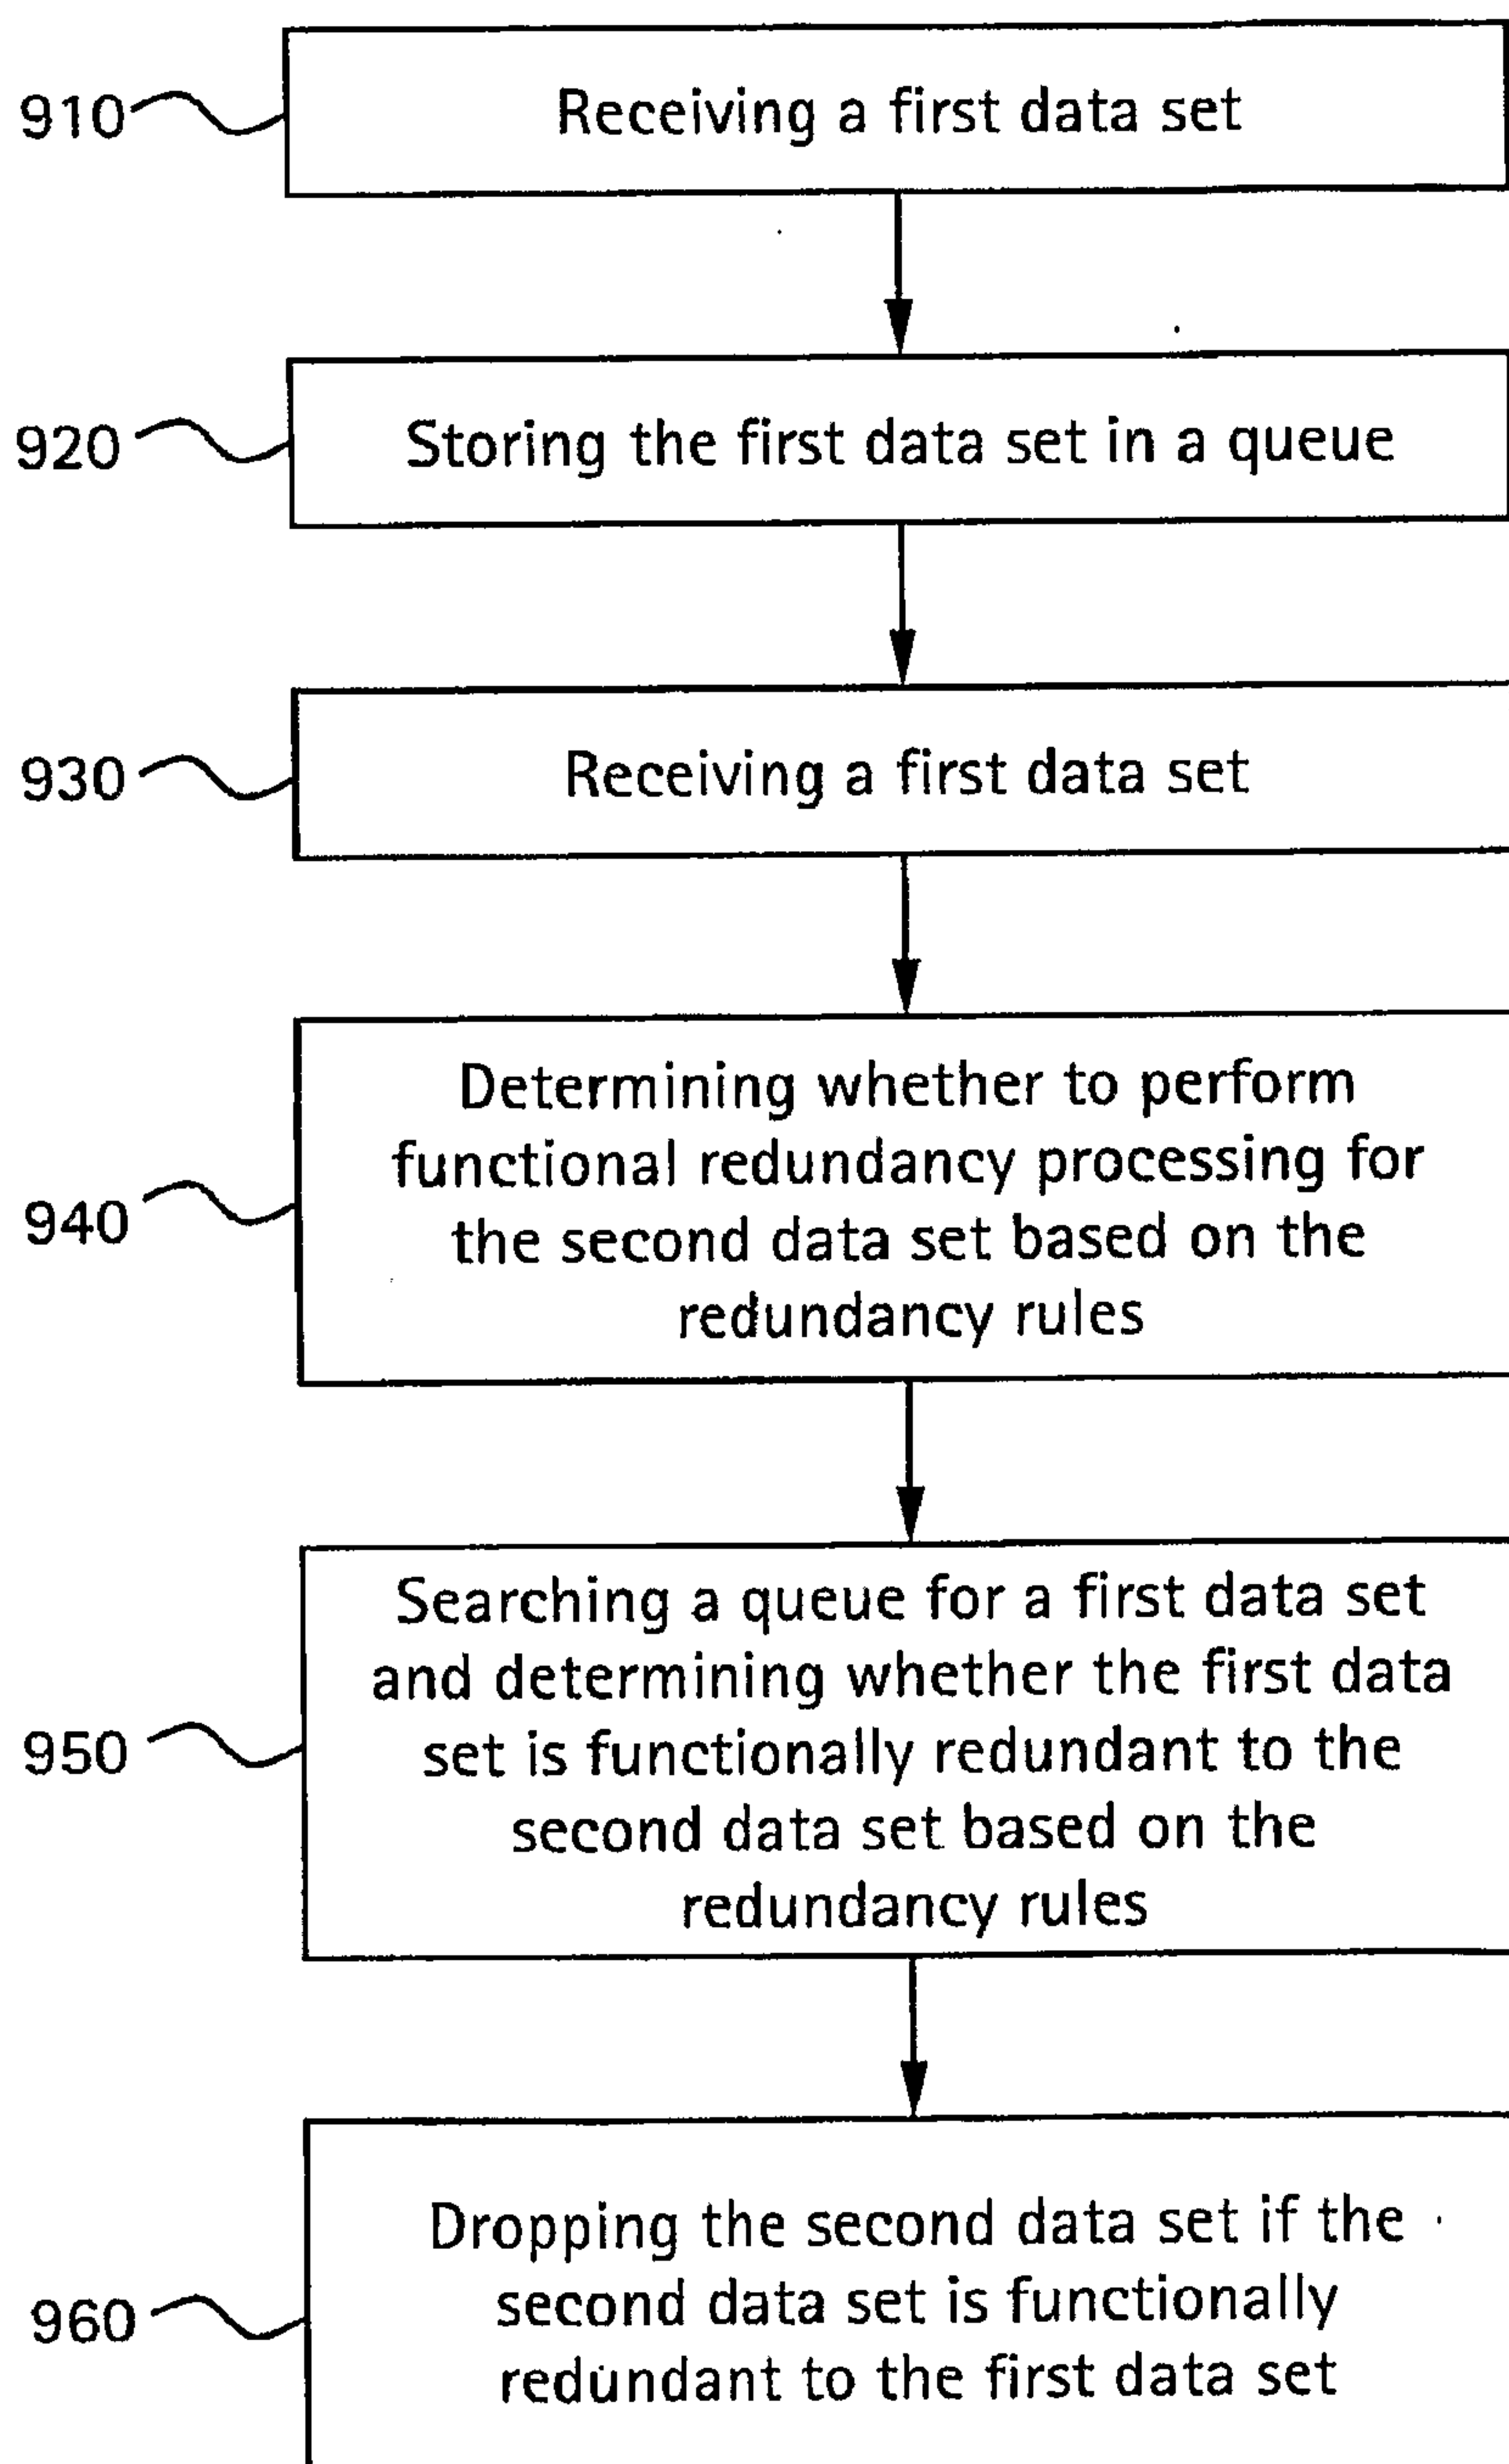


FIG. 8

800

