

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

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2013/116456 A1

(43) International Publication Date

8 August 2013 (08.08.2013)

(51) International Patent Classification:

G06F 15/173 (2006.01)

(21) International Application Number:

PCT/US2013/024039

(22) International Filing Date:

31 January 2013 (31.01.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/592,746 31 January 2012 (31.01.2012) US
13/754,398 30 January 2013 (30.01.2013) US

(71) Applicant: MASSACHUSETTS INSTITUTE OF TECHNOLOGY [US/US]; 77 Massachusetts Avenue, Cambridge, Massachusetts 02139 (US).

(72) Inventors: CALMON, Flavio, du Pin; 235 Albany Street, Room 3120, Cambridge, Massachusetts 02139 (US). CLOUD, Jason, M.; 524 Putnam Avenue, Apt. 11, Cambridge, Massachusetts 02139 (US). MÉDARD, Muriel; 45 Evergreen Way, Belmont, Massachusetts 02478 (US). ZENG, Weifei; 75 5th Street, Floor 2, Cambridge, Massachusetts 02141 (US).

(74) Agents: SCOTT, John, C. et al.; Daly, Crowley, Mofford & Durkee, LLP, 354A Turnpike Street, Suite 301A, Canton, Massachusetts 02021 (US).

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

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

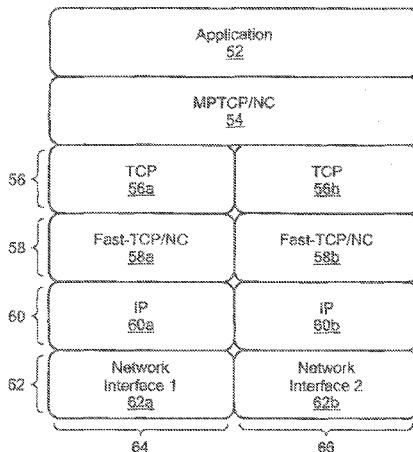
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

(54) Title: MULTI-PATH DATA TRANSFER USING NETWORK CODING

FIG. 3



(57) Abstract: Techniques, devices, systems, and protocols are disclosed herein that relate to data transfer between communication nodes via multiple heterogeneous paths. In various embodiments, network coding may be used to improve data flow and reliability in a multiple path scenario. Transmission control protocol (TCP) may also be used within different paths to further enhance data transfer reliability. In some embodiments, multiple levels of network coding may be provided within a transmitter in a multiple path scenario, with one level being applied across all paths and another being applied within individual paths.

MULTI-PATH DATA TRANSFER USING NETWORK CODING

CROSS REFERENCE TO RELATED APPLICATION

[0001] The present application claims the benefit of U.S. Provisional Patent Application No. 61/592,746 filed on January 31, 2012, which is incorporated by reference herein in its entirety.

GOVERNMENT RIGHTS

[0002] This invention was made with government support under Contract No. FA8721-05-C-0002 awarded by the U.S. Air Force and under Grant No. FA9550-09-1-0196 awarded by the Air Force Office of Scientific Research. The government has certain rights in the invention.

FIELD

[0003] Subject matter disclosed herein relates generally to data communications and, more particularly, to techniques, systems, and protocols for use in transferring data between locations via multiple different paths.

BACKGROUND

[0004] Modern communication devices (e.g., smart phones and other handheld communicators, tablet computers with communication functionality, etc.) often possess multiple network interfaces for use with different network technologies having different connectivity characteristics. The different network technologies may each have, for example, different delay, throughput, and reliability characteristics. A mobile device that is connected to a remote source node through both a 3G network and an IEEE 802.11 wireless LAN may, for example, observe different usage costs and quality of service through each interface. It would be desirable to be able to utilize multiple available network resources to carry out a data transfer operation for a communication device to, for example, increase throughput, decrease data transit time, and/or make efficient use of available resources. It would also be desirable if this could be done in a reliable and

efficient manner that takes advantage of already existing communication techniques and protocols.

SUMMARY

[0005] Techniques, devices, systems, and protocols are described herein that support efficient and reliable data transfer between communication nodes via multiple different paths involving different network technologies. A novel transport layer protocol is proposed (i.e., the Multi-Path Transfer Control Protocol with Network Coding protocol or MPTCP/NC) that may be used to coordinate several parallel connections over multiple interfaces/networks. The proposed protocol incorporates some features of the multi-path TCP (MPTCP) protocol introduced by the Internet Engineering Task Force (IETF), but builds on the protocol to allow many of the benefits of network coding to be achieved in a multi-path scenario. It has been shown through simulation that the techniques described herein are capable of achieving a significant improvement in data transfer completion times and average transfer rates over other known multi-path approaches. Simulations have also shown that the techniques are more robust and reliable to link failure and variation than other multi-path schemes.

[0006] In accordance with one aspect of the concepts, systems, circuits, and techniques described herein, a machine implemented method for use in transferring data to a destination node comprises: obtaining a plurality of original data packets to be transferred to the destination node; generating first coded packets by linearly combining original data packets using network coding; distributing the first coded packets among multiple available paths leading to the destination node; generating second coded packets by linearly combining first coded packets distributed to a first path of the multiple available paths, using network coding; and transmitting the second coded packets associated with the first path to the destination node via a network associated with the first path.

[0007] In one embodiment, the method further comprises: generating second coded packets by linearly combining first coded packets distributed to a second path of the multiple available paths, using network coding; and transmitting the second coded packets associated with the second path to the destination node via a network associated with the second path.

[0008] In one embodiment, generating first coded packets includes generating first coded packets by linearly combining original data packets that are within a sliding coding window. In one embodiment, the sliding coding window has a variable width. In one embodiment, the method further comprises: receiving acknowledgement messages from the destination node that each indicate that a new degree of freedom has been received by the destination node in connection with the data transfer; and adjusting the width of the sliding coding window based, at least in part, on received acknowledgement messages.

[0009] In one embodiment, generating second coded packets includes generating second coded packets by linearly combining first coded packets that are within a sliding coding window. In one embodiment, generating second coded packets includes generating redundant second coded packets for each set of first coded packets within the sliding coding window.

[0010] In one embodiment, each of the multiple available paths includes a transmission control protocol (TCP) layer that adds sequence numbers to first coded packets distributed to the path, wherein generating second coded packets by linearly combining first coded packets masks the sequence numbers added to the first coded packets associated with a path.

[0011] In one embodiment, each of the multiple available paths includes a transmission control protocol (TCP) layer; and distributing the first coded packets includes distributing the first coded packets among the multiple available paths based, at least in part, on TCP congestion control window dynamics.

[0012] In one embodiment, the multiple available paths are associated with different network technologies.

[0013] In one embodiment, obtaining a plurality of original data packets includes receiving the original data packets from an application layer.

[0014] In accordance with another aspect of the concepts, systems, circuits, and techniques described herein, a machine implemented method for use in processing coded packets received from a source node via multiple different paths comprises: receiving coded packets associated with a first connection via multiple different paths, the coded packets each including a linear combination of original data packets; for each coded packet associated with the first connection that is successfully received, sending an acknowledgement message to the source node without first determining whether the coded

packet is linearly independent of previously received coded packets associated with the first connection; forwarding all coded packets associated with the first connection, received from all paths, to a common processing layer without decoding the coded packets; and for each coded packet associated with the first connection forwarded to the common processing layer: (a) determining whether the coded packet is linearly independent of coded packets associated with the connection that were previously forwarded to the common processing layer; and (b) sending an acknowledgement message to the source node acknowledging that a new degree of freedom has been received for the first connection if the coded packet is determined to be linearly independent.

[0015] In one embodiment, determining whether the coded packet is linearly independent includes performing a Gauss-Jordan elimination operation on a coefficient matrix.

[0016] In one embodiment, receiving coded packets associated with a first connection via multiple different paths includes receiving coded packets from paths that use different network technologies, wherein each of the different paths uses transmission control protocol (TCP) to control a transfer of packets through the path. In one embodiment, sending an acknowledgement message to the source node without first determining whether the coded packet is linearly independent from previously received coded packets is performed as part of a TCP layer. In one embodiment, sending an acknowledgement message to the source node without first determining whether the coded packet is linearly independent of previously received coded packets is performed within a first portion of the TCP layer for coded packets received via a first path and within a second portion of the TCP layer for coded packets received via a second path that is different from the first path.

[0017] In one embodiment, the method further comprises discarding a received coded packet if it is determined that the packet is not linearly independent.

[0018] In one embodiment, receiving coded packets associated with a first connection via multiple different paths includes receiving first coded packets via a first path following an IEEE 802.16 wireless networking standard and receiving second coded packets via a second path following an LTE wireless cellular standard.

[0019] In one embodiment, forwarding all coded packets associated with the first connection to a common processing layer includes forwarding the coded packets to a

multiple path transfer control protocol with network coding (MPTCP-NC) layer that is higher than a TCP layer in a corresponding protocol architecture. In one embodiment, the protocol architecture includes separate protocol stacks associated with each of the multiple different paths.

[0020] In accordance with still another aspect of the concepts, systems, circuits, and techniques described herein, a communication device comprises: a first network interface unit configured for communication in a first network; a second network interface unit configured for communication in a second network that is different from the first network; and at least one processor to manage data transfer between the communication device and a destination node using multiple different paths, the at least one processor to: (a) obtain a plurality of original packets representative of data to be transferred to the destination node; (b) generate first coded packets by linearly combining original data packets using network coding; (c) distribute first coded packets to the multiple different paths to the destination node; and (d) generate second coded packets within individual paths by linearly combining first coded packets distributed to the path, using network coding.

[0021] In one embodiment, the at least one processor is configured to transmit second coded packets generated within individual paths to the destination node via corresponding networks.

[0022] In one embodiment, the at least one processor is configured to generate first coded packets by linearly combining original data packets within a sliding coding window.

[0023] In one embodiment, the at least one processor is configured to generate second coded packets for a first path by linearly combining first coded packets within a sliding coding window. In one embodiment, the at least one processor is configured to generate redundant second coded packets for the first path for each set of first coded packets within the sliding coding window.

[0024] In one embodiment, the at least one processor is configured to distribute the first coded packets to the multiple different paths based, at least in part, on TCP congestion control window dynamics.

[0025] In one embodiment, a first of the multiple different paths is associated with the first network interface unit and a second of the multiple different paths is associated with the second network interface unit.

[0026] In accordance with a further aspect of the concepts, systems, circuits, and techniques described herein, a communication device comprises: a first network interface unit configured for communication in a first network; a second network interface unit configured for communication in a second network that is different from the first network; and at least one processor to manage processing of coded data packets received from a source node via multiple different paths as part of a first connection, the coded data packets each including a linear combination of original data packets, the at least one processor to: (a) send an acknowledgement message to the source node for each coded data packet associated with the first connection that is successfully received; (b) forward coded data packets associated with the first connection, received from all paths, to a common processing layer without first decoding the coded packets; and (c) for each coded packet associated with the first connection forwarded to the common processing layer: (i) determine whether the coded packet is linearly independent of coded packets associated with the connection that were previously forwarded to the common processing layer; and (ii) send an acknowledgement message to the source node acknowledging that a new degree of freedom has been received for the first connection if the coded packet is determined to be linearly independent.

[0027] In one embodiment, the at least one processor is configured to determine whether the coded packet is linearly independent by performing a Gauss-Jordan elimination operation on a coefficient matrix.

[0028] In one embodiment, the at least one processor is configured to discard a received coded packet if it is determined that the packet is not linearly independent.

[0029] In one embodiment, a first of the multiple different paths is associated with the first network interface unit and a second of the multiple different paths is associated with the second network interface unit.

BRIEF DESCRIPTION OF THE DRAWINGS

[0030] The foregoing features may be more fully understood from the following description of the drawings in which:

[0031] Fig. 1 is a diagram illustrating an exemplary multiple path communication scenario that can utilize one or more of the techniques or features described herein;

- [0032] Fig. 2 is a diagram illustrating another exemplary multiple path communication scenario that can utilize one or more of the techniques or features described herein;
- [0033] Fig. 3 is a diagram illustrating an exemplary protocol architecture that may be used within a communication system having heterogeneous links in accordance with an embodiment;
- [0034] Fig. 4 is a diagram illustrating an exemplary communication scenario involving a transfer of data between a source node and a destination node in accordance with an embodiment;
- [0035] Fig. 5 is a diagram illustrating operation of an exemplary process for encoding data within a source node in accordance with an embodiment;
- [0036] Fig. 6 is a flowchart illustrating an exemplary method for use in transmitting data from a source node to a destination node via multiple different paths in accordance with an embodiment;
- [0037] Figs. 7a and 7b are portions of a flowchart illustrating a method for use in processing coded packets received at a destination node via multiple different paths in accordance with an embodiment;
- [0038] Fig. 8 is a diagram illustrating a round duration approach used for analyzing operation of an MPTCP/NC protocol;
- [0039] Fig. 9 is a diagram illustrating round trip time and packet loss probability parameters for several different wireless standards;
- [0040] Figs. 10a-10c are diagrams illustrating analysis results for three considered models using two sub-flows having parameters $RTT_1 = 15\text{ms}$, $p_1 = 10^{-3}$, $RTT_2 = 80\text{ms}$, and $p_2 = 10^{-2}$;
- [0041] Figs. 11a-11c are diagrams illustrating analysis results for three considered models using two sub-flows having parameters $RTT_1 = 1500\text{ms}$, $p_1 = 10^{-2}$, $RTT_2 = 80\text{ms}$, and $p_2 = 10^{-1}$;
- [0042] Figs. 12a-12c are diagrams illustrating analysis results for three considered models using two sub-flows having parameters $RTT_1 = 1500\text{ms}$, $p_1 = 10^{-3}$, $RTT_2 = 15\text{ms}$, and $p_2 = 0.3$;

[0043] Figs. 13a-13c are diagrams illustrating analysis results for three considered models using two sub-flows having parameters $RTT_1 = 80\text{ms}$, $p_1 = 10^{-1}$, $RTT_2 = 250\text{ms}$, and $p_2 = 10^{-1}$; and

[0044] Fig. 14 is a block diagram illustrating an exemplary node device architecture that may be used in a node device incorporating features described in the present disclosure in one or more embodiments.

DETAILED DESCRIPTION

[0045] Techniques, devices, systems, protocols described herein relate to the transfer of data between communication nodes via multiple heterogeneous paths. In various embodiments, network coding is used to improve data flow and reliability within a multiple path scenario. Transmission control protocol (TCP) may also be implemented within each of the different paths to further enhance data transfer reliability. In some embodiments, multiple levels of network coding may be provided within a transmitter in a multiple path networking scenario, one level that is applied across all of the different paths (or sub-flows) and another level that is applied within each individual path. As will be described in greater detail, a communication protocol (MPTCP/NC) is also described for use in implementing and standardizing these techniques in some embodiments.

[0046] Fig. 1 is a diagram illustrating an exemplary multiple path communication scenario 10 that may utilize one or more of the techniques or features described herein. As shown, a first node 12 (e.g., a smart phone) may wish to exchange information with a second node 14 (e.g., a server) via a network cloud 16. As is known, a network cloud 16 may provide multiple different routes or paths that data packets can traverse between the first and second nodes 12, 14. In addition, some of the different routes through the network cloud 16 may involve different network technologies in some cases. For example, in the illustrated embodiment, a first path 18 may exist between the first and second nodes 12, 14 that is associated with a WiFi-based network 22 (i.e., based on the IEEE 802.11 family of wireless networking standards) and a second path 20 may exist that is associated with a long term evolution (LTE) based 4G cellular network 24. To use either of these different paths 18, 20, the first and second nodes 12, 14 would have to include (or have access to) appropriate hardware and software to be able to communicate

according to the corresponding standards. Many modern communication devices (e.g., cell phones, smart phones, handheld communicators, computers or servers with communication functionality, network gateways, wireless access points, audio/video devices with communication functionality, appliances with communication functionality, wearable communication devices, etc.) include network adapters for multiple different network technologies.

[0047] Fig. 2 is a diagram illustrating another exemplary multiple path communication scenario 30 that may utilize one or more of the techniques or features described herein. As shown, an aircraft 32 may act as a communication platform or node that is capable of supporting wireless communication in accordance with a number of different technologies, specifications, and/or standards (military and/or commercial). Aircraft 32 may desire to perform a data transfer operation with a remote terrestrial server 34. One or more communication satellites 36, 38 may be within communication range of aircraft 32. Each of these satellites 36, 38 may provide a corresponding path 42, 44 through which aircraft 32 can communicate with server 34. One or more ground stations 40 may also be within communication range of aircraft 32. Ground station 40 may also provide a path 48 through which aircraft 32 can establish communication with remote server 34. The satellites 36, 38 and the ground station 40 may each operate in accordance with a different wireless standard.

[0048] In each of the above-described scenarios, it may be possible to use multiple different network technologies to support a single data transfer operation (or a single connection) between communication nodes. The use of multiple different network technologies associated with multiple different paths to support a single data transfer operation or other function may be referred to as “heterogeneous networking.” As used herein, phrases such as “network technologies” and “network access technologies” encompass both wireless and wired technologies and can involve various different wireless and/or wire line standards, specifications, and/or protocols. The different wireless standards may include, for example, wireless networking standards, wireless cellular standards, wireless satellite communication standards, infrared communications standards, point to point microwave link technologies, LMDS and MMDS technologies, and/or other forms of wireless standards. Although specific heterogeneous network scenarios are shown in Figs. 1 and 2, it should be appreciated that the number and type of

network technologies used in particular heterogeneous network implementations may vary widely. As described above, various techniques, devices, and systems are described herein that allow heterogeneous network arrangements to be used to facilitate data transfer between two communicating nodes in a system in an efficient and reliable manner.

[0049] Transmission control protocol (TCP) is a communication protocol that may be used to provide reliable data stream service between, for example, a source node and a destination node in a network. TCP makes use of positive acknowledgment with retransmission to improve the reliability of data transfer in both accuracy and sequence. A TCP source node may keep a record of each packet transmitted as part of a data transfer operation. After transmitting a packet, the source node may automatically retransmit the packet if an acknowledgement message is not received from the destination node within a predetermined time period. TCP is considered a part of the TCP/IP protocol suite and provides a communication service at an intermediate layer between an application program and the Internet Protocol (IP).

[0050] Internet Protocol (IP) typically uses “IP packets” as a vehicle to transfer data between nodes on a physical medium. An IP packet is a sequence of octets that includes both a header and a body. Among other things, the header of an IP packet may include destination node address information to identify a destination of the IP packet. The body of an IP packet may include the data being transferred. When an application program within a first node desires to send a large amount of data to a remote node, the program can split up the data into smaller portions and then issue a series of IP requests to an IP layer to transfer the data. Alternatively, the program can issue a single TCP request and let TCP deal with the IP layer. For various reasons, IP packets can be lost, duplicated, or delivered out of sequence in a network. TCP can detect such problems, request retransmission of lost data, and rearrange packets in a receiver that were delivered out of sequence. TCP can also operate to reduce congestion in a network. Once a TCP-enabled receiver has received and reassembled a number of octets, it may pass them to an application program for use in performing a particular function. The operation of TCP can be relatively transparent to an application.

