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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
6 May 2010 (06.05.2010)

PCT

(10) International Publication Number
WO 2010/051232 A1

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

(21) International Application Number:
PCT/US2009/061857

(22) International Filing Date:
23 October 2009 (23.10.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/262,498 31 October 2008 (31.10.2008) US

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

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **DHESIKAN, Subhasri** [US/US]; 1024 Cynthia Lane, San Jose, California 95129 (US). **POLK, James** [US/US]; 3913 Treemont Circle, Colleyville, Texas 76034 (US).

(74) Agent: **FRAME, Thomas J.**; Patent Capital Group, 6119 McCommas Blvd., Dallas, Texas 75214 (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, 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, PE, 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, SM,

[Continued on next page]

(54) Title: EXECUTING AND SUPPORTING A MULTIPLE BANDWIDTH RESERVATION REQUEST

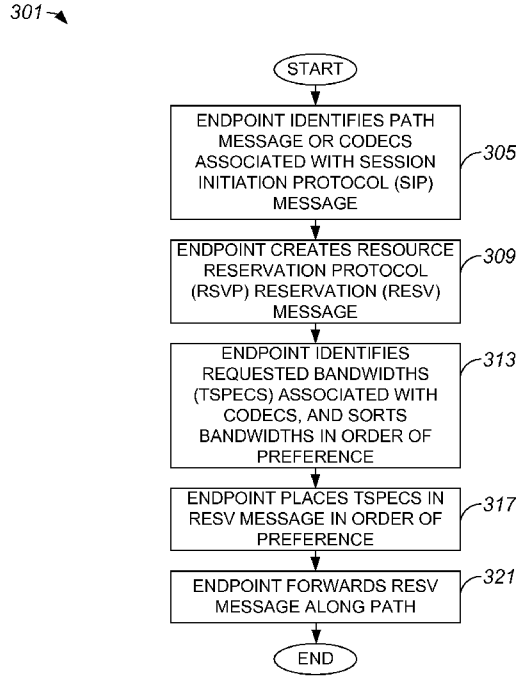

FIG. 3

(57) Abstract: In one embodiment, a method includes obtaining a first message that includes at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth. The first bandwidth is a preferential bandwidth. The method also includes determining whether the first bandwidth may be allocated, and storing the first bandwidth and the second bandwidth in a stored reservation state if the first bandwidth may be allocated. If the first bandwidth may not be allocated, the method includes determining whether the second bandwidth may be allocated. The second bandwidth in the stored reservation state if it is determined that the second bandwidth may be allocated. In one embodiment, if the second bandwidth may be allocated, the first bandwidth is removed during process prior to sending the message to a subsequent node upstream.

WO 2010/051232 A1

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

EXECUTING AND SUPPORTING A MULTIPLE BANDWIDTH RESERVATION REQUEST

BACKGROUND OF THE INVENTION

[0001] The present invention relates generally to communications systems.

[0002] The Resource Reservation Protocol (RSVP) is a transport layer protocol which allows resources to be reserved across a network. The resources are generally reserved for a flow. Typically, a sender of a request will send an RSVP, or Path, message downstream to a receiver. The receiver will send an RSVP reservation request (RESV) message upstream to a sender of the request in response to the Path message. Some hops in a path between the receiver and the sender is arranged to create and maintain a reservation state that is associated with the resources reserved for data packets which are to be sent from the sender.

[0003] Typically, a RESV message specifies a bandwidth request which effectively indicates the amount of resources to be reserved for a flow. Each hop in a path that receives the RESV message ascertains whether it has sufficient resources to accommodate the bandwidth request. If a hop is able to accommodate the bandwidth request, the hop effectively reserves the appropriate amount of bandwidth for the flow. If the hop is unable to accommodate the bandwidth request, the hop generates an RSVP reservation error (RESVERR) message and sends the RESVERR message to the endpoint that is the source of the RESV message.

[0004] When an endpoint that is the source of the RESV message receives the RESVERR message, the endpoint may send a new RESV message which specifies a request for a smaller amount of bandwidth. As will be appreciated by those skilled in the art, some applications are only able to use certain bandwidths which match with the codecs supported by the applications. Hence, a request for a smaller amount of bandwith is generally a request for a smaller amount of bandwidth which matches the appropriate codes. In general, multiple RESV

messages may be sent by an endpoint in an effort to reserve resources for a flow. By way of example, if a RESV message that requests eighty kilobytes per second (kbps) may not be accommodated, a subsequent RESV message that requests forty kbps may be sent. If the RESV message that request forty kbps also may not be accommodated, a subsequent RESV message that requests twenty kbps may be sent, and so forth. Typically, the amount of bandwidth within a RESV message substantially matches the amount of bandwidth within a chosen audio or video codec for a voice or video call, for example. An endpoint may support more than one codec, and may want to reserve the amount of bandwidth its preferred codec requires to operate within.

[0005] The sending and processing of multiple RESV messages may be relatively time consuming. As a result, it may sometimes not be possible for resources to be reserved for a flow in an efficient manner.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The invention will be readily understood by the following detailed description in conjunction with the accompanying drawings in which:

[0007] FIG. 1A is a block diagram representation of one system which may support Resource Reservation Protocol (RSVP) reservation (RESV) messages that specify multiple bandwidths at a time t1in which a PATH message is sent by an endpoint in accordance with an embodiment of the present invention.

[0008] FIG. 1B is a block diagram representation of one system which may support RESV messages, *e.g.*, system 100 of FIG, 1A, at a time t2 at which a RESV message that specifies multiple bandwidths is sent by an endpoint in accordance with an embodiment of the present invention.

[0009] FIG. 1C is a block diagram representation of one system which may support RESV messages, *e.g.*, system 100 of FIG, 1A, at a time t3 at which a RESV message that specifies

multiple bandwidths is processed by a router in accordance with an embodiment of the present invention.

[0010] FIG. 2 is a diagrammatic representation of a device which is capable of processing an overall request that includes multiple bandwidth requests in accordance with an embodiment of the present invention.

[0011] FIG. 3 is a process flow diagram which illustrates a method of specifying multiple acceptable bandwidths in a RESV request arranged to reserve a bandwidth in accordance with an embodiment of the present invention.

[0012] FIGS. 4A and 4B are a process flow diagram which illustrates a method of reserving a bandwidth selected from multiple acceptable bandwidths specified in an RESV message received at a node or a hop along a path in accordance with an embodiment of the present invention.

[0013] FIG. 5 is a process flow diagram which illustrates a method of processing a received RESV message at an end host, *e.g.*, an upstream end host, in accordance with an embodiment of the present invention.

[0014] FIG. 6 is a process flow diagram which illustrates a method of downgrading a bandwidth allocation at a node along a path in accordance with an embodiment of the present invention.

[0015] FIG. 7 is a process flow diagram which illustrates a method of downgrading a bandwidth allocation at an end host, *e.g.*, a downstream end host, in response to a received RESVERR message in accordance with an embodiment of the present invention.

[0016] FIG. 8 is a process flow diagram which illustrates a method of responding to a high priority request when substantially all bandwidth is in use in accordance with an embodiment of the present invention.

DESCRIPTION OF EXAMPLE EMBODIMENTS

GENERAL OVERVIEW

[0017] According to one aspect of the present invention, a method includes obtaining a first message that includes at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth. The first bandwidth is a preferential bandwidth. The method also includes determining whether the first bandwidth may be allocated, and storing the first bandwidth and the second bandwidth in a stored reservation state if the first bandwidth may be allocated. If the first bandwidth may not be allocated, the method includes determining whether the second bandwidth may be allocated. The second bandwidth in the stored reservation state if it is determined that the second bandwidth may be allocated.

DESCRIPTION

[0018] In some situations, multiple Resource Reservation Protocol (RSVP) reservation (RESV) request messages are sent along a path, or an accumulation of routers, between endpoints before a bandwidth allocation may be reserved with respect to that path for a desired flow. Sending multiple RESV requests messages may often be inefficient. By including multiple bandwidth requests within the same reservation request in a single RSVP message, a bandwidth allocation may be reserved relatively efficiently. If an RESV request message carries more than one bandwidth request, if a relatively high bandwidth request contained in the RESV request message may not be accommodated, it may be possible to accommodate a lower bandwidth request contained in the RESV request message. Hence, a lower than desired bandwidth may be allocated for a particular flow in response to a single RESV request message, *i.e.*, the same RESV request message used to request a desired bandwidth.

[0019] Allowing RESV request messages to carry multiple traffic specification (TSPEC) objects enables a single RESV request message to effectively provide multiple acceptable

bandwidths allocations for a desired flow. For example, a single RESV request message may specify first, second, and third bandwidth requests such that if a first bandwidth request may not be accommodated, the second bandwidth is considered and, if necessary, the third bandwidth is considered. It should be appreciated that more than three bandwidth requests may be specified. Further, when a single RESV request message contains more than one TSPEC object, a node along a path which receives the RESV request message may maintain more than one potential reservation state. Hence, if a first potential reservation state is associated with a bandwidth that may no longer be accommodated, the node may switch to a different reservation state, *e.g.*, a reservation state associated with a lower bandwidth..

[0020] In one embodiment, an application within a requesting node or endpoint may effectively ask for a desired bandwidth allocation and at least one backup bandwidth allocation, *e.g.*, a bandwidth allocation that would be acceptable in the absence of the desired bandwidth allocation. The desired bandwidth allocation and at least one backup bandwidth allocation may be provided in TSPEC objects contained in an RSVP request message such as an RESV request message. If the accumulation of routers that the RESV request message passes through, *i.e.*, a network, is able to accommodate the desired bandwidth allocation, the desired bandwidth may be reserved. Otherwise, if a backup bandwidth allocation may be accommodated, the backup bandwidth allocation may be reserved.

[0021] Referring initially to FIG. 1A, one system which may support RSVP RESV request messages that specify multiple bandwidths will be described at a time t1 in accordance with an embodiment of the present invention. A system 100 includes an upstream endpoint host 104, a downstream endpoint host 108, and an accumulation of routers 112a-c or, more generally, intermediate hops or nodes. Endpoint 104 may be a server or a server application, and may be arranged to send a PATH message downstream to endpoint 108 via routers 112a-c at a time t1.

[0022] Upon receiving a PATH message from endpoint 104, endpoint 108 may generate an RESV request message and send the RESV request message upstream through routers 112a-c at

a time t2, as shown in FIG. 1B. The RESV request message may include two or more TSPEC objects which specify bandwidths that are requested by endpoint 108. TSPEC objects may be ordered within the RESV request such that the TSPEC object which specifies a preferred bandwidth is effectively ordered in a first position. That is, TSPEC objects associated with the RESV request message may be arranged in an order of preference with the preferential bandwidth being the highest bandwidth requested. In one embodiment, TSPEC objects may be arranged in an order of preference similar to how a session description protocol (SDP) identifies an order of preference associated with codecs in a media description line.

[0023] At least one router 112a-c includes functionality which enables multiple reservation states to be identified and stored. That is, one or more of routers 112a-c is configured to process an RESV request message that includes a plurality of TSPEC objects. For ease of discussion, each router 112a-c will be described as being arranged to process an RESV request message that includes a plurality of TSPEC objects. It should be appreciated, however, that network 100 may generally support the use of an RESV request message that includes a plurality of TSPEC objects as long as at least one router 112a-c is arranged to process such an RESV request message. Processing an RESV request message that includes a plurality of TSPEC objects on a router 112a-c may include, but is not limited to including, identifying a TSPEC object with the highest bandwidth that may be supported by the router 112a-c, storing the highest bandwidth that may be supported, and storing substantially all bandwidths identified in TSPEC objects that are lower than the highest bandwidth that may be supported. Such bandwidths may be stored on a routers 112a-c as reservation states.

