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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 August 2009 (13.08.2009)

PCT

(10) International Publication Number
WO 2009/099847 A2

(51) International Patent Classification:
H04L 1/00 (2006.01) *H04L 1/06* (2006.01)

(21) International Application Number: PCT/US2009/032305

(22) International Filing Date: 28 January 2009 (28.01.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/027,483 10 February 2008 (10.02.2008) US
12/101,796 11 April 2008 (11.04.2008) US

(71) Applicant *(for all designated States except US)*: **CISCO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, CA 95134-1706 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **BEGEN, Ali** [TR/US]; 170 W. Tasman Drive, San Jose, CA 95134 (US). **ORAN, David** [US/US]; 170 W. Tasman Drive, San Jose, CA 95134 (US).

(74) Agents: **FORD, Stephen, S.** et al.; Stolowitz Ford Cowger LLP, American Bank Building, 621 SW Morrison Street, Suite 600, Portland, OR 97205 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

(54) Title: FORWARD ERROR CORRECTION BASED DATA RECOVERY WITH PATH DIVERSITY

FIG. 1

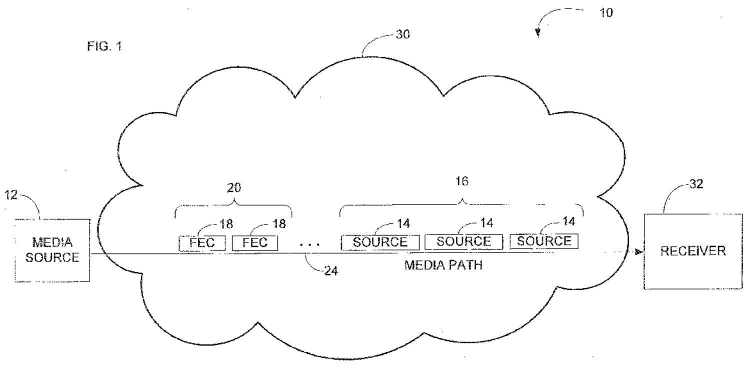

(57) Abstract: A media source sends media packets over a first media path. Repair packets are encoded from the media source packets and sent over a second different media path. Sending the source packets and repair packet over different media paths is referred to as Forward Error Correction (FEC) spatial diversity and reduces the amount of repair packet overhead required for repairing the media source packets in case of a network outage or packet loss. To provide load balancing, a first set of media streams may be sent over the first media path and a second set of media streams may be sent over the second media path. If a fault is detected on one of the media paths, then the repair packets may no longer be transmitted and the one or more media streams from the disabled media path are transmitted over the working media path.

WO 2009/099847 A2

FORWARD ERROR CORRECTION BASED DATA RECOVERY WITH PATH DIVERSITY

TECHNICAL FIELD

The present disclosure relates generally to networking.

BACKGROUND

This application claims priority to U.S. Application Ser. No. 12/101,796, filed April 11, 2008 and U.S. Provisional Patent Application Ser. No. 61/027,483, filed February 10, 2008 which are incorporated by reference in their entirety.

Video and other types of media are sensitive to packet loss and any entertainment-caliber video service should provide essentially loss-free video delivery from the media source to the media receiver(s). Packet loss can be due to congestion, link errors, and re-routing events. Individual losses or short burst losses can be adequately repaired with Forward Error Correction (FEC) or selective retransmission techniques, depending on the exact nature of the error and the delay in the network.

Selective retransmission is workable only where there is a very short round-trip time between the receivers and the transmitter. In addition, it is difficult and complex to limit the duration of certain outages in packet networks through techniques like Multi-Protocol Label Switching (MPLS) or IP Fast ReRoute (FRR).

Outages in a packet-switched or label-switched core network are usually due to a link or path failure or device/interface failure. Measurements from real deployments show that it usually takes between 50 and 500 milliseconds (ms) to restore the data path, or converge to a new one, and resume packet transmission. The packets that are sent or in flight during the outage are usually lost. In order to provide a robust video delivery, these losses should be repaired within the time frame that would satisfy the real-time requirements of the video application.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of a network that uses temporal diversity to send both source packets and FEC packets over the same media path.

FIGS. 2A and 2B are diagrams explaining how repair packet overhead limitations are associated with temporal diversity.

FIG. 3 is a diagram showing how spatial diversity is used to reduce some of the repair packet overhead shown in FIGS. 2A and 2B.

FIG. 4 is a graph that compares overhead for different media stream repair schemes.

FIG. 5A is a diagram of a network that uses Forward Error Correction (FEC) with spatial diversity.

FIG. 5B shows how FEC spatial diversity is performed by intermediary network devices in a network.

FIG. 6 is a diagram of a network that provides load balancing and FEC spatial diversity.

FIG. 7 is a diagram of a network that redirects source packets over an FEC network path when a fault condition is detected.

FIG. 8 is a flow diagram that further explains how a media source redirects the source packets in FIG. 7.

FIG. 9 is a diagram of two different networks that are used for providing FEC spatial diversity.

FIG. 10 is a flow diagram showing how FEC encoding can be changed according to detected network outage duration.

INTRODUCTION

A media source sends media packets over a first media path. Repair packets are encoded from the media source packets and sent over a second different media path. Sending the source packets and repair packet over different media paths is referred to as Forward Error Correction (FEC) spatial diversity and reduces the amount of repair packet overhead required for repairing the media source packets in case of a network outage or packet loss. To provide load balancing, a first set of media streams may be sent over the first media path and a second set of media streams may be sent over the second media path. If a fault is detected on one of the media paths, then the repair packets may no longer be transmitted and the one or more media streams from the disabled media path are transmitted over the working media path.

The foregoing and other objects, features and advantages will become more readily apparent from the following detailed description of a preferred embodiment which proceeds with reference to the accompanying drawings.

DETAILED DESCRIPTION

Referring to FIG. 1, a network 10 includes a media source 12 that transmits a media stream of source packets 14 over a packet switched network 30 to a receiver 32. The network 10 is any combination of Wide Area Networks (WANs) and Local Area Networks (LANs) that together form the Internet. The subset of this network 30 over which the correction scheme operates can be any combination of routers, switches and/or other network devices that form a portion of the Internet.

The media source 12 is any device that can transmit a media stream over a media path 24 in network 30 and the receiver 32 can be any device that receives the media stream. For example, the media source 12 can be a network server and the receiver 32 can be a Personal Computer (PC), set-top box, cable modem, Digital Subscriber Loop (DSL) modem, or any other type of wired or wireless device that received packet data. For example, the receiver 32 can also be a wireless Personal Digital Assistant (PDA) or cellular telephone. In another embodiment shown in FIG. 5B, media source 12 and/or receiver 32 may instead be intermediary devices in the network 30, such as routers, switches, gateways, etc.

Source packets 14 form the media stream sent over media path 40 and can contain any type of real-time media. For example, the source packets 14 can contain video data and/or audio data. The source packets can also include digital data. Groups of source packets 14 are referred to as a source block 16.

Forward Error Correction (FEC) packets 18 are generated from the source packets 14 by the media source 12 or another device. The FEC packets 18 are alternatively referred to as repair packets and a group of repair packets 18 used for repairing packets in source block 16 are together referred to as a repair block 20. The repair packets 18 are used to repair lost or corrupted source packets 14 in source block 16.

In this example, the source packets 14 and the FEC repair packets 18 are transmitted to the receiver(s) 32 over the same media path 24. Sending both the source packets 14 and the repair packets 18 over the same media path 24 but at different times is referred to as temporal diversity. Temporal diversity is alternatively referred to as interleaving because both the source packets 14 and FEC packets 18 are interleaved together on the same media path 24 but transmitted at different times so that their occupancy periods on the network 30 are disjoint.

At the receiver 32, missing source packets 14 are recovered by erasure decoding provided that a sufficient number of source packets 14 and repair packets 18 are received.

Recovery probability increases with the number of repair packets 18 provided per source block 16. One way to reduce the bandwidth overhead required for packet repair is to increase the size of source block 16. In other words, overhead can be decreased by increasing the ratio of the number of transmitted source packets (or bytes) 14 to the number of transmitted repair packets (or bytes) 18. However, as the size of source block 16 increases so does the playout delay relative to the encoding time or ingestion time. Thus, the size of source block 16 cannot be increased arbitrarily. The maximum tolerable time difference between the first source packet in a source block and the last repair packet in the repair block that protects this source block is referred to as the latency budget. In real-time applications, the value for the latency budget is often preferred to be small, and its value may be adjusted based on various application, user and network requirements.

Sending both the source packets 14 and FEC packets 18 on the same media path 24 causes the FEC packets 18 to likely arrive after all of the source packets 14 in the source block 16. If the source block 16 is too large, the time from the first source packet 14 in source block 16 to the last FEC packet 18 in repair block 20 may be longer than the latency budget. Violating this latency budget means that the FEC packets 18 in repair block 20 will not be able to repair all of the source packets 14 in source block 16.

For example, some source packets 14 lost at the beginning of the source block 16 may be dropped by the receiver 32 prior to receiving the necessary FEC packets 18 from repair block 20. Even if the buffer size in the receiver 32 were increased, longer latency budgets increase both the end-to-end reception delay and consequently the time required for the receiver 32 to wait before playing out the content initially or to resume the playout after a fast-forward/rewind operation in video-on-demand, or to switch among different media streams, e.g., channel change in IP television.

To explain further, FIG. 2A shows a series of source packets 14 that are transmitted over the media path 24 in FIG. 1. A value Q denotes a maximum outage duration 40 that is intended to be repairable. Any outages longer than Q are not completely repairable. For example, Q may be associated with the amount of time required for the network 30 in FIG. 1 to reconverge to another media path when there is an outage. The value of Q in one example is specified in milliseconds (ms).

The source block 16A is a group of source packets 14 that are repairable with the FEC packets 18 in repair block 20A. When an outage 40 hits the source block 16A, the number of FEC packets 18 in repair block 20A should be sufficient to recover the maximum outage duration 40. Recall that the packets 14 and 18 are sent over the same media path 24. Thus, any source packets 14 or FEC repair packets 18 transmitted during that Q-ms outage 40 are lost.

FEC codes such as Reed-Solomon codes can recover the packets 14 dropped during the outage duration 40 provided that a minimum amount of repair information is available in the repair block 20A. Some sub-optimal FEC codes may require more repair information for full recovery. Thus, at least Q-ms worth of source data needs to be contained in the repair block 20A.

One definition of overhead is the number of bytes in repair block 20A over the number of bytes in source block 16A. The maximum latency budget X refers to the maximum amount of acceptable latency for error-repair (e.g., FEC) operations in the network 10 shown in FIG. 1. For example, the maximum latency budget X may be chosen based on the amount of delay tolerable by the receiver 32 in FIG. 1 to switch between media streams. The value of X is specified in milliseconds (ms).

The latency budget could also be bounded by the amount of available buffer space in the receiver 32 in FIG. 1. More buffer space may result in a longer amount of time available to repair lost source packets 14. The lower bound for the amount of overhead required to repair a media stream exists when the size of source block 16A is equal to the maximum latency budget X as shown in FIG. 2A.

