



Network Control Systems for Large-Scale Constellations

Jeroen Cappaert and Sreeja Nag

Contents

Introduction	2
Smallsat Constellations	3
Network Control System Functionality	4
Figures of Merit	6
Constraints	6
Challenges	7
Network Control Subsystems	9
Satellites	10
Ground Stations	12
Control Center	12
Operations Interfaces	16
Network Control with Intersatellite Links	17
Machine Learning and AI Applications	18
Conclusion	20
Cross-References	20
References	20

Abstract

Network control systems for large-scale constellations of small satellites are of critical importance to the successful and safe operation of smallsat systems – especially in the case of remote sensing networks but also for other applications. As the level of complexity of the space segment and the ground segment of these networks increases so does the level of complexity of the network control systems

J. Cappaert (✉)
Spire Global, Glasgow, UK
e-mail: jeroen@spire.com

S. Nag
Nasa Ames Research Center and Baer Institute, Mountain View, CA, USA
e-mail: sreejanag@alum.mit.edu

that are used. Today small satellite constellations such as those represented by Planet and Spire represent the two largest small satellite constellations in operation, but in the near future, there are many new constellations that may require network control of many thousands of smallsats. There are challenges associated with initially deployment, configuration of spare satellites, operation of large constellation, and de-commissioning and de-orbit of smallsats that require a high level of sophistication and technical expertise. All of these levels of complexities and methods to create effective network control systems for large-scale constellations are addressed in this chapter.

Keywords

Artificial intelligence (AI) · Machine learning · Mission control · Scheduling · Planning · Tasking · Figures of merit · Delay-tolerant networking · Software-defined networking (SDN)

Introduction

In recent years, there has been a proliferation of smallsat constellations. There are now over a dozen companies now having fully operational constellations and commercializing their telecommunications, networking, automatic identification system (AIS), remote sensing, and RF geolocation products. These small satellite constellation operators now have a special focus on how to achieve the most efficient and optimal use of their spacecraft assets. Most modern smallsat constellations are no longer operated on an on-demand basis. Instead most of these smallsat systems collect or communicate data on the basis of software programming. Thus most of these spacecraft operate at their own pace. Further in the case of remote sensing and other systems, there is a trend toward processing of data on-board the satellite to greatest extent possible. If it is possible to pre-process the data on-board before bringing it to the ground, it can lead to greater efficiencies at several levels.

The goal of a network control system for a smallsat network thus typically becomes to translate direct customer demands into the operational plans for the satellite networks without commands or specific requests for data coming from the ground. This process is called, especially for remote sensing satellite systems, asset tasking. This method of optimizing the usage of the space and ground segment assets via network control software is key to the successful operation of smallsat networks (see Fig. 1). As such, network control is now central to achieving mission and business success. In short optimized network control systems are key to determining achieving the highest return on investment and the most efficient asset allocation in orbit and on the ground.




Fig. 1 Network control as part of business process

Smallsat Constellations

A Euroconsult study has projected that, over the next 10 years, about 8600 smallsats will be launched, at an average of 835 per year as of 2023. They also project that this will grow to an average of 880 launched per year by 2028.

According to the Euroconsult study, 322 smallsats were launched in 2018. Of these launches about 40% were for constellations (Euroconsult 2019).

Proposed constellation sizes range from 10s of satellites to 100s and 1000s. A few operational constellations already exist today in the range of 10s and 100s of satellites (e.g., Planet, Spire).

This trend toward the use of more efficient software, optimized network control systems, and the use of pre-processing where possible is not restricted to smallsat networks. These same dynamics are also driving fundamental change in the way large networks of satellites are operated.

In the past it was feasible to manage and control a few satellites by means of an on-the-ground team of operators. It was even for larger networks of perhaps 20 in number by means of simple automation and a scripting process followed by manual operators on the ground. This approach however becomes increasingly unwieldy for larger constellations. Human errors – sometimes called cockpit errors – become more likely as network operations grow in size. There is a threshold – of perhaps a constellation of 50 satellites – that forces a change to a completely different level of automation.

A constellation of 50 or more, such as was first demonstrated by the Iridium and Globalstar constellations, requires a different approach. At this level of complexity, the satellites can no longer be thought of as individual assets. This is because the

complexity of interaction is increasing exponentially, which can be shown analytically (Alderson and Doyle 2010). A network of this size has to be considered as a holistic combine. It becomes necessary for space assets to be continually optimized. Automated network control can then be programmed to provide maximum product value.

One must start from a purely practical perspective to consider what satellite operators are able to achieve. It is very challenging for operators to constantly monitor and in real time resolve anomalies that occur within a constellation larger than perhaps 20 satellites operating in low Earth orbit (LEO) and truly impossible at a level of 50 satellites. This is because the complexity of interaction is increasing exponentially. If one considers the case of a modern smallsat constellation, the complexity becomes very great indeed. Such a constellation could have:

- Greater than 100 satellites. These might be distributed in many orbital planes.
- There could be ground station at 50 or more locations.
- There might be well over 1000 satellite contacts per day.
- The level of communications throughput and data download might be collectively measured in gigabytes per second.
- Satellites might have varying data requirements and have been programmed to provide multiple services or data products.

With this amount of complexity, manual control essentially becomes impossible. Automated and programmed systems are needed for network control. This then allows operators to focus on providing the needed responses and inputs to the system. The role of human operators thus becomes that of balancing business priorities and re-acting to specific issues or anomalies that might require attention.

Network Control System Functionality

All these elements of optimized systems operations are shown in Fig. 2 with the many parts of a network control system depicted. There are many aspects to consider. These elements include but are not limited to the following (Cho et al. 2018).