[0024] When endpoint 108 sends an RESV request message that includes a plurality of TSPEC objects upstream to endpoint 104, router 112c is the first hop in the overall path to endpoint 104. FIG. 1C shows system 100 at a time t3 at which a RESV request message that includes a plurality of TSPEC objects is processed by router 112c in accordance with an embodiment of the present invention. It should be appreciated that each router 112a-c may process the RESV request message if the RESV request message is to successfully reach

endpoint 104. At a time t3, when router 112c receives an RESV request message that includes a plurality of TSPEC objects, *e.g.*, specifies multiple requested bandwidths, router 112c determines which of the requested bandwidths it may support or otherwise accommodate.

[0025] Typically, router 112c identifies the highest bandwidth associated with the plurality of TSPEC objects that it may support. Then, router 112c sends an indication of a supportable bandwidth, *i.e.*, the highest bandwidth associated with the plurality of TSPEC objects that router 112c is able to support. It should be appreciated that the highest bandwidth that router 112c is able to support may be lower than the preferential bandwidth specified by the TSPEC objects. If the supportable bandwidth is the preferential bandwidth, or the highest bandwidth requested, then router 112c propagates the RESV request message to the next hop, or router 112b. If the supportable bandwidth is not the preferential bandwidth, a modified RESV request message which provides an indication of the supportable bandwidth is propagated upstream to router 112b in a modified RESV request message, while a message such as an RSVP reservation error (RESVERR) message which provides an indication of the supportable bandwidth is propagated downstream to endpoint 108. In one embodiment, a modified RESV request message that is propagated upstream by router 112c may be modified such that TSPEC objects corresponding to bandwidths which are higher than the supportable bandwidth are effectively removed. Generally, if no bandwidth is supportable, *e.g.*, if router 112c has no available bandwidth to allocate or to reserve, router 112c may send an appropriate indication downstream and/or upstream.

[0026] A supportable bandwidth may be stored by router 112c as a reservation state. In addition to storing the supportable bandwidth, any bandwidths lower than the supportable bandwidth may also be stored by router 112a as a reservation state. That is, router 112c may store the highest bandwidth indicated in the TSPEC objects of a RESV request that router 112c is capable of supporting, as well as any lower bandwidths indicated in the TSPEC objects.

[0027] As previously mentioned, at least one of routers 112a-c has the capability to process an overall request which includes more than one bandwidth request, *e.g.*, a RESV request message which is associated with more than one TSPEC object. With reference to FIG. 2, a router or, more generally, a device which is capable of processing an overall request that includes multiple bandwidth requests in accordance with an embodiment of the present invention. A routing device or intermediate node 212 is generally arranged to process an incoming message, as for example a request message 232, and propagate an outgoing message, as for example a request message 232'. In general, node 212 includes hardware and/or software logic that is embodied on a tangible medium. The logic is typically executable and, when executed, is operable to process an incoming message and to propagate an outgoing message. The incoming message is not limited to being a request message such as request message 232. By way of example, node 212 may process incoming messages which are error messages such as error messages.

[0028] Node 212 includes an input/output arrangement 216 which may include ports that allow node 212 to receive incoming messages such as request message 232 and to send outgoing messages such as request message 232'. Input/output arrangement 216 enables node 212 to be in communication with a network. Node 212 also includes routing functionality 228 that allows node 212 to route messages to appropriate devices, *e.g.*, next hops.

[0029] Bandwidth allocation functionality 224 allows node 212 to allocate bandwidth or, more generally, to allocate resources for calls that are routed through node 212. Bandwidth allocation functionality 224 may identify bandwidth requests specified in request message 232, and, using multiple bandwidth processing functionality 226, may process the bandwidth requests to determine the highest or most preferred bandwidth request that node 212 may support. Multiple bandwidth processing functionality 226 is configured to process request messages 232 which include a plurality of TSPEC objects 236a-c. By way of example, if node 212 has sufficient bandwidth available to reserve 100 kbps of bandwidth in response to request message 232, multiple bandwidth processing functionality 226 may store the 100 kbps

corresponding to TSPEC object 236a as a reservation state in a storage arrangement 220 such that the 100 kbps is effectively identified as the bandwidth allocation corresponding to request message 232. Additionally multiple bandwidth processing functionality 226 may also store bandwidths corresponding to TSPEC objects 236b, 236c in a reservation state of storage arrangement 220. Multiple bandwidth processing functionality 226 is also configured to down-size or otherwise lower a bandwidth allocation when necessary, *e.g.*, when notified that other nodes in a path that includes node 212 are substantially only able to accommodate a down-sized bandwidth allocation.

[0030] In one embodiment, request message 232' may be substantially the same as request message 232. That is, node 212 may be arranged to effectively forward request message 232 within a network. It should be appreciated, however, that request message 232' may instead be a modified version of request message 232.

[0031] Request message 232 includes multiple TSPEC objects 236a-c that are each arranged to specify a desired bandwidth. As shown, for purposes of illustration, TSPEC object 236a specifies a bandwidth of approximately 100 kilobytes per second (kbps), TSPEC object 236b specifies a bandwidth of approximately 80 kbps, and TSPEC object 236c specifies a bandwidth of approximately 8 kbps. It should be appreciated that the number of TSPEC objects 236a-c, as well as the bandwidths associated with TSPEC objects 236a-c, may vary widely. Within request 232, TSPEC objects 236a-c may be ordered such that TSPEC object 236a, which corresponds to a highest desired bandwidth, is a preferred or preferential bandwidth whereas TSPEC object 236c, which corresponds to a lowest desired bandwidth, is a least preferred bandwidth.

[0032] When node 212 processes request 232 using bandwidth allocation functionality 224, node 212 determines which bandwidth specified in request message 232 may be accommodated, *e.g.*, stored as reservations states in storage 220 and forwarded in request 232'. Request 232' effectively includes a TSPEC object which specifies a highest possible bandwidth 240 which node 212 may support or accommodate. Request 232' optionally

includes a TSPEC object or objects which specify lower possible bandwidths 240 which node 212 may support or accommodate, *i.e.*, lower possible bandwidths 240 may be specified when they exist. Depending upon the bandwidth availability associated with node 212, any of TSPEC objects 236a-c may be associated with highest possible bandwidth 240. For example, if bandwidth allocation functionality 224 processes request 232 and determines that the preferential bandwidth associated with TSPEC object 236a may not be supported, then highest possible bandwidth 240 may be the bandwidth associated with TSPEC object 236b, namely approximately 80 kbps. As such, lower bandwidth 244 is the bandwidth associated with TSPEC object 236b, namely 8 kbps.

[0033] FIG. 3 is a process flow diagram which illustrates a method of specifying multiple acceptable bandwidths in a RESV message arranged to reserve a bandwidth in accordance with an embodiment of the present invention. A process 301 of specifying multiple acceptable bandwidths in a RESV message begins at step 305 in which an endpoint, *e.g.*, a downstream endpoint such as endpoint 108 of FIG. 1A, identifies a PATH message or codecs associated with a received session initiation protocol (SIP) message. The identified PATH message, which is an RSVP message as will be appreciated by those skilled in the art, may be received from an upstream endpoint such as endpoint 104 of FIG. 1A.

[0034] In step 309, the endpoint creates a RESV message. Then, in step 313, the endpoint identifies requested bandwidths, *e.g.*, TSPEC objects, associated with codes, and sorts the bandwidths in order of preference. That is, the endpoint may order TSPEC objects in an order that enables the TSPEC object corresponding to a preferential bandwidth to be identified as being preferred. Typically, the order of preference may be such that the TSPEC object associated with a highest desired bandwidth has the highest preference whereas the TSPEC object associated with a lowest acceptable bandwidth has the lowest preference.

[0035] Once the bandwidths are sorted in order of preference, the endpoint places the TSPEC objects in the RESV message in step 317. Typically, the TSPEC objects are specified in the

RESV message in the order of preference. After the TSPEC objects are placed in the RESV message, the endpoint forwards the RESV message along a path in step 321. That is, the endpoint sends the RESV message to the sender of the PATH message on the path that was traversed by the PATH message. Forwarding the RESV message along the path includes providing the RESV message to the next hop in the path. The process of specifying multiple acceptable bandwidths in a RESV message.

[0036] When a hop receives an RESV message that includes a plurality of TSPEC objects and, hence, includes a plurality of bandwidths, the hop generally processes the RESV message to identify the highest bandwidth that the hop is capable of supporting. FIGS. 4A and 4B are a process flow diagram which illustrates a method of reserving a bandwidth selected from multiple acceptable bandwidths specified in an RESV message received at a node or a hop along a path in accordance with an embodiment of the present invention. A process 401 of reserving a bandwidth on a node begins at step 405 in which a node receives an RESV message that includes a plurality of TSPEC objects or, more generally, a plurality of bandwidths. The RESV message may be received directly from an endpoint, or may be received indirectly from an endpoint via another hop in a path that includes the endpoint and the node. The node is typically a hop in the path that the RESV message is traversing.

[0037] After the RESV message is received by the node, the node processes the plurality of TSPEC objects contained in the RESV message in step 409. That is, the node processes the plurality of bandwidths associated with the RESV message. The TSPEC objects or bandwidths are typically processed to identify the highest bandwidth that the node may accommodate.

[0038] A determination is made in step 413 as to whether the node has identified the highest priority TSPEC object or bandwidth that it is able to accommodate. If it is determined that the node has not identified a highest priority TSPEC object that it is able to accommodate, the indication is that the node is unable to support any of the bandwidths specified in the RESV message. As such, process flow moves to step 417 in which the node generates a RESVERR

message which indicates that the node is unable to allocate any bandwidth in response to the RESV message, and sends the RESVERR message downstream. Then, in step 421, the node prevents the RESV message from being forwarded, and the process of reserving a bandwidth on a node is completed.

[0039] Returning to step 413, if the determination is that the node has identified a highest priority TSPEC object or bandwidth that it may accommodate, then the implication is that the node may be able to accommodate the preferential bandwidth specified in the RESV message. As such, a determination is made in step 425 regarding whether the highest priority TSPEC object or bandwidth may be accommodated. In other words, it is determined whether the preferential bandwidth, *i.e.*, the highest priority bandwidth specified in the RESV message, may be accommodated.

[0040] If it is determined in step 425 that the highest priority TSPEC object may be accommodated, then the node forwards the RESV message "as is" upstream in step 429. Then, in step 433, the node adds substantially all bandwidths associated with the RESV message into its stored reservation state, and the process of reserving a bandwidth on a node is completed. When the node adds the bandwidths into its stored reservation state, the preferential bandwidth may effectively be identified as the bandwidth that is currently accommodated, while the other bandwidths may be identified substantially as "alternatives" in the event that the currently accommodated bandwidth is subsequently no longer accommodatable.

[0041] Alternatively, if the determination is that the highest priority TSPEC object may not be accommodated, then process flow moves to step 437 in which the node removes TSPEC objects from the RESV message that correspond to bandwidth requests which may not be accommodated. That is, the node removes requested bandwidths from the RESV message that are higher than the highest priority TSPEC object that may be accommodated.

