

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 April 2012 (12.04.2012)

(10) International Publication Number
WO 2012/047880 A2

- (51) **International Patent Classification:**
H04W 40/20 (2009.01)
- (21) **International Application Number:**
PCT/US2011/054755
- (22) **International Filing Date:**
4 October 2011 (04.10.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/389,459 4 October 2010 (04.10.2010) US
- (71) **Applicant (for all designated States except US):** TELCORDIA TECHNOLOGIES, INC.; One Telcordia Drive 5G116, Piscataway, NJ 08854-4157 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** NATARAJAN, Narayanan [US/US]; 4 Heron Drive, Marlboro, NJ 07746 (US). BAGCHI, Anindo [US/US]; 42 Meadows Lane, Ocean, NJ 07712 (US). STEPHENS, William, E. [US/US]; 23536 Haas Avenue, Torrance, CA 90501 (US).
- (74) **Agents:** FEIG, Philip, J. et al.; TELCORDIA TECHNOLOGIES, INC., One Telcordia Drive 5G116, Piscataway, NJ 08854-4157 (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, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

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



WO 2012/047880 A2

(54) **Title:** A METHOD AND SYSTEM FOR DETERMINATION OF ROUTES IN LEO SATELLITE NETWORKS WITH BANDWIDTH AND PRIORITY AWARENESS AND ADAPTIVE REROUTING

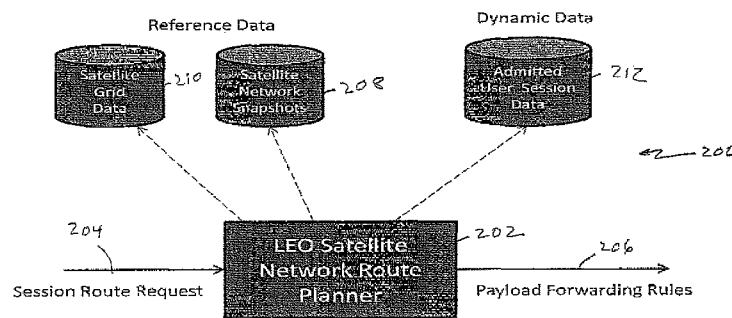


Figure 2.

(57) **Abstract:** In Low Earth Orbit (LEO) satellite networks data routes for point to point and point to multipoint communication sessions in LEO satellite networks are determined by considering session bandwidth, priority, and the constantly changing terminal-satellite and satellite-satellite connectivity. Multiple routes are computed for a session to facilitate automatic adaptive rerouting by satellite payloads when they encounter network failures or congestion conditions.

Method and System for Determination of Routes in LEO Satellite Networks with Bandwidth and Priority Awareness and Adaptive Rerouting

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 61/389,459, filed on October 4, 2010, which is incorporated by reference herein in its entirety.

[0002] This application is related to U.S. Patent Application Serial No. 13/149,106 filed May 31, 2011, entitled "Context Aware Adaptive Switching in Reconfigurable Low Earth Orbit Satellite Networks" which is incorporated by reference herein in its entirety.

FIELD OF THE INVENTION

[0003] The present invention relates generally to Low Earth Orbit (LEO) satellite networks, and specifically to the determination of routes for point to point and point to multipoint communication sessions in LEO satellite networks considering session bandwidth, priority, and the constantly changing terminal-satellite and satellite-satellite connectivity. Multiple routes are computed for a session to facilitate automatic adaptive rerouting by satellite payloads when they encounter network failures or congestion conditions.

BACKGROUND OF THE INVENTION

[0004] Use of Low Earth Orbit (LEO) satellite networks is a promising approach for global data communications (including broadband) in view of their low delay and bit error characteristics compared to Geostationary (GEO) satellites. Further, they do not require high power antennas in ground terminals due to their low distance from the ground. However, due to their orbit, the communication link between a ground terminal and an LEO satellite will be available periodically only for a few minutes. In order to provide continuous communication between ground terminals, the LEO satellites need to be interconnected via inter-satellite links (called crosslinks) and terminal-terminal data traffic will be routed over multiple satellites using these crosslinks. Typically, an LEO satellite network consists of multiple orbit planes, and each plane consists of multiple satellites. The satellite topology (neighbor relationship) within a plane remains invariant while inter-plane topology will change

constantly; interplane crosslinks will be dynamically set up and removed. Terminal-satellite associations (uplink and downlink connectivity) will also be changing constantly. Traffic routing in such a dynamically changing network considering the bandwidth and priority requirements of user traffic, network failures, and network congestion is a complex problem.

[0005] Several proposals have been made for distributed routing in LEO satellite networks; i.e., each payload router determines routing decisions independently for each packet.

[0006] E. Ekici, I.F. Akyildiz, M.D. Bender, A Distributed Routing Algorithm for Datagram Traffic in LEO Satellite Networks, IEEE/ACM Transactions on Networking, VOL.9, NO.2, April 2001 propose a distributed routing algorithm for routing datagram traffic in an LEO satellite network. The algorithm generates minimum propagation delay paths. The algorithm is based on the concept that the entire earth is covered by logical locations of satellites. These logical locations do not move and are filled by the nearest satellite. This concept is similar to the grid concept in the present invention. In Ekici et al routing is performed by considering these logical locations as hops. Unlike classical routing algorithms used in terrestrial routers, the algorithm does not require satellites to exchange topology information, but computes routes based on some mathematical properties of minimum delay paths between logical locations taking into account the different distances between neighbor locations in polar region and non-polar regions. Based on the mathematical properties, a routing table is generated at design time that prescribes on which direction (up, down, right, or left) a satellite should route a packet depending on which region the destination is in (polar or non-polar). The table also includes a secondary direction that should be used if the link on the primary direction is congested, i.e., the output queue is longer than some threshold. Thus, the routing algorithm is capable of rerouting packets around congestion and failure points with low degradation in performance. This routing algorithm is intended for pure datagram traffic. It does not consider bandwidth (QoS) and priority of traffic.

[0007] Ekici et al., have extended this algorithm to multicast routing in E. Ekici, I.F. Akyildiz, M.D. Bender, A Multicast Routing Algorithm for LEO Satellite IP Networks, IEEE/ACM Transactions on Networking, VOL.10, NO.2, April 2002. The multicast extension requires routing information exchange between payloads.

