Temporospatial SDN for Aerospace Communications

Brian J. Barritt*

Google, Mountain View, CA, 94043, USA

Alanax Technologies, San Jose, CA, 95125, USA

Case Western Reserve University, Cleveland, OH, 44106, USA

Wesley M. Eddy†

Alanax Technologies, Avon, OH, 44011, USA

This paper describes the development of new methods and software leveraging Software Defined Networking (SDN) technology that has become common in terrestrial networking. We are using SDN to improve the state-of-the-art in design and operation of aerospace communication networks. SDN enables the implementation of services and applications that control, monitor, and reconfigure the network layer and switching functionality. SDN provides a software abstraction layer that yields a logically centralized view of the network for control plane services and applications. Recently, new requirements have led to proposals to extend this concept for Software-Defined Wireless Networks (SDWN), which decouple radio control functions, such as spectrum management, mobility management, and interference management, from the radio data-plane. By combining these concepts with high-fidelity modeling of predicted mobility patterns and wireless communications models, we can enable SDN applications that optimally and autonomously handle aerospace network operations, including steerable beam control, RF interference mitigation, and network routing updates. This approach is specifically applicable to new constellation designs for LEO relay networks that include hundreds or thousands of spacecraft, serving millions of users, and exceed the ability of legacy network management tools.

I. Introduction

The mission complexity expected for most new flagship space systems, with multiple on-board instruments, flight computers, and other avionics systems, is driving the industry towards networked platforms, which utilize packet or frame structures for their mission-internal communications. Networking technologies are essential to enabling automation and scalable mission control in complex missions architectures where communications may need to be relayed, and management of the communications paths is necessary in addition to management of the mission platforms.

Additionally, NewSpace companies and philosophy requires building cheaper and easier-to-operate systems. Existing technology and standards need to be leveraged wherever possible in order to control costs. For telemetry, command, and control applications, this generally implies software systems built on the Internet protocols. These can be successfully used for communications with (and between) CubeSats, hosted payloads, or between larger constellations of spacecraft.

Beyond this, the large LEO constellation plans that have been recently announced by multiple companies are aiming to provide Internet services. These constellations will have networking as their reason for being, and require improvements in the state of space network management due to the magnitude of their size, geometry, and handling of packet-communications rather than circuits. These differ enormously in comparison to current GEO Internet services and even other LEO constellations (like Iridium). Operating these constellations will require improved automation of functions including multi-hop relay configuration, relay failover, selection of optimal network routes, multicast distribution, and store-and-forward capabilities.

*brian@alanax.com

†wes@alanax.com

II. Software-Defined Networking

Software-Defined Networking (SDN) decouples the control and data planes of networking devices.[1] SDN enables the implementation of services and applications that manage and control the network through a software abstraction layer on a centralized controller, which provides a holistic view of the entire underlying network infrastructure. The SDN Controller handles the control logic of the network and interacts with network elements through a southbound interface called the Control-to-Data-Plane Interface (CDPI). OpenFlow is a well-known example of a CDPI protocol, and has become widely supported across the current generation of terrestrial router and switch hardware, as well as in software switch platforms.

Beyond the capabilities of legacy control plane protocols (such as intradomain routing protocols), SDN gives network owners and operators a more flexible and programmable control of their infrastructure. This allow them to use standard cross-vendor interfaces in order pursue customization and optimization, and reducing the overall capital and operational costs.

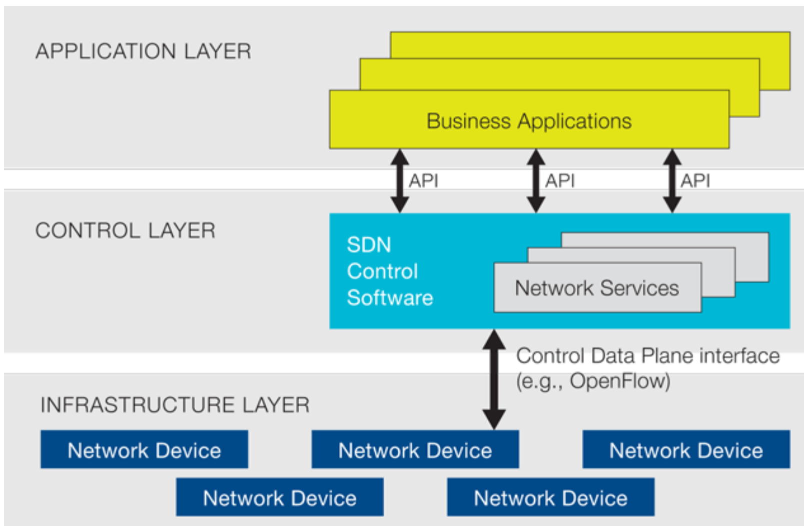

Figure 1. A high-level overview of the software-defined networking architecture[2]

Recently, new requirements have led to proposals to extend this concept to Software-Defined Wireless Networks (SDWN),[3] which decouple radio control functions, such as spectrum management, mobility management, and interference management, from the radio data-plane.