[0042] Once TSPEC objects are removed from the RESV message, the node adds substantially all bandwidths associated with accommodatable TSPEC objects to its stored

reservation state. The accommodatable TSPEC objects or bandwidths include the highest priority bandwidth that may be accommodated and substantially all bandwidths with a lower priority than the highest priority bandwidth that may be accommodated.

[0043] After the node adds bandwidths to its stored reservation state, the node forwards the RESV message, from which TSPEC objects corresponding to bandwidths which the node is unable to support have been removed, upstream in step 445. Then, in step 449, the node generates a RESVERR message which indicates that the highest priority bandwidth in the RESV message, as received in step 405, may not be accommodated. The node also sends the RESVERR message downstream. It should be appreciated that the RESVERR message may identify the highest priority bandwidth that may be accommodated by the node. Upon generating and sending the RESERR message downstream, the process of reserving a bandwidth on a node is completed.

[0044] When an endpoint, *e.g.*, an upstream endpoint, receives an RESV message, the endpoint may process the RESV message to determine the bandwidth that is supported along a path through which the RESV message was received. FIG. 5 is a process flow diagram which illustrates a method of processing a received RESV message at an endpoint host, *e.g.*, an upstream endpoint, in accordance with an embodiment of the present invention. A method 501 of processing an RESV message received at an endpoint begins at step 505 in which an endpoint receives an RESV message which includes one or more TSPEC objects. After receiving the RESV message, the endpoint identifies the highest priority TSPEC object and, hence, highest priority bandwidth, contained in the RESV message in step 509. The highest priority bandwidth contained in the RESV message is generally indicative of the bandwidth that is allocated or allotted along the path through which the RESV message is received.

[0045] In step 513, the endpoint notes or otherwise records the highest allocated bandwidth. By noting the highest allocated bandwidth, the endpoint may send messages on the path that do not exceed the highest allocated bandwidth. Once the highest allocated

bandwidth is noted or otherwise recorded, the method of processing an RESV message received at an endpoint is completed.

[0046] When a node along a path receives a RESVERR message, that node is generally downsizes or downgrades its current bandwidth allocation to a bandwidth consistent with the bandwidth specified in the RESVERR message. It should be appreciated that a RESVERR message is just one example of a message that may effectively provide a reduced bandwidth specification. FIG. 6 is a process flow diagram which illustrates a method of downgrading a bandwidth allocation at a node along a path in response to receiving a reduced bandwidth specification in accordance with an embodiment of the present invention. A process 601 of downgrading a bandwidth allocation begins at step 605 in which a node, *e.g.*, a hop along a path, receives a RESVERR message that includes a reduced bandwidth specification. The node has previously stored a reservation state which includes at least one bandwidth allocation, *i.e.*, the node has previously stored at least a currently allotted bandwidth. In addition to the currently allotted bandwidth, or allotted reservation, the node may also have stored additional reservation states, or bandwidths to which the currently allotted bandwidth may be downsized to.

[0047] A determination is made in step 609 as to whether the reduced bandwidth specification specifies no bandwidth. In other words, it is determined if the RESVERR message indicates that there is no available bandwidth upstream from the node. If it is determined that the reduced bandwidth specification specifies no bandwidth, the node adjusts the allotted bandwidth in step 613 to indicate that substantially no bandwidth may be supported. The node also substantially removes the allotted reservation and any previous reservations states. Once the node adjusts the allotted reservation, the node forwards the RESVERR message downstream in step 617, and the process of downgrading a bandwidth allocation is completed.

[0048] Returning to step 609, if the determination is that the reduced bandwidth specification specifies a bandwidth, then process flow moves to step 621 in which the node

adjusts the allotted reservation, or the current bandwidth allocation, the specified reduced bandwidth. Then, in step 625, the node removes substantially all bandwidths that are higher than the specified reduced bandwidth from its reservation state. Removing substantially all bandwidths that are higher than the specified reduced bandwidth effective downgrades the bandwidth allocation to the specified reduced bandwidths. After the node removes substantially all bandwidths that are higher than the specified reduced bandwidth from its reservation state, the node forwards the RESVERR message downstream in step 629, and the process of downgrading a bandwidth allocation is completed.

[0049] Eventually, a RESVERR message sent in response to an RESV message initiated by an endpoint or end host will receive the RESVERR message. Referring next to FIG. 7, a process of downgrading a bandwidth allocation at an endpoint, *e.g.*, a downstream endpoint, in response to a received RESVERR message will be described in accordance with an embodiment of the present invention. A process 701 of downgrading a bandwidth allocation at an endpoint begins at step 705 in which an endpoint receives a RESVERR message that includes a reduced bandwidth specification. The reduced bandwidth specification effectively specifies the bandwidth which has been allotted along the path on which the RESV message was forwarded. Upon receiving the RESVERR message, the endpoint typically notes or otherwise records the allotted bandwidth in step 709.

[0050] In step 713, the endpoint makes adjustments to reflect the reduced bandwidth specification. Such adjustments may include effectively ensuring that messages forwarded by the endpoint along the path will utilize at most the reduced bandwidth. The process of downgrading a bandwidth allocation at an endpoint is completed once the endpoint makes adjustments to maintain consistency with the reduced bandwidth specification.

[0051] The inclusion of multiple bandwidth requests within the same reservation request in a single RSVP message, *e.g.*, a single RESV request message, allows applications to efficiently reserve substantially the exact bandwidth that may be used in a call. With multiple bandwidth

request specified in a single message, a post-dial delay budget is generally not overrun. As a result, an application may efficiently select and/or change the codec of a call based on the reservation information up to the maximum that may be supported, as indicated by the single RSVP message.

[0052] As discussed above, a single RSVP message which includes multiple bandwidth requests within the same reservation allows multiple bandwidth request to be stored as reservation states on a node. Hence, if a relatively high priority call is received on a node which is experiencing congestion on an interface, the node may downsize or downgrade the bandwidth allocated in response to the single RSVP message such that the relatively high priority call may be accommodated substantially without discontinuing a lower priority call. FIG. 8 is a process flow diagram which illustrates a method of responding to a high priority call or request when substantially all bandwidth is in use on a node in accordance with an embodiment of the present invention. A process 801 of responding to a high priority call begins at step 805 in which a node, *e.g.*, a router, which is congested receives a high priority call. A congested node generally has substantially no available bandwidth or insufficient bandwidth to accommodate the high priority call.

[0053] After the node receives a high priority call that it is effectively incapable of accommodating, it is determined in step 809 whether there are one or more lower priority calls which are utilizing bandwidth on the node. If it is determined that there are no lower priority calls, the indication is that substantially all calls which are utilizing bandwidth on the node have a high priority. As such, in step 829, notification is provided, *e.g.*, to the initiator of the high priority call, that the high priority call may not be accommodated. Once the notification is provided, the process of responding to a high priority call is completed.

[0054] Alternatively, if it is determined in step 809 that there are one or more lower priority calls utilizing bandwidth on the node, then process flow moves to step 813 in which lower priority calls with bandwidth reservations that may be downsized are identified. That is, lower

priority calls with reservation states that specify preferred bandwidths and lower bandwidths are identified.

[0055] A determination is made in step 817 as to whether there are any lower priority calls with bandwidth reservations that may be downsized. If the determination is that there are lower priority calls with bandwidth reservations that may be downsized, then the bandwidth reservations of one or more lower priority calls are downsized in step 821. In one embodiment, if downsizing one lower priority call is sufficient to accommodate the high priority call, then substantially only one lower priority call will be downsized. It should be appreciated, however, that the policies and algorithms used to determine how to downsize reservations may vary widely. By way of example, in lieu of specifying that substantially only a certain number of lower priority calls may have their bandwidth reservations downsized, a policy may specify that the bandwidth reservations of the oldest lower priority calls are downsized, or a policy may specify that the bandwidth reservations of substantially all lower priority calls are downsized.

[0056] After the bandwidth reservations of one or more of the lower priority calls are downsized in step 821, the high priority call may be accommodated, *e.g.*, put through or otherwise connected, in step 825. The process of responding to a high priority call is completed once the high priority call is accommodated.

[0057] Returning to step 817, if the determination is that none of the lower priority calls have bandwidth reservations that may be downsized, then at least one lower priority call is preempted or discontinued in step 833 such that the high priority call may be accommodated. After at least one lower priority call is preempted, the process of responding to a high priority call is completed.

[0058] Although only a few embodiments of the present invention have been described, it should be understood that the present invention may be embodied in many other specific forms without departing from the spirit or the scope of the present invention. By way of example, an application such as the Cisco Unified Communications Manager (available

commercially from Cisco Systems, Inc. of San Jose, California) may effectively ask an RSVP agent to book a highest preferred bandwidth as well as other bandwidths that may be used in the even that the highest preferred bandwidth is not available. The highest preferred bandwidth request may be accommodated, *e.g.*, if endpoints are in substantially the same location, or a reduced bandwidth may be accepted by the application if the highest preferred bandwidth request may not be accommodated. Depending upon the bandwidth allotted, the application may identify an appropriate codec to invoke.

[0059] While multiple bandwidths have been described as being specified in a RESV message through the use of TSPECs or TSPEC objects, the bandwidths are not limited to being specified using TSPECs or TSPEC objects. For instance, multiple bandwidths may be specified in service-request specification (RSPEC) objects associated RESV messages.

[0060] In one embodiment, if none of the bandwidths specified in a RESV message may be supported, no bandwidth may be allocated in response to the RESV message. It should be appreciated, however, that some bandwidth may instead be allocated even if none of the specified bandwidths may be supported. For example, if a RESV message specifies a lowest preferred bandwidth, if the lowest preferred bandwidth may not be accommodated, rather than allocating substantially no bandwidth in response to the RESV message, an amount of bandwidth that is less than the lowest preferred bandwidth may be allocated.

[0061] The steps associated with the methods of the present invention may vary widely. Steps may be added, removed, altered, combined, and reordered without departing from the spirit of the scope of the present invention. By way of example, when a high priority request is received on a congested router that has no available bandwidth and no reservations for lower priority calls, in lieu of not accommodating the high priority request, the router may instead downsize the bandwidths of calls with substantially the same bandwidth as the request such that the request may be accommodated. Therefore, the present examples are to be considered

as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope of the appended claims.

WHAT IS CLAIMED IS:

1. A method comprising:

obtaining a first message, the first message including a plurality of bandwidth requests, wherein the plurality of bandwidth requests includes at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth, the first bandwidth being a preferential bandwidth;

determining whether the first bandwidth may be allocated;

storing the first bandwidth and the second bandwidth in a stored reservation state if it is determined that the first bandwidth may be allocated;

determining whether the second bandwidth may be allocated if it is determined that the first bandwidth may not be allocated; and

storing the second bandwidth in the stored reservation state if it is determined that the second bandwidth may be allocated.

2. The method of claim 1 wherein the first message is a resource reservation protocol (RSVP) message associated with a reservation request.

3. The method of claim 2 wherein the first message is an RSVP reservation (RESV) request message.

4. The method of claim 3 wherein the first bandwidth request is included in a first traffic specification (TSPEC) object and the second bandwidth request is included in a second TSPEC object.

5. The method of claim 1 further including:

removing the first bandwidth request from the first message if it is determined that the first bandwidth may not be allocated; and

forwarding the first message after removing the first bandwidth request.

