



- (51) **International Patent Classification:**
H04L 29/08 (2006.01) *H04L 12/70* (2013.01)
- (21) **International Application Number:**
PCT/IB2012/056877
- (22) **International Filing Date:**
30 November 2012 (30.11.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/311,018 5 December 2011 (05.12.2011) US
- (71) **Applicant:** LIVEQOS INC. [CA/CA]; 535 Legget Drive, Suite 500, Ottawa, Ontario K2K 3B8 (CA).
- (72) **Inventors:** WILLIAMS, Matthew, Robert; 7 Insmill Crescent, Kanata, Ontario K2T 1G5 (CA). VEMULAPALI, Mohan, Krishna; 144 Rodeo Drive, Nepean, Ontario K2J 5K7 (CA).
- (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(H))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(in))*

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

- (54) **Title:** METHODS FOR ACHIEVING TARGET LOSS RATIO




FIG. 1

- (57) **Abstract:** A method of forwarding data transmissions from a first network to a third network via a second network comprises receiving packets of a first type from the first network, segmenting each packet into packets of a second type that are transmitted to the second network, and producing and transmitting at least one encoded duplicate of each of the packets of a second type to the second network to allow a packet of the first type to be recreated in the event that not all the packets of the second type are received. In the event that a sufficient number of the packets of a second type and the encoded duplicate packets are not received to recreate the packet of a first type, the method determines a loss ratio that represents the number of packets not recreated relative to the number of packets transmitted during a selected time interval.



METHODS FOR ACHIEVING TARGET LOSS RATIO

FIELD OF THE INVENTION

[0001] The present invention relates to communication data networks. More specifically, the present invention relates to systems and methods for increasing the throughput of data transmissions through a network as seen from the edge of the network.

BACKGROUND OF THE INVENTION

[0002] To enable data communication from a source end-station or network to a destination end-station or network, packet-based networks break-up data streams into smaller packets of data. As these packets traverse a network, some of these packets can be lost due to congestion or other network limitations. This loss can have a tremendous impact on the applications leveraging the communication channel between the source and destination end-stations. Ideally, a network, from the point of view of many applications, must provide perfect performance with deterministic packet latency and no packet loss. However, the capital and operational cost to achieve perfect network performance is not practical for most service and enterprise network providers.

[0003] Accordingly, systems and methods are required which can be used with low cost networks to provide applications with a high network performance. One approach is to create new encoding protocol stacks which are installed at the end-stations to improve the response to loss and latency. However, this approach is non-trivial since all end-stations in the source and destination networks must be upgraded to use the new encoding protocol stacks.

[0004] Another approach uses network devices that intercept standard protocols, and an encoding protocol between the intercepting network devices, to recover from network loss. These devices are deployed in areas of the network where resident applications require better network performance than what is generally available in the network itself. Such devices are described in issued Patent 7,953,114, which is a continuation-in-part of U.S. Patent 8,009,696, which is a continuation-in-part of U.S. Patent 7,742,501.

[0005] The encoding protocol is intended to reduce loss. To achieve this goal, the encoding protocol increases the overall bandwidth required within the network. The increase in bandwidth can actually result in increasing loss instead due to network constraints. The intercepting network device needs to detect this issue and react to ensure that the communicating end-stations achieve the desired application performance.

SUMMARY OF THE INVENTION

[0006] In accordance with one embodiment, a method of forwarding data transmissions from a first network to a third network via a second network comprises receiving packets of a first type from the first network, segmenting each packet of a first type into packets of a second type, transmitting the packets of a second type to the second network, and producing and transmitting at least one encoded duplicate of each of the packets of a second type to the second network. The encoded duplicate packet allows a packet of the first type to be recreated in the event that not all the packets of the second type are received at the second network. In the event that a sufficient number of the packets of a second type and the encoded duplicate packets are not received to recreate the packet of a first type, the method determines a loss ratio that represents the number of packets not recreated relative to the number of packets transmitted during a selected time interval, dynamically adjusts the overhead required to produce the encoded duplicate packets to attempt to minimize the loss ratio, detects a significant increase in the determined loss ratio over a plurality of the time intervals, successively reduces the overhead required to produce the encoded duplicate packets in response to the detection of a significant increase in the determined loss ratio, and determines whether the reduced overhead requirement stabilizes the determined loss ratio and, if the answer is affirmative, terminates the successive reduction of the overhead requirement.

[0007] In one implementation, the overhead required to produce the encoded duplicate packets is reduced by adjusting an encoding level in the production of the encoded duplicate packets in response to the detection of a significant increase in the determined loss ratio over the plurality of time intervals. For example, the encoding level may be decreased sufficiently to stop increases in said loss ratio. The significant increase in the determined loss ratio over a plurality of the time intervals may be detected by determining the change in the loss ratio, or an average of the loss ratio over the time intervals, and comparing the change with a preselected threshold.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Embodiments of the present invention will now be described, by way of example only, with reference to the attached Figures, wherein:

[0009] FIG. 1 is a block diagram of an environment in which the invention may be practiced.

[0010] FIG. 2 is a block diagram illustrating the components in an intercepting network device used in FIG. 1.

[0011] FIG. 3 illustrates an example segmentation and encoding of a standard packet.

[0012] FIG. 4 illustrates a first example algorithm to evaluate loss.

[0013] FIG. 5 illustrates an example embodiment considering average loss.

[0014] FIG. 6 illustrates an example of encoding level back-off

[0015] FIG. 7 is a flow chart illustrating the steps executed when a significant increase in loss is detected and encoding level back-off is enabled.

[0016] FIG. 8 is a flow chart illustrating the steps executed to detect loss ratio stabilization after encoding level back-off is enabled.

[0017] FIG. 9 is a flow chart illustrating an alternate algorithm which is executed when a significant increase in loss is detected and encoding level back-off is enabled.

[0018] FIG. 10 is a flow chart of an example of restricting an encoding level due to burst loss.

[0019] FIG. 11 is a flow chart of an example of handling a restricted encoding level for the remaining time the encoding channel remains open.

[0020] FIG. 12 is a flow chart of an example of handling a restricted encoding level for a period of time d .

[0021] FIG. 13 is a flow chart of an example of handling a restricted encoding level. In this case, the encoding level is increase over a period of time to match the restricted encoding level.

[0022] FIG. 14 is a flow chart of an example of interleaved mode according to an embodiment of the present invention.

[0023] FIG. 15 is a flow chart of an example of multi-level interleaved mode according to an embodiment of the present invention.

[0024] FIG. 16 is a flow chart illustrating the steps executed when the intercepting network device is handling burst loss.

[0025] FIG. 17 is a flow chart illustrating the steps executed when the intercepting network device is handling burst loss and interleaved mode is enabled.

[0026] FIG. 18 is a flow chart illustrating the steps executed when the intercepting network device is leaving interleaved mode and random mode is enabled.

[0027] FIG. 19 is a flow chart illustrating an alternate algorithm which is executed when the intercepting network device is leaving interleaved mode and random mode is enabled.

[0028] FIG. 20 is a flow chart illustrating another alternate algorithm which is executed when the intercepting network device is leaving interleaved mode and random mode is enabled.

[0029] FIG. 21 is a flow chart illustrating another alternate algorithm which is executed when the intercepting network device is leaving interleaved mode and random mode is enabled.

[0030] FIG. 22 is a flow chart illustrating the steps executed when the intercepting network device is handling burst loss using interleaved mode and encoding level back-off.

[0031] FIG. 23 is a flow chart illustrating the steps to determine if encoding level back-off is required when handling burst loss.

[0032] FIG. 24 is a ladder diagram demonstrating TCP traffic flow between the network elements in FIG. 1.

[0033] FIG. 25 is a ladder diagram demonstrating the loss of a TCP segment in the TCP session in FIG. 24.

[0034] FIG. 26 is a flow chart defining how TCP segments are handled at the Decoder to determine goodput and throughput.

[0035] FIG. 27 is a flow chart which provides the algorithm for an IND encoder for splitting up a single TCP ACK into multiple ACKs.

[0036] FIG. 28 defines the mechanism to split a TCP ACK into multiple TCP ACKs in the form of a flowchart.

[0037] FIG. 29 continues the flowchart started in FIG. 27.

[0038] FIG. 30 is a ladder diagram breaking down the functions for calculating Round Trip Time (RTT) using a proprietary protocol embedded in the encoded segments.

[0039] FIG. 31 is a flow chart documenting the procedure to generate the RTT Query message required in FIG. 30 Step 3000.

[0040] FIG. 32 is a flowchart of how an IND handles a received RTT Query packet.

[0041] FIG. 33 is a ladder diagram breaking down the functions required to calculate Round Trip Time (RTT) using the information collected from a TCP session.

[0042] FIG. 34 demonstrates an example of the handling of the loss report associated with goodput and RTT measurements for TCP sessions.

[0043] FIG. 35 demonstrates an example of the handling of the loss report associated with loss and RTT measurements for UDP sessions.

[0044] FIG. 36 demonstrates an example handling of an encoding backoff.

[0045] FIG. 37 demonstrates an example of unlocking encoding levels.

DETAILED DESCRIPTION OF THE ILLUSTRATED EMBODIMENTS

[0046] Although the invention will be described in connection with certain preferred embodiments, it will be understood that the invention is not limited to those particular embodiments. On the contrary, the invention is intended to cover all alternatives, modifications, and equivalent arrangements as may be included within the spirit and scope of the invention as defined by the appended claims.

[0047] FIG. 1 illustrates a network environment in which a source end-station 10 and a destination end-station 30 have applications which communicate using a standard protocol (generally, packet-based). This protocol is carried through a network 20 which can be, for example, the Internet or some other communications network. The end-stations 10 and 30 can be terminals, such as personal computers connected through a modem, local or wide-area networks, or other devices or systems having access to the network 20. For certain applications resident on the end-stations 10 and 30, better network performance is required, and packets related to these applications are intercepted by network devices 40 and 50. Referred to as "intercepting" network devices ("INDs"), the network devices 40 and 50 encode the standard packets using an encoding protocol to increase network performance. To illustrate how this network operates, standard packet flow is assumed from the source end-station 10 to the destination end-station 30 (i.e., left to right in FIG. 1). The intercepting network device 40 encodes the standard packets from the source end-station 10 and forwards the encoded packets to the intercepting network device 50 through the network 20. The intercepting network device 50 decodes the packets and forwards the reconstituted standard packet to the destination end-station 30. In the reverse direction, packets flowing from the destination end-station 30 to the source end-station 10 are encoded by the intercepting network device 50 and forwarded to the intercepting network device 40 over the network 20. The intercepting network device 40 decodes these packets, and the reconstituted standard packets are forwarded to the source end-station 10.

[0048] FIG. 2 depicts an exemplary embodiment of the intercepting network device 40 that includes modules necessary to implement the encoding protocol to improve network performance. The modules can be resident in a single device, or can be distributed amongst several devices. In the illustrative embodiment of FIG. 2, the intercepting network device 40 includes a native interface 60, an encoded interface 70, a decoder 80, and an encoder 90. The

native interface 60 sends and receives standard packets to and from the source end-station 10, and the encoded interface 70 sends and receives encoded packets to and from the network 20.

[0049] The encoder 80 receives standard packets from the native interface 60 and produces encoded packets by dividing the packet into one or more segments which then are ready for transmission to the network 20 by way of the encoded interface 70. The decoder 90 receives encoded packets from the encoded interface 70 and produces standard packets for transmission to the source end-station 10 by way of the native interface 60.

[0050] To assist in recreating the standard packets from the encoded packets, the encoder 90 may also create extra encoded packets. These extra encoded packets are derived from the standard packets and encoded packets. The extra encoded packets assist the decoder in the destination intercepting network device 50 to recreate or reassemble the original standard packets in the event one or more encoded packets are lost during their transmission through the network 20.

[0051] The encoding of a standard packet according to one embodiment is illustrated in FIG. 3. To simplify the description, headers have been omitted from the original standard packet and the encoded packets. A standard packet 400 received at an encoder has a packet payload of 1 byte, and is segmented into n segments 402. Where n is chosen as an integer factor of 1, each of the n segments will have Vn bytes, as shown. In addition, m extra encoded packets are created. In the example shown in FIG. 3, $m=1$ and the additional or extra segment 404 is a parity segment created by applying a logical function (e.g., XOR function) to all of the n segments.

[0052] Once the encoder has completed the process of creating the segments and the extra encoded segment, headers are added to the packets to create encoded packets and extra encoded packets. The headers use a standard protocol such that no proprietary operational functions are expected in the network. Therefore, the encoded and extra encoded packets are carried through the network to the decoding intercepting network element. The encoder can take into account the size of the packet and can increase n automatically to avoid transmitting an encoded packet of a size that would exceed the Maximum Transfer Unit (MTU) of the network when the standard protocol header is added to the encoded packet. This capability can also be useful for splitting jumbo frames prior to entering a network that cannot handle them.

[0053] As a decoder reassembles standard packets from the encoded and extra encoded packets, it can deal with loss of encoded packets. If enough information is contained in the extra encoded packets, missing segments caused by lost encoded packets can be recreated. If

the loss is too great, the encoder has options. If the standard packet has a limited lifespan (such as a video packet), then the affected standard packet is discarded. If the application natively provides resiliency to packet loss, then again, the standard packet may be discarded. If the standard packet has a reasonable lifespan and the application does not provide resiliency natively, the decoder can request a retransmission of the missing segment from the encoder. This assumes that the encoder buffers the standard packet until it has been successfully decoded on the other end of the network.

[0054] To enable higher network performance for communicating applications, an encoding channel is established between intercepting network devices. Examples of how these channels are created are described in U.S. Application Serial No. 12/193,345. For the purpose of this application, an encoded channel between the intercepting network devices 40 and 50 is assumed to be successfully negotiated. The encoded channel is used to carry both data and control information, and encoded packets and extra encoded packets can now be sent and received between the intercepting network devices.

[0055] Applications residing on end-stations in the network above can tolerate different levels of lost packets within the communication network. The intercepting network devices can differentiate these applications based upon standard protocol signaling elements. File transfers using FTP are signaled using a TCP port. This will open TCP ports in the 64000 range to transport the files between the end-stations. TCP is resilient to loss and variation of latency within the network. The key goal of the transfer is to ensure that the file arrives intact and in as timely a manner as the network will allow. This application is in contrast to a video conference application. These can be signaled using the Session Initiation Protocol on UDP port 5060 and contains the information for the video, voice and data sessions associated with the conference. Video sessions have a low tolerance to loss; the loss of an eye frame (a complete picture of the screen) can greatly disrupt video compression algorithms. The greater the bandwidth required for a video stream, the lower the tolerance for loss since each standard packet is carrying more information related to the video stream. Voice sessions are relatively more tolerant to loss due to the predictive nature of human speech. Voice codecs are able to fill in gaps to ensure a continuous voice stream. In either case, video and voice sessions need to have less than 1% standard packet loss to avoid video blocking or voice clipping, and most providers attempt to keep standard packet loss at less than 0.5%.

[0056] Within the context on a given session, the different streams can be treated with different loss objectives. Deep Packet Inspection (DPI) can be used in the intercepting

network devices to detect the type of application (video/voice/data) and set the loss objectives according to a configured network policy for that type of application. Ultimately each standard packet can have its own loss objective based on information provided in the header that identifies the level of protection required (e.g., SVC —scalable video codec). This loss objective is defined as the Target Loss Ratio (TLR).

[0057] The TLR can be used to determine the encoding rate which is determined by both n and m . The encoding rate can be adjusted based on observed and/or reported network performance for one or all encoded channels terminating at an encoding component associated with a given intercepting network device. The goal is to set the encoding level such that the encoded channel achieves the TLR for the session or better.

[0058] Referring to FIG. 4, the determination of the loss ratio observed is described in relation to the encoded channel between intercepting network device A 1106 and intercepting network device B 1107, more specifically looking at the channel in the direction from intercepting network device A to intercepting network device B. Under a "no loss", or acceptable loss, condition, the encoding can be set to $n=\backslash$ and $m=0$ to avoid using extra bandwidth in the network when it is not necessary. The decoder 1109 of the intercepting network device B 1107 counts the number of packets received P_x and the number of packets lost LP_x over an interval of W units of time (for example $W=8$ seconds) using the subroutine illustrated in FIG. 4. The interval period can also be defined as the reception of a particular number of packets. Alternatively, LP_x can represent the number of retransmission requests, such that the successfully recovered loss is not counted as part of the loss ratio. The end of each interval is detected at step 1100, and the loss ratio L_x is calculated at step 1101 using, for example, the following ratio:

$$L_x = \frac{LP_x}{(LP_x + P_x)} \quad (1)$$

[0059] P_x in equation 1 can be the number of standard packets transmitted over the time period W . In this case, LP_x is the number of lost standard packets. This is referred to the SPL_x . Another option for P_x is to count the number of encoded segments and extra encoded segments which should have arrived at the decoder to re-assemble the standard packets over the time period W . LP_x then is the number of segments and extra segments which were lost in the network over the time period W . This is referred to the EPL_x . In general, SPL_x is lower than EPL_x due to the fact that the encoding scheme helps the decoder recover from loss, allowing

standard packets to be reproduced even though encoded packets related to the standard packet were lost. From that perspective, EPL_x is a truer indication of the performance of the network.

[0060] The current loss ratio (L_x) can refer to loss calculated from measurement of standard packets (SPL_x). It can refer to current loss calculated from encoded and extra encoded segments (EPL_x). It can also refer to a current loss ratio calculated from the combination of SPL_x and EPL_x . One example of this combination is:

$$CLR_x = \frac{\omega_s SPL_x + \omega_E EPL_x}{\omega_s + \omega_E} \quad (2)$$

where ω_s is the weight given to the loss of standard packets and ω_E is the weight given to the loss of encoded segments.

[0061] To avoid reacting to random events, step 1102 computes an average loss over the last z loss measurements. In order to take into account the most recent network status, a weighted average can be performed to weight the most recent current loss ratio measurements, e.g., using the equation:

$$WL_x = \frac{\sum_{i=l..z} CLR_{iXai}}{\sum_{i=l..z} \omega_i} \quad (3)$$

where WL_x represents the weighted average of the most recent loss measurements for interval x . The weights are such that $\omega_i < \omega_j$ for $i < j \leq x$.

[0062] To avoid keeping track of many loss measurements, the previous weighted loss can also be used to evaluate the new weighted loss, e.g., using the equation:

$$WL_x = \frac{(WL_{x-1Xmold}) + (CLR_{xXmnew})}{\omega_{old} + \omega_{new}} \quad (4)$$

where ω_{old} and ω_{new} are weights that are set in general such that $\omega_{old} < \omega_{new}$.

[0063] Optionally the weighted loss ratio can be normalized at step 1103 to simplify the lookup using integers only. The normalized loss NL_x can be computed using the following equation:

$$NL_x = \frac{NxWL_x}{P_x} \quad (5)$$

where N is the normalization factor (e.g., $N=10000$).

[0064] NL_x is then used to index an encoding level table at step 1104 to extract the current encoding level. An example of an encoding level table is shown below, providing 8 levels, where INT_max represents the largest integer or a large value (e.g., 10000).

NL_xmin	NL_xmax	Encoding Level
0	0	0
1	75	1
75	145	2
146	185	3
186	210	4
211	260	5
261	370	6
371	INT _{max}	7

[0065] The encoder 1111 of the intercepting network device B 1107 then inserts the extracted encoding level 1105 in the header of each encoded packet sent for that session to the intercepting network device A 1106. The decoder 1112 in the intercepting network device A 1106 reads the encoding level from the received packets. It then indexes a new parameter table with the encoding level to obtain the value of I and Max_n , which represents the largest value to which n should be set given the current encoding level. Multiple new parameter tables can be used to achieve different TLRs based on the application requirements and acceptable overhead. An example of such a new parameter table using 8 loss ratios is shown below:

Encoding Level	Max_n	m
0	INT _{max}	0
1	INT _{max}	1
2	INT _{max}	1
3	3	1
4	2	1
5	2	1
6	1	1
7	1	1

where INT_{max} represents the largest integer (infinity).

[0066] Before the encoder 1108 of the intercepting network device A 1106 encodes a packet of size s , it indexes a preconfigured packet size table with s , to obtain Rec_n, which represents the recommended value of n given the packet size. An example packet size table is shown below:

Packet Size (bytes)	Rec_n
< 88	1
89 to 264	2
265 to 528	3
> 528	4

[0067] Using this table, a packet of size $s < 88$ bytes returns $\text{Rec_}n=1$. If $s > 528$ bytes, then $\text{Rec_}n=4$. The value of n used to encode the packet is then determined as $n=\min(\text{Rec_}n, \text{Max_}w)$.

[0068] In another embodiment, the decoder 1109 calculates the weighted loss ratio WL_x using equation (3) or (4) above. It can optionally normalize to compute NL_x using equation (5). The exact loss value (WL_x or NL_x) is sent at regular intervals (e.g., every second) to the decoder 1112 in a control message that is inserted in the encoded channel for the session. The decoder 1112 extracts the exact loss values from the control messages. The encoder 1108 uses the exact loss value to index a new parameter table to obtain $\text{Max_}n$ and m .

NL_x min	NL_x max	$\text{Max_}n$	m
0	0	INT_max	0
1	75	INT_max	1
76	145	INT_max	1
146	185	3	1
186	210	2	1
211	260	2	1
261	370	1	1
371	10000	1	1

[0069] The value of n is derived as $n=\min(\text{Rec_}n, \text{Max_}w)$. Multiple new parameter tables can be configured to reflect different Target Loss Ratios, and the encoder 1108 uses the appropriate table based on the Target Loss Ratios of the application. Sending the actual loss ratio, instead of the current loss ratio, allows configuration the parameter tables at the encoding side, thus simplifying the configurations.

[0070] The previous equations demonstrate how the loss ratio is calculated for individual streams. Using FIG. 5 as an example, an intercepting network device D 1203 calculates the loss ratios of three flows 1204, 1205 and 1206 originating from intercepting network devices A 1202, B 1207 and C 1208, respectively. In addition to returning the loss ratio measured for the flow, the receiving intercepting network device can compute the average of the weighted loss ratio AL_x from the three flows that are sent to it over an interval period x .