With a maximum outage duration of Q and a latency budget of X, the best case overhead required to repair outage duration Q is:

$$\text{Overhead} \geq Q/X$$

Equation 1.0

In the example shown in FIG. 2A, the latency budget X is proportional to 10 source packets and the maximum outage duration Q is proportional to 5 source packets. Thus, the minimum amount of overhead required to repair a maximum outage duration Q is:

$$\text{Overhead} = 5 \div 10 = 0.5 = 50\%$$

However, in real-time applications where temporal diversity is used, it is not possible to achieve this lower bound. For example, not all source packets 14 may be available at the beginning of the repair block 20A. Further, once the source packets 14 in source block 16A are generated after X-ms, all FEC packets 18 in repair block 20A would need to be transmitted instantaneously so that both the source packets 14 and repair packets 20 arrive within the latency budget X. Remember, that any FEC packet 18 received outside of latency budget X might not be usable for repairing lost source packets 14 in source block 16A.

FIG. 2A shows that the minimum overhead exists when the source block 16A is close to the duration of the maximum latency budget X. However, FIG. 2A also shows that the source block 16A cannot be extended to the maximum latency budget X in a temporal diversity scheme since additional time is needed during the latency period X to transmit FEC packets 18.

FIG. 2B shows one solution that addresses the physical limitations of a temporal diversity scheme where both the source packets 14 and repair packets 18 are sent over the same media stream 24 (FIG. 1). As mentioned above, the source block size 16A cannot be substantially the same size as the maximum latency budget X. This would cause any subsequently arriving FEC packets 18 in repair block 20A to be outside of the time period required to repair any of the source packets 14 in source block 16A.

Thus, (X-Q)-ms of source packets 16B are sent along with Q-ms of repair packets in repair block 20B. In other words, a smaller source block 16B is used and the remaining available time in the latency budget X is used for transmitting the FEC packets 18 in repair block 20B. The size of repair block 20B corresponds to the number of FEC packets 18 required to repair the same maximum outage duration Q. Also note that the source packets 14 and repair packets 18 do not overlap to avoid simultaneous loss during a loss event or a network outage. In this case, the overhead is as follows:

$$\text{Overhead} = Q \div (X-Q)$$

Equation 2.0

The example in FIG. 2B again has a latency budget X proportional to 10 packets and a maximum outage duration Q proportional to 5 packets. Using equation 2.0, the repair overhead is:

$$\text{Overhead} = 5 \div (10-5) = 1 = 100\%.$$

Thus, the overhead is twice as much as the theoretical lower bound achieved in FIG. 2A.

Another loss-recovery technique for network outages uses path (spatial) diversity where source packets are injected into the network 30 in FIG. 1. The same source data is encapsulated into a retransmission payload format for transmission over the other path. This spatial diversity scheme is described in co-pending application Ser. No. 11/686,321, filed March 14, 2007, entitled: UNIFIED TRANSMISSION SCHEME FOR MEDIA STREAM REDUNDANCY, which is herein incorporated by reference.

This dual source transmission technique does not require FEC encoding/decoding operations. Furthermore, extra delay is not induced in the media stream and works regardless of the outage duration as long as there are no simultaneous losses on the different paths.

Another temporal diversity technique transmits a copy of each source packet 14 Q ms after transmission of the original source packet. Here, Q is still the maximum outage duration that is intended to be repaired. This approach does not require FEC operations but introduces a delay of Q ms in the media stream and also requires 100% repair bandwidth overhead.

FEC Spatial Diversity

Referring to FIG. 3, to address some of the limitations of the repair schemes described above, the source packets 14 are sent on a first media path 52 and the FEC repair packets 18 are sent on a second different media path 54. This is referred to generally as FEC spatial diversity. Spatial or path diversity is compared with the temporal diversity shown in FIGS. 1, 2A, and 2B above where the source packets 14 and FEC packets 18 are sent over the same media path 24 but at different times.

Sending the source packets 14 and associated FEC packets 18 on different media paths 52 and 54, respectively, provides several advantages. First, less network overhead is required than either the temporal or spatial redundancy schemes described above. The spatial diversity also reduces delay introduced by the repair packets 18 when compared with temporal diversity.

The requirements on the latency budget X are also relaxed when the source packets 14 and repair packets 18 are transmitted on different media paths 52 and 54. This is because the source packets 14 and the repair packets 18 can be transmitted at the same time. Further, there is little likelihood that both the source packets 14 and the repair packets 18 will be simultaneously lost even when transmitted at the same time since the source and repair packet are transmitted on links with mutually uncorrelated error patterns.

Referring still to FIG. 3, an example latency budget X is still proportional to 10 source packets and the example maximum outage duration Q is still proportional to 5 source packets. Also note that the source block size 16A is substantially equal to the maximum latency budget X. Also note that in this example X = 2Q ms. Refer first to the prior overhead equation 2.0 associated with temporal diversity.

$$\text{Overhead} = Q \div (X\text{-}Q).$$

With spatial FEC diversity, the source block size 16A can be the same as the maximum latency budget X. Therefore, the repair overhead equation is:

$$\text{Overhead} = Q/X,$$

<u>Equation 3.0</u>

which is equal to the theoretical lower bound.

In the example shown in FIG. 3, the repair overhead is therefore:

$$\text{Overhead} = 5 \div 10 = 0.5 = 50\%.$$

Note the large difference between Equation 2.0 where the denominator = (X-Q) and Equation 3.0 where the denominator = X. This difference is particularly significant when Q is comparable to X in size. Thus, using the spatial diversity scheme shown in FIG. 3 reduces the repair overhead by 50% compared to the repair overhead required by the temporal diversity scheme shown in FIG. 2B.

The graph in FIG. 4 compares the repair overhead for different packet correction schemes. The y-axis shows the overhead introduced by FEC. The x-axis shows the latency budget X in terms of Q. The latency budget is the time difference between the first source packet in a particular source block 16 and the last FEC packet that belongs to the associated repair block 20. Each increment on the horizontal x-axis represents a ratio of the amount of repair data required to repair an outage duration Q and the amount of source data that can be sent with a latency budget X. For example, the first ratio value on the horizontal x-axis represents a ratio of 1:1 between the latency budget X and the outage duration Q. For example, a source block size of X=5 packets and repair block size of Q = 5 packets. The second value on the horizontal x-axis represents a ratio between the latency budget X and the outage duration Q of 2:1. For example, a source block size of X=10 packets and a repair block size of Q = 5 packets, etc.

The line 60 represents the overhead for path diversity without using FEC where two sets of the same source packets 14 are sent over two different media paths. The overhead is always 100% regardless of the ratio between the latency budget X and the outage duration Q.

The line 62 shows the overhead using FEC repair without path diversity. For example, line 62 represents the repair overhead associated with equation 2.0. Notice that at a ratio X:Q = 1, FEC repair without path diversity is not feasible since there is no time available during the latency budget X for sending any FEC packets. At a ratio of X:Q=2, line 62 at horizontal location 2:1 shows the overhead using FEC repair without temporal diversity is 100% as previously described above in FIG. 2B.

Line 64 shows the repair overhead associated with the FEC repair with path diversity scheme described above in FIG. 3. Line 64 shows that at a ratio of X:Q = 2:1 (second location on the horizontal x-axis), the overhead is 50%. This is compared to an overhead of 100% for FEC repair without path diversity as shown by line 62. As the ratio of X:Q increases, the latency budget X becomes substantially greater than the outage duration Q. Accordingly, the overhead required for FEC repair without path diversity as represented by line 62 starts to converge with the overhead required for FEC repair with path diversity as represented by line 64.

Thus, temporal diversity FEC requires Q<<X to have a decent overhead performance. However, spatial diversity FEC can tolerate much larger Q values at the same overhead cost. Spatial diversity FEC can also tolerate the cases Q=X where the temporal diversity FEC approach cannot. As shown in FIG 4, there is no feasible solution using the temporal diversity approach when Q=X. Accordingly, it can be seen that FEC with path diversity provides substantial improvements in repair overhead efficiency, especially when Q is comparable to X in size.

FIG. 5A shows one example of how FEC repair with path diversity is implemented. The media source 12 sends the source packets 14 over the first media path 52 in the network 30. The FEC packets 18 used for repairing the source packets 14 are sent over the second media path 54 in the network 30. It should be understood that different media paths 52 and 54 can refer to either a different logical links between two or more of the same network devices or can alternatively refer to different device media paths where a first set of network devices are used to transmit the packets 14 and one or more other different network devices are used to transmit packets 18.

In this example, the source 12 and receiver 32 operate as Real Time Protocol (RTP) mixer/translators (or equivalent for non-RTP media streams) taking a single media stream and splitting it into the main+FEC stream at the "source" and then at the receiver the main+FEC is reconstructed and either consumed locally or passed further as a non-redundant media stream.

Referring to FIGS. 5A, the media source 12 uses a first destination Internet Protocol (IP) address or Multi-Protocol Label Switching (MPLS) label 14A with the media 14B in the source packets 14. The destination IP address or other label 14A causes the source packets 14 to be sent along the first media path 52 between routers 70 and 72. Of course, there may be additional links between routers and/or switches 70 and 72 in network 30 that together provide media path 52.

Similarly, the media source 12 uses a second different destination IP address or other MPLS label 18B with the FEC data 18B in repair packets 18. The destination IP address or other label 18A causes the repair packets 18 to be sent along the second media path 54 between routers/switches 74 and 76. Recall, that for logical link diversity, the same routers/switches may be used but the interfaces associated with the source packets 14 and FEC packets 18 may be different. For node diversity, the two packet streams must not converge at any router/switch along the path.

As described above, by using path diversity as shown in FIG. 5A, repair packet overhead is reduced by allowing the source block duration to be substantially equal to the latency budget X. Thus, the ratio of source packets 14 to repair packets 18 can be increased.

The spatial FEC diversity scheme can also be used with multicast packets as described in co-pending U.S. patent application Ser. No. 11/736,463 filed on April 17, 2007, entitled MONITORING AND CORRECTING UPSTREAM PACKET LOSS, which is herein incorporated by reference. In this embodiment, the FEC packets 18 are multicast over the second media path 54. Multicast addressing used for transmitting multicast FEC packets 18 is also described in the above referenced co-pending patent application.

As another example, suppose the maximum outage duration (Q) is 500 ms and only a 1 second media delay (latency budget X) is allowable in the network 30. By combining FEC with path diversity, the repair overhead is reduced from 100% to 50%. If a 50% overhead is desirable, the delay is reduced from 1.5 seconds to 1 second. Delay reduction is beneficial in reducing the memory requirements for the network devices such as a Digital Content Manager (DCM) that generates the FEC packets 18.

In an alternate embodiment, multiple different servers 12 are used to send the different media streams (source diversity) and to improve the resiliency against packet losses. The source video is also protected by FEC at each server 12. Before each server starts streaming video, the receiver 32 first runs a rate allocation algorithm to determine the rate for each server 12, and then runs a packet partitioning algorithm to ensure that no packet is sent by more than one server 12. In another embodiment, a first server 12 may be used for sending the source packets 14 and a second different server may be used for sending the FEC repair packets 18 associated with the source packets 14.

FIG. 5B shows an alternative embodiment where two intermediary network devices 78 and 79, such as routers, switches, gateways, etc. operate as the two devices that establish spatial FEC diversity between the source packets 14 and FEC packets 18 over network 30. It should be understood, that the network 30 represents any portion of the Internet network where it may be advantageous to use temporal FEC diversity. For example, the initial link between a media source and the network 30 may have a highly reliable link or may have bandwidth restrictions that prevent or do not warrant temporal FEC diversity. Thus, temporal

FEC diversity may be implemented in another portion of network 30 where link conditions and network bandwidth warrant spatial FEC diversity.

