
Stream: Internet Engineering Task Force (IETF)
RFC: [9657](#)
Category: Informational
Published: September 2024
ISSN: 2070-1721
Authors: E. Birrane, III N. Kuhn Y. Qu
JHU/APL Thales Alenia Space Futurewei Technologies
R. Taylor L. Zhang
Aalyria Technologies Huawei

RFC 9657

Time-Variant Routing (TVR) Use Cases

Abstract

This document introduces use cases where Time-Variant Routing (TVR) computations (i.e., routing computations that take into consideration time-based or scheduled changes to a network) could improve routing protocol convergence and/or network performance.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9657>.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Resource Preservation	3
2.1. Assumptions	4
2.2. Routing Impacts	4
2.3. Example	5
3. Operating Efficiency	6
3.1. Assumptions	6
3.2. Routing Impacts	7
3.3. Example: Cellular Network	7
3.4. Another Example: Tidal Network	9
4. Dynamic Reachability	9
4.1. Assumptions	10
4.2. Routing Impacts	10
4.3. Example: Mobile Satellites	11
4.4. Another Example: Predictable Moving Vessels	13
5. Security Considerations	13
6. IANA Considerations	13
7. Informative References	13
Acknowledgments	14
Authors' Addresses	14

1. Introduction

There is a growing number of use cases where changes to the routing topology are an expected part of network operations. In these use cases, the pre-planned loss and restoration of an adjacency, or formation of an alternate adjacency, should be seen as a nondisruptive event.

Expected changes to topologies can occur for a variety of reasons. In networks with mobile nodes, such as unmanned aerial vehicles and some orbiting spacecraft constellations, links are lost and re-established as a function of the mobility of the platforms. In networks without

reliable access to power, such as networks harvesting energy from wind and solar, link activity might be restricted to certain times of day. Similarly, in networks prioritizing green computing and energy efficiency over data rate, network traffic might be planned around energy costs or expected user data volumes.

This document defines three categories of use cases where a route computation might beneficially consider time information. Each of these use cases are included as follows:

1. An overview of the use case describing how route computations might select different paths (or subpaths) as a function of time.
2. A set of assumptions made by the use case as to the nature of the network and data exchange.
3. Specific discussion on the routing impacts of the use case.
4. Example networks conformant to the use case.

The use cases that are considered in this document are as follows:

1. Resource Preservation (described in [Section 2](#)), where there is information about link availability over time at the client level. Time-Variant Routing (TVR) can utilize the predictability of the link availability to optimize network connectivity by taking into account endpoint resource preservation.
2. Operating Efficiency (described in [Section 3](#)), where there is a server cost or a path cost usage varying over time. TVR can exploit the predictability of the path cost to optimize the cost of the system exploitation. The notion of a path cost is extended to be a time-dependent function instead of a constant.
3. Dynamic Reachability (described in [Section 4](#)), where there is information about link availability variation between nodes in the end-to-end path. TVR can exploit the predictability of the link availability to optimize in-network routing.

The document does not intend to represent the full set of cases where TVR computations could beneficially impact network performance -- new use cases are expected to be generated over time. Similarly, the concrete examples within each use case are meant to provide an existence proof of the use case and not to present any exhaustive enumeration of potential examples. It is likely that multiple example networks exist that could be claimed as instances of any given use case.

The document focuses on deterministic scenarios. Non-deterministic scenarios, such as vehicle-to-vehicle communication, are out of the scope of the document.

2. Resource Preservation

Some nodes in a network might operate in resource-constrained environments or otherwise with limited internal resources. Constraints, such as available power, thermal ranges, and on-board storage, can all impact the instantaneous operation of a node. In particular, resource management on such a node can require that certain functionality be powered on (or off) to extend the ability of the node to participate in the network.

When power on a node is running low, noncritical functions on the node might be turned off in favor of extending node life. Alternatively, certain functions on a node may be turned off to allow the node to use available power to respond to an event, such as data collection. When a node is in danger of violating a thermal constraint, normal processing might be paused in favor of a transition to a thermal safe mode until a regular operating condition is reestablished. When local storage resources run low, a node might choose to expend power resources to compress, delete, or transmit data off the node to free up space for future data collection. There might also be cases where a node experiences a planned offline state to save and accumulate power.