[0051] Network coding is a coding technique that may be used to, among other things, improve information flow in a network. Network coding usually involves the generation of linear combinations of data packets to form “coded packets” for transmission in a

network. One form of network coding, known as Random Linear Network Coding (RLNC), uses randomly generated coefficients to form linear combinations of packets (or coded packets) to be transmitted in the network. In some cases, information describing the coefficients used to make the linear combination may be appended to the coded packet before transmission. RLNC has been shown to be a powerful technique for achieving robust, high throughput packet distribution in certain network environments. One example system and method for implementing RLNC is described in U.S. Patent No. 7,706,365 to Effros et al. entitled "Randomized Distributed Network Coding," which is hereby incorporated by reference in its entirety.

[0052] In a network that uses RLNC, a source node that has a number of original data packets to transfer to a destination node may combine the packets together into a "coded packet" that is then transmitted into the network. A destination node in the network that receives the coded packet may then store the packet for eventual decoding. The destination node will typically require a particular number of coded packets before it is able to decode the coded packets to extract the original data packets. The decoding process may involve, for example, the solution of a system of linear equations having a number of independent variables (i.e., the data packets) and a number of independent equations. Before a destination node can reliably solve for the original data packets, the number of "degrees of freedom" (or DOFs) received must equal or exceed the number of original packets. The DOFs missing may in some cases be defined as the number of unknown independent variables of a system of equations less the number of independent equations. In some implementations, each coded packet successfully received at a destination node may provide another equation for use in decoding and thus reduce the number of DOFs needed by the node by one (assuming the coded packet is linearly independent of previously received coded packets). In some networks, a source node may transmit a stream of coded packets until a destination node is able to decode the original data packets.

[0053] In developing the techniques, devices, systems, and protocols described herein, it was determined that improved throughput performance and decreased communication delay could be achieved in a network by intelligently combining TCP, network coding, and heterogeneous networking. The techniques, devices, systems, and protocols are also capable of increasing data transfer reliability and decreasing management overhead when

transmitting a single source-destination flow over multiple interfaces/networks. Although described herein in connection with network coding, it should be appreciated that some or all of the techniques described herein may use other similar codes in place of, or in addition to, network coding in some implementations (e.g., fountain codes, systematic codes, structured codes, etc.).

[0054] Fig. 3 is a diagram illustrating an exemplary protocol architecture 50 that may be used within a communication system having heterogeneous links in accordance with an embodiment. The corresponding protocol will be referred to herein as multi-path TCP with network coding (MPTCP-NC). Protocol architecture 50 may be used to split a data transfer operation involving a source node and a destination node between multiple different heterogeneous communication links (or paths or sub-flows) in certain situations. As will be described in greater detail, protocol architecture 50 may also be used to provide two separate levels of network coding to signals being transferred to a destination node via multiple heterogeneous links. The phrases “data flow” and “connection” may be used herein to refer to a data transfer operation that seeks to transfer a certain amount of data (e.g., a data file of a particular size, etc.) from a source node to a destination node.

[0055] As shown in Fig. 3, protocol architecture 50 may include: an application layer 52, an MPTCP/NC layer 54, a TCP (or transport) layer 56, a fast-TCP/NC layer 58, an IP layer 60, and a network interface layer 62. Application layer 52, TCP layer 56, IP layer 60, and network interface layer 62 may perform similar functions to corresponding layers in the basic TCP/IP protocol stack. The MPTCP/NC layer 54 may sit directly above TCP layer 56 in protocol architecture 50. Fast-TCP/NC layer 58 may sit directly below TCP layer 56. Application layer 52 is the highest layer in protocol architecture 50 and is the layer where one or more application programs may be executed that may require one or more communications-related processes to be performed in a corresponding system. MPTCP/NC layer 54 is operative for performing one level of network coding operations for the system and for dividing data transfer operations between the different paths, or sub-flows 64, 66. TCP layer 56 is operative for implementing transmission control protocol within the communication system. Fast-TCP/NC layer 58 is operative for performing a second level of network coding operations for the system and for simplifying the implementation of network coding within the TCP framework. IP layer 60 is operative for providing internet protocol (IP) related services in the communication system. Network

interface layer 62 is operative for providing an interface to a physical network medium for the system.

[0056] As illustrated, the lower portion of protocol architecture 50 is divided into a number of different “sub-flows” 64, 66 that may each correspond to a different networking technology and a different path between a source node and a destination node. Each sub-flow 64, 66 within protocol architecture 50 may include a corresponding protocol stack. For example, first sub-flow 64 may include a stack with a TCP layer 56a, Fast-TCP/NC layer 58a, IP layer 60a, and network interface layer 62a. Likewise, second sub-flow 66 may include a stack with a TCP layer 56b, Fast-TCP/NC layer 58b, IP layer 60b, and network interface layer 62b. The protocol stack associated with each sub-flow 64, 66 may be used to process communications flowing through the sub-flow 64, 66 for a corresponding networking technology. Thus, if the first sub-flow 64 corresponds to an IEEE 802.11 based network, TCP layer 56a will be used to provide TCP services for the IEEE 802.11-based communications and network interface 62a will be used to interface to a corresponding IEEE 802.11 wireless medium, and so on. Although illustrated in Fig. 3 with only two sub-flows, it should be appreciated that protocol architecture 50 may be implemented with any number of different sub-flows (greater than one) in different implementations.

[0057] In a typical communication scenario, an application program within application layer 52 may determine that a large amount of data (e.g., a large data file, etc.) needs to be transferred to a remote destination node. The application program may then transfer some or all of the data to MPTCP/NC layer 54 with instructions to transfer the data to the destination node. MPTCP/NC layer 54 may then apply network coding to the data and then divide the coded data amongst the various sub-flows 64, 66 for transmission to the destination node. Thus, instead of seeing multiple TCP links, application layer 52 only needs to interface with MPTCP/NC layer 54.

[0058] In one implementation, the data transferred from application layer 52 to MPTCP/NC layer 54 will include data packets that will be referred to herein as “original packets.” As will be described in greater detail, MPTCP/NC layer 54 may process original packets received from application layer 52 using network coding to generate “coded packets.” As described previously, coded packets may be generated by forming linear combinations of original packets using, for example, randomly generated

coefficients. In one approach, MPTCP/NC layer 54 may generate a new coded packet by linearly combining all original packets that are within a coding buffer of MPTCP/NC layer 54 (i.e., all original packets received so far that are associated with the corresponding connection). In another approach, MPTCP/NC layer 54 may generate a new coded packet by linearly combining all original packets that are within a sliding coding window of MPTCP/NC layer 54. Each new original packet added to the coding window may drive an older original packet out of the window. The coded packets generated by MPTCP/NC layer 54 will be referred to herein as “first coded packets” to distinguish them from subsequently generated coded packets. The sliding coding window of MPTCP/NC layer 54 may have a fixed or variable length in different implementations.

[0059] In the discussion that follows, it will be assumed that two or more TCP sub-flows are properly established and managed under a common connection. In at least one implementation, the procedure for initiating and managing multiple TCP sub-flows will be the same as, or similar to, the procedures described in the MPTCP protocol proposed by the IETF (e.g., *MPTCP Application Interface Considerations*, by Scharf et al., draft-ietf-mptcp-api-06, 2012-10-22, which is incorporated by reference herein in its entirety), although other processes may alternatively be used. MPTCP/NC layer 54 may monitor the different TCP sub-flows 64, 66 of protocol architecture 50 to determine when a particular sub-flow is able to receive one or more first coded packets. As TCP transmits packets over the network, its window slides or expands to include new packets. When there is a chance to inject a new coded packet to one of the TCP sub-flows 64, 66, the MPTCP/NC layer 54 may produce a new random linear network coded packet (i.e., a new “first coded packet”) and inject the new coded packet to the identified TCP sub-flow.

[0060] As described above, in one possible scenario, the MPTCP/NC layer 54 may decide to inject a new coded packet into, for example, sub-flow 64. The TCP layer 56a of sub-flow 64 may operate in a conventional TCP manner as all operations in MPTCP/NC layer 54 and Fast-TCP/NC layer 58a may be transparent to the TCP protocol. However, it should be noted that all packets in the TCP window of TCP layer 56a of protocol architecture 50 will be coded packets (i.e., first coded packets). These packets are pushed from TCP layer 56a to Fast-TCP/NC layer 58a, which maintains another coding window. Fast-TCP/NC layer 58a may then apply a second level of network coding to first coded packets in its coding window to generate second coded packets. For each first coded

packet TCP layer 56a sends to Fast-TCP/NC 58a, the latter may produce R second coded packets that are random linear combinations of all first coded packets in its coding window. The second coded packets are eventually passed to IP layer 60a, after which they may be transmitted over a corresponding network (via network interface layer 62a) to a destination node or nodes. Thus, the transmitted packets are twice-encoded packets.

[0061] Although packets received by TCP layer 56 are previously coded by MPTCP/NC layer 54, TCP will not be aware of this and will assign sequence numbers to the packets. In a conventional system, the sequence numbers will allow the destination node to determine whether any of the transmitted packets were lost in transit. Due to the nature of network coding, assuming the network coding window is designed correctly, sequence number are no longer necessary to identify packets lost in transit. That is, once enough packets have been successfully received in the destination node, the node will be able to decode the original packets, regardless of the sequence in which the coded packets were transmitted. Thus, the second layer of coding provided by Fast-TCP/NC layer 58a will serve to mask the sequence numbers of the packets added by TCP layer 56a as they are not needed.

[0062] Fig. 4 is a diagram illustrating an exemplary communication scenario 80 involving a transfer of data between a source node 82 and a destination node 84 that both utilize protocol architecture 50 of Fig. 3 in accordance with an embodiment. As shown, a first network 86 provides a first link (or path) between the nodes 82, 84 and a second network 88 provides a second link (or path) between the nodes 82, 84. These links may be heterogeneous. Referring now to Fig. 4, at the destination node 84, if a twice-encoded packet is successfully received at the Fast-TCP/NC layer 58 (i.e., it is not erased in transit), the layer may immediately send an acknowledgement (ACK) message to the TCP/NC layer 58 of the source node 82. The Fast-TCP/NC layer 58 of destination node 84 may then pass the received packet up the protocol stack. All of the packets delivered by the various TCP sub-flows will eventually reach the MPTCP/NC layer 54 of destination node 84. The MPTCP/NC layer 54 may then perform Gaussian-Jordan elimination (or a similar procedure) progressively as it receives new packets from the different sub-flows.

[0063] A received packet is considered “innovative” if it is linearly independent of all previously received packets associated with the connection. If innovative, the packet will

contribute a new “degree of freedom” and induce a new “seen packet” after Gaussian-Jordan elimination. Once the MPTCP/NC layer 54 receives an innovative packet from any of the sub-flows, it sends an ACK message (i.e., DATA_ACK) to acknowledge the new degree of freedom to the MPTCP/NC layer 54 of source node 82. Eventually, the MPTCP/NC layer 54 of the destination node 84 collects enough linearly independent packets to allow it to solve the linear system and recover all original packets. The original packets may then be delivered to the application level 52 of the destination node 84.

[0064] As described above, the Fast-TCP/NC layer 58 is directly below the transport layer 56. In the Fast-TCP/NC layer 58 of destination node 84, received packets are not checked for linear independence. That is, Gaussian-Jordan elimination is not performed and the seen packet concept is not employed. Instead, for each packet successfully received, regardless of whether it is linearly independent of previously received packets, Fast-TCP/NC layer 58 will automatically send an acknowledgement message back to source node 82. This is because in any particular sub-flow, TCP layer or below, there is no global knowledge of the linear space spanned by the packets received from all sub-flows. Therefore, a linear dependency check will not be efficient at the Fast-TCP/NC layer 58. Instead, Fast-TCP/NC layer 58 may deliver packets upwards to MPTCP/NC layer 54, which does have global knowledge of all packets, to handle the linear independency check and decoding. Furthermore, assuming that the coefficients in the packet header are properly organized and adjusted, the results of the second encoding are just other random linearly encoded packets. These packets can be decoded at the MPTCP/NC layer 54 of destination node 84, without intermediate decoding at the lower level. Since the Fast-TCP/NC layer 58 does not handle linear dependency issues, it may send an ACK every time a packet is received. This is reasonable as packet linear dependency is a problem associated with the coding operation, rather than with link congestion. The fast ACK allows the window management to be more efficient.

[0065] The MPTCP/NC layer 54 may maintain a coding window which decides the number of original packets to be coded and pushed to the TCP sub-flows 64, 66. As mentioned above, MPTCP/NC layer 54 has global knowledge of all packets received from all TCP sub-flows. Therefore, Gaussian-Jordan elimination may be carried out here, whenever a packet is received, to check linear dependency. The MPTCP/NC layer 54 of source node 82 receives a DATA_ACK for each degree of freedom received at destination

node 84 and may adjust its coding window accordingly. The design of the coding window is very flexible, as is not constrained by the TCP congestion window. For example, the left boundary of the window could slide according to the slowest TCP sub-flow, while the right side of the window can slide when the fastest sub-flow contributes a degree of freedom. This will allow packets delivered by any sub-flow to be innovative with high probability.

[0066] Fig. 5 is a diagram illustrating operation of an exemplary process 100 for encoding data within a source node in accordance with an embodiment. As shown, a data file to be transferred to a remote destination node may be comprised of (or divided into) a number of original packets 102 (i.e., p_1 to p_s). In the illustrated scenario, it is assumed that the MPTPC/NC coding window is currently four packets wide and the fast-TCP/NC coding window is 2 packets wide. As shown, the first four original packets (p_1 - p_4) may first be shifted into the MPTPC/NC coding window and used to generate a coded packet 104 (c_1) by forming a linear combination using randomly generated coefficients. As shown, it may be decided to inject c_1 into a first sub-flow 112 (based on, for example, TCP congestion control window dynamics or some other criterion). A new original packet p_5 may then be shifted into the coding window, thus pushing the first original packet p_1 out of the window. A next coded packet 106 (c_2) may then be generated. It may be decided to inject c_2 into a second sub-flow 114. A similar process may then be used to generate a coded packet 108 (c_3) and coded packet 110 (c_4). As shown, it may be decided to inject coded packet c_3 into first sub-flow 112 and coded packet c_4 into second sub-flow 114.

[0067] Within the first sub-flow 112, a number of redundant coded packets 116, 118 may be generated using coded packet c_1 and coded packet c_3 (and randomly generated coefficients). The redundant packets are generated to overcome potential packet losses in transit. Likewise, within the second sub-flow 114, a number of redundant coded packets 120, 122, 124 may be generated using coded packet c_2 and coded packet c_4 (and randomly generated coefficients). Once generated, the coded packets 116, 118, 120, 122, 124 may be transmitted to a remote destination node via corresponding network media. It should be noted that the number of redundant coded packets generated within first sub-flow 112 does not have to be the same as the number generated within second sub-flow 114 (although it could be the same in some scenarios).

[0068] Fig. 6 is a flowchart illustrating a method 130 for use in transmitting data from

a source node to a destination node via multiple different paths (or sub-flows) in accordance with an embodiment. As described previously, the multiple different paths may be associated with different communication technologies in some implementations to form a heterogeneous networking arrangement. First, a number of original data packets may be obtained that are representative of the data to be transferred (step 132). For example, the data to be transferred may comprise a large data file and the original data packets may be generated by splitting up the file into smaller components. In some embodiments, the division of a larger file into original packets may be performed in an application layer or an MPTCP/NC layer of an underlying protocol. First coded data packets may then be generated by combining some or all of the original data packets using network coding (block 134). In some implementations, a first coding window (e.g., a MPTCP/NC coding window, as described previously) may be used to generate the first coded packets. The first coding window may have a particular packet width (e.g., four packets wide, etc.) and a first coded packet may be generated by forming a linear combination of the original packets currently within the first coding window, using randomly generated coefficients. A random number generator (or similar functionality) may be used to generate the random coefficients. As each new original packet is added to the first coding window, an oldest original packet in the window may be shifted out. A new first coded packet may be generated for each new original packet added to the first coding window in some implementations. The size of the first coding window may be fixed or variable.

[0069] The multiple different paths may be monitored to identify one or more paths that can receive a new first coded packet at a particular point in time (136). This monitoring process may continue until all of the desired data has been successfully transferred to and decoded at the destination node. In some embodiments, the monitoring process may be based upon the dynamics of a TCP congestion control window. When a path is identified, one or more first coded packets may be injected into the identified path (block 138). The injected packet may then be added to a second coding window (e.g., a Fast TCP/NC coding window, etc.) associated with the identified path. Over time, different first coded packets may be injected into different paths using the above-described technique. For each path, one or more second coded packets may be generated by linearly combining the first coded packets that currently reside within a corresponding second coding window using network coding (block 140). Randomly generated coefficients may

be used to perform the linear combination, as described above.

[0070] In some implementations, a number of redundant second coded packets may be generated for each set of first coded packets within the second coding window associated with a particular path. As before, for each new first coded packet added to a second coding window, an older first coded packet may be shifted out of the window. A new set of redundant packets may be generated for each new first coded packet added to a second coding window in some implementations. The redundant packets may be transmitted to the destination node via a corresponding network medium after they have been generated (block 142). The process of generating second coded packets within each path may be ongoing until all of the desired data has been successfully transferred to, and decoded at, the destination node. As described previously, the generation of second coded packets using first coded packets may mask any sequence numbers added to the first coded packets by corresponding TCP functionality.

[0071] Figs. 7a and 7b are portions of a flowchart illustrating a method 160 for use in processing coded packets received at a destination node from a source node via multiple different paths in accordance with an embodiment. The method 160 may be used, for example, to process second coded packets generated in a source node using method 130 of Fig. 6 or a similar method. Referring now to Fig. 7a, coded packets that are associated with a first connection (or a first data transfer operation) are received at a destination node via multiple different paths (block 162). For each coded packet successfully received, an acknowledgement (ACK) message may be sent to the source node acknowledging the receipt of the packet (block 164). The ACK message may be sent without first determining whether the coded packet is linearly independent of previously received coded packets associated with the first connection. All coded packets associated with the first connection that are received via all active paths may be forwarded to a common processing layer (e.g., an MPTCP/NC layer or other layer common to all paths, etc.) (block 166). The common processing layer may watch for newly forwarded coded packets associated with the first connection (block 168).

[0072] Referring now to Fig. 7b, when a new coded packet has been forwarded to the common processing layer, it may be next be determined whether the new coded packet is linearly independent of previously received coded packets associated with the first connection (block 170). To determine whether the new packet is linearly independent, a

Gaussian-Jordan elimination operation (or similar process) may be performed on a coefficient matrix associated with previously received packets. If the packet is not linearly independent (block 170-N), the packet may be discarded (block 180) and the common processing layer may proceed to look for a new coded packet (block 168). If the packet is linearly independent (block 170-Y), then the common processing layer may send an ACK message (e.g., DATA_ACK, etc.) to the source node indicating that a new degree of freedom has been received for the first connection (block 172). The source node may use these degree of freedom related ACK messages to, for example, modify a coding window size within the source node.