6. The method of claim 1 wherein if it is determined that the first bandwidth may be allocated, the method further including:

obtaining a second message, the second message being arranged to indicate that the first bandwidth may no longer be allocated; and

removing the first bandwidth from the stored reservation state.

7. The method of claim 1 wherein if it is determined that the second bandwidth may not be allocated, the method further includes:

generating an error message, the error message being arranged to indicate that none of the plurality of bandwidth request may be accommodated.

8. An apparatus comprising:

means for obtaining a first message, the first message including a plurality of bandwidth requests, wherein the plurality of bandwidth requests includes at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth, the first bandwidth being a preferential bandwidth;

means for determining whether the first bandwidth may be allocated;

means for storing the first bandwidth and the second bandwidth in a stored reservation state if it is determined that the first bandwidth may be allocated;

means for determining whether the second bandwidth may be allocated if it is determined that the first bandwidth may not be allocated; and

means for storing the second bandwidth in the stored reservation state if it is determined that the second bandwidth may be allocated.

9. The apparatus of claim 8 wherein the first message is a resource reservation protocol (RSVP) message associated with a reservation request.

10. The apparatus of claim 9 wherein the first message is an RSVP reservation (RESV) request message.

11. The apparatus of claim 10 wherein the first bandwidth request is included in a first traffic specification (TSPEC) object and the second bandwidth request is included in a second TSPEC object.

12. Logic encoded in one or more tangible media for execution and when executed operable to:

obtain a first message, the first message including a plurality of bandwidth requests, wherein the plurality of bandwidth requests includes at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth, the first bandwidth being a preferential bandwidth;

determine whether the first bandwidth may be allocated;

store the first bandwidth and the second bandwidth in a stored reservation state if it is determined that the first bandwidth may be allocated;

determine whether the second bandwidth may be allocated if it is determined that the first bandwidth may not be allocated; and

store the second bandwidth in the stored reservation state if it is determined that the second bandwidth may be allocated.

13. The logic of claim 12 wherein the first message is a resource reservation protocol (RSVP) message associated with a reservation request.

14. The logic of claim 13 wherein the first message is an RSVP reservation (RESV) request message.

15. The logic of claim 14 wherein the first bandwidth request is included in a first traffic specification (TSPEC) object and the second bandwidth request is included in a second TSPEC object.

16. The logic of claim 12 further operable to:

remove the first bandwidth request from the first message if it is determined that the first bandwidth may not be allocated; and

forward the first message after removing the first bandwidth request.

17. The logic of claim 12 wherein if it is determined that the first bandwidth may be allocated, the logic is further operable to:

obtain a second message, the second message being arranged to indicate that the first bandwidth may no longer be allocated; and

remove the first bandwidth from the stored reservation state.

18. The logic of claim 12 wherein if it is determined that the second bandwidth may not be allocated, the logic is further operable to:

generate an error message, the error message being arranged to indicate that none of the plurality of bandwidth request may be accommodated.

19. An apparatus comprising:

an input/output arrangement, the input/output arrangement being arranged to obtain a first message that includes a plurality of bandwidth requests, the plurality of bandwidth requests including at least a first bandwidth request that specifies a first bandwidth and a second bandwidth request that specifies a second bandwidth, the first bandwidth being a preferential bandwidth;

a storage arrangement; and

a bandwidth processing arrangement, the bandwidth processing arrangement being configured to determine whether the first bandwidth may be allocated, the bandwidth processing arrangement further being configured to store the first bandwidth and the second bandwidth into the storage arrangement as a stored reservation state if it is determined that the first bandwidth may be allocated.

20. The apparatus of claim 20 wherein the bandwidth processing arrangement is further arranged to determine whether the second bandwidth may be allocated if it is determined that the first bandwidth may not be allocated, and to store the second bandwidth into the storage arrangement as the stored reservation state if it is determined that the second bandwidth may be allocated.

21. The apparatus of claim 20 wherein the bandwidth processing arrangement is still further arranged to remove the first bandwidth request from the first message if it is determined that the first bandwidth may not be allocated, and wherein the input/output arrangement is configured to forward the first message after the first bandwidth request is removed.

22. The apparatus of claim 20 wherein the apparatus is a routing device, and wherein the apparatus further includes a routing arrangement.

23. The apparatus of claim 19 wherein the input/output arrangement is configured to obtain a second message if it is determined that the first bandwidth may be allocated, the second message being arranged to indicate that the first bandwidth may no longer be allocated, and wherein the bandwidth processing arrangement is further arranged to remove the first bandwidth from the stored reservation state.

24. The apparatus of claim 19 wherein the first message is a resource reservation protocol (RSVP) message associated with a reservation request.

25. The apparatus of claim 24 wherein the first message is an RSVP reservation (RESV) request message.

26. The apparatus of claim 25 wherein the first bandwidth request is included in a first traffic specification (TSPEC) object and the second bandwidth request is included in a second TSPEC object.

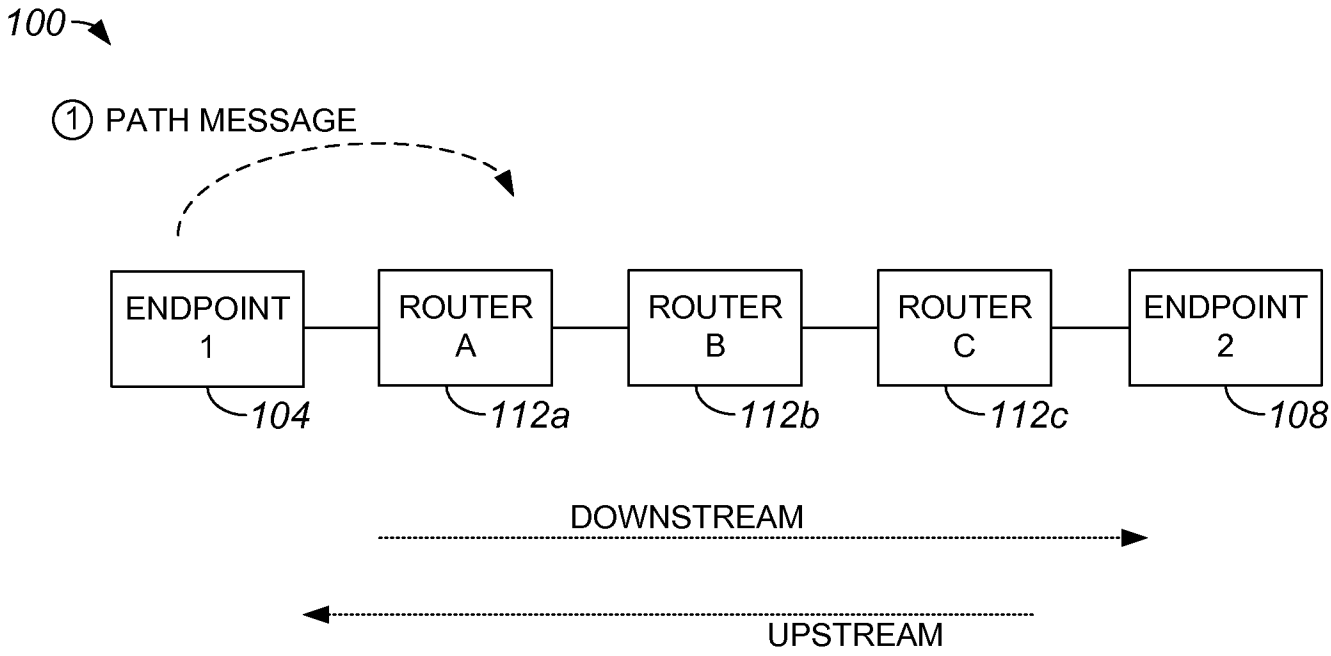
100
1 PATH MESSAGE
ENDPOINT 1
ROUTER A
ROUTER B
ROUTER C
ENDPOINT 2
104
112a
112b
112c
108
DOWNSTREAM
UPSTREAM

FIG. 1A

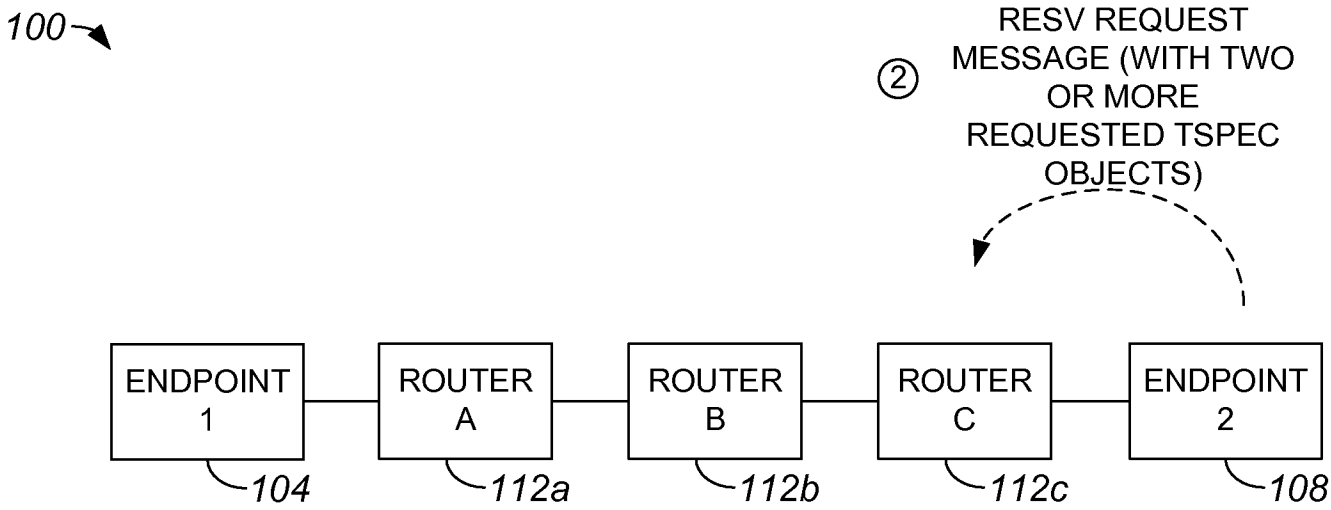
100
2 RESV REQUEST MESSAGE (WITH TWO OR MORE REQUESTED TSPEC OBJECTS)
ENDPOINT 1
ROUTER A
ROUTER B
ROUTER C
ENDPOINT 2
104
112a
112b
112c
108