In addition to power, thermal, and storage, other resource constraints may exist on a node such that the preservation of resources is necessary to preserve the existence (and proper function) of the node in the network. Nodes operating in these conditions might benefit from TVR computations as the connectivity of the node changes over time as part of node preservation.

2.1. Assumptions

To effectively manage on-board functionality based on available resources, a node must comprehend specific aspects concerning the utilization and replenishment of resources. It is expected that patterns of the environment, device construction, and operational configuration exist with enough regularity and stability to allow meaningful planning. The following assumptions are made with this use case:

1. Known resource expenditures. It is assumed that there exists some determinable relationship between the resources available on a node and the resources needed to participate in a network. A node would need to understand when it has met some condition for participating in, or dropping out of, a network. This is somewhat similar to predicting the amount of battery life left on a laptop as a function of likely future usage.
2. Predictable resource accumulation. It is assumed that the accumulation of resources on a node are predictable such that a node might expect (and be able to communicate) when it is likely to next rejoin a network. This is similar to predicting the time at which a battery on a laptop will be fully charged.
3. Consistent cost functions. It is assumed that resource management on a node is deterministic such that the management of a node as a function of resource expenditure and accumulation is consistent enough for link planning.

2.2. Routing Impacts

Resource management in these scenarios might involve turning off elements of the node as part of on-board resource management. These activities can affect data routing in a variety of ways.

1. Power Savings. On-board radios may be turned off to allow other node processing. This may happen on power-constrained devices to extend the battery life of the node or to allow a node to perform some other power-intensive task.
2. Thermal Savings. On-board radios may be turned off if there are thermal considerations on the node, such as an increase in a node's operating temperature.

3. Storage Savings. On-board radios may be turned on with the purpose of transmitting data off the node to free local storage space to collect new data.

Whenever a communications device on a node changes its powered state there is the possibility (if the node is within range of other nodes in a network) that the topology of the network is changed, which impacts route calculations through the network. Additionally, whenever a node joins a network there may be a delay between the joining of the node to the network and any discovery that may take place relating to the status of the node's functional neighborhood. During these times, forwarding to and from the node might be delayed pending some synchronization.

2.3. Example

An illustrative example of a network necessitating resource preservation is an energy-harvesting wireless sensor network. In such a network, nodes rely exclusively on environmental sources for power, such as solar panels. On-board power levels may fluctuate based on various factors including sensor activity, processing demands, and the node's position and orientation relative to its energy source.

Consider a simple three-node network where each node accumulates power through solar panels. Power available for radio frequency (RF) transmission is shown in [Figure 1](#). In this figure, each of the three nodes (Node 1, Node 2, and Node 3) has a different plot of available power over time. This example assumes that a node will not power its radio until available power is over some threshold, which is shown by the horizontal line on each plot.




Figure 1: Node Power over Time

The connectivity of this three-node network changes over time in ways that may be predictable and are likely able to be communicated to other nodes in this small sensor network. Examples of connectivity are shown in [Figure 2](#). This figure shows a sample of network connectivity at three times: t1, t2, and t3.

- At time t1, Node 1 and Node 2 have their radios powered on and are expected to communicate.
- At time t2, it is expected that Node 1 has its radio off but that Node 2 and Node 3 can communicate.
- Finally, at time t3, it is expected that Node 1 may be turning its radio off, that Node 2 and Node 3 are not powering their radios, and there is no expectation of connectivity.