[0073] After the ACK message has been transmitted, it may next be determined whether enough linearly independent coded packets have been successfully received for the first connection to be able to decode the received coded packets to extract the original packets (block 174). If enough linearly independent coded packets have not been received to allow decoding (block 174-N), the common processing layer may proceed to look for a new coded packet (block 168). If enough linearly independent packets have been received (block 174-Y), then a decoding procedure may be performed to extract the original packets from the received coded packets (block 176). The extracted original packets may then be forwarded to a corresponding application for use thereby (block 178). The above-described process may be repeated for each data transfer operation to be carried out between the source node and the destination node.

[0074] In Table 1 below, the proposed protocol is compared with some existing methods that use TCP connections for data transmission. In Table 1, the TCP/NC protocol is referring to that proposed in "Network Coding Meets TCP: Theory and Implementation," by Sundararajan et al., *Proceedings of the IEEE*, vol. 99, no. 3, pp. 490–512, Mar. 2011.

TCP	TCP/NC	MPTCP	MPTCP/NC
One TCP per connection	One TCP per connection	Multiple TCP per connection	Multiple TCP per connection
Split a file or a stream offline into different parts, one for each TCP flow	Split a file offline into different parts, one for each TCP flow	Online and dynamically split file based on sub-flow TCP window evolutions	Online and dynamically assign coded packets to TCP sub-flows
Each part treated as a separate file transmission	Coding occurs within each flow for disjoint parts	Original data packets assigned to sub-flows	

TABLE 1

[0075] An analytical model will now be presented to evaluate the performance of MPTCP/NC. The model will use several key elements presented in previous sections to provide some of the relevant abstractions. First, the model will assume that the MPTCP/NC layer provides a method that eliminates the need to track specific packets sent over a sub-flow. Because degrees of freedom are delivered to a sub-flow, rather than individual uncoded packets, the analysis presented will only need to track each degree of freedom through the network. Second, the model will assume that the congestion control algorithm used is similar to that of TCP Reno. For example, the congestion window increases by the number of received acknowledgements divided by the current window size, decreases by half upon receipt of a triple-duplicate acknowledgement, and completely closes if no acknowledgement is received after a given time-out period. Third, the model will assume that the network coding operations performed by the Fast-TCP/NC layer eliminate the need to consider the effects of triple-duplicates on TCP's window size. Instead, network coding simplifies the analysis. Since each received acknowledgement is an indication of a degree of freedom obtained by the receiver, only the effects of coded packet losses need to be considered. As a result, packet losses are interpreted as an increase in the round-trip time (RTT) between the source and destination. Finally, the concept of sending redundant coded packets will be used. For every packet contained in

the TCP contention window, R linearly independent packets will be transmitted. This added redundancy significantly reduces the probability of a time-out and it will be assumed that R is large enough to reduce the probability of a time-out to zero. Each of the above abstractions contributes to the ability to analyze the performance of MPTCP/NC.

[0076] MPTCP/NC will be analyzed below using two metrics, the average throughput T and the expected MPTCP/NC congestion window evolution $E[W]$. Since the instantaneous throughput of TCP is directly related to $E[W]$, the analysis will focus on finding $E[W]$ and then take an average of $E[W]$ over time to find T . MPTCP/NC's behavior will be modeled in terms of rounds. Within a round i , MPTCP/NC's expected congestion window size W_i will be determined, which is defined as the sum of each sub-flow's congestion window size $W_i^{(j)}$ for sub-flows $j = \{1,2\}$:

$$W_i = \sum_j W_i^{(j)} = W_i^{(1)} + W_i^{(2)}. \quad (1)$$

[0077] The natural choice for determining the duration of a round is to use the round trip time (RTT) from the sender to the receiver (i.e., $\text{duration}(i) = RTT$), where equation (1) assumes that $RTT_1 = RTT_2$. While this may work if there is a single TCP connection, each sub-flow in MPTCP/NC will typically have different round trip times making it difficult to determine which RTT to use. In order to account for this, the duration of each round will be set equal to the least common multiple (LCM) of the RTT for sub-flow 1 and sub-flow 2 (i.e., $\text{duration}(i) = LCM(RTT_1, RTT_2)$). Fig. 8 provides an illustration of this concept. The TCP congestion window size will now need to be determined for each sub-flow, given the duration of a round. This may be done by defining α and β as follows:

$$\alpha := \frac{RTT_1}{LCM(RTT_1, RTT_2)} \text{ and } \beta := \frac{RTT_2}{LCM(RTT_1, RTT_2)}. \quad (2)$$

The window size for sub-flow j in round i is then $W_{\{ik\}}^{(j)}$ for $k = \{\alpha, \beta\}$. Equation (1) now becomes:

$$W_i = W_{\{\alpha\}}^{(1)} + W_{\{\beta\}}^{(2)}. \quad (3)$$

[0078] The expected TCP congestion window size will now be analyzed for each sub-flow. The results will then be extended so that MPTCP/NC's end-to-end throughput can be determined. The analysis for the j th sub-flow in MPTCP/NC will be the same for each sub-flow. The most basic implementation of TCP will be used in the analysis and initially it will be assumed that each round's duration is equal to the RTT of sub-flow j . It will also be assumed that the window size during round i is determined by the number of acknowledgements a obtained during round $i - 1$. Specifically, the window size on each round is:

$$W_i^{(j)} = W_{i-1}^{(j)} + \frac{a_{i-1}^{(j)}}{W_{i-1}^{(j)}}. \quad (4)$$

Assuming that R_j linearly independent packets are sent for each packet contained the TCP congestion window and a packet loss rate of p_j , the expectation of the window size, $\mathbb{E}[W_i^{(j)}]$, may be expressed as:

$$\mathbb{E}[W_i^{(j)}] = \mathbb{E}\left[W_{i-1}^{(j)} + \frac{a_{i-1}^{(j)}}{W_{i-1}^{(j)}}\right] \quad (5)$$

$$= \mathbb{E}[W_{i-1}^{(j)}] + \mathbb{E}\left[\frac{a_{i-1}^{(j)}}{W_{i-1}^{(j)}}\right] \quad (6)$$

$$= \mathbb{E}[W_{i-1}^{(j)}] + \mathbb{E}\left[\mathbb{E}\left[\frac{a_{i-1}^{(j)}}{W_{i-1}^{(j)}} \mid W_{i-1}^{(j)}\right]\right] \quad (7)$$

$$= \mathbb{E}[W_{i-1}^{(j)}] + \mathbb{E}\left[\frac{1}{W_{i-1}^{(j)}} \mathbb{E}[a_{i-1}^{(j)} | W_{i-1}^{(j)}]\right] \quad (8)$$

$$= \mathbb{E}[W_{i-1}^{(j)}] + \mathbb{E}\left[\frac{1}{W_{i-1}^{(j)}} ((1 - p_j) R_j W_{i-1}^{(j)})\right] \quad (9)$$

$$= \mathbb{E}[W_{i-1}^{(j)}] + (1 - p_j) R_j. \quad (10)$$

Since the window size of TCP can increase by a maximum of one packet per round,

$$\mathbb{E}[W_i^{(j)}] = \mathbb{E}[W_{i-1}^{(j)}] + \min(1, (1 - p_j) R_j) \quad (11)$$

$$= \mathbb{E}[W_1^{(j)}] + (i - 1) \min(1, (1 - p_j) R_j), \quad (12)$$

where equation (12) is obtained by iterating equation (11).

[0079] Now that the expected window size in round i has been found, the throughput $T_i^{(j)}$ during the round can be calculated as follows:

$$T_i^{(j)} = \frac{\mathbb{E}[W_i^{(j)}]}{RTT_j} \min(1, (1 - p_j) R_j), \quad (13)$$

where the minimization is necessary to account for packets that are received that do not deliver new degrees of freedom. Since the Fast-TCP/NC layer codes all packets within the TCP congestion window, delivered packets 1 through $W_i^{(j)}$ contain new degrees of freedom. If more than $W_i^{(j)}$ packets are received in the round, the MPTCP/NC layer will disregard them since they obtain no new degrees of freedom. Therefore, the throughput is adjusted to ensure that only new degrees of freedom are taken into account.

[0080] The fact that the round duration is not equal to sub-flow j 's RTT (i.e., the duration of a round is equal to the LCM(RTT_1, RTT_2)) will now be taken into account.

The expected window size in equation (12) can be adjusted to account for the shorter rounds by substituting $|l/\gamma|$ for i , where l is the round number when using the shorter round duration and $\gamma = (\alpha, \beta)$:

$$\mathbb{E}[W_{|l/\gamma|}^{(j)}] = \mathbb{E}[W_1^{(j)}] + (|l/\gamma| - 1)\min(1, (1 - p_j)R_j). \quad (14)$$

Substituting the above equation into equation (13), the per-round throughput for a TCP sub-flow j is obtained as follows:

$$\tau_l^{(j)} = 1/\gamma \mathbb{E}[W_{|l/\gamma|}^{(j)}] \quad (15)$$

$$\begin{aligned} &= \frac{1}{\gamma \cdot RTT_j} \left(\mathbb{E}[W_1^{(j)}] \right. \\ &\quad \left. + (|l/\gamma| - 1)\min(1, (1 - p_j)R_j) \right) \min(1, (1 - p_j)R_j). \end{aligned} \quad (16)$$

This equation can be reduced if a large enough redundancy factor R_j is considered. As will be demonstrated later, the value chosen for R_j can be critical in achieving the maximum throughput. If it is assumed that the network capacity between the source and the destination is larger than the maximum window size $W_{\max}^{(j)} \times R_j$ for values of $R_j > \frac{1}{1-p_j}$, then equation (16) becomes,

$$\tau_l^{(j)} = \frac{1}{\gamma \cdot RTT_j} \left(\mathbb{E}[W_1^{(j)}] + (|l/\gamma| - 1) \right). \quad (17)$$

Finally, the fact that the window size of a TCP connection is limited by a maximum value $W_{\max}^{(j)}$ needs to be addressed. This changes the expected window size during round l to be

the minimum of equation (14) and $W_{\max}^{(j)}$. Furthermore, the per-round throughput becomes:

$$T_i^{(j)} = \frac{1}{\gamma \cdot RTT_j} \left(\min \left(W_{\max}^{(j)}, \mathbb{E}[W_i^{(j)}] + (l/\gamma - 1) \right) \right). \quad (18)$$

[0081] The model used in the above analysis of MPTCP/NC sub-flow performance makes several assumptions that, in practice, should be considered. First, it is assumed that packet losses are independent and identically distributed (i.i.d.) with loss probability p_i . When the network is not congested and none of the links experience outages, this assumption is valid; but it fails when packet losses are correlated. A possible extension to the analysis presented here is to take correlated packet losses into account by modeling them using Gilbert-Elliott model. Second, it is assumed that the number of redundant packets R_j sent for every packet contained in the TCP congestion window was sufficiently large to ignore the possibility of a time-out, which would close the TCP window to $\mathbb{E}[W_i^{(j)}]$. When $R_j = 1/(1 - p_i)$, this assumption is not necessarily valid. As will be shown in subsequent sections, time-outs occur frequently with i.i.d. packet losses when $R_j = 1/(1 - p_i)$. Specifically, a time-out occurs when the sum of received acknowledgements over two rounds, i and $i + 1$, is less than the window size during round i with probability $\Pr(a_i + a_{i+1} < W_i)$. One possible direction is to model time-outs as a random walk with a threshold crossing (although time-outs are not included in the present analysis).

[0082] Third, it is assumed that RTT_j remains constant. In practice, this is not true. Implementations of TCP generally use an averaged round-trip time often referred to as the “smoothed” round-trip time $SRTT$. Because network coding masks i.i.d. packet losses by extending the duration of RTT and network congestion or transmission delay is not taken into account, the RTT used in our analysis is effectively equal to $SRTT$.

[0083] It is desirable to determine average end-to-end throughput for MPTCP/NC. As shown above, the throughput for TCP/NC is dependent on the TCP congestion window size. Likewise, the average MPTCP/NC throughput is dependent on the joint window size

of all sub-flows (i.e., $W_i = \sum_j W_j^{(j)}$). While the approach shown here can be used for any number of TCP/NC sub-flows, results will only be provided for two sub-flows designated as sub-flow 1 and sub-flow 2 with packet loss probabilities and round-trip times of p_1 , RTT_1 and p_2 , RTT_2 , respectively.

[0084] Using the least common multiple of RTT_1 and RTT_2 as the duration of each round and the values of α and β defined in (2), the expected congestion window size for MPTCP/NC is:

$$\mathbb{E}[W_i] = \mathbb{E}\left[W_{\lceil i/\alpha \rceil}^{(1)} + W_{\lceil i/\beta \rceil}^{(2)}\right] \quad (19)$$

$$= \mathbb{E}\left[W_{\lceil i/\alpha \rceil}^{(1)}\right] + \mathbb{E}\left[W_{\lceil i/\beta \rceil}^{(2)}\right] \quad (20)$$

$$\begin{aligned} &= \mathbb{E}\left[W_1^{(1)}\right] + (\lceil i/\alpha \rceil - 1)\min(1, (1 - p_1)R_1) \\ &+ \mathbb{E}\left[W_1^{(2)}\right] + (\lceil i/\beta \rceil - 1)\min(1, (1 - p_2)R_2). \end{aligned} \quad (21)$$

One thing that was not taken into account is the maximum window size of each TCP/NC sub-flow. Incorporating the maximum window size of sub-flow 1, $W_{\max}^{(1)}$, and sub-flow 2, $W_{\max}^{(2)}$, the expected MPTCP/NC window size becomes:

$$\begin{aligned} \mathbb{E}[W_i] &= \max\left(W_{\max}^{(1)}, \mathbb{E}\left[W_1^{(1)}\right] + (\lceil i/\alpha \rceil - 1)\min(1, (1 - p_1)R_1)\right) \\ &+ \max\left(\mathbb{E}\left[W_1^{(2)}\right] + (\lceil i/\beta \rceil - 1)\min(1, (1 - p_2)R_2)\right). \end{aligned} \quad (22)$$

The equation above shows that the expected MPTCP/NC window size is monotonically increasing. As a result, it would be desirable to determine how fast it increases. This may be investigated by finding the number of rounds R_W it takes to get to a window size of W .

The intuition behind finding the expected number of rounds it will take to reach a window size of W is that each sub-flow is geometrically distributed. Letting $0 \leq x \leq W$:

$$\mathbb{E}[\mathcal{R}_W] = \mathbb{E}\left[\mathcal{R}_W^{(1)}\right] + \mathbb{E}\left[\mathcal{R}_W^{(2)}\right] \quad (23)$$

$$= \frac{\alpha}{1 - p_1} \sum_{i=1}^{W_x^{(1)}-1} i + \frac{\beta}{1 - p_2} \sum_{i=1}^{W_{W-x}^{(2)}-1} i \quad (24)$$

$$= \frac{\alpha \left(W_x^{(1)}(W_x^{(1)} - 1) \right)}{2(1 - p_1)} + \frac{\beta \left(W_{W-x}^{(2)}(W_{W-x}^{(2)} - 1) \right)}{2(1 - p_2)} \quad (25)$$

[0085] Now that the expected MPTCP/NC congestion window size has been determined, the average end-to-end throughput for MPTCP/NC may be found, which is a function of the above window size:

$$\mathcal{T} = \frac{1}{n} \sum_{i=1}^n \mathcal{T}_i \quad (26)$$

$$= \frac{1}{n} \sum_{i=1}^n (\mathcal{T}_i^{(1)} + \mathcal{T}_i^{(2)}) \quad (27)$$

$$= \frac{1}{n \cdot \alpha \cdot RTT_1} \sum_{i=1}^n \mathbb{E}\left[W_i^{(1)}\right] + \frac{1}{n \cdot \beta \cdot RTT_2} \sum_{i=1}^n \mathbb{E}\left[W_i^{(2)}\right] \quad (28)$$

Assuming that $n\alpha, n\beta \in \mathbb{Z}$,

$$\mathcal{T} = \frac{1}{n \cdot RTT_1} \sum_{i=1}^{n\alpha} \mathbb{E}\left[W_i^{(1)}\right] + \frac{1}{n \cdot RTT_2} \sum_{i=1}^{n\beta} \mathbb{E}\left[W_i^{(2)}\right] \quad (29)$$

$$= \frac{1}{n \cdot RTT_1} \left(\frac{n}{\alpha} \mathbb{E}\left[W_1^{(1)}\right] + \frac{n\alpha(n\alpha - 1)}{2} \min(1, (1 - p_1)R_1) \right)$$

$$+ \frac{1}{n \cdot RTT_2} \left(\frac{n}{\beta} \mathbb{E}[W_1^{(2)}] + \frac{n\beta(n\beta - 1)}{2} \min(1, (1 - p_2)R_2) \right) \quad (30)$$

$$\begin{aligned} &= \frac{1}{n \cdot \alpha \cdot RTT_1} \left(n \mathbb{E}[W_1^{(1)}] + \frac{n(n - \alpha)}{2\alpha} \min(1, (1 - p_1)R_1) \right) \\ &\quad + \frac{1}{n \cdot \beta \cdot RTT_2} \left(n \mathbb{E}[W_1^{(2)}] + \frac{n(n - \beta)}{2\beta} \min(1, (1 - p_2)R_2) \right). \end{aligned} \quad (31)$$

If $n\alpha, n\beta \notin \mathbb{Z}$, the above equation will contain additional terms that contain packets sent in the rounds from $\lfloor ny \rfloor$ to ny for $y = \{\alpha, \beta\}$. While these additional packets increase the throughput, their contribution is negligible for large enough n . Finally, the maximum window size of each sub-flow $W_{\max}^{(j)}$ is considered. First, let:

$$\begin{aligned} r^{(1)} &= \alpha \left(W_{\max}^{(1)} - \mathbb{E}[W_1^{(1)}] \right) \text{ and } r^{(2)} \\ &= \beta \left(W_{\max}^{(2)} - \mathbb{E}[W_1^{(2)}] \right). \end{aligned} \quad (32)$$

Using equation (31) and assuming that $R_1 > 1/(1 - p_1)$ and $R_2 > 1/(1 - p_2)$, four cases for the average end-to-end throughput are generated:

\mathcal{T}

$$\begin{aligned}
 &= \left\{ \begin{array}{l} \frac{1}{n \cdot \alpha \cdot RTT_1} \left(n \mathbb{E}[W_1^{(1)}] + \frac{n(n-\alpha)}{2\alpha} \right) \\ + \frac{1}{n \cdot \beta \cdot RTT_2} \left(n \mathbb{E}[W_1^{(2)}] + \frac{n(n-\beta)}{2\beta} \right) \end{array} \right. && \text{for } n \leq r^{(1)}, r^{(2)}, \\ &\quad \frac{1}{n \cdot \alpha \cdot RTT_1} \left(n \mathbb{E}[W_1^{(1)}] + \frac{n(n-\alpha)}{2\alpha} \right) \\ &\quad + \frac{1}{n \cdot \beta \cdot RTT_2} \left(r^{(2)} \mathbb{E}[W_1^{(2)}] + \frac{r^{(2)}(r^{(2)}-\beta)}{2\beta} + W_{\max}^{(2)}(n-r^{(2)}) \right) && \text{for } r^{(2)} < n \leq r^{(1)}, \\ &= \frac{1}{n \cdot \alpha \cdot RTT_1} \left(r^{(1)} \mathbb{E}[W_1^{(1)}] + \frac{r^{(1)}(r^{(1)}-\beta)}{2\alpha} + W_{\max}^{(1)}(n-r^{(1)}) \right) && (33) \\ &\quad + \frac{1}{n \cdot \beta \cdot RTT_2} \left(n \mathbb{E}[W_1^{(2)}] + \frac{n(n-\beta)}{2\beta} \right) && \text{for } r^{(1)} < n \leq r^{(2)}, \\ &\quad \frac{1}{n \cdot \alpha \cdot RTT_1} \left(r^{(1)} \mathbb{E}[W_1^{(1)}] + \frac{r^{(1)}(r^{(1)}-\beta)}{2\alpha} + W_{\max}^{(1)}(n-r^{(1)}) \right) \\ &\quad + \frac{1}{n \cdot \beta \cdot RTT_2} \left(r^{(2)} \mathbb{E}[W_1^{(2)}] + \frac{r^{(2)}(r^{(2)}-\beta)}{2\beta} + W_{\max}^{(2)}(n-r^{(2)}) \right) && \text{for } r^{(1)}, r^{(2)} < n. \end{array} \right.$$

[0086] An analysis has thus been provided to determine the sub-flow performance and overall MPTCP/NC performance using the expected window size W and the end-to-end average throughput \mathcal{T} as metrics. As noted above, results were provided for two sub-flows; but network coding makes it possible to use the same technique for any number of sub-flows. In the analysis, assumptions were made regarding the congestion control algorithm, time-outs and the amount of redundancy introduced into the network to account for packet losses. In the discussion that follows, numerical results are provided which show that the amount of redundancy is critical in eliminating time-outs.

