RFC 9717
A Routing Architecture for Satellite Networks

Abstract

Satellite networks present some interesting challenges for packet networking. The entire topology is continually in motion, with links far less reliable than what is common in terrestrial networks. Some changes to link connectivity can be anticipated due to orbital dynamics.

This document proposes a scalable routing architecture for satellite networks based on existing routing protocols and mechanisms that is enhanced with scheduled link connectivity change information. This document proposes no protocol changes.

This document presents the author's view and is neither the product of the IETF nor a consensus view of the community.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Satellite networks present some interesting challenges for packet networking. The entire topology is continually in motion, with links far less reliable than what is common in terrestrial networks. Some changes to link connectivity can be anticipated due to orbital dynamics.

This document proposes a scalable routing architecture for satellite networks based on existing routing protocols and mechanisms that is enhanced with scheduled link connectivity change information. This document proposes no protocol changes.

Large-scale satellite networks are being deployed, presenting an unforeseen application for conventional routing protocols. The high rate of intentional topological change and the extreme scale are unprecedented in terrestrial networking. Links between satellites can utilize free-space optics technology that allows liberal connectivity. Still, there are limitations due to the range of the links and conjunction with the sun, resulting in links that are far less reliable than network designers are used to. In addition, links can change their endpoints dynamically, resulting in structural changes to the topology.

Current satellite networks are proprietary, and little information is generally available for analysis and discussion. This document is based on what is currently accessible.

This document proposes one approach to provide a routing architecture for such networks utilizing current standards-based routing technology and to provide a solution for the scalability of the network while incorporating the rapid rate of topological change. This document intends to provide some initial guidance for satellite network operators, but without specific details, this document can only provide the basis for a more complete analysis and design.

This document presents the author's view and is neither the product of the IETF nor a consensus view of the community.

1.1. Related Work

A survey of related work can be found in [Westphal]. Link-state routing for satellite networks has been considered in [Cao] and [Zhang].

1.2. Terms and Abbreviations

Constellation: A set of satellites.

Downlink: The half of a ground link leading from a satellite to an Earth station.

Earth station: A node in the network that is on or close to the planetary surface and has a link to a satellite. This includes ships, aircraft, and other vehicles below LEO [ITU].

Gateway: An Earth station that participates in the network and acts as the interconnect between satellite constellations and the planetary network. Gateways have a much higher bandwidth than user stations, have ample computing capabilities, and perform traffic engineering duties, subsuming the functionality of a network controller or Path Computation Element (PCE) [RFC4655]. Multiple gateways are assumed to exist, and each serves a portion of the network.

GEO: Geostationary Earth Orbit. A satellite in GEO has an orbit that is synchronized to planetary rotation, so it effectively sits over one spot on the planet.

Ground link: A link between a satellite and an Earth station, composed of a downlink and an uplink.

IGP: Interior Gateway Protocol. A routing protocol that is used within a single administrative domain. Note that 'gateway' in this context is semantically equivalent to 'router' and has no relationship to the 'gateway' used in the rest of this document.

IS-IS: Intermediate System to Intermediate System. An IGP that is commonly used by service providers [ISO10589] [RFC1195].

ISL: Inter-Satellite Link. Frequently implemented with free-space optics that allow signaling using photons without any intervening medium [Bell].

L1:   IS-IS Level 1

L1L2:   IS-IS Level 1 and Level 2

L2:   IS-IS Level 2

LEO: Low Earth Orbit. A satellite in LEO has an altitude of 2,000 km or less.

Local gateway: Each user station is associated with a single gateway in its region.

LSP: Link State Protocol Data Unit. An IS-IS LSP is a set of packets that describe a node's connectivity to other nodes.

MEO: Medium Earth Orbit. A satellite in MEO is between LEO and GEO orbits and has an altitude between 2,000 km and 35,786 km.

SID: Segment Identifier [RFC8402]

Stripe: A set of satellites in a few adjacent orbits. These form an IS-IS L1 area.

SR: Segment Routing [RFC8402]

Uplink: The half of a link leading from an Earth station to a satellite.