FIG. 1B

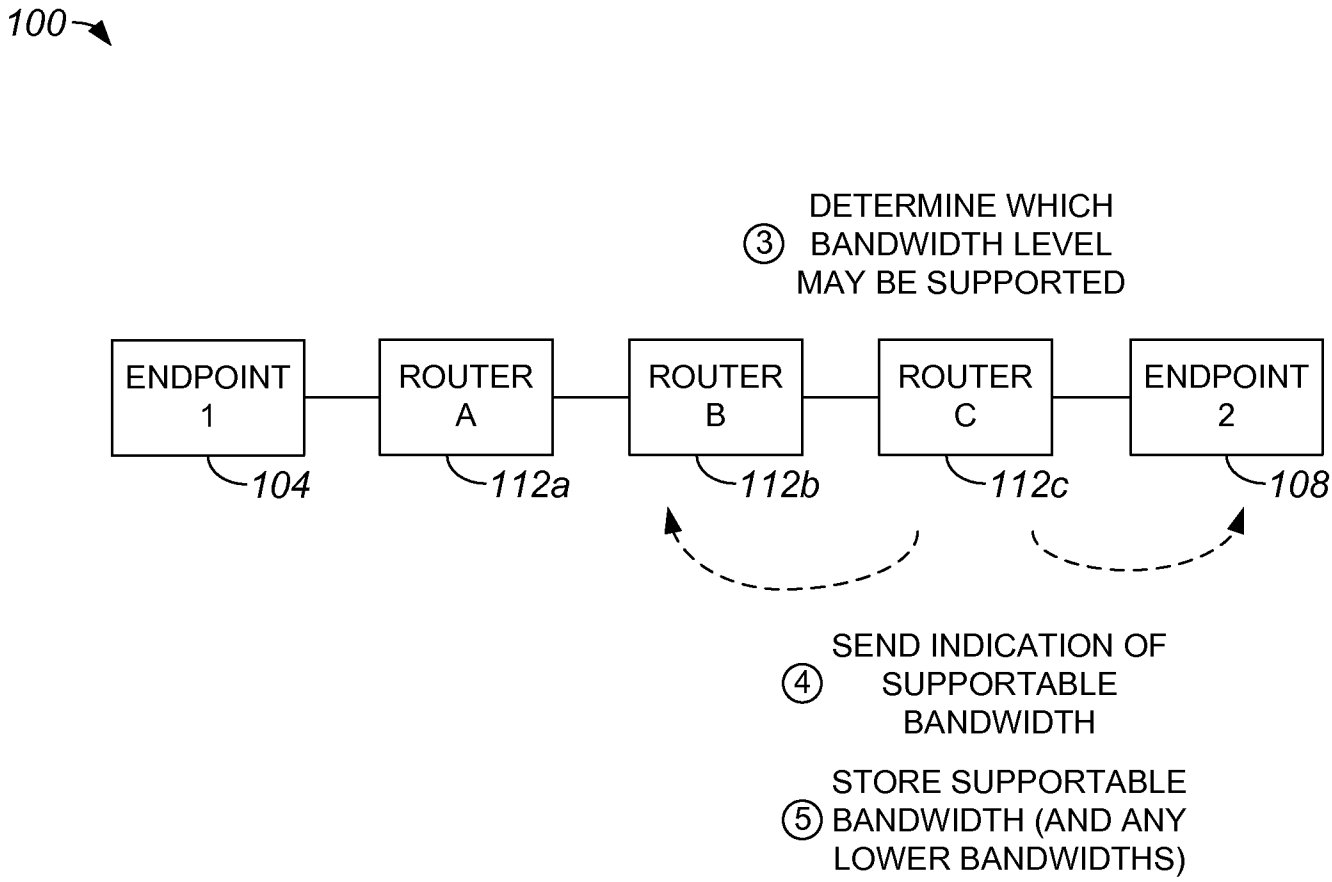
100
③ DETERMINE WHICH BANDWIDTH LEVEL MAY BE SUPPORTED
ENDPOINT 1
ROUTER A
ROUTER B
ROUTER C
ENDPOINT 2
104
112a
112b
112c
108
④ SEND INDICATION OF SUPPORTABLE BANDWIDTH
⑤ STORE SUPPORTABLE BANDWIDTH (AND ANY LOWER BANDWIDTHS)

FIG. 1C

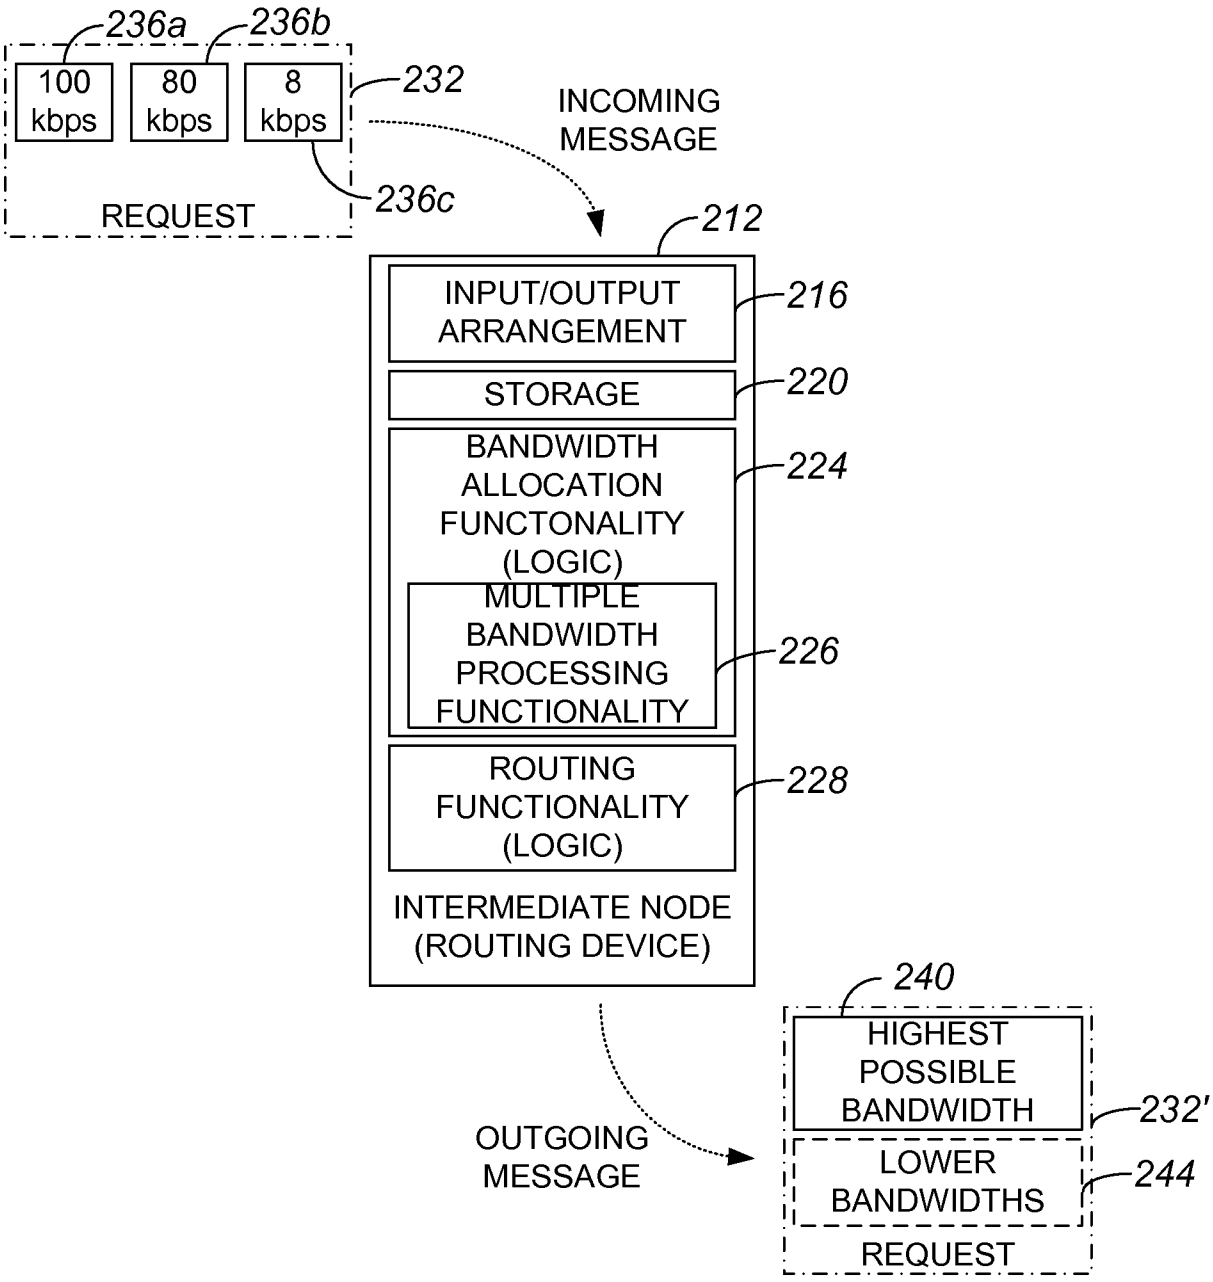
236a
236b
100 kbps
80 kbps
8 kbps
232
REQUEST
236c
INCOMING MESSAGE
212
INPUT/OUTPUT ARRANGEMENT
216
STORAGE
220
BANDWIDTH ALLOCATION FUNCTONALITY (LOGIC)
224
MULTIPLE BANDWIDTH PROCESSING FUNCTIONALITY
226
ROUTING FUNCTIONALITY (LOGIC)
228
INTERMEDIATE NODE (ROUTING DEVICE)
240
HIGHEST POSSIBLE BANDWIDTH
232'
LOWER BANDWIDTHS
244
REQUEST
OUTGOING MESSAGE

FIG. 2

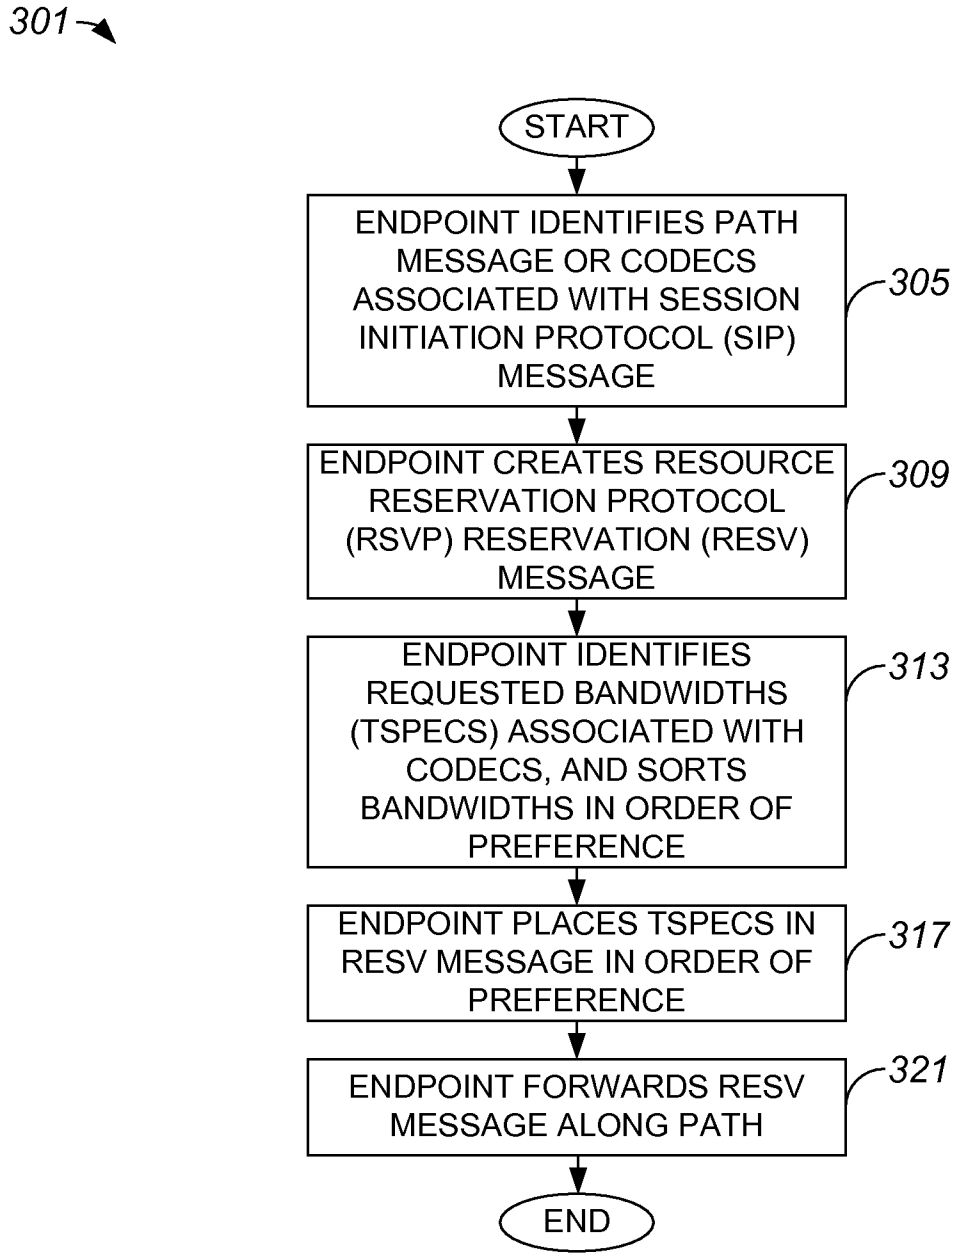
301
START
ENDPOINT IDENTIFIES PATH MESSAGE OR CODECS ASSOCIATED WITH SESSION INITIATION PROTOCOL (SIP) MESSAGE
305
ENDPOINT CREATES RESOURCE RESERVATION PROTOCOL (RSVP) RESERVATION (RESV) MESSAGE
309
ENDPOINT IDENTIFIES REQUESTED BANDWIDTHS (TSPECS) ASSOCIATED WITH CODECS, AND SORTS BANDWIDTHS IN ORDER OF PREFERENCE
313
ENDPOINT PLACES TSPECS IN RESV MESSAGE IN ORDER OF PREFERENCE
317
ENDPOINT FORWARDS RESV MESSAGE ALONG PATH
321
END