- **Satellites:** a typical smallsat constellation and the satellites deployed in it are built up over time. This may include the use of rideshare launches. In such cases the constellation can be less “managed” than traditional constellations that can utilize dedicated launches.
- **Ground stations:** Smallsat operators typically rely on a ground station network that is not monolithic in nature. Rather these ground stations may well be constituted by multiple owners and service providers. The result of this diversity of providers and terminals is a ground network that has capabilities that might widely differ. For example, on some stations there might just be a downlink available, while other stations can send commands as well as receive data.

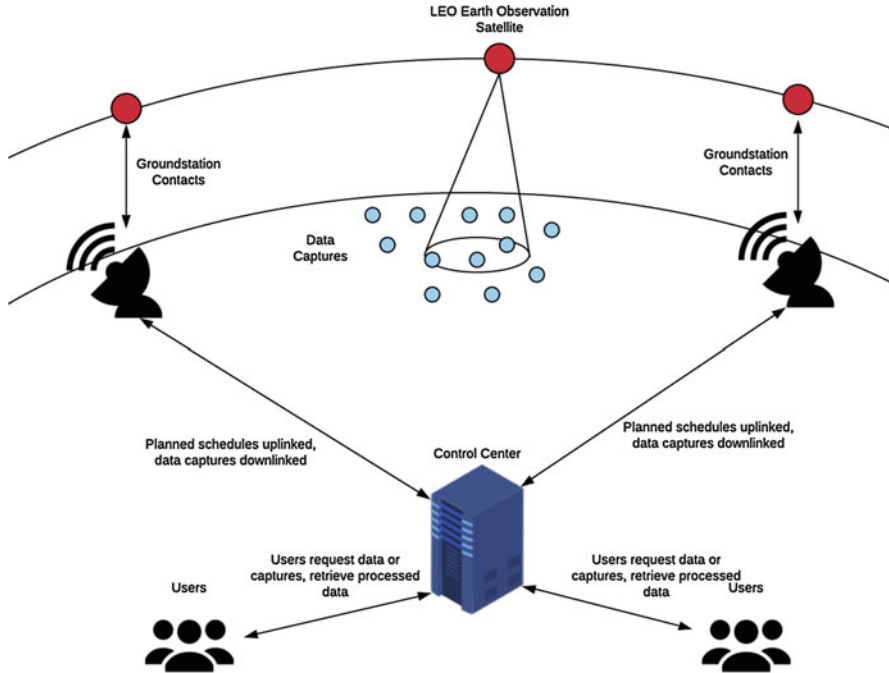


Fig. 2 Illustration of the elements of a network control system and need for flexible network control

- **Mission control:** The mission control centers can also differ. One facility could be a traditional physical mission control center. Others might be designed as completely cloud-based distributed control center. Such a facility could have limited on-shift operators.
- **Schedules:** Scheduling is also key concept that drives asset allocation. Transit and payload schedules are traditionally generated centrally from the mission control center. In modern systems, however, network control varies depending on whether spacecraft are programmed with on-board satellite autonomy.
- **Customers/users:** Internal users and customers have varying needs. These may vary by season or other variables. Factors that might vary include changing number of measurements per day or once per day, week, month, etc. There can be some form of a one-off on-demand requirement such as a measurement at a specific time and place. Users could also specify multiple interfaces as to where and when to retrieve data.

Smallsat constellations come in various forms and sizes and now comprise many applications, but generally the goal of network control can be generalized to the following:

- Optimize both the ground and space segment assets so as to provide the maximum return on investment.
- Ability to exploit the satellite system to allow maximum flexibility to respond to customer demands and changing market needs.

In the case of a communications constellation, the goal is typically to determine how to route given volumes of throughput optimally. Alternatively, the objective for an Earth observation constellation is to schedule data collections efficiently as well as to find the best route to relay data to the ground. Both problems are highly interconnected to routing efficiency. The use of network control systems to achieve routing efficiency is considered in the rest of this analysis.

Figures of Merit

A figure of merit provides a way to best gauge the performance of a network control system and is especially key for Earth observation services. The rationale for establishing a figure of merit typically considers one or more of the following criteria.

- **Revisit time:** This is the time between subsequent observation and/or coverage of the same geographical area. This is key for both networking and Earth Observation services. It's an important aspect as it defines the intervals during which coverage or service is not provided.
- **Refresh time:** This is the time between subsequent observations of the same asset (e.g., in the case of asset tracking like automatic identification systems (AIS), automatic dependent surveillance (ADS-B), or machine-to-machine (M2M) communications).
- **Latency:** This represents the time between an observation and data delivery to customer or the time to provide data networking connectivity. Sometimes this also includes the time required to command a spacecraft to take a certain measurement.
- **Coverage:** This provides a precise measure of the geographical coverage of a satellite constellation. This defines the percentage of the Earth's surface (or an area of interest) that is covered. Similarly, asset coverage can be defined as the ratio of assets that is covered (ships, buildings, airplanes).

Constraints

A network control system thus needs to be flexible and responsive in order to take into account various operational constraints. The principal constraints include:

- **Energy constraints:** It is important to ensure that the stored energy in the spacecraft battery is always sufficient to meet power needs of the mission and

the spacecraft bus as well as meet the associated needs during the times of solar eclipse. All scheduling of spacecraft activities needs to take into account the satellite's power budget. There also needs to be some reasonable margin. In short scheduling of service must always respect power availability plus some margin.