In FIG. 5B, the first intermediate network device 78 may receive one or more media streams 77A from a media source or from another section of the network. Similarly as described above in FIG. 5A, the intermediate network device 78 attaches a destination IP address or MPLS label 14A to the source packets 14 from the media stream 77A that cause the source packets 14 to travel over the first media path 52 in network 30 to a second intermediate network device 79.

The intermediate network device 78 may receive the FEC packets 18 over the same link that carries media stream 77A. Alternatively, the intermediate network device 78 may generate the FEC packets 18 from the media stream 77A. Either way, an IP address or MPLS label 18A is attached to the received or derived FEC packets 18 that cause the FEC packets to travel over the second media path 54 to the second intermediate network device 77B. Again the two media paths 52 and 54 may be different logical links or different physical paths between different network devices.

The second intermediate network device 79 may use the FEC packets 18 to repair any of the missing or corrupted source packets 14 received from network device 78. The second intermediate network device 79 then forwards the corrected media stream 77B to the one or more receiver(s) 32 shown in FIG. 5A.

Distributing Source Packets

Referring to FIG. 6, when multiple media streams (say 500) are transmitted between the media source 12 and different, or the same, receiver(s) 32. The FEC streams 80B generated from a first set of source streams 80A and the set of FEC streams 82B generated from a second set of source streams 82A are equally split among the two different media paths (data planes) 52 and 54 such that each media path 52 and 54 carries approximately 50% of the source packets and 50% of the FEC packets. This provides load balancing for both planes 52 and 54 regardless of the number of source streams 80A and 82A.

For example, the load is balanced on both media paths 52 and 54 by transmitting half of the source streams 80A and half of the FEC streams 82B on the first media path 52. The other half of the media streams 82A and the other half of the FEC streams 80B derived from media stream 80A are transmitted on the second media path 54. That is, the source data for 250 of the media channels 80A and the FEC data 82B for the other 250 channels 82A are transmitted on the first media path 52. The second media path 54 carries the rest of the source data for media channels 82A and the FEC data 80B for the first group of media channels 80A. This provides equal load on both data planes 52 and 54 regardless of the FEC overhead.

Referring to FIG. 7, when the duration of repair block 20 is large enough relative to a failure detection time, failures can be detected on either media path 52 or 54. In this case, the media source 12 can switch to single stream operation and start sending source data only on the working data plane 52 or 54 and still achieve zero loss at the receivers. This is again a lower-overhead redundancy approach compared to full stream redundancy.

A fault 86 may be detected on the media path 52 when a Negative ACKnowledgement (NACK) or any other routing algorithm status message 88 is sent back to the media source 12 by one of the routers 70 or 72 or from the receiver 32. After receiving the NACK message 88, the media source 12 stops sending FEC packets 18 over media path 54 and redirects the source packets 14 from media path 52 to media path 54.

Example types of fault detection and notification can include a Multi-Protocol Label Switching (MPLS)-capable network, receiving a Label-Switched Path (LSP) teardown signal or a LSP reroute request signal in protocols like Label Distribution Protocol (LDP) or Resource Reservation Protocol for Traffic Engineering (RSVP-TE). A link-state update can be used in the Open Shortest Path First (OSPF) or Intermediate System (IS) ISIS Interior Gateway Protocol (IGP) or a Border Gateway Protocol (BGP) route withdrawal is received covering the media destination on one of the two paths. A notification of an interface failure can be received, or adjacency failure via Bi-directional Forwarding Detection (BFD) protocol can be used at the source system.

Referring to both FIGS. 7 and 8, the media source 12 sends source packets 14 over the first media path 52 in operation 100. The FEC packets 18 for the source packets 14 are encoded in operation 102 and sent over the second media path 54 in operation 104. The source packets 14 could also be sent over media path 54 and the FEC packets 18 could be sent over media path 52 at the same time.

In operation 106, the media source 12 monitors the media paths 52 and 54 for outages 86. Upon detecting a long-duration failure 86 on media path 52 in operation 108, the media source 12 switches to a single stream operation and starts sending source packets 14 only (no FEC) on the working media path 54 in operation 110.

No longer bounded by the repair block size, the media source 12 can perform a full recovery by transmitting the source packets 14 over media path 54 that were lost on the media path 52. This provides a lower-overhead redundancy approach compared to full stream redundancy. Conventional FEC schemes continue to send FEC packets and, hence, are still required to recover missing data within the latency budget. Thus, full recovery probability is significantly lower. If the disabled media path 52 comes back up, media source

12 may start sending source packets 14 again on media path 52 and start sending the FEC packets 18 again on media path 54. Alternatively, the media source 12 may continue to send the source packets 14 over media path 54 and start sending the associated FEC packets 18 on the recovered media path 52. The recovery of the media path 52 can be detected via routing protocol status messages.

If the media path 54 carrying the FEC packets 18 goes down, the media source 12 does not need to take any action. However, when the media path 52 normally carrying the source packets 14 goes down, the media source 12 stops sending the FEC packets 18 and starts transmitting the source packets 14 on the working media path 54.

The switching scheme described in FIGS. 7 and 8 also works in conjunction with the load balancing scheme previously shown in FIG. 6. For example, FIG. 6 shows how multiple media streams 80A and 82A were sent on the different media paths 52 and 54, respectively. In FIG. 6, the repair packets 80B for media streams 80A are normally transmitted over media path 54 and the repair packets 82B for media streams 82A are normally transmitted over media path 52.

If a long-duration failure is detected on media path 52, the media source 12 might stop sending the repair packets 80B on the media stream 54 and instead starts transmitting media streams 80A on media path 54. Similarly, a long-duration failure may be detected on media path 54. Accordingly, the media source 12 stops sending the repair packets 82B on the media stream 52 and instead starts transmitting media streams 82A on media path 52.

FIG. 9 shows how the FEC spatial diversity schemes described above can be used with different networks. For example, a first cable network 154 may include a Cable Modem Termination System (CMTS) 152 located at a cable headend. The CMTS 152 sends packets over a media path 156 in the cable network 154 to one or more cable modems 158. The cable modems 158 then convert the packets into signaling used by a media endpoint 32, such as a television, set-top box, personal computer, or other wired or wireless device.

A Digital Subscriber Line (DSL) network 162 may include a Digital Subscriber Line Access Multiplexer (DSLAM) 166 located at a telephone company central office. The DSLAM 166 sends packets over a media path 164 in the DSL network 162 to one or more DSL modems 160. The DSL modems 160 then convert the packets into digital signals used by the same media endpoint 32.

The media source 12 can send the different source packets 14 over the cable network 154 or DSL network 162 and send the associated FEC repair packets 18 over the other cable or DSL network. In this example, the media source 12 sends the source packets 14 over Internet network 150 to the CMTS 152. The CMTS 152 forwards the source packets 14 over the media path 156 in cable network 154 to the cable modem 158. The cable modem 158 then converts the source packets 14 into signaling compatible with media endpoint 32.

In this example, the media source 12 sends the source packets 14 over the media path 156 of cable network 154 and sends the FEC packets 18 over the media path 164 in DSL network 162. If there is a failure in the media path 156 of cable network 154, then the media source 12 starts sending the source packets 14 over Internet network 150 to the DSLAM 166. The DSLAM 166 then forwards the source packets 14 over media path 164 to the DSL modem 160. The media endpoint 32 upon failing to receive source packets from the cable modem 158 then switches to receiving the source packets 14 from the DSL modem 160.

FIG. 10 shows another embodiment where the amount of FEC data is dynamically adapted according to the monitored outage duration Q and/or according to a different latency budget for the network 30 in FIG. 5. If the X and/or Q values change over time, the source packet encoder and FEC encoder can be reconfigured to adjust the size of the source blocks 16 and/or the amount of FEC protection to correspond with the new X and/or Q values.

Referring both to FIGS. 5 and 10, X and Q values for network 30 are obtained in operation 170. The X and Q values may be obtained through empirical data measured for the network 30 and/or may be set by the network operator. The source and FEC encoders in media source 12 are then configured in operation 172 according to the X and Q values. For example, the source block size 16 is generated according to the latency budget X and the repair block size 20 is generated to correct the maximum outage duration 40 shown in FIG. 3.

If a dynamic latency budget operation is enabled in operation 174, then the media source 12 in operation 176 monitors for any changes to the latency budget. For example, the media source 12 may receive a new latency budget value X from a network administrator.

If dynamic FEC is enabled in operation 174, then the media source 12 in operation 176 tracks all further outage durations. For example, the media source 12 may receive NACK messages back from the receivers 32 that indicate the number of source packets 14 that are typically dropped on media path 52. The feedback might be in NACKs, if retransmission is configured for the stream. However, reception reports can also be used such when RTP is used as the transport protocol, RTCP receiver reports could be used.

Alternatively, the network operator may manually modify the Q value for particular media streams. For example, during a football telecast, the network administrator may want to generate the highest possible repair block size 20 to increase the reliability of the transmitted media stream. If the new latency budget X is different from the previously configured latency budget X in operation 178, the encoder in the media source is adjusted in operation 180 according to the new X value. For example, if the new latency budget X is larger, the encoder in the media source 12 may be adjusted to generate larger source blocks 16. If the new latency budget X is smaller, smaller source blocks 16 may be generated.

If the new outage duration Q is different from the previously configured outage duration Q in operation 178, the FEC encoder is adjusted in operation 180 according to the new Q value. For example, if the new Q value is smaller than the previous Q value, then the FEC encoder may be adjusted to generate a smaller FEC repair block 20. If the new Q value is larger than the previous Q value, then the FEC encoder may be adjusted to generate a larger FEC repair block 20. It should understood that any of the examples shown and described in FIGS. 6-10 could be different parts of the network where the media source 12 and the receiver 32 are instead other intermediate nodes 78 and 79 as shown in FIG. 5B. For example, any of the examples described in FIGS. 6-10 could be used in the portion of the network described above in FIG. 5B.

The system described above can use dedicated processor systems, micro controllers, programmable logic devices, or microprocessors that perform some or all of the operations. Some of the operations described above may be implemented in software and other operations may be implemented in hardware.

For the sake of convenience, the operations are described as various interconnected functional blocks or distinct software modules. This is not necessary, however, and there may be cases where these functional blocks or modules are equivalently aggregated into a single logic device, program or operation with unclear boundaries. In any event, the functional blocks and software modules or features of the flexible interface can be implemented by themselves, or in combination with other operations in either hardware or software.

Having described and illustrated the principles of the invention in a preferred embodiment thereof, it should be apparent that the invention may be modified in arrangement and detail without departing from such principles. Claim is made to all modifications and variation coming within the spirit and scope of the following claims.

Claims

1. An apparatus, comprising:

one or more network devices configured to send source packets over a first link or device media path, the one or more network devices further configured to encode repair packets from the source packets and send the repair packets over a second different link or device media path.

2. The apparatus according to claim 1 further comprising determining a latency budget X and configuring the one or more network devices to send source blocks of source packets having a duration approximately equal to the latency budget X.

3. The apparatus according to claim 2 further comprising determining a maximum outage duration Q and configuring the one or more network devices to encode and send repair blocks of repair packets on the second link or device media path large enough to repair the media packets for the maximum outage duration Q.

