NOAA GPD Final Report
Focus Area 1: Constellation Mission Operation

Cognitive Space

November 27, 2024

This material is based upon work supported by the Dept of Commerce under Contract No. 1332KP24C0004. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Dept of Commerce.

Delivered to the U.S. Government with Unlimited Rights, as defined in FAR 52.227-14. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

Contents

1 Introduction **5**

2 Methodology **6**

2.1 Definitions . . . 7

2.2 Assumptions . . . 7

2.3 Optimization Method . . . 8

2.3.1 Graph Representation - Latency as a Constraint . . . 8

2.3.2 Simpler Representations - Latency as Objective . . . 9

2.4 Base Scenario Configuration . . . 10

2.5 Approach to Analysis . . . 11

3 28 Satellite Results **12**

3.1 Constellation Configuration . . . 12

3.2 Scenarios . . . 13

3.2.1 Base + KSAT Results . . . 14

3.2.2 X-band Conversion Results . . . 16

3.2.3 Phased Array Results - Fairbanks . . . 18

3.2.4 Phased Array Results - McMurdo . . . 20

3.2.5 Phased Array Results - Svalbard . . . 22

3.3 Summary . . . 24

4 100 Satellite Results **25**

4.1 Constellation Configuration . . . 25

4.2 Scenarios - Supporting One Contact per Revolution . . . 26

4.2.1 Scenario Overlap . . . 27

4.2.2 Supplementing a Phased Array at FCDAS . . . 30

4.3 Scenarios - Supporting Two Contacts per Revolution . . . 34

4.3.1 Scenario Overlap. . . . 34

4.3.2 Dual Phased Array Results . . . 35

4.4 Summary . . . 36

5 Discussion **37**

5.1 Summary of Results . . . 37

5.2 CNTIENT.Optimize . . . 37

5.2.1 Performance . . . 37

List of Figures

1 CNTIENT.Optimize schedule optimization process . . . 11

2 28 Satellite Constellation Configuration . . . 13

3 Antenna Schedule for Base + KSAT Scenario with 3 min Reconfiguration Time . . . 15

4 Antenna Schedule for Base + KSAT Scenario with 10 min Reconfiguration Time . . . 15

5 Breakdown of Antenna Utilization for Base Scenario with 3 min Reconfiguration Time . . . 16

6 Antenna Schedule for Base Scenario with 3 min Reconfiguration Time . . . 17

7 Antenna Schedule for X-band Conversion Scenario with 3 min Reconfiguration Time . . . 17

8 Antenna Schedule for Phased Array at Fairbanks with 3 min Reconfiguration Time . . . 19

9 Antenna Schedule for Phased Array at McMurdo with 3 min Reconfiguration Time . . . 21

10 Antenna Schedule for Phased Array at Svalbard with 10 min Reconfiguration Time . . . 23

11 100 Satellite Constellation Configuration . . . 25

12 Antenna Schedule Using KSAT with 1 min Reconfiguration Time . . . 28

13 Antenna Schedule for Svalbard Phased Array with 10 min Reconfiguration Time . . . 29

14 Antenna Schedule for McMurdo Phased Array with 10 min Reconfiguration Time . . . 29

15 Antenna Schedule for Fairbanks Phased Array and X-band Conversion with 10 min Reconfiguration Time . . . 31

16 Antenna Schedule with Fairbanks Phased Array and Extra Antennas at Svalbard (10 min Reconfiguration Time) . . . 32

17 Antenna Schedule with Fairbanks Phased Array and Extra Antennas at McMurdo (10 min Reconfiguration Time) . . . 33

18 Antenna Schedule for Dual Phased Array . . . 35

List of Tables

1 28 Satellite Constellation Configuration . . . 12

2 Antenna Reconfiguration Time vs. Number of KSAT Contacts . . . 14

3 Antenna Utilization with X-band Conversion . . . 16
4 Antenna Utilization with Fairbanks Phased Array . . . 18
5 Antenna Utilization with McMurdo Phased Array . . . 20
6 Antenna Utilization with Svalbard Phased Array . . . 22
7 Antenna Utilization with KSAT . . . 27
8 Antenna Utilization with Svalbard/McMurdo Phased Array . . . 28
9 Antenna Utilization with Fairbanks Phased Array and X-band Conversion . . . 30
10 Antenna Utilization with Fairbanks Phased Array and Extra Capacity at Svalbard . . . 31
11 Antenna Utilization with Fairbanks Phased Array and Extra Capacity at McMurdo . . . 31
12 Antenna Utilization with Fairbanks Phased Array and Extra Capacity at Svalbard . . . 33
13 Antenna Utilization for Dual Phased Array . . . 35