[0008] Vijay Kapoor, Directional Routing of Packets in a Satellite Network, U S Patent No. 6,404,769, issued June 11, 2002 proposes a method for routing packets across a satellite

constellation that is made up of multiple orbital planes, each plane carrying multiple satellites that are equally spaced in the plane. Satellites in the same orbital plane are connected via north-south crosslinks. Satellites in different orbital planes are connected via east-west crosslinks. When a connection is needed between two satellites that are in different planes, a controller station in the ground determines a route for the connection taking into account the traffic load on various crosslinks. That is, the cost for route optimality is traffic load on the path rather than delay. The route so determined consists of a "band" which denotes a logical region in the constellation that is orthogonal to orbital planes. The controller provides this route information including the band information, the destination plane, and the destination satellite id to the source satellite which includes it in all packets belonging to that connection. Thus, this scheme employs source routing. Satellites forward packets using this route information by forwarding packets on the source plane until the specified band is reached, then on the band towards the destination plane, and finally on the destination plane towards the destination satellite. In this method, the satellites do not have any route adaptation capabilities. All rerouting decisions are made only by the controller on the ground. Another limitation of this scheme is that it does not consider traffic priority and does not support preemption. Also, the "band" based routes are unlikely to be optimal, since they constrain shortest path selection.

[0009] V.V.Gounder, R. Prakash, H. Abu-Amara, Routing in LEO-Based Satellite Networks, IEEE Symposium on Wireless Communications and Systems, April 1999 propose a virtual connection-oriented approach based on tag switching to route data in an LEO satellite network. Routing tables are computed by ground stations and uploaded to satellites at predetermined intervals. Routes are computed by ground stations for different network snapshots. Only forwarding tables for a limited set of snapshots are uploaded to LEO satellites. The payload switch does not have any route adaptation capabilities. All rerouting computations and decisions are made only by ground stations. The ground stations generate routes for each snapshot by determining k shortest paths between the source and the destination. Each one of the k paths is used for routing a specific type of traffic; e.g., delay sensitive traffic is routed on the first shortest path and other types of traffic are routed on the remaining k-1 paths. This scheme also does not consider traffic priority and does not support preemption.

SUMMARY OF THE INVENTION

[0010] The present invention presents a scheme for the determination of routes between a source ground terminal and one or more destination ground terminals across an LEO satellite network. The scheme is intended to be used by a network operations center on the ground that receives requests from users of the satellite network to establish communication sessions between two or more ground terminals with specific bandwidth, priority, and duration and generates network routes for the requested sessions.

[0011] The scheme has the following innovative features:

1. Routes are determined by the operations center considering session bandwidth, priority, and duration requirements.
2. If necessary, lower priority sessions are preempted to accommodate an incoming high priority session. The number of sessions preempted is minimized.
3. Routes are generated considering changes in terminal-satellite connectivity and satellite-satellite connectivity that may occur during the session lifetime. Thus, data between two ground terminals may be routed via different paths (crosslinks/uplink/downlink) in the satellite network at different time instants.
4. Multiple routes are computed for a session to facilitate automatic adaptive rerouting by the payloads when they encounter crosslink failures or congestion. Payloads do not require knowledge of the entire network topology or status. A payload needs to know only the failure or congestion status of its crosslinks to perform adaptive rerouting upon crosslink failures or congestion. Based on the determined routes for a session, the operations center on the ground generates data forwarding rules for payloads and uploads them to the payloads. The forwarding rules specify for each payload the crosslink on which a session traffic should be routed under normal conditions, and the crosslink on which the session traffic should be routed in case of congestion or failure.
5. Routes are generated for both point-to-point sessions and point-to-multipoint sessions.

[0012] Some of the prior solutions to LEO satellite routing, such as the work of Ekici et al., present distributed routing algorithms that are executed in payloads. The invention described herein presents a centralized routing solution that is executed in a planning node on the ground.

[0013] Some other prior solutions to LEO satellite routing, such as the work of Kapoor and Grounder et.al., present a centralized routing solution. However, they do not consider session priority and routes are computed without considering preemption of lower priority sessions. Further, they do not support adaptive rerouting by payloads. When network failures or congestions occur, new routes are computed by the planning node on the ground only after the events occur. This results in a slow reaction to network events and causes disruption of communications. Further, neither Kapoor nor Grounder et.al. provides point-to-multipoint routing.

[0014] A novel feature of the described invention is the clear separation of route determination from payload forwarding rule generation. Routes are computed based on the Grid concept that represents a logical satellite constellation topology that is invariant over time and available bandwidth on logical crosslinks that varies over time. This concept facilitates route computation without considering satellite orbits. After the routing algorithm has determined the logical paths (primary and alternative paths), payload forwarding rules are generated. Only in this phase, the satellite orbits are considered and forwarding rules are generated considering the satellite geolocations at different time instants during the session lifetime.

[0015] The primary uniqueness of the present invention is the LEO Satellite system context considered. The system context is shown in Figure 1. The payloads being considered are resource limited and do not have the computing and memory resources to implement a routing algorithm. Despite this resource limitation, one of the major design goals is satellite autonomy that enables a payload to rapidly react to network anomalies. The routing scheme described herein addresses the payload resource constraint and the satellite autonomy objectives. Hence, the routing algorithm is centralized and performed on a ground node, but yet generates multiple routes that facilitate payloads to reroute traffic when they encounter a failure or congestion. A payload needs to know only the failure and congestion status of its crosslinks, and does not need to know network-wide global conditions.

[0016] The invention will be better understood when the following description is read in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] Figure 1 illustrates a LEO satellite system context.

[0018] Figure 2 is a schematic diagram of LEO satellite route planning.

[0019] Figure 3 is a schematic representation of an LEO satellite grid.

[0020] Figures 4A and 4B are a flowchart for determining the primary path and alternative paths for a session.

[0021] Figures 5A and 5B are a flowchart for determining a path between two nodes for a session.

[0022] Figure 6 is a flowchart for determining an alternative path between two nodes for a session.

[0023] Figure 7 is a flowchart for the generation of payload forwarding rules for a session.

DETAILED DESCRIPTION

[0024] Figure 1 illustrates a LEO satellite system context 100 for the route determination scheme described hereinafter. As shown in the figure, the Network Operations Center 102 plans routes for user sessions and disseminates the determined routes to Distributed Management Centers 104A, 104 B, etc.. The route information is conveyed in the form of forwarding rules for each payload involved in the route. The Distributed Management Centers upload the forwarding rules to the payloads 106A, 106B, 106C, 106D, and 106E within its domain. Payloads forward data packets in accordance with these forwarding rules.

[0025] The Network Operations Center includes centralized functions for Network Management. The Network Operations Center plans routes for services, generated payload forwarding rules, disseminates payload forwarding rules to management centers, generates SLA reports, and performs network trouble resolution.

[0026] The Distributed Management Center uploads network configuration and policies to payloads and terminals, monitors payload and terminal status and performance, and reports aggregated performance data to the Network Operations Center.

[0027] The payload manages resources and forwards packets in accordance with policies, dynamically reroutes packets upon link failure or congestion based on policies, dynamically establishes crosslinks based on policies, and reports status and performance to a designated management center.

[0028] The concept and details of forwarding rules have been described in the related patent application 13/149,106 and they are summarized here with some simplifications for easy reference.