4. The apparatus according to claim 3 wherein the outage duration Q can be comparable or equal to the latency budget X.

5. The apparatus according to claim 3 wherein:

the duration of the source block is approximately equal to the latency budget X;

the duration of the repair block is approximately equal to the maximum outage duration Q; and

a bandwidth overhead for sending the repair block is approximately equal to Q/X.

6. The apparatus according to claim 1 wherein the one or more network devices are further configured to:

send a first set of media streams over the first link or device media path;

second a second set of media streams over the second link or device media path;

encode a first group of repair packets for the first set of media streams and send the first group of repair packets over the second link or device media path; and

encode a second group of repair packets for the second set of media streams and send the second group of repair packets over the first link or device media path.

7. The apparatus according to claim 6 wherein the one or more network devices are further configured to:

detect a failure on the first link or device media path;

stop sending the second group of repair packets on the second link or device media path; and

start sending the first set of media streams on the second link or device media path.

8. The apparatus according to claim 1 wherein the one or more network devices are further configured to:

detect a failure on the first link or device media path;

stop sending the repair packets on the second link or device media path when the failure is detected on the first link or device media path; and

redirect the source packets from the first link or device media path to the second link or device media path.

9. The apparatus according to claim 8 wherein the first link or device media path is established over a first type of access network and the second link or device media path is established over another different type of access network.

10. The apparatus according to claim 1 wherein the source packets and repair packets are multicast over the first and second link or device media path.

11. The apparatus according to claim 1 wherein the source packets and repair packets are unicast over the first and second link or device media path.

12. A method, comprising:

generating Forward Error Correction (FEC) packets from media for repairing the source packets;

transmitting the source packets on a first network path or link to one or more network devices; and

transmitting the FEC packets on a second different network path or link to the same one or more network devices.

13. The method according to claim 12 wherein the first and second network path or link may be either different logical links or different network device paths.

14. The method according to claim 12 further comprising:

detecting a network failure on the first network path or link;

discontinuing transmission of the FEC packets on the second network path or link; and

redirecting the transmission of the source packets over the second network path or link.

15. The method according to claim 12 further comprising:

determining a latency budget duration X;

determining an outage duration Q;

generating source blocks of source packets proportional to the latency budget duration X;

sending the source blocks of source packets over the first network path or link;

generating repair blocks of FEC packets proportional to the outage duration Q; and

sending the repair blocks of FEC packets over the second network path or link.

16. The method according to claim 15 further comprising adaptively changing the value of Q based on empirical data, administrative policies, or an importance of the media in the source packets.

17. The method according to claim 15 further comprising dynamically changing the value of X based on empirical data, administrative policies, or an importance of the media in the source packets.

18. The method according to claim 15 wherein the duration of the source blocks and the duration of the repair blocks have an associated repair overhead of approximately Q/X.

19. The method according to claim 15 wherein the outage duration Q is comparable or equal to the latency budget X.

20. The method according to claim 12 further comprising:

using a first Internet Protocol (IP) destination address or MPLS label in the source packets associated with the first network path or link; and

using a second IP destination address or MPLS label in the FEC packets associated with the second network path or link.

21. The method according to claim 12 further comprising:

sending a first set of media streams over the first network path or link;

sending a second set of media streams over the second network path or link;

sending a first set of FEC packets for the first set of media streams over the second network path or link; and

sending a second set of FEC packets for the second set of media streams over the first network path or link.

22. An apparatus, comprising:

a receiving network device configured to receive one or more media streams from a transmitting network over a first network path or link, the receiving network device further configured to receive Forward Error Correction (FEC) packets over a second network path or link for repairing the one or more media streams received from the transmitting network device over the first network path or link.

23. The apparatus according to claim 22 wherein the media streams are load balanced over the first and second network path or link with a first set of media streams being received over the first network path or link, a second set of media streams being received over the second network path or link, a first set of FEC packets for the first set of media streams being received over the second network path or link, and a second set of FEC packets for the second set of media streams being received over the first network path or link.

24. The apparatus according to claim 23 wherein the receiving network device is further configured to:

receive the first set of media streams over the first network path or link and receive the second set of media streams over the second network path or link until a fault is detected on the first or second network path or link; and

receive both the first set of media steams and the second set of media streams over the first or second network path that does not have the detected fault; and

no longer receive any FEC packets on the first or second network path or link that has the detected fault.

25. The apparatus according to claim 23 further comprising:

receiving the one or more media streams on the first network path or link from the transmitting network device until a fault is detected on the first network path or link; and

no longer receiving the FEC packets from the transmitting network device on the second network path or link when the fault is detected on the first network path or link; and

receiving the one or more media streams previously received on the first network path or link from the transmitting network device on the second network path or link.

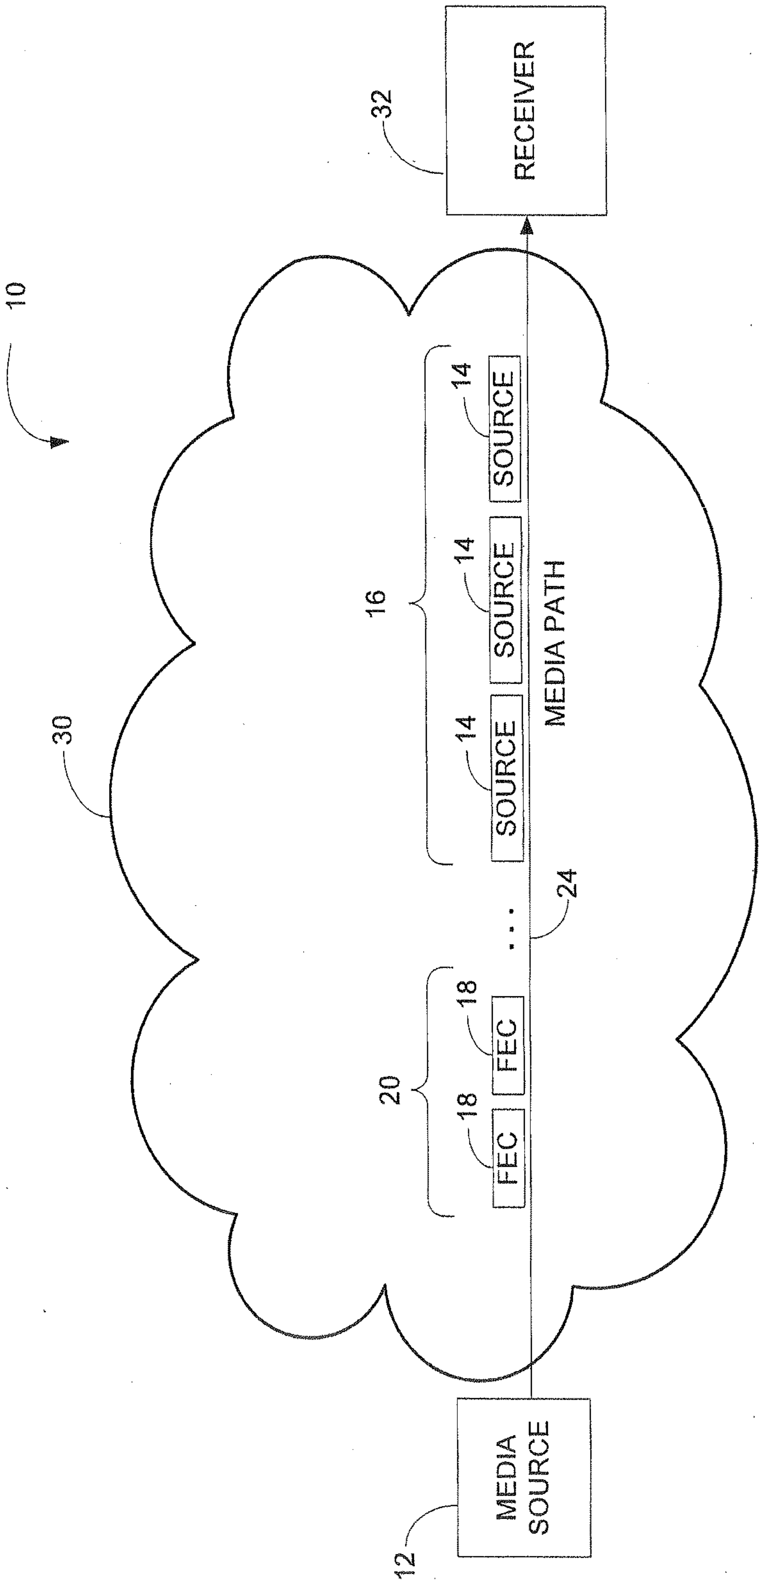
10
RECEIVER
32
30
16
SOURCE
14
SOURCE
14
SOURCE
14
MEDIA PATH
24
20
FEC
18
FEC
18
MEDIA SOURCE
12

FIG. 1

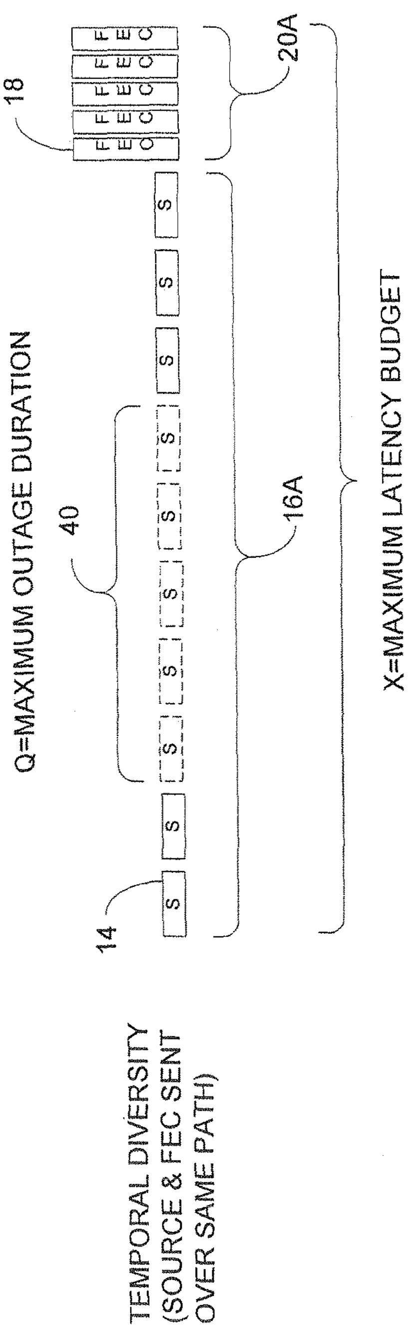
TEMPORAL DIVERSITY (SOURCE & FEC SENT OVER SAME PATH)
Q=MAXIMUM OUTAGE DURATION
40
14
S
18
FEC
16A
20A
X=MAXIMUM LATENCY BUDGET