Acronyms

AoS Acquisition of Signal

CDAS Command and Data Acquisition Station

FCDAS Fairbanks Command and Data Acquisition Station

GEO Geosynchronous Orbit

GNN Graph Neural Network

JPSS Joint Polar Satellite System

LEO low-Earth Orbit

LTAN Local Time of the Ascending Node

NEON Near Earth Orbit Network

NESDIS National Environmental Satellite, Data, and Information Service

SMD Science Mission Data

SSO Sun-Synchronous Orbit

TT&C Telemetry, Tracking, and Command

1 Introduction

NOAA has been operating environmental satellites in low-Earth Orbit (LEO) for over 50 years. The next generation of these constellations, the Near Earth Orbit Network (NEON), is being planned with a shift in design paradigm that emphasizes quick deployment and scalability, aiming to produce a larger, more resilient constellation of smaller satellites. This new paradigm requires simultaneous investment in scaling ground systems to support ever-increasing demand as constellations scale, necessitating intelligent planning to maximize network performance and reduce the burden on operators.

Traditional methods of planning ground contacts rely on heuristic approaches or manual scheduling, solutions that do not scale as the decision space of the problem increases with more satellites and ground stations. The proliferation of artificial intelligence in computing provides new alternatives to solving this problem. CNTIENT.Optimize leverages this AI-enabled approach to mission planning in order to efficiently scale alongside growing constellations while producing optimal schedules that meet all mission requirements at a fraction of the cost and manual effort currently required.

In this demonstration CNTIENT.Optimize was used to conduct a study of future NOAA ground system capacity needs in fulfillment of National Environmental Satellite, Data, and Information Service (NESDIS) requirements by generating schedules for a variety of simulated operational scenarios. The goal of this exercise was to demonstrate product scalability and performance in supporting large scale multi-mission management.

2 Methodology

The objective of this analysis was to demonstrate CNTIENT.Optimize's ground contact schedule optimization by simulating increasingly complex operational scenarios. Scenarios were designed to provide a trade space analysis for potential future ground system design in support of two satellite constellation configurations:

28 Satellites - This constellation was designed to reflect the potential future real-time operational environment in support of the NEON mission. The constellation is heterogeneous comprising multiple missions including the 14 NEON satellites in addition to supporting 14 additional satellites from partner missions.

100 Satellites - This constellation was designed to stress test the ground system by fully saturating NOAA Command and Data Acquisition Stations (CDASs). The constellation is comprised of a homogeneous fleet modelled after the NEON spacecraft.

Both constellations are based on similar orbital geometry intended to reflect typical mission design for environmental satellites. All spacecraft are arranged in Sun-Synchronous Orbits (SSOs) across different planes, each having a repeat cycle of approximately 16 days. All scenarios were evaluated on a 16 day scheduling period to reflect any day-to-day differences within this cycle.

A variety of operational scenarios were simulated for each constellation to provide insight into different trades that could be made in ground system configuration. The analysis focused on four main questions in evaluating ground capacity needs:

1. How well do the current NOAA-owned antennas support the constellation?
2. How could commercial antenna vendors, such as KSAT, be used to augment the current NOAA system to provide additional capacity?
3. As an alternative to buying commercial contacts, how could the current NOAA ground system be updated to provide additional capacity?
4. How much capacity could the newer technology of phased arrays buy for the system? Phased arrays represent a huge potential improvement in performance over traditional parabolic dishes by allowing multiple simultaneous contacts with none of the overhead due to antenna reconfiguration.

CNTIENT.Optimize was configured to reflect each scenario, each of which represents a different ground network configuration, and used to generate viable mission schedules that reflect all operational constraints for each mission. These schedules were then post-processed to generate metrics on data latency and antenna utilization as a means of quantifying system performance. Further analysis then focused on evaluating trade-offs between each scenario to provide a picture of which configurations performed best.

2.1 Definitions

Antenna Reconfiguration Time - This term refers to the length of time required for a parabolic dish antenna to reconfigure ahead of taking the next satellite contact. This reconfiguration involves actions such as re-pointing of the dish and changing frequency band or protocol. This reconfiguration time is enforced in the optimized schedule as a post-contact blackout for a particular antenna after accepting a pass, during which no other satellites can schedule time on that antenna.