User station: An Earth station interconnected with a small end-user network.

2. Overview

2.1. Topological Considerations

Satellites travel in specific orbits around their parent planet. Some of them have their orbital periods synchronized to planetary rotation, so they are effectively stationary over a single point. Other satellites have orbits that cause them to travel across regions of the planet either gradually or quite rapidly. Respectively, these are typically known as the Geostationary Earth Orbit (GEO), Medium Earth Orbit (MEO), or Low Earth Orbit (LEO), depending on the altitude. This discussion is not Earth-specific; as we get to other planets, we can test this approach's generality.

Satellites may have data interconnections with one another through Inter-Satellite Links (ISLs). Due to differences in orbits, ISLs may be connected temporarily with periods of potential connectivity computed through orbital dynamics. Multiple satellites may be in the same orbit but separated in space with a roughly constant separation. Satellites in the same orbit may have ISLs that have a higher duty cycle than ISLs between different orbits, but they are still not guaranteed to be always connected.

```
  User             +-----------------+    Local
  Stations   --- |   Satellites    |--- Gateway --- Internet
                   +-----------------+
```

Figure 1: Overall Network Architecture

Earth stations can communicate with one or more satellites in their region. User stations are Earth stations with a limited capacity that communicate with only a single satellite at a time. Other Earth stations that may have richer connectivity and higher bandwidth are commonly called "gateways" and provide connectivity between the satellite network and conventional wired networks. Gateways serve user stations in their geographic proximity and are replicated globally as necessary to provide coverage and to meet service density goals. User stations are associated with a single local gateway. Traffic from one Earth station to another may need to traverse a path across multiple satellites via ISLs.

2.2. Link Changes

Like conventional network links, ISLs and ground links can fail without warning. However, unlike terrestrial links, there are predictable times when ISLs and ground links can potentially connect and disconnect. These predictions can be computed and cataloged in a schedule that can be distributed to relevant network elements. Predictions of a link connecting are not guaranteed: A link may not connect for many reasons. Link disconnection predictions due to orbital dynamics are effectively guaranteed, as the underlying physics will not improve unexpectedly.

2.3. Scalability

Some proposed satellite networks are fairly large, with tens of thousands of proposed satellites [CNN]. A key concern is the ability to reach this scale and larger, as useful networks tend to grow.

As we know, the key to scalability is the ability to create hierarchical abstractions, so a key question of any routing architecture will be about the abstractions that can be created to contain topological information.

Normal routing protocols are architected to operate with a static but somewhat unreliable topology. Satellite networks lack the static organization of terrestrial networks, so normal architectural practices for scalability may not apply, and alternative approaches may need consideration.

In a traditional deployment of a link-state routing protocol, current implementations can be deployed with a single area that spans a few thousand routers. A single area would also provide no isolation for topological changes, causing every link change to be propagated throughout the entire network. This would be insufficient for the needs of large satellite networks.

Multiple areas or multiple instances of an Interior Gateway Protocol (IGP) can be used to improve scalability, but there are limitations to traditional approaches.

Currently, the IETF actively supports two link-state IGPs: OSPF [RFC2328] [RFC5340] and IS-IS.

OSPF requires that the network operate around a backbone area, with subsidiary areas hanging off of the backbone. While this works well for traditional terrestrial networks, this does not seem appropriate for satellite networks, where there is no centralized portion of the topology.

IS-IS has a different hierarchical structure, where Level 1 (L1) areas are connected sets of nodes, and then Level 2 (L2) is a connected subset of the topology that intersects all of the L1 areas. Individual nodes can be L1, L2, or both (L1L2). Traditional IS-IS designs require that any node or link that is to be used as transit between L2 areas must appear as part of the L2 topology. In a satellite network, any satellite could end up being used for L2 transit, and so every satellite and link would be part of L2, negating any scalability benefits from IS-IS's hierarchical structure.

We elaborate on considerations specific to IS-IS in Section 4.

2.4. Assumptions

In this section, we discuss some of the assumptions that are the basis for this architectural proposal.

