

AALYRIA

Overview and FAQ

SPACETIME

Orchestrating the Next Generation of Aerospace Networks



The information in this document is proprietary to Aalyria Technologies. This data shall not be duplicated, used, or disclosed -- in whole or in part -- without the express written consent of Aalyria Technologies. The data subject to this restriction are contained in all sheets of this whitepaper.

A Preface on Aerospace Communication Networks

Aerospace communication networks are evolving from static, point-to-point systems into dynamic, networked architectures. Traditionally, geostationary (GEO) satellites have provided basic coverage using “bent pipe” models, where satellites act as passive relays, merely reflecting signals back to Earth without processing or intelligent routing. This simple design works well for broadcasting or other low-interactivity applications, but struggles to meet the demands of modern data-driven applications. Latency and capacity-density limitations inherent in this model increasingly limit real-time connectivity requirements, particularly as data-intensive and interactive applications proliferate. To meet these demands, satellite communications are now moving toward an active, networked model in which satellites and ground stations work together to route data flexibly, creating more resilient, responsive, and intelligent networks.

To understand the leap from bent-pipe satellites to a fully networked satellite system, it helps start with basic networking principles. Networks connect devices so they can share data, often using a structure defined by the OSI (Open Systems Interconnection) model. The OSI model divides communication into layers, ranging from the physical connections (cables, radios) at the base layer to the application layer, where data appears in a form users interact with. Layering provides a clear framework for communication and enables complex networking functions to be managed and understood in parts. Traditional routing, for instance, operates at the network layer (Layer 3), where routers manage the paths data takes across interconnected devices. Routers make these decisions autonomously based on protocols that analyze the network’s current state, often by assessing link performance and topology.

This is where Software-Defined Networking (SDN) enters the scene. Unlike traditional routers, which independently make routing decisions, SDN separates the decision-making “control plane” from the “data plane,” where data flows. SDN logically centralizes control, allowing a software-based controller to direct all network routing decisions leveraging its holistic view of the network. By decoupling these planes, SDN makes networks far more adaptable and manageable. For dynamic or large-scale environments like space, where topology is continuously changing as satellites orbit, SDN offers a critical advantage: it provides centralized control that can proactively adapt the network routing around changes in the environment, topology, or mission needs, improving resilience and performance.

In satellite networks, especially those comprising large constellations in low-Earth orbit (LEO), maintaining optimal connectivity is complex. Traditional routing protocols struggle in these dynamic environments because they depend on frequent updates between individual nodes to make routing decisions. This dependency can create lag or even prevent routing convergence in large networks with rapidly changing topologies. In contrast, SDN’s centralized control allows for more stable, efficient routing across the network, a principle that has driven the development of Aalyria’s Spacetime platform and its Temporospatial SDN (TS-SDN) model.

Spacetime’s implementation of a TS-SDN goes beyond typical SDN/SD-WAN by incorporating control over the wireless interconnections and the physical layer itself. It enables the software-based

controller to manage not only how data moves through the network, but also the structure of the network in real-time. Through a detailed digital model, or “digital twin,” Spacetime continuously monitors and predicts the positions and orientations of all networked assets, such as satellites and ground stations, and environmental conditions like weather. This predictive control is essential for seamless connectivity and high data quality, even as satellites move or face potential disruptions. When a connection is anticipated to degrade due to a weather event or orbital dynamics, Spacetime can proactively reroute data through entirely new alternate paths, making changes in rather than solely waiting for a disruption to trigger reactive routing changes.

With this approach, Spacetime enables continuous network adaptation across domains—land, sea, air, and space—while ensuring the resilience needed in aerospace and defense. Its temporospatial control allows resilient, high-performance connectivity, essential for applications where even minor disruptions could have significant consequences. As you explore this white paper, you will see how Spacetime’s innovative approach to network orchestration transforms the potential for satellite networks, delivering unprecedented adaptability and continuity in connectivity.