[0087] The performance of the proposed MPTCP/NC scheme was analyzed through numerical simulations. As part of the analysis, the same two-link model used in the mathematical analysis was considered, where the packet losses on a given link are i.i.d. with probability p_i , $i = 1, 2$ and fixed throughout the use of the channel. In addition, the losses are independent between links, packets are not received out of order, the round trip

time on each link is fixed and acknowledgments are not lost. This proposed model is suitable for a wireless scenario where congestion is not a major factor for packet losses.

[0088] In the simulations, TCP is restricted to operate in two modes: slow start and congestion avoidance. The slow start threshold (*SST*) for each TCP window is 64 packets, and the window can have a maximum size of $W_{max} = 128$ packets. A timeout is declared by TCP when a packet is not acknowledged (remains in the window) for two round trip times. When a timeout occurs, the window is reset to one. Furthermore, the window size is reduced to half when a triple duplicate ACK is received. Note that this never occurs in MPTCP/NC, since every received package is acknowledged. In addition, it is assumed that the transmitter sends a file of a given (fixed) size, to which it has access before transmission starts.

[0089] Three MPTCP schemes were considered:

- (1) Uncoded MPTCP: this is traditional MPTCP operation where packets are dynamically allocated to a given TCP subflow when requested.
- (2) MPTCP/NC with low redundancy: MPTCP/NC operation is modeled with redundancy on each link equal to $R_i = (1 - p_i)^{-1}$, $i = 1,2$, as discussed in the mathematical analysis above. It is assumed that the entire transmitted file is coded.
- (3) MPTCP/NC with high redundancy: MPTCP/NC is modeled with redundancy 5% above the average loss, i.e., $R_i = 1.05(1 - p_i)^{-1}$, $i = 1,2$.

[0090] Three parameters were measured for each of these schemes considering different file sizes, packet loss probabilities and round trip times: (i) completion time, (ii) average rate, and (iii) total window size (i.e., sum of the windows of the two TCP subflows). The round trip time and packet loss probability parameters are shown in Fig. 9, together with a correspondence between the parameters and different wireless standards.

[0091] The simulated model has several limitations. For example, it assumed that each link always supports up to 128 packets being transmitted simultaneously, regardless of the round trip time and link capacity. As mentioned previously, this model also does not consider losses due to congestion or lost acknowledgments. Furthermore, the assumption that the MPTCP/NC layer has access to the whole file being transmitted is unrealistic for several different applications, such as video streaming. Nevertheless, as

will be shown in the following analysis, this model does capture some of the fundamental characteristics of MPTCP/NC's operation, clearly illustrating the benefits of coding over two TCP flows in regards using uncoded MPTCP.

[0092] In the discussion that follows, numerical results are provided for four different simulated scenarios, each with different values of RTT and link loss probability. The average results presented were taken over a set of 1000 different simulation instances done in MATLAB.

[0093] In the first scenario, the upper left hand side corner of the table in Fig. 3 is considered, modeling one link as very reliable and with low delay, with parameters $RTT_1 = 15\text{ms}$, $p_1 = 10^{-3}$, and the other link as slightly worse, but still reliable with average delay, with parameters $RTT_2 = 80\text{ms}$ and $p_2 = 10^{-2}$. The goal is to illustrate the performance and window size in a scenario with small loss rate and low round trip times. Out of all the scenarios considered, this is the one that achieves the highest transmission rates.

[0094] The average completion time and average rate for different file sizes are illustrated in Figs. 10a and 10b. For this scenario, the average rate of uncoded MPTCP is considerably lower than the coded cases. In addition, the extra 5% redundancy introduced by MPTCP/NC also significantly increases the rate. Notice also that the average rate with MPTCP/NC with 5% extra redundancy increases with file size. This is due to the fact that each TCP subflow rarely times out in this low loss scenario, with the window staying at its maximum size during virtually all of the transmission time after it achieves its maximum value. The effect of time out events is more noticeable in the uncoded case and in MPTCP/NC with just average redundancy, with the average rate achieving its maximum size at files of size 10^5 packets. On the other hand, for the case with coding and extra redundancy, the rate steadily increases with file size.

[0095] The behavior of each of the three schemes can be better understood by Fig. 10c. Here "snapshots" of the joint window size are presented for specific instances of the initial operation of each scheme together with the averages calculated over different simulations of the channel, considering a file size of 10^5 packets. Observe that the extra 5% redundancy has two key contributions to MPTCP/NC: (i) the window size achieves its maximum value before the case where we use MPTCP/NC with average redundancy and (ii) each TCP subflow suffers less time outs. The first contribution follows naturally from

the fact that extra redundancy increases the number of acknowledgments received from the Fast-TCP/NC layer. The second point stems from the fact that, for the proposed MPTCP/NC scheme, time out events have an extremely low probability of occurring. When an additional 5% redundancy is added, these events become even rarer, with a time out not being observed in this particular instance.

[0096] Even for this low loss scenario, uncoded MPTCP/NC performs poorly, with the joint window size never achieving the maximum value. Notice how for this particular instance uncoded TCP quickly “averages out”, already varying over its average joint window size after only the first second of transmission time.

[0097] The second considered scenario is more realistic, with one link representing, for example, a WiMax connection, with $RTT_2 = 80\text{ms}$ and $p_2 = 10^{-1}$, and the second link having a high round trip time and low packet loss probability, given by $RTT_2 = 80\text{ms}$ and $p_2 = 10^{-1}$. The second link could represent, for example, a satellite connection. Similarly to case 1, Figs. 11a and 11b present the average completion time and average rate for different file sizes, and Fig. 11c depicts specific instances of the joint window size for each of the considered transmission schemes.

[0098] The advantage of MPTCP/NC over uncoded MPTCP becomes clear in this scenario, with an average rate increase of over 10x for the investigated file sizes. In the uncoded MPTCP scheme, the link with low RTT is underutilized due to the high number of losses. This can be seen by the small variations of the line corresponding to uncoded MPTCP in Fig. 11c. As a consequence, the window size is mostly determined by the slow link which, due to the high round trip time, increases slowly.

[0099] On the other hand, MPTCP/NC is able to compensate for the packet loss rate in the link with low RTT through coding. When adding redundancy according to the average packet loss probability, time outs still occur periodically in the faster, less reliable link, as can be observed in the lowest curve in Fig. 11c. However, the corresponding Fast-TCP/NC subflow is still able to quickly recover from the time out events.

[0100] When an additional 5% redundancy is added, time out events become very rare in the low RTT link. Consequently, this link is fully utilized, with the corresponding TCP subflow achieving its maximal window during the beginning of the transmission. Time out events still occur, as shown in Fig. 11c, but the spacing between time outs increases

significantly, and Fast-TCP/NC manages to quickly recover from these events. As a result, the increase of the joint window size is determined by the link with high RTT, and slowly grows towards its maximum value.

[0101] In the previous case, it was shown how uncoded MPTCP/NC tends to rely on the link with the lowest packet loss probability, regardless of the round trip time. In order to make this phenomena even more apparent, a third case will now be considered where one link is very reliable but with a very long round trip time, with $RTT_1 = 1500\text{ms}$ and $p_1 = 10^{-3}$, and the second link has a low round trip time (100 times smaller) but is very unreliable, with $RTT_2 = 15\text{ms}$ and $p_2 = 0.3$. This corresponds to two opposite extreme points of Fig. 9. The corresponding results are illustrated in Figs. 12a, 12b, and 12c.

[0102] The average rate obtained by MPTCP/NC with 5% extra redundancy is nearly two orders of magnitude larger than uncoded MPTCP. The dependence of MPTCP on the slower link becomes obvious in Fig. 12c, with the fast unreliable link corresponding to variations in the joint window size of only a few packets, and the coarse changes due to the link with high RTT. Observe also that, for a file of size 10^5 packets, after over 25 seconds of transmission the joint window size of uncoded MPTCP is still approximating its final average values, whereas the MPTCP/NC has already concluded transmission.

[0103] Once again it is noted that MPTCP/NC manages to reduce the number of time outs with coding. In addition, when a time out occurs, it is able to quickly recover due to the Fast-TCP/NC mechanism. As noted previously, the extra 5% redundancy also contributes significantly towards less timeouts, resulting in a significantly larger throughput in this noise scenario.

[0104] In the final scenario, both links are assumed to have the same error probability and different but comparable round trip times, with $RTT_1 = 80\text{ms}$, $p_1 = 10^{-1}$, $RTT_2 = 250\text{ms}$ and $p_2 = 10^{-1}$. This scenario could represent, for example, a 3G and a WiFi link in parallel. The results for this case are depicted in Figs. 13a, 13b, and 13c.

[0105] Due to the high packet loss probability in both links, the MPTCP scheme quickly “averages out”, and does not manage to increase significantly the joint window size of the TCP sub-flows. Furthermore, for large file sizes, adding 5% redundancy to MPTCP/NC above the average throughput of the channel almost doubles the achievable average rate. As discussed previously, this is due to the fact that the extra redundancy

makes time out events very rare, and a joint time out for the two links will almost never occur for most file sizes. On the other hand, when $R = (1 - p)^{-1}$, time outs happen more frequently and can occur jointly for both links, reducing the window size to only a few packets. Nevertheless, even in this case, the Fast-TCP/NC layer manages to quickly increase the joint window size.

[0106] Based on the above analyses, it is apparent that MPTCP/NC can provide an average transmission rate that is an order of magnitude or more better than MPTCP in different cases of interest. Some conclusions that may be drawn include: (a) a little redundancy in MPTCP/NC can be very helpful, reducing time outs (by $\approx 2x$ average rate in some noisy scenarios); (b) MPTCP performs poorly when compared to the coded case, even with low error probabilities; (c) MPTCP's performance is determined by the link with lower packet loss, nearly independent of RTT; (d) using coding masks packet losses as higher RTTs, as expected; and (e) MPTCP/NC is effective in leveraging links with significantly different RTTs.

[0107] Fig. 14 is a block diagram illustrating an exemplary node device architecture 200 that may be used in a node device that incorporates features described in the present disclosure in one or more implementations. As illustrated, the node device architecture 200 may include: one or more digital processors 204, memory 206, a first wireless transceiver 208, a second wireless transceiver 210, a wired network interface 212, a user interface 214, and a network encoder/decoder 216. A bus 220 and/or other structure(s) may be provided for establishing interconnections between various components of device architecture 200. Digital processor(s) 204 may include one or more digital processing devices that are capable of executing programs or procedures to provide functions and/or services for a user. Memory 206 may include one or more digital data storage systems, devices, and/or components that may be used to store data and/or programs for other elements of node device architecture 200. User interface 214 may include any type of device, component, or subsystem for providing an interface between a user and a corresponding node device. First and second wireless transceivers 208, 210 may include any type of transceivers that are capable of supporting wireless communication with one or more remote wireless entities. Wired network interface 212 may include any type of transceiver that is capable of supporting wired communication with one or more external communication entities.

[0108] Digital processor(s) 204 may include, for example, one or more general purpose microprocessors, digital signals processors (DSPs), controllers, microcontrollers, application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), programmable logic arrays (PLAs), programmable logic devices (PLDs), reduced instruction set computers (RISCs), and/or other processing devices or systems, including combinations of the above. Digital processor(s) 204 may be used to, for example, execute an operating system for a corresponding node device. Digital processor(s) 204 may also be used to, for example, execute one or more application programs for a node device. In addition, digital processor(s) 204 may be used to implement, either partially or fully, one or more of the communications related processes or techniques described herein in some implementations.

[0109] As described above, first and second wireless transceivers 208, 210 may include any type of transceivers that are capable of supporting wireless communication with one or more remote wireless entities. In various implementations, these transceivers 208, 210 may be configured in accordance with one or more wireless standards (e.g., wireless networking standards, wireless cellular standards, etc.) for use in communicating with corresponding wireless networks, systems, or devices. In this fashion, these transceivers may be used to support heterogeneous communication as described above. Although illustrated with two wireless transceivers in Fig. 14, it should be appreciated that any number of transceivers may be used in various embodiments. In some implementations, one or both of wireless transceivers 208, 210 may be capable of communicating with peer devices in a peer-to-peer, ad-hoc, or wireless mesh network arrangement. In addition, in some implementations, one or both of wireless transceivers 208, 210 may be capable of communicating with a base station or access point of an infrastructure-type wireless communication system or network. As illustrated in Fig. 14, wireless transceivers 208, 210 may each be coupled to one or more antennas 222, 224 and/or other transducers, to facilitate the transmission and/or reception of communication signals. In some implementations, wireless transceivers 208, 210 may be used to implement or facilitate, either partially or fully, one or more of the communications related processes or techniques described herein.

[0110] Wired network interface 212 may be configured in accordance with one or more wired communication standards. Although illustrated with a single wired network interface

212, it should be appreciated that any number of wired interfaces may be used in other implementations. In some embodiments, wired network interface 212 may be used to support heterogeneous networking in conjunction with one or more other transceivers (wired or wireless).

[0111] Memory 206 may include any type of system, device, or component, or combination thereof, that is capable of storing digital information (e.g., digital data, computer executable instructions and/or programs, etc.) for access by a processing device or other component. This may include, for example, semiconductor memories, magnetic data storage devices, disc based storage devices, optical storage devices, read only memories (ROMs), random access memories (RAMs), non-volatile memories, flash memories, USB drives, compact disc read only memories (CD-ROMs), DVDs, Blu-Ray disks, magneto-optical disks, erasable programmable ROMs (EPROMs), electrically erasable programmable ROMs (EEPROMs), magnetic or optical cards, and/or other digital storage suitable for storing electronic instructions and/or data. In some embodiments, computer executable instructions may be stored in memory 206 that, if executed within a digital processor, can result in the performance of one or more of the methods described herein, or portions thereof.

[0112] Network encoder/decoder 216 may include a device or system for performing network encoding and/or decoding for a node device. In a source node device that will generate and transmit network coded packets, network encoder/decoder 216 may include network encoding functionality. Likewise, in a destination node device that will receive and decode network encoded packets, network encoder/decoder 216 may include decoding functionality. In a node that may serve as both a source node and a destination node, both network encoding and decoding functionality may be provided. It should be appreciated that, although illustrated as a separate unit, network encoder/decoder 216 may be implemented, at least partially, within another component or device of a node in some implementations (e.g., within processor(s) 204 of Fig. 14, etc.).

[0113] It should be appreciated that the node device architecture 200 of Fig. 14 represents one possible example of an architecture that may be used in an implementation. Other architectures may alternatively be used. As used herein, the terms "node device," "node," or "communication device" are used to describe any type of digital electronic device or system that includes some form of communication capability (wireless and/or

wired) that can act as a source or destination of data flow. This may include, for example, a laptop, desktop, notebook, or tablet computer; a personal digital assistant (PDA); a personal communication service (PCS) device; a personal navigation assistant (PNA); a cellular telephone, smart phone, or other wireless communication device; a pager; a wireless sensor device; a satellite communication device; a media player having communication capability; a digital storage device with communication capability; an integrated circuit or multi-chip module, and/or other devices. It should be appreciated that all or part of the various devices, processes, or methods described herein may be implemented using any combination of hardware, firmware, and/or software. In some embodiments, methods described herein, or portions or variations thereof, may be implemented as computer executable instructions stored on one or more computer readable media. Any form of non-transitory computer readable media may be used in various embodiments.

[0114] Having described exemplary embodiments of the invention, it will now become apparent to one of ordinary skill in the art that other embodiments incorporating their concepts may also be used. The embodiments contained herein should not be limited to disclosed embodiments but rather should be limited only by the spirit and scope of the appended claims. All publications and references cited herein are expressly incorporated herein by reference in their entirety.

What is claimed is:

1. A machine implemented method for use in transferring data to a destination node, the method comprising:

obtaining a plurality of original data packets to be transferred to the destination node;

generating first coded packets by linearly combining original data packets using network coding;

distributing the first coded packets among multiple available paths leading to the destination node;

generating second coded packets by linearly combining first coded packets distributed to a first path of the multiple available paths, using network coding; and

transmitting the second coded packets associated with the first path to the destination node via a network associated with the first path.

2. The method of claim 1, further comprising:

generating second coded packets by linearly combining first coded packets distributed to a second path of the multiple available paths, using network coding; and

transmitting the second coded packets associated with the second path to the destination node via a network associated with the second path.

3. The method of claim 1, wherein:

generating first coded packets includes generating first coded packets by linearly combining original data packets that are within a sliding coding window.

4. The method of claim 3, further comprising:

receiving acknowledgement messages from the destination node that each indicate that a new degree of freedom has been received by the destination node in connection with the data transfer; and

adjusting the width of the sliding coding window based, at least in part, on received acknowledgement messages.

5. The method of claim 1, wherein:

generating second coded packets includes generating second coded packets by linearly combining first coded packets that are within a sliding coding window.