The data payload is IP packets.

Satellites are active participants in the control and data planes for the network, participating in protocols and forwarding packets.

There may be a terrestrial network behind each gateway that may interconnect to the broader Internet. The architecture of the terrestrial network is assumed to be a typical IS-IS and BGP deployment [RFC4271] and is not discussed further in this document.

The satellite network interconnects user stations and gateways. Interconnection between the satellite network and the satellite networks of other network operators is outside the scope of this document.

2.4.1. Traffic Patterns

We assume that the primary use of the satellite network is to provide access from a wide range of geographic locations. We also assume that providing high-bandwidth bulk transit between peer networks is not a goal. It has been noted that satellite networks can provide lower latencies than terrestrial fiber networks in [Handley]. This proposal does not preclude such applications but does not articulate the mechanisms necessary for user stations to perform the appropriate traffic engineering computations. Low-latency, multicast, and anycast applications are not discussed further in this document.

As with most access networks, we assume that there will be bidirectional traffic between the user station and the gateway but that the bulk of the traffic will be from the gateway to the user station. We expect the uplink from the gateway to the satellite network to be the bandwidth bottleneck and that gateways will need to be replicated to scale the uplink bandwidth, as the satellite capacity reachable from a gateway will be limited.

We assume that it is not essential to provide optimal routing for traffic from user station to user station. If this traffic is sent to a gateway first and then back into the satellite network, it might be acceptable to some operators as long as the traffic volume remains very low. This type of routing is not discussed further in this document.

We assume that traffic for a user station should enter the satellite network through a gateway that is in some close geographic proximity to the user station. This is to reduce the number of ISLs used by the path to the user station. Similarly, we assume that user station traffic should exit the satellite network through the gateway that is in the closest geographic proximity to the user station. Jurisdictional requirements for landing traffic in certain regions may alter these assumptions, but such situations are outside the scope of this document.

This architecture does not preclude gateway-to-gateway traffic across the satellite constellations, but it does not seek to optimize it.

2.4.2. User Station Constraints

The user station is an entity whose operation is conceptually shared between the satellite constellation operator and the operator of the cluster of end stations it serves. For example, the user station is trusted to attach MPLS label stacks to end-user packets. It gets the information to do so from some combination of its direct satellite and its local gateway via protocols outside the scope of this document. Equally, it bootstraps communication via an exchange with the current local satellite so that it can find and communicate with its local gateway -- again with the details of how that is done being outside the scope of this document.

User stations that can concurrently access multiple satellites are not precluded by this proposal but are not discussed in detail.

2.4.3. Stochastic Connectivity

We assume that links in general will be available when scheduled. As with any network, there will be failures, and the schedule is not a guarantee, but we also expect that the schedule is mostly accurate. We assume that at any given instant, there are enough working links and aggregate bandwidth to run the network and support the traffic demand. If this assumption does not hold, no routing architecture can magically make the network more capable.

Satellites that are in the same orbit may be connected by ISLs. These are called "intra-orbit" ISLs. Satellites that are in different orbits may also be connected by ISLs. These are called "inter-orbit" ISLs. Generally, we assume that intra-orbit ISLs have higher reliability and persistence than inter-orbit ISLs.

We assume that the satellite network is connected (in the graph theory sense) almost always, even if some links are down. This implies that there is almost always some path to the destination. In the extreme case with no such path, we assume that it is acceptable to drop the payload packets. We do not require buffering of traffic when a link is down. Instead, traffic should be rerouted.

2.5. Problem Statement

The goal of the routing architecture is to provide an organizational structure to protocols running on the satellite network such that topology information is conveyed through relevant portions of the network, paths are computed across the network, and data can be delivered along those paths so that the structure can scale without any changes to the organizational structure.