[0029] A forwarding rule specifies how a data packet received on a crosslink/uplink should be forwarded to another crosslink/downlink.

[0030] A forwarding rule is of the form

[0031] <Source, Destination, Start Time, End Time, Priority, Output Port 1, Output Port 2, ..., Output Port n>

where

Source is the Source Address in the packet received;

Destination is the Destination address in the packet received;

Start Time and End Time specify the time interval within which the rule is effective; the time interval may be a recurring interval.

Priority is the priority level assigned to the packet (in the range 0-7, 0 being the lowest and 7 being the highest priority) in accordance with IEEE 802.1pQ Tag standard;

Output Port 1, ..., Output Port n is an ordered sequence of preferred crosslinks/downlink to which the received packet is to be switched; an Output Port in this sequence is selected if and only if either all previous Output Ports in the sequence have either failed or congested (as determined by a preconfigured threshold)

[0032] A forwarding rule specifies the following forwarding behavior. A packet received with source and destination address set to Source and Destination respectively between the time interval Start Time and End Time and that has priority level Priority should be switched to the first found crosslink/downlink, if any, in the specified Output Port sequence, that has neither failed nor congested. If no such crosslink/downlink is found, the packet should be dropped.

[0033] Figure 2 shows a block schematic diagram 200 for a functional component called the LEO Satellite Network Route Planner 202 that incorporates the route determination scheme. This component takes as input 204 a Session Route Request of the form <Session Id, Source Terminal Address, Source Terminal Location, Addresses and Locations of Destination Terminals, Start Time, End Time, Bandwidth, Priority> where Start Time and End Time together specify the session duration. The Route Planner produces as output 206 the payload forwarding rules for the routes determined for the session. The topology of the satellite network may change during the lifetime of the session due to satellite orbits. Such topology changes are captured using a sequence of snapshots 208, where each snapshot represents the network topology for a specific time interval. If the duration of a session encompasses multiple snapshots, the Route Planner generates forwarding rules for each snapshot. In addition, the Route Planner also computes multiple routes for a session to facilitate adaptive rerouting by the payloads in the following manner. For each session, a primary path and multiple alternative paths are determined. Each alternative path is a path that branches at some payload involved in the primary path, and is to be used by the payload if it encounters a failure or congestion in the crosslink traversed by the primary path egressing from the payload.

[0034] The satellite network snapshots are represented using the concept of LEO Satellite Grid (hereafter called the Grid) as shown in Figure 3. The Grid 300 represents a fixed logical satellite network (constellation) consisting of Nodes 302 connected via vertical edges 304 and horizontal edges 306. A Node represents a geolocation. The geolocations are chosen such that at any time, the LEO satellites at these locations together provide global coverage. A vertical edge represents a Logical Intra-Plane Crosslink, and an horizontal edge represents a Logical Inter-Plane Crosslink. Nodes in the top row and the bottom row are the same and this is denoted by the dotted lines between these nodes. Similarly, nodes in the left most column and right most column are the same and this is also denoted using dotted lines. In addition to the logical topology, the Grid representation also includes delay information for each Logical Inter-Plane and Intra-Plane Crosslink. Delays are minimum near the polar region and maximum in the equator region. This attribute is the cost metric in route computation.

[0035] The Satellite Grid Data 210 shown in Figure 2 has the following information:

- a. Geolocation of each node

b. Coverage area of each node, i.e., the area in the ground that is covered by a satellite when it occupies the node

c. Crosslink delay for each Logical Intra-Plane Crosslink and Logical Inter-Plane Crosslink

[0036] A snapshot is an instance of the grid for a specific time interval and is specified by a set of <Node, Satellite Id> mappings identifying which satellite is occupying which node. This is predetermined based on satellite orbits. From these mappings, the mappings of Logical Crosslinks to Physical Crosslinks (i.e., crosslinks between satellite pairs) in a snapshot are derived. Thus, a snapshot represents the actual satellite network topology for a specific time interval. The satellite network topology changes in a time deterministic and recurring manner. An epoch is the shortest cyclic time interval such that the same sequence of topology changes is repeated in each epoch. All topology changes are entirely captured by the sequence of snapshots within an epoch. Since the sequence of snapshots is the same in each epoch, it is possible to uniquely identify each snapshot by a number in the range 1 to n where n is the number of snapshots in an epoch.

[0037] As shown in Figure 2, the Satellite Grid data 210 and the Satellite Network Snapshot 208 data are reference data for the Route Planner. These are static data and are provided to Route Planner prior to servicing session route requests.

[0038] In addition to this reference data, the Route Planner uses another database, called Admitted User Session Data 212 that has dynamic data. This database stores session and route information for each user session that has been admitted into the satellite network. For each session, the following information is stored: Session Id, Start Time, End Time, Bandwidth, Priority, Primary Route, and Alternative Routes. A route is represented as a sequence of logical crosslinks (intra-plane or inter-plane) traversed by the route. This database is updated by the Route Planner when it completes route determination for a session.

[0039] Using information in these databases, the Route Planner determines routes for a session in the following manner. The basic concept is that the Route Planner determines routes based on the network topology and delay information in the Grid Data and available bandwidth on each crosslink for the duration of the requested session that is derived from data stored in Admitted User Session Data 212. After the routes are determined for a session,

the Planner 202 generates payload forwarding rules 206 for the routes based on Satellite Network Snapshots data 210.

[0040] First is a description of the scheme for point-to-point routing, i.e., finding routes for a session with one source terminal and one destination terminal. Extensions for routing with one source and multiple destinations (point-to-multipoint routing) are described later.

[0041] The routing problem for determining paths between two terminals is reduced to the problem of finding paths between two nodes in a grid as follows. The source node is determined based on Source Terminal Location and the coverage area of each node. A node whose coverage area contains the Source Terminal Location is selected as the source node. In case of multiple candidate nodes, one node is chosen based some predetermined criteria. The destination node is determined in a similar manner.

[0042] The algorithm for finding a Primary Path and Alternative Paths for a point to point session is described in the flowcharts shown in Figure 4A and Figure 4B. This algorithm uses two other algorithms: one for finding a path between two nodes for a session which is shown in Figure 5A and Figure 5B; and another for finding an alternative path between two nodes for a session given the primary path which is shown in Figure 6. Figure 7 illustrates the flow chart for the generation of forwarding rules for a session.

[0043] The process for determining the primary path and alternative paths for a sessions starts 402 and Victims is set to an empty set 404, where Victims represents the set of sessions that are preempted to accommodate the requested session.

[0044] Find the primary path for the session by invoking Find Path (Source Node, Destination Node, Start Time, End Time, Bandwidth, Priority, Victims) 406.

[0045] A determination is made whether a primary path is found 408. If a primary path is not found, the system returns a failure 410 and the process ends 412.