FIG. 3

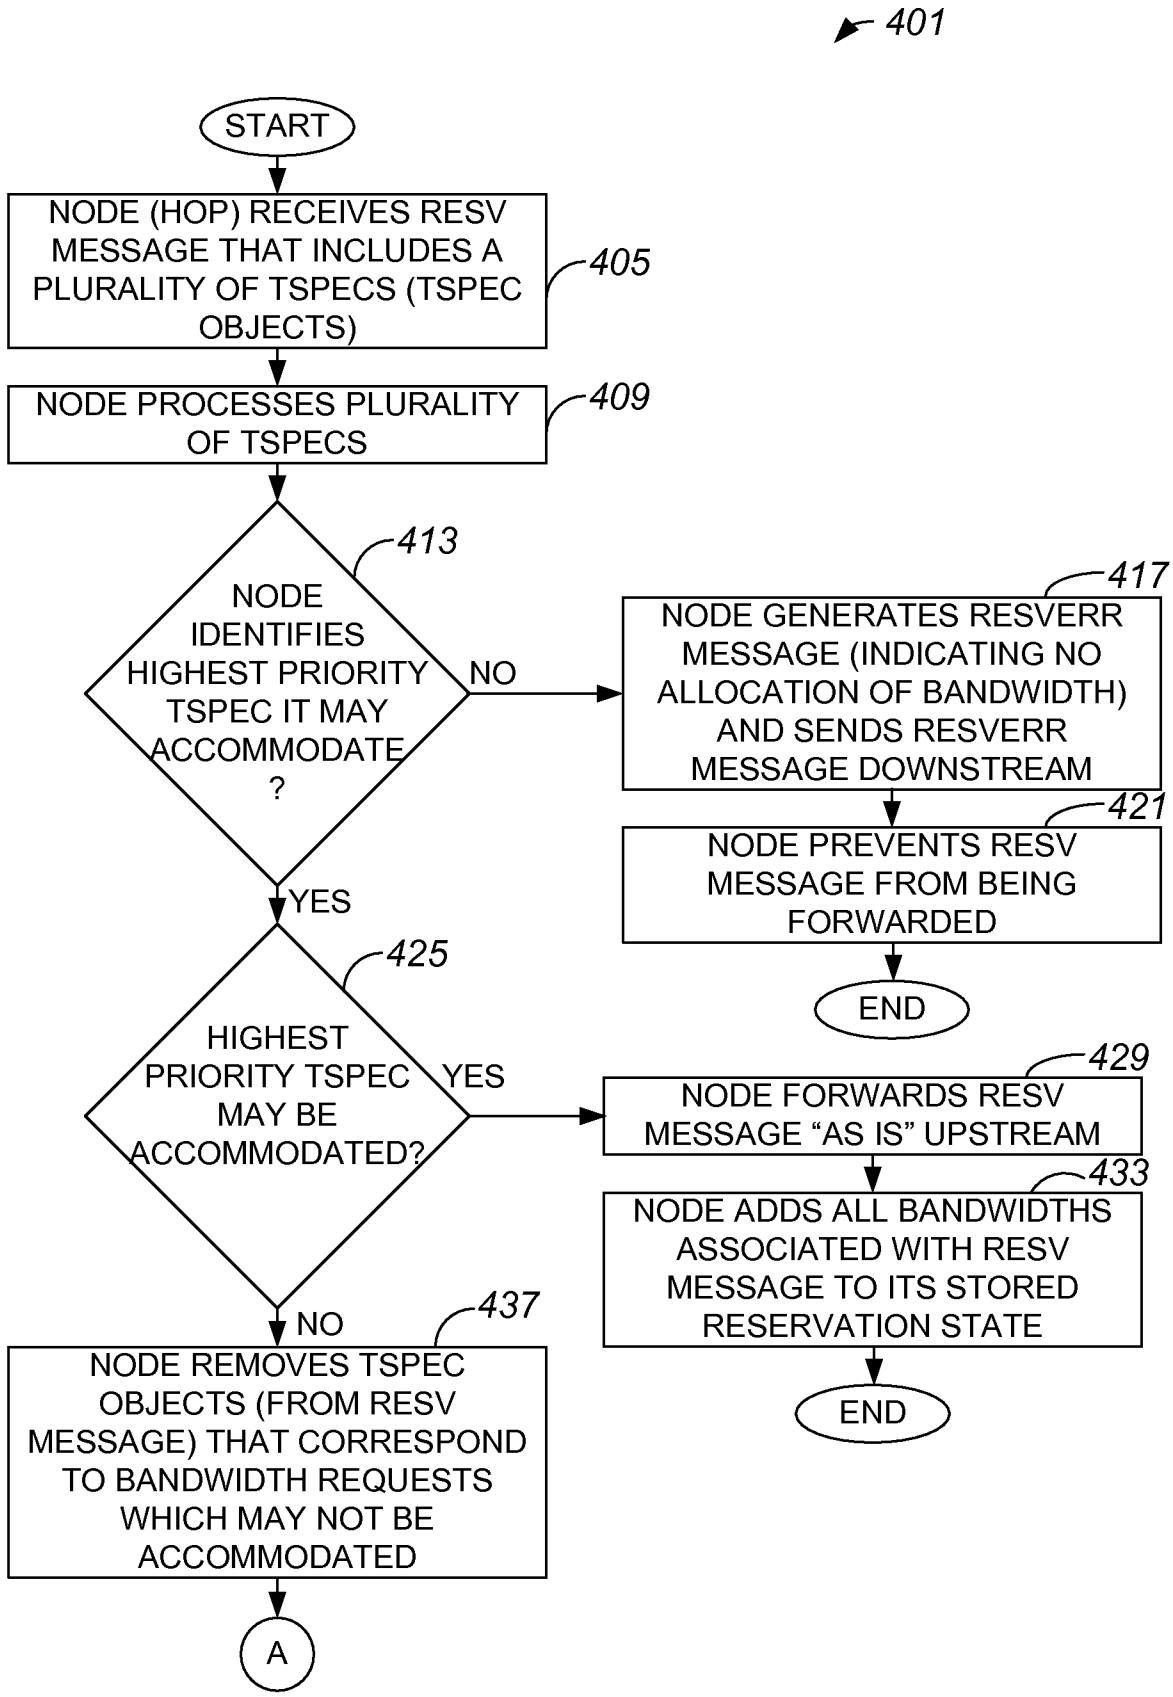
401
START
NODE (HOP) RECEIVES RESV MESSAGE THAT INCLUDES A PLURALITY OF TSPECS (TSPEC OBJECTS)
405
NODE PROCESSES PLURALITY OF TSPECS
409
NODE IDENTIFIES HIGHEST PRIORITY TSPEC IT MAY ACCOMMODATE ?
413
NO
NODE GENERATES RESVERR MESSAGE (INDICATING NO ALLOCATION OF BANDWIDTH) AND SENDS RESVERR MESSAGE DOWNSTREAM
417
NODE PREVENTS RESV MESSAGE FROM BEING FORWARDED
421
END
YES
HIGHEST PRIORITY TSPEC MAY BE ACCOMMODATED?
425
YES
NODE FORWARDS RESV MESSAGE “AS IS” UPSTREAM
429
NODE ADDS ALL BANDWIDTHS ASSOCIATED WITH RESV MESSAGE TO ITS STORED RESERVATION STATE
433
END
NO
NODE REMOVES TSPEC OBJECTS (FROM RESV MESSAGE) THAT CORRESPOND TO BANDWIDTH REQUESTS WHICH MAY NOT BE ACCOMMODATED
437
A

FIG. 4A

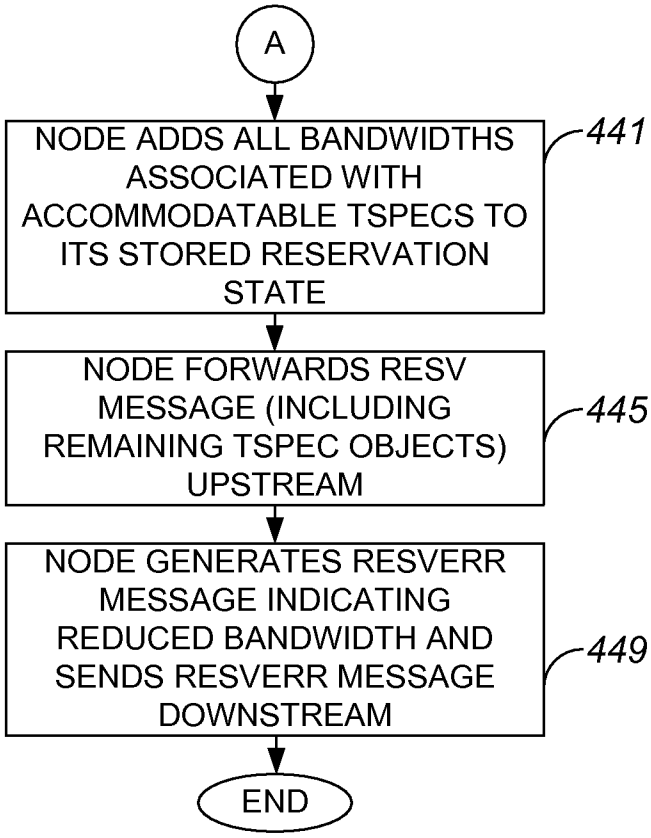
A
NODE ADDS ALL BANDWIDTHS ASSOCIATED WITH ACCOMMODATABLE TSPECS TO ITS STORED RESERVATION STATE
441
NODE FORWARDS RESV MESSAGE (INCLUDING REMAINING TSPEC OBJECTS) UPSTREAM
445
NODE GENERATES RESVERR MESSAGE INDICATING REDUCED BANDWIDTH AND SENDS RESVERR MESSAGE DOWNSTREAM
449
END