3. Forwarding Plane

The end goal of a network is to deliver traffic. In a satellite network where the topology is in a continual state of flux and the user stations frequently change their association with the satellites, having a highly flexible and adaptive forwarding plane is essential. Toward this end, we propose using MPLS as the fundamental forwarding plane architecture [RFC3031]. Specifically, we propose using an approach based on Segment Routing (SR) [RFC8402] with an MPLS data plane [RFC8660], where each satellite is assigned a node Segment Identifier (SID). This allows the architecture to support both IPv4 and IPv6 concurrently. A path through the network can then be expressed as a label stack of node SIDs. IP forwarding is not used within the internals of the satellite network, although each satellite may be assigned an IP address for management purposes. Existing techniques may be used to limit the size of the SR label stack so that it only contains the significant waypoints along the path [Giorgetti]. The label stack operates as a loose source route through the network. If there is an unexpected topology change in the network, the IGP will compute a new path to the next waypoint, allowing packet delivery despite ISL failures. While the IGP is converging, there may be micro-loops in the topology. These can be avoided by using Topology Independent Loop-Free Alternate (TI-LFA) paths [SR-TI-LFA]; otherwise, traffic will loop until discarded based on its TTL.

We assume that there is a link-layer mechanism for a user station to associate with a satellite. User stations will have an IP address assigned from a prefix managed by its local gateway. The mechanisms for this assignment and its communication to the end station are not discussed herein but might be similar to DHCP [RFC2131]. User station IP addresses change infrequently and do not reflect their association with their first-hop satellite. Gateways and their supporting terrestrial networks advertise prefixes covering all its local user stations into the global Internet.

User stations may be assigned a node SID, in which case MPLS forwarding can be used for all hops to the user station. Alternatively, if the user station does not have a node SID, then the last hop from the satellite to the end station can be performed based on the destination IP address of the packet. This does not require a full longest-prefix-match lookup, as the IP address is merely a unique identifier at this point.

Similarly, gateways may be assigned a node SID. A possible optimization is that a single SID value could be assigned as a global constant to always direct traffic to the topologically closest gateway. If traffic engineering is required for traffic that is flowing to a gateway, a specific path may be encoded in a label stack that is attached to the packet by the user station or by the first-hop satellite.

Gateways can also perform traffic engineering using different paths and label stacks for separate traffic flows. Routing a single traffic flow across multiple paths has proven to cause performance issues with transport protocols, so that approach is not recommended. Traffic engineering is discussed further in Section 6.

4. IGP Suitability and Scalability

As discussed in Section 2.3, IS-IS is architecturally the best fit for satellite networks but does require some novel approaches to achieve the scalability goals for a satellite network. In particular, we propose that all nodes in the network be L1L2 so that local routing is done based on L1 information and then global routing is done based on L2 information.

An interesting property of IS-IS is that it does not require interface addresses. This feature is commonly known as "unnumbered interfaces". This is particularly helpful in satellite topologies because it implies that ISLs may be used flexibly. Sometimes an interface might be used as an L1 link to another satellite, and a few orbits later, it might be used as an L1L2 link to a completely different satellite without any reconfiguration or renumbering.

Scalability for IS-IS can be achieved through a proposal known as "Area Proxy" [RFC9666]. With this proposal, all nodes in an L1 area combine their information into a single L2 Link State Protocol Data Unit (LSP). This implies that the size of the L1 Link State Database (LSDB) scales as the number of nodes in the L1 area and the size of the L2 LSDB scales with the number of L1 areas.

With Area Proxy, topological changes within an L1 area will not be visible to other areas, so the overhead of link-state changes will be greatly reduced.

The Area Proxy proposal also includes the concept of an Area SID. This is useful because it allows traffic engineering to construct a path that traverses areas with a minimal number of label stack entries.

For example, suppose that a network has 1,000 L1 areas, each with 1,000 satellites. This would mean that the network supports 1,000,000 satellites but only requires 1,000 entries in its L1 LSDB and 1,000 entries in its L2 LSDB, which are easily achievable numbers today. The resulting MPLS label table would contain 1,000 node SIDs from the L1 (and L2) LSDB and 1,000 area SIDs from the L2 LSDB. If each satellite advertises an IP address for management purposes, then the IP routing table would have 1,000 entries for the L1 management addresses and 1,000 area proxy addresses from L2.