[0046] If a primary path is found, consider each logical crosslink L in the primary path, initially the first crosslink 414. The Branch Node is made source node of L in step 416. Next Find an alternative path for the session branching at Branch Node by invoking Find Alternative Path (Primary Path, Branch Path, Destination Node, Start Time, End Time, Bandwidth, Priority, Victims) 418.

[0047] Determine if there are more logical crosslinks to consider 420. If yes, return to step 414. If no, Update Admitted User Session Data: Record the attributes of the new session including its Primary Path and Alternative Paths; Remove entries of the sessions included in Victims 422. Return Success, Primary Path and Alternative Path 424 and then end the process 426.

[0048] The process for determining a path between two nodes for a session is shown in Figures 5A and 5B. The process starts 502 and construct a reduced grid G for the session duration by removing logical crosslinks whose Available BW < Bandwidth sometime during the session duration after reclaiming bandwidth allocated to Victims 504. Find a shortest path between Source Node and Destination Node in the reduced grid using Dijkstra's algorithm and cross-link delay as cost metric 506. Determine if a path found 508. If yes, set Route to the path found 510 and Return Success, Route, Victims 512.

[0049] If no, update the grid G for the session duration by reclaiming bandwidth allocated to all overlapping sessions that have priority < Priority 516. Find a shortest path between Source Node and Destination Node in the updated grid G using Dijkstra's algorithm with cross-link delay as cost metric 518.

[0050] If no path is found 520 return failure 522 and end process 524. If a path is found 520 set Route to the path found 526. Consider each logical crosslink L in the path found; initially consider the first crosslink 528.

[0051] Is Sum of bandwidth of all sessions that is not in Victims but traverse L including the new session > Capacity of L sometime during session duration? 530. If no, return to step 528. If yes, consider in the order of increasing priority, each lower priority session S that traverse L and that is not in Victim 532. Reclaim the bandwidth of S, 534. Add S to Victims 536. Is Sum of bandwidth of all sessions that is not in Victims but traverse L including new session > Capacity of L sometime during session duration? 538. If yes, return to step 532. If no, Are there more logical crosslinks to consider for the path? 540. If yes return to step 528. If no, Return Success, Route, Victims 542 and end the process 544.

[0052] Figure 6 is a flowchart for determining an alternative path between two nodes for a session. The process starts 602 and constructs a reduced grid G for the session duration by removing from G the logical crosslink in Primary Path egressing from Branch Node 604.

Next, invoke Find Path (Branch Node, Destination Node, Start Time, EndTime, Bandwidth, Priority, Victims) 606. End the process 608.

[0053] The process for generating of forwarding rules for a session is shown in Figure 7. Start the process 700 by identifying all snapshots encompassed by the session 702. Consider each snapshot of the session, initially the first snapshot 704.

[0054] The details of the step "Construct forwarding rules for each payload in the primary path and alternative paths for the current snapshot" 706 in Figure 7 are as follows:

[0055] For each edge from a node n_1 to a node n_2 in the path, the forwarding rule in constructed as follows:

Let A be the satellite occupying n_1 (302_1) in the snapshot;

Let B be the satellite occupying n_2 (302_2) in the snapshot;

Let P be the output port in A for the crosslink connected to B;

Let t_s be either the start time of the session or the start time of the snapshot whichever is later;

Let t_e be either the end time of the session or the end time of the snapshot whichever is earlier;

Then, the forwarding rule for the satellite A is <Source Terminal Address, Destination Terminal Address, t_s , t_e , p, P> where p is the session priority.

[0056] The details of the step "Merge primary path and alternative path forwarding rules for each payload into a single forwarding rule for the current snapshot" 708 in Figure 7 are as follows:

Let <Source Terminal Address, Destination Terminal Address, t_s , t_e , p, P1> be the forwarding rule for a satellite A in the primary path in a snapshot.

Let <Source Terminal Address, Destination Terminal Address, t_s , t_e , p, P2> be the forwarding rule for an alternative path branching at satellite A in the same snapshot.

Then, the merged forwarding rule for satellite A in the same snapshot is <Source Terminal Address, Destination Terminal Address, t_s , t_e , p, P1, P2>.

Then determine if there are more snapshots to be considered 710. Is yes, then consider the next snapshot 704 and follow the steps above. If no, return forwarding rules 712 and then

end the process 714.

[0057] Point-to-Multipoint routing

[0058] The point-to-point routing algorithm described above is extended for multipoint routing in the following manner:

[0059] The point-to-point routing algorithm is repeatedly invoked by considering each destination in turn. The order in which the destinations are considered is unimportant.

[0060] At the end of the initial iteration, the multipoint tree is the path determined for the first destination.

[0061] At the end of each subsequent iteration, the path determined for the current destination (called current path) is merged with the multipoint tree which is the cumulative result of previous iterations. This merge is accomplished by scanning the nodes on the current path starting from the current destination until a node *n* is found such that *n* is on the current path as well as in the multipoint tree. In the worst case, the source node may be *n*. Then, the subpath starting from *n* in the current path is attached as a subtree of *n* in the multipoint tree.

[0062] To avoid double counting of bandwidth allocated on common edges (logical crosslinks) in the paths for different destinations in the multipoint tree, the link bandwidth allocated to the multipoint session is not considered while computing paths (Figures 5A and 5B). That is, the Available Bandwidth information for logical crosslinks is computed without considering the bandwidth allocated on the links for the multipoint session.

[0063] As will be appreciated by one skilled in the art, the present invention may be embodied as a system, method or computer program product. Accordingly, the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system."

[0064] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly

indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0065] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements, if any, in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

[0066] Various aspects of the present disclosure may be embodied as a program, software, or computer instructions stored in a computer or machine usable or readable storage medium or device, which causes the computer or machine to perform the steps of the method when executed on the computer, processor, and/or machine. A computer readable storage medium or device may include any tangible device that can store a computer code or instruction that can be read and executed by a computer or a machine. Examples of computer readable storage medium or device may include, but are not limited to, hard disk, diskette, memory devices such as random access memory (RAM), read-only memory (ROM), optical storage device, and other recording or storage media.

[0067] The system and method of the present disclosure may be implemented and run on a general-purpose computer or special-purpose computer system. The computer system may be any type of known or will be known systems and may typically include a processor, memory device, a storage device, input/output devices, internal buses, and/or a communications interface for communicating with other computer systems in conjunction with communication hardware and software, etc.

[0068] The terms “computer system” and “computer network” as may be used in the present application may include a variety of combinations of fixed and/or portable computer hardware, software, peripherals, and storage devices. The computer system may include a plurality of individual components that are networked or otherwise linked to perform collaboratively, or may include one or more stand-alone components. The hardware and software components of the computer system of the present application may include and may be included within fixed and portable devices such as desktop, laptop, and server. A module may be a component of a device, software, program, or system that implements some “functionality”, which can be embodied as software, hardware, firmware, electronic circuitry, or etc.