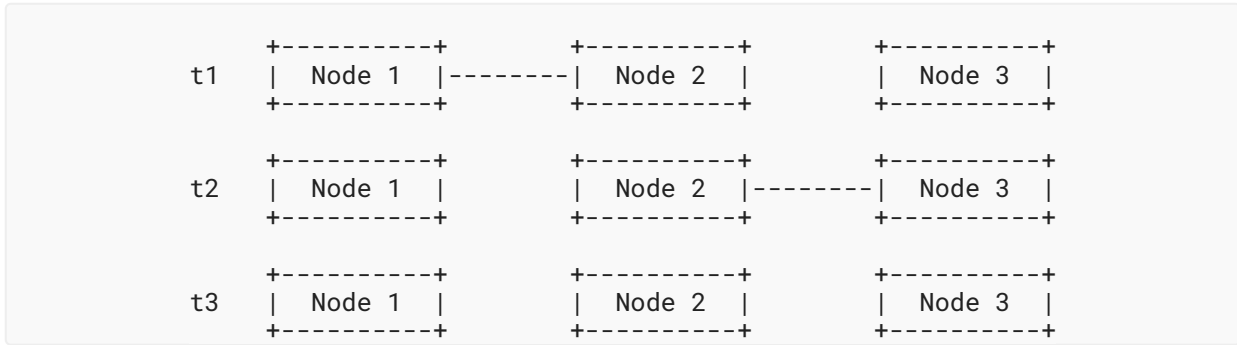


Figure 2: Topology over Time

3. Operating Efficiency

Some nodes in a network might alter their networking behavior to optimize metrics associated with the cost of a node's operation. While the resource preservation use case described in [Section 2](#) addresses node survival, this use case discusses non-survival efficiencies such as the financial cost to operate the node and the environmental impact (cost) of using that node.

When a node operates using some preexisting infrastructure, there is typically some cost associated with the use of that infrastructure. Sample costs are included as follows:

1. Nodes that use existing wireless communications, such as a cellular infrastructure, must pay to communicate to and through that infrastructure.
2. Nodes supplied with electricity from an energy provider pay for the power they use.
3. Nodes that cluster computation and activities might increase the temperature of the node and incur additional costs associated with cooling the node (or collection of nodes).
4. Beyond financial costs, assessing the environmental impact of operating a node may also be modeled as a cost associated with node operation, to include achieving carbon credits or other incentives for green computing.

When the cost of using a node's resources changes over time, a node can benefit from predicting when data transmissions might optimize costs, environmental impacts, or other metrics associated with operation.

3.1. Assumptions

The ability to predict the impact of a node's resource utilization over time presumes that the node exists within a defined environment (or infrastructure). Some characteristics of these environments are listed as follows:

1. Cost Measurability. The impacts of operating a node within its environment can be measured in a deterministic way. For example, the cost-per-bit of data over a cellular network or the cost-per-kilowatt of energy used are known.

2. **Cost Predictability.** Changes to the impacts of resource utilization are known in advance. For example, if the cost of energy is less expensive in the evening than during the day, there exists some way of communicating this change to a node.
3. **Cost Persistent.** Changes to the cost of operating in the environment persist for a sufficient amount of time such that behavior can be adjusted in response to changing costs. If costs change too rapidly, it is likely not possible to meaningfully react to their change.
4. **Cost Magnitude.** The magnitude of cost changes is such that a node experiences a minimum threshold cost reduction through optimization. A specified time period is designated for measuring the cost reduction.

3.2. Routing Impacts

Optimizing resource utilization can affect route computation in ways similar to those experienced with resource preservation. The route computation may not change the available path, but the topology as seen by an endpoint would be different. Cost optimization can impact route calculation in a variety of ways, some of which are described as follows:

1. **Link Filtering.** Data might be accumulated on a node waiting for a cost-effective time for data transmission. Individual link costs might be annotated with cost information such that adjacencies with a too high cost might not be used for forwarding. This effectively filters which adjacencies are used (possibly as a function of the type of data being routed).
2. **Burst Planning.** In cases where there is a cost savings associated with fewer longer transmissions (versus many smaller transmissions), nodes might refuse to forward data until a sufficient data volume exists to justify a transmission.
3. **Environmental Measurement.** Nodes that measure the quality of individual links can compute the overall cost of using a link as a function of the signal strength of the link. If link quality is insufficient due to environmental conditions (such as clouds on a free-space optical link or long distance RF transmission in a storm) the cost required to communicate over the link may be too much, even if access to infrastructure is otherwise in a less expensive time of day.