- **Thermal constraints:** There are also thermal constraints that must be considered in the provision of services. Each of the spacecraft's subsystems must be used only if the temperature expected in that interval is within an acceptable range. If a particular payload observation or communication windows are found out of range, then this service capability must be cancelled.
- **Data balance or storage capacity:** There must be care taken to ensure that the storage space remaining on the spacecraft is more than adequate to meet mission objectives. This includes the capability to always be aware of the storage capacity and downlinking capabilities in terms of data rate and windows that are available for such data transfers.
- **Subsystem compatibility constraints:** There can be competing power, thermal, or other constraints that must be considered. This means that network control systems must be programmed so that non-compatible subsystems are not used at the same time or that all operational requests are compatible for the times they are in use. For example, high-power transmitters and a sensitive receiver might operate effectively at the same time. Other payload elements might have competing pointing, thermal, or other requirements.
- **Geographical constraints:** In some cases there are constraints related to operating over certain geographical areas. For example, one might want to operate payloads only over land, deserts, oceans, or arctic regions. Additionally, an operator might want to restrict service in polar regions so as to recharge batteries. Thus the geographic constraint would hinge on whether is above highly trafficked areas or not.
- **Regulatory constraints:** It is also important to make sure that all operations are conducted within the limits of space operation frameworks and applicable licenses. Satellite communication services are restricted in many ways, but there are also constraints on Earth observation satellites and other applications. Those who are launching communications satellite systems might have to coordinate with other operators on a non-interference basis. This can result in the need to duty cycle radios, implement communications blackout windows, or point away for geostationary satellites, which have protected status, to reduce the risk of interference. Additionally, operators must take care not to exceed power flux density limits set by the ITU since many satellites might be communicating at the same time. Any network control systems need to take these constraints into account.

Challenges

Some of the main challenges for a satellite network control system are examined below.

Heterogeneity

Smallsat constellations, which include hundreds or more satellites, are deployed in different generations. Thus these satellites may have different capabilities. Thus these networks might contain satellites that are in some ways heterogeneous in design. This is due to the iterative development process over time. Constellations are built up piecewise as hardware is developed and improved. Geosynchronous satellites might have lifetimes of 15 to 18 years, but constellations with LEO satellites may be replenished much more often than traditional satellite systems. This results in a constellation that can be composed of dozens of hardware versions, some of which might be minor, while others might be more significant.

Another reason for differences in hardware of satellite design, or even operational frequency bands, can be a satellite operator purchasing satellites from other networks. Sometimes acquisitions can come from completely different satellite networks (e.g., Planet and the initial SkySat constellation).

Smallsat systems can also be different, not just in terms of hardware differences, but there can also be a variety of software versions across the constellation. A single satellite might be composed of subsystems each running their own software version and each being software-upgradeable. This is a manageable issue, but it requires considerable discipline to make it work. Unique software configurations might arise as well as satellites experience key issues and/or major anomalies. There may be the need to create workaround solutions that are unique to one satellite or a particular cluster of smallsats. Every satellite tends to develop a unique “personality,” based on the specific hardware and software features and potential anomalies or workarounds in place (Open Networking Foundation 2012).

Any network control system needs to take into account these differences when scheduling its assets.

Orbital Patterns

While some smallsat constellations are actively managed and station is kept to special parameters, many are not. Other constellations are more irregular and rely on available rideshare opportunities. For example, see Fig. 3a–c. These figures provide a comparison between the Spire, Planet, and Iridium orbits. The difference in launch strategy results in satellite to ground station transit patterns that are quite different for these three constellations. Satellites might see a ground station a few times an hour regularly for a while, only to have hours-long contact gaps later.

This usually results in figure of merit distributions with a long tail. For example, Fig. 4 shows the latency distribution of the Spire constellation and ground network. While latency is normally very good, due to the orbital configuration, sometimes there are large contact gaps between satellites, resulting in a few periods with long latency that skew the average latency.

Multi-customer Nature

Many modern constellations no longer serve one type of customer or might even operationally produce more than one type of data from multiple payloads. The need




Fig. 3 (a) Spire's LEMUR constellation. (b) Planet's Flock constellation. (c) Iridium

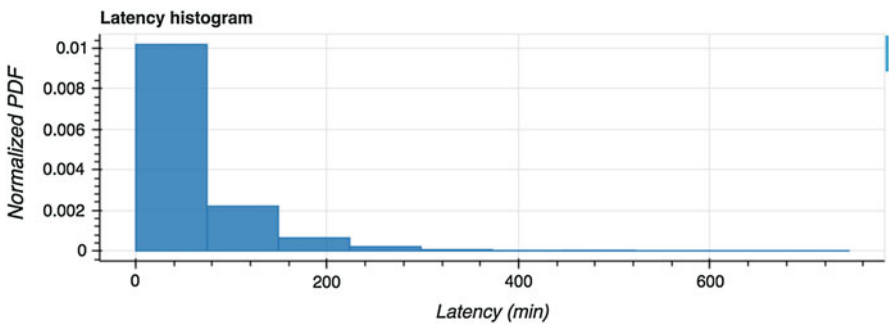


Fig. 4 Latency patterns for the Spire constellation (Graphic courtesy of Spire)

to manage and control the competing priorities between various data types can be difficult. These challenges become even greater as various data types are added, the volume of data transfer grows, and latency requirements differ. Additionally, within a single data type, multiple different service levels could and often do exist. For example, captures relating to emergencies (e.g., imagery for natural disasters, ship data following collisions, etc.) could have higher priority needs than “standard” captures. Any network control system thus needs to be flexible so it can consider various data types and different tiers of data. One might end up with a latency-data volume curve as shown in Fig. 5, where various tiers of data volumes are down-linked at various latencies as well as different data speeds.

Network Control Subsystems

- On Fig. 6 an example is shown of what a constellation network control system might look like. The various components are discussed (Cho et al. 2018) below.