[0071] The interval period can be the same as the interval period used for calculating the loss ratio for a single flow. The average loss ratio can be calculated, for example, as:

$$\text{AL}_x = \frac{\sum_{i=1..f} \text{WL}_i}{f} \quad (6)$$

where WLi can be calculated using equation (3) or (4) above. The AL_x can then be normalized using equation (5) above to create a normalized average loss ratio, NAL_x . The NAL_x is then indexed in an encoding level table to obtain an average loss ratio at the far end ($ALFE_x$). The encoding level table can be the same as illustrated above or preconfigured with different numbers. The $ALFN_x$ is added to the packet header of the encoded packets transmitted by the decoder of the intercepting network device D 1203 to the respective intercepting network devices 1202, 1207 and 1208 along with the per flow current loss ratio computed as described above.

[0072] If the aggregate information is included in the packet header, the intercepting network device 1202, 1207 or 1208 can use the information to decide whether to change its encoding level. The decoder of the intercepting network device A 1202 also calculates the average loss at the near end ($ALNE_x$) which is the average of the current loss ratio received in the encoding channel from the g encoded sessions 1211, 1212 and 1204 active during a measurement period x ,

$$ALNE_x = \frac{\sum_{i=1..g} CurrentLossRatio_i}{8} \quad (7)$$

[0073] In the example of FIG. 5, the $ALNE_x$ calculated at the intercepting network device A 1202 represents the average of the current loss ratio received from intercepting network devices D 1203, E 1209 and F 1210.

[0074] If the difference between the current loss ratio for the session and the $ALFE_x$ is below a predetermined threshold, the current loss ratio is used to set the encoding level, as described above. In this case, it is assumed that it is unlikely that the sessions are congesting the upstream network since the current loss ratio is better, or slightly worse than, the computed $ALFE_x$.

[0075] If the difference between the current loss ratio and the $ALFE_x$ is greater or equal to a predetermined threshold AND the current loss ratio minus the $ALNE_x$ is greater than or equal to a predetermined threshold, then the current loss ratio is ignored, and the encoding levels are set according to the packet size table only to minimize bandwidth usage by choosing the most bandwidth-efficient encoding method for the given packet. The predetermined thresholds could be different and set according to network policies and topology.

[0076] Irrespective of how n is determined (current loss or average loss), the increase in the value of n can be performed gradually to avoid a big step in the increase of overhead. When a higher value of n is recommended by the table, it can be applied to only a subset of the

subsequent standard packets: only v standard packets out of the next incoming w standard packets use the increased value of n , while the other $w-v$ packets use the previous lower value for n . The values of v and w can also change if the measured loss ratio continues to increase or as the measured loss ratio approaches the next level.

[0077] For example, if the measured loss (weighted or normalized) is 0%, then $n=l$, $m=0$ and $v=w=\backslash$. Therefore all packets are encoded with $n=\backslash$, $m=0$. If the measured loss increase to greater than 0% but lower than 0.05%, then $n=4$, $m=l$, but $v=\backslash$ and $w=3$, such that only one out of three packets is encoded with $n=4$, while the others use the previous encoding level $n=\backslash$, $m=0$. When the measured loss exceeds 0.05% but is below 0.1%>, then change to $v=\backslash$ and $w=2$, such that every second packet is encoded with $n=4$, while the others use the previous encoding level $n=\backslash$, $m=0$. When the measured loss exceeds 0.1%> but is below or equal to 0.2%>, then use $v=\backslash$ and $w=l$, such that every packet is encoded with $n=4$. Different values of v and w can be configured to smooth out the overhead increase at different predetermined loss ratios. This capability can significantly smooth out the transfer function between different loss ratios.

[0078] Intercepting network devices operate on the theory that network loss will decrease as the encoding level is increased. However, as n and/or m are increased, the number of packets generated by the intercepting network devices and, traversing the network, is increased. This results in an increase in the bandwidth required to achieve application communication.

[0079] The additional bandwidth may in some circumstances lead to an increase in loss. If one or more links in the path are bandwidth constrained, the additional packets actually increase congestion at this point. When a packet arrives at a link for transmission, it may or may not be able to be transmitted immediately. If the packet cannot be sent, most devices implement a queue to buffer the packet so that it can be transmitted later. In some devices, as the queues associated with this link grow to a particular level, a network congestion control mechanism (e.g., RED, WRED, etc.) may be invoked to handle the congestion on this link. As the number of packets arriving at this link is increased, the likelihood of one of the encoded packets being discarded is increased. The probability is also increased due to the additional overhead added by the encoding channel. The additional overhead increases the depth of the queue more quickly thus increasing the likelihood that a packet will be discarded.

[0080] Another cause of congestion can be switching context. A network device in the path between the source end-station and the destination end-station may only be able to forward a limited number of packets per period of time. Therefore, even though enough link

bandwidth is available to transport the encoded packets, the network device cannot forward all the frames, leading to loss.

[0081] Therefore, by increasing the encoding level which increases the number of packets in the path and the amount of overhead, the loss ratio may increase. The intercepting network device requires an algorithm to react to this situation. First, the intercepting device needs to detect a significant increase in loss ratio on an encoding channel. This can be done, for example, by tracking the results of the loss ratio over a period of time. This can be the current loss ratio as calculated by equation (1) or (2), the weighted loss ratio calculated by equation (3) or (4), the normalized loss ratio calculated by equation(5), or any other method to obtain loss.

[0082] One method of detecting a significant increase is to use the instantaneous change in the loss ratio. This looks at the difference between the current loss ratio and the previous loss ratio'.

$$ACLR_x = CLR_x - CLR_{x-1} \quad (8)$$

If this difference exceeds a significant increase threshold, defined as SIT, then a significant increase in loss ratio is detected. To avoid spuriously declaring that a significant increase has occurred, the number of encoded packets sent between the source intercepting network device and the destination intercepting network device must be statistically relevant. This is captured in the case equation below:

$$significant_increase = \begin{cases} true & \rightarrow packets \geq \min_relevant + ACIT \geq SIT \\ false & \rightarrow packets < \min_relevant \end{cases} \quad (9)$$

[0083] Instead of using an instantaneous ratio, the previous equations can use the average loss ratio (like equations (3) and (4)). The change can be calculated using equation (10):

$$AWLX_x = WLX_x - WLX_{x-1} \quad (10)$$

Then the significant increase can be determined as follows:

$$significant_increase = \begin{cases} true & \rightarrow AWL_x \geq SIT \\ false & \rightarrow AWL_x < SIT \end{cases} \quad (11)$$

[0084] Another approach is to record the loss ratio when an increase in the encoding level is enacted. The change in loss ratio ($CLR_{original}$) is recorded and used as the basis of comparison to the current instantaneous loss ratio. If the current loss ratio is, for example, ten times the original loss ratio, a significant increase in loss has been detected:

$$significant_increase = \begin{cases} true & \rightarrow CLR_x \geq 10 \times CLR_{original} \\ false & \rightarrow CLR_x < 10 \times CLR_{original} \end{cases} \quad (12)$$

[0085] Once a significant increase in the loss ratio has been detected, the intercepting network device reacts using an encoding level back-off scheme. In this case, the intercepting network device reacts to loss by decreasing the encoding level instead of increasing it. FIG. 6 is a flow chart of one technique for decreasing the encoding level in response to the detection of a significant increase in the loss ratio. First, step 1300 attempts to find an encoding level that stops the increase in loss ratio. Once this level is achieved, step 1301 ensures that the loss ratio is stable before returning the intercepting network device to normal processing of loss ratio.

[0086] FIG. 7 depicts one implementation of step 1300 in FIG. 6. This algorithm defines the steps to find an encoding level that stops the increase in loss ratio. The algorithm starts by recording the current loss ratio at step 1400, and then the encoding level of the channel is reduced at step 1401. The reduction of the encoding level can be effected in several ways. For example, the encoding level can be reduced linearly for each iteration of the algorithm, or it can be decreased exponentially. One implementation supports both methods and allows one of the methods to be selected when the encoding channel is established. Another example uses a configuration policy that is set before the encoding channel is created.

[0087] Once the encoding level has been reduced at step 1401, step 1402 determines whether the encoding level is greater than zero. A negative answer at step 1402 indicates that the current loss ratio is still higher than the starting loss ratio, which indicates that the encoding back-off algorithm is not working, and thus the algorithm is terminated at step 1403. The encoding channel then returns to normal loss ratio processing or reverts to $n=\backslash$ and $m=0$ and notifies the user and/or the system.

[0088] If the answer at step 1402 is affirmative, the system advances to step 1404 where the intercepting network device waits for a period of time d before sampling the current loss ratio again at step 1405. Step 1406 then determines whether the current loss ratio is greater than the starting loss ratio at step 1406, i.e., the loss ratio is still increasing by some predetermined amount, the algorithm returns from step 1406 to steps 1401 and 1402 to reduce the encoding level for the channel again. If the answer at step 1406 is negative, the system proceeds to step 1407 to determine if the current encoding level results in a stable loss ratio.

[0089] FIG. 8 depicts one implementation of step 1301 in FIG. 6. This algorithm defines the steps to determine if the current encoding level has resulted in a stable loss ratio. The algorithm runs in a loop which is initialized in step 1500. The intercepting network device then waits for a period d at step 1501 before sampling the current loss ratio again at step 1502. If the current loss ratio is significantly greater than the start loss ratio, the algorithm exits and returns to

backing off the encoding level at step 1503. Step 1503 determines whether the current loss ratio remains below the starting loss ratio (step 1503), then the algorithm increments the loop counter (step 1504) and continues to sample the current loss ratio until the maximum number of iterations has been reached (step 1505). At this point, the algorithm exits since the significant increase in loss ratio has been handled.

[0090] There are other implementations of the algorithm of FIG. 6. For example, FIG. 9 depicts an alternate implementation of the encoding level back-off scheme. In this case, the algorithm sets the starting loss ratio to the current loss ratio value at step 1600, and then step 1601 initially increases the encoding level to determine whether the loss ratio will be reduced. If step 1604 detects a reduction in the loss ratio, the algorithm exits and proceeds to ensure that the loss ratio is stable before returning to normal loss ratio processing. If step 1604 does not detect a reduction in the loss ratio, the algorithm proceeds to steps 1605 and 1606 to reduce the encoding level by 2 levels, and then to step 1602 to wait for a period before checking the current loss ratio again. If this succeeds (steps 1603 and 1604), again the algorithm ensures that the loss ratio is stable before returning to normal loss ratio processing. If step 1604 determines that the current loss ratio is still greater than the starting loss ratio and step 1605 determines this is not the first iteration, then the algorithm continues backing off the encoding level using the algorithm of FIG. 7.

[0091] Another approach to handling a significant increase in loss ratio (in step 1300 of FIG. 6) is to restrict the encoding level that is causing the significant increase, using the algorithm of FIG. 10. In step 1700, the current encoding level is recorded as the restricted encoding level, and then the current encoding level is reduced at step 1701. Then step 1702 ensures that the current encoding level remains below the restricted encoding level.

[0092] There are a number of ways to determine whether the restriction of the encoding level stabilizes the loss ratio, in step 1301 of FIG. 6. One way depicted in FIG. 11 is to simply exit the algorithm so that the encoding level is never used again while the encoding channel is active. FIG. 12 depicts an alternative technique that delays for a period of time d at step 1800 before re-enabling the encoding level for the encoding channel at step 1801. FIG. 13 depicts another alternative in which the encoding level is slowly adjusted back to the restricted encoding level. The period of time over which this occurs depends upon the delay period d set in step 1900, and the increment used to increase the encoding level in step 1902. Step 1901 compares the current encoding level with the restricted encoding level to determine whether the current encoding level is less than the restricted coding level. As long as the answer is

affirmative, the system advances to step 1902 to increase the coding level. Once the restricted level is achieved, the algorithm exits at step 1903.

[0093] Another strategy for handling a high level of network loss is to configure the encoder to transmit encoded packets in an interleaved manner (referred to as "interleaved mode"), thereby spreading the risk of losing multiple encoded packets across different standard packets and providing a higher probability of encoded packet re-assembly. Thus, instead of sequentially sending groups of encoded packets such that each group corresponds to a single standard packet, encoded packets from different standard packets are interleaved with one another. FIG. 14 illustrates an example of the interleaving mode. In this example, it is assumed that standard packets A, B, and C are respectively segmented into encoded packets A1, A2; B1, B2, B3; and C1, C2, C3 and C4. Encoded packets A2, B3 and C4 are extra encoded packets. Instead of sending these encoded packets grouped according to their respective standard packets, they are interleaved by groups of k standard packets. In the example of FIG. 14, $k=3$. The encoder stores the encoded packets corresponding to k standard packets into a buffer 1303. Alternatively, the encoder stores a number of packets that arrive during a predetermined time period or until a packet with time sensitive information is detected (e.g., via deep packet inspection). To avoid reordering, the set of encoded packets is aligned to the left such that the last second data of each standard packet is at the tail and is sent last, and the first encoded packet of each standard packet is at the head and is transmitted first. The encoder transmits on the physical interface 1304 the data units at the head in order of top to bottom, or any order that minimizes sending two consecutive encoded packets from the same standard packet. In the example of FIG. 14, the encoded packets can be sent in the order C1, B1, C2, A1, B2, C3, A2, B3, and C4. Sending all the extra encoded packets last can minimize delay. The interleaving can also be done randomly. The group of packets to interleave can include all the encoded packets transmitted from the interface, or the interleaving group can include encoded packets of the same application (e.g., same channels of a video conference), the same destination or a preconfigured grouping.

[0094] In FIG. 15, another level of interleaving is achieved, by interleaving the encoded packets of different interleaved groups 1401, 1402, 1403 of interleaved encoded packets. By such interleaving, the impact of any large loss of data units can be minimized, and, depending on the coding and strategy employed, this type of loss may be recoverable.

[0095] The "interleaved mode" can also be used to address burst loss. Burst loss can be detected if the current loss ratio (i.e., the result of equation (1), (2), (3) or (4)) exceeds the

Target Loss Ratio by a predetermined amount (e.g., double). One technique for handling burst loss is illustrated in FIG. 16. When burst loss is initially detected, the loss ratio is debounced. Assuming that the burst loss persists, then a transition to the interleaved mode is executed at step 2000. If the interleaved mode is enabled, the conditions to return to the random mode are checked and verified at step 2001.

[0096] An example of how to transition to the interleaved mode is illustrated in FIG. 17. This algorithm debounces the current loss ratio to ensure that the issue persists before enabling interleaved mode. The implementation uses a loop that is initialized at step 2100 and then delays for a time period of d at step 2101. The current loss ratio is retrieved at step 2102 and checked at step 2103 to determine whether the current loss ratio exceeds twice the Target Loss Ratio, which is acceptable for this standard protocol type. An affirmative answer at step 2103 causes the loop counter to be incremented at step 2104, and then step 2105 determines whether the loop count has reached the maximum iteration. If the answer is negative, the system returns to step 2101 and repeats steps 2102 and 2103. A negative answer at step 2103 ends the routine, and an affirmative answer at step 2105 causes the interleaved mode to be enabled at step 2106.

[0097] Thus, assuming that the loss ratio is maintained for delay period d times the max iterations, the queuing is set to the interleaving mode, and the algorithm continues on to handle the return to the random mode. If the loss ratio does not continue to exceed twice the TLR for the application, the algorithm exits and the interleaved mode is disabled.

[0098] Returning to the random mode can be implemented by any of several different methods illustrated in FIGs. 18-21. The method actually used for an encoding channel can be selected in several ways. The intercepting network device may be configured with the method to use for all encoding channels. Also, the intercepting network device may have a policy that matches the standard protocol being encoded. Alternatively, the intercepting network devices may negotiate the method to be used when the encoding channel is established.

[0099] FIG. 18 illustrates the most straightforward method for returning to the random mode. In this case, a timer is started at step 2200 when the encoded channel has entered the interleaved mode. When the timer expires, the channel returns to the random mode at step 2201.

[00100] In FIG. 19, the algorithm ensures that the interleaved mode has an effect on the current loss ratio before queuing is returned to the random mode. The current loss ratio must be lower than twice the Target Loss Ratio for a time period of d times the max iterations before

the random mode is enabled. To complete this, a loop is used to sample the current loss ratio. The loop is initialized at step 2300 and then waits for a delay period of d at step 2301. The current loss ratio is retrieved at step 2302 and checked at step 2303 to determine whether the current loss ratio exceeds twice the Target Loss Ratio. An affirmative answer at step 2303 causes the loop counter to be incremented at step 2304, and then step 2305 determines whether the loop count has reached the maximum iteration. If the answer is negative, the system returns to step 2301 and repeats steps 2302 and 2303. An affirmative answer at step 2305 causes the random mode to be enabled at step 2306.

[00101] FIG. 20 is similar to FIG. 19 with the exception that instead of using the current loss ratio, the average loss ratio is used. The loop is initialized at step 2400 by setting the loop counter to zero and setting the Start Average Loss to the Current Average Loss and, the system then waits for a delay period of d at step 2401. The current average loss ratio is retrieved at step 2402 and checked at step 2403 to determine whether the current average loss ratio exceeds the Start Average Loss. An affirmative answer at step 2403 causes the loop counter to be incremented at step 2404, and then step 2405 determines whether the loop count has reached the maximum iteration. If the answer is negative, the system returns to step 2401 and repeats steps 2402 and 2403. An affirmative answer at step 2405 causes the random mode to be enabled at step 2406.

[00102] FIG. 21 is similar to FIG. 20, except that changing the encoding level is also allowed. The loop is initialized in step 2500 by setting the loop counter to zero and setting the Start Average Loss to the Current Average Loss and, the system then waits for a delay period of d at step 2501. The current average loss ratio is retrieved at step 2502 and checked at step 2503 to determine whether the current average loss ratio exceeds the Start Average Loss. A negative answer at step 2503 adjusts the encoding level at step 25607 to handle the current average loss ratio and then returns to step 2500. An affirmative answer at step 2503 causes the loop counter to be incremented at step 2504, and then step 2505 determines whether the loop count has reached the maximum iteration. If the answer is negative, the system returns to step 2501 and repeats steps 2501-2503. An affirmative answer at step 2505 causes the random mode to be enabled at step 2506.

[00103] FIG. 22 illustrates an algorithm that combines the handling of both significant changes in the current loss ratio or burst loss exceeding twice the TLR. In this embodiment, the algorithm starts at step 3000 by determining if the encoding channel needs to transition to the interleaved mode. Assuming this is required, the algorithm then checks to see if encoder level

back-off is required at step 3001. Regardless of whether back-off is required, the algorithm then determines when the encoding channel transitions back to the random mode at step 3002.

[00104] Handling the transition to the interleaved mode in step 3000 can be handled in the same way as in FIG. 18. Step 3002 looks at the transition back to the random mode, using any of the algorithms illustrated in FIGs. 18-21.

[00105] Determining whether back-off is required at step 3002 is illustrated in FIG. 23. This algorithm debounces the effect of transitioning from the random to the interleaved mode. This measurement should be done after an increase in encoding level. It loops for a period of delay d times max iterations and checks the current loss ratio. The loop, initialized at step 3100, has a delay period of d at step 3101. When this delay period expires, the current loss ratio is sampled at step 3102 and compared to twice the TLR at step 3103. If the current loss ratio is greater than twice the TLR, the exceed counter is incremented at step 3104. Otherwise, the loop counter is incremented at step 3105. This process continues until the loop counter reaches max iterations at step 2506. At this point, the exceed counter is checked at step 2507, and if it equals max iterations, encoding level back-off is enabled at step 2508. Otherwise, the algorithm proceeds on to handle the return to the random mode.

[00106] Some protocols, such as TCP, natively provide reliability and retransmission capabilities. In this case, the INDs can use native protocol information to determine loss and appropriate encoding levels. In the case of TCP, the stream of bytes travelling from a source end-station to a destination end-station is numbered using a sequence number. As these bytes are received by the destination end-station, the sequence of bytes is acknowledged. FIG. 24 provides an example of one such flow leveraging the network elements provided in FIG. 1. This flow begins with the destination end-station 30 advertising a receive window size of 3500 bytes to the source end-station 10 in message 2400. Message 2400 is part of the flow control mechanism provided by TCP to ensure that the destination end-station 30 is not overwhelmed with data from the source end-station 10. When the message 2400 arrives at the source end-station 10, three TCP segments 2410 to 2412 are sent to the destination end-station 30 to fill the destination end-station 30's open receive window. Source end-station will hold onto the segment 2410 to 2412 in case these segments need to be retransmitted to the destination end-station 30. These segments are held until they are acknowledged by the destination end-station 30 as is shown in FIG. 24 in TCP ACK messages 2420 to 2422. These TCP ACK messages indicate that the receive window in the destination end station 30 is being filled with the TCP segments as they are received. When the source end-station 10 receives these acknowledgements, the associated

segment is released and the source end-station is free to send more segments as the receive window opens again starting this process all over again.

FIG.25 is based upon FIG. 24 and explores how the destination end-station 30 reacts when a segment is lost. In this case, the first TCP segment sent by the source end-station 10 (2410) is lost. When the TCP segment 2411 is received by the destination end-station 30, it realizes that it has not received TCP segment 2410 and responds by sending a TCP ACK 2501 with a sequence number of 1. The destination end-station repeats this process with every segment it receives until the missing segment is recovered. When the TCP ACK 2500 is received by source end-station 10, it retransmits the first TCP segment in message 2510. The INDs in this figure, network element 40 and 50, can monitor the transactions between source end-station 10 and destination end-station 30 and infer the network state through the TCP segment and TCP ACK messages. In one embodiment, the Decoder 80 in the IND 50 analyzes the TCP state of the connection between the source end-station 10 and the destination end-station 30. Using the sequence number in the TCP segments, a measurement of the throughput and the goodput of the TCP connection can be ascertained. Throughput is the bandwidth used to transmit all the segments received to the destination end-station 30. Goodput is the bandwidth used to transmit unique segments to the destination end-station 30; this bandwidth does not include retransmitted segments. These two measurements can be used to derive a loss measurement such as the example shown in Equation 13.

$$CLR_x = 1 - \frac{Goodput_x}{throughput_x} \quad (13)$$

Once a loss measurement L_x is derived, the subsequent equations and algorithms defined in this patent can be applied to this measurement to define an encoding level for this TCP session.

[00107] Once encoding has been applied, GOODPUT needs to be watched. If the GOODPUT is reduced due to the introduction of encoding, then the encoding level is reduced to see if the change improves GOODPUT. This value provides natural feedback to the encoding algorithm to ensure that the correct level is applied to the session. This assessment is independent of the encoding level algorithm.