	SDN/SD-WAN	Spacetime (TS-SDN)
Control Plane	Orchestrate routing and traffic flow at Layer 3 (network layer) over pre-existing, static network topologies (SDN) or overlay networks (SD-WAN).	Orchestrates the network control plane across Layers 1 (physical link topology), 2 lower MAC (transceiver/air interface configuration), and 3 (routing). It establishes and manages the actual physical connections between nodes, not just the paths data takes through them.
Physical Topology Management	N/A. Relies on sensing link health and performance to automate switching between networks or paths.	Employs a “digital twin” of the network, modeling the positions, orientations, and motion of all platforms (satellites, aircraft, ships, etc.). This allows it to predict link availability and proactively task steerable antennas or laser terminals to establish new connections and create new paths.
Handling Dynamic Environments	Limited to react to connectivity issues after they occur. Reaction times at lower layers can have ripple effects at layer 4+ ¹ .	Excels in dynamic environments by predicting future connectivity opportunities and potential disruptions. This allows it to proactively reconfigure, ensuring seamless connectivity even in the face of unexpected events.
Use Cases	Primarily used for terrestrial or geostationary networks with fixed infrastructure — or fixed coverage	Ideally suited for large-scale, dynamic networks, especially those incorporating steerable directional antennas/optics.

¹ Brian J. Barritt and Wesley Eddy. "SDN Enhancements for LEO Satellite Networks," AIAA 2016-5755. 34th AIAA International Communications Satellite Systems Conference. October 2016.

Table of Contents

Background.....	4
The Motivating Trends.....	4
Commercial market.....	4
Government & defense market.....	5
The Problem that Spacetime Solves.....	6
Spacetime Overview.....	8
Deployment and Authorization to Operate.....	11
About the Federation.....	11
Frequently Asked Questions (FAQ).....	15
Is this a network management system? A network traffic protocol?.....	15
Who are the intended customers that would deploy Spacetime?.....	15
What if some or all satellite payloads are out of our control?.....	15
How is it better than SD-WAN over multiple COMSAT providers?.....	16
Is it production ready? What's the flight heritage? TRL?.....	18
Are any of the major commercial SATCOM operators using it?.....	19
Does it support manual scheduling by human operators?.....	19
Does the mission/user data flow through Spacetime?.....	20
Is there a limit to the size of networks Spacetime can orchestrate?.....	20
Can it maneuver satellites or aircraft to improve the network?.....	20
Does Spacetime command the actual satellites?.....	20
Does Spacetime need to be installed on the individual nodes?.....	21
Does it support store & forward (intermittent) scheduling?.....	21

Background

The Motivating Trends

Commercial market

The compound annual growth rate (CAGR) of global consumer data traffic continues to grow by double digits annually. However, it is often uneconomic to deploy fiber or cell towers outside urban areas. Geostationary orbiting (GEO) satellites have provided global coverage for decades. However, from an altitude of 35,786 km, they are very limited in the capacity density they can support per unit area (e.g., Mbps/km²). The round trip latency through a GEO satellite is also more than 500 ms, creating a poor experience for interactive and other latency-sensitive applications. The unmet demand for commercial data traffic is driving satellite operators to deploy new satellite constellations into non-geostationary orbits (NGSO). As the altitude of satellites is reduced by as much as an order of magnitude to medium-earth orbit (MEO) and another order of magnitude to low-earth-orbit (LEO), or perhaps another order of magnitude with High-Altitude Platform Station (HAPS) aircraft, the network operator is sacrificing the coverage area provided by each satellite for orders of magnitude greater capacity density – dramatically improving the data traffic that can be delivered. However, the significantly increased number of satellites now required for global coverage, and the non-geostationary motion, create many new operational challenges.