In each of these cases, some consideration of the efficiency of transmission is prioritized over achieving a particular data rate. Waiting until data rate costs are lower takes advantage of platforms using time-of-use rate plans -- both for pay-as-you-go data and associated energy costs. Accumulating data volumes and choosing more opportune times to transmit can also result in less energy consumption by radios and, thus, less operating cost for platforms.

3.3. Example: Cellular Network

One example of a network where nodes might seek to optimize operating cost is a set of nodes operating over cellular connections that charge both peak and off-peak data rates. In this case, individual nodes may be allocated a fixed set of "peak" minutes such that exceeding that amount of time results in expensive overage charges. Generally, the concept of peak and off-peak minutes exists to deter the use of a given network at times when the cellular network is likely to encounter heavy call volumes (such as during the workday).

Just as pricing information can act as a deterrent (or incentive) for a human cellular user, this pricing information can be codified in ways that also allow machine-to-machine (M2M) connections to prioritize off-peak communications for certain types of data exchange. Many M2M traffic exchanges involve schedulable activities, such as nightly bulk file transfers, pushing software updates, synchronizing datastores, and sending noncritical events and logs. These activities are usually already scheduled to minimize impact on businesses and customers but can also be scheduled to minimize overall cost.

Consider a simple three-node network, similar to the one pictured in [Figure 1](#), except that in this case the resource that varies over time is the cost of the data exchange. This case is illustrated below in [Figure 3](#). In this figure, a series of three plots are given, one for each of the three nodes (Node 1, Node 2, and Node 3). Each of these nodes exists in a different cellular service area that has different peak and off-peak data rate times. This is shown in each figure by times when the cost is low (off-peak) and when the cost is high (peak).

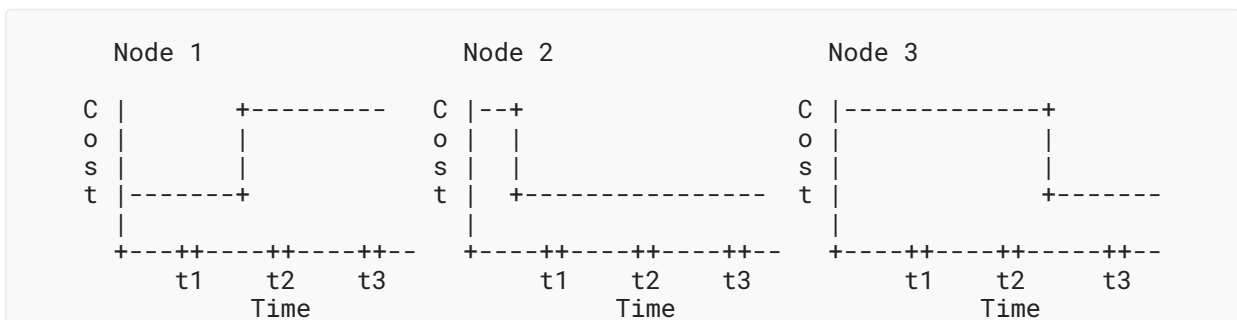


Figure 3: Data Cost over Time

Given the presumption that peak times are known in advance, the cost of data exchange from Node 1 through Node 2 to Node 3 can be calculated. Examples of these data exchanges are shown in [Figure 4](#). From this figure, both times t1 and t3 result in a smaller cost of data exchange than choosing to communicate data at time t2.