In this proposal, IS-IS does not carry IP routes other than those in the satellite topology. In particular, there are no IP routes for user stations or the remainder of the Internet.

5. Stripes and Areas

A significant problem with any link-state routing protocol is that of area partition. While there have been many proposals for automatic partition repair, none has seen notable production deployment. It seems best to avoid this issue and ensure areas have an extremely low probability of partitioning.

As discussed above, intra-orbit ISLs are assumed to have higher reliability and persistence than inter-orbit ISLs. However, even intra-orbit ISLs are not sufficiently reliable to avoid partition issues. Therefore, we propose to group a small number of adjacent orbits as an IS-IS L1 area, called a "stripe". We assume that for any given reliability requirement, there is a small number of orbits that can be used to form a stripe that satisfies the reliability requirement.

Stripes are connected to other adjacent stripes using the same ISL mechanism, forming the L2 topology of the network. Each stripe should have multiple L2 connections and never become partitioned from the remainder of the network.

By using a stripe as an L1 area, in conjunction with Area Proxy, the overhead of the architecture is greatly reduced. Each stripe contributes a single LSP to the L2 LSDB, completely hiding all the details about the satellites within the stripe. The resulting architecture scales proportionately to the number of stripes required, not the number of satellites.

Groups of MEO and GEO satellites with interconnecting ISLs can also form an IS-IS L1L2 area. Satellites that lack intra-constellation ISLs are better as independent L2 nodes.

6. Traffic Forwarding and Traffic Engineering

Forwarding in this architecture is straightforward. A path from a gateway to a user station on the same stripe only requires a single node SID for the satellite that provides the downlink to the user station.

```
            \
Gateway -->  X
              \
               \
                X
                 \
                  \
                   X ---> x  User Station
                    \
```

Figure 2: On-Stripe Forwarding

Similarly, a user station returning a packet to a gateway need only provide a gateway node SID.

For off-stripe forwarding, the situation is a bit more complex. A gateway would need to provide the area SID of the final stripe on the path plus the node SID of the downlink satellite. For return traffic, user stations or first-hop satellites would want to provide the area SID of the stripe that contains the satellite that provides access to the gateway as well as the gateway SID.
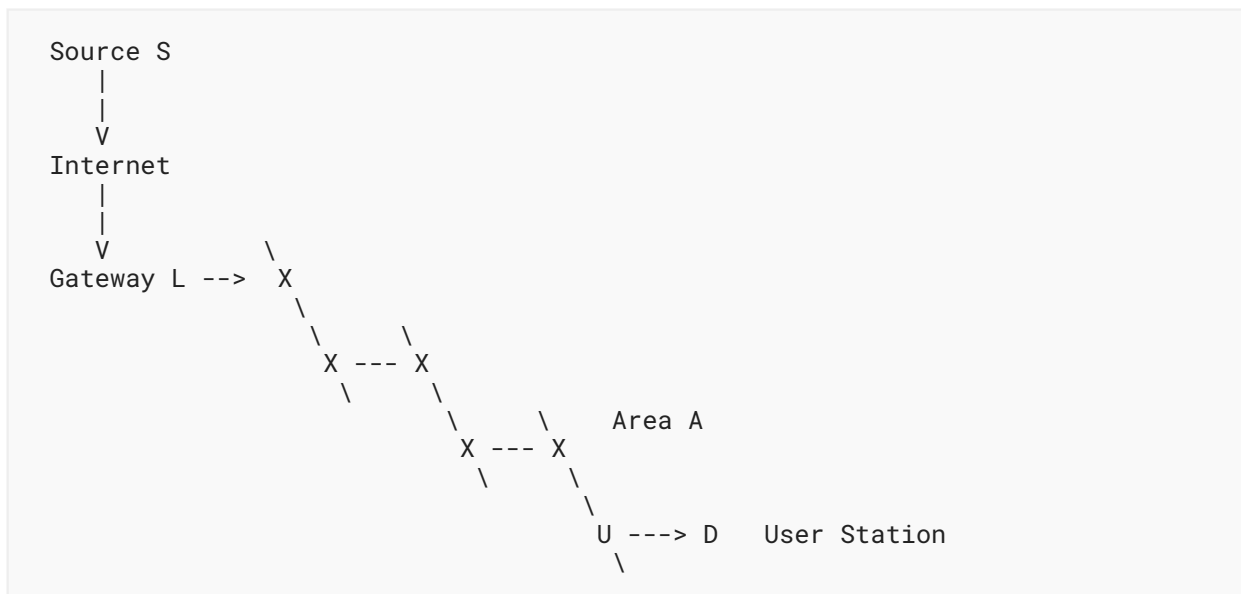
```
    Source S
       |
       |
       V
    Internet
       |
       |
       V           \
    Gateway L -->   X
                     \
                      \     \
                       X --- X
                        \     \
                         \     \     Area A
                          X --- X
                           \     \
                            \     \
                             U ---> D   User Station
                              \
```