Population densities (and income distribution of people able to afford satellite user terminals) are not uniformly distributed across the earth (in fact, they are [quite clumpy](#)). And NGSO satellites are constantly moving. So these satellite and HAPS operators employ steerable directional antennas to focus the capacity to support the data traffic demands of the most populated areas. And, to backhaul the capacity from each satellite or HAPS platform to the Internet, the network operators further leverage steerable directional antennas and free-space optical (FSO) communications for the Inter-Satellite Links (ISLs) and Gateway (GW) links to the Points of Presence (POPs) that connect back to the Internet.




Figure 1: Characteristics of different orbital regimes and domains

Government & defense market

As conflict theaters become increasingly sophisticated and contested, the need for reliable, resilient, self-healing, and secure communication has never been more critical *in all domains*. Near-peer adversaries seek to degrade communications and are increasingly developing capabilities to jam and degrade traditional communication channels. A solution lies within the domain of low probability of intercept (LPI), detect (LPD), and jam (LPJ) technologies – collectively known as LPX. LPX technologies offer the essential attributes needed in such environments. They minimize the likelihood of signal interception and detection by adversarial systems, protect against jamming attempts, and ensure the location of the transmitter is challenging to discern. This combination of properties is critical in preserving the integrity and availability of communication networks in hostile environments.

LPX attributes are also possible by harnessing the power of steerable directional antennas, advanced frequency bands – often millimeter-wave bands, and free space optical communication systems. Steerable directional antennas limit the spatial region within which reception of the signal is possible, which enhances signal security and reduces the probability of being located. Antenna directivity also allows wireless links to close (with higher data rates) at longer distances, creating additional self-healing options to ensure mission success.

The Problem that Spacetime Solves

Deploying steerable directional communications in real-world environments is not without its challenges. Moving vehicles or other assets that host steerable directional communications infrastructure can encounter obstructions, and the terrain itself may be a barrier. Changing weather can impact free space optics (FSO) communications or radio wave propagation in certain bands, especially at higher frequencies. The factors that make steerable directional communication high-throughput and secure also require constant attention to the state and orientation of network nodes, requiring a dynamic and responsive network structure or link topology.

The choice of physical link structure in such a network has profound implications on network performance, resilience, and adaptability. Some wireless link structures might excel in a specific context and mission requirements – but fall short in others. Moreover, most modern high-throughput radios are equipped with adaptive coding and modulation (ACM) capabilities. ACM allows the radio to adjust its transmission characteristics in response to changes in link quality or distance. This means that depending on the link geometry, range, weather conditions, and other factors, the effective data rate of the link can vary, affecting network performance.

The task of managing these diverse, complex, and continually changing network conditions is immense and requires advanced network orchestration capabilities. Such orchestration should be capable of not only *maintaining* network connectivity, but also *optimizing* it based on continually changing end-to-end service requirements. For commercial networks, these end-to-end requirements should be expressed by operations support systems and business support systems (OSS/BSS) via software interfaces (e.g., MEF APIs or 5G/6G slicing). For defense networks, these end-to-end requirements should be expressed by Battle Management and Decision Support systems (e.g., Anduril's Lattice OS, Lockheed's DIAMONDShield, etc.) – taking into account the full range of constraints and possibilities offered by the dynamically changing set of available assets, the environment, and the mission's needs. That is where Aalyria's Spacetime comes into play.

Unlike contemporary SDN and SD-WAN solutions, Spacetime can orchestrate the physical link topology of steerable directional beams – across all domains. Instead of performing path computation and route selection over a topology you don't control, it flips the script: it enables

path computation and route selection over the multiverse of topologies that can exist across space and time. From this set of all possible wireless links, it ensures that *the chosen links are optimal for the needs of the end-to-end network services and their associated requirements*. It can bring entirely new LPX link topologies into existence, configuring their spectrum resources, and configuring the network routing over them to reconstitute networks or meet new mission requirements in tactically responsive timeframes.