FIG. 4B

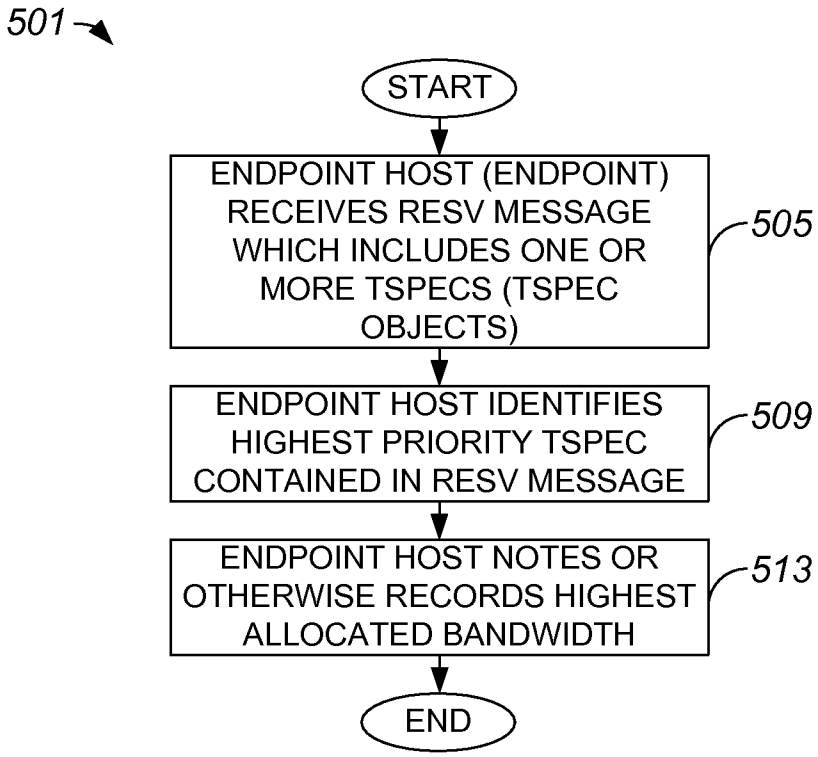
501
START
ENDPOINT HOST (ENDPOINT) RECEIVES RESV MESSAGE WHICH INCLUDES ONE OR MORE TSPECS (TSPEC OBJECTS)
505
ENDPOINT HOST IDENTIFIES HIGHEST PRIORITY TSPEC CONTAINED IN RESV MESSAGE
509
ENDPOINT HOST NOTES OR OTHERWISE RECORDS HIGHEST ALLOCATED BANDWIDTH
513
END

FIG. 5

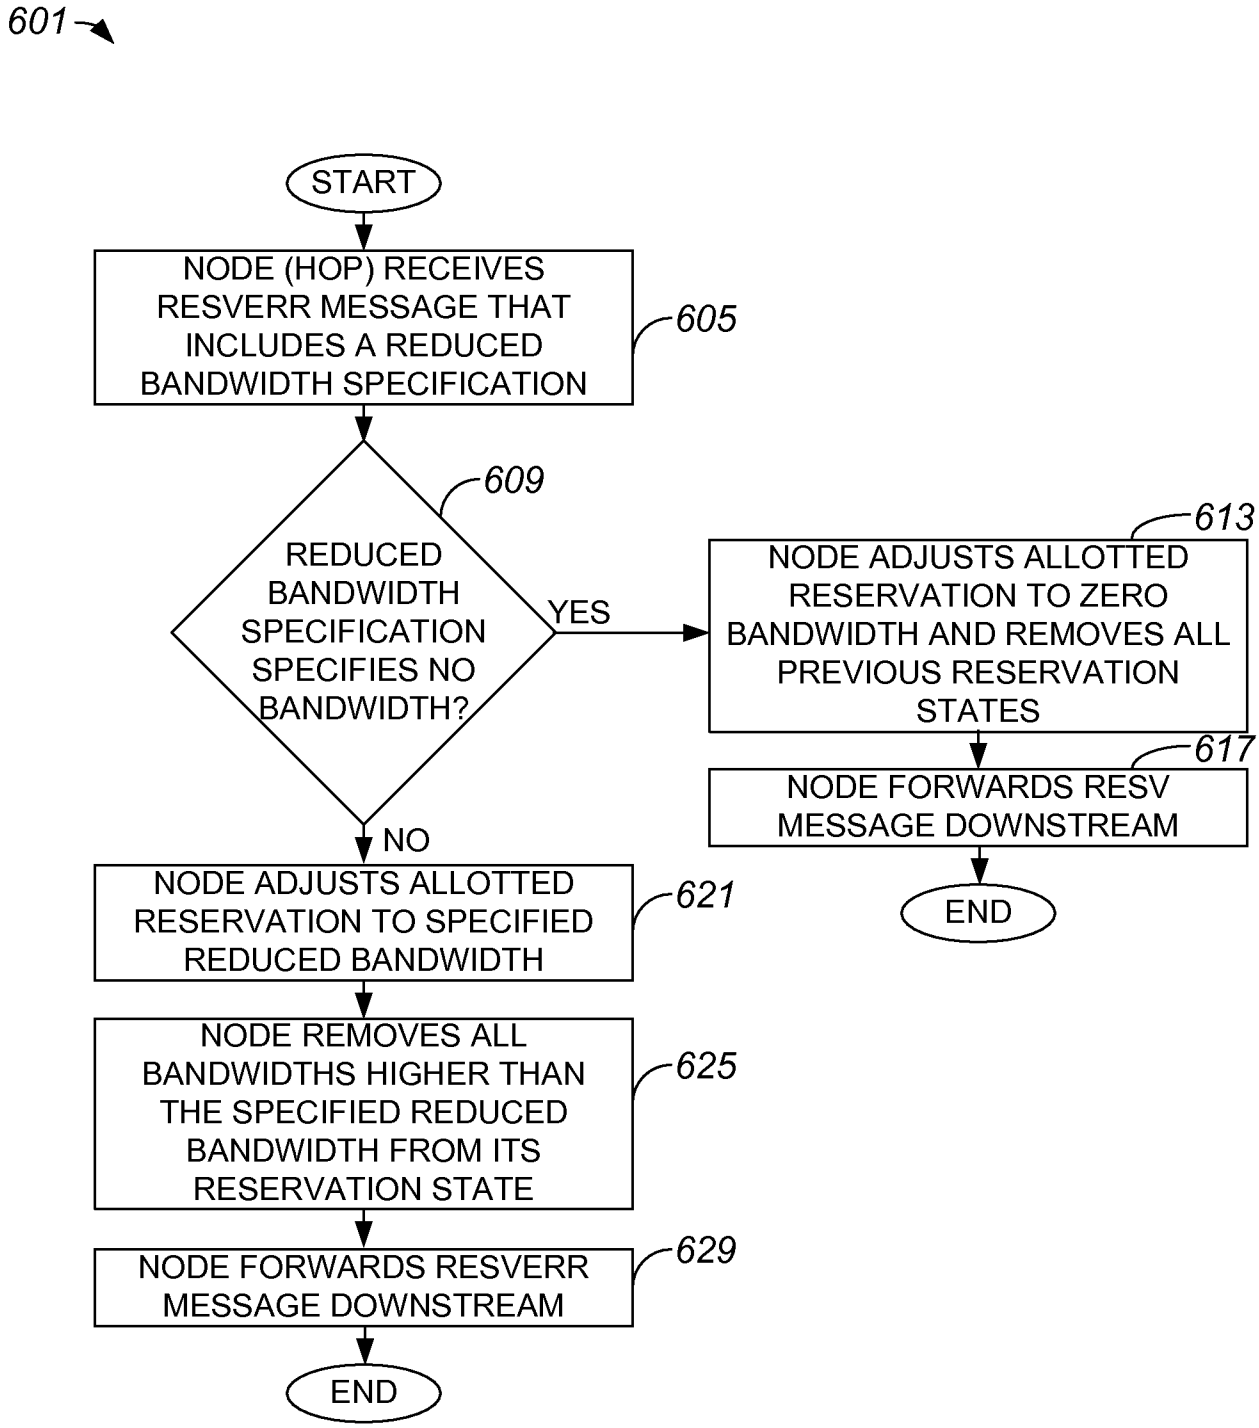
601
START
NODE (HOP) RECEIVES RESVERR MESSAGE THAT INCLUDES A REDUCED BANDWIDTH SPECIFICATION
605
609
REDUCED BANDWIDTH SPECIFICATION SPECIFIES NO BANDWIDTH?
YES
613
NODE ADJUSTS ALLOTTED RESERVATION TO ZERO BANDWIDTH AND REMOVES ALL PREVIOUS RESERVATION STATES
617
NODE FORWARDS RESV MESSAGE DOWNSTREAM
END
NO
NODE ADJUSTS ALLOTTED RESERVATION TO SPECIFIED REDUCED BANDWIDTH
621
NODE REMOVES ALL BANDWIDTHS HIGHER THAN THE SPECIFIED REDUCED BANDWIDTH FROM ITS RESERVATION STATE
625
NODE FORWARDS RESVERR MESSAGE DOWNSTREAM
629
END

FIG. 6

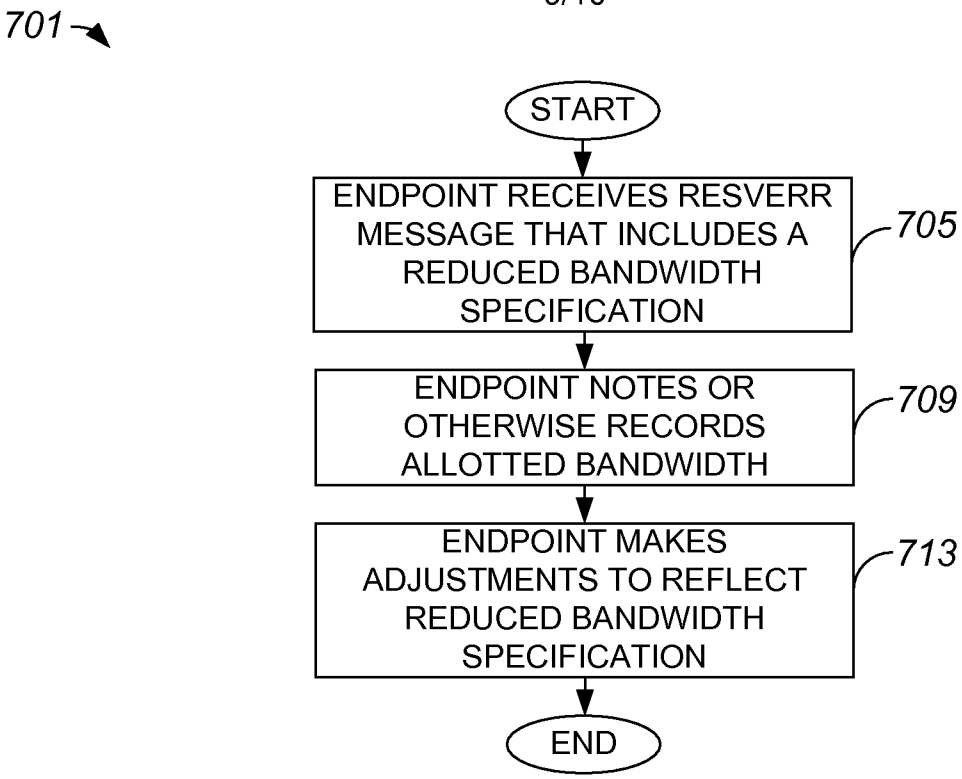
701
START
ENDPOINT RECEIVES RESVERR MESSAGE THAT INCLUDES A REDUCED BANDWIDTH SPECIFICATION
705
ENDPOINT NOTES OR OTHERWISE RECORDS ALLOTTED BANDWIDTH
709
ENDPOINT MAKES ADJUSTMENTS TO REFLECT REDUCED BANDWIDTH SPECIFICATION
713
END