[0069] While there has been described and illustrated a method and system for determination of routes for point to point and point to multipoint communication sessions in LEO satellite networks, it will be apparent to those skilled in the art that variations and modifications are possible without deviating from the broad principles and teachings of the present invention which shall be limited solely by the scope of the claims appended hereto.

CLAIMS

What is claimed is:

1. A method of determining a primary path and multiple alternative paths for a point-to-point communication session between two ground terminals connected across by a Low Earth Orbit (LEO) satellite network comprising:
 - determining routes using the following information:
 - changes in terminal – satellite connectivity due to satellite orbits;
 - changes in satellite – satellite connectivity due to satellite orbits;
 - locations of Source and Destination terminals;
 - communication session duration;
 - session traffic bandwidth; and
 - session traffic priority;
 - determining routes for the primary path and alternative paths based on a LEO Satellite Grid where the Grid represents a fixed logical satellite constellation topology comprising logical nodes representing geolocations, vertical edges representing logical intra-plane crosslinks, horizontal edges representing logical inter-plane crosslinks, and delay information on crosslinks;
 - wherein the routes for the primary path and alternative paths are determined by computing shortest paths from a Source Node whose coverage area includes a Source terminal location to a Destination Node whose coverage area includes a Destination terminal location, taking into account available bandwidth on logical crosslinks for the duration of the communication session.
2. The method of claim 1 wherein the geolocations of nodes are chosen such that at any time, the LEO satellites at these locations together provide global coverage.
3. The method of claim 1 wherein a shortest path is a path that has the minimum delay across the Grid.

4. The method of claim 3 wherein the primary path is determined first without any preemption by constructing a reduced Grid by eliminating all crosslinks that do not have enough available bandwidth for some time during the session duration, and then invoking Dijkstra's algorithm to find a shortest path.
5. The method of claim 4 wherein a path is found by preempting lower priority sessions to accommodate the requested session if a path could not be found without preemption
6. The method of claim 5 wherein the route determination with preemption comprises:
 - checking if preemption is of help by determining if a route exists if bandwidth allocated to all lower priority sessions is reclaimed;
 - considering in turn each crosslink traversed by the route and checking if preemption is required on that crosslink to accommodate the session; and
 - considering each lower priority session traversing such above said crosslink in order of increasing priority, and reclaiming the bandwidth allocated to the lower priority session until enough bandwidth is accumulated to accommodate the high priority session.
7. The method of claim 6 wherein multiple alternative paths are generated for the primary path by:
 - considering in turn each node traversed by the primary path;
 - removing the crosslink in the primary path egressing from the above said Branch Node; and
 - determining a shortest path to the Destination Node starting from the Branch Node taking into account bandwidth reclaimed from all lower priority sessions that have been preempted thus far and preempting additional lower priority sessions as needed.

8. The method of claim 7 wherein the point-to-point route determination is extended for point-to-multipoint routing by:
- considering each destination node in some order, and computing a path to the Destination Node avoiding double counting of bandwidth allocated on common crosslinks in the paths for different destinations;
 - at the end of the initial iteration, the multipoint tree is the path determined for the first destination; and
 - merging the current path determined at each iteration with the multipoint tree which is the cumulative result of previous iterations by scanning the nodes on the current path starting from the current destination until a node N is found such that N is on the current path as well as in the multipoint tree; and attaching the subpath starting from N in the current path as a subtree of N in the multipoint tree.
9. The method of claim 8 wherein forwarding rules are generated for each payload for the routes determined for the primary path and alternative paths for a session by considering each snapshot spanned by the session, where a snapshot is an instance of the Grid for a specific time interval as predetermined by satellite orbits, and consists of a set of <Node, Satellite Id> mappings identifying which satellite is occupying which Node.
10. The method of claim 9 wherein the forwarding rule generation for a snapshot for an edge from a node N1 to a node N2 in a primary path or alternative path comprises:
- identifying the satellite A that occupies N1 in the snapshot;
 - identifying the satellite B that occupies N2 in the snapshot;
 - identifying the output port P in A for the crosslink connected to B in the snapshot;
 - identifying the time TS as being the later of either the start time of the session or the start time of the snapshot;
 - identifying the end time TE as being the earlier of the end time of the session or the end time of the snapshot; and

generating the forwarding rule for satellite A as <Source terminal address, Destination terminal address, TS, TE, p, P>, where p is the session priority.

11. The method of claim 10 wherein forwarding rules generated for the primary path and the alternative path branching at a payload for a snapshot are combined into a single forwarding rule by taking the output port element of each rule and constructing a single output sequence {P0, P1} where P0 is the Port element in the forwarding rule for the primary path and P1 is the Port element in the forwarding rule for the alternative path.
12. A system for determination of routes for a communication session in an LEO satellite network comprising of:
 - a LEO Satellite Network Route Planner that performs the computational aspects of the point-to-point and point-to-multipoint route determination.
 - a first database that stores satellite Grid data that contains the following information: geolocation of each node; coverage area of each node; and crosslink delay for each logical intra-plane crosslink and inter-plane crosslink;
 - a second database that stores the satellite network snapshots data in the form of <Node, Satellite Id> mappings for each time interval; and
 - a third database that stores session and route information for each session that is admitted into the satellite network.
13. A LEO satellite system comprising of:
 - a Network Operations Center that has the LEO Route Planning System that plans routes for user communication sessions, generates payload forwarding rules, disseminates the forwarding rules to the payloads via Distributed Management Centers, generates SLA satisfaction reports, and performs network trouble resolution;
 - one or more Distributed Management Centers that upload forwarding rules to the payloads within their domains, monitor the status and performance

of payloads within their domains, and report the status and performance data to the Network Operations Center; and LEO satellite payloads that forward data packets in accordance with these forwarding rules, dynamically reroute packets based on crosslink failure or congestion, and report their status and performance to designated management centers.

14. A non-transitory computer readable medium having computer readable program for operating on a computer for determining a primary path and multiple alternative paths for a point-to-point communication session between two ground terminals connected across by a Low Earth Orbit (LEO) satellite network, said program comprising instructions that cause the computer to perform the steps of:

determining routes using the following information:

- changes in terminal – satellite connectivity due to satellite orbits;
- changes in satellite – satellite connectivity due to satellite orbits;
- locations of Source and Destination terminals;
- communication session duration;
- session traffic bandwidth; and
- session traffic priority;

determining routes for the primary path and alternative paths based on a LEO Satellite Grid where the Grid represents a fixed logical satellite constellation topology comprising logical nodes representing geolocations, vertical edges representing logical intra-plane crosslinks, horizontal edges representing logical inter-plane crosslinks, and delay information on crosslinks; wherein the routes for the primary path and alternative paths are determined by computing shortest paths from a Source Node whose coverage area includes a Source terminal location to a Destination Node whose coverage area includes a Destination terminal location, taking into account available bandwidth on logical crosslinks for the duration of the communication session.