FIG. 2A

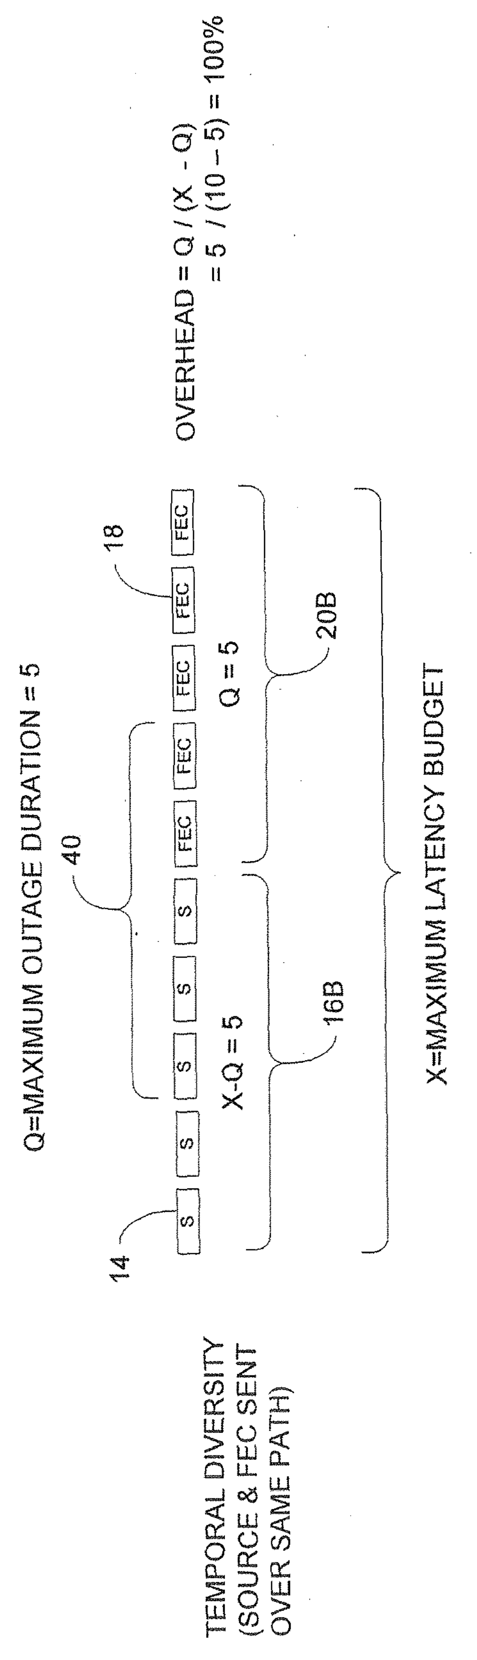
TEMPORAL DIVERSITY (SOURCE & FEC SENT OVER SAME PATH)
Q=MAXIMUM OUTAGE DURATION = 5
40
14
S
18
FEC
X-Q = 5
Q = 5
16B
20B
OVERHEAD = Q / (X - Q)
= 5 / (10 – 5) = 100%
X=MAXIMUM LATENCY BUDGET

FIG. 2B

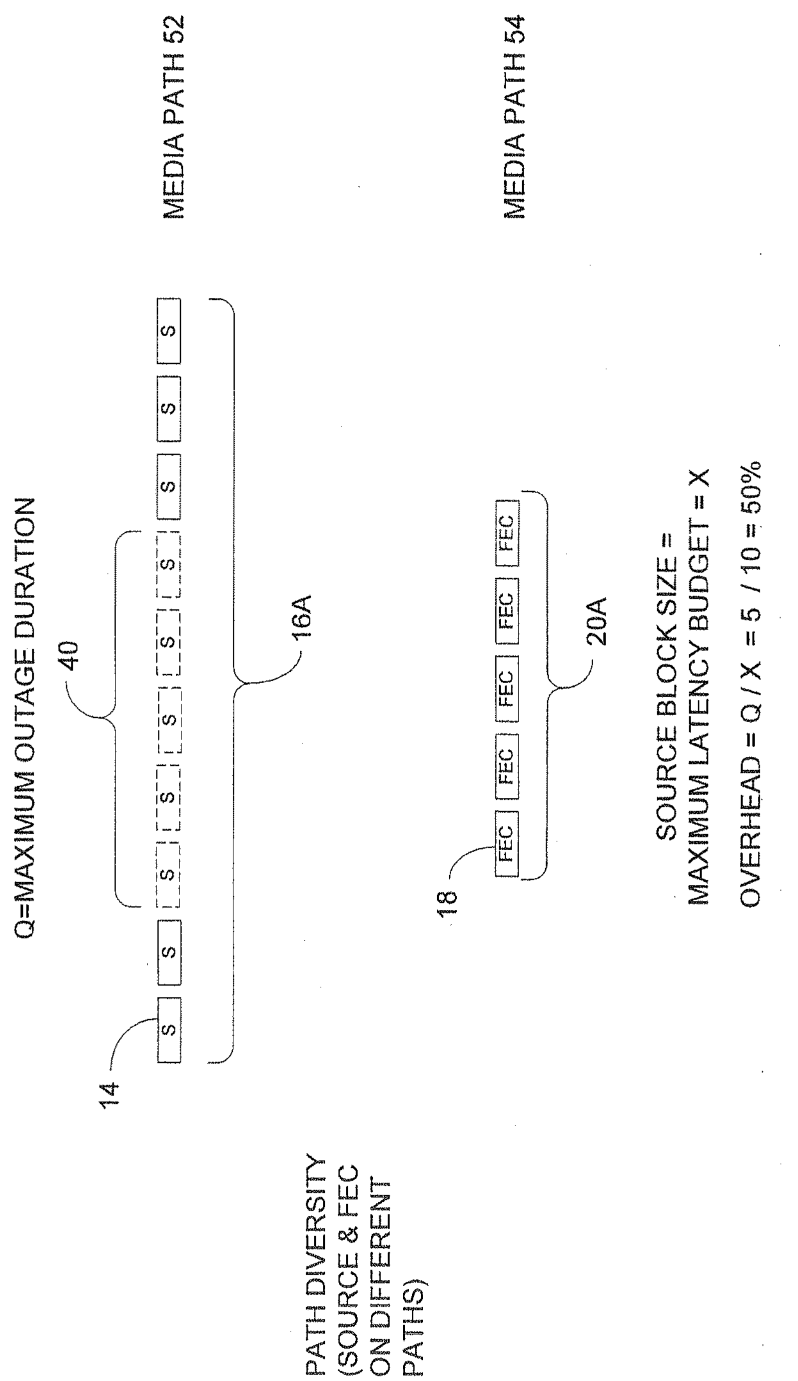
Q=MAXIMUM OUTAGE DURATION
40
14
S
16A
MEDIA PATH 52
18
FEC
20A
MEDIA PATH 54
SOURCE BLOCK SIZE =
MAXIMUM LATENCY BUDGET = X
OVERHEAD = Q / X = 5 / 10 = 50%
PATH DIVERSITY
(SOURCE & FEC
ON DIFFERENT
PATHS)

FIG. 3

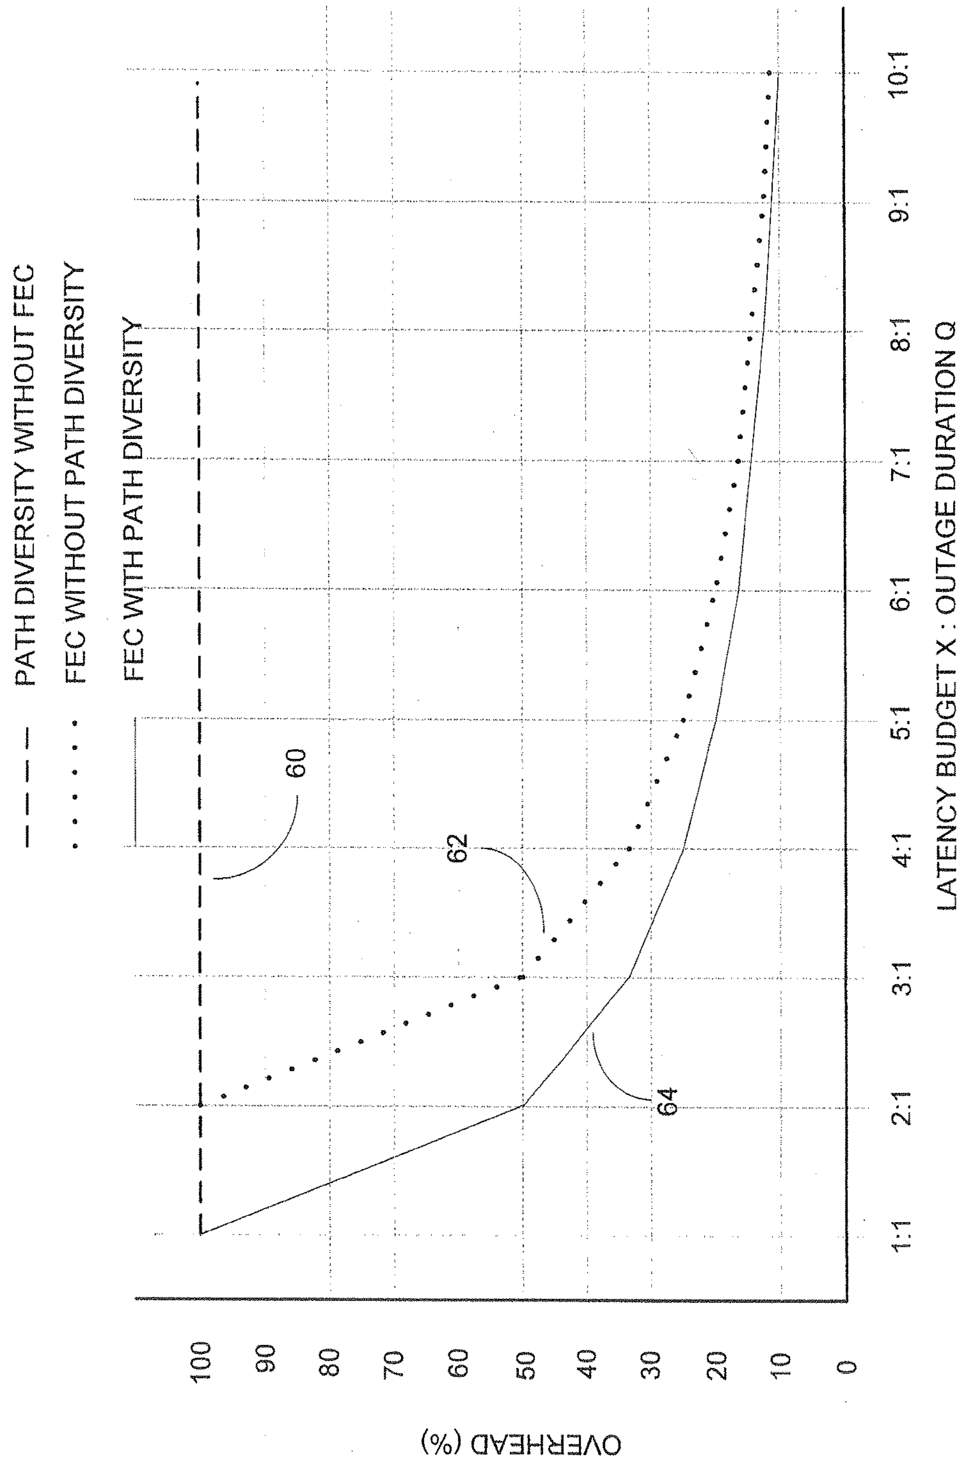
PATH DIVERSITY WITHOUT FEC
FEC WITHOUT PATH DIVERSITY
FEC WITH PATH DIVERSITY
60
62
64
OVERHEAD (%)
100
90
80
70
60
50
40
30
20
10
0
1:1
2:1
3:1
4:1
5:1
6:1
7:1
8:1
9:1
10:1
LATENCY BUDGET X : OUTAGE DURATION Q