6. The method of claim 5, wherein:

generating second coded packets includes generating redundant second coded packets for each set of first coded packets within the sliding coding window.

7. The method of claim 1, wherein:

each of the multiple available paths includes a transmission control protocol (TCP) layer

that adds sequence numbers to first coded packets distributed to the path, wherein generating second coded packets by linearly combining first coded packets masks the sequence numbers added to the first coded packets associated with a path.

8. The method of claim 1, wherein:

each of the multiple available paths includes a transmission control protocol (TCP) layer; and

distributing the first coded packets includes distributing the first coded packets among the multiple available paths based, at least in part, on TCP congestion control window dynamics.

9. The method of claim 1, wherein:

the multiple available paths are associated with different network technologies.

10. The method of claim 1, wherein:

obtaining a plurality of original data packets includes receiving the original data packets from an application layer.

11. A machine implemented method for use in processing coded packets received from a source node via multiple different paths, the method comprising:

receiving coded packets associated with a first connection via multiple different paths, the coded packets each including a linear combination of original data packets;

for each coded packet associated with the first connection that is successfully received, sending an acknowledgement message to the source node without first determining whether the coded packet is linearly independent of previously received coded packets associated with the first connection;

forwarding all coded packets associated with the first connection, received from all paths, to a common processing layer without decoding the coded packets; and

for each coded packet associated with the first connection forwarded to the common processing layer:

determining whether the coded packet is linearly independent of coded packets associated with the connection that were previously forwarded to the common processing layer; and

sending an acknowledgement message to the source node acknowledging that a new degree of freedom has been received for the first connection if the coded packet is determined to be linearly independent.

12. The method of claim 11, wherein:

determining whether the coded packet is linearly independent includes performing a Gauss-Jordan elimination operation on a coefficient matrix.

13. The method of claim 11, wherein:

receiving coded packets associated with a first connection via multiple different paths includes receiving coded packets from paths that use different network technologies, wherein each of the different paths uses transmission control protocol (TCP) to control a transfer of packets through the path.

14. The method of claim 13, wherein:

sending an acknowledgement message to the source node without first determining whether the coded packet is linearly independent from previously received coded packets is performed as part of a TCP layer.

15. The method of claim 14, wherein:

sending an acknowledgement message to the source node without first determining whether the coded packet is linearly independent of previously received coded packets is performed within a first portion of the TCP layer for coded packets received via a first path and within a second portion of the TCP layer for coded packets received via a second path that is different from the first path.

16. The method of claim 11, wherein:

receiving coded packets associated with a first connection via multiple different paths includes receiving some coded packets via a first path following a first communication standard and receiving some other coded packets via a second path following a second communication standard that is different from the first communication standard.

17. The method of claim 11, wherein:

forwarding all coded packets associated with the first connection to a common processing layer includes forwarding the coded packets to a multiple path transfer control protocol with network coding (MPTCP-NC) layer that is higher than a TCP layer in a corresponding protocol architecture.

18. The method of claim 17, wherein:

the protocol architecture includes separate protocol stacks associated with each of the multiple different paths.

19. A communication device comprising:

a first network interface unit configured for communication in a first network;

a second network interface unit configured for communication in a second network that is different from the first network; and

at least one processor to manage data transfer between the communication device and a destination node using multiple different paths, the at least one processor to:

obtain a plurality of original packets representative of data to be transferred to the destination node;

generate first coded packets by linearly combining original data packets using network coding;

distribute first coded packets to the multiple different paths to the destination node; and

generate second coded packets within individual paths by linearly combining first coded packets distributed to the path, using network coding.

20. The communication device of claim 19, wherein the at least one processor is configured to transmit second coded packets generated within individual paths to the destination node via corresponding networks.

21. The communication device of claim 19, wherein the at least one processor is configured to generate first coded packets by linearly combining original data packets within a sliding coding window.

22. The communication device of claim 19, wherein the at least one processor is configured to generate second coded packets for a first path by linearly combining first coded packets within a sliding coding window.

23. The communication device of claim 22, wherein the at least one processor is configured to generate redundant second coded packets for the first path for each set of first coded packets within the sliding coding window.

24. The communication device of claim 19, wherein the at least one processor is configured to distribute the first coded packets to the multiple different paths based, at least in part, on TCP congestion control window dynamics.

25. The communication device of claim 19, wherein a first of the multiple different paths is associated with the first network interface unit and a second of the multiple different paths is associated with the second network interface unit.

26. The communication device of claim 25, further comprising:
at least one additional network interface unit configured for communication within at least one additional network.

27. A communication device comprising:
a first network interface unit configured for communication in a first network;
a second network interface unit configured for communication in a second network that is different from the first network; and
at least one processor to manage processing of coded data packets received from a source node via multiple different paths as part of a first connection, the coded data packets each including a linear combination of original data packets, the at least one processor to:
send an acknowledgement message to the source node for each coded data packet associated with the first connection that is successfully received;
forward coded data packets associated with the first connection, received from all paths, to a common processing layer without first decoding the coded packets; and
for each coded packet associated with the first connection forwarded to the common processing layer:
determine whether the coded packet is linearly independent of coded packets associated with the connection that were previously forwarded to the common processing layer; and

send an acknowledgement message to the source node acknowledging that a new degree of freedom has been received for the first connection if the coded packet is determined to be linearly independent.

28. The communication device of claim 27, wherein the at least one processor is configured to discard a received coded packet if it is determined that the packet is not linearly independent.

29. The communication device of claim 27, wherein a first of the multiple different paths is associated with the first network interface unit and a second of the multiple different paths is associated with the second network interface unit.

30. The communication device of claim 27, further comprising:
at least one additional network interface unit configured for communication within at least one additional network.

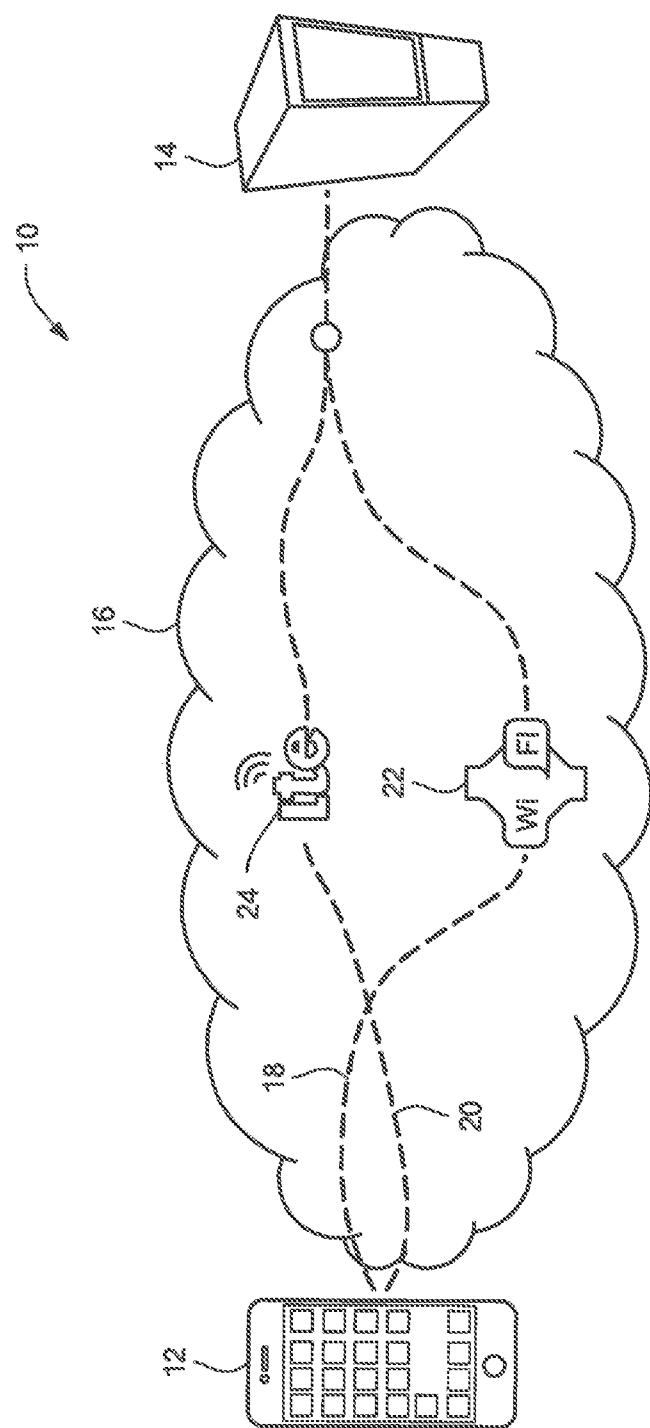


FIG. 1

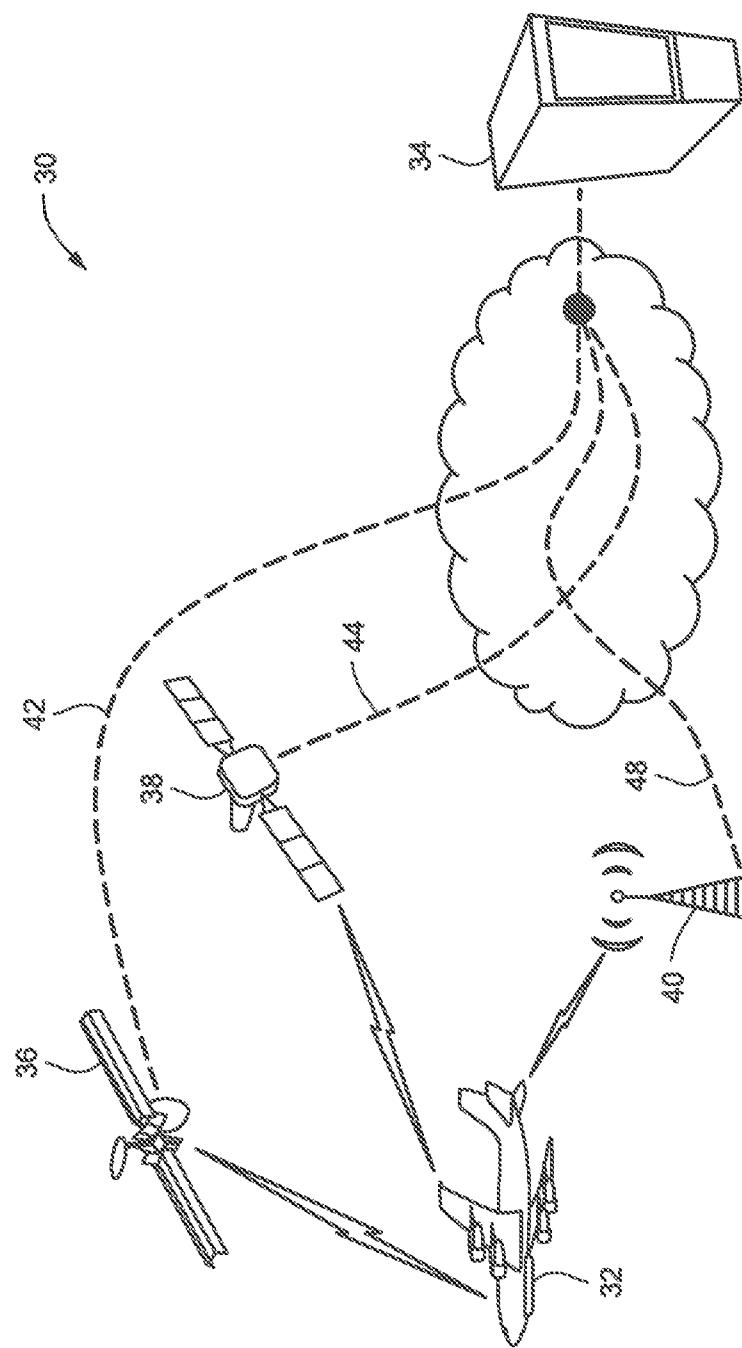
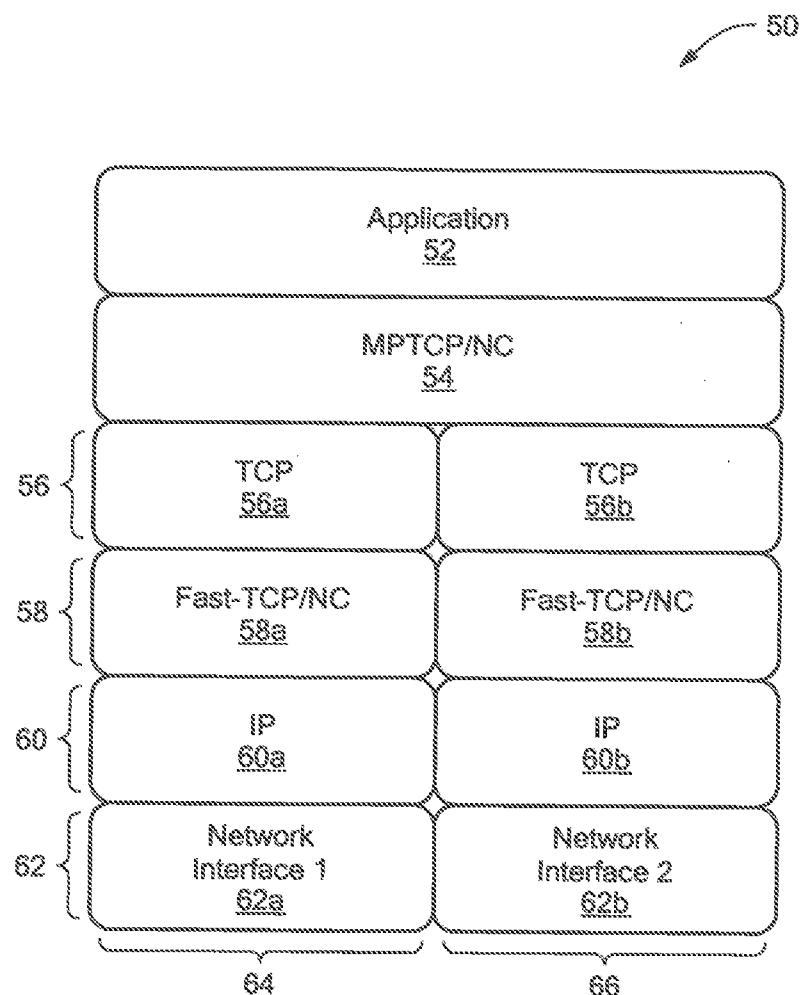
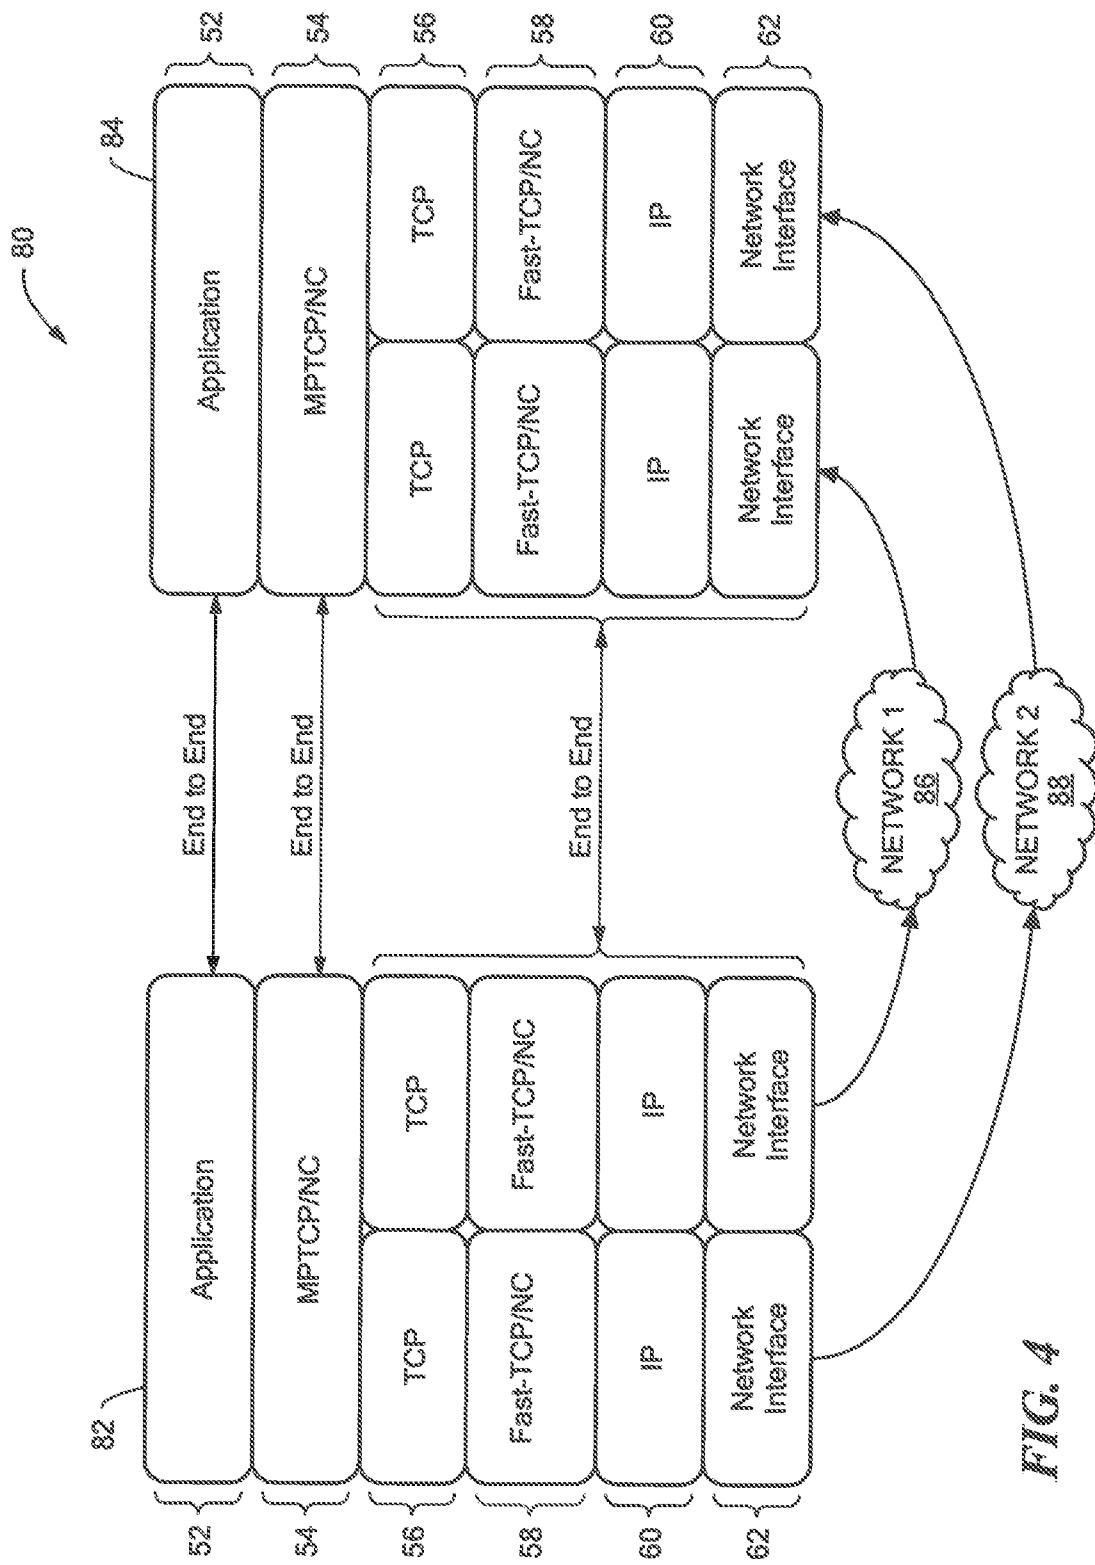