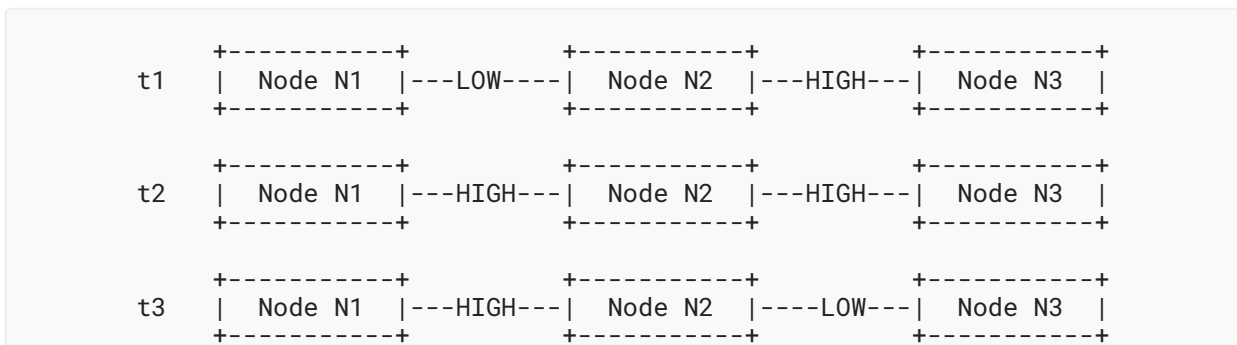


Figure 4: Data Exchange Cost over Time

While not possible in every circumstance, a highly optimized plan could be to communicate from Node 1 to Node 2 at time t1 and then queue data at Node 2 until time t3 for delivery to Node 3. This case is shown in [Figure 5](#).

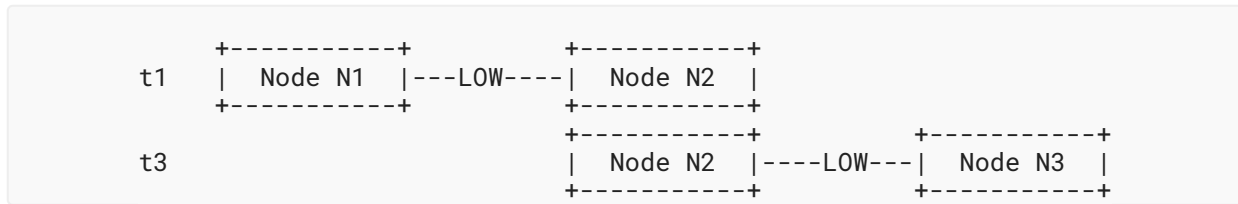


Figure 5: Data Cost Using Storage

3.4. Another Example: Tidal Network

Another example related to operating efficiency is often referred to as a "tidal network," in which traffic volume undergoes significant fluctuations at different times. Take, for instance, a campus network, where thousands of individuals go to classrooms and libraries during the daytime and retire to the dormitories at night. This results in a regular oscillation of network traffic across various locations within the campus.

In the context of a tidal network scenario, energy-saving methods may include the deactivation of some or all components of network nodes. These activities have the potential to alter network topology and impact data routing in a variety of ways. Ports on network nodes can be selectively disabled or enabled based on traffic patterns, thereby reducing the energy consumption of nodes during periods of low network traffic.

More information on tidal networks can be found in [TIDAL].

4. Dynamic Reachability

When a node is placed on a mobile platform, the mobility of the platform (and thus the mobility of the node) may cause changes to the topology of the network over time. The impacts on the dynamics of the topology can be very important. To the extent that the relative mobility between and among nodes in the network and the impacts of the environment on the signal propagation can be predicted, the associated loss and establishment of adjacencies can also be planned for.

Mobility can cause the loss of an adjacent link in several ways, such as that which follows:

1. Node mobility can cause the distance between two nodes to become large enough that distance-related attenuation causes the mobile node to lose connectivity with one or more other nodes in the network.
2. Node mobility can also be used to maintain a required distance from other mobile nodes in the network. While moving, external characteristics may cause the loss of links through occultation or other hazards of traversing a shared environment.
3. Node mobility can cause the distance between two nodes to vary quickly over time, making it complicated to establish and maintain connectivity.
4. Nodes equipped with communication terminals capable of adjusting their orientation or moving behind and emerging from barriers will also establish and lose connectivity with other nodes as a function of that motion.

Mobile nodes, like any node, may encounter issues regarding resource preservation and cost efficiency. In addition, they may face unique challenges associated with their mobility. The intermittent availability of links can lead to dynamic neighbor relationships at the node level. This use case aims to examine the routing implications of motion-induced changes to network topology.