Figure 3: Off-Stripe Forwarding

As an example (Figure 3), consider a packet from an Internet source (Source S) to a user station (D). A local gateway (Gateway L) has injected a prefix covering D into BGP and has advertised it globally. The packet is forwarded to L using IP forwarding. When L receives the packet, it performs a lookup in a pre-computed forwarding table. This contains a SID list for the user station that has already been converted into a label stack. Suppose the user station is currently associated with a different stripe so that the label stack will contain an area label (A) and a label (U) for the satellite associated with the user station, resulting in a label stack (A, U).

The local gateway forwards this into the satellite network. The first-hop satellite now forwards based on the area label (A) at the top of the stack. All area labels are propagated as part of the L2 topology. This forwarding continues until the packet reaches a satellite adjacent to the destination area. That satellite pops label A, removing that label and forwarding the packet into the destination area.

The packet is now forwarded based on the remaining label U, which was propagated as part of the L1 topology. The last satellite forwards the packet based on the destination address (D) and forwards the packet to the user station.

The return case is similar. The label stack, in this case, consists of a label for the local gateway's stripe/area (A') and the label for the local gateway (L), resulting in the stack (A', L). The forwarding mechanisms are similar to the previous case.

Very frequently, access networks congest due to over-subscription and the economics of access. Network operators can use traffic engineering to ensure that they get higher efficiency out of their networks by utilizing all available paths and capacity near any congestion points. In this particular case, the gateway will have information about all of the traffic it is generating and can use all of the possible paths through the network in its topological neighborhood. Since we're already using SR, this is easily done by adding more explicit SIDs to the label stack. These can be additional area SIDs, node SIDs, or adjacency SIDs. Path computation can be performed by Path Computation Elements (PCEs) [RFC4655].

Each gateway or its PCE will need topological information from the areas it will route through. It can do this by participating in the IGP directly, via a tunnel, or through another delivery mechanism such as BGP-LS [RFC9552]. User stations do not participate in the IGP.

Traffic engineering for packets flowing into a gateway can also be provided by an explicit SR path. This can help ensure that ISLs near the gateway do not congest with traffic for the gateway. These paths can be computed by the gateway or PCE and then distributed to the first-hop satellite or user station, which would apply them to traffic. The delivery mechanism is outside the scope of this document.

7. Scheduling

The most significant difference between terrestrial and satellite networks from a routing perspective is that some of the topological changes that will happen to the network can be anticipated and computed. Both link and node changes will affect the topology, and the network should react smoothly and predictably.

The management plane is responsible for providing information about scheduled topological changes. The exact details of how the information is disseminated are outside the scope of this document but could be shown through a YANG model [YANG-SCHEDULE]. Scheduling information needs to be accessible to all of the nodes that will make routing decisions based on the topological changes in the schedule (i.e., data about an L1 topological change will need to be circulated to all nodes in the L1 area and information about L2 changes will need to propagate to all L2 nodes) and to the gateways and PCEs that carry the related topological information.