Fig. 5 Example of latency-volume curve (Graphic courtesy of Spire)

Satellites

On-Board Automation

The main goals of the on-board automation include at least the following four objectives:

- Checkout and commissioning of satellites:** The efficient deployment and commissioning of multiple satellites in a constellation is a demanding task. This can mean the deployment, checkout, and commissioning of multiple satellites that are deployed from one launch vehicle. Thus the managing satellite commissioning for many, many satellites across dozens of vehicles over a wide span of time is cumbersome. For example, Planet moved to a fully automated commissioning system which was probably essential since they were involved with the launch, checkout, and commissioning of 88 satellites associated with a single launch by a Polar Satellite Launch Vehicle (Doan et al. 2017).
- Tasking payloads and communication radios.** For many small satellite constellations, there is a formal process to create “tasking schedules.” These are typically generated by the mission control system on the ground and synchronized periodically with the spacecraft for a specific time interval (e.g., 24 h). Various levels of automation might exist on the spacecraft. Such automation

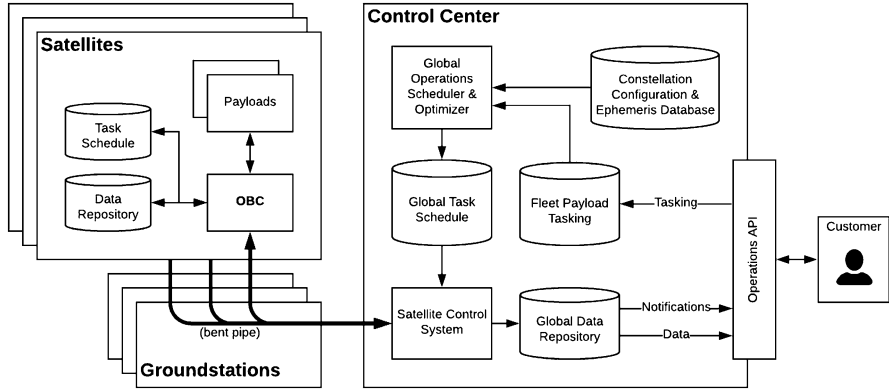


Fig. 6 Components of typical network control system

serves to make this task easier and more efficient. In the case of automation, a task list might be quite simple. There might just be a list of commands with timestamps (e.g., “turn on payload one”). Alternatively there might be a higher-level list of subtasks (e.g., “complete a payload collection on payload one and prepare the data for downlink”). In some of the more advanced network control systems, more autonomy might be given to the spacecraft. In this case there might be a more flexible task provided. This might result in a tasking such as “Complete 10 payload windows with optimal geographic distribution in the next 24hrs.”

- **Managing spacecraft data.** In the case of different payload data types, there might be a need for different types of data handling. Data might exist in various forms and conditions. These might include ring buffers, databases, and plain text or binary files. For example, native data type will be allowed to determine the best storage tool. In the case of IoT message, it might be best to create a database. In the case of raw radio frequency collection, this process might work better with binary files. Data size and data latency are among the factors that contribute to these decisions. In the case of data with low latency requirements, it is best to avoid large files that would take a long time to download. The on-board computer can also apply data compression where appropriate and pre-process the data to prepare it for downlinking. Thus it is best to encode large files in advance of downlink.
- **Collect telemetry, monitor system health, and recover where needed.** Lastly, it is the task of the on-board computer to monitor health for all subsystems. It is also tasked with the responsibility to collect system telemetry and perform fault detection and isolation and recovery as required. Traditionally the on-board computer would be tasked to “flag” any faults and transmit them to the ground in order to alert operators. In more advanced systems, especially in the form of 0 recurring and simpler faults, this might be programmed to be handled through an automated recovery process. The latest systems employ modern machine

learning telemetry systems. This allows longer-term trending of telemetry data and a greater ability to identify indicators of impending system failures.

Ground Stations

In most systems, the ground stations act simply in bent-pipe mode. This means that data is transmitted directly to the control center without buffering or processing. The advantages of this architecture would be in the form of enhanced security. This is to say that no data would ever be left unencrypted on a ground station. It would also decrease latency in that data can be streamed directly to the control center without having to wait to complete a satellite pass.

Control Center

Configuration and Ephemeris Databases

As indicated above, each satellite usually ends up having what might be called a unique personality. This then results in the need for a per-satellite configuration database. Such a detailed database can keep track of things like (i) satellite frequency configuration, (ii) licensing jurisdiction, (iii) status of subsystems, (iv) status of watchdogs, (v) timestamps of the last time maintenance procedures were executed, (vi) software interface versions, (vii) ADCS control modes, (viii) telemetry alerting limits, and other data peculiar to an individual satellite. Whenever tasks are scheduled and executed, the satellite configuration database is used to determine how to interact with a specific satellite. This would include decision as to what software interfaces to use. This key database would also contain the ephemeris information for satellites.

In addition, the database can contain ground station characteristics. This would alert the scheduler as to which satellites are compatible with which ground stations.

Operations Scheduler and Optimizer

To make the best use of all assets, a globally optimal ground system transit and on-board tasking schedule needs to be generated.

The block diagram shown in Fig. 7 shows how an optimizer fits into the large satellite control process. The input for the optimization process is derived from business-level goal functions. These might take the form of “Collect X MB of data at Y min of latency” or “optimally cover a specific area of interest.” Business objectives can also change over time. Thus, the goal functions are not necessarily static. High-level functions need to be translated into actionable satellite operations interface variables. This is where figures of merit are employed. Thus figure of merit conditions are then fed into the optimization process. The output from this process will be translated into constellation control parameters. These then become specific tasking schedules, the setting of telemetry limits, etc. This ultimately becomes a circular process after these parameters are executed and telemetry data collected.

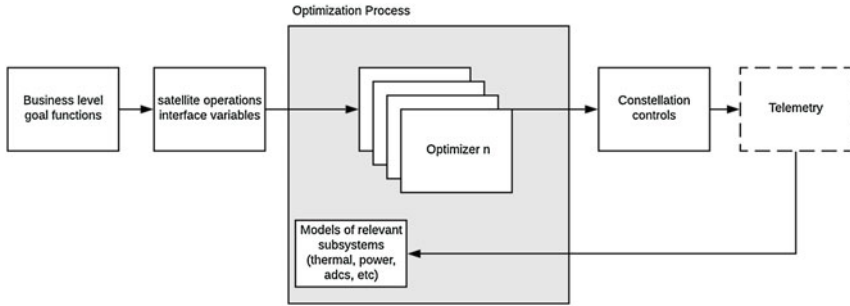


Fig. 7 Example of optimization process

This informs the control center on the level of success achieved compared to the initial goals.