4.1. Assumptions

Predicting the impact of node mobility on route computation requires some information relating to the nature of the mobility and the nature of the environment being moved through. Some information presumed to exist for planning is listed as follows:

1. **Path Predictability.** The path of a mobile node through its environment is known (or can be predicted) as a function of (at least) time. It is presumed that mobile nodes using TVR algorithms would not exhibit purely random motion.
2. **Environmental Knowledge.** When otherwise well-connected mobile nodes pass through certain elements of their environment (such as a storm, a tunnel, or the horizon), they may lose connectivity. The duration of this connectivity loss is assumed to be calculable as a function of node mobility and the environment itself.

4.2. Routing Impacts

Changing a network topology affects the computation of paths (or subpaths) through that topology. In particular, the following features can be implemented in a network with mobile nodes such that different paths might be computed over time:

1. **Adjacent Link Expiration.** A node might be able to predict that an adjacency will expire as a function of that node's mobility, the other node's mobility, or some characteristic of the environment. Determining that an adjacency has expired allows a route computation to plan for that loss rather than default to an error recovery mechanism.
2. **Adjacent Link Resumption.** Just as the loss of an adjacency can be predicted, it may be possible to predict when an adjacency will resume.
3. **Data Rate Adjustments.** The achievable data rate over a given link is not constant over time and may vary significantly as a function of both relative mobility between a transmitter and receiver as well as the environment being transmitted through. Knowledge of both mobility and environmental state may allow for prediction of data rates, which may impact path computation.
4. **Adjacent Link Filtering.** Separate from the instantaneous presence or absence of an adjacency, a route computation might choose to not use an adjacency if that adjacency is likely to expire in the near future or if it is likely to experience a significant drop in predicted data rate.

4.3. Example: Mobile Satellites

A relatively new type of mobile network that has emerged over the past several years is the Low Earth Orbit (LEO) networked constellation. There are a number of such constellations being built by both private industry and governments. While this example describes LEO satellite systems, the mobility events can be applied to satellite systems orbiting at different altitudes (including Very LEO (V-LEO) or Medium Earth Orbit (MEO)).

Many LEO networked constellations have a similar operational concept of hundreds to thousands of inexpensive spacecraft that can communicate both with their orbital neighbors as well as down to any ground station that they happen to be passing over. A ground station is a facility used to communicate with satellites in LEO. The relationship between an individual spacecraft and an individual ground station becomes somewhat complex as each spacecraft may only be over a single ground station for a few minutes at a time. Moreover, as a function of the constellation topology, there are scenarios where (1) the inter-satellite links need to be shut down for interference avoidance purposes or (2) the network topology changes, which modifies the neighbors of a given spacecraft.

A LEO networked constellation represents a good example of planned mobility based on the predictability of spacecraft in orbit. While other mobile vehicles may encounter unpredictable fluctuations in velocity, spacecraft operate in an environment with relatively stable velocity conditions. This determinism makes them an excellent candidate for TVR computations. However, inter-satellite link failures could still introduce unpredictability in the network topology.

Consider three spacecraft (N1, N2, and N3) following each other sequentially in the same orbit. This is sometimes called a "string of pearls" configuration. Spacecraft N2 always maintains connectivity to its two neighbor spacecraft: N1, which is behind in the orbit, and N3, which is ahead in the orbit. This configuration is illustrated in Figure 6. While these spacecraft are all mobile, their relative mobility ensures continuous contact with each other under normal conditions.