Page furnished blank

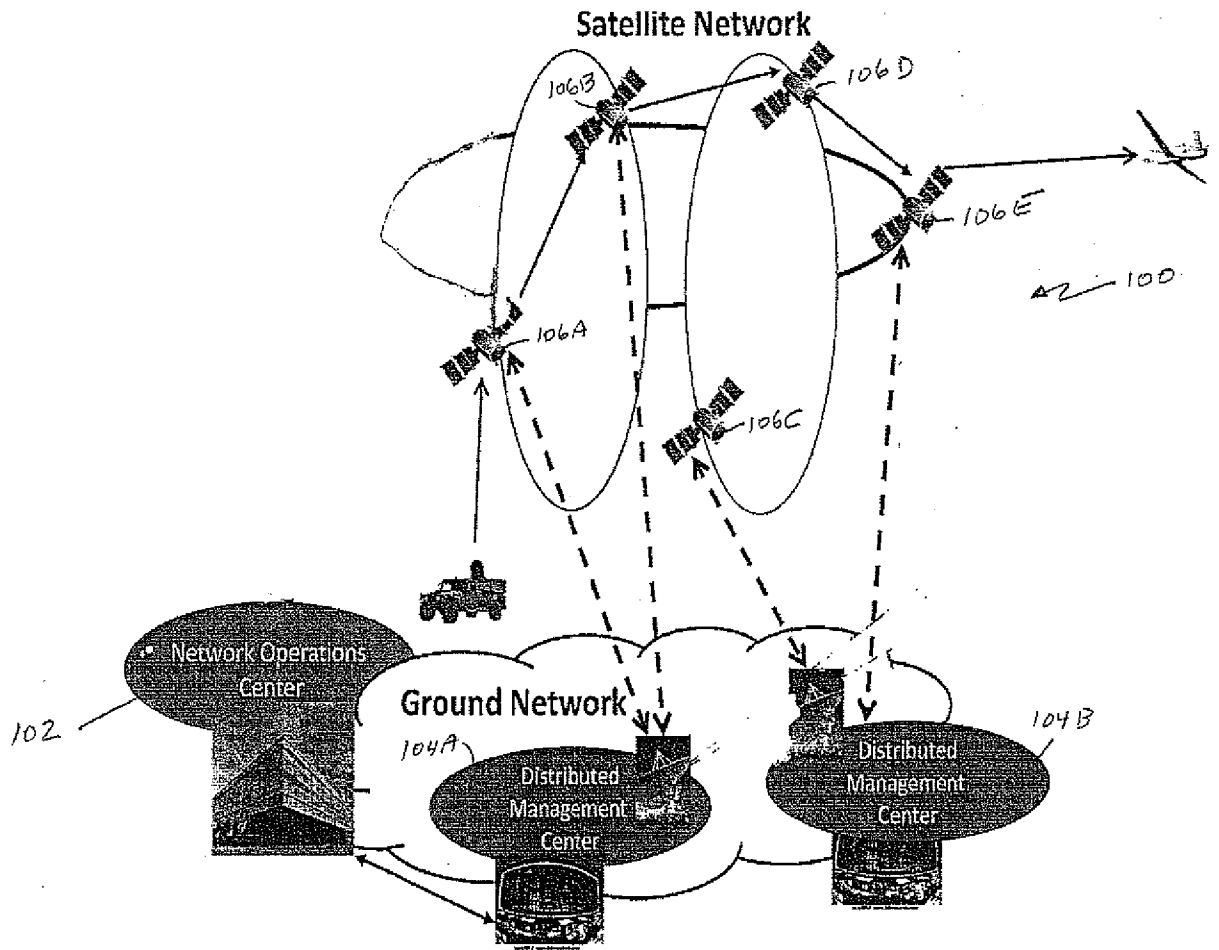


Figure 1.

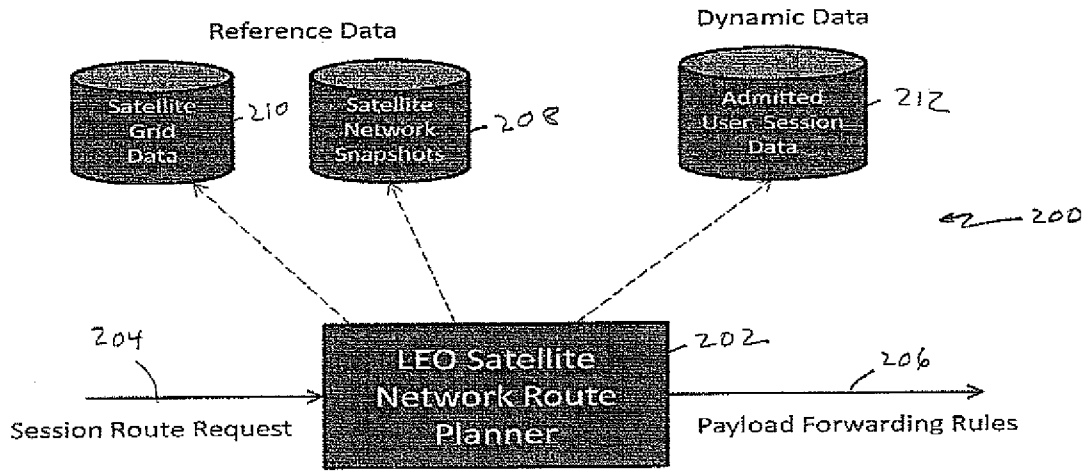


Figure 2.




Figure 3.

- Node (Geo-location)
- | Logical Crosslink (Intra-plane)
- Logical Crosslink (Inter-plane)

Flow name: Find Session Paths
Inputs: Source Node, Destination Node, Start Time, End Time, Bandwidth, Priority
Outputs: Success/Failure, Primary Path, Alternative Paths