FIG. 2

3/19

**FIG. 3**



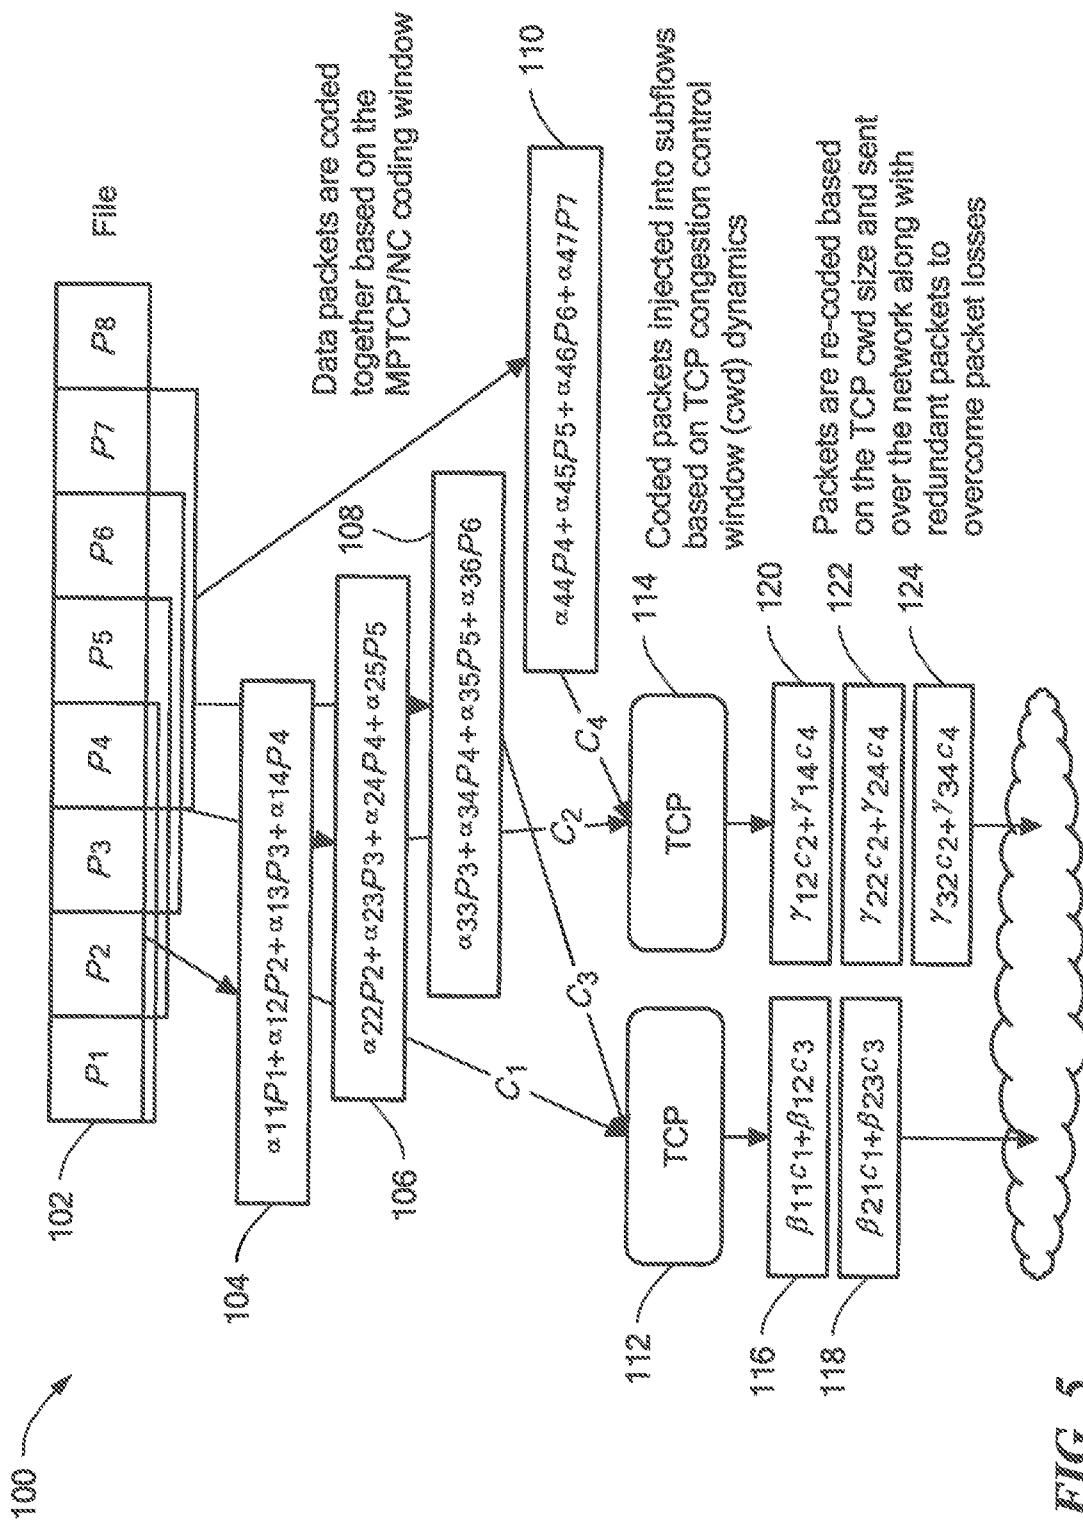
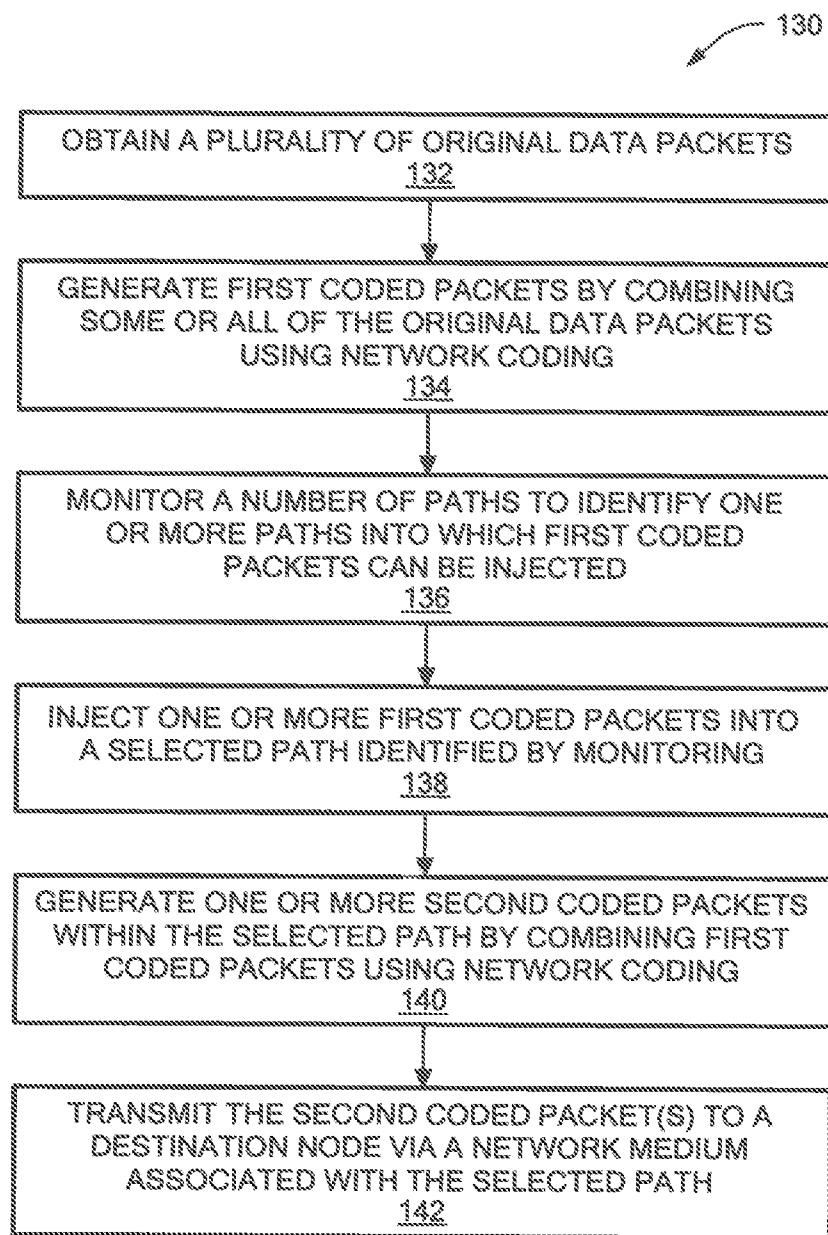