Antenna Percent Utilization - This metric represents the fraction of time an antenna is busy in the schedule. Usually expressed as an average across all antennas in the system, it is intended to provide a qualitative picture of how much stress a particular scenario imposes on the system as compared to other configurations. Antenna percent utilization is calculated following Equation 1:

$$utilization = \frac{total\ contact\ time + total\ reconfiguration\ time}{total\ access\ time} \tag{1}$$

Where:

- $utilization$ = Antenna percent utilization
- $total\ contact\ time$ = Total time spent in communication with any spacecraft
- $total\ reconfiguration\ time$ = Total time spent re-configuring the antenna ahead of the next pass
- $total\ access\ time$ = Total time that any spacecraft is in view of the station

2.2 Assumptions

Assumption 1. The transmission of Science Mission Data (SMD) is the main bottleneck in meeting latency. Other overhead from actions such as Telemetry, Tracking, and Command (TT&C) and Acquisition of Signal (AoS) are much smaller in magnitude and are accounted for in the conservatism in our analysis. CNTIENT.Optimize also supports direct modelling of this kind of overhead with each link if it represents a significant constraint on operations, and can be configured on a case-by-case basis.

Assumption 2. All satellites require the use of X-band for transmission of SMD. This is the anticipated frequency band that the NEON constellation will use, and any partner missions supported will be assumed to use the same. Although CNTIENT.Optimize can support a mixed use of frequency bands for multi-mission management, requiring all satellites to compete for the same resources will increase the likelihood of scheduling conflicts and add conservatism to any results.

Assumption 3. The only traffic on the network comes from the constellation being supported. Any traffic from other sources on the network being planned would have to be explicitly configured in the product. External providers can be integrated into the product to provide antenna unavailabilities to be accounted for in scheduling. For example, KSAT integration is already supported in the product.

Assumption 4. The Geosynchronous Orbit (GEO) and deep space satellites have deterministic schedules on their own antennas that do not conflict with scheduling for the LEO satellites. Due to their deterministic nature there is not a need for schedule optimization as in the case for LEO satellites, and are therefore left out of the scope of this analysis.

2.3 Optimization Method

2.3.1 Graph Representation - Latency as a Constraint

For the purposes of the trade study a hard performance boundary was sought to test the limits of the satellite and ground antenna configuration. By introducing a graph representation of the contact planning problem we were able to encode latency as a constraint which must be satisfied in any feasible solution, before considering improvements in an objective function as a secondary consideration.

Graph construction begins by generating and enumerating ground contact opportunities, which are taken to represent nodes in the graph. Edges are then drawn between any two nodes representing a valid transition in the mission schedule. Edges are the primitive decision variables in this formulation, with a valid schedule being generated by selecting a route through the graph consisting of a sequence of edges. As a quick example, consider 2 contact opportunities A, B which occur in time order $t(A) < t(B)$. There must also be virtual options at the start and end of the route, V_s for the start and V_e for the end. In the general case there would be then be 6 possible edges $(V_s, A), (V_s, B), (A, B), (A, V_e), (B, V_e), (V_s, V_e)$. On each edge there would be a binary decision variable x_i. One can immediately see that this formulation introduces many variables, which can lead to large graph sizes in basic implementations. However, this approach provides very powerful modeling capabilities that guarantee latency requirements and support finding a true optimum solution that make it ideal for this analysis.

Objective function. The objective function of the graph representation is

$$\Phi = \sum_{i=1}^{N} x_i \cdot r_i \tag{2}$$

where N is the number of edges in the graph, $x_i \in \{0, 1\}$ is the binary decision variable representing selecting the ith edge, and r_i is the corresponding reward. To facilitate application to ground contact planning let us enumerate all contact options in a planning period as g_J, $J = 1, ..., N_{\text{contacts}}$, i.e. we use capital indexes for the contact options to distinguish the subscripts from the edge variables. For notation specificity suppose the ith edge in the graph representation is a route segment starting at contact g_J and proceeding to contact g_K. Note that a specific contact g_J corresponds a unique combination of satellite, ground antenna, and conjunction index, so that we know properties such as antenna priority at a given contact option. The reward factor for that edge is decomposed as

$$r_i = w_i \cdot d_i \cdot \hat{b}(\Delta t_i) \tag{3}$$

where

$$\Delta t_i = t_s(g_K) - t_e(g_J), \tag{4}$$

is the time difference between an edge's contacts: $t_s(g_K)$ is the start time of the destination contact g_K and $t_e(g_J)$ is the end time of the start contact g_J, w_i is the weight based on antenna priority corresponding to the ith edge, and d_i is the corresponding contact duration. Both of these values come from the properties of the destination contact g_K. The final factor $\hat{b}(\Delta t)$ is a function which favors ground pass intervals close to an optimal time spacing. For instance to optimize plans to reduce antenna utilization, we use an objective function which rewards contact spacing close to the max latency gap allowed, e.g.

$$\hat{b}(\Delta t(g_J, g_K); L_{\max}) = \frac{|t_s(g_K) - t_e(g_J)|}{L_{\max}} \tag{5}$$