FIG. 4

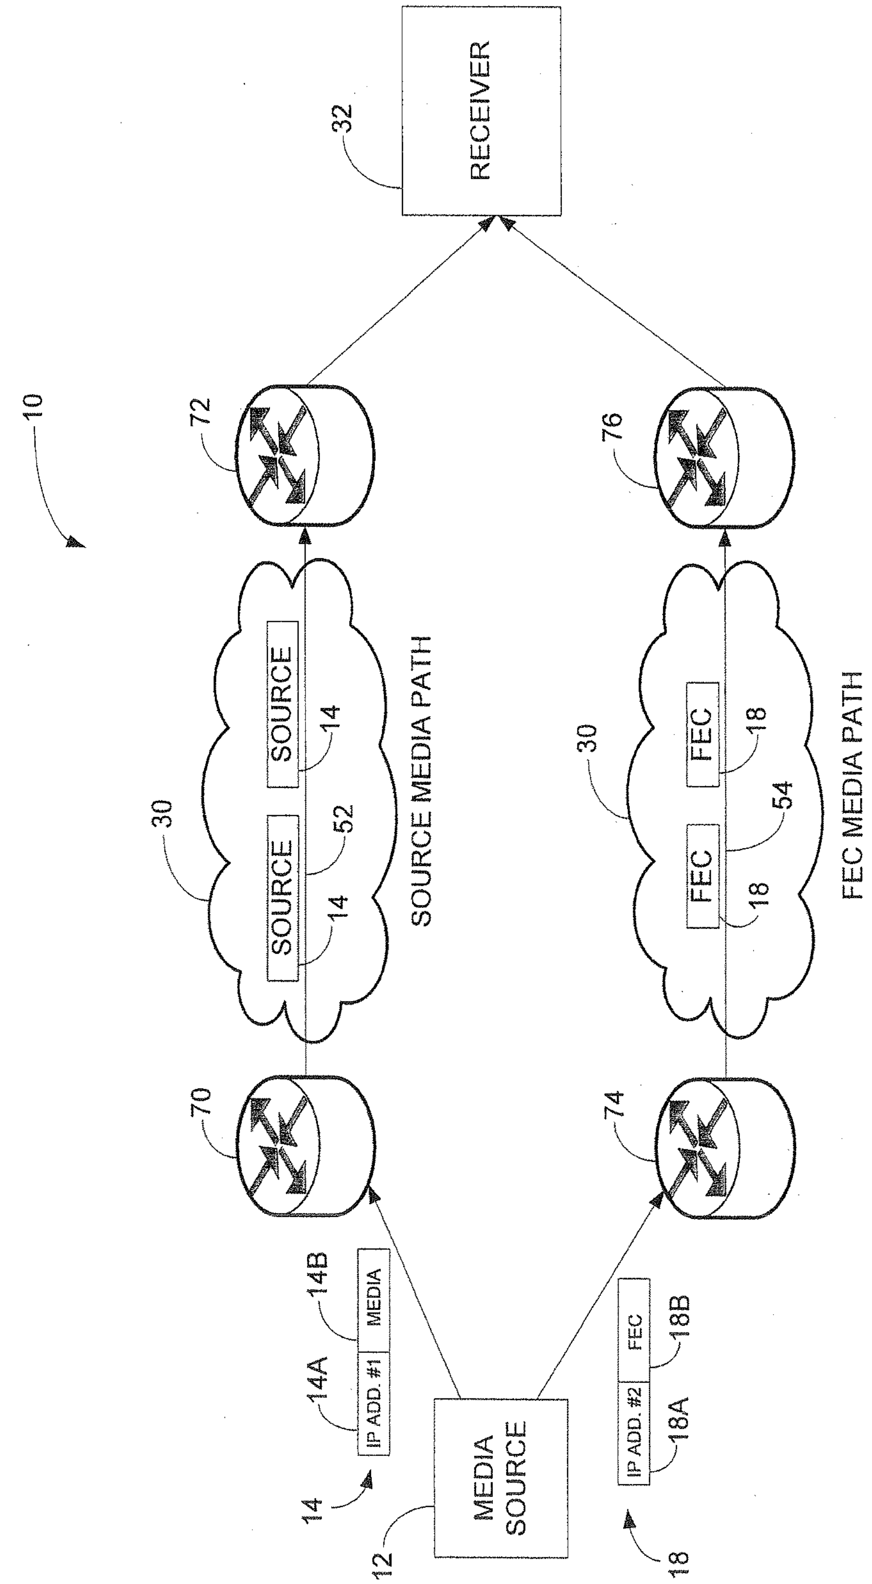
RECEIVER
32
10
72
76
70
74
SOURCE
14
SOURCE
52
14
30
SOURCE MEDIA PATH
FEC
18
FEC
54
18
30
FEC MEDIA PATH
14
IP ADD. #1
14A
MEDIA
14B
MEDIA SOURCE
12
IP ADD. #2
18A
FEC
18B
18

FIG. 5A

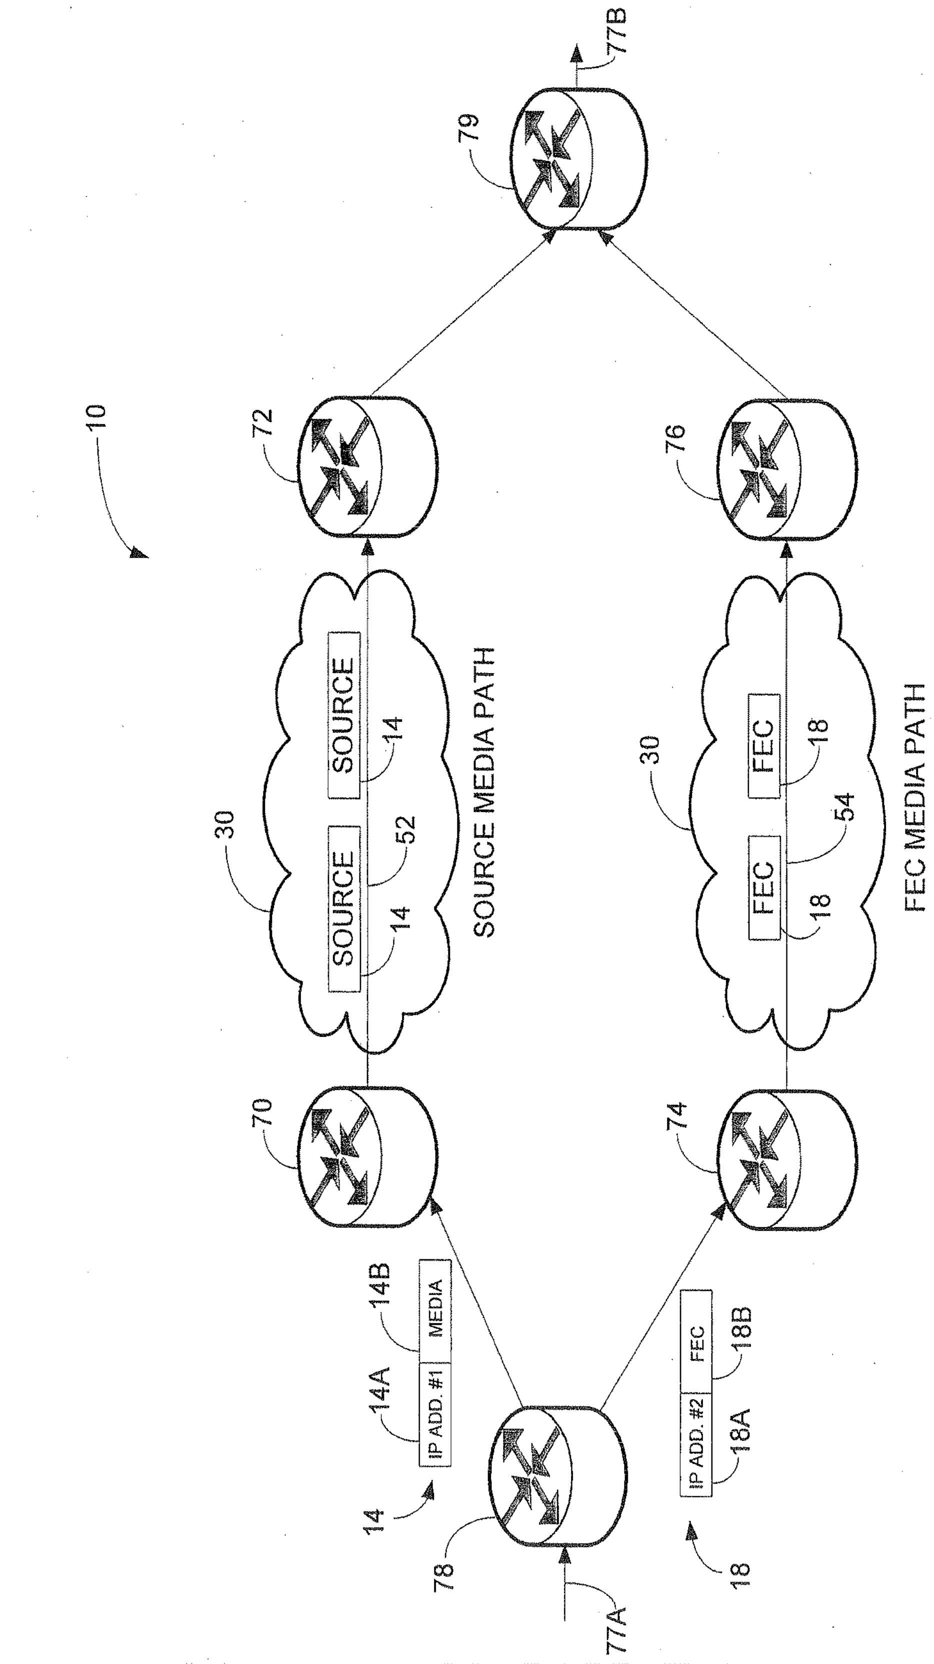
10
77B
79
72
76
30
SOURCE
SOURCE
14
52
14
SOURCE MEDIA PATH
30
FEC
FEC
18
54
18
FEC MEDIA PATH
70
74
14
14A
IP ADD. #1
14B
MEDIA
78
77A
18
18A
IP ADD. #2
18B
FEC