There is very little that the network should do in response to a topological addition. A link coming up or a node joining the topology should not have any functional change until the change is proven to be fully operational based on the usual IS-IS liveness mechanisms. Nodes may pre-compute their routing table changes but should not install them before all relevant adjacencies are received. The benefits of this pre-computation appear to be very small. Gateways and PCEs may also choose to pre-compute paths based on these changes but should not install paths using the new parts of the topology until they are confirmed to be operational. If some path pre-installation is performed, gateways and PCEs must be prepared for the situation where the topology fails to become operational and may need to take alternate steps instead, such as reverting any related pre-installed paths.

The network may choose not to pre-install or pre-compute routes in reaction to topological additions, at a small cost of some operational efficiency.

Topological deletions are an entirely different matter. If a link or node is to be removed from the topology, the network should act before the anticipated change to route traffic around the expected topological loss. Specifically, at some point before the topology change, the affected links should be set to a high metric to direct traffic to alternate paths. This is a common operational procedure in existing networks when links are taken out of service, such as when proactive maintenance needs to be performed. This type of change does require some time to propagate through the network, so the metric change should be initiated far enough in advance that the network converges before the actual topological change. Gateways and PCEs should also update paths around the topology change and install these changes before the topology change occurs. The time necessary for both IGP and path changes will vary depending on the exact network and configuration.

Strictly speaking, changing to a high metric should not be necessary. It should be possible for each router to exclude the link and recompute paths. However, it seems safer to change the metric and use the IGP methods for indicating a topology change, as this can help avoid issues with incomplete information dissemination and synchronization.

8. Future Work

This architecture needs to be validated by satellite operators, both via simulation and operational deployment. Meaningful simulation hinges on the exact statistics of ISL connectivity; currently, that information is not publicly available.

Current available information about ISLs indicates that links are mechanically steered and will need to track the opposite end of the link continually. The angles and distances that can be practically supported are unknown, as are any limitations about the rate of change.

It is expected that intra-orbit and inter-orbit ISL links will have very different properties. Intra-orbit links should be much more stable but still far less stable than terrestrial links. Inter-orbit links will be less stable. Links between satellites that are roughly parallel should be possible but will likely have a limited duration. Two orbits may be roughly orthogonal, resulting in a limited potential for connectivity. Finally, in some topologies there may be parallel orbits where the

satellites move in opposite directions, giving a relative speed between satellites around 34,000 mph (55,000 kph). Links in this situation may not be possible or may be so short-lived that they are impractical.

The key question to address is whether the parameters of a given network can yield a stripe assignment that produces stable, connected areas that work within the scaling bounds of the IGP. If links are very stable, a stripe could be just a few parallel orbits, with only a few hundred satellites. However, if links are unstable, a stripe might have to encompass dozens of orbits and thousands of satellites, which might be beyond the scaling limitations of a given IGP's implementation.

9. Deployment Considerations

The network behind a gateway is expected to be a normal terrestrial network. Conventional routing architectural principles apply. An obvious approach would be to extend IS-IS to the terrestrial network, applying L1 areas as necessary for scalability.

The terrestrial network may have one or more BGP connections to the broader Internet. Prefixes for user stations should be advertised to the Internet near the associated gateway. If gateways are not interconnected by the terrestrial network, then it may be advisable to use one autonomous system per gateway as it might simplify the external perception of the network and subsequent policy considerations. Otherwise, one autonomous system may be used for the entire terrestrial network.

10. Security Considerations

This document discusses one possible routing architecture for satellite networks. It proposes no new protocols or mechanisms and thus has no new security impact. Security for IS-IS is provided by [RFC5304] and [RFC5310].

User stations will interact directly with satellites, potentially using proprietary mechanisms, and under the control of the satellite operator, who is responsible for the security of the user station.

11. IANA Considerations

This document has no IANA actions.

12. References

12.1. Normative References

[ISO10589]   ISO/IEC, "Information technology - Telecommunications and information exchange between systems - Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the

protocol for providing the connectionless-mode network service (ISO 8473)",
ISO/IEC 10589:2002, November 2002, <https://www.iso.org/standard/
30932.html>.