Spacetime Overview

The Spacetime software platform was designed to orchestrate very large and constantly changing communications networks of flying, orbiting, floating, and driving vehicles, carrying a wide array of wireless radio and optical links and delivering data to a widely scattered set of users and endpoints across land, sea, air and space. It has years of heritage performing such orchestration in production operations for real-world aerospace networks. Vendor-, asset-, frequency-, orbit-, data- and waveform-adaptive, Spacetime uses any available links across those platforms – along with any accessible terrestrial fixed or mobile wireless, wired, and fiber assets – to form a continuous and highly-redundant mesh. It orchestrates this mesh to constantly and autonomously adapt to changing weather, motion/location/speed of vehicles, user requests for transport, and link obstructions (e.g., wingtips, adversary action, horizon masks, rain-fade, sun outages).




Figure 2:

To ensure a perpetually reliable network, Spacetime continuously updates and exercises a detailed physics model of all current and potential links. Using a planet-scale digital twin, the system predicts the motion of platforms to maintain understanding of connectivity opportunities, across space and time, for every possible pair of compatible transceivers. Using nowcast and

forecast weather data, and the location and characteristics of known obstructions (on and surrounding the platforms), relevant terrain, and all known and detected interferers (environmental (e.g., solar), unintentional or adversarial), Spacetime rapidly predicts impacts on link quality and throughput and the performance and reliability of future possible links. Our propagation modeling assesses how much attenuation or degradation of wireless signals would be encountered between all pairs of compatible transceivers from real-time motion, weather, atmosphere, geometric or field-of-regard obstructions, and terrain impacts; predicts wireless access, backhaul, and inter-satellite link (ISL) opportunities with full consideration of signal strength and quality based on near-term attenuation/degradation constraints; and models and precomputes for every possible link, how much energy would be received by the intended receiver, and also how much interference that choice would impart upon other receivers, or potential adversaries that want to detect transmissions, or potential victims of interference in a spectrum coordination environment. With this information, Spacetime intelligently and proactively evolves the network topology, making new connections and re-routing traffic before links disappear or degrade, all without dropping a packet. Spacetime also monitors traffic flow, accommodating increases by activating additional links to maximize new routes, or adjusting to reduced traffic flow by adding redundant routes or increasing coding rates to maximize resilience and reliability.




Figure 3: Dynamic, Autonomous Reconstitution. In less than 100 milliseconds, Spacetime automatically evolves the network and re-routes all impacted traffic in response to the loss (or isolation) of a satellite - shown in blue.

There is no need to tell Spacetime through what satellite to traverse or what communication band to use, unless desired. Instead, business or mission operations systems (machines) or operators (humans) express end-to-end goals for transport, including committed and expected information rates, latency constraints, and priority. As a continuous self-healing, dynamic network control plane, the system solves the rest – optimizing and evolving the schedule for all beam tasking (pointing directions for steerable antennas and optical terminals), encapsulation and routing functions (channelizers, IPv4, IPv6, segment routing, MPLS, etc.), and radio & optical transceiver resource management (frequencies/wavelengths, and bandwidth to use).

Aalyria designed Spacetime to not only evolve the network in seconds, as it orchestrates reactively around the environment in real-time, but to also *proactively* schedule the *predictive* tasking of the infrastructure over minutes, days, and weeks according to continually changing weather forecasts, motion forecasts, jamming threats, and scheduled dates for when the network operator will bring new network resources and transport requirements online. Spacetime immediately discovers and incorporates newly deployed or reactivated assets into the network (with secure provisioning and utilization in seconds) to add redundancy, decrease latency, or increase throughput across the network.

Spacetime is built on an open architecture that provides simple, powerful “northbound” APIs for use by commercial OSS/BSS or defense Battle Management Decision Support applications, human operators, or any other authorized application to request end-to-end data services matching their throughput, priority, latency, and other requirements. Spacetime uses this information to constantly evolve and optimize the topology to ensure all data gets to its endpoint as fast as (or faster than) requested, even as weather, motion, and adversaries work to stop it. Common, open source “southbound” APIs also allow the on-ramp of new technologies at any time – freeing the commercial industry and public sector from costly vendor lock-in and supporting the integration of proliferated architecture components developed by multiple vendors, as well as interoperability with external systems and users.