FIG. 5B

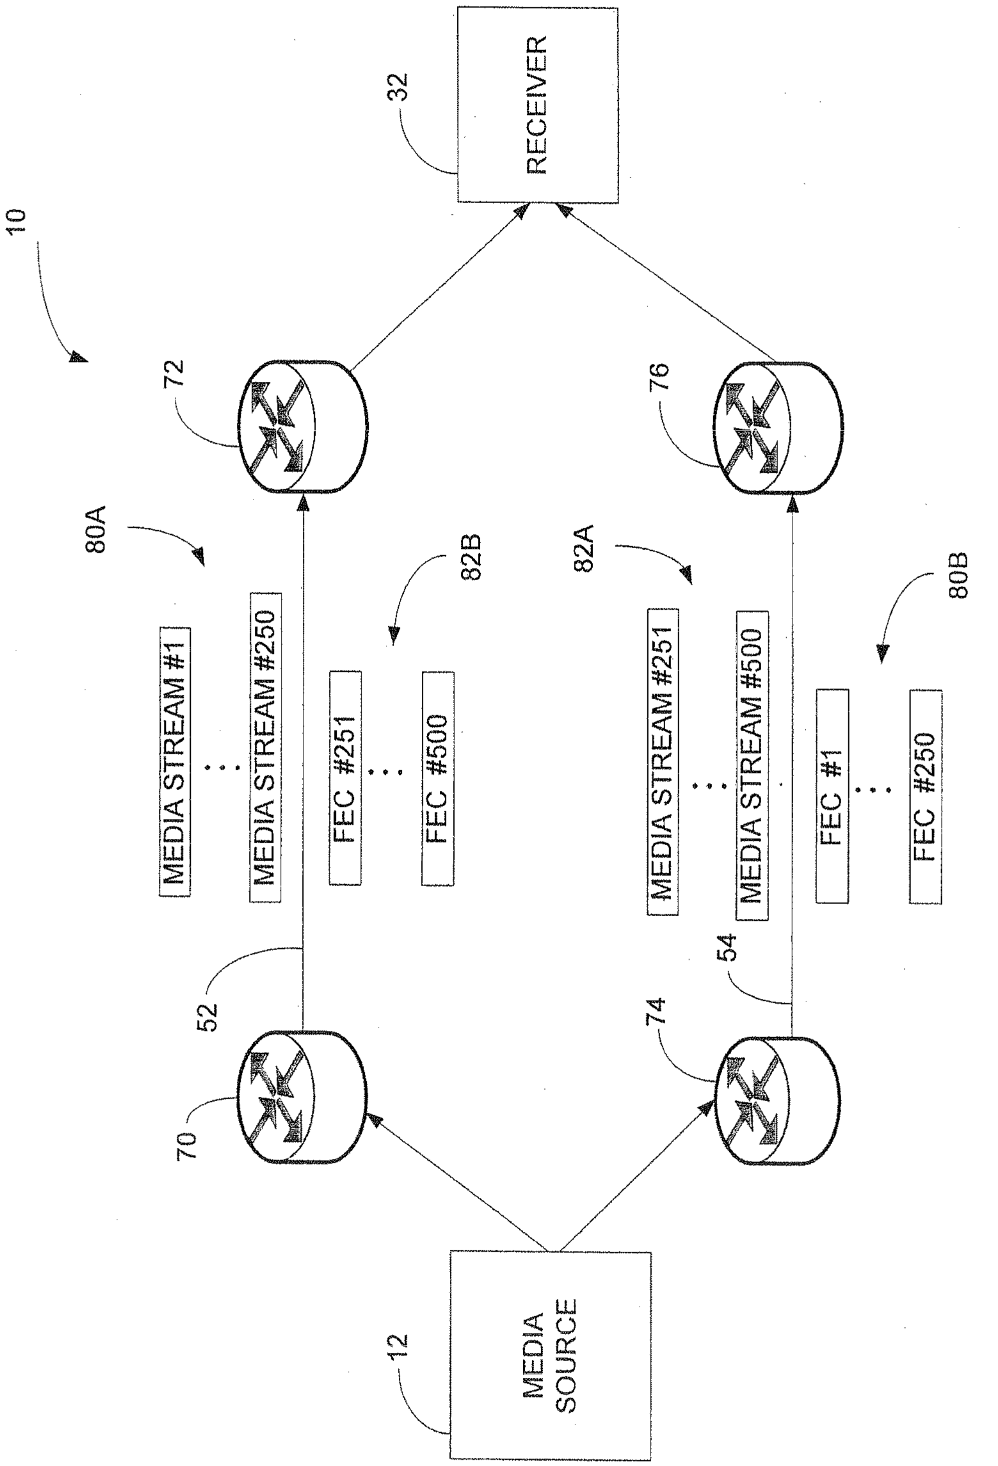
10
MEDIA SOURCE
12
70
52
74
54
72
76
RECEIVER
32
80A
80B
82A
82B
MEDIA STREAM #1
MEDIA STREAM #250
FEC #251
FEC #500
MEDIA STREAM #251
MEDIA STREAM #500
FEC #1
FEC #250

FIG. 6

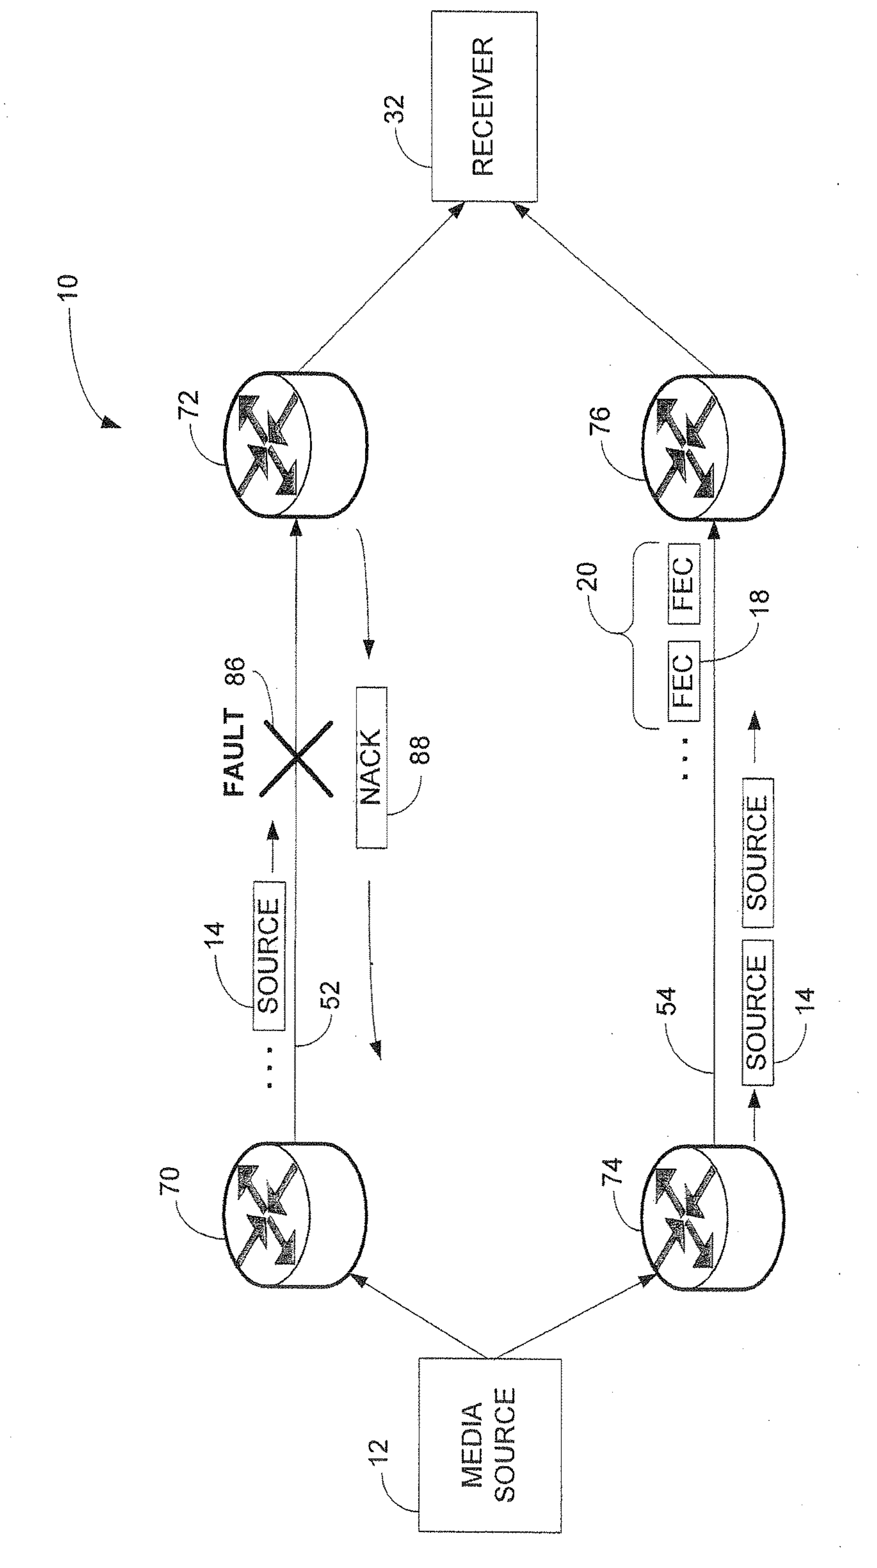
RECEIVER
32
10
72
76
20
FEC
FEC
18
FAULT
86
NACK
88
14
SOURCE
52
54
SOURCE
SOURCE
14
70
74
MEDIA
SOURCE
12

FIG. 7

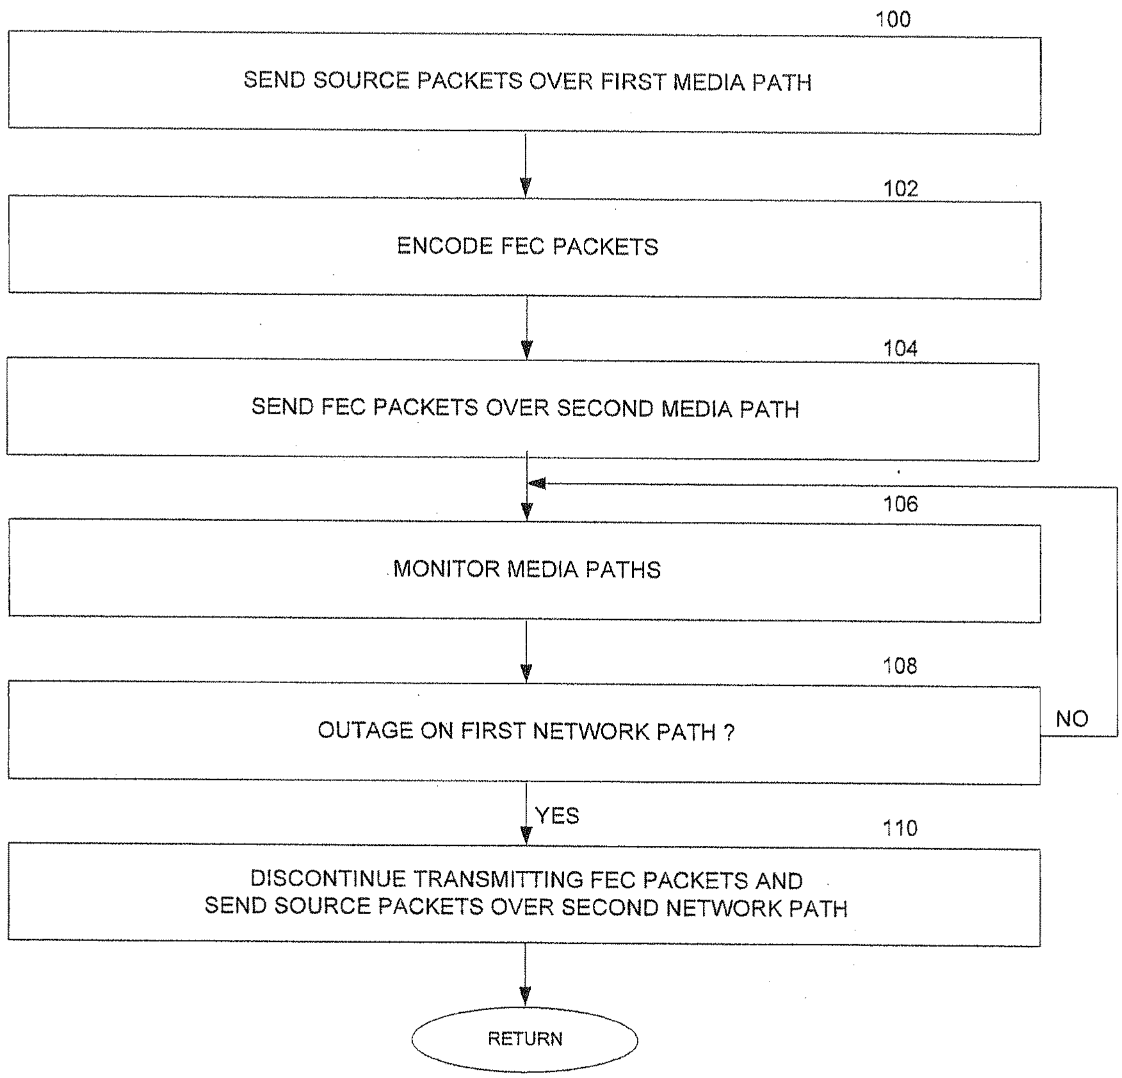
100
SEND SOURCE PACKETS OVER FIRST MEDIA PATH
102
ENCODE FEC PACKETS
104
SEND FEC PACKETS OVER SECOND MEDIA PATH
106
MONITOR MEDIA PATHS
108
OUTAGE ON FIRST NETWORK PATH ?
NO
YES
110
DISCONTINUE TRANSMITTING FEC PACKETS AND SEND SOURCE PACKETS OVER SECOND NETWORK PATH
RETURN