Optimization Strategy

On-Board Optimization

One option is for the ground to output contact schedules as well as a set of objectives for each spacecraft which in theory it can autonomously pursue. Each spacecraft would in this case need its own optimizer. This would include its own thermal and power models, priority maps, etc. This optimizer would need to computationally be programmed to run efficiently on the on-board computer.

Some advantages with this approach are relatively simple ground-based scheduling; more resiliency against unpredictable hardware behavior that was not foreseen during scheduling, such as an instrument failing to turn on; and ability to react to evolving phenomena to observe.

On the other hand, there are disadvantages with this approach: less fine-grained control with a less deterministic outcome that might potentially result in performance further away from a global optima and the higher levels of on-board autonomy, requiring more complex satellite side code.

Ground-Based Optimization

This approach entails an entire schedule (as optimized on the ground) to be uplinked to each satellite. This would be expected to be followed under nominal operations. This has the advantage that it provides very fine-grained control over what the constellation is asked to do with a deterministic outcome.

Hybrid Optimization

A third option would be a hybrid optimization that bridges the above two approaches. This approach generates a globally optimal schedule on the ground. Yet there is also enough intelligence on the satellite to intervene when operational conditions deviate from assumptions made during the generation of the ground-based schedule.

Such an intervention might, for example, allow satellite to throttle back instrument usage if the amount of data throughput is backing up. It might also activate an instrument to expend energy stored in excess of what was expected.

Algorithms in Literature

A number of planning and scheduling algorithms have been published, and at a very high level, in open literature for single large satellite missions. Key examples are represented by Automated Scheduling and Planning Environment (ASPEN) for EO-1 satellite. Another example is the Advanced Spaceborne Thermal Emission and Reflectance Radiometer (ASTER) that is utilized on the Terra satellite as well as the high-resolution imagery from the IKONOS commercial satellite. It is also used for scheduling observations for the geostationary GEO-CAPE satellite or image strips over Taiwan by ROCSAT-2.

Stochastic algorithms have been proposed and computationally simulated for single spacecraft (Xhafa et al. 2012) and multiple payloads (Jian and Cheng 2008) and comparative merits documented for satellite fleets (Globus et al. 2002). While they offer accurate solutions, the cost of initial condition dependence, exponential time to converge, and large training sets make them very limited in mission applications. (Abramson 2012; Robinson 2017) have developed a coordinated planner that can handle a continuous stream of image requests from users by finding opportunities of collection and scheduling air or space assets to maximize collected utility. Agent-based autonomous scheduling has been implemented on NASA's Deep Space 1 (Schetter et al. 2003).

Some analysts have (Bunkheila et al. 2016) demonstrated the ability to provide for the scheduling of the scan of single LEO orbiting satellite by dividing the areas of interest and choose the sequence of strips of coverage by this method. This is based on various distance functions for an ideal orbit. It fixed scanning times based on roll and pitch angles and is therefore able to compute the temporal feasibility of pointing without modeling the Attitude Determination and Control (ADC) system or its uncertainties.

Simulation studies have optimized the scheduling for single Cubesat downlink to a network of ground stations (Spangelo and Cutler 2012) or multiple payloads' downlink to existing stations (Jian and Cheng 2008), within available storage and energy and access time constraints. Studies have also combined single satellite scheduling with information sharing across satellites for a weak consensus on feasible charging, downlink, and observation schedules (Kennedy et al. 2015) using fixed view imagers.

Scheduling observations for constellations of large satellites with payload re-pointing has been formulated for the French Pléiades constellation (Lemaitre et al. 2002) (Damiani et al. 2005) and COSMO-SkyMed constellation of synthetic aperture radars (Bianchessi and Righini 2008). Schedulers for Cubesat constellations such as the 200+ Dove spacecraft fleet operated by Planet Labs, Inc. (Boshuizen et al. 2014) assume static orientation of the sensor in orbit and only schedule duty cycles for payload power. Accounting for full reorientation in multi-spacecraft missions imposes computationally expensive constraints on scheduling spacecraft slews

between payload operations. It is only recently that scheduling with slew-time variations has shown reasonable convergence using hierarchical division of assignment (He et al. 2019) and step-and-stare approaches using matrix imagers (Shao et al. 2018) and dynamic programming to account for utility maximization with ADC modeling (Nag et al. 2018). Planet Labs has published a preliminary scheduler used to operate their agile Skybox spacecraft fleet (Augenstein et al. 2016). Spire uses a similar approach for their RF sensing fleet. Agile observation schedulers that can recompute science value real time and autonomously reschedule observations across a constellation have demonstrated utility in simulation (Nag et al. 2019).

EOS is well suited for the mixed integer linear programming (MILP) approach because to perform any activity at any time instant (or not) can be modeled as a binary integer (e.g., ROCSAT-2 scheduler), and CPLEX can solve such formulations efficiently. However, MILP allows for only linear constraints and a single objective function, and there is no guarantee of reaching an optimum in linear time. The first drawback can be addressed using linear bounds to otherwise nonlinear variables (under-constrained formulation). The lack of a single, linear objective can be addressed using a Lagrangian sum of multiple objectives or a convex function representation of the objective. However, such approximations take away from accuracy and cutting plane bounds need not always be reliable. As examples, MILP has been successfully formulated for EOS for PLEIADES and the SPOT series, GEO-CAPE, and adapted as Constraint-based Interval Planning in the EUROPA planner for Deep Space 1. Constraint programming is not restricted to integer variables and linear equations/inequalities. Variables can be intervals or anything in the finite domain and constraints can be arithmetic or symbolic.