[00108] FIG. 26 provides an outline of an example algorithm at the Decoder 80 in IND 50. This algorithm uses three variables. The first variable is the expected sequence number (EXP_SEQ_#) which holds the current progress of the TCP byte stream as seen by the IND. The second and third variables, GOODPUT and THROUGHPUT, maintain byte

counts for calculating the goodput and throughput for the TCP session during this interval. The algorithm starts by checking for a sequence number wrapping in steps 2601, 2602 and 2603. As long as the result of the step 2601 is greater than the result of the step 2602, the sequence number received in the TCP segment (SEQ #) is later than the expected sequence number (EXP_SEQ_#). This is checked in step 2603. If the SEQ # comes later, then the expected sequence number is set to SEQ # (step 2604) and GOODPUT and THROUGHPUT are incremented by the segment size (steps 2605/2606). If the packet is a retransmission as identified in step 2603, only THROUGHPUT are incremented by the segment size (step 2606). Once the interval completes, L_x is calculated using Equation 13. Then the variables for GOODPUT and THROUGHPUT are reset to zero and the process of counting bytes using the algorithm in FIG. 26 is resumed.

[00109] The INDs, embodying the algorithms, have a view of the current state of the network 20 which can be leveraged to improve network utilization. As network resources become free or network performance significantly improves, the INDs can influence the end-stations to increase their TCP windows. Stimulus for this activity can be the closure of TCP sessions that were using a significant amount of THROUGHPUT or experienced a significant increase in GOODPUT by the TCP sessions between the INDs 40 and 50. INDs can only perform this function when encoding is not being applied to the TCP session. If the TCP session begins experience a higher level of loss and encoding is applied, the INDs stops trying to open the end-stations TCP window since this can lead to an increase in loss.

[00110] It may be beneficial to stimulate the congestion window on the source end-station through its additive increase / multiplicative decrease algorithm to improve the performance of the transfer. One mechanism to stimulate an end-station to increase its TCP window is to split TCP ACKs into multiple TCP ACKs. TCP implementations increase the congestion window by a fraction of the TCP maximum segment size for every TCP ACK which is received. By increasing the number of TCP ACKs, the congestion window opens faster putting more TCP segments in flight.

[00111] An algorithm that can be used to implement this function is shown in FIGs. 27, 28 and 29. These algorithms are implemented to control the amount network 20 resources used to stimulate the source end-station. One mechanism which is used in FIG. 27 to control bandwidth is to define a duty cycle for the algorithm. This means that the encoding algorithm is enabled and disabled for defined periods. One method is to define these periods in terms of time defining a total time period and a duty factor in percent. The duty factor

defines the percentage of the total time period, which the encoding algorithm is operating. Another option is to define the time periods in terms of packets. The total period is defined as `PACKET_PERIOD`. Within the `PACKET_PERIOD`, `ENCODE_LIMIT` defines the number of TCP ACK packets which are split into multiple TCP ACK packets. The packet method is used in FIG. 27.

[00112] FIG. 27 covers the generic algorithm for splitting TCP ACKs. It starts by inspecting the `LAST_VALID_ACK` parameter (step 2701). This is used as the basis for splitting up the TCP ACK packet before it is transmitted and ensures that each TCP ACK packet is carrying valid information. Therefore, if this parameter has not been initialized, it is set to the current sequence number, `CURRENT_SEQ` (step 2706), the packet period counter, `ACK_COUNT` is initialized to zero (step 2709) and the TCP ACK packet(s) are transmitted (step 2707). Assuming the `LAST_VALID_ACK` parameter has been initialized, the current TCP ACK packet sequence number is compared to the `LAST_VALID_ACK` (step 2702). If the current TCP ACK packet sequence number is equal to the `LAST_VALID_ACK` parameter, the associated TCP connection has experienced loss and that situation is being signaled to the back to the source end-station 10. To ensure that network utilization does not suffer as result of this segment loss, the TCP ACK packet needs to get back to the source end-station to initiate the TCP fast recovery algorithms. Therefore, this TCP ACK packet is sent immediately (step 2707). If `CURRENT_SEQ` is not equal to the `LAST_VALID_ACK`, the `ACK_COUNT` is incremented (2703) and compared to the `ENCODE_LIMIT` (step 2704). If the `ACK_COUNT` is less than or equal to the `ENCODE_LIMIT`, the TCP ACK packet is split up (step 2705), the `LAST_VALID_ACK` is set to the `CURRENT_SEQ` (step 2706) and the TCP ACK packets are sent (step 2707). If the `ACK_COUNT` is greater than `ENCODE_LIMIT` (step 2704), then the `ACK_COUNT` is compared to the `PACKET_PERIOD` (step 2708). If the `ACK_COUNT` is less than the `PACKET_PERIOD`, then `LAST_VALID_ACK` is set to the `CURRENT_SEQ` (step 2706) and the TCP ACK packets are sent (2707). Otherwise, the packet period is reset by setting `ACK_COUNT` to 0 (step 2709).

[00113] A method of splitting a TCP ACK packet is described in FIG. 28 and is continued in FIG. 29. These algorithms also control bandwidth utilization by controlling the number of TCP ACK packets which are produced from a single TCP ACK packet. The first parameter is the `MAXIMUM_ACKs`; this parameter defines the maximum number of TCP ACKs which can be generated from a single TCP ACK. The `MINIMUM_SEGMENT_SIZE`

determines the minimum difference between the sequence numbers of sequential TCP ACKs. For example, MAXIMUM_ACKs can be defined as 3 and MINIMUM_SEGMENT_SIZE can be defined as 500. If a packet arrives when the sequence difference is between the current TCP ACK and the previous TCP ACK seen by the encoder is 2100 bytes, the encoder produces three TCP ACKs each with a sequential difference of 700 bytes each. This encoding is limited by the MAXIMUM_ACK parameter. If a TCP ACK arrives with a sequential difference of 400 bytes, it is transmitted unchanged since it is lower than the MINIMUM_SEGMENT_SIZE.

[00114] FIG. 28 starts by calculating the difference between the current sequence number, CURRENT_SEQ, and LAST_VALID_ACK (step 2801). If the DIFFERENCE is less than the MINIMUM_SEGMENT_SIZE (step 2802), only 1 TCP ACK packet is produced (step 2803); this is achieved by setting NUM_ACKS to 1 and SEQUENCE_INTERVAL to DIFFERENCE. If the DIFFERENCE is greater than MINIMUM_SEGMENT_SIZE (step 2802), the number of TCP ACK packets is determined by dividing the DIFFERENCE by the MINIMUM_SEGMENT_SIZE (step 2804). If NUM_ACKS exceeds the MAXIMUM_ACKs (step 2805), then NUM_ACKS is set to MAXIMUM_ACKs and the SEQUENCE_INTERVAL is set to the DIFFERENCE divided by MAXIMUM_ACKs (step 2806). Otherwise, the SEQUENCE_INTERVAL is set to MINIMUM_SEGMENT_SIZE (step 2807).

[00115] Now that the number of TCP ACK packets and the sequence interval between the TCP ACKs has been determined, the TCP ACK packet can be segmented. This algorithm is described in FIG. 29. To separate the TCP ACK packet into multiple TCP ACK packets, a loop counter, CURRENT_ACK is initialized to 1 and the LAST_SEQ is set to the LAST_VALID_ACK (step 2901). Then the loop condition is checked by comparing CURRENT_ACK to NUM_ACKS, the total number of TCP ACKs to be generated (step 2902). If CURRENT_ACK does not equal NUM_ACKS, then the sequence number of the current TCP ACK packet is calculated using the LAST_SEQ number plus the SEQUENCE_INTERVAL (step 2903) and this TCP ACK packet is generated (step 2904). Then the TCP ACK packet is added to the list of TCP ACK packets to be sent (step 2905) and the CURRENT_ACK counter is incremented (step 2906). This procedure repeats until CURRENT_ACK equals NUM_ACKS (step 2902). At this point, the loop exits and the final TCP ACK packet is generated by setting the sequence number of this packet to the

CURRENT_SEQ (step 2907) and adding this packet to the list of TCP ACK packets to be transmitted (step 2908).

[00116] The stimulation of the end stations in a TCP session may continue while the measured goodput and throughput meet specified targets. If the loss measurements derived from these measurements exceed a specified threshold or the GOODPUT measurement decreases, the algorithms defined in FIGs. 27, 28 and 29 are discontinued until network conditions return to a level, which allow for the INDs to start stimulating the end-stations again.

[00117] Another indicator of network performance is the Round Trip Time (RTT) of a piece of information. RTT is the length of time it takes for a piece of information to be sent plus the length of time it takes for an acknowledgment of that information to be received. How this time changes over the lifetime of a communication session reflects the networks ability to deliver information (packets or segments) related to a session. Jitter is another network performance measure of interest. The jitter can be defined as the variation in RTT or in one-way delay, that can affect play-back buffer for real-time applications.

[00118] An IND can measure RTT is several ways. For real-time sessions such as VoIP or IPTV, no acknowledgement is provided for information, which passes from the source end station to the destination end station. Therefore, to achieve a measurement, the INDs need to inject packets, henceforth referred to as RTT Query packets, into the session to achieve a measurement. FIG. 30 provides a sample network where such measurements may occur. This figure builds upon the network in FIG.1 to examine a particular flow of packets for RTT measurement. In step 3000, IND 40, henceforth referred to as the originating IND, measures the RTT for a particular session and, to start the measurement, it injects a RTT Query packet towards its peer IND 50 in step 3001. The RTT Query packet is detected by the IND 50 (step 3002) and, since the source of the query is the IND 40, the packet is returned to the originating IND. When the originating IND 40 detects the packet, it can calculate the current RTT associated with this query.

[00119] FIG. 31 shows how an originating IND (such as originating IND 40) decides to send an RTT query packet. In this embodiment, a delay timer is used and an RTT Query is sent every time the delay timer expires for the lifetime of the encoding session. This procedure starts by executing the delay timer in step 3100. Once this timer expires, the originating IND prepares to send the RTT Query. First, it allocates a sequence identifier in step 3101 such that the specific query and its associated information can be located when the

RTT Query returns. Then, the originating IND gets a timestamp to record when the RTT Query is sent in step 3102 and then it proceeds to send the packet 3103. The originating IND then records all the data associated with this query and stores this for processing when the RTT Query returns in step 3104. Alternatively, the timestamp is embedded in the RTT Query packet and returned by then receiving IND. The RTT value is stored in memory 3105 for future processing (trending or averaging). The final step in this process is to check to see if the session is still valid or if there are enough RTT computations (step 3106). If the session is still valid, but RTT need to be computed, the process returns to step 3100 and the process starts all over again. If not, the process exits.

[00120] Other embodiments to handle the generation of RTT Queries are envisioned. For example, a query can be generated after a number of packets have been transmitted. Another option is to generate a query after number of bytes has been transmitted by the communication session. Another embodiment can use a combination of the delay timer and the byte threshold to measure RTT. This method ensures that measurements are taken at the minimum frequency of the delay timer if a direction of a session becomes silent. This method can be used for VoIP sessions, which support silence suppression.

[00121] FIG. 32 demonstrates how INDs handle the reception of an RTT Query. This logic is implemented in the decoder 80 of the receiving IND. If the received packet is not from a valid encoding stream 3201, it is ignored at 3203, otherwise if the packet is an RTT Query (step 3200) the packet is returned to the originating IND (step 3202) including all data and meta-data included in the packet, otherwise, the packet is decoded and forwarded to the end station.

[00122] If the communication session includes identification data that allows correlation of information in each direction (e.g., TCP/IP), the originating IND can leverage this identification to measure RTT. For example, the originating IND can leverage the TCP sequence numbers to perform measurements. It can correlate the sequence numbers in the segments sent in the encoding direction with the TCP ACKs which are returned in the decoding direction. FIG. 33 demonstrates how the originating IND can perform this function. This ladder diagram starts with the originating IND 40 detecting a TCP segment (step 3301) being transmitted from the source end station 10 to the destination end station 30 through the network 20 (step 3300). When this is detected, the originating IND 40 records information related to this segment, such as the original TCP timestamp provided by the TCP stack when the TCP packet is created and then forwards the segment towards the IND 50 and the

destination end station 30 (step 3302). The forwarding time is stored along with the original TCP timestamp for later processing. End station 30 acknowledged the receipt of this segment using a TCP ACK (step 3303) which is detected by the originating IND 40 (step 3305). Originating IND 40 can now calculate the RTT time for the segment identified by the original TCP timestamp stored in the TCP ACK along with the forwarding time previously stored and forwards this packet onto the source end station 10.

[00123] This embodiment calculates the RTT for every TCP segment. Other embodiments could limit the number of RTT calculations by using a delay timer to hold RTT calculations until the timer expires. Another embodiment could only take an RTT calculation after a set number of segments or a set number of bytes have been transferred.

[00124] Whichever method is used to gather the samples of RTT, a predetermined amount of RTT measurements r are collected by the originating IND to determine of the network performance between the IND 40 and the IND 50. If the network performance is deemed acceptable, the encoding can be set to $n=1$ and $m=0$ to avoid using extra bandwidth in the network when it is not necessary. The new sample is averaged with $(f-1)$ previous samples to determine the recent RTT for the session :

$$AverageCurrentRTT_x = \frac{\sum_{i=1}^r RTT_i}{r} \quad (14)$$

The current RTT sample is also added to calculate the average RTT over the lifetime of the communication session (step 3602):

$$AverageRTT_x = \frac{\omega_{new} RTT_x + \omega_{old} AverageRTT_{x-1}}{\omega_{new} + \omega_{old}} \quad (15)$$

where ω_{old} and ω_{new} are weights that are set in general such that $\omega_{old} > \omega_{new}$.

[00125] Fig. 34 provides an example of an algorithm that can be used to change the encoding level of a TCP session when a new encoding level is recommended as described above and in related applications. If the recommended level is lower than the current level 3402, the current encoding level is changed and set for a period of z units of time 3413, otherwise, if the recommended encoding level is locked 3403, then the recommended

encoding level is set to the highest unlocked level 3404. Goodput measurements, referred to herein as "pre-goodput" measurements (as exemplified in Fig. 25) are performed for w units of time 3405. The encoding level is then changed to the recommended level 3406. Goodput measurements, referred to herein as "post-goodput" measurements (as exemplified in Fig. 25) along with RTT measurements (as per example in Figs. 31 and 33) are performed for y units of time 3405. The pre-goodput, post-goodput and RTT values are processed to determine an average pre-goodput, average post-goodput and trending on the RTT 3408. If the RTT increases 3409 and if the average post-goodput is greater or equal to the average pre-goodput 3410, then more post-goodput and RTT measures are taken 3407. If the RTT is decreasing or constant 3409 and if the average post-goodput is greater or equal to the average pre-goodput 3411 then the encoding level is set to the recommended encoding level for a period of z units of time. If the post-goodput is lower than the pre-goodput 3410, 3411 then the backoff algorithm (Fig. 36) is invoked.

[00126] Fig. 35 provides an example of an algorithm that can be used to change the encoding level of a non-TCP session (e.g. UDP). When a new encoding level is recommended as described above and in related applications. If the recommended level is lower than the current level 3502, the current encoding level is changed and set for a period of z units of time 3513, otherwise, if the recommended encoding level is locked 3503, then the recommended encoding level is set to the highest unlocked level 3504. Loss measurements, referred to herein as "pre-loss" measurements (as exemplified in Fig. 25) are performed for w units of time 3505. The encoding level is then changed to the recommended level 3506. Loss measurements, referred to herein as "post-loss" measurements (as exemplified in Fig. 25) along with RTT measurements (as per example in Figs. 31 and 33) are performed for y units of time 3505. The pre-loss, post-loss and RTT values are processed to determine an average pre-loss, average post-loss and trending on the RTT 3508. If the RTT increases 3509 and if the average post-loss is smaller than the average pre-loss 3510, then more post-loss and RTT measures are taken 3507. If the RTT is decreasing or constant 3509 and if the average post-loss is smaller to the average pre-loss 3511 then the encoding level is set to the recommended encoding level for a period of z units of time. If the post-loss is higher or equal than the pre-loss 3510, 3511 then the backoff algorithm (Fig. 36) is invoked.

[00127] Fig. 36 exemplifies a backoff mechanism invoked 3600 when a recommended level can't be used for network performance reasons (e.g. Fig. 34 and 35). In

this case the current encoding for the session is set to a lower level (e.g. $n=0$, $m=1$) and the recommended level and all levels above are banned for that session 3602. A timer is set to a pre-determined value 3603 and when it expires 3604, the lowest locked encoding level is unlocked 3605 until all levels are unlocked 3606.

[00128] Fig. 37 provides an example mechanism used to unlock banned encoding levels for a session. When the encoding device detects that a flow or session has stopped 3701, the encoding device selects one active session, which has one or more encoding levels locked or banned (e.g., Fig. 36) 3702. The selection can be based on a policy, priority or randomly. The encoding device unlocks one or more levels for the selected session 3703, starting from the lowest locked level.

[00129] In the following description, for purposes of explanation, numerous details are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention. In other instances, well-known electrical structures and circuits are shown in block diagram form in order not to obscure the present invention. For example, specific details are not provided as to whether the embodiments of the invention described herein are implemented as a software routine, hardware circuit, firmware, or a combination thereof.

[00130] Embodiments of the invention may be implemented in a network having endpoints, such as servers or other computing devices, and associated encoding components. The encoding components, and the described methods, can be implemented in hardware, software or a combination thereof. Those portions that are implemented in software can be represented as a software product stored in a machine-readable medium (also referred to as a computer-readable medium, a processor-readable medium, or a computer usable medium having a computer readable program code embodied therein). The machine-readable medium may be any suitable tangible medium, including magnetic, optical, or electrical storage medium including a diskette, compact disk read only memory (CD-ROM), memory device (volatile or non-volatile), or similar storage mechanism. The machine-readable medium may contain various sets of instructions, code sequences, configuration information, or other data, which, when executed, cause a processor to perform steps in a method. Those of ordinary skill in the art will appreciate that other instructions and operations necessary to implement the described embodiments may also be stored on the machine-readable medium. Software running from the machine readable medium may interface with circuitry to perform the described tasks.

[00131] While particular embodiments and applications of the present invention have been illustrated and described, it is to be understood that the invention is not limited to the precise construction and compositions disclosed herein and that various modifications, changes, and variations may be apparent from the foregoing descriptions without departing from the spirit and scope of the invention as defined in the appended claims.

WHAT IS CLAIMED IS:

1. A method of transmitting TCP data segments via an encoded TCP channel from a source end-station to a destination end-station, said method comprising
establishing an encoded TCP channel between first and second intercepting network devices for a TCP session between said end-stations,
intercepting a TCP data segment related to said TCP session at one of said intercepting network devices,
segmenting and encoding said intercepted TCP data segment at a selectable encoding level in one of said intercepting network devices and transmitting the encoded segments to the other of said intercepting network devices,
calculating a network performance measurement for said TCP session between said end-stations, and
selecting said encoding level on the basis of said calculated network performance measurement for said TCP session.
2. The method of claim 1 which includes decoding said encoded segments at said other intercepting network device, and transmitting said TCP data segment from said other intercepting network device to said destination end-station.
3. The method of claim 1 in which said network performance measurement is at least one measurement selected from a group consisting of throughput and goodput in a session, wherein throughput is the bandwidth used to transmit all the segments received to said other intercepting network device, goodput is the bandwidth used to transmit unique segments to said other intercepting network device, and a session is the duration of transmitting a TCP data segment from said source end-station to said destination end-station.
4. The method of claim 3 in which said network performance measurement includes both throughput and goodput measurements, a change in loss ratio is determined from said throughput and goodput measurements, and said encoding level is defined in response to said change in loss ratio.
5. The method of claim 4 in which said performance measurement includes a measurement of goodput, and said encoding level is modified in response to a change in goodput.

6. A data transmission network for transmitting TCP data segments via an encoded TCP channel from a source end-station to a destination end-station, said network comprising

said first and second intercepting network devices establishing an encoded TCP channel for a TCP session between said end-stations,

one of said intercepting network devices being adapted to

intercept a TCP data segment related to said TCP session,

segment and encode an intercepted TCP data segment at a selectable encoding level, transmit the encoded segments to the other of said intercepting network devices,

calculate a network performance measurement for a TCP session between said end-stations, and

select said encoding level on the basis of said calculated network performance measurement for said TCP session.

7. The data transmission network of claim 6 in which said other intercepting network device decodes said encoded segments at said other intercepting device, and transmits said TCP data segment to said destination end-station.

8. The data transmission network of claim 6 in which said network performance measurement is at least one measurement selected from a group consisting of throughput and goodput in a session, wherein throughput is the bandwidth used to transmit all the segments received to said other intercepting network device, goodput is the bandwidth used to transmit unique segments to said other intercepting network device, and a session is the duration of transmitting a TCP data segment from said source end-station to said destination end-station.

9. The data transmission network of claim 8 in which said network performance measurement includes both throughput and goodput measurements, a variation in loss ratio is determined from said throughput and goodput measurements, and said encoding level is defined in response to said variation in loss ratio.

10. The data transmission network of claim 9 in which said performance measurement includes a measurement of goodput, and said encoding level is reduced in response to a reduction in goodput.

11. An intercepting network device for use in establishing an encoded TCP channel for a TCP session between a source end-station to a destination end-station in a data

transmission network for transmitting TCP data segments, said intercepting network device comprising a processor programmed to

intercept a TCP data segment related to said TCP session,

segment and encode an intercepted TCP data segment at a selectable encoding level, transmit the encoded segments to the other of said intercepting network devices,

calculate a network performance measurement for a TCP session between said end-stations, and

select said encoding level on the basis of said calculated network performance measurement for said TCP session.

12. The intercepting network device of claim 11 in which said network performance measurement is at least one measurement selected from a group consisting of throughput and goodput in a session, wherein throughput is the bandwidth used to transmit all the segments received to said other intercepting network device, goodput is the bandwidth used to transmit unique segments to said other intercepting network device, and a session is the duration of transmitting a TCP data segment from said source end-station to said destination end-station.

13. The intercepting network device of claim 12 in which said network performance measurement includes both throughput and goodput measurements, a variation in loss ratio is determined from said throughput and goodput measurements, and said encoding level is defined in response to said variation in loss ratio.

14. The intercepting network device of claim 13 in which said performance measurement includes a measurement of goodput, and said encoding level is reduced in response to a reduction in goodput.

15. A method of stimulating a TCP congestion window of a source end-station for transmitting TCP data segments in a data transmission network, said method comprising

generating a TCP ACK segment in response to the receipt of a TCP data segment at a destination end-station and transmitting said TCP ACK segment to said source end-station,

intercepting said TCP ACK segment and generating extra encoded TCP ACK segments corresponding to said intercepted TCP ACK segment, and

transmitting said extra TCP ACK segments to said source end-station to stimulate said TCP congestion window.