Figure 4:

Deployment and Authorization to Operate

Each customer of Spacetime is provided with one or more dedicated private instances (i.e., development, staging, and production environments) for *each* network administrative domain. For example, a commercial satellite network operator like Telesat will use their private, dedicated production instance of Spacetime to operate and orchestrate the antenna tasking & scheduling, radio resource management, and routing rules for all of their satellite payloads, user terminals, gateway antennas, and POPs. Some customers may have multiple production instances to support different administrative domains within their organization or enterprise.

Spacetime can be deployed across cloud, on-prem, and edge environments, and it is agnostic to the underlying compute infrastructure or the specific cloud provider due to Aalyria's use of Kubernetes-managed, container-based deployment models.

Spacetime's integration with Second Front's trusted [Game Warden](#) platform provides accredited deployment and risk management procedures in accordance with Department of Defense Impact Levels 4 and 5 (IL4 and IL5). Adopting a [Continuous Authorization to Operate \(cATO\)](#) approach means the ability to focus on maintaining security and risk management frameworks without impacting product development velocity. Expansions into IL6 and Top Secret Cloud environments are additionally underway with accreditation submissions targeted in late Q1 2024 and customer evaluations to begin immediately thereafter.

Customers retain full control and governance authority over their Spacetime instances and what external connectivity they wish to establish (if any). In highly controlled environments, a customer may elect to incorporate a Cross Domain Solution (CDS) to inspect and validate any message traffic to or from Spacetime's APIs in accordance with their security models. Coordination of interconnections *between* network administrative domains (e.g., between the DoD and a commercial satellite operator, between different MILSATCOM networks within the US Government, or between different business units of a commercial enterprise) is facilitated by Spacetime's East/West or "Federation" APIs.

About the Federation

Spacetime's [Federation API](#) addresses the growing need to allow peer networks to request and to supply network resources and interconnections between partners' networks (across administrative domains). This results in far greater network resilience and cost effectiveness than can be achieved by any one network alone; the increased coordination facilitates dynamic, real-time inter-network connections, which allows operators to automatically and quickly supplement gaps in network coverage or advertise unused capacity to make full use of underutilized assets.

Spacetime, acting as a temporospatial SDN controller, exposes and utilizes the Federation API between partners' networks, providing full federation orchestration capabilities. This is particularly powerful given Spacetime's ability to bridge multiple domains, enabling it to orchestrate meshes across different networks in land, sea, air, and space.




Figure 5: Federation API

Imagine you are a provider that serves user terminals on the ground through your network of satellites and ground stations. Normally, your network is fairly stable; but there are many reasons why a connection to a user terminal might fail – for example, a surge in usage in one part of the network might over-task the ground stations in that area, and unfortunately you have no extra capacity to serve those users because an ongoing weather event has severely attenuated all satellite links into that region. Or, maybe there aren't any problems with broken connections at all; but, you require lower latency to serve a new customer / mission requirement that just arrived.

This is exactly what the federation is built for. Networks within the federation receive tailored, dynamic information about different options from a variety of providers across terrestrial, HAPS, LEO, MEO, and GEO networks. The options may include interconnections and transit options that would require changes to the orchestration of those partner networks' underlying physical link topology to enact. You can use these to temporarily replace a broken or high-latency link. These options come with details such as latency, bandwidth, availability time windows, and so on, enabling you to make the most informed decisions possible.




Figure 7: Any organization can add support for the Federation API to facilitate dynamic interconnection and transit

Frequently Asked Questions (FAQ)

Is this a network management system? A network traffic protocol?