EOS can also be framed as an orienteering problem because it is a selective traveling salesman problem (TSP) where agents are required to visit as many checkpoints as possible within a time frame, each associated with a weight. The target captures can be assigned weights and orbital or subsystem constraints set up to represent travel times between captures and the objective is to maximize the weighted sum. The prize-collecting TSP not only minimizes travel time but also penalizes for unvisited cities. However, the interplay between orbital access to required captures by fast-moving LEO satellites and ACS-dependent slewing times causes the time for the traveling salesman to go from one capture to another to be highly dependent on the absolute time either capture is accessed. Variable time further adds to TSP solution complexity.

While dynamic scheduling across the full state space is known to generate an astronomical number of paths, unsolvable in polynomial time, branch-and-bound-like approximations to the dynamic programming (DP) approach have demonstrated a practically suboptimal resolution of the NP-hard non-restricted problem. DP has also been applied to the dynamic scheduling problem after it is decomposed into several simple, static problems defined by a rolling horizon, which is automatically appended with waiting tasks as current tasks are completed.

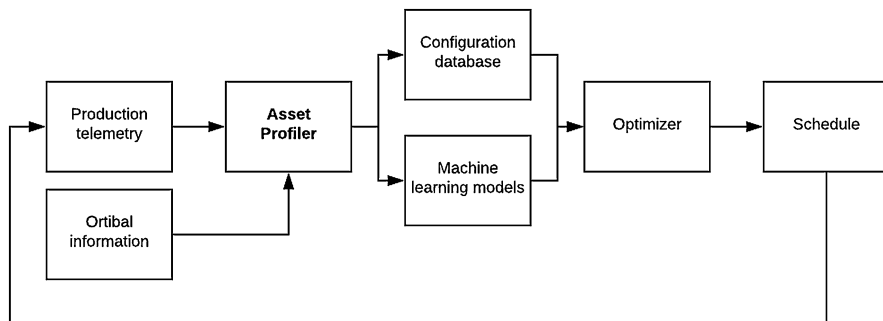


Fig. 8 Example of asset profiling process

Asset Profiling

To be able to optimize assets, one must know the capabilities of these assets. While a good first guess is to use an analytical approach based on theoretical performance of assets, the asset performance changes over time (e.g., degradation), there can be variability in the actual performance between satellites, and different satellites might have different constraints (e.g., performance differs between different orbital planes). Because of these issues, it's possible to integrate automated asset profiling into the optimization process. Based on real satellite telemetry, expected figures of merit are determined that can then be fed into an optimizer. Using machine learning techniques, it's also possible to anticipate asset performance over time. It's also possible to alert on sudden changes in performance characteristics (e.g., a drop in downlink volume or collected observations). Figure 8 shows an example of asset profiling integration.

An example of a plot that could be obtained by an asset profiler is shown in Fig. 9. This example shows a model of the payload data generation rates on-board a satellite. As time goes by, the data volume on board increases until the satellite is in contact with a ground station and the data buckets are emptied. These models can be used to determine which satellite is most in need of a ground station contact, to optimize for data latency.

Operations Interfaces

While in a large-scale operations system, satellite operators cannot directly interface with each satellite all of the time, certain interfaces are required to ensure they can input priorities into the system, react to issues, and monitor health.

Priority and Request Management

Typically, any network control system will have a method of interfacing with the inputs of the optimization process. It is the task of operators to translate business priorities into actionable goal functions for the constellation. This may take the form




Fig. 9 Example of asset profiling model

of defining constellating wide figures of merit (e.g., system data throughput, or data collected under a certain latency target). When ad hoc requests come in, operators normally have override mechanism to command the satellites as needed outside the process.

Health Monitoring and Incident Management

Given the level of automation present, the main job of satellite operators is not to directly command or task assets but rather to monitor the constellation for any anomalies that might occur and manage those appropriately. To be able to do this effectively, an incident management system is required that can link back to operational data, on-ground test results, and any other information that can help resolve the issues at hand. Operators can then feed this information back to engineering development teams as new satellite versions are being developed.

Network Control with Intersatellite Links

If intersatellite links are present on a constellation, the concept of operations is changed quite drastically. It should be noted that the presence of intersatellite links doesn't necessarily imply a real-time connection between satellites and the ground. One particularly interesting concept is the delay- or disruption-tolerant network. This type of system has been designed for networks lacking continuous connectivity (Nag et al. 2019). This approach significantly decreases the latency associated with data getting from source to destination while tolerating intermittent connectivity.

Similarly, the advent of software-defined networks (concept shown in Fig. 10) has provided an abstraction toolset that can simplify the design of networked constellations. Some of the benefits include isolating multiple virtual networks, isolating various data flows (e.g., TT&C vs. data), more efficient sharing of spectrum, and

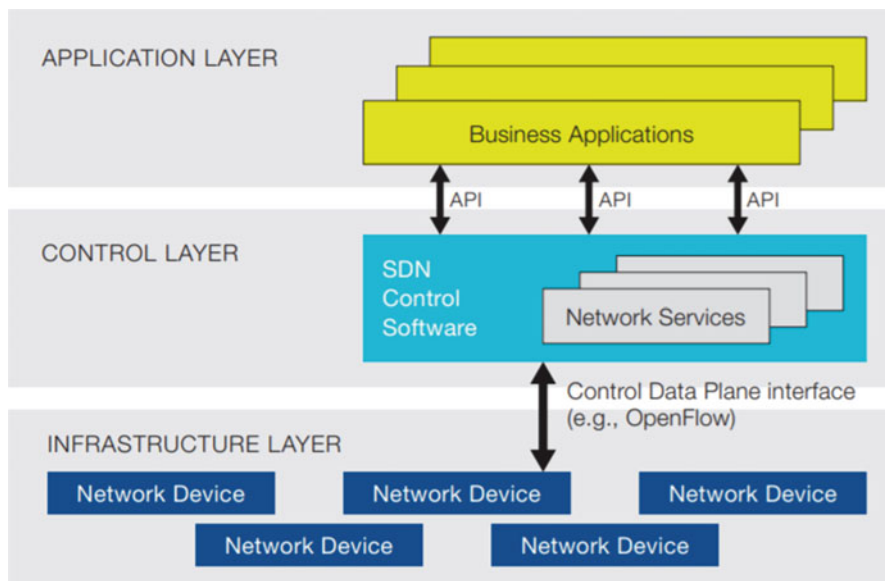


Fig. 10 Software-defined networking concept (Open Networking Foundation 2012)

enabling a more flexible network control structure. Additionally temporospatial SDNs enable additional benefits such as advanced packet routing and usage of orbit and attitude knowledge in data routing (Barrit and Eddy 2015).