16. The method of claim 15 in which said generating of said extra encoded TCP ACK segments is continued for a predetermined time period.

17. The method of claim 15 which includes controlling the number of said extra TCP ACK segments based on the number of bytes contained in TCP ACK segments received after said intercepted TCP ACK segment.

18. The method of claim 15 which includes calculating a network performance measurement, and generating said extra TCP ACK segments when said network performance measurement exceeds a predetermined threshold.

19. The method of claim 18 in which said network performance measurement is at least one measurement selected from the group consisting of goodput, throughput, and loss measurements.

20. The method of claim 15 in which said extra TCP ACK segments are generated from an original TCP ACK segment based on a predetermined minimum segment size and maximum number of TCP ACK segments allowed, said method comprising

comparing the byte size of said original TCP ACK segment to the minimum segment byte size,

if said original TCP ACK segment size is greater than said minimum segment size, generating said extra TCP ACK segments, containing the same information as the original TCP ACK segment, the size of each of the said extra TCP ACK segments being at least the minimum segment size, and the number of total said extra TCP ACK segments generated not exceeding the maximum number of TCP ACK segments allowed, and

transmitting the generated extra TCP ACK segments to the source end-station.

21. A method for calculating round trip time RTT for a flow of TCP or non-TCP data segments from a first intercepting network device to a second intercepting network device, said method comprising

creating an RTT query segment and saving a local time stamp and a specific sequence identifier associated with said query segment,

transmitting said RTT query segment to said second intercepting network device,

receiving said RTT query segment at said second intercepting network device and transmitting said RTT query segment back to said first intercepting network device for computation of RTT,

retrieving said RTT query segment by detecting said sequence identifier, and

computing an RTT based on saved timestamp information and the current local time.

22. A data transmission network for calculating round trip time RTT for a flow of TCP or non-TCP data segments in an encoded channel from a source end-station to a destination end-station, said network comprising

first and second intercepting network devices establishing an encoded channel for a communication session between said end-stations,

one of said intercepting network devices being adapted to

create an RTT query segment and store a local time stamp and a specific sequence identifier associated with said RTT query segment,

transmit said RTT query segment to said second intercepting network device,

return said RTT query segment to first intercepting network device,

retrieve said RTT query segment by detecting said sequence identifier, and

compute an RTT based on stored timestamp information and the current local time.

23. An intercepting network device for use in calculating round trip time RTT for a flow of TCP or non-TCP data segments in an encoded channel from a source end-station to a destination end-station in a data transmission network, said intercepting network device comprising a processor programmed to

create an RTT query segment and save a local time stamp and a specific sequence identifier associated with said RTT query segment,

transmit said RTT query segment to said second intercepting network device,

return said RTT query segment to first intercepting network device,

retrieve said RTT query segment by detecting said sequence identifier, and

compute an RTT based on saved timestamp information and the current local time.

24. A method for calculating round trip time RTT for a flow of TCP data segments from a first intercepting network device to a second intercepting network device, said method comprising

intercepting a TCP data segment and storing a local time stamp and a unique TCP identifier associated with the intercepted TCP data segment,

detecting a TCP ACK segment having a unique TCP identifier corresponding to said stored TCP timestamp,

retrieving said stored local timestamp, and

computing an RTT based on said retrieved local timestamp and the current local time.

25. A data transmission network for calculating round trip time RTT for a flow of TCP data segments in an encoded TCP channel from a source end-station to a destination end-station, said network comprising

first and second intercepting network devices establishing an encoded TCP channel for a communication session between said end-stations,

one of said intercepting network devices being adapted to

intercept a TCP data segment and storing a local time stamp and a unique TCP identifier associated with the intercepted TCP data segment,

detect a TCP ACK segment having a TCP timestamp corresponding to said stored TCP timestamp,

retrieve said stored local timestamp, and

compute an RTT based on said retrieved local timestamp and the current local time.

26. An intercepting network device for use in calculating round trip time RTT for a flow of TCP data segments in an encoded TCP channel from a source end-station to a destination end-station in a data transmission network, said intercepting network device comprising a processor programmed to

intercept a TCP data segment and storing a local time stamp and a unique TCP identifier associated with the intercepted TCP data segment,

detect a TCP ACK segment having a TCP timestamp corresponding to said stored TCP timestamp,

retrieve said stored local timestamp, and

compute an RTT based on said retrieved local timestamp and the current local time.

27. A method of transmitting TCP data segments via an encoded TCP channel from a source end-station to a destination end-station, said method comprising

establishing an encoded TCP channel between first and second intercepting network devices for a TCP session between said end-stations,

intercepting a TCP data segment related to said TCP session at one of said intercepting network devices,

segmenting and encoding an intercepted TCP data segment at a selectable encoding level in one of said intercepting network devices and transmitting the encoded segments to the other of said intercepting network devices,

receiving a recommended encoding level for said TCP session,

setting the encoding level to the minimum level between said recommended encoding level and the highest encoding level that is not locked,

performing a first sequence of performance measurements for the TCP session between said end-stations,

setting the encoding level to said recommended encoding level,

performing a second sequence of performance measurements for said TCP session between said end-stations,

comparing said first and second sequences of performance measurements, and

if said second sequence of performance measurements is same or better than said first sequence of performance measurements, setting the encoding level to said recommended encoding level for a predetermined amount of time, or

if said second sequence of performance measurements is not better than said first sequence of performance measurements, locking said recommended encoding level.

28. The method of claim 27 in which said performance measurements include round trip time RTT and goodput.

29. The method of claim 27 which includes unlocking the lowest locked encoding level for said communication session after the expiration of a predetermined time period.

30. The method of claim 27 which include unlocking at least one of said locked encoding levels of a communication session when another communication session is completed.

31. A method of transmitting non-TCP data segments via an encoded channel from a source end-station to a destination end-station, said method comprising

establishing an encoded channel between first and second intercepting network devices for a non-TCP session between said end-stations,

intercepting a non-TCP data segment related to said non-TCP session at one of said intercepting network devices,

segmenting and encoding an intercepted non-TCP data segment at a selectable encoding level in one of said intercepting network devices and transmitting the encoded segment to the other of said intercepting network devices,

receiving a recommended encoding level for said non-TCP session,

setting the encoding level to the minimum level between said recommended encoding level and the highest encoding level that is not locked,

performing a first sequence of performance measurements for non-TCP session between said end-stations,

setting the encoding level to said recommended encoding level,

performing a second sequence of performance measurements for said non-TCP session between said end-stations,

comparing said first and second sequences of performance measurements, and

if said second sequence of performance measurements is same or better than said first sequence of performance measurements, setting the encoding level to said recommended encoding level for a predetermined amount of time, or

if said second sequence of performance measurements is not better than said first sequence of performance measurements, locking said recommended encoding level.

32. The method of claim 31 which include unlocking the lowest locked encoding level for said communication session after the expiration of a predetermined time period.

33. The method of claim and 31 which include unlocking at least one of said locked encoding levels of a communication session when another communication session is completed.