FIG. 7

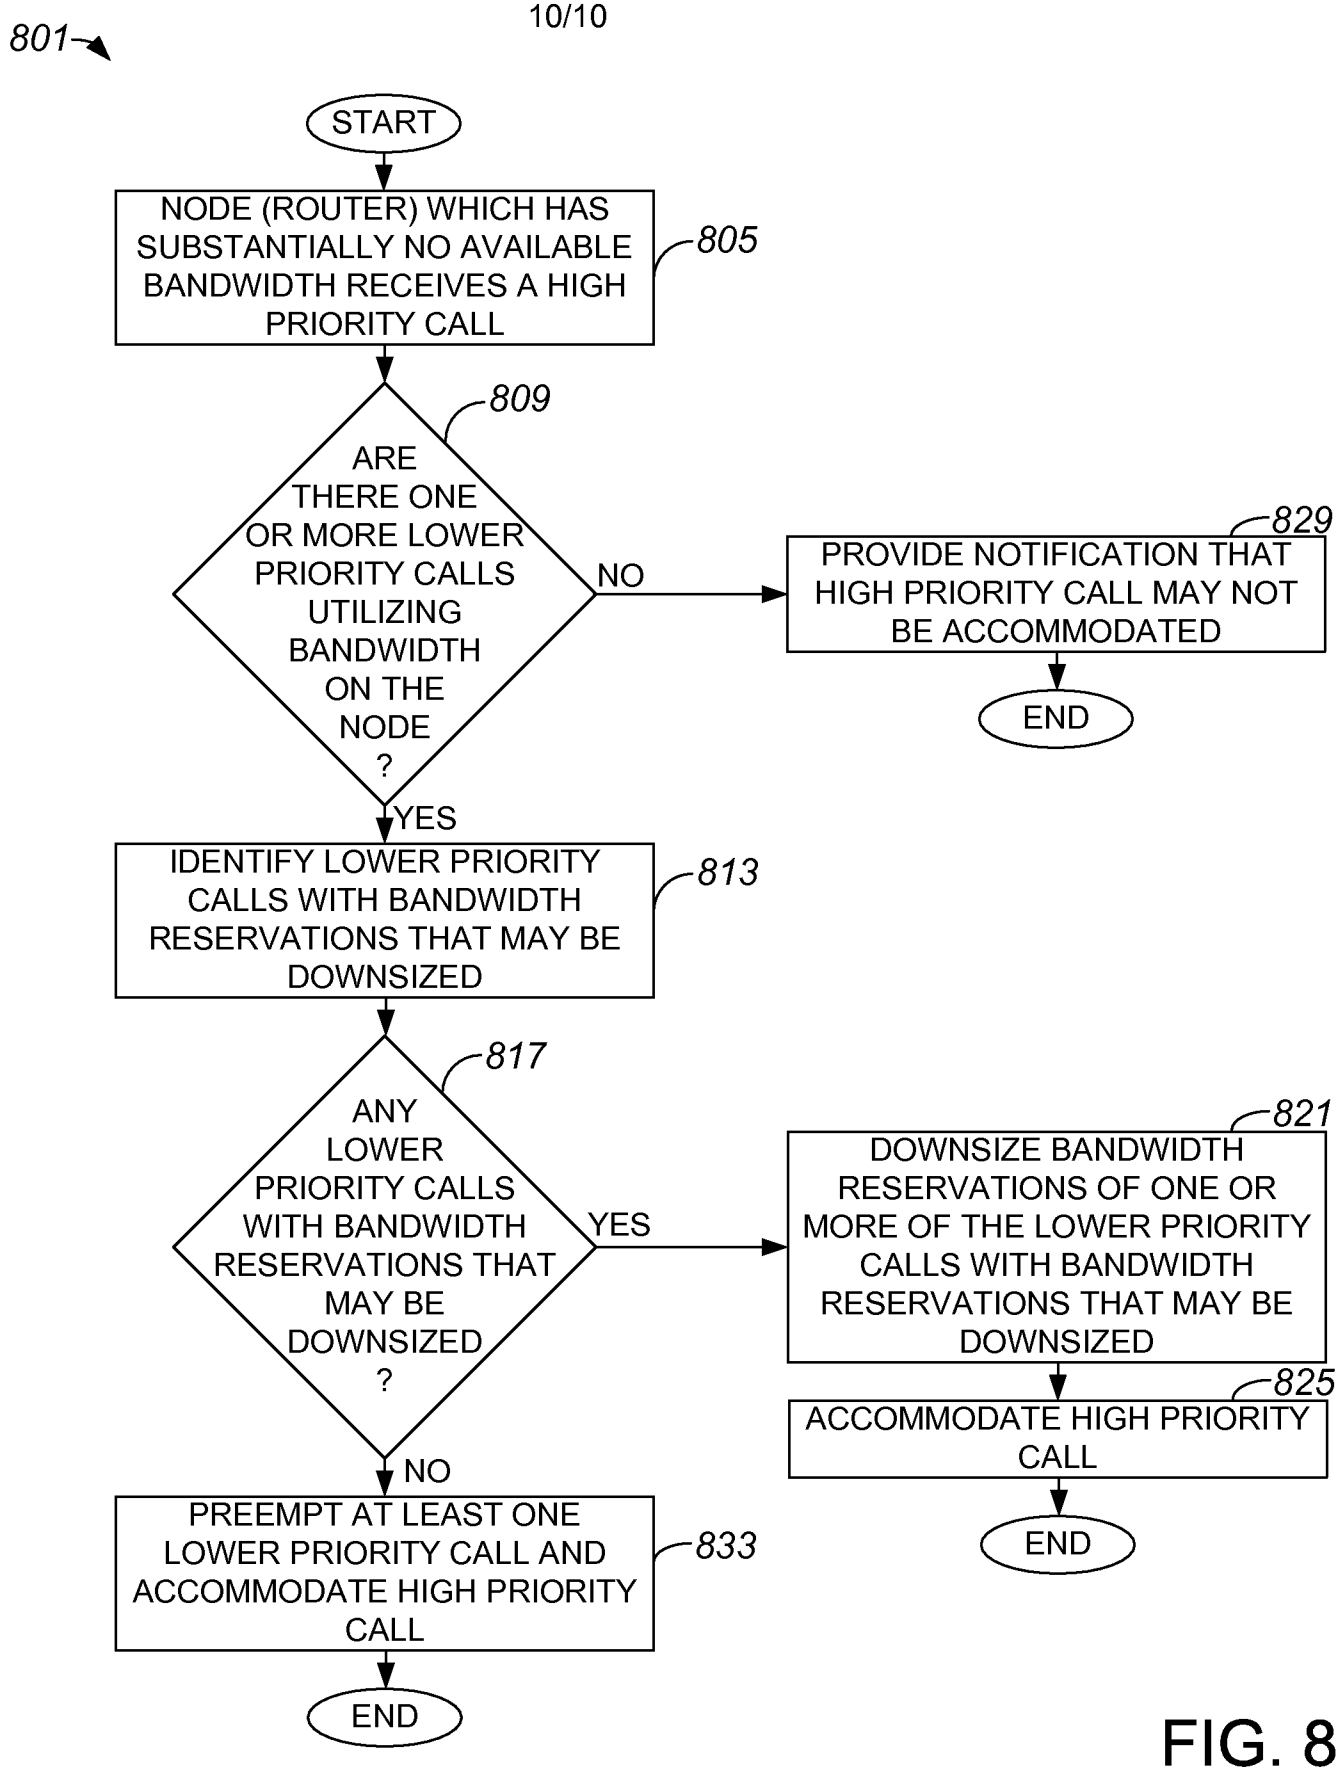
801
START
NODE (ROUTER) WHICH HAS SUBSTANTIALLY NO AVAILABLE BANDWIDTH RECEIVES A HIGH PRIORITY CALL
805
809
ARE THERE ONE OR MORE LOWER PRIORITY CALLS UTILIZING BANDWIDTH ON THE NODE ?
NO
829
PROVIDE NOTIFICATION THAT HIGH PRIORITY CALL MAY NOT BE ACCOMMODATED
END
YES
IDENTIFY LOWER PRIORITY CALLS WITH BANDWIDTH RESERVATIONS THAT MAY BE DOWNSIZED
813
817
ANY LOWER PRIORITY CALLS WITH BANDWIDTH RESERVATIONS THAT MAY BE DOWNSIZED ?
YES
821
DOWNSIZE BANDWIDTH RESERVATIONS OF ONE OR MORE OF THE LOWER PRIORITY CALLS WITH BANDWIDTH RESERVATIONS THAT MAY BE DOWNSIZED
825
ACCOMMODATE HIGH PRIORITY CALL
END
NO
PREEMPT AT LEAST ONE LOWER PRIORITY CALL AND ACCOMMODATE HIGH PRIORITY CALL
833
END

FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2009/061857

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/56

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)
H04L

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

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2008/008886 A1 (QUALCOMM INC [US]; GINDE SAMIR VIJAY [US]; BALASUBRAMANIAN SRINIVASAN) 17 January 2008 (2008-01-17)	1,4-8, 11-12, 15-21, 23,26
Y	paragraph [0026] - paragraph [0063]; figures 1,2	2-3, 9-10, 13-14, 22,24-25

-/--

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

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
18 December 2009	04/01/2010

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	García Bolós, Ruth

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2009/061857

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	POLK J. ET AL.: "A Resource Reservation Protocol (RSVP) Extension for the Reduction of Bandwidth of a Reservation Flow" REQUEST FOR COMMENTS (RFC) 4495, May 2006 (2006-05), pages 1-21, XP002560136 Retrieved from the Internet: URL:http://tools.ietf.org/pdf/rfc4495.pdf> the whole document	2-3, 9-10, 13-14, 22,24-25
A	BOS S LEROY J-C ROJAS L THIEBAUT J VANDENAMEELE ALCATEL P VEENSTRA KPN L: "A Framework for End-to-End User Perceived Quality of Service Negotiation; draft-bos-mmusic-sdpqos-framework-00.txt" IETF STANDARD-WORKING-DRAFT, INTERNET ENGINEERING TASK FORCE, IETF, CH, 1 November 2001 (2001-11-01), XP015011049 ISSN: 0000-0004 the whole document	1-26

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2009/061857

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2008008886 A1	17-01-2008	EP 2052518 A1	29-04-2009
		KR 20090031776 A	27-03-2009
		US 2008013544 A1	17-01-2008