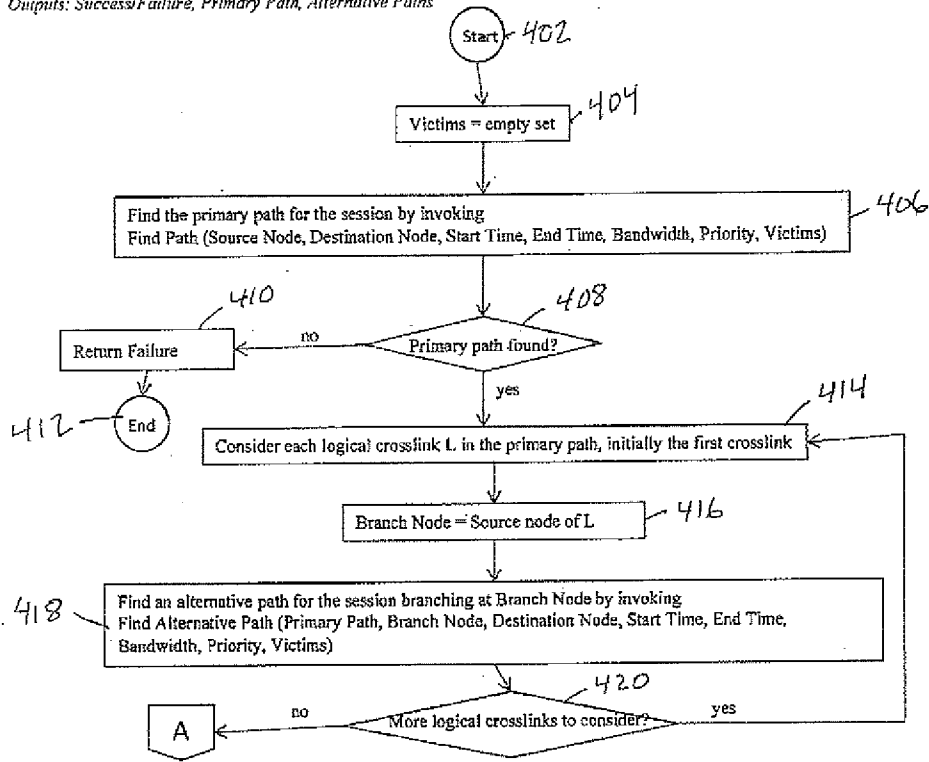


Figure 4A

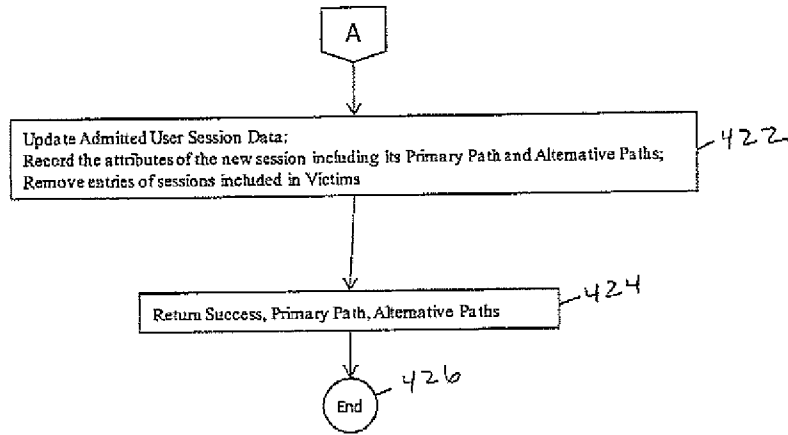


Figure 4B

Flow name: Find Path
 Inputs: Source Node, Destination Node, Start Time, End Time, Bandwidth, Priority, Victims
 Outputs: Success/Failure, Route, Victims