34. The method of claim 31 in which said performance measurements include round trip time RTT and loss.

1/36




FIG. 1

2/36




FIG. 2

3/36


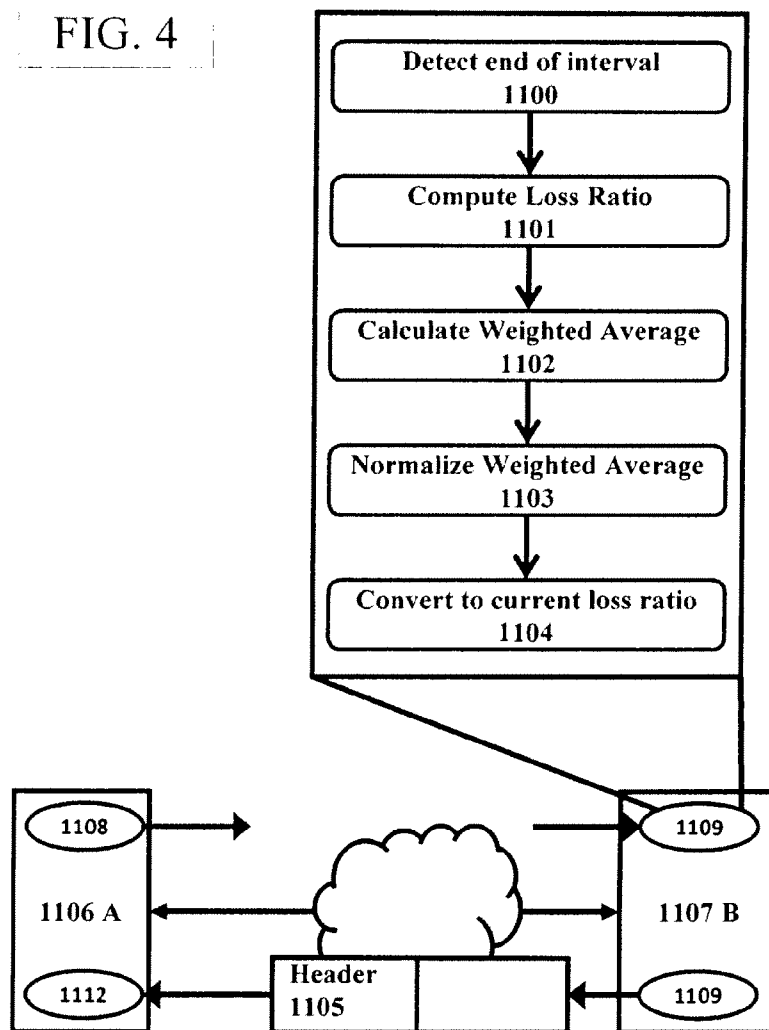



FIG. 3

4/36

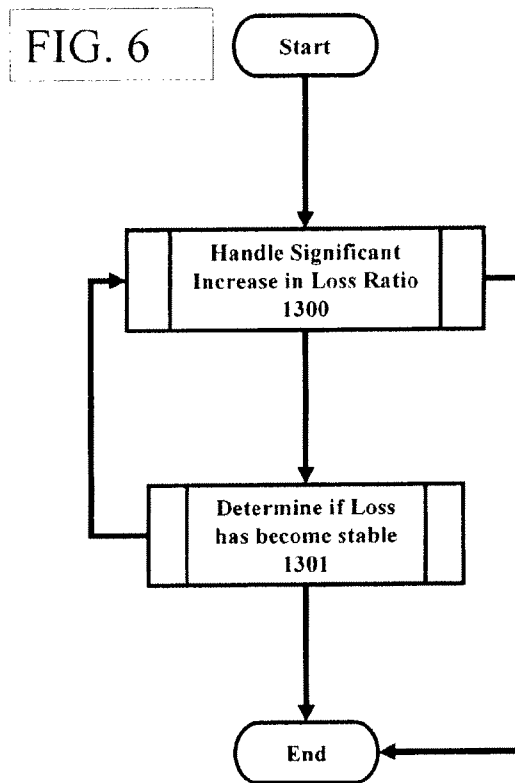
FIG. 4



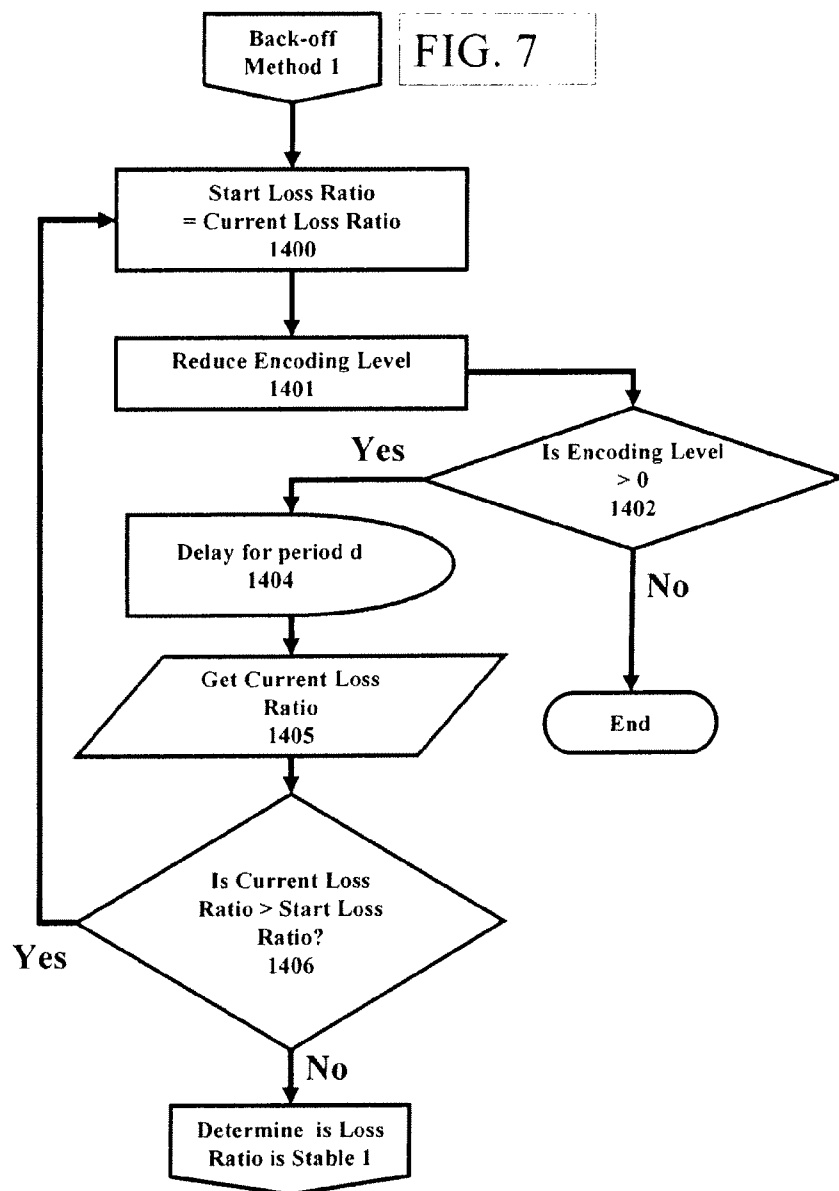
5/36



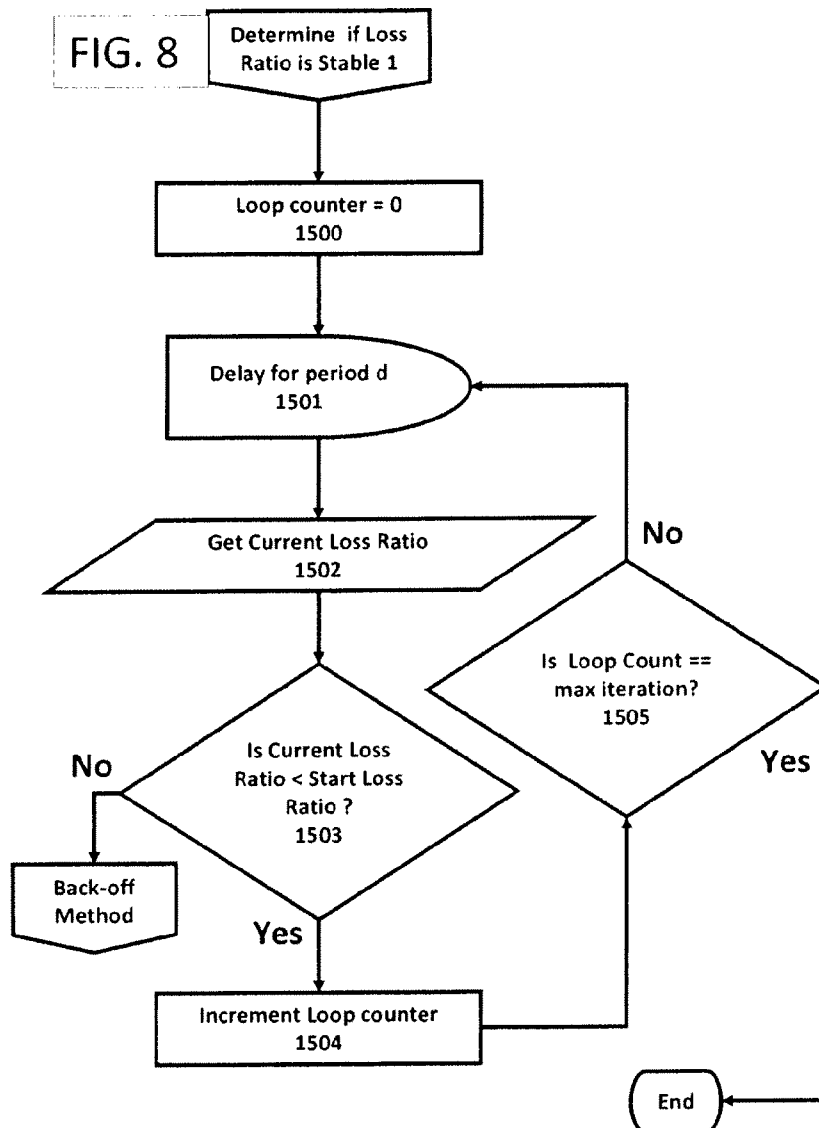
6/36



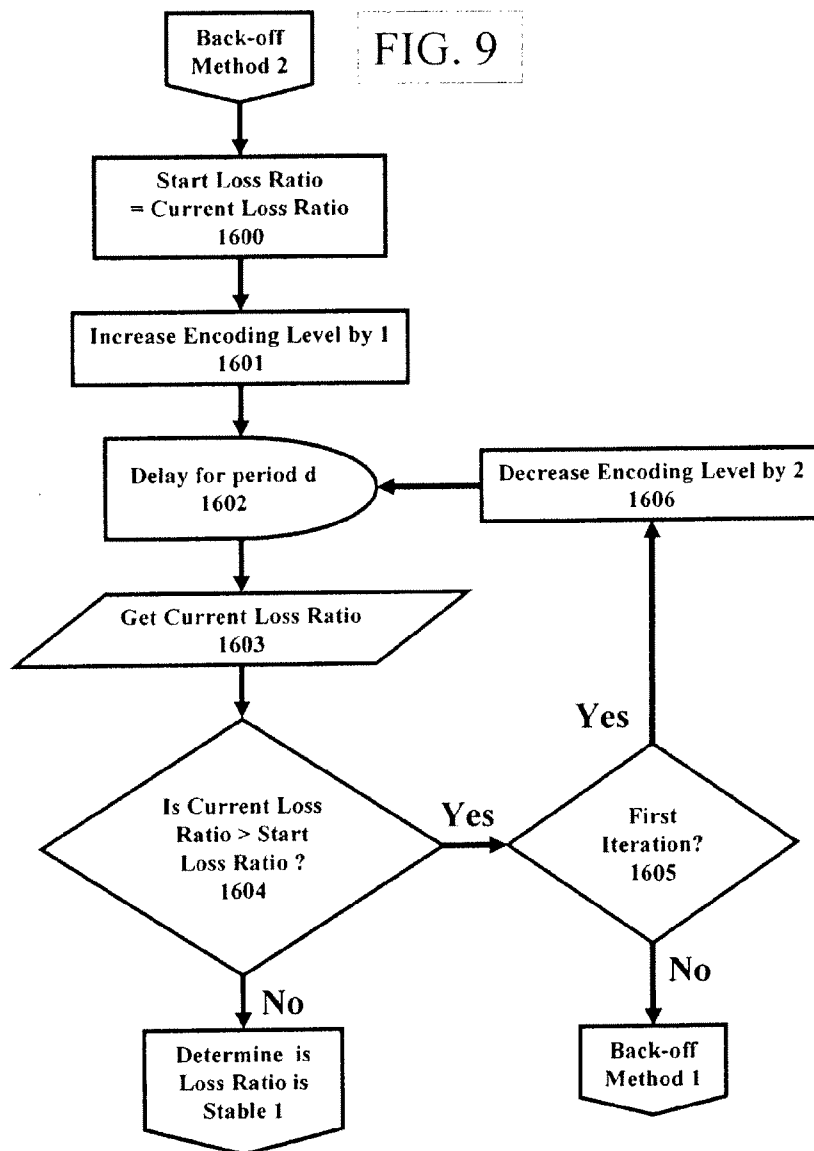
7/36




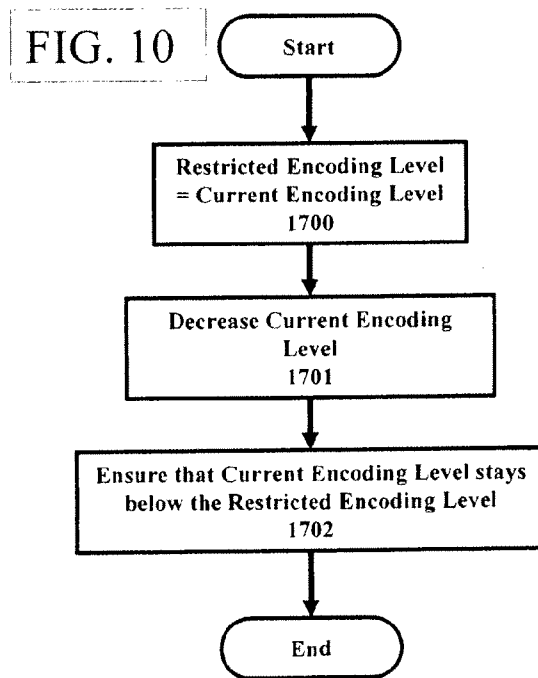
8/36




9/36



10/36

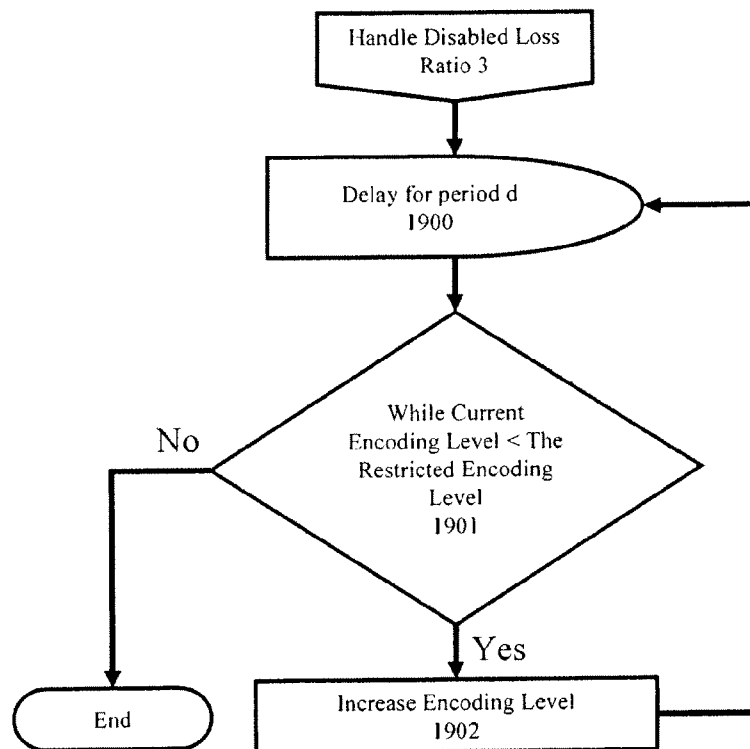


11/36



12/36

FIG. 13



13/36



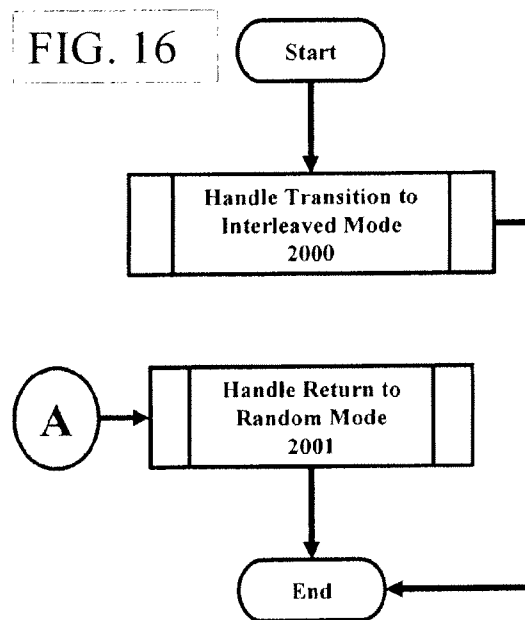


FIG. 14

14/36

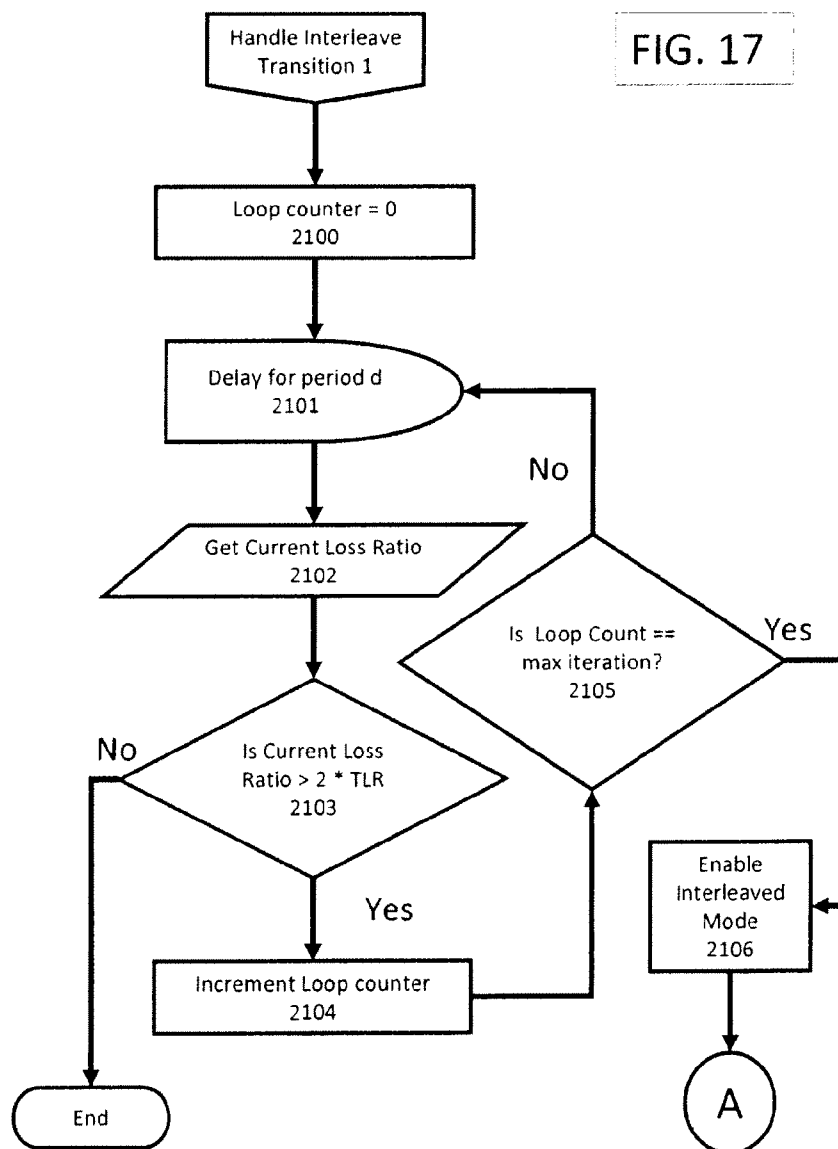


15/36




16/36