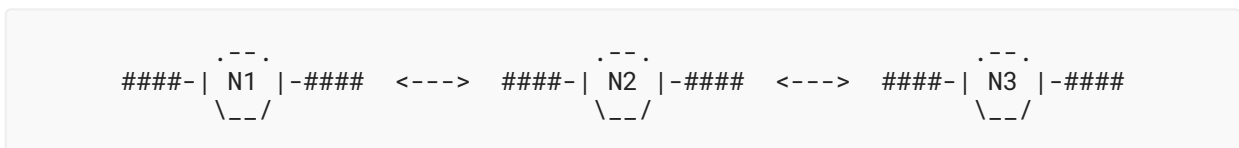


Figure 6: Three Sequential Spacecraft

Flying over a ground station imposes a non-relative motion between the ground and the spacecraft -- namely that any given ground station will only be in view of the spacecraft for a short period of time. The times at which each spacecraft can see the ground station is shown in the plots in Figure 7. In this figure, ground contact is shown when the plot is high, and a lack of ground contact is shown when the graph is low. From this, we see that spacecraft N3 can see ground at time t1, N2 sees ground at time t2, and spacecraft N1 sees ground at time t3.




Figure 7: Spacecraft Ground Contacts over Time

Since the ground station in this example is stationary, each spacecraft will pass over it, resulting in a change to the network topology. This topology change is shown in Figure 8. At time t1, any message residing on N3 and destined for the ground could be forwarded directly to the ground station. At time t2, that same message would need to, instead, be forwarded to N2 and then forwarded to ground. By time t3, the same message would need to be forwarded from N2 to N1 and then down to ground.