Various proposed data routing algorithms are also available in literature for intersatellite-link constellations, e.g., Fig. 11 (Bunkheila et al. 2016).

Machine Learning and AI Applications

A few uses of machine learning and AI applications in network control systems have already been mentioned above, such as telemetry trending or asset profiling, but these are far from the only applications, and as larger constellations are deployed, more and more useful applications appear. This is particularly true in the area of remote sensing and Earth observation.

Machine learning and AI applications are driving innovations in on-board data processing or simply pre-processing. Mechanically, this explosive growth in launches of private and public satellites has induced an exceedingly rapid growth for space-generated data. ESA observation satellites reached requirements of close to 150 terabytes per day in 2018. This volume and definition of data has contributed to the across-the-board rise in users for space-generated data. The volume and reliability of this data have tremendously improved. Further the speed of delivery has been reduced to almost real-time capabilities for certain data categories. This process has created considerable issues in the management of the data flows, with delays,

Routing Algorithm	Routing Metric	Advantages
Handover optimized routing algorithm [88]	Connection matrix	Identifies the presence of inter-satellite links
Bandwidth delay satellite routing [89]	Delay and bandwidth	Ideal for LEO satellite networks
Destruction resistant routing algorithm [90]	Link state	Survivability of the network is enhanced
Steiner tree routing [91]	Number of hops	Limited overhead, supports a large number of satellites
Distributed multi-path routing [92]	N/A	Better end-to-end delay, instantaneous tracking of the changing topology of LEO satellite networks
Dynamic routing algorithm based on MANET [93]	N/A	Provides high autonomy, compatible functionality, limited overhead

Fig. 11 Available routing algorithms for intersatellite links (Radhakrishnan et al. 2016)

defaults, and crucial data reaching the user a long time after it ceased to be operationally useful.

However, the latest technological developments in chip-making have made possible the pre-processing of data directly at the satellite level. Thus data analytics in space has progressively reduced the amount of data to be downloaded back to ground stations.

This factor process of on-board processing of data improves constellation efficiency greatly. It also reduces the demands on staff, equipment, and other the infrastructure on the ground and also reduced the power speeds required for the satellite transmissions. Moreover, it allows the raising of the speed at which critical information is transmitted to users by allowing the satellite to prioritize independently which data will be downloaded first, thus decreasing latency for truly relevant information, as well as allowing the satellite to autonomously direct sensors and make time-critical decisions.

In parallel, many new machine vision chips have stimulated interest by allowing smaller platforms to accomplish many more analytics. This increased capability includes the ability to run complex pattern determination and identification programs. There is an increasing consensus across the space ecosystem that this development is potentially a game changer for many current space-based products. Many feel it is likely to create profoundly disruptive applications that will largely exceed the capabilities of current space-based platforms in generating large sets of “smart,” directly actionable data.

For example, the ESA-supported HyperScout-2 mission carried a Myriad 2 neural networking chip (Esposito et al. 2019). Spire’s LEMUR2 satellite is carrying a Nvidia Jetson AI computing device in an ESA-supported demonstration mission. The range of potential applications is wide and includes cloud detection and removal from imagery, RF fingerprinting, RF IQ analysis and de-interference, pattern detection, and many others.

Conclusion

Network control systems are often-neglected systems but are becoming more and more important as operators move from manual operation of a few satellites to operating large constellations where optical asset allocation is very non-transparent.

As complexity increases, more tools are required to offload operations from a team of operators to a true network control system. Part of the complexity of finding the optimal asset utilization for a constellation is the heterogeneity of a lot of smallsat networks. Modern network control systems leverage the appropriate tools to distribute the right levels of automation and autonomy between space and ground assets.

For constellations with intersatellite links, existing tools such as disruption-tolerant network (DTN) and software-defined networking (SDN) can be leveraged to enable packet routing, as compared to algorithm development from scratch.

Lastly, the application of machine learning and the use of newly developed AI-oriented chipsets provide a large opportunity to increase the on-board autonomy of spacecraft.

Cross-References

- ▶ [CleanSats and Efforts to Develop New Techniques for Space Debris Removal](#)
- ▶ [Planet's Dove Satellite Constellation](#)
- ▶ [Remove Debris Small Satellite Project](#)
- ▶ [The Brazilian Small Satellite Programme](#)
- ▶ [The Capella Satellite System](#)
- ▶ [The Kepler Satellite System](#)
- ▶ [The OneWeb Satellite System](#)

References

- M. Abramson, *Earth Phenomena Observation System (EPOS) for Coordination of Asynchronous Sensor Webs*. NASA Contract Report (2012). [https://esto.nasa.gov/2012test/conferences/estc-2002/Papers/ASP1\(Abramson\).pdf](https://esto.nasa.gov/2012test/conferences/estc-2002/Papers/ASP1(Abramson).pdf)
- D.L. Alderson, J.C. Doyle, Contrasting views of complexity and their implications for network-centric infrastructures. *IEEE Trans. Syst. Man Cybern. Part A Syst. Hum.* **40**(4), 839–852 (2010)
- S. Augenstein, A. Estanislao, E. Guere, S. Blaes, Optimal scheduling of a constellation of earth-imaging satellites, for maximal data throughput and efficient human management, in *Twenty-Sixth International Conference on Automated Planning and Scheduling*, 2016
- B.J. Barrit, W. Eddy, *Temporospatial SDN for Aerospace Communications*. Semantic Scholar (2015). <https://pdfs.semanticscholar.org/7da3/6c18aafb29b1620dc78a89d3e7e6115b5570.pdf>
- N. Bianchessi, G. Righini, Planning and scheduling algorithms for the COSMO-SkyMed constellation. *Aerosp. Sci. Technol.* **12**(7), 535–544 (2008)
- C. Boshuizen, J. Mason, P. Klupar, S. Spanhake, Results from the planet labs flock constellation. Presented at the AIAA/USU Small Satellite Conference, Logan, Utah, 2014