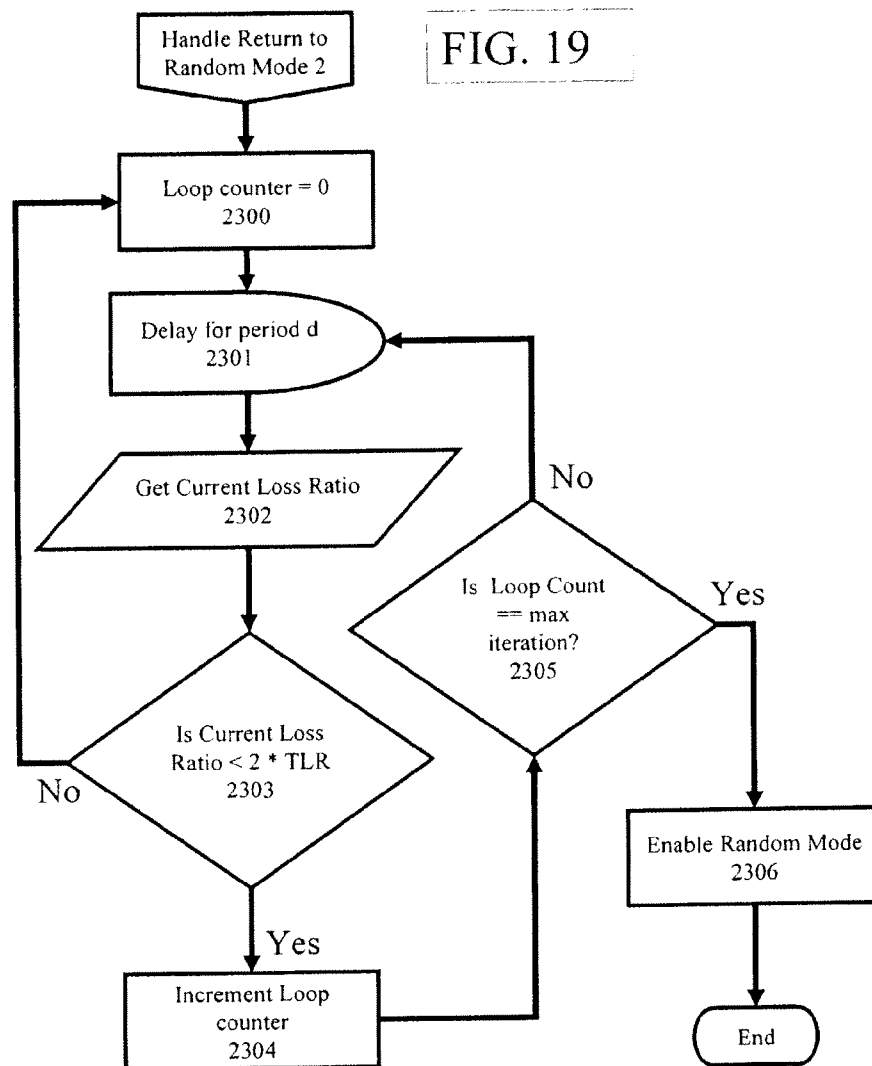
FIG. 17



17/36

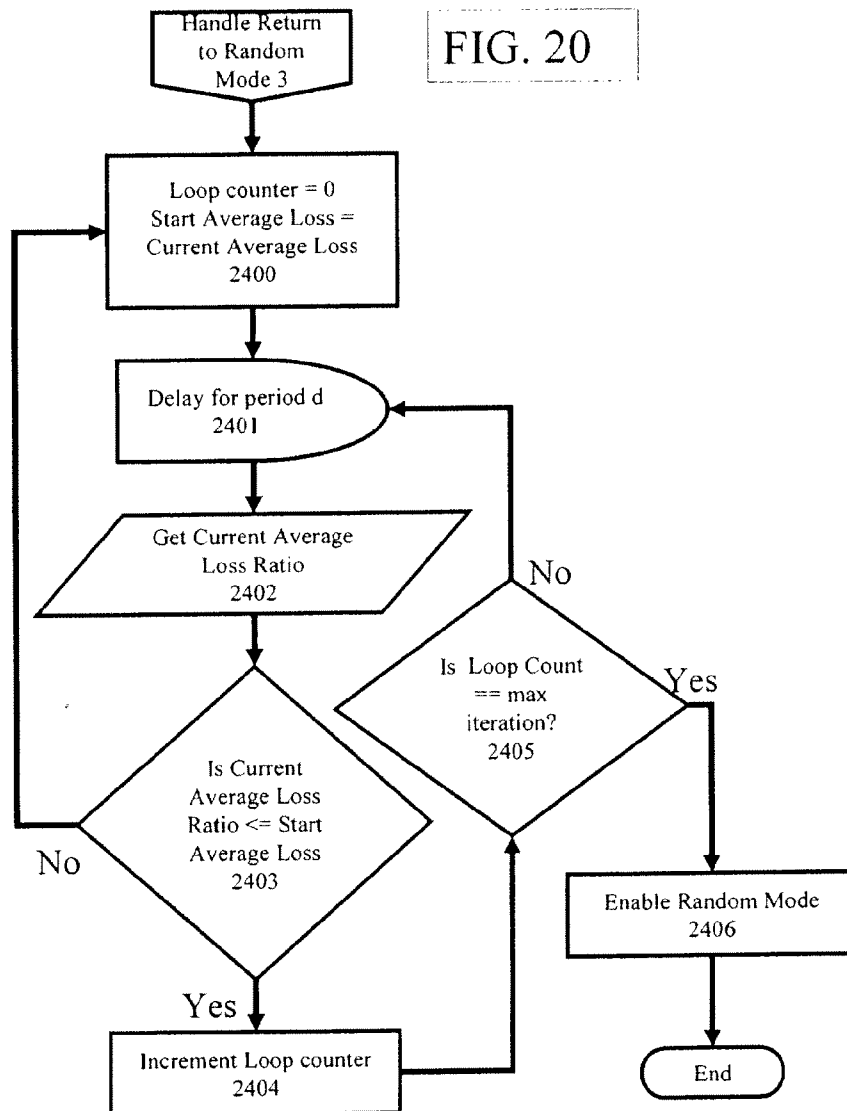


18/36

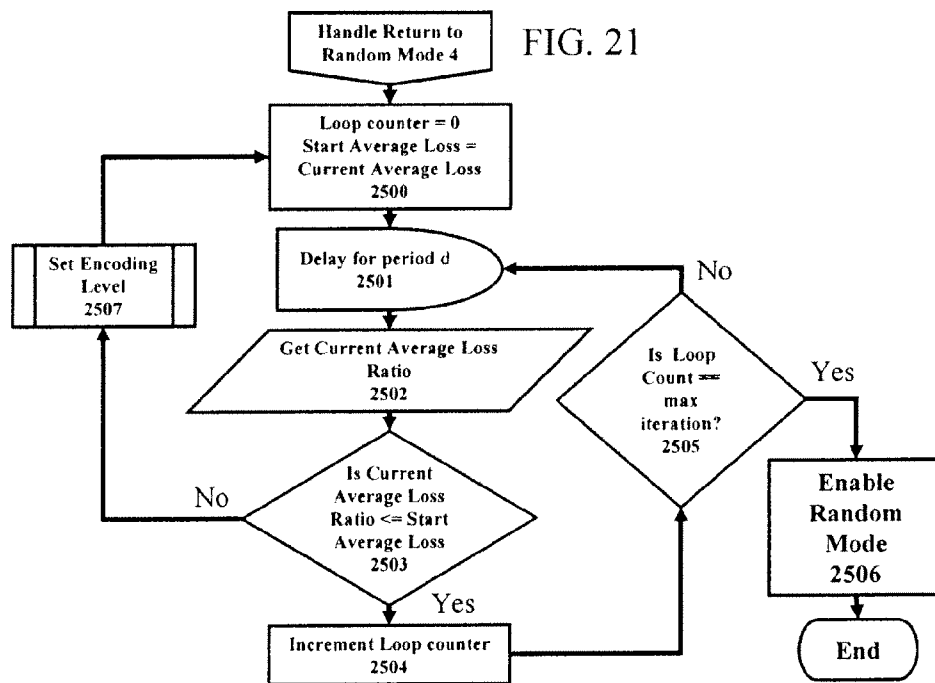


19/36

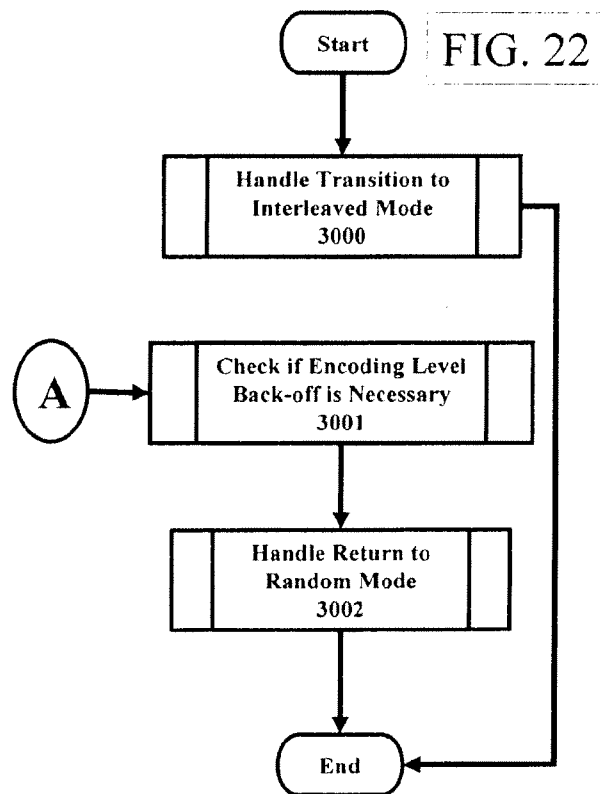
FIG. 20



20/36

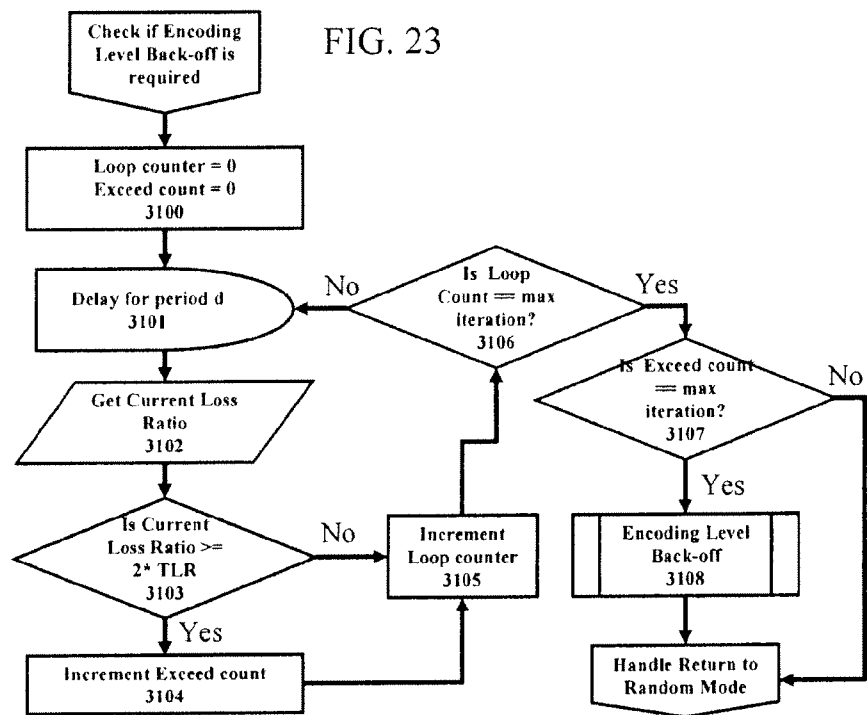


21/36



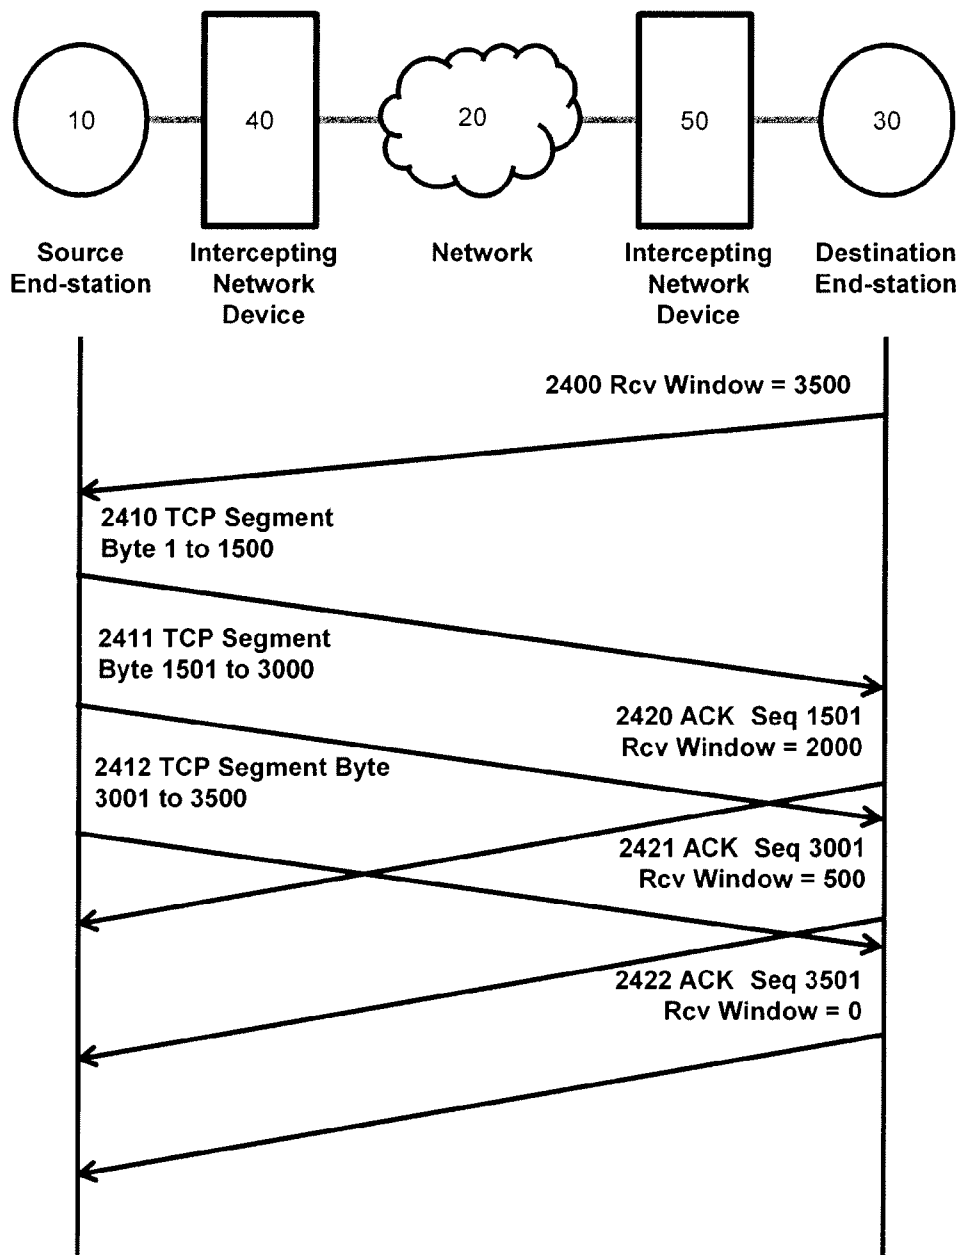
22/36

FIG. 23



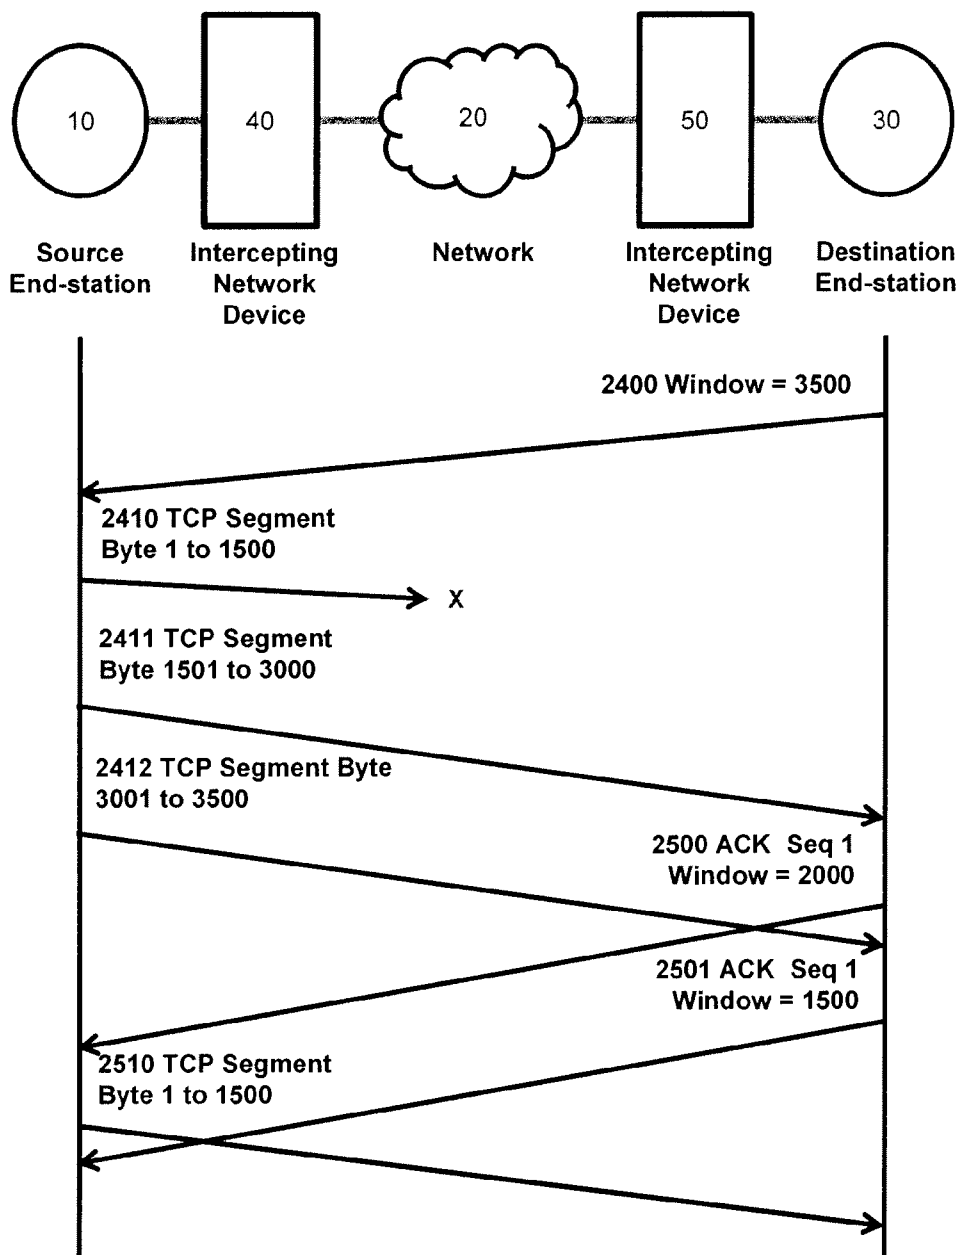
23/36

FIG. 24



24/36

FIG. 25



25/36

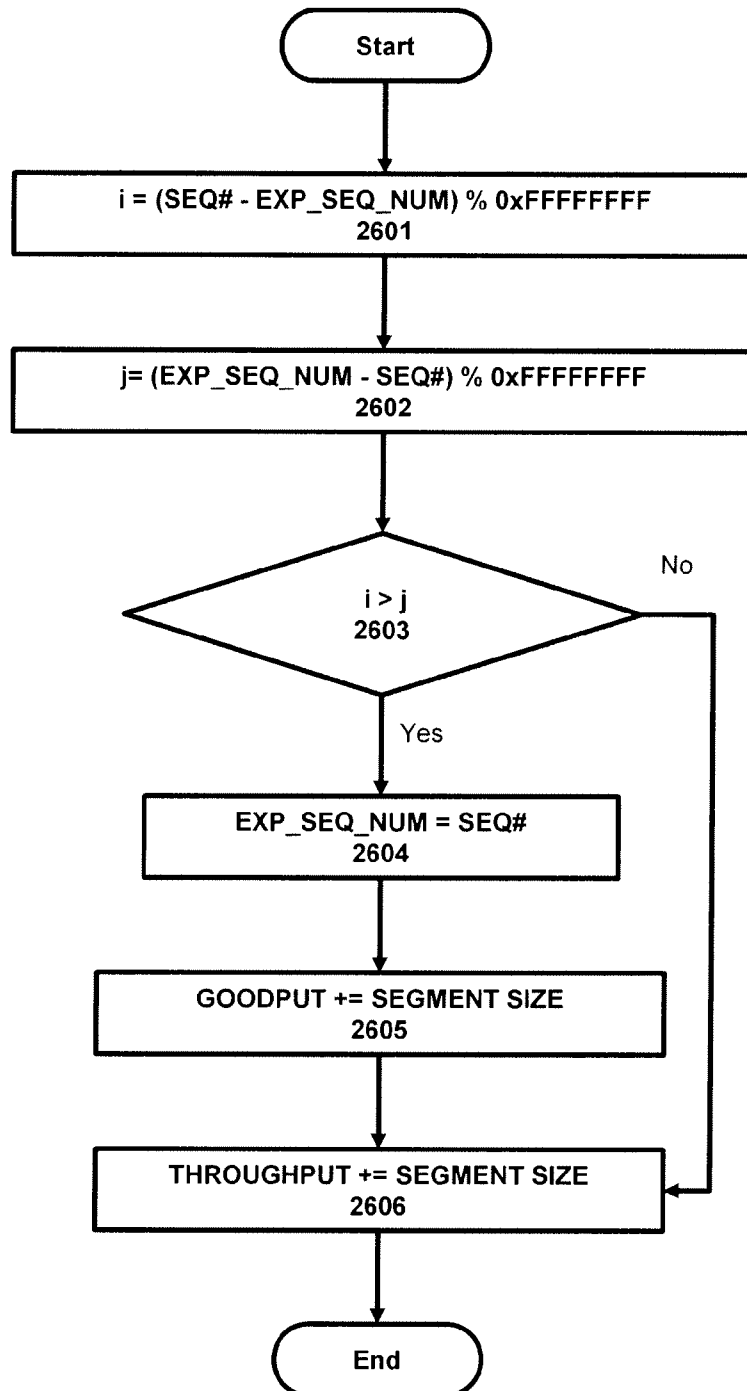


FIG. 26

26/36

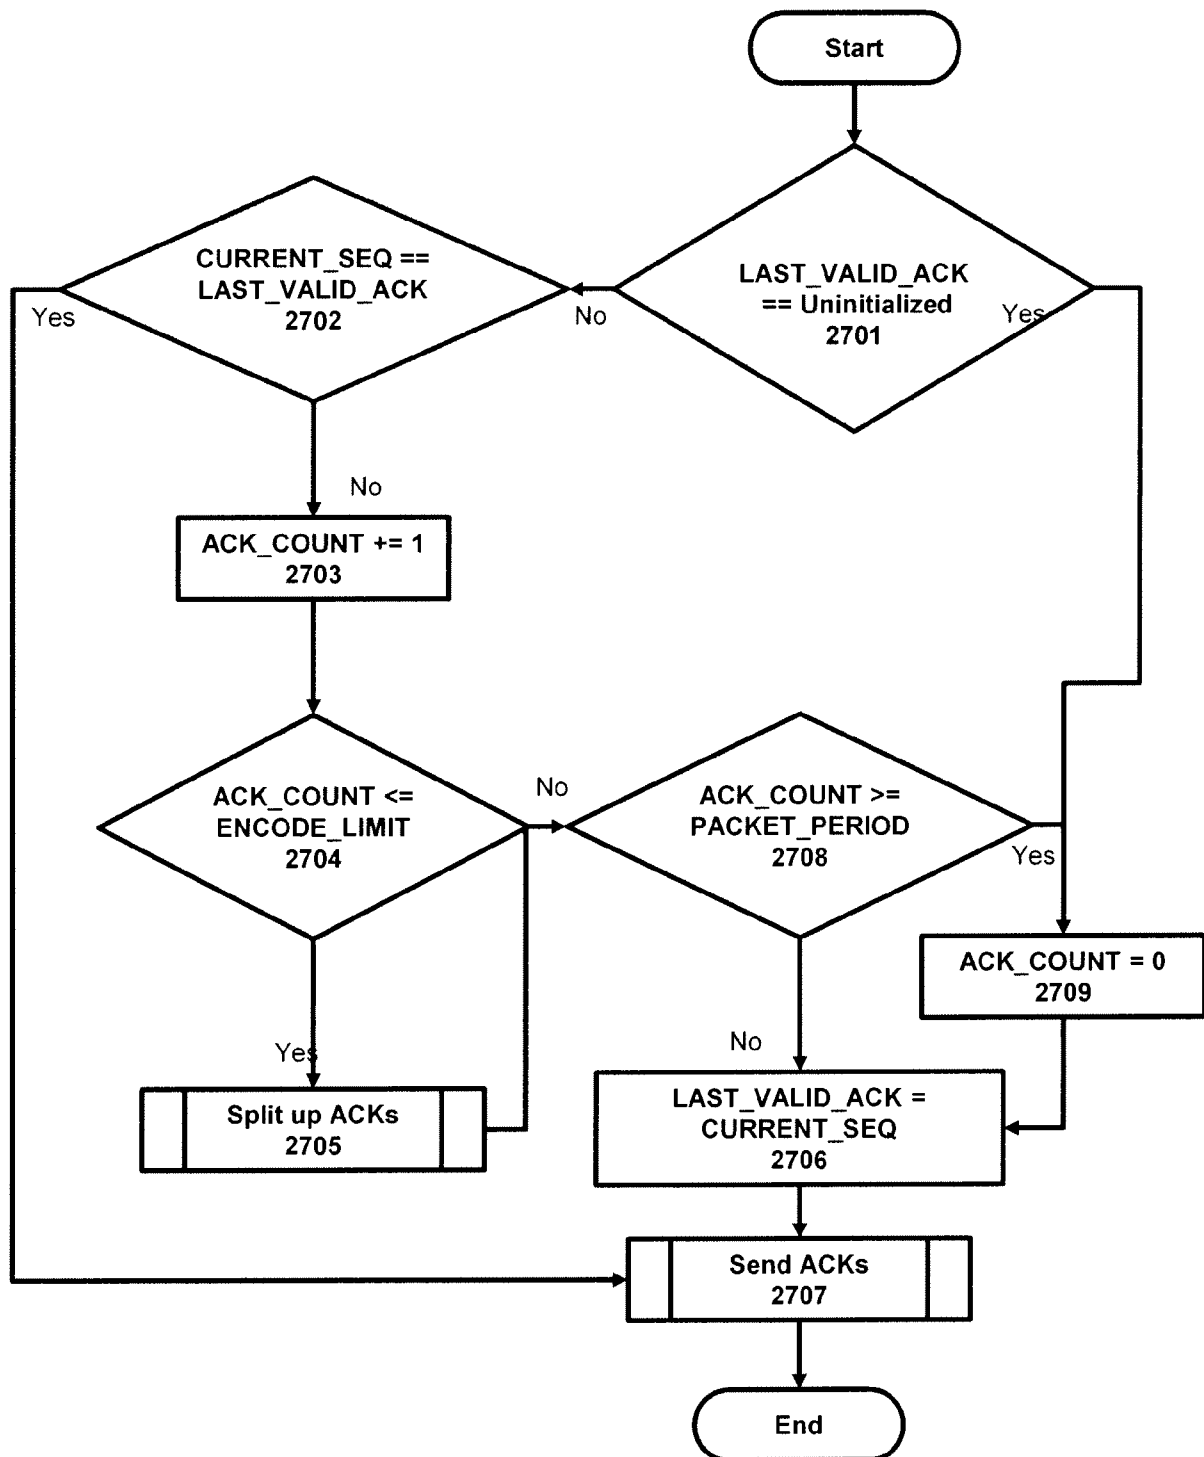
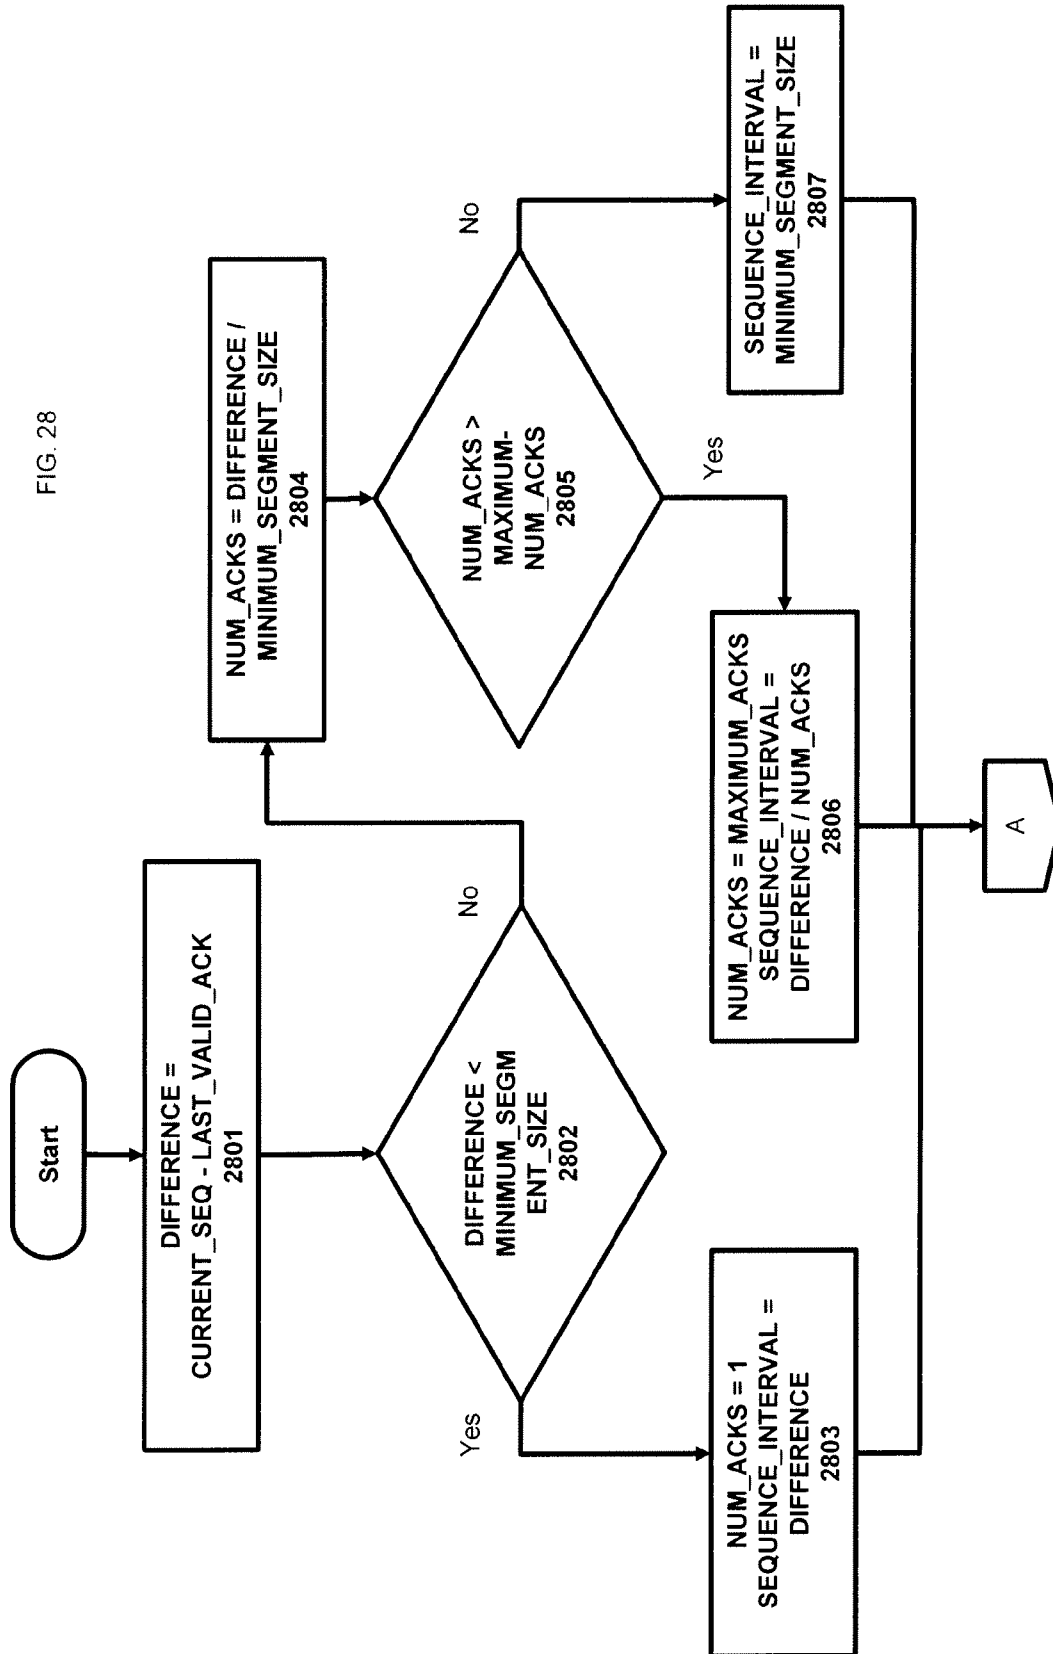


FIG. 27

27/36

FIG. 28



28/36

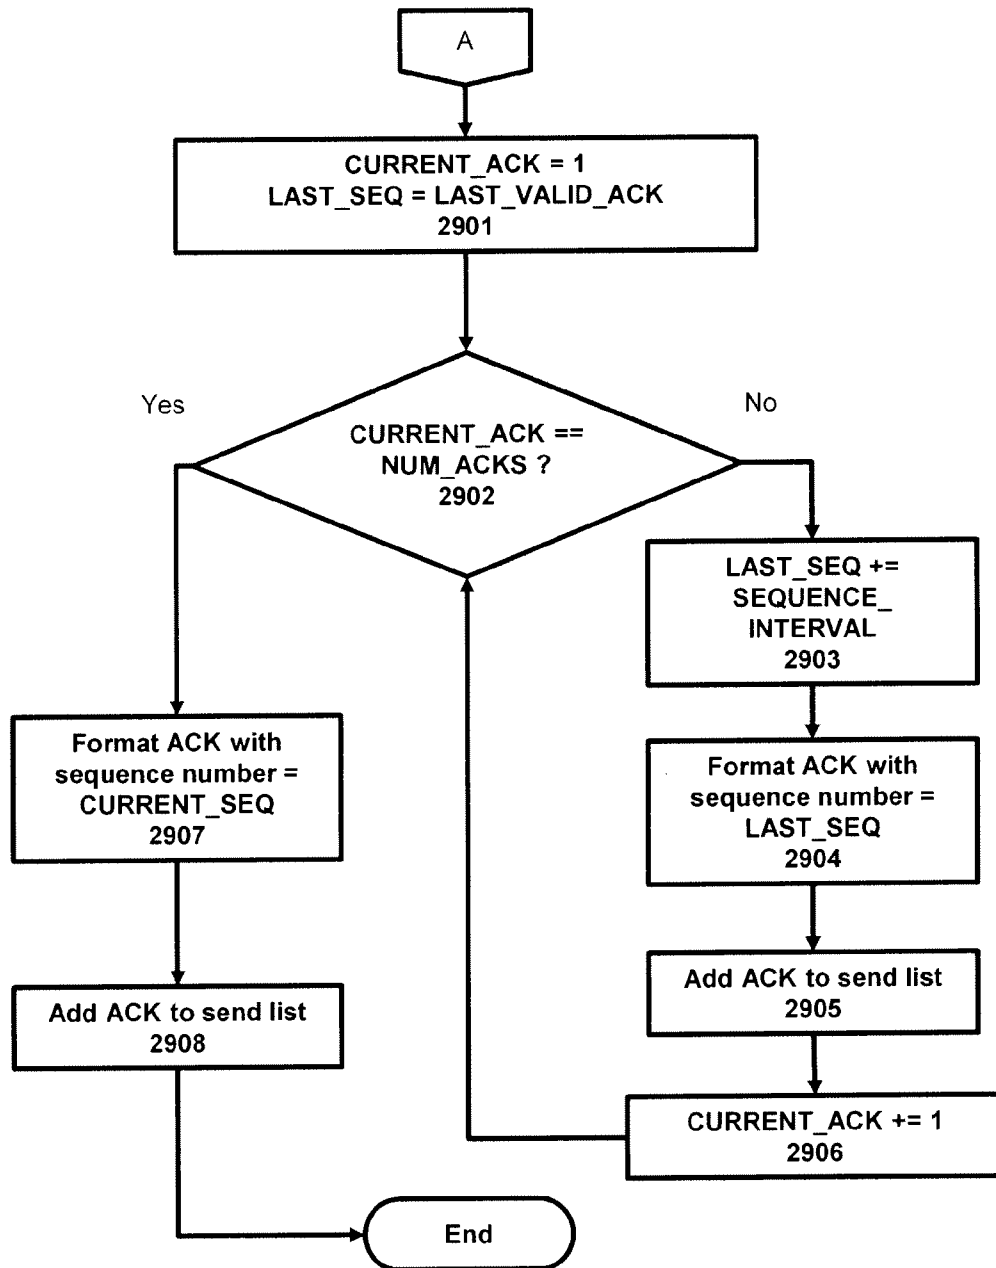
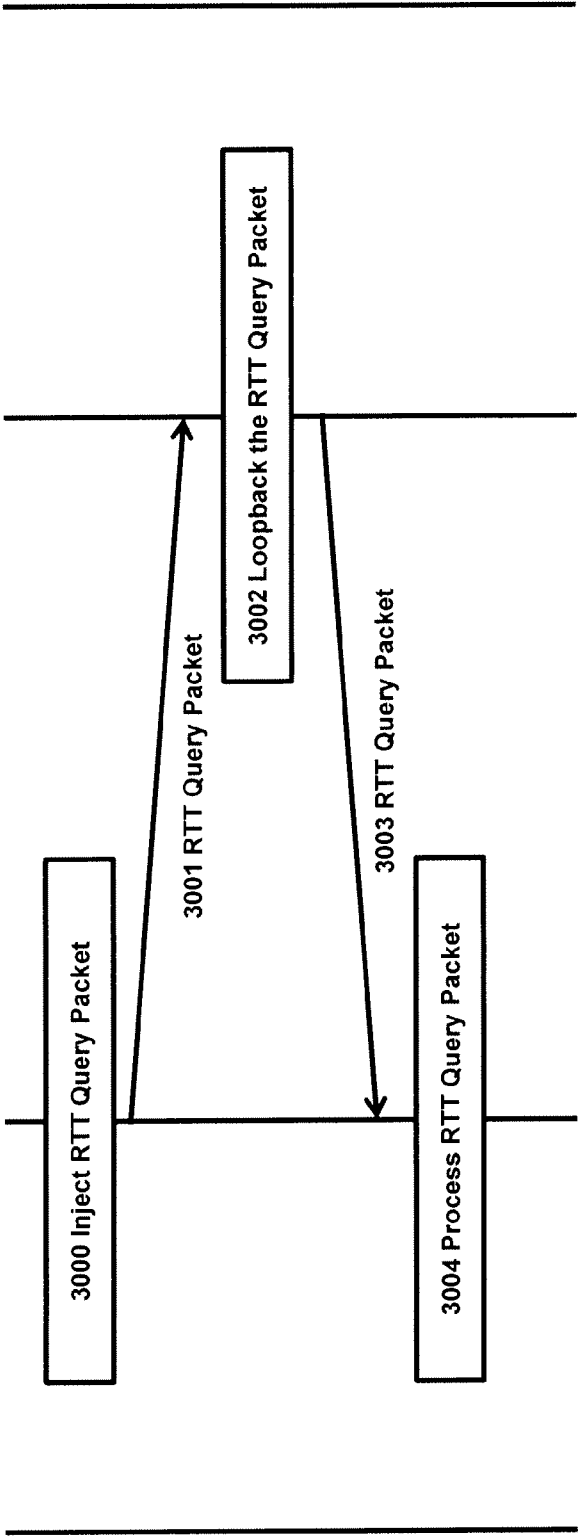