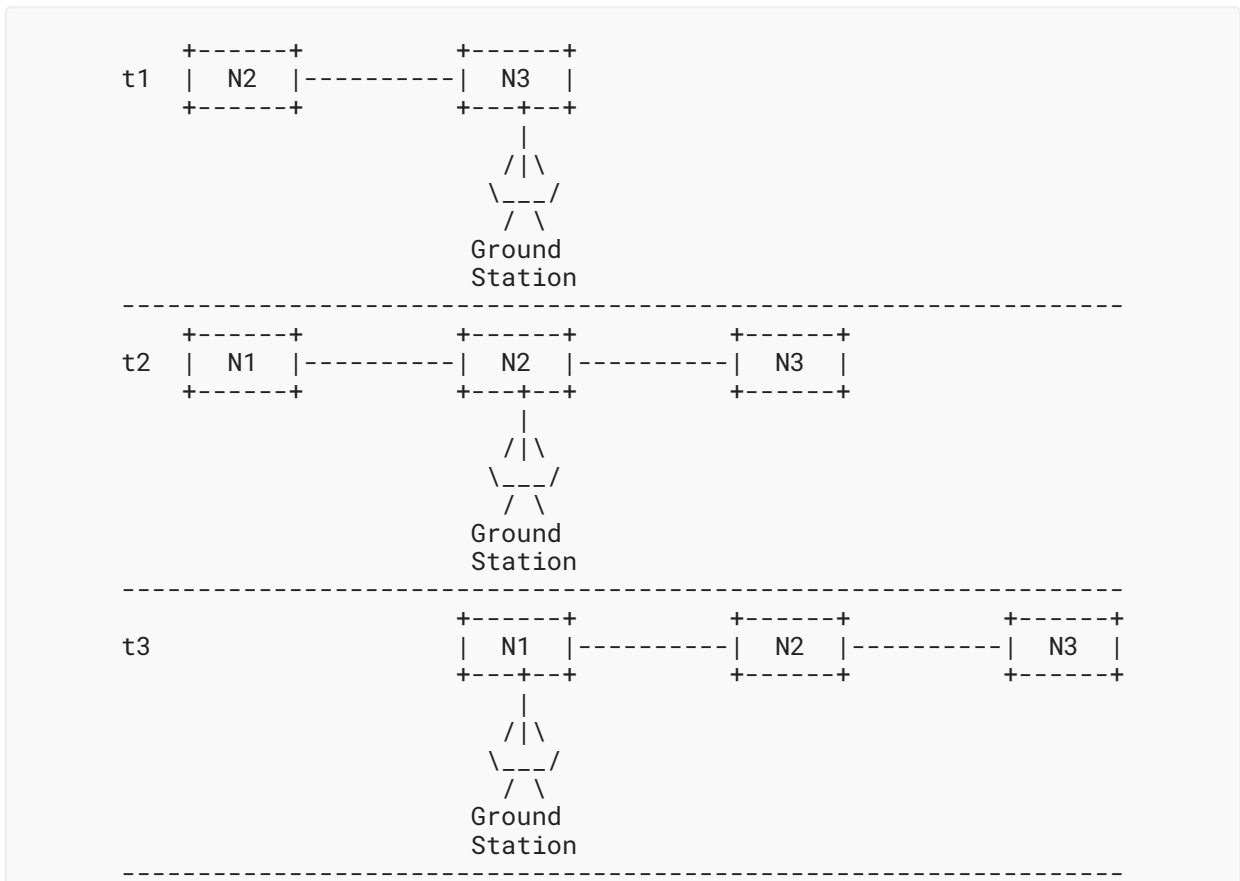


Figure 8: Constellation Topology over Time

This example focuses on the case where the spacecrafts fly over a ground station and introduce changes in the network topology. There are also scenarios where the in-constellation network topology varies over time following a deterministic time-driven operation from the ground system. More information on in-constellation network topology can be found in [SAT-CONSTELLATION] and [SCN]. For this example, and in particular for within constellation network topology changes, the TVR approach is important to avoid the Interior Gateway Protocol (IGP) issues mentioned in [SAT-CONSTELLATION].

4.4. Another Example: Predictable Moving Vessels

Another relevant example for this use case involves the movement of vessels with predictable trajectories, such as ferries or planes. These endpoints often rely on a combination of satellite and terrestrial systems for Internet connectivity, capitalizing on their predictable journeys.

This scenario also covers situations where nodes employ dynamic pointing solutions to track the mobility of other nodes. In such cases, nodes dynamically adjust their antennas and application settings to determine the optimal timing for data transmission along the path.

5. Security Considerations

While this document does not define a specific mechanism or solution, it serves to motivate the use of time-based validation and revocation strategies. Therefore, security considerations are anticipated to be addressed elsewhere, such as within a TVR schedule definition or through a protocol extension utilizing a TVR schedule. However, it's important to note that time synchronization is critical within a network employing a TVR schedule. Any unauthorized changes to network clocks can disrupt network functionality, potentially leading to a Denial of Service (DoS) attack.

6. IANA Considerations

This document has no IANA actions.

7. Informative References

[SAT-CONSTELLATION] Han, L., Li, R., Retana, A., Chen, M., Su, L., and T. Jiang, "Problems and Requirements of Satellite Constellation for Internet", Work in Progress, Internet-Draft, draft-lhan-problems-requirements-satellite-net-06, 4 January 2024, <<https://datatracker.ietf.org/doc/html/draft-lhan-problems-requirements-satellite-net-06>>.

[SCN] Wood, L., "Satellite Constellation Networks", *Internetworking and Computing over Satellite Networks*, pp. 13-34, DOI 10.1007/978-1-4615-0431-3_2, April 2003, <https://link.springer.com/chapter/10.1007/978-1-4615-0431-3_2>.

[TIDAL] Zhang, L., Zhou, T., Dong, J., and N. Nzima, "Use Case of Tidal Network", Work in Progress, Internet-Draft, draft-zzd-tvr-use-case-tidal-network-02, 28 July 2023, <<https://datatracker.ietf.org/doc/html/draft-zzd-tvr-use-case-tidal-network-02>>.

Acknowledgments

Many thanks to Tony Li, Peter Ashwood-Smith, Abdussalam Baryun, Arashmid Akhavain, Dirk Trossen, Brian Sipos, Alexandre Petrescu, Haoyu Song, Hou Dongxu, Tianran Zhou, Jie Dong, Nkosinathi Nzima, and Vinton Cerf for their useful comments that helped improve the document.

Authors' Addresses

Edward J. Birrane, III

JHU/APL

Email: edward.birrane@jhuapl.edu

Nicolas Kuhn

Thales Alenia Space

Email: nicolas.kuhn.ietf@gmail.com

Yingzhen Qu

Futurewei Technologies

Email: yingzhen.ietf@gmail.com

Rick Taylor

Aalyria Technologies

Email: rtaylor@aalyria.com

Li Zhang

Huawei

Email: zhangli344@huawei.com