- F. Bunkheila et al., A new algorithm for agile satellite-based acquisition operations. *Acta Astronomica* (2016). <https://www.sciencedirect.com/science/article/abs/pii/S0094576515302629>
- D.-H. Cho et al., Optimization-based scheduling method for agile earth-observing satellite constellation. *J. Aerosp. Inf. Syst.* **15**(11), 1–16 (2018). <https://www.researchgate.net/publication/327982972>
- S. Damiani, G. Verfaillie, M.-C. Charmeau, An earth watching satellite constellation: How to manage a team of watching agents with limited communications, in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 2005, pp. 455–462
- D. Doan, R. Zimmerman et al., *Commissioning the World's Largest Satellite Constellation*. Semantic Scholar (2017). <https://www.semanticscholar.org/paper/Commissioning-the-World%E2%80%99s-Largest-Satellite-Doan-Zimmerman/77884ab008f482712c4a3d9457f6611f1006b391>
- M. Esposito et al., In-orbit demonstration of artificial intelligence applied to hyperspectral and thermal sensing from space, in *SPIE Conference*, Aug 2019. <https://spie.org/Publications/Proceedings/Paper/10.1117/12.2532262?SSO=1>. Last accessed 4 March 2020; <https://nebula.esa.int/content/delay-tolerant-network-flexible-communication-eo-satellites>. Last accessed 5 March 2020
- Euroconsult, *Prospects for the Small Satellite Market* (2019). http://www.euroconsult-ec.com/shop/index.php?id_product=116&controller=product. Last accessed 4 March 2020
- A. Globus, J. Crawford, J. Lohn, R. Morris, *Scheduling Earth Observing Fleets Using Evolutionary Algorithms: Problem Description and Approach* (Academia, 2002). https://www.academia.edu/28374353/Scheduling_earth_observing_fleets_using_evolutionary_algorithms_Problem_description_and_approach
- L. He, X.-L. Liu, Y.-W. Chen, L.-N. Xing, K. Liu, Hierarchical scheduling for real-time agile satellite task scheduling in a dynamic environment. *Adv. Space Res.* **63**(2), 897–912 (2019). <https://doi.org/10.1016/j.asr.2018.10.007>
- L. Jian, W. Cheng, Resource planning and scheduling of payload for satellite with genetic particles swarm optimization, in *IEEE Congress on Evolutionary Computation, 2008. CEC 2008. (IEEE World Congress on Computational Intelligence)*, 2008, pp. 199–203
- A. Kennedy, A. Marinan, K. Cahoy, J. Byrne, T. Cordeiro, Z. Decker, W. Marlow, S. Shea, Automated Resources-Constrained Science Planning for the MiRaTa Mission, 29th Annual AIAA/USU Conference on Small Satellites (2015)
- M. Lemaitre, G. Verfaillie, F. Jouhaud, J.-M. Lachiver, N. Bataille, Selecting and scheduling observations of agile satellites. *Aerosp. Sci. Technol.* **6**(5), 367–381 (2002)
- W.-C. Lin, D.-Y. Liao, C.-Y. Liu, Y.-Y. Lee, Daily imaging scheduling of an earth observation satellite. *IEEE Trans. Syst. Man Cybern. Part A Syst. Hum.* **35**(2), 213–223 (2005)
- S. Nag, A.S. Li, J.H. Merrick. Scheduling algorithms for rapid imaging using agile Cubesat constellations. *Advances in Space Research* **61**(3), 891–913 (2018)
- S. Nag, A. Li, V. Ravindra, M. Sanchez Net, K.M. Cheung, R. Lammers, B. Bledsoe, Autonomous scheduling of agile spacecraft constellations with delay tolerant networking for reactive imaging, in *International Conference on Planning and Scheduling (ICAPS) SPARK Applications Workshop*, Berkeley, 2019. <https://openreview.net/pdf?id=rJeDNZq5FE>
- Open Networking Foundation, *Software-Defined Networking: The New Norm for Networks*. 13 April 2012. White Paper from the Open Networking Foundation. <https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf>
- R. Radhakrishnan et al., *Survey of Inter-satellite Communication for Small Satellite Systems: Physical Layer to Network Layer View IEEE Explore* (Fourth Quarter 2016). <https://ieeexplore.ieee.org/document/7466793>
- E. Robinson, *Optimized Stochastic Coordinated Planning of Asynchronous Air and Space Assets*, 11 Jan 2017. <https://arc.aiaa.org/doi/abs/10.2514/1.1010415>
- T. Schetter, M. Campbell, D. Surka. Multiple agent-based autonomy for satellite constellations. *Artificial Intelligence* **145**(1-2), 147–180 (2003)

-
- E. Shao et al., Area coverage planning with 3-axis steerable, 2D framing sensors, in *Scheduling and Planning Applications Workshop, International Conference on Automated Planning and Scheduling*, Delft, 2018
- S. Spangelo, J. Cutler, *Optimization of Single-Satellite Operational Schedules Towards Enhanced Communication Capacity*, 16 Aug 2012. <https://arc.aiaa.org/doi/abs/10.2514/6.2012-4610>. Last accessed 5 March 2020
- F. Xhafa et al., *Genetic Algorithms for Satellite Scheduling Problems*. Semantic Scholars (2012). <https://pdfs.semanticscholar.org/70b1/afcf68533e712802d6cd5c1e42a0c2951438.pdf>