FIG. 29

FIG. 30



30/36

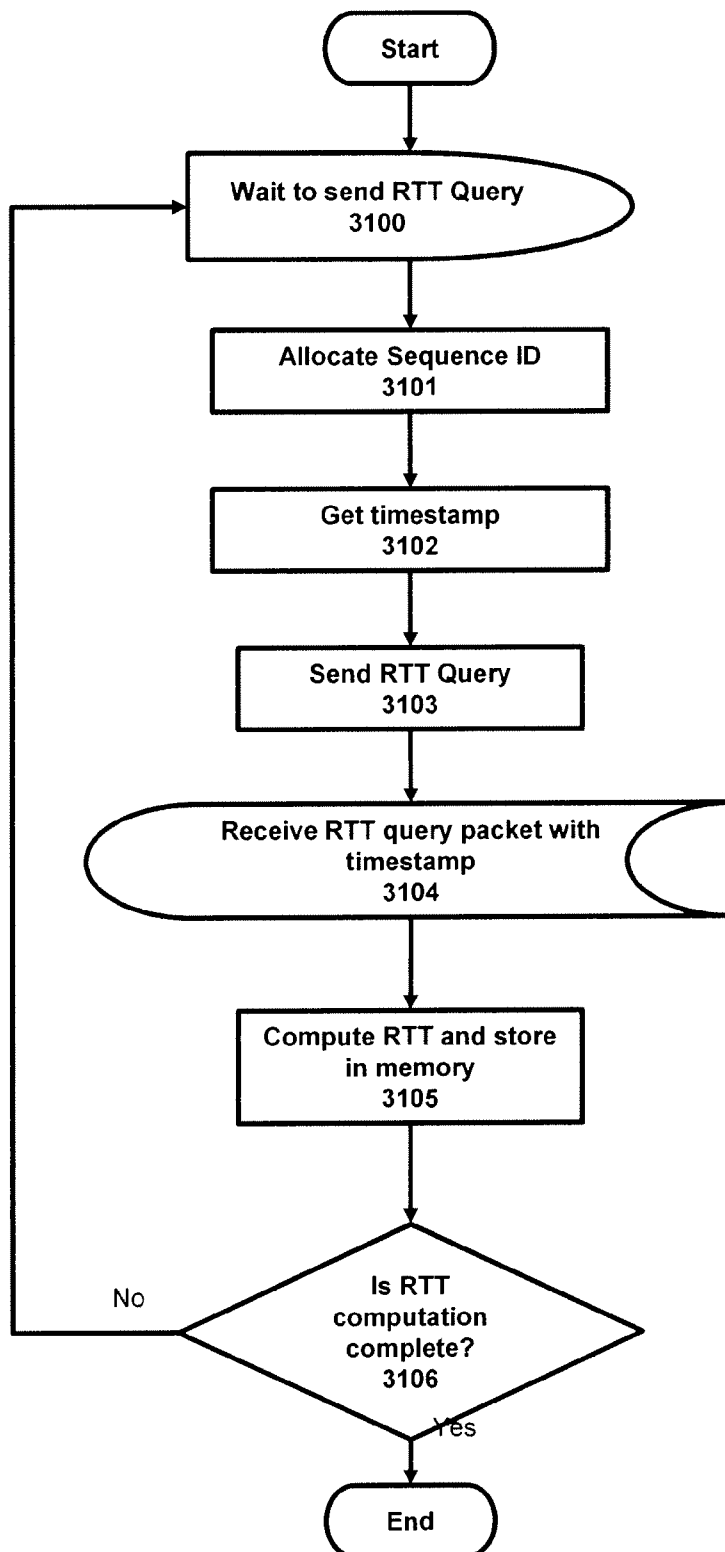


FIG. 31

31/36




FIG. 33

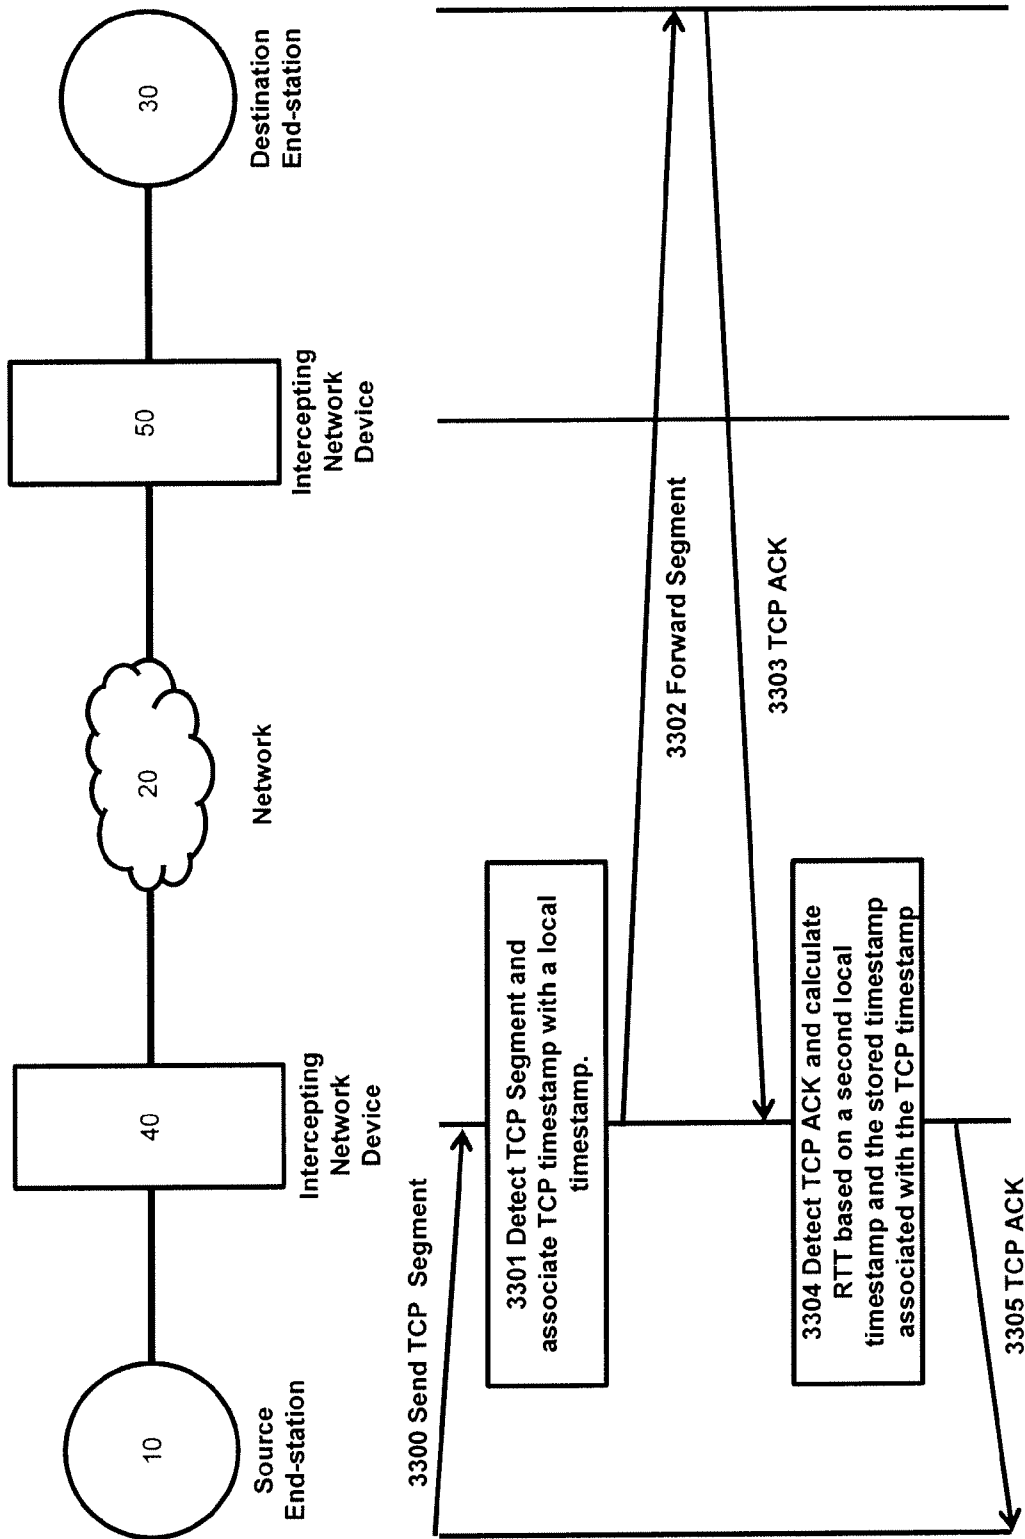


FIG. 34

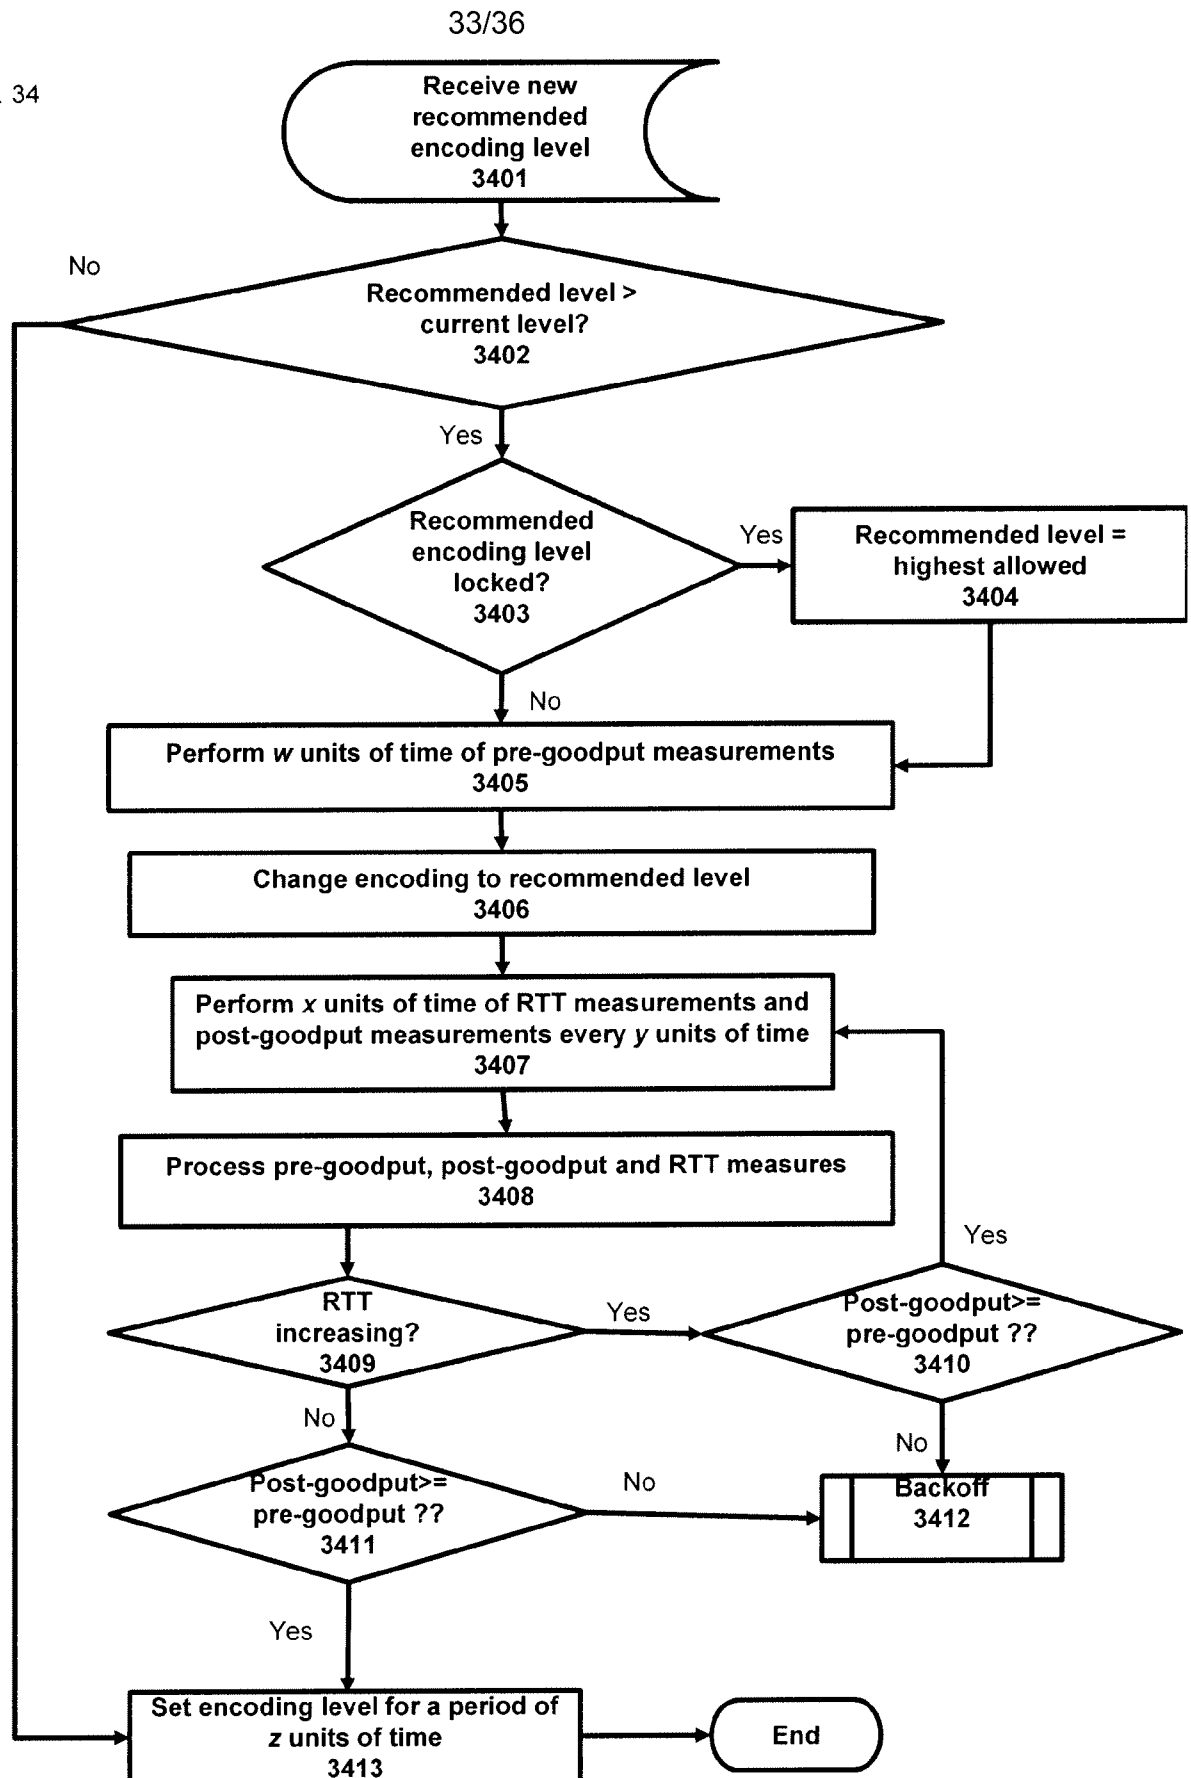
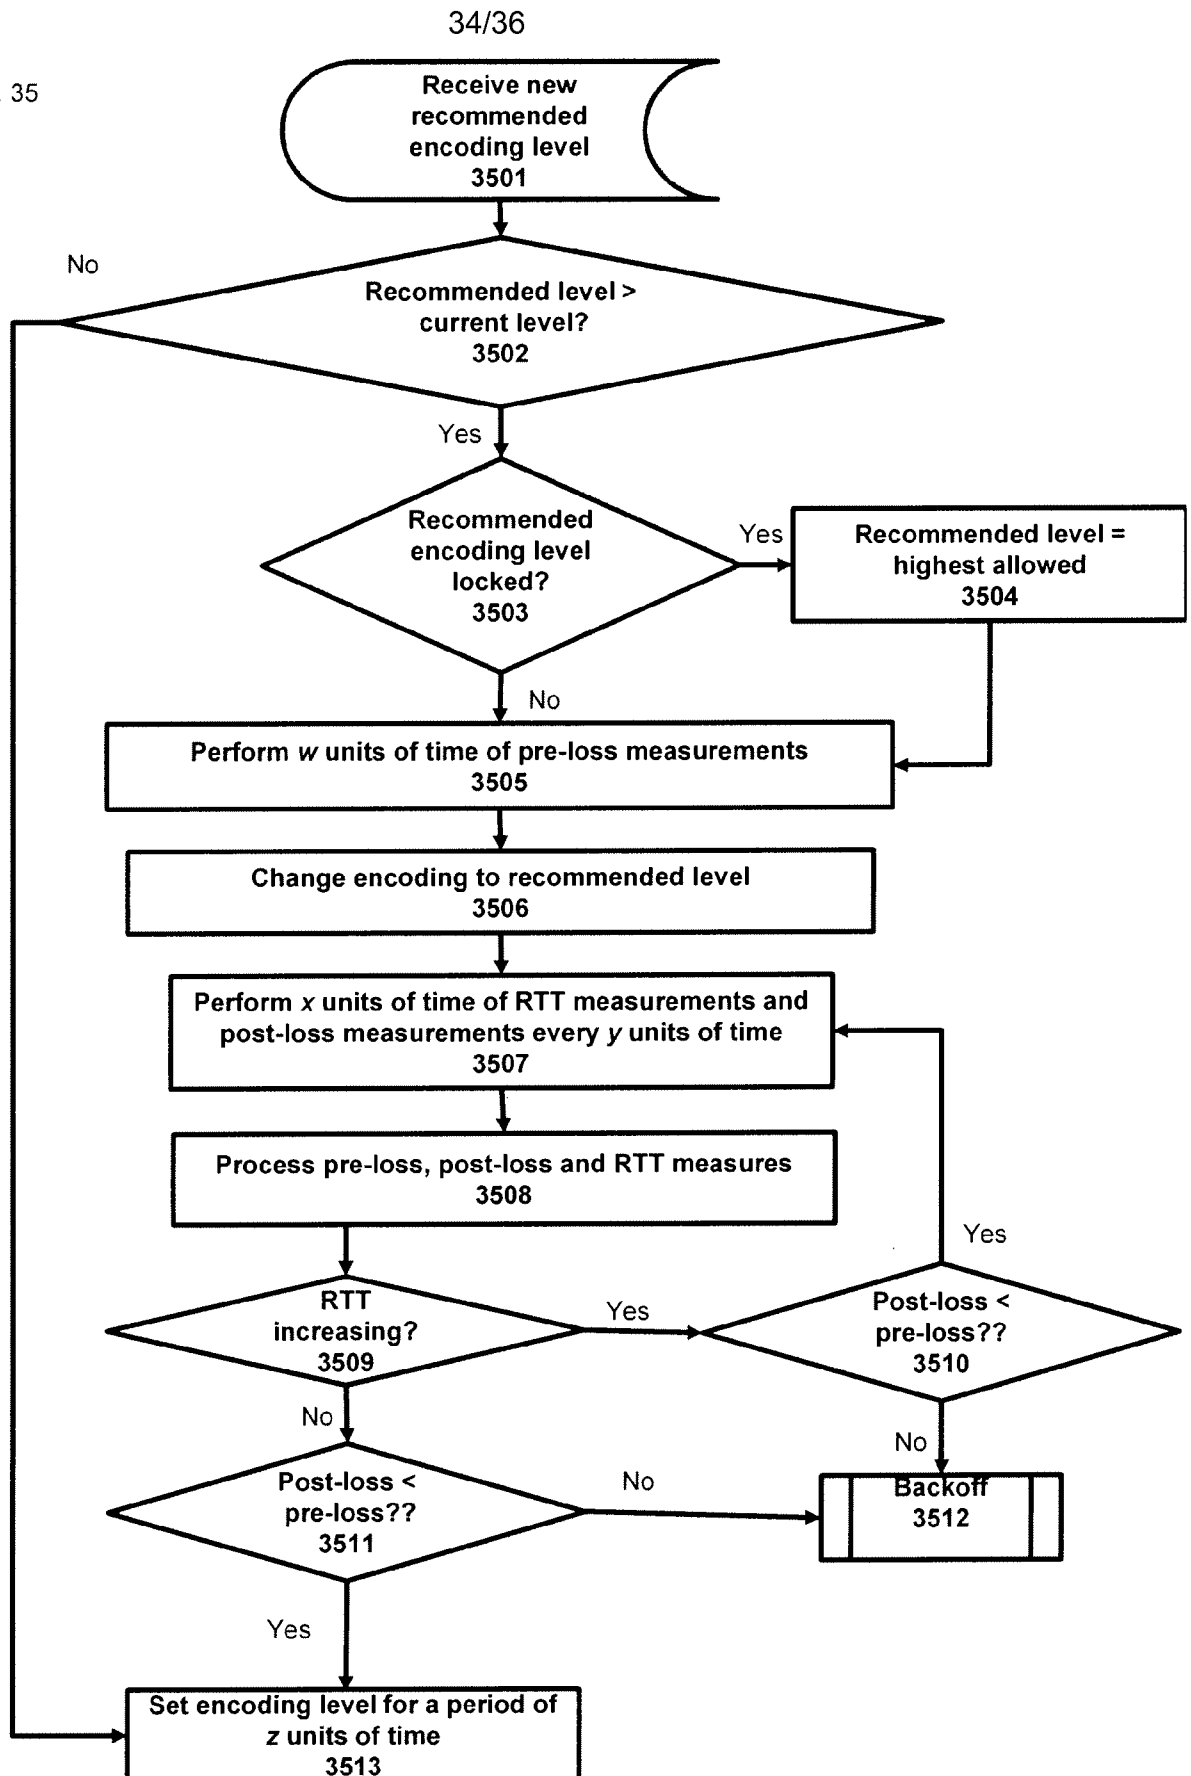
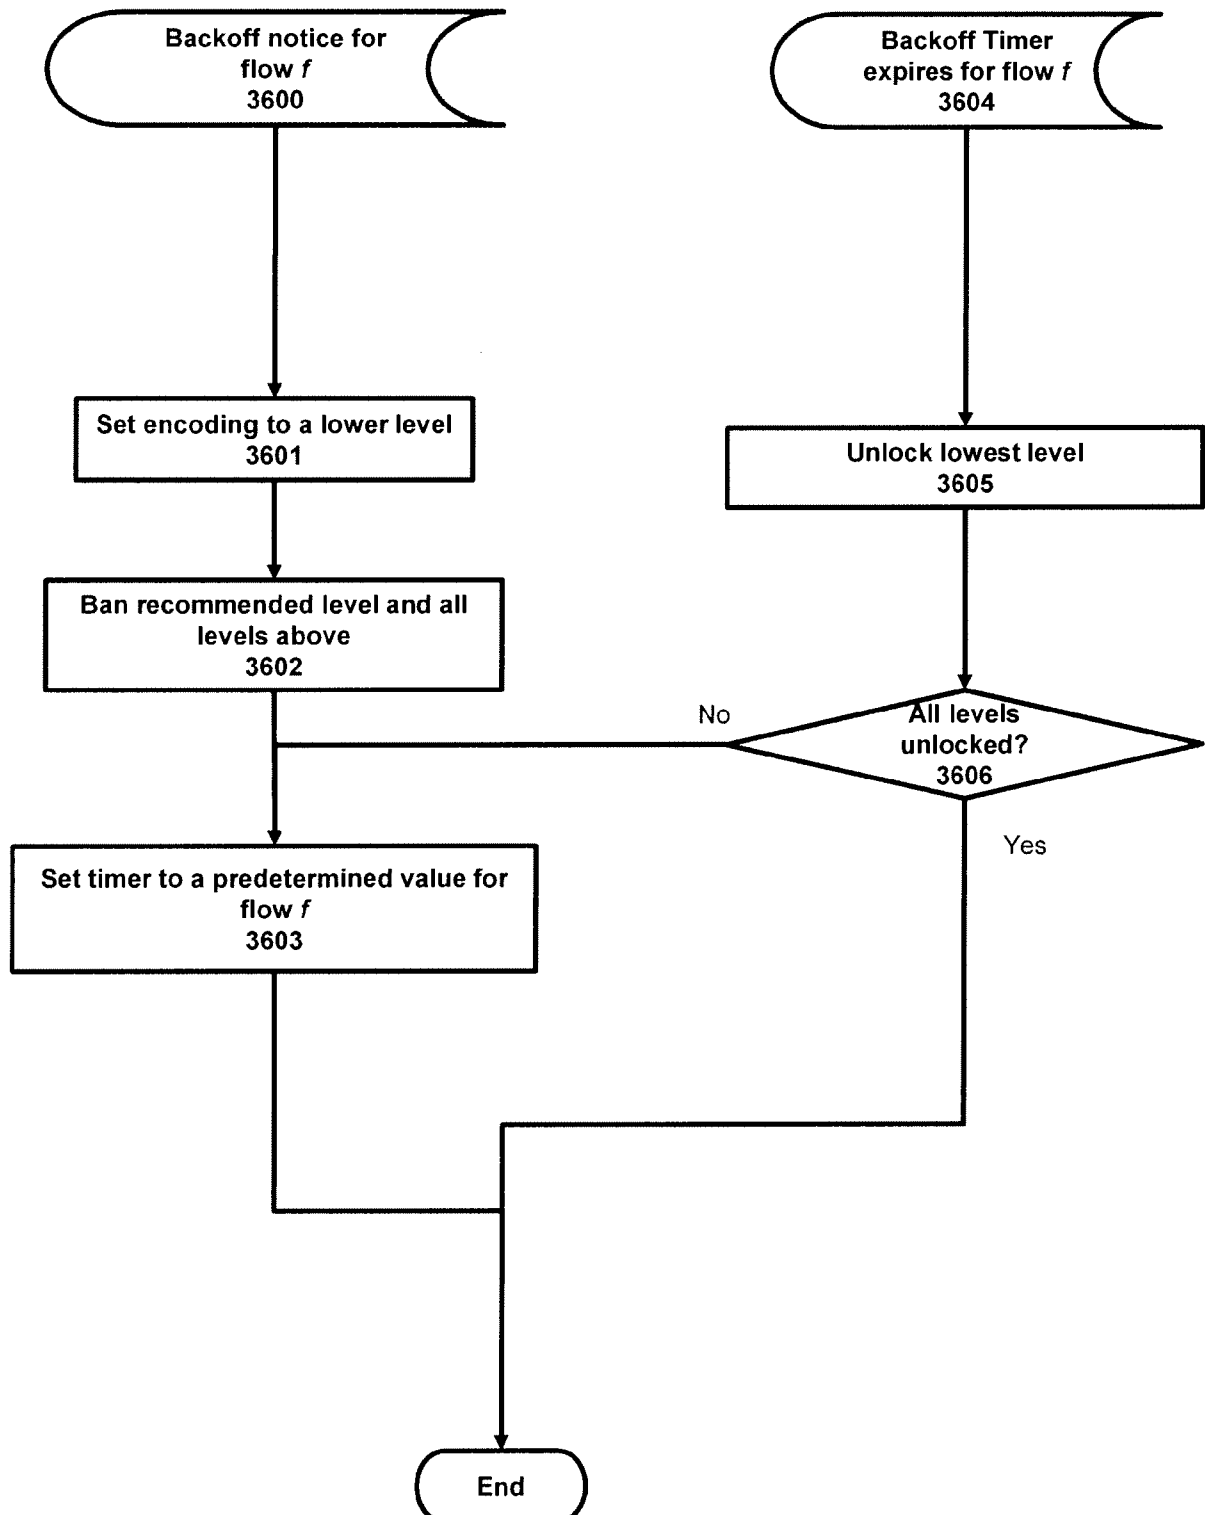