Advantages of SDN for wireline networks that can be extended to wireless networks through SDWN techniques include:

- Slicing of flows and network resources in order to implement multiple isolated virtual networks on top of the same physical infrastructure. For spacecraft onboard networks, this can permit isolating experiment data and instrument flows from spacecraft control and housekeeping data, while acheiving SWaP efficiencies by sharing the same RF links, communications systems, and onboard buses. For wireless space networks, this can permit radio system reuse and sharing between multiple payloads (e.g. in a hosted payload). Within a constellation or mission involving multiple platforms, it can permit more efficient and aggressive sharing of spectrum among missions.

- Enabling a more flexible network control structure, involving either single or distributed controllers in multiple arrangements. Adjustments to the control plane composition are possible during and throughout operations and allow operators to trade-off link efficiencies, fault-tolerance, and other factors dynamically – and to bring-in or remove additional control-plane systems for diagnosis, debugging, or other temporary tasks without impacting the rest of the network.

- Traffic engineering within the wireless network, in order to perform analogous functions as supported by traffic engineering in wireline networks. This includes optimizing cost of paths selected, loading of individual resources, and meeting other policy objectives. This could also support easy shifting of flows and and spectrum-agile radios between different spectrum areas (e.g. Ka-band to Ku-band), for instance.

Beyond the basic SDWN features applicable to space networks, however, there are additional benefits that can be obtained by using SDN techniques to control a network with additional information about the physical platforms and their environment. In this paper, we propose *Temporospatial SDN*, which extends the SDN/SDWN paradigm even further to enable SDN applications to make network control decisions based on the location, motion, and orientation of assets in space (i.e. position vector, velocity vector, and attitude), the relationships between those assets and their constraints (e.g. pointing angle limits, planetary, structural, or other occlusions), and the quality of wireless communications as assets move through space and time (e.g. sources of RF interference, mutual interference, space weather, charged particle wells, etc).

III. Temporospatial SDN Applications

Work on Temporospatial SDN has led to identification of several types of SDN applications that are fundamentally enabled by the addition of temporospatial knowledge, not possible with vanilla SDN and SDWN technology. These include

A. Topology Management

The use of centralized controllers in aerospace network operations is not new. In satellite communications, scheduling controllers are commonly used to reserve ground stations and task antenna gimbals to track satellites. These controllers rely on satellite orbit propagation models to predict and deconflict access intervals for communications. Establishing wireless links between mechanically steered and highly directional antennas is difficult using decentralized, distributed protocols, and this motivates the use of a controller with a holistic view of the network to direct the establishment of directional, point-to-point links.

In the Temporospatial SDN paradigm, scheduling controllers are implemented as *SDN applications*; they have access to a holistic view of the time-dynamic modeled mobility of network elements (satellites, rovers, aircraft, balloons, ships, etc) in order to drive operational decisions about the time-dynamic topology of the network (antenna tasking schedule). Computing feasible schedules may involve models that include many aspects of the physical systems, including gimbal rates for slewing antennas, attitude control system parameters, ability to propagate orbits with tolerable accuracy, etc. None of these capabilities are part of existing SDN systems, but are fundamental to Temporospatial SDN, and enable controllers to produce and distribute network directives that match asset capabilities.

In addition to modeling time-dynamic mobility, time-dynamic prediction of network link accessibility and performance metrics through high-fidelity antenna pattern modeling, wireless signal propagation, and network protocol behavior is also possible.[4] Based on trajectory predictions and mission planning data, software can be used to model the time-dynamic position and orientation of aerospace vehicles; and, given these dynamic positions and orientations, along with the modeled characteristics and pointing of sensor, communications, and other payloads aboard the vehicles, the time-dynamic spatial relationships between all of the objects can be determined.[5] This includes predictions of line-of-sight visibility times, propagation delay, predicted signal-to-noise measures, and other communications system performance impacts. Environmental effects, such as weather, thermal noise due to beta angle, galactic noise, and atmospheric density may also be incorporated in assessing certain relationships, such as in determining communications link budgets.[5] By incorporating this entire corpus of information and predictions into the holistic view of the network provided by the Temporospatial SDN control layer, we can better optimize topology management and use automation to intelligently plan which of the subset of the many possible paths through the potential future networks should be utilized.