Spacetime is a new twist on a Software Defined Network controller. Whereas contemporary SDN controllers can only orchestrate the layer 3 network control plane (routing), Spacetime can orchestrate the present and scheduled network control plane across layer 1 (physical link topology), layer 2 lower MAC (transceiver/air interface configuration), and layer 3 (routing). We refer to this concept as Temporospatial SDN in academic literature. It reflects the fact that the controller's information base must be able to propagate the future motion of the physical platforms to model the candidate wireless link topology, orchestrate that topology, and compute routing paths across space and time.

In addition to these layer 1-3 control plane functions, Spacetime is also capable of monitoring network telemetry (like signal quality and traffic congestion) for the purposes of network resource optimization. It also exposes key metrics (like present and forecast fulfillment of service level agreements, comparisons of signal quality vs digital twin estimates, etc.) to facilitate automated alerting by external incident response and management systems.

Who are the intended customers that would deploy Spacetime?

Aalyria's primary target customers for Spacetime are those commercial and public sector organizations charged with operating communications network infrastructure – especially ones with mesh networks that incorporate steerable directional wireless communications. These include commercial and military satellite network operators, earth observation and intelligence community satellite constellation operators, HAPS and airborne mesh network operators, maritime mesh network operators, terrestrial network operators employing millimeter wave and optical technology in integrated access and backhaul networks, and civil space exploration agencies.

However, the network effect of Spacetime and [the Federation](#) model provides secondary benefits to organizations and enterprises, even when they do not have any such infrastructure under their control. See [“What if some/all of the satellite payloads are out of our control?”](#) in the FAQ.

What if some or all satellite payloads are out of our control?

An instance of Spacetime can incorporate networks outside your control, supports automated network switching, and can even support the automated brokering of dynamic interconnections with and across these networks – while simultaneously performing the complete control and

optimization of any assets you fully control. Each instance can mix and match assets and networks across the following tiers, which are listed and described in increasing order of capability afforded to the customer:

(1) **Basic:** Like contemporary SD-WAN solutions, Spacetime can monitor key performance metrics for network services that transit outside your control and reactively switch networks when necessary. It does not need any information for this basic tier.

(2) **Enhanced:** By adding any available information about the satellite orbits, transmitter and receiver chains, and antenna gain patterns, Spacetime can predict impacts caused by weather, terrain obstructions, potential jamming threats, and other disruptions – proactively switching connections for seamless continuity. It can also proactively forecast disruptions to mission / customer end-to-end service requirements before they even occur.

(3) **Federation:** Buyers of satellite capacity can compel commercial satellite operators to support the open-source Federation API, even if they aren't using Spacetime to orchestrate their networks. A number of the largest commercial satellite network operators have already selected Spacetime to orchestrate their production networks, which natively supports the Federation API. When these networks choose to share information with your network via these open APIs, Spacetime can dynamically broker connections across their networks in real-time, offering unparalleled flexibility and responsiveness.

(4) **Full Control:** For first party network infrastructure, Spacetime can directly orchestrate everything – satellites, ground stations, user terminals, etc. – for complete control and optimization.

How is it better than SD-WAN over multiple COMSAT providers?

Contemporary SD-WAN solutions ignore that the physical wireless links and paths across the underlying networks are constantly changing. In reality, the commercial satellite providers have provisioned their network around certain expectations about the geographic locations and statistical distribution of demand – all based on the static procurement contracts and service level agreements they've signed with their customers. These expectations determine the operators' static schedule of beam handovers, management of satellite payload resources, and paths across their network. SD-WAN solutions only attempt to sense the health and current performance of these paths and automate switching between commercial providers.

This approach has its limitations. To illustrate them, consider a LEO constellation like SpaceX Starlink/Starshield configured to provide services to their static prediction of the statistical demand from a population of user terminals registered in known locations (labeled UT in the diagram below).