FIG. 35



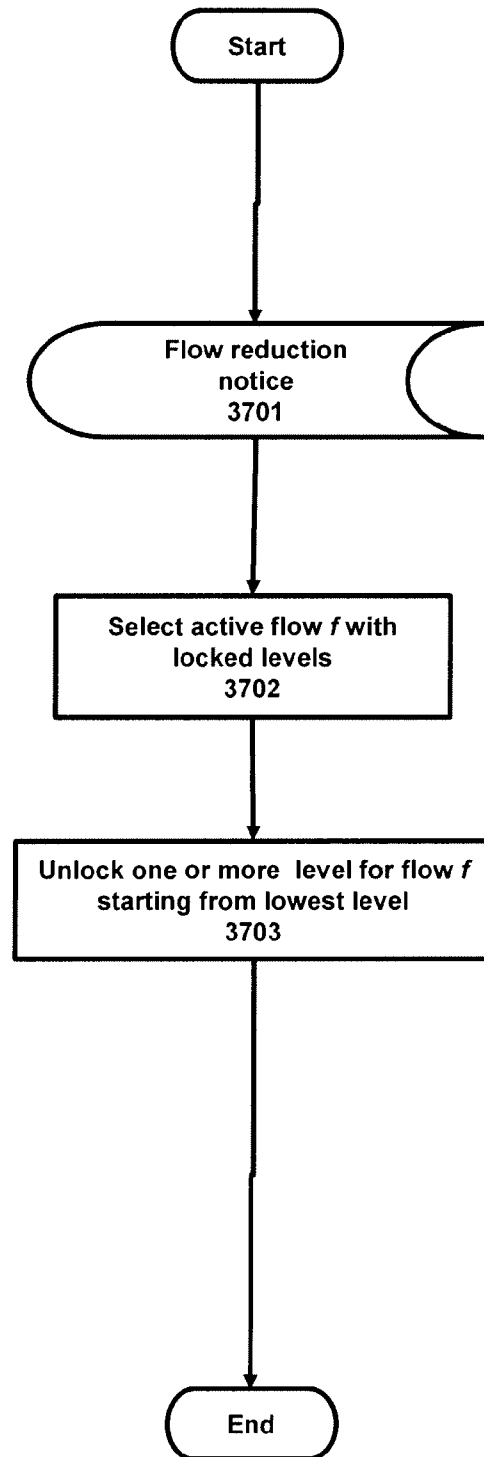
35/36

FIG. 36



36/36

FIG. 37



INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB20 12/056877

A. CLASSIFICATION OF SUBJECT MATTER

IPC: **H04L 29/08** (2006.01) , **H04L 12/70** (2013.01)

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)

IPC (2006.01) H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)

Canadian patent database, JEE Xplore, TotalPatent: TCP, data, segment, encoded, channel, source, destination, end-station, endpoint, intercepting, network, device, TCP session, segment, selectable, dynamic, encoding level, encoding rate, calculate, determine, network performance measurement, target, loss ratio, decode, throughput, goodput, bandwidth, variation, response, reduced and all such related

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 201 1/107963 A1 WILLIAMS et al. 09 September 201 1 (09-09-201 1) Entire document	1-5, 6-10, 1 1-14 and 27-34
X	US 2011/0206043 A1 WILLIAMS 25 August 201 1 (25-08-20 1 1) Entire document	1-5, 6-10, 1 1-14 and 27-34
X	US 201 1/0103388 A1 WILLIAMS et al. 05 May 201 1 (05-05-201 1) Entire document	1-5, 6-10, 1 1-14 and 27-34
X	US 2010/0220728 A1 WILLIAMS et al. 02 September 2010 (02-09-2010) Entire document	1-5, 6-10, 1 1-14 and 27-34
X	US 2008/0304483 A1 WILLIAMS et al. 11 December 2008 (11-12-2008) Entire document	1-5, 6-10, 1 1-14 and 27-34

[] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search

17 May 2013 (17-05-2013)

Date of mailing of the international search report

21 May 2013 (21-05-2013)

Name and mailing address of the ISA/CA

Canadian Intellectual Property Office
Place du Portage I, CI 14 - 1st Floor, Box PCT
50 Victoria Street
Gatineau, Quebec K1A 0C9
Facsimile No.: 001-819-953-2476

Authorized officer

Salvatore Ginese (819) 934-4888

INTERNATIONAL SEARCH REPORTInternational application No.
PCT/IB2012/056877**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of the first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons :

1. ☐ Claim Nos. :
because they relate to subject matter not required to be searched by this Authority, namely :

2. ☐ Claim Nos. :
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically :

3. ☐ Claim Nos. :
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows :

REFER TO PAGE 7

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claim Nos. :
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claim Nos. : 1-5, 6-10, 11-14 and 27-34

- Remark on Protest** ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

Box No. III**Observations where unity of invention is lacking**

This International Search Authority found multiple inventions in this international application, as follows:

GROUP A: Claims 1-5, 6-10, 11-14 and 27-34

A method of transmitting TCP data segments via an encoded TCP channel from a source end-station to a destination end-station (in addition to the associated data transmission network and intercepting network device), said method comprising:

establishing an encoded TCP channel between first and second intercepting network devices for a TCP session between said end-stations,
intercepting a TCP data segment related to said TCP session at one of said intercepting network devices,
segmenting and encoding said intercepted TCP data segment at a selectable encoding level in one of said intercepting network devices and transmitting the encoded segments to the other of said intercepting network devices,
calculating a network performance measurement for said TCP session between said end-stations, and
selecting said encoding level on the basis of said calculated network performance measurement for said TCP session.

GROUP B: Claims 15-20

A method of stimulating a TCP congestion window of a source end-station for transmitting TCP data segments in a data transmission network, said method comprising:

generating a TCP ACK segment in response to the receipt of a TCP data segment at a destination end-station and transmitting said TCP ACK segment to said source end-station,
intercepting said TCP ACK segment and generating extra encoded TCP ACK segments corresponding to said intercepted TCP ACK segment, and
transmitting said extra TCP ACK segments to said source end-station to stimulate said TCP congestion window.

GROUP C: Claims 21-26

A method for calculating round trip time RTT for a flow of TCP or non-TCP data segments from a first intercepting network device to a second intercepting network device (in addition to the associated data transmission network and intercepting network device), said method comprising:

creating an RTT query segment and saving a local time stamp and a specific sequence identifier associated with said query segment,
transmitting said RTT query-segment to said second intercepting network device,
receiving said RTT query segment at said second intercepting network device and transmitting said RTT query segment back to said first intercepting network device for computation of RTT,
retrieving said RTT query segment by detecting said sequence identifier, and
computing an RTT based on saved timestamp information and the current local time.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/IB20 12/056877

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
WO201 1107963A1	09 September 201 1 (09-09-201 1)	AT513392T CA2576038A1 CA2576038C CA2675S66A1 CN102859921A EP1779606A1 EP1779606A4 EP1779606B1 EP2157749A1 EP2157749B1 EP2528289A1 EP2543156A1 HK1 106639A1 JP2008509580A JP4712039B2 JP2010063094A JP5164123B2 KR20070065872A KR101 160056B1 KR20100021990A US2006029101A1 US7742501 B2 US2010220728A1 US79531 14B2 US2008304483A1 US8009696B2 US2010272122A1 US201 1103388A1 US201 1206043A1 US2012213232A1 US2012287806A1 WO2006012743A1	15 July 201 1 (15-07-201 1) 09 February 2006 (09-02-2006) 16 November 2010 (16-1 1-2010) 18 February 2010 (18-02-2010) 02 January 2013 (02-01-2013) 02 May 2007 (02-05-2007) 14 April 2010 (14-04-2010) 15 June 201 1 (15-06-201 1) 24 February 2010 (24-02-2010) 10 October 2012 (10-10-2012) 28 November 2012 (28-1 1-2012) 09 January 2013 (09-01-2013) 09 September 201 1 (09-09-201 1) 27 March 2008 (27-03-2008) 29 June 201 1 (29-06-201 1) 18 March 2010 (18-03-2010) 13 March 2013 (13-03-2013) 25 June 2007 (25-06-2007) 26 June 2012 (26-06-2012) 26 February 2010 (26-02-2010) 09 February 2006 (09-02-2006) 22 June 2010 (22-06-2010) 02 September 2010 (02-09-2010) 31 May 201 1 (31-05-201 1) 11 December 2008 (11-12-2008) 30 August 201 1 (30-08-201 1) 28 October 2010 (28-10-2010) 05 May 201 1 (05-05-201 1) 25 August 201 1 (25-08-201 1) 23 August 2012 (23-08-2012) 15 November 2012 (15-1 1-2012) 09 February 2006 (09-02-2006)
US201 1206043A1	25 August 201 1 (25-08-201 1)	AT513392T CA2576038A1 CA2576038C CA2675866A1 CN102859921A EP1779606A1 EP1779606A4 EP1779606B1 EP2157749A1 EP2157749B1 EP2528289A1 EP2543156A1 HK1 106639A1 JP2008509580A JP4712039B2 JP2010063094A JP5164123B2 KR20070065872A KR101 160056B1 KR20100021990A US2006029101A1 US7742501 B2 US2010220728A1 US79531 14B2 US2008304483A1 US8009696B2 US2010272122A1 US201 1103388A1 US2012213232A1 US2012287806A1 WO2006012743A1 WO201 1107963A1	15 July 201 1 (15-07-201 1) 09 February 2006 (09-02-2006) 16 November 2010 (16-1 1-2010) 18 February 2010 (18-02-2010) 02 January 2013 (02-01-2013) 02 May 2007 (02-05-2007) 14 April 2010 (14-04-2010) 15 June 201 1 (15-06-201 1) 24 February 2010 (24-02-2010) 10 October 2012 (10-10-2012) 28 November 2012 (28-1 1-2012) 09 January 2013 (09-01-2013) 09 September 201 1 (09-09-201 1) 27 March 2008 (27-03-2008) 29 June 201 1 (29-06-201 1) 18 March 2010 (18-03-2010) 13 March 2013 (13-03-2013) 25 June 2007 (25-06-2007) 26 June 2012 (26-06-2012) 26 February 2010 (26-02-2010) 09 February 2006 (09-02-2006) 22 June 2010 (22-06-2010) 02 September 2010 (02-09-2010) 31 May 201 1 (31-05-201 1) 11 December 2008 (11-12-2008) 30 August 201 1 (30-08-201 1) 28 October 2010 (28-10-2010) 05 May 201 1 (05-05-201 1) 23 August 2012 (23-08-2012) 15 November 2012 (15-1 1-2012) 09 February 2006 (09-02-2006) 09 September 201 1 (09-09-201 1)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB20 12/056877

US201 1103388A1	05 May 201 1 (05-05-201 1)	AT513392T CA2576038A1 CA2576038C CA2675866A1 CN102859921A EP1779606A1 EP1779606A4 EP1779606B1 EP2157749A1 EP2157749B1 EP2528289A1 EP2543156A1 HK1 106639A1 JP2008509580A JP4712039B2 JP2010063094A JP5164123B2 KR20070065872A KR101 160056B1 KR20100021990A US2006029101A1 US7742501 B2 US2010220728A1 US79531 14B2 US2008304483A1 US8009696B2 US2010272122A1 US201 1206043A1 US2012213232A1 US2012287806A1 WO2006012743A1 WO201 1107963A1	15 July 201 1 (15-07-201 1) 09 February 2006 (09-02-2006) 16 November 2010 (16-1 1-2010) 18 February 2010 (18-02-2010) 02 January 2013 (02-01-2013) 02 May 2007 (02-05-2007) 14 April 2010 (14-04-2010) 15 June 201 1 (15-06-201 1) 24 February 2010 (24-02-2010) 10 October 2012 (10-10-2012) 28 November 2012 (28-1 1-2012) 09 January 2013 (09-01-2013) 09 September 201 1 (09-09-201 1) 27 March 2008 (27-03-2008) 29 June 201 1 (29-06-201 1) 18 March 2010 (18-03-2010) 13 March 2013 (13-03-2013) 25 June 2007 (25-06-2007) 26 June 2012 (26-06-2012) 26 February 2010 (26-02-2010) 09 February 2006 (09-02-2006) 22 June 2010 (22-06-2010) 02 September 2010 (02-09-2010) 3 1 May 201 1 (31-05-201 1) 11 December 2008 (1 1-12-2008) 30 August 201 1 (30-08-201 1) 28 October 2010 (28-10-2010) 25 August 201 1 (25-08-201 1) 23 August 2012 (23-08-2012) 15 November 2012 (15-1 1-2012) 09 February 2006 (09-02-2006) 09 September 201 1 (09-09-201 1)
US2010220728A1	02 September 2010 (02-09-20 10)	AT513392T CA2576038A1 CA2576038C CA2675866A1 CN102859921A EP1779606A1 EP1779606A4 EP1779606B1 EP2157749A1 EP2157749B1 EP2528289A1 EP2543156A1 HK1 106639A1 JP2008509580A JP4712039B2 JP2010063094A JP5164123B2 KR20070065872A KR101 160056B1 KR20100021990A US2006029101A1 US7742501 B2 US79531 14B2 US2008304483A1 US8009696B2 US2010272122A1 US201 1103388A1 US201 1206043A1 US2012213232A1 US2012287806A1 WO2006012743A1 WO201 1107963A1	15 July 201 1 (15-07-201 1) 09 February 2006 (09-02-2006) 16 November 2010 (16-1 1-2010) 18 February 2010 (18-02-2010) 02 January 2013 (02-01-2013) 02 May 2007 (02-05-2007) 14 April 2010 (14-04-2010) 15 June 201 1 (15-06-201 1) 24 February 2010 (24-02-2010) 10 October 2012 (10-10-2012) 28 November 2012 (28-1 1-2012) 09 January 2013 (09-01-2013) 09 September 201 1 (09-09-201 1) 27 March 2008 (27-03-2008) 29 June 201 1 (29-06-201 1) 18 March 2010 (18-03-2010) 13 March 2013 (13-03-2013) 25 June 2007 (25-06-2007) 26 June 2012 (26-06-2012) 26 February 2010 (26-02-2010) 09 February 2006 (09-02-2006) 22 June 2010 (22-06-2010) 3 1 May 201 1 (31-05-201 1) 11 December 2008 (1 1-12-2008) 30 August 201 1 (30-08-201 1) 28 October 2010 (28-10-2010) 05 May 201 1 (05-05-201 1) 25 August 201 1 (25-08-201 1) 23 August 2012 (23-08-2012) 15 November 2012 (15-1 1-2012) 09 February 2006 (09-02-2006) 09 September 201 1 (09-09-201 1)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB20 12/056877

US2008304483A1

11 December 2008 (11-12-2008) AT513392T

CA2576038A1
CA2576038C
CA2675866A1
CN102859921A
EP1779606A1
EP1779606A4
EP1779606B1
EP2157749A1
EP2157749B1
EP2528289A1
EP2543156A1
HK1 106639A1
JP2008509580A
JP4712039B2
JP2010063094A
JP5164123B2
KR20070065872A
KR101 160056B1
KR20100021990A
US2006029101A1
US7742501 B2
US2010220728A1
US79531 14B2
US8009696B2
US2010272122A1
US201 1103388A1
US201 1206043A1
US2012213232A1
US2012287806A1
WO2006012743A1
WO201 1107963A1

15 July 201 1 (15-07-201 1)
09 February 2006 (09-02-2006)
16 November 2010 (16-1 1-2010)
18 February 2010 (18-02-2010)
02 January 2013 (02-01-2013)
02 May 2007 (02-05-2007)
14 April 2010 (14-04-2010)
15 June 201 1 (15-06-201 1)
24 February 2010 (24-02-2010)
10 October 2012 (10-10-2012)
28 November 2012 (28-1 1-2012)
09 January 2013 (09-01-2013)
09 September 201 1 (09-09-201 1)
27 March 2008 (27-03-2008)
29 June 201 1 (29-06-201 1)
18 March 2010 (18-03-2010)
13 March 2013 (13-03-2013)
25 June 2007 (25-06-2007)
26 June 2012 (26-06-2012)
26 February 2010 (26-02-2010)
09 February 2006 (09-02-2006)
22 June 2010 (22-06-2010)
02 September 2010 (02-09-2010)
3 1 May 201 1 (31-05-201 1)
30 August 201 1 (30-08-201 1)
28 October 2010 (28-10-2010)
05 May 201 1 (05-05-201 1)
25 August 201 1 (25-08-201 1)
23 August 2012 (23-08-2012)
15 November 2012 (15-1 1-2012)
09 February 2006 (09-02-2006)
09 September 201 1 (09-09-201 1)