Links that are physically feasible but not needed can be pruned, and power, spectrum, momentum, fuel or other resources can be conserved. Failures can be proactively modeled or predicted (e.g. based on trending and analysis data) and their impact can be assessed automatically, with alternative routes pre-computed and ready in case a failure occurs. The combination of the physical information provided through Temporospatial SDN control software and the network programmability of SDN hardware fundamentally enables topology management functionality far beyond what existing satellite network control and scheduling systems support.

B. Packet Routing

A number of aerospace systems are currently being developed for the purposes of providing Internet access. In contrast to previous GEO systems and even existing LEO constellations, some of these new systems are attempting to operate with lower orbits or even altitudes within the atmosphere as a means to avoid the unacceptable end-to-end packet latency that is known to plague existing systems. However, lower altitudes usually come at the expense of greater mobility and a higher rate of change in the network topology. For example, the coverage area of an individual spot-beam on a Low-Earth Orbiting (LEO) satellite may traverse a municipality in a matter of seconds or minutes. If network routing is unable to keep up with the rate of change in the network topology, unacceptably long periods of packet loss could occur, resulting in exponential back-off or broken sockets for Transmission Control Protocol (TCP) application flows, and dropped calls or extreme delays in voice and video applications.

Temporospatial SDN can solve such problems. By predicting the near-term trajectory of aerospace network platforms (i.e., using satellite orbit propagators) in the SDN control layer and providing this information to the SDN application layer, an SDN routing application could anticipate topology changes and route breakages before they occur. It can then utilize modeled link metrics to proactively select a new route and and modify the flow rules accordingly, throughout the network, at the appropriate time.

This is possible even in large systems, such as hundreds or thousands of moving platforms, because their positions and communications capabilites can typically be predicted in advance based on propagating the orbital elements for spacecraft, flight paths for aircraft and balloons, etc., and combining the knowledge of their future positions with other environmental data.

Network control via routing protocols (such as OSPF, ISIS, or any of the mobile ad-hoc networking protocols) is generally reactive to changes in the physical environment after they have occured and the occurrence is detected. They do not take into account known degradations that will occur and proactively re-route traffic ahead of the disturbances or changes to the topology. Temporospatial SDN enables this, and is also more efficient than relying on routing protocols because there is no signalling ”chatter” between the network elements in the data plane in order to probe and detect connectivity properties. There are only low-bandwidth control messages providing direction from the Temporospatial SDN control system.

C. Interference Avoidance

Aerospace communication systems often rely on regulated RF spectrum. In order to be permitted to operate a new aerospace communication system, the operator must demonstrate that the new system will not interfere with incumbent systems. For example, OneWeb’s priority spectrum allocated from the International Telecommunications Union (ITU) levied the constraint that their use of the spectrum must not cause interference with GEO satellites.

Temporospatial SDN can be used to respect spectrum constraints and to avoid interference. The SDN control layer can model constraints on link accessibility and designate hypothetical links as inaccessible during windows of time in which their use would otherwise violate RF spectrum allocation rules. Any other SDN application, such as those responsible for topology and route management, would then react to reorganize the network topology and routing around the inaccessible link intervals while still meeting business objectives.

Additional sources of interference, such as cellular infrastructure, radar sites, solar and cosmic objects can be inferred as well, based on their presense in the line-of-sight cones of antenna objects. Additionally, since their actions are coordinated by the Temporospatial SDN control software, interference that might occur due to orientation of the aerospace network objects among themselves can be predicted and avoided, since most assets within a constellation or other cooperative system will be operating within the same range of spectrum.

IV. Conclusions

In this paper, we introduced Temporospatial SDN as an extension to SDN/SDWN that is intended to address the challenges of designing and operating complex aerospace communication networks. We proposed new SDN applications for aerospace network topology management, packet routing, and interference avoidance. As for future work, the authors are exploring the feasibility of implementing a Temporospatial SDN controller in Java based on the Open Network Operating System (ONOS) platform[1] and STK Components technology.[5]

References

[1]Berde, P., Gerola, M., Hart, J., Higuchi, Y., Kobayashi, M., Koide, T., Lantz, B., O'Connor, B., Radoslavov, P., Snow, W., et al., "ONOS: towards an open, distributed SDN OS," *Proceedings of the third workshop on Hot topics in software defined networking*, ACM, 2014, pp. 1–6.

[2]Committee, O. et al., "Software-defined networking: The new norm for networks," *Open Networking Foundation*, 2012.

[3]Costanzo, S., Galluccio, L., Morabito, G., and Palazzo, S., "Software Defined Wireless Networks: Unbridling SDNs," Oct 2012.

[4]Barritt, B. J. and Eddy, W. M., *Astrolink for Modeling, Simulation, and Operation of Aerospace Communication Networks*, American Institute of Aeronautics and Astronautics, 2014.

[5]AGI, "STK Components," August 2015.