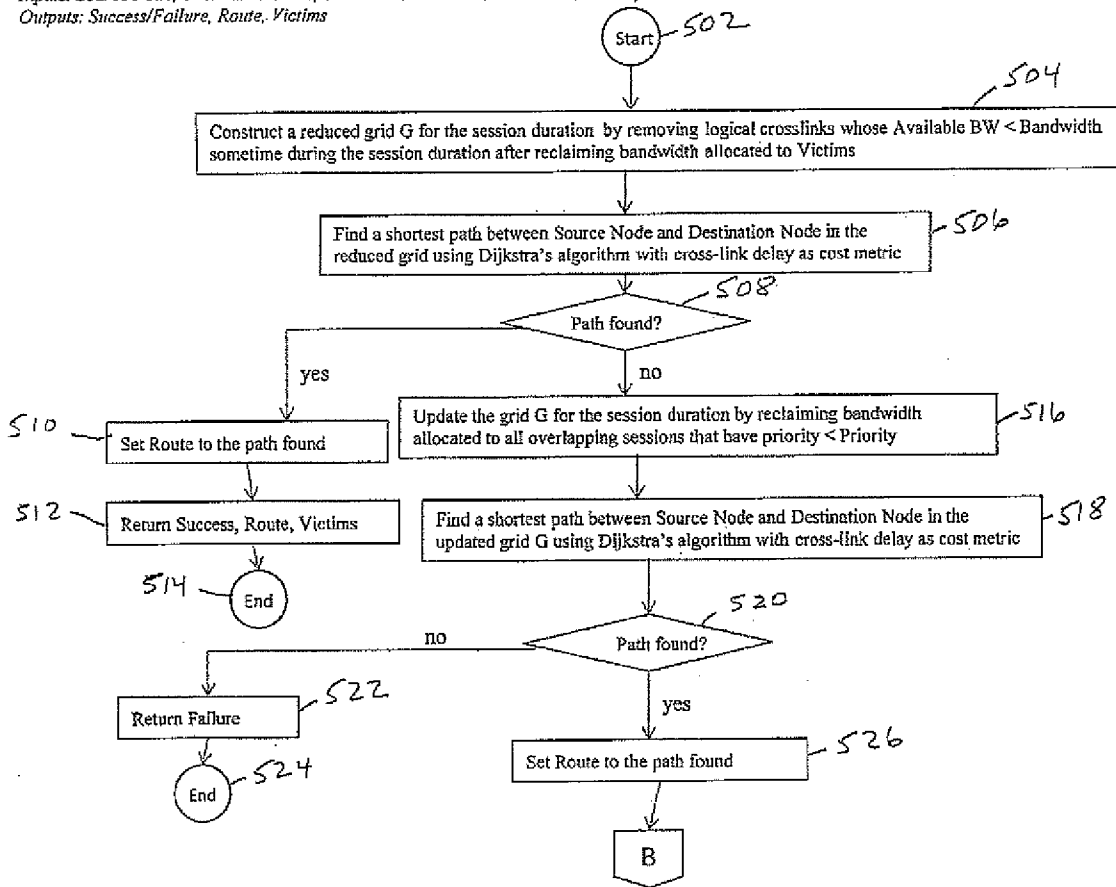


Figure 5A

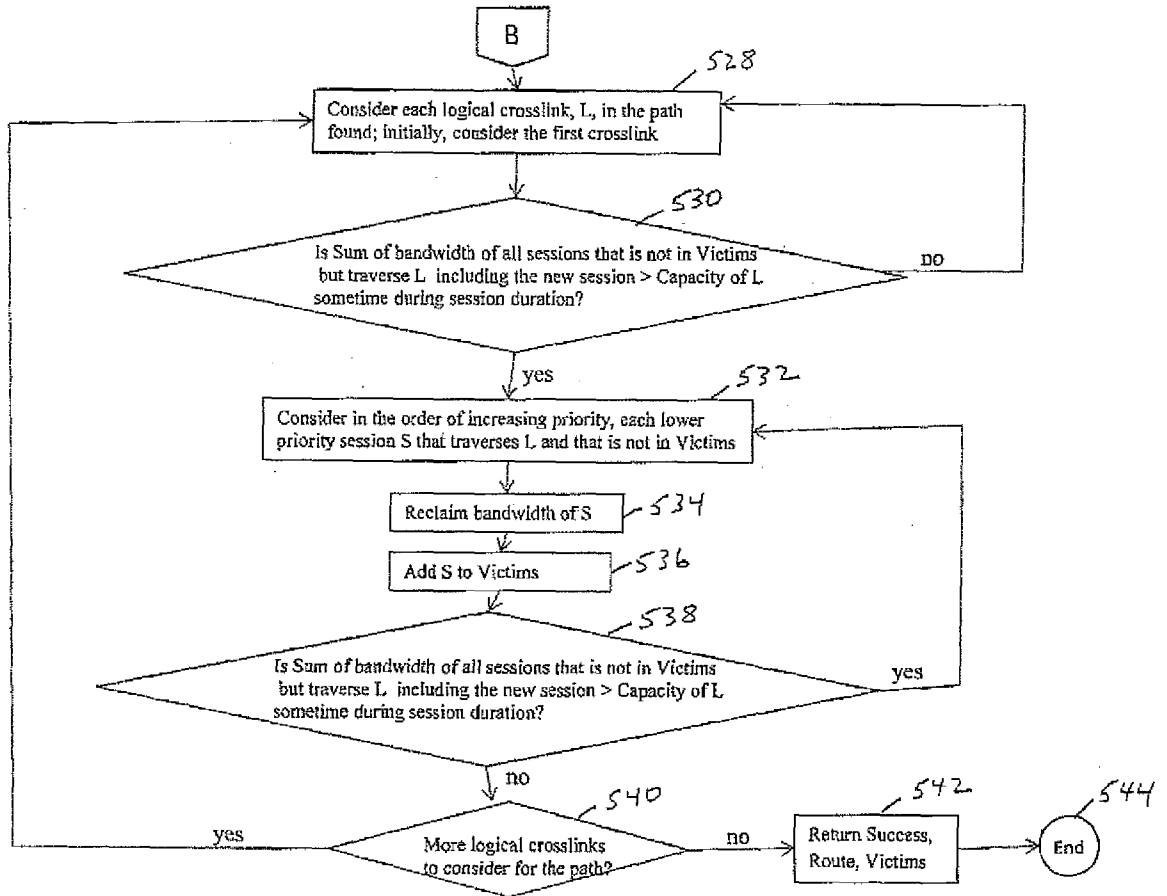


Figure 5B

FlowName: Find Alternative Path
Inputs: Primary Path, Branch Node, Destination Node, Start Time, End Time, Bandwidth, Priority, Victims
Outputs: Success/Failure, Route, Victims

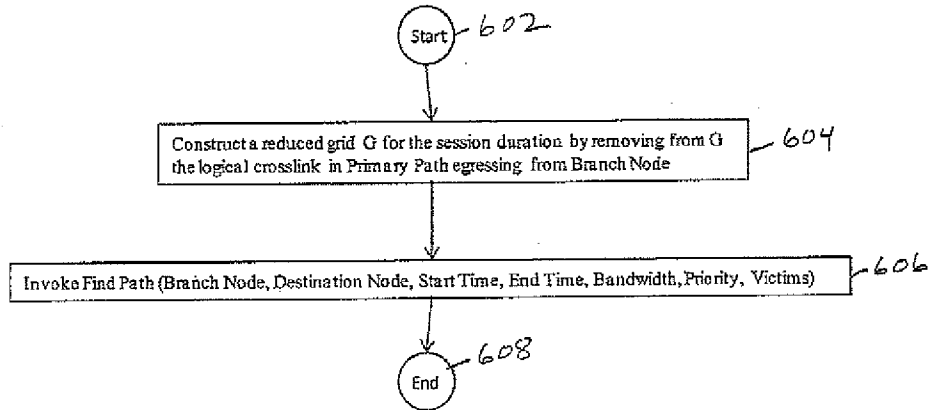


Figure 6

Flow name: Generate Forwarding Rules

Inputs: Source Terminal Address, Destination Terminal Address, Start Time, End Time, Priority, Primary Path, Alternative Paths
Outputs: Forwarding Rules

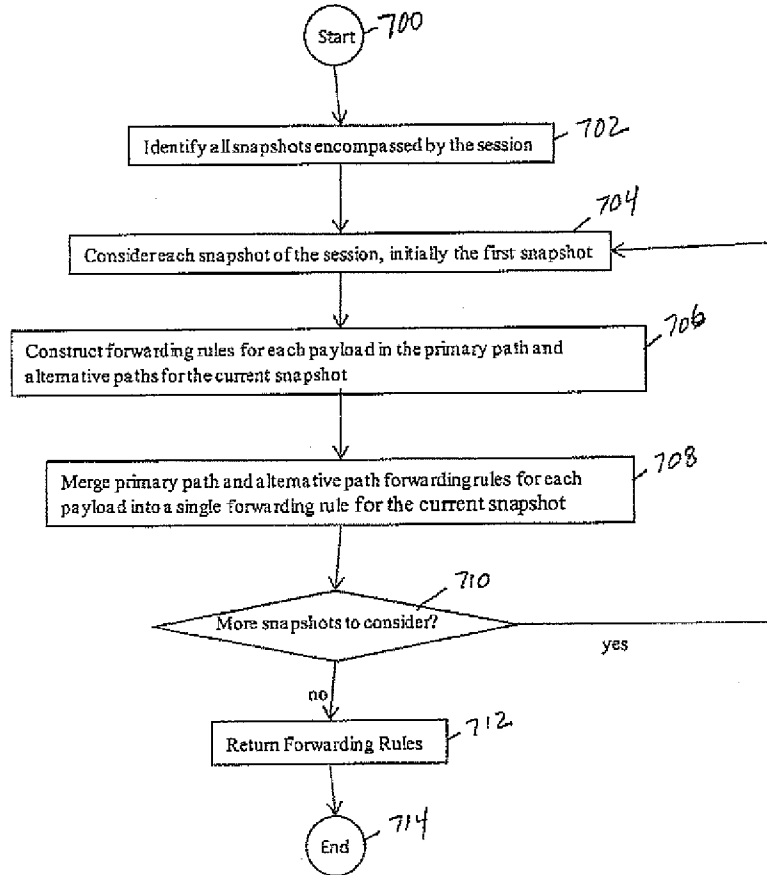


Figure 7