FIG. 5

6/19

**FIG. 6**

7/19

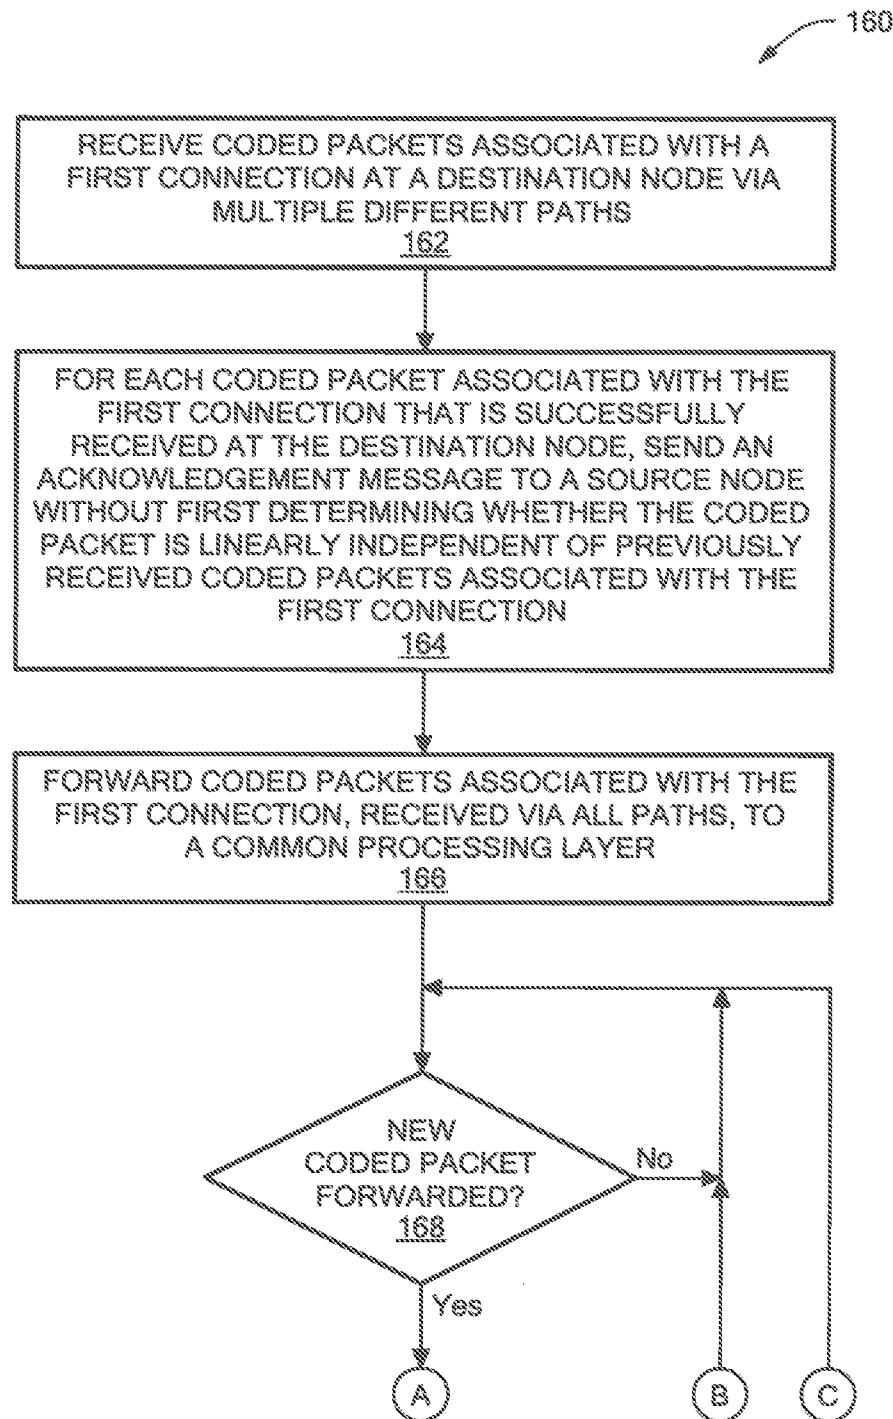


FIG. 7A

8/19

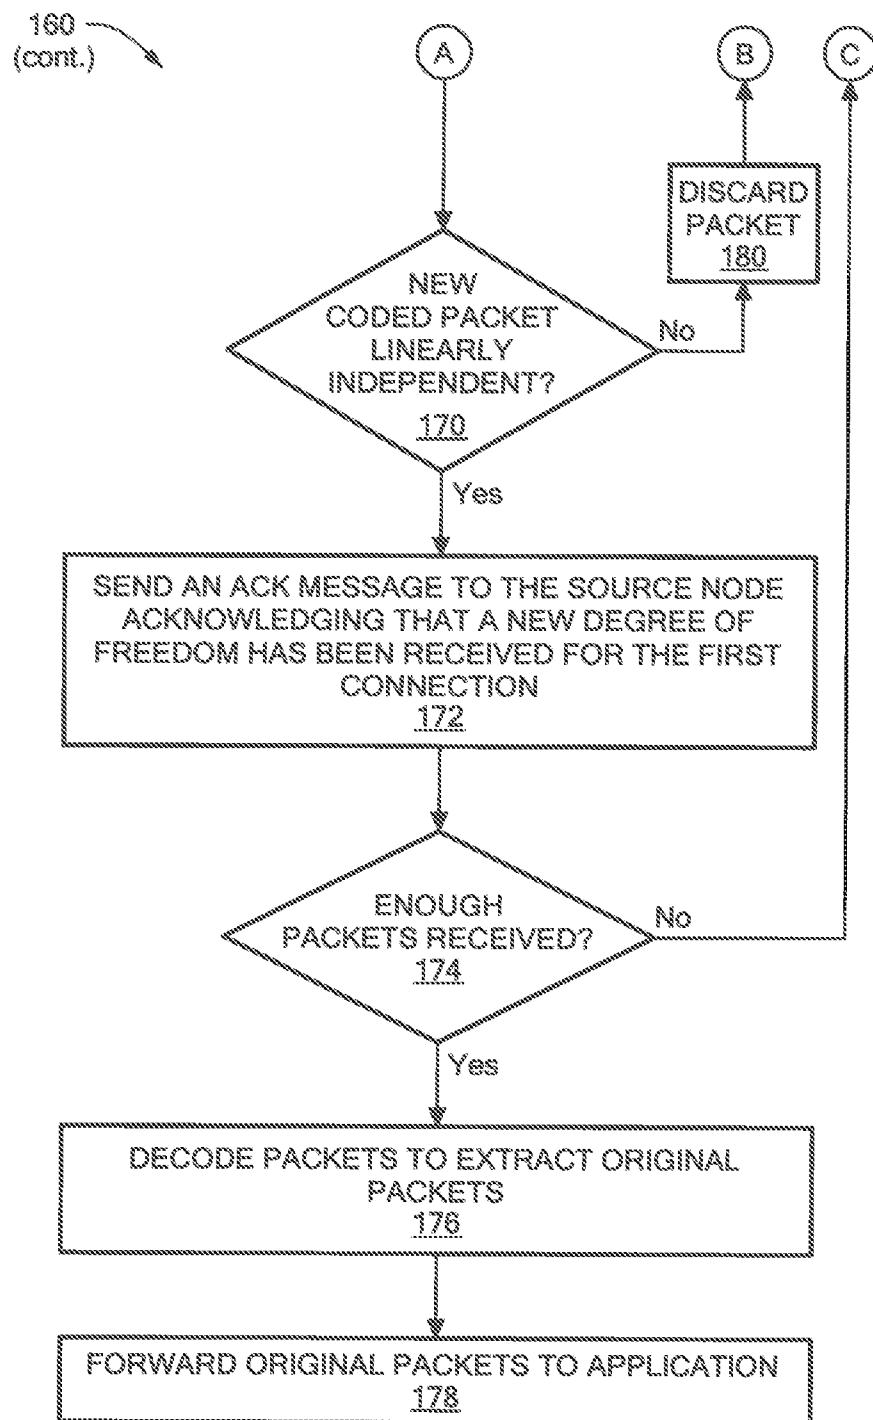


FIG. 7B

9/19

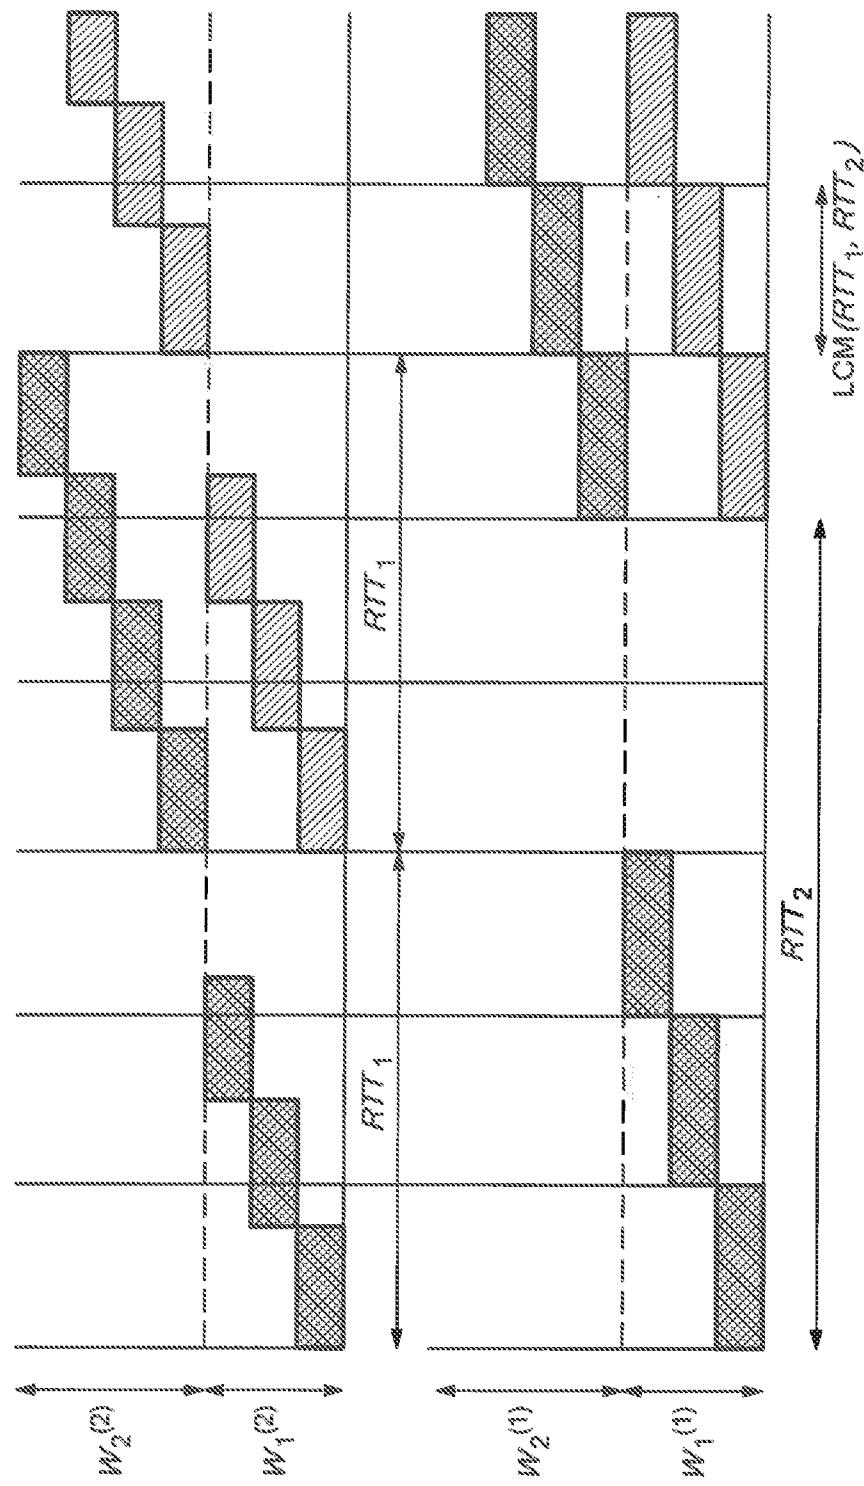
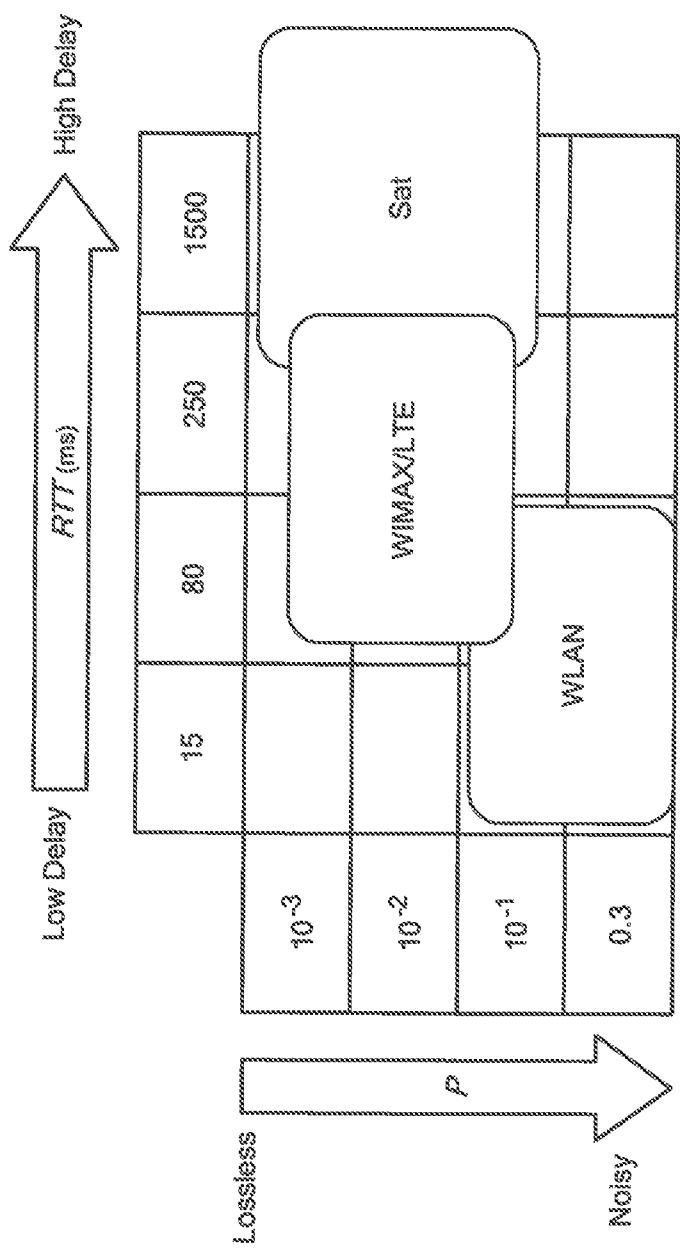
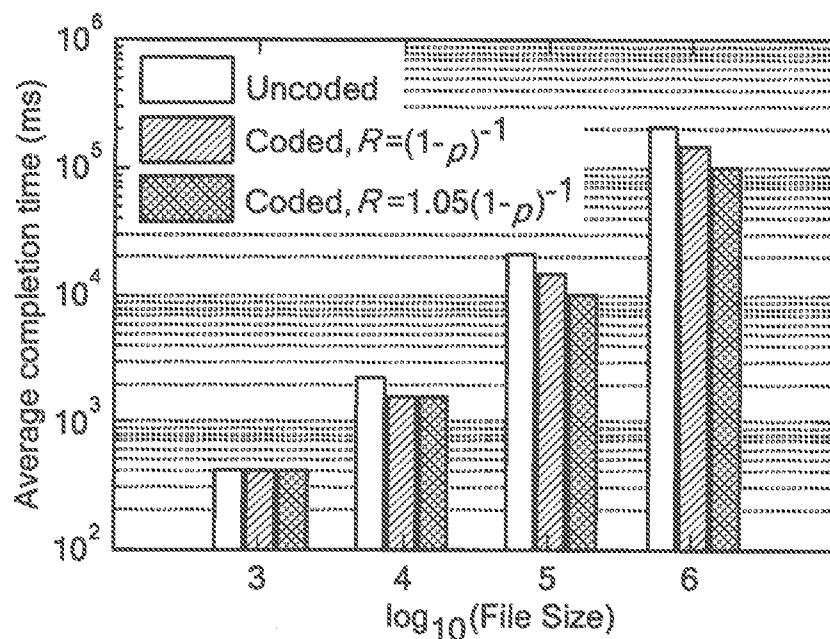
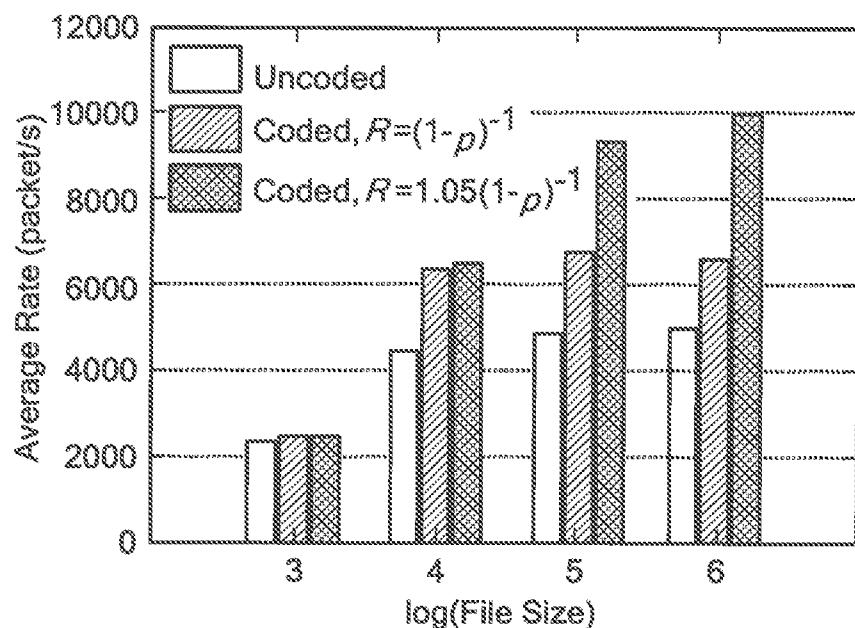


FIG. 8



9

11/19

**FIG. 10A****FIG. 10B**

12/19

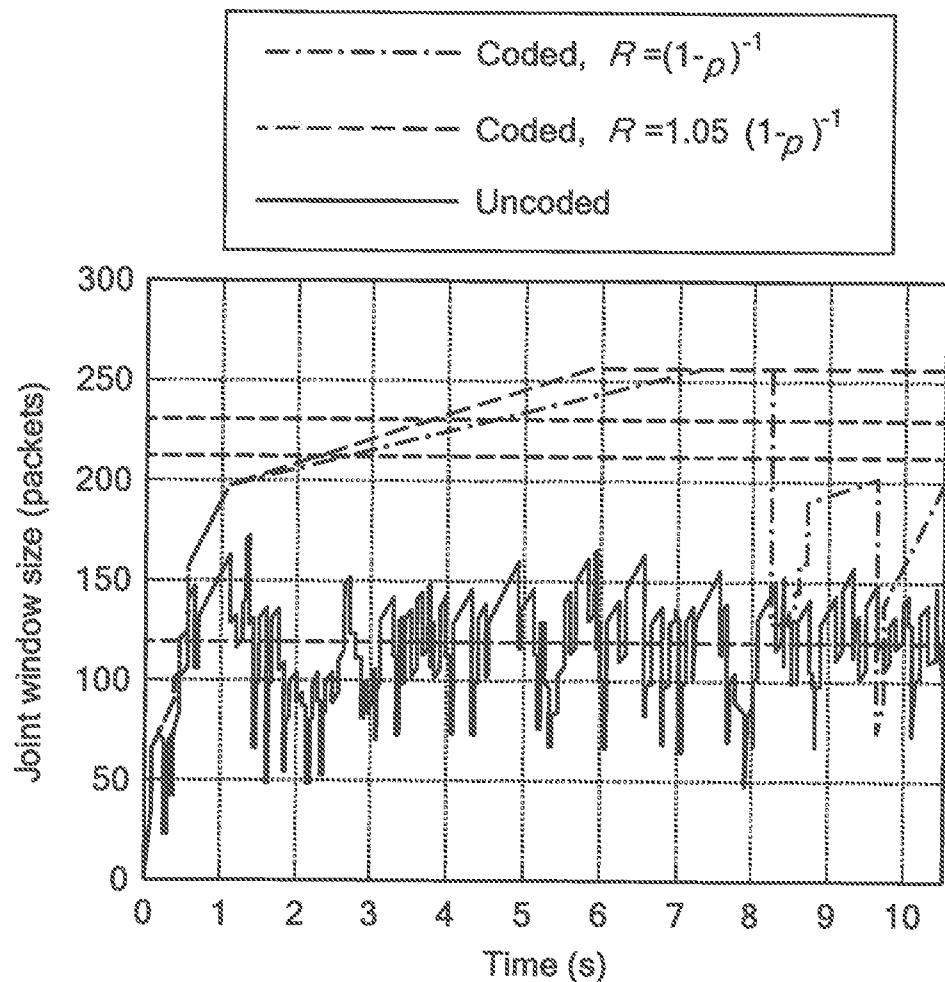
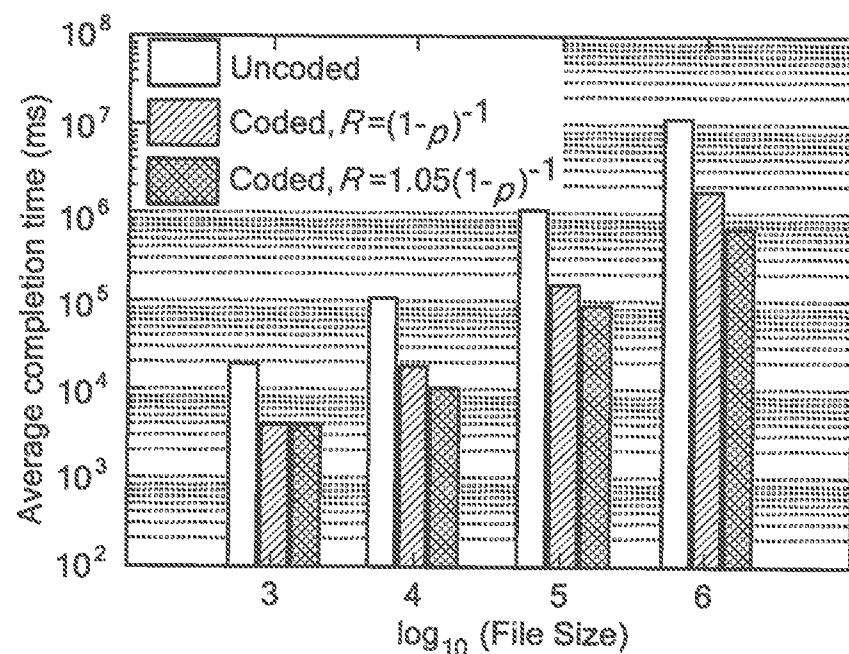
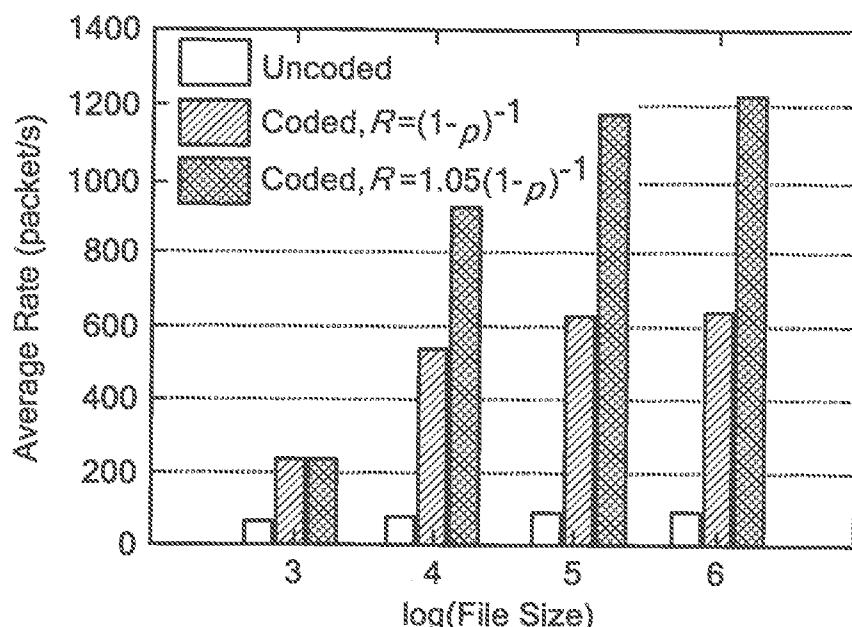


FIG. 10C

13/19

**FIG. 11A****FIG. 11B**

14/19

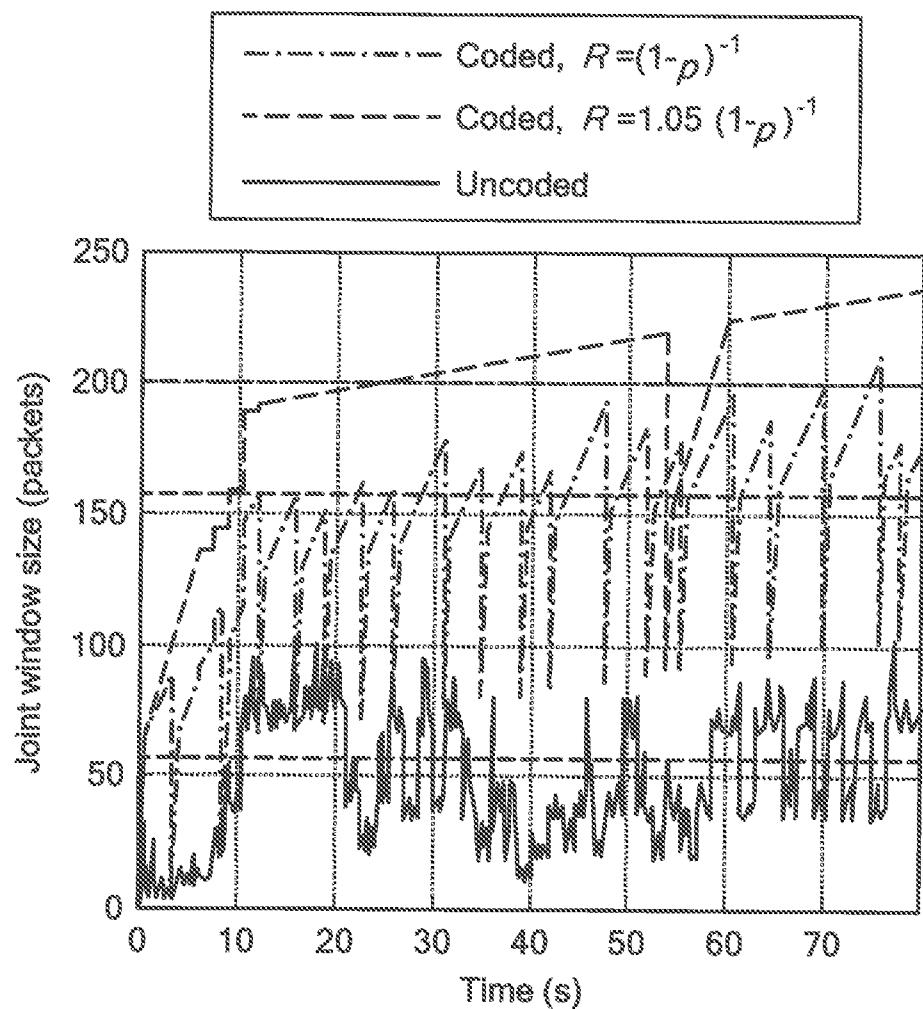
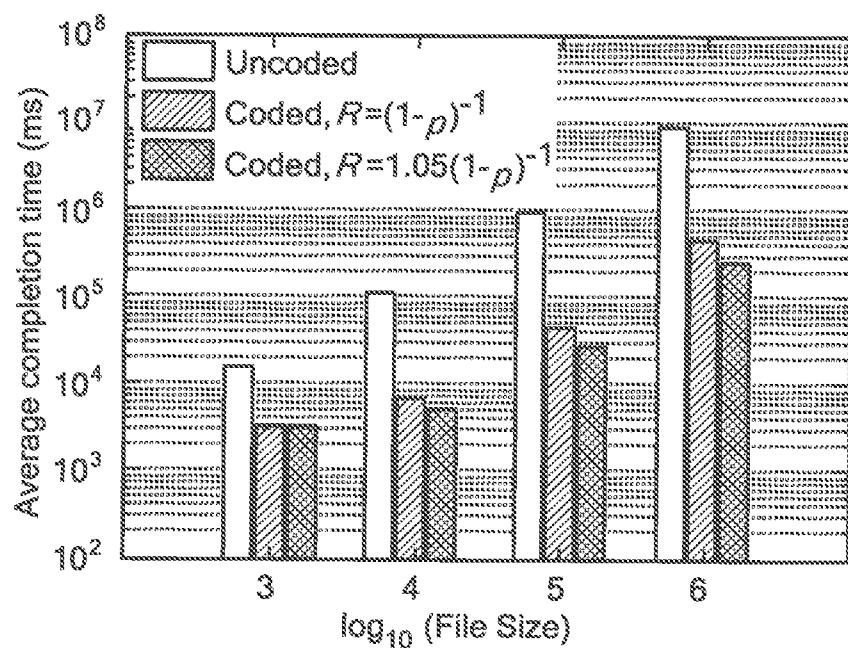
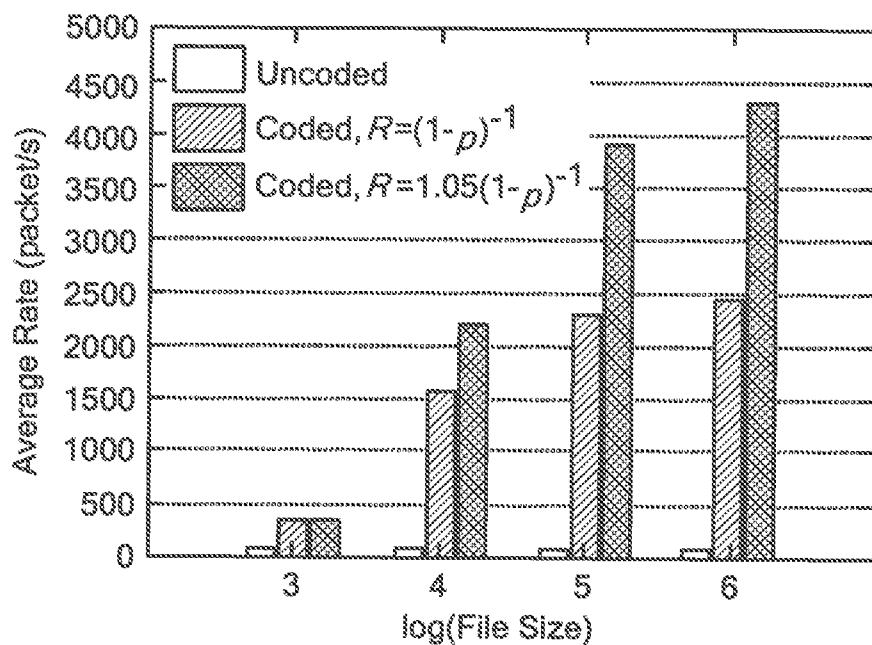