FIG. 8

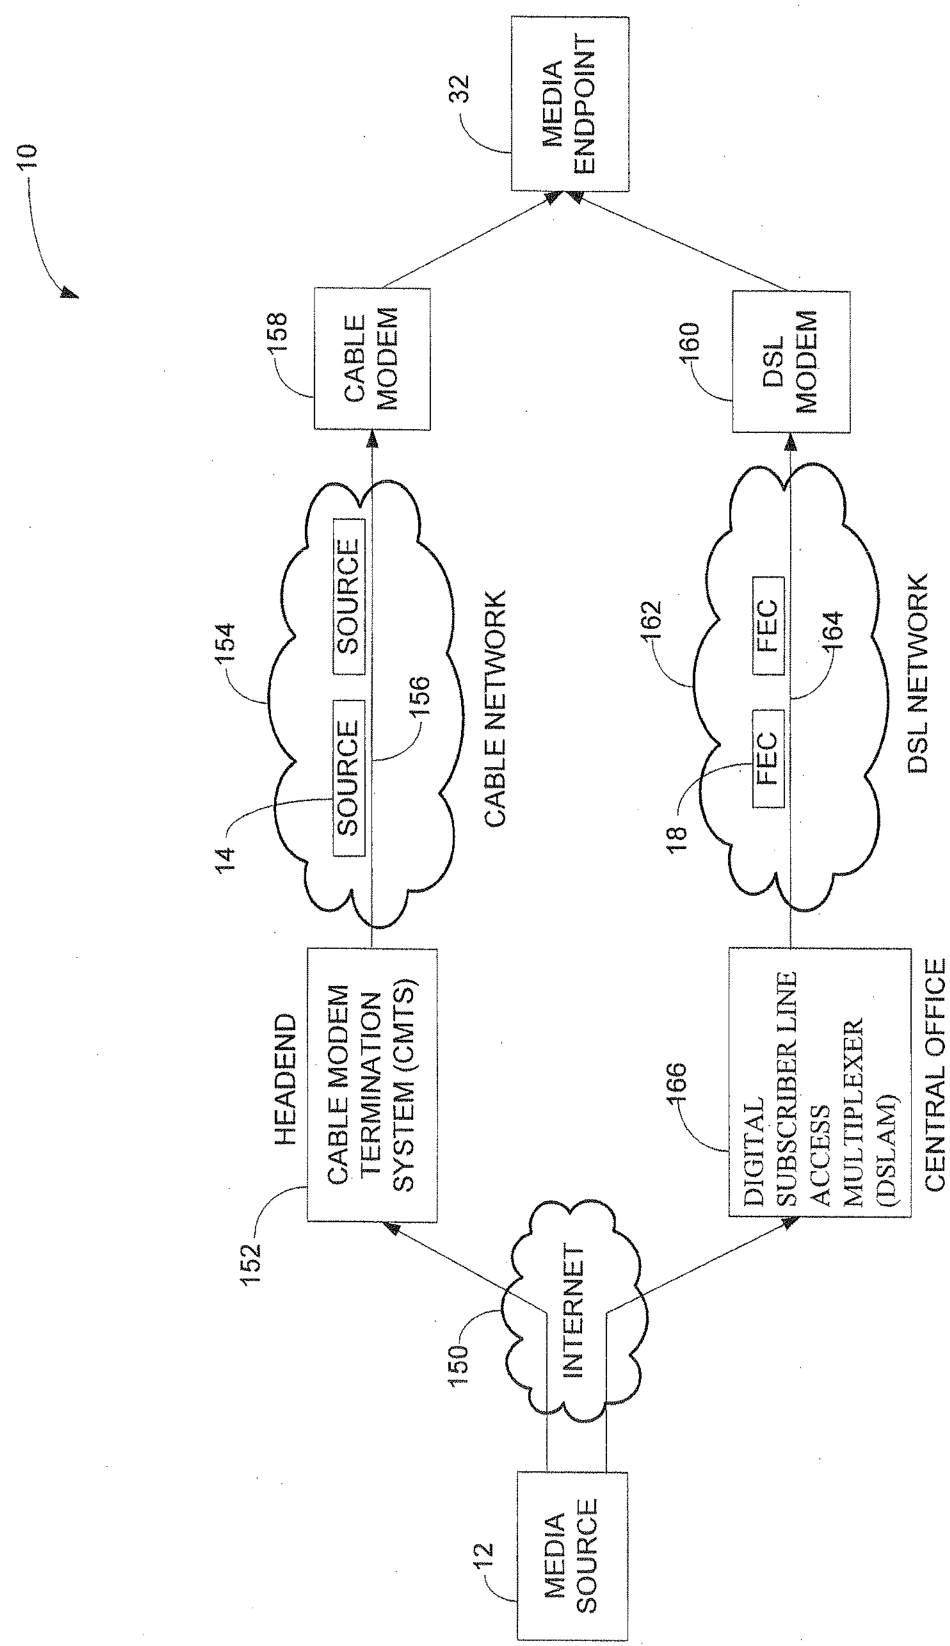
10
MEDIA ENDPOINT
32
CABLE MODEM
158
DSL MODEM
160
154
SOURCE
SOURCE
156
14
CABLE NETWORK
162
FEC
FEC
164
18
DSL NETWORK
HEADEND
CABLE MODEM TERMINATION SYSTEM (CMTS)
152
166
DIGITAL SUBSCRIBER LINE ACCESS MULTIPLEXER (DSLAM)
CENTRAL OFFICE
INTERNET
150
MEDIA SOURCE
12

FIG. 9

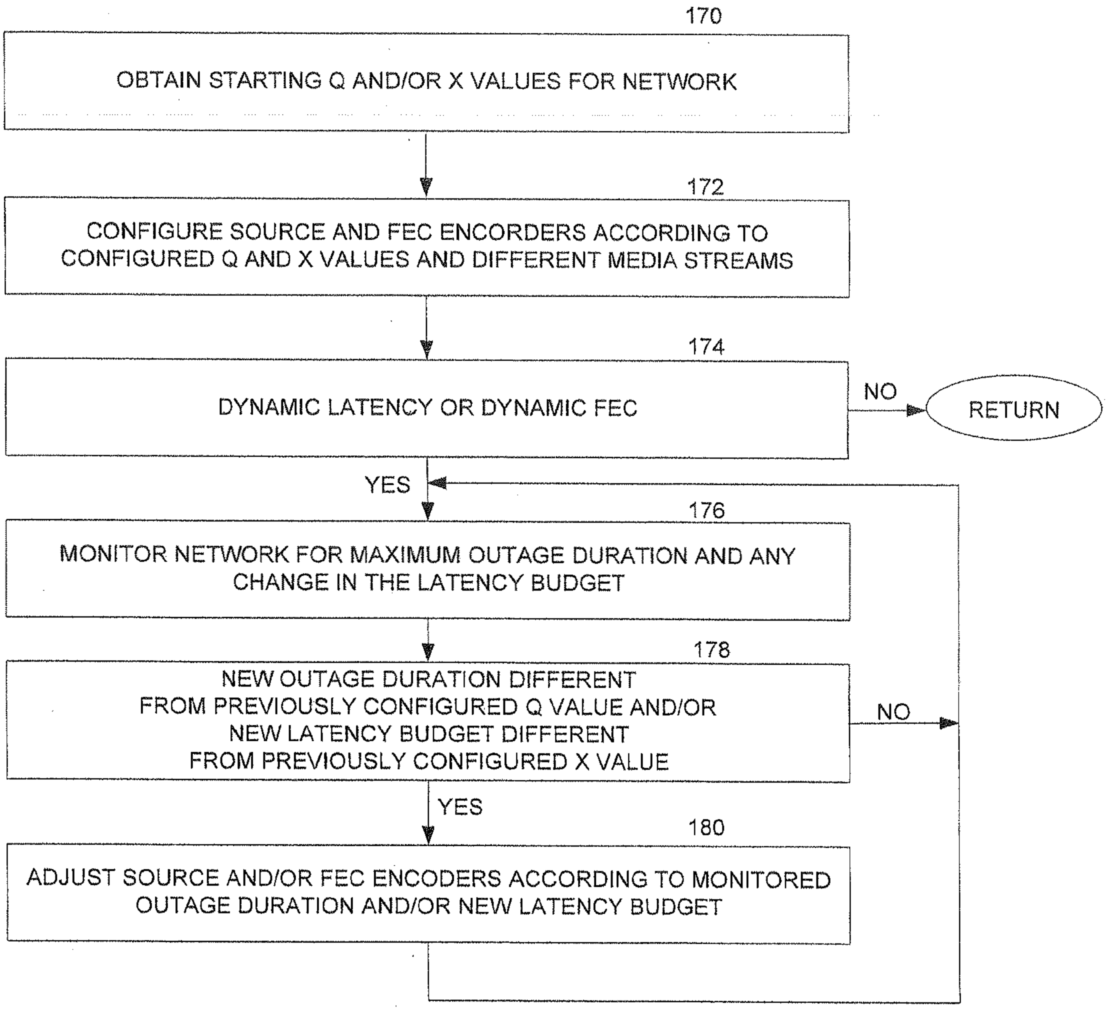
170
OBTAIN STARTING Q AND/OR X VALUES FOR NETWORK
172
CONFIGURE SOURCE AND FEC ENCORDERS ACCORDING TO CONFIGURED Q AND X VALUES AND DIFFERENT MEDIA STREAMS
174
DYNAMIC LATENCY OR DYNAMIC FEC
NO
RETURN
YES
176
MONITOR NETWORK FOR MAXIMUM OUTAGE DURATION AND ANY CHANGE IN THE LATENCY BUDGET
178
NEW OUTAGE DURATION DIFFERENT FROM PREVIOUSLY CONFIGURED Q VALUE AND/OR NEW LATENCY BUDGET DIFFERENT FROM PREVIOUSLY CONFIGURED X VALUE
NO
YES
180
ADJUST SOURCE AND/OR FEC ENCODERS ACCORDING TO MONITORED OUTAGE DURATION AND/OR NEW LATENCY BUDGET

FIG. 10