Figure 8: <figure name>. Credit Mike Puchol's [blog](#)

In this example, SpaceX has configured their network expecting some statistical demand from US DoD customer terminals in configured geographies. But, what if a US DoD customer terminal has a sudden mission requirement to log in from an unexpected cell / geography – one where SpaceX has not configured their network to expect a log-in? Or what if it's in an expected location; but, the customer has a mission constraint that requires transit through a certain POP to a private network instead of the Internet? Or what if the US DoD mission requirements exceed the link layer capacity of the spectrum resources assigned to the cells – or the available capacity through ISLs and GWs, given competing flows? SD-WAN cannot even attempt to solve such problems.

Supporting these requirements requires the SATCOM provider to modify the way they orchestrate their network. Until now, typical turnaround times for a customer to ask a SATCOM provider to support such needs (for example, to provision a new beam in a new geography) would typically involve a human <> human request between the customer <> provider and then, once understood and approved, a human-in-the-loop network operations task to reconfigure the provider network. This would take hours or *days*. **Those are not tactically responsive timeframes.**

To reduce this timeframe to mere *seconds*, three things are needed:

1. The SATCOM provider needs access to a technology capable of intent-driven, automated orchestration of the network topology, radio resource optimization, and routing within seconds of new requests.
2. The customer network (in this example, the US DoD) needs to compel the provider to implement open APIs that expose the candidate points of interconnection (ingress) and candidate reachable destinations (egress) and accept requests for on-demand network

transit. Note that the provider does *not* need to expose anything about the inner workings of their network.

3. The customer network (again, in this example, the US DoD) needs access to a technology capable of modeling the time periods when interconnection with the provider's ingress points is possible. For wired interconnections, this is trivial. It's also fairly straightforward for anyone to build software to model the time periods when a government-managed terminal has an unobstructed field of view of provider satellites. For interconnections from and between airborne terminals, maritime terminals, and in-space interconnections (e.g., between SDA PWSA and commercial LEO across SDA-compliant optical communication terminals), this is non-trivial.

Spacetime fulfills all three of these needs.

Is it production ready? What's the flight heritage? TRL?

Yes, it is production-ready TRL9 infrastructure with a strong flight heritage.

Spacetime ran a dynamic mesh network of high-altitude platforms with motion and conditions far less predictable than traditional satellite provider platforms, and provided resilient connectivity to real-world commercial users in dynamic real-world conditions for three continuous years.

In that period of production operation, Spacetime's core technology overcame challenging technical hurdles in this space. Orchestrating the Loon network was in many respects more complex than even the largest proliferated LEO constellations: the motion of its platforms was far less predictable than the orbital ephemerides of traditional satellites, and the frequency bands of Loon's mesh and backhaul network (E-band) are some of the most susceptible to rain and atmospheric conditions. In spite of those challenges, Spacetime (known previously as "Minkowski" or the "Loon SDN" while in earlier stages of development at Google) autonomously operated and orchestrated these networks in production for years, serving hundreds of thousands of users across multiple continents.

The following references are available to substantiate the flight heritage:

- [SDN in the Stratosphere: Loon's Aerospace Mesh Network](#) [peer-reviewed paper]
- [The Loon Library: Lessons from Building Loon's Stratospheric Communications Service](#) ["Minkowski" is discussed on pp. 207-215]
- A [comprehensive dataset](#) consisting of internal state from Spacetime and network telemetry gathered from years of serving commercial traffic and R&D experiments.

Space Systems Command (SSC), in collaboration with the Defense Innovation Unit (DIU) and the Naval Research Laboratory (NRL), successfully completed on **Dec 7, 2023** a significant live on-orbit demonstration of Aalyria's Spacetime network orchestration software platform. This

event represented a pivotal moment in the evaluation of the Spacetime platform's operational readiness, by including familiar equipment and service providers (OneWeb, Viasat, iDirect, Kymeta, SpaceX, Viasat, and Comtech) as part of a comprehensive assessment of use cases where satellite payloads were outside its control.