FIG. II C

15/19

**FIG. 12A****FIG. 12B**

16/19

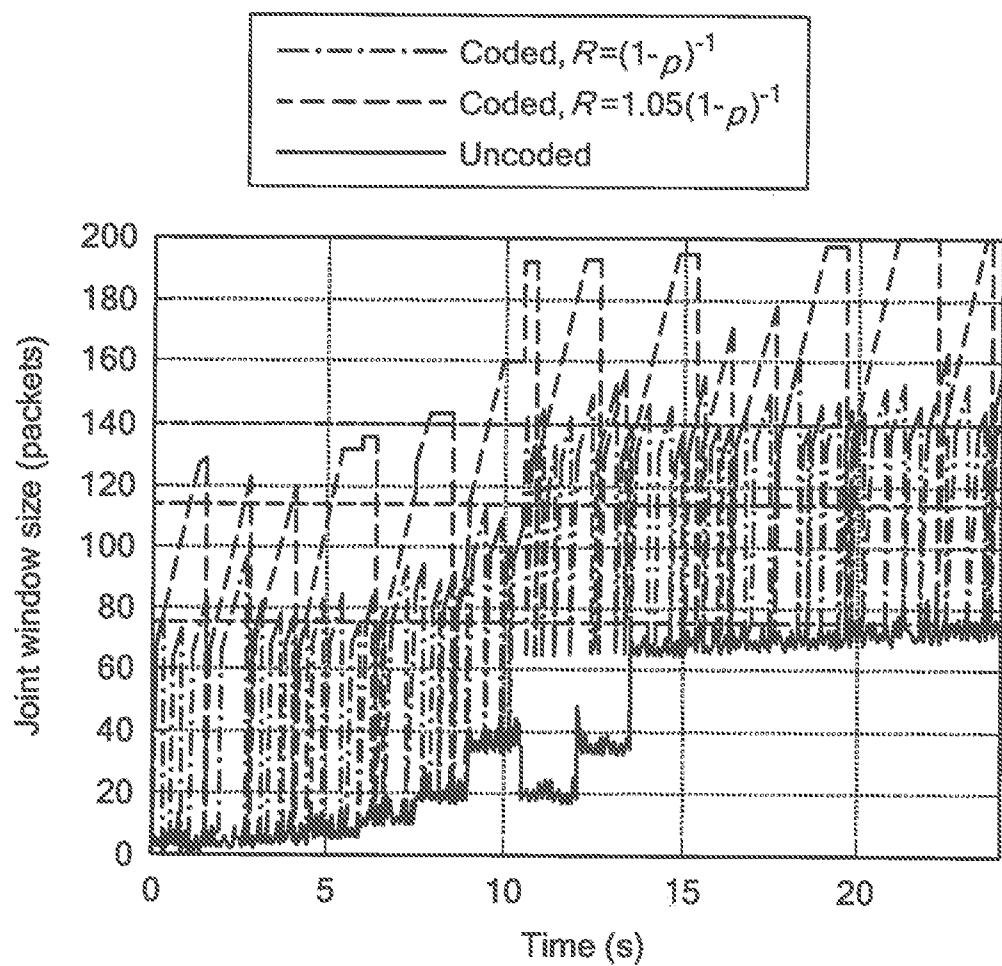
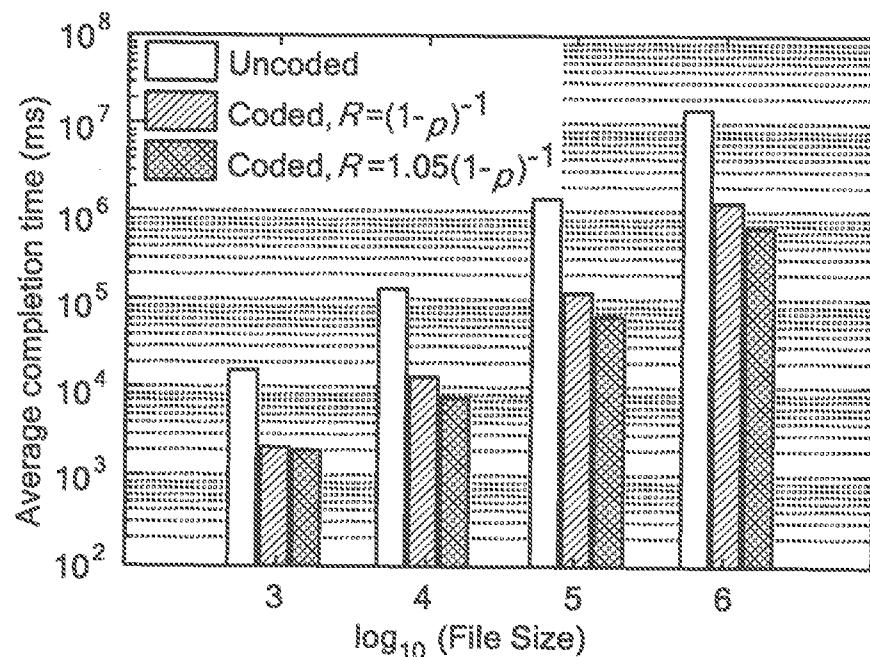
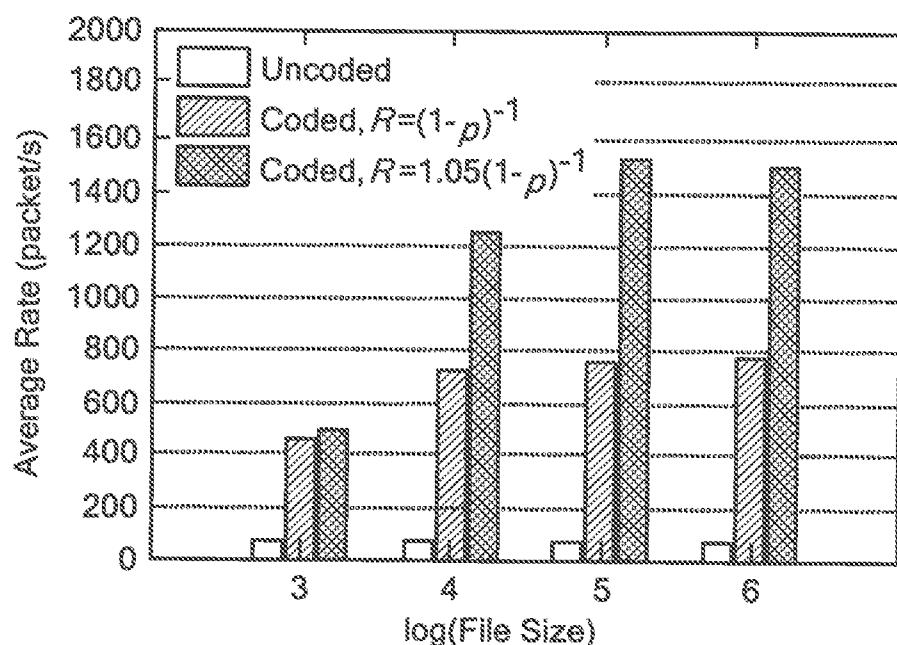


FIG. 12C

17/19

**FIG. 13A****FIG. 13B**

18/19

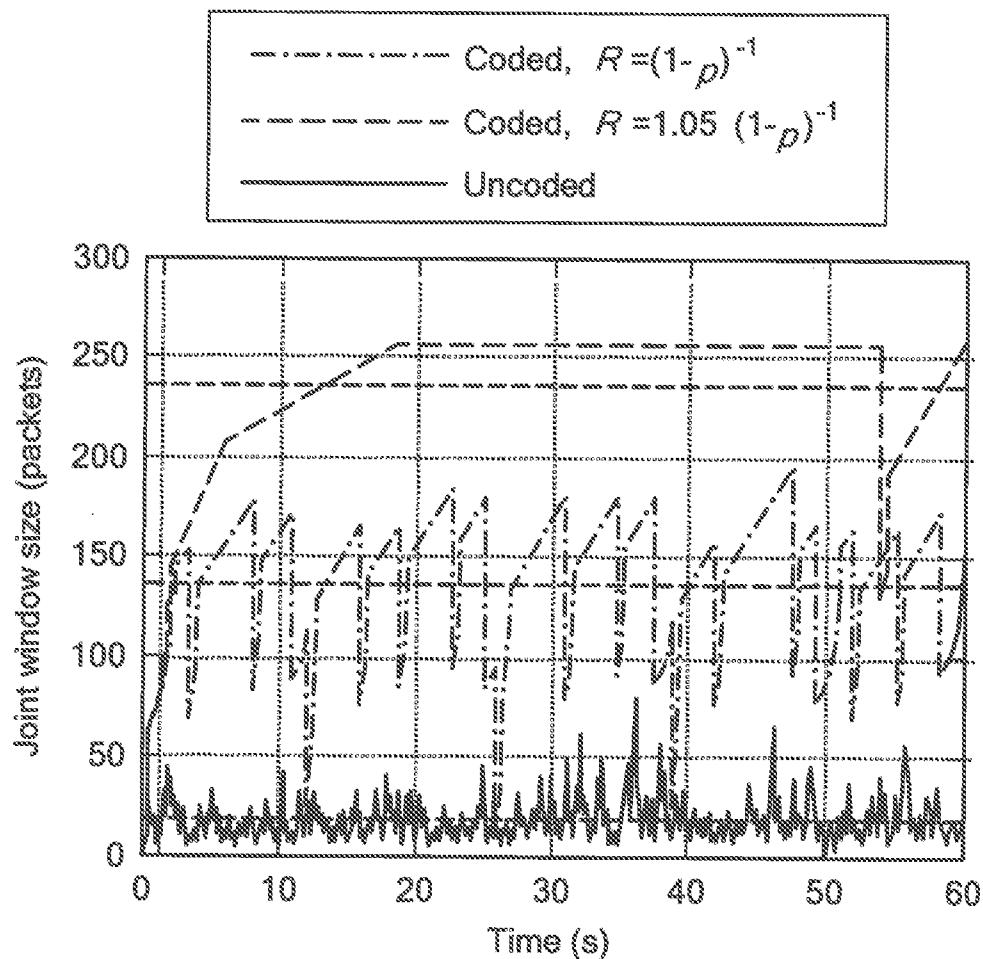


FIG. 13C

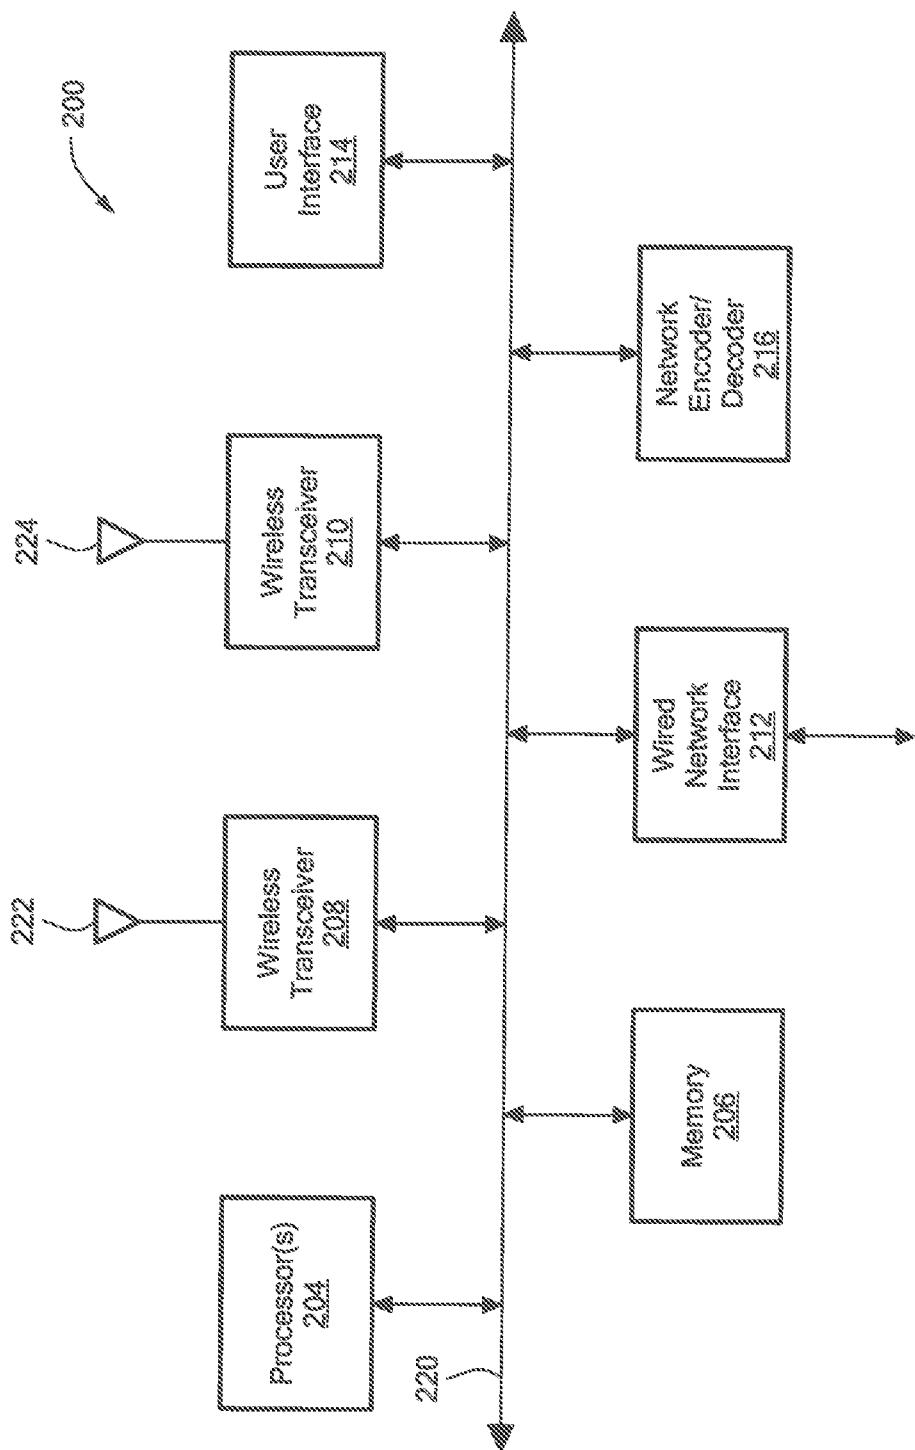


FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 13/24039

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(8) - G06F 15/173 (2013.01)
 USPC - 709/238

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 USPC: 709/238

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 USPC: 709/220, 227, 238 (keyword limited - see terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 PatBase; GOOGLE; GoogleScholar; GooglePatents; PubWEST(PGPB,USPT,EPAB,JPAB)
 Search Terms: network, packet, encoded, coded, congested, traffic, sequence, hide, mask, TCP source, destination, transmit, send, receive, transfer, data, connection, layer, path, route, multiple, different, plurality, message, confirm, acknowledge

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2009/0003216 A1 (Radunovic et al.) 01 January 2009 (01.01.2009), entire document, especially; abstract, para. [0024]-[0026], [0034], [0066], [0088], [0094]-[0096]	1 - 30
Y	US 2008/0049746 A1 (Morrill et al.) 28 February 2008 (28.02.2008), entire document, especially; abstract, para. [0008]-[0010], [0032], [0283]-[0286], [0322] [0326], [0341], [0342], [0385], [0391], [0416]	1 - 30
A	US 2009/0310582 A1 (Beser) 17 December 2009 (17.12.2009), entire document	1 - 30
A	US 2008/0043676 A1 (Mousseau et al.) 21 February 2008 (21.02.2008), entire document	1 - 30

Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel, or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search 05 March 2013 (05.03.2013)	Date of mailing of the international search report 12 APR 2013
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774