[RFC3031]   Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching
Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <https://www.rfc-
editor.org/info/rfc3031>.

[RFC5304]   Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI
10.17487/RFC5304, October 2008, <https://www.rfc-editor.org/info/rfc5304>.

[RFC5310]   Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic
Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009,
<https://www.rfc-editor.org/info/rfc5310>.

[RFC8402]   Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R.
Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July
2018, <https://www.rfc-editor.org/info/rfc8402>.

[RFC8660]   Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R.
Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/
RFC8660, December 2019, <https://www.rfc-editor.org/info/rfc8660>.

[RFC9666]   Li, T., Chen, S., Ilangovan, V., and G. Mishra, "Area Proxy for IS-IS", RFC 9666, DOI
10.17487/RFC9666, October 2024, <https://www.rfc-editor.org/info/rfc9666>.

12.2. Informative References

[Bell]   Bell, A. G., "On the Production and Reproduction of Sound by Light", American
Journal of Science, vol. S3-20, no. 118, pp. 305-324, DOI 10.2475/ajs.s3-20.118.305,
October 1880, <https://ajsonline.org/article/64037>.

[Cao]   Cao, X., Li, Y., Xiong, X., and J. Wang, "Dynamic Routings in Satellite Networks:
An Overview", Sensors, vol. 22, no. 12, pp. 4552, DOI 10.3390/s22124552, 2022,
<https://www.mdpi.com/1424-8220/22/12/4552/pdf?version=1655449925>.

[CNN]   Wattles, J., "Elon Musk's SpaceX now owns about a third of all active satellites in
the sky", CNN Business, 11 February 2021, <https://www.cnn.com/2021/02/11/
tech/spacex-starlink-satellites-1000-scn/index.html>.

[Giorgetti] Giorgetti, A., Castoldi, P., Cugini, F., Nijhof, J., Lazzeri, F., and G. Bruno, "Path
Encoding in Segment Routing", 2015 IEEE Global Communications Conference
(GLOBECOM), DOI 10.1109/GLOCOM.2015.7417097, December 2015, <https://
ieeexplore.ieee.org/document/7417097>.

[Handley]   Handley, M., "Delay is Not an Option: Low Latency Routing in Space", HotNets
'18: Proceedings of the 17th ACM Workshop on Hot Topics in Networks, pp.
85-91, DOI 10.1145/3286062.3286075, November 2018, <https://dl.acm.org/doi/
10.1145/3286062.3286075#>.

[ITU]   ITU, "Radio Regulations - Articles", 2016, <https://search.itu.int/history/HistoryDigitalCollectionDocLibrary/1.43.48.en.101.pdf>.

[RFC1195]   Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <https://www.rfc-editor.org/info/rfc1195>.

[RFC2131]   Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC2328]   Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC4271]   Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC4655]   Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC5340]   Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC9552]   Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <https://www.rfc-editor.org/info/rfc9552>.

[SR-TI-LFA]   Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-lfa-19, 22 November 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-19>.

[Westphal]   Westphal, C., Han, L., and R. Li, "LEO Satellite Networking Relaunched: Survey and Current Research Challenges", arXiv:2310.07546v1, DOI 10.48550/arXiv.2310.07646, October 2023, <https://arxiv.org/pdf/2310.07646.pdf>.

[YANG-SCHEDULE] Qu, Y., Lindem, A., Kinzie, E., Fedyk, D., and M. Blanchet, "YANG Data Model for Scheduled Attributes", Work in Progress, Internet-Draft, draft-ietf-tvr-schedule-yang-03, 20 October 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-tvr-schedule-yang-03>.

[Zhang]   Zhang, X., Yang, Y., Xu, M., and J. Luo, "ASER: Scalable Distributed Routing Protocol for LEO Satellite Networks", 2021 IEEE 46th Conference on Local Computer Networks (LCN), DOI 10.1109/LCN52139.2021.9524989, 2021, <https://doi.org/10.1109/LCN52139.2021.9524989>.

Acknowledgements

Author's Address

Tony Li
Juniper Networks
Email: tony.li@tony.li