Are any of the major commercial SATCOM operators using it?

Yes. Several of the largest commercial SATCOM operators have already publicly announced long-term agreements to use Spacetime in their production networks. See official press releases from [Intelsat](#) (for GEO and future NGSO), [Telesat](#) (LEO), and [Rivada](#) (LEO).

The screenshot below shows one of the LEO mega constellations being operated at full scale in a real time emulation, as visualized through the Spacetime NetOps User Interface.




Figure 9:

Does it support manual scheduling by human operators?

Spacetime allows human network operators to manually override its optimization engine to administratively “pin” use of a link or path. In our own use of Spacetime in production networks, this feature was rarely used (for example, to prevent switching to a more optimal link to allow engineers time to debug a rare reoccurrence of a difficult to reproduce issue).

Spacetime was, however, designed to fully *automate* the tasking & scheduling of non-terrestrial networks. The growing adoption and incorporation of non-geostationary orbits, weather-sensitive communications bands, and multi-hop communication paths (inter-satellite links and in-space relays) are making it wholly impractical for human operators to manually schedule such networks on tactically responsive timeframes.

Our model of network operations is based around a human-on-the-loop, rather than a human-in-the-loop, operational support model. Spacetime's web based Network Operations User Interface (NetOps UI) provides human operators with situational awareness and a common operating picture of the network across all domains. It also provides contextual mapping between geospatial (physical) and logical (networking) views across the past, present and future scheduled states. URLs embedded into the alerts generated in external incident response and management systems allow human operators to jump directly into the relevant views in the NetOps UI to rapidly gain the context necessary to troubleshoot operational issues.

Does the mission/user data flow through Spacetime?

No. Spacetime only exchanges network control plane messages with software “agents” that anyone can build and install on the network nodes – or via adapters to heritage ICDs (e.g., element managers). None of the data plane traffic flows through Spacetime.

Is there a limit to the size of networks Spacetime can orchestrate?

No; we're able to elastically scale Spacetime to support large networks (e.g. Telesat LEO Lightspeed, Rivada Space LEO). Aalyria provides dedicated instances of Spacetime as a hosted, managed Software as a Service (SaaS) on the customer's preferred Kubernetes-based compute cluster, which can support auto-scaling (i.e., on Google Cloud, AWS GovCloud, or AWS Secret and Top Secret Region) or be dimensioned based on the size of the network (for bare metal servers installed in private enclaves).

Can it maneuver satellites or aircraft to improve the network?

Spacetime does not control the navigation, orbital maneuvers, or motion of any of the physical platforms or vehicles. It only subscribes to changes to the motion parameters for the physical platforms relevant to network control and orchestration (including 3rd party platforms like jamming threats, interception threats, or potential victims of interference).

Does Spacetime command the actual satellites?

Spacetime does not command the satellite bus; it only communicates with satellite payloads (and other nodes in the network) to orchestrate the network topology, radio resource management, and network traffic routing.

Does Spacetime need to be installed on the individual nodes?

No. All of Spacetime's APIs are modular and open source, which allow *anyone* to adapt individual nodes (or heritage element managers) to work with Spacetime. These APIs use simple service definitions and support automatic generation of idiomatic client and server stubs in various languages to facilitate rapid integration.

Does it support store & forward (intermittent) scheduling?

Yes; Spacetime can orchestrate the intermittent storing and forwarding of data across a network. This capability is especially relevant for sensing and earth observation.

Spacetime's APIs allow any node to be designated as store-and-forward capable at the Application Layer (where an application on the originating node manages the queuing of data) or at the Network Layer (e.g., Bundle protocol). The total memory storage capability of nodes can also be specified. The latency field can optionally be used to limit the delivery timeline.

Aalyria is creating, organizing and managing the world's most advanced networks to enable connectivity everywhere at the speed of discovery.

Contact us at info@aalyria.com or visit us at www.aalyria.com.