Graph constraints. The graph constraints may be thought of as flow conservation equations for possible routes through the graph. At each interior node, the sum of edge values into the node must equal the outflow, and this value must be less than or equal to 1:

$$\sum_{i=1}^{N} x_i(\delta_{\mathrm{s}})_J^i = \sum_{i=1}^{N} x_i(\delta_{\mathrm{d}})_J^i \quad \forall \quad J \in 1, ..., N_{\mathrm{contacts}} \tag{6}$$

where $(\delta_{\mathrm{s}})_J^i = 1$ if contact g_J is the source of the ith edge and is 0 otherwise, and $(\delta_{\mathrm{d}})_J^i = 1$ if contact g_J is the destination of the ith edge and is 0 otherwise. The flow out of the virtual nodes must also be 1, e.g.

$$\sum_{i=1}^{N} x_i(\delta_{\mathrm{s}})_{V_s}^i = 1 \tag{7}$$

and

$$\sum_{i=1}^{N} x_i(\delta_{\mathrm{d}})_{V_e}^i = 1 \tag{8}$$

where V_s, V_e are the indexes of the virtual nodes.

Parabolic dish constraints. Due to the directional nature and tuning requirements of a parabolic dish antenna, a ground antenna can maintain an effective communication link with only one satellite at a time. In the problem representation this physical assumption is formulated as a constraint by first finding the contact options which overlap (suppose we identify specific instances x_{i^*} and x_{j^*} from 2 different satellites) and then applying

$$x_{i^*} + x_{j^*} \leq 1 \tag{9}$$

This process is done for all such contact conflicts.

Latency. Finally, for this formulation to guarantee data latency is met, we pre-filter the possible edges of the graph by considering forming an edge between g_J and g_K only if $t_s(g_K) - t_e(g_J) < L_{\max}$. In this way, a feasible solution to the problem automatically must satisfy the latency requirement. Otherwise the flow constraints Equation 6-Equation 8 would be violated and an infeasible solution would be returned from the solve step.

2.3.2 Simpler Representations - Latency as Objective

To prevent infeasible solutions during real-time operations, the latency constraint would be relaxed to a latency objective. This formulation would seek to minimize latency during optimization, with conflicts being sorted out by contact priority. While this solution would not guarantee latency, as in the graph representation, it would still eventually converge on a schedule that satisfies latency requirements, if such a schedule exists. This formulation guarantees a schedule, and is shown to converge much faster than the graph formulation.

2.4 Base Scenario Configuration

Ground station configuration. In order to provide a standard baseline to compare against for each scenario, a base scenario was defined to best reflect the current configuration of the NOAA ground system, as available to the Joint Polar Satellite System (JPSS). Since not all the antennas available in the current network support X-band, which is required by Assumption 2, it further limits the number of available antennas:

Fairbanks - 1 Ka-band + 3 S/X-band antennas. Allows for 3 simultaneous contacts.

Svalbard - 1 Ka-band + 1 S/X-band antennas. Allows for 1 contact at a time.

McMurdo - 3 S/X/Ka-band antennas. Allows for 3 simultaneous contacts.

Svalbard (KSAT) - 1 S/X/Ka-band antenna. Allows for 1 contact at a time. Lower priority than NOAA-owned stations (unpreferred).

Troll (KSAT) - 1 S/X/Ka-band antenna. Allows for 1 contact at a time. Lower priority than NOAA-owned stations (unpreferred).

Operational constraints. Throughout this analysis there were three main operational constraints that made a significant impact on system performance:

Latency Maximum data latency was imposed as a constraint to provide a definitive feasible/infeasible solution for each scenario. Maximum latency was set at 80 min for primary NEON sensors and 140 min for backup NEON sensors and partner missions. As mentioned in subsubsection 2.3.1, for this analysis latency was enforced as a schedule constraint to provide a definitive answer on whether a particular scenario is feasible or infeasible. While this is ideally suited for a trade space analysis, this latency constraint would be relaxed to a latency object in real-time operations to avoid returning an infeasible schedule.

Antenna Reconfiguration Antenna reconfiguration time (section 2.1) only applies to parabolic dish antennas. It represents the total amount of overhead required at an antenna level to support a satellite contact. This constraint was varied in each scenario as an independent variable to feel out the amount of capacity provided by a particular antenna configuration. The longer it takes to reconfigure an antenna, the greater the likelihood of a scheduling conflict. Therefore, maximum allowable antenna reconfiguration time qualitatively represents the amount of resiliency a particular antenna configuration has to satellite scheduling conflicts.

Contact Duration The lengths of all contact were made 33% longer than the ones seen for JPSS to provide additional conservatism in the schedules and account for future unknowns. The length of a contact varied depending on the time between contacts.

2.5 Approach to Analysis

All scenarios simulated in this analysis provide some modification of the base scenario configuration laid out in subsection 2.4. The following steps were performed in execution of each scenario:

1. **Modify configuration** to reflect different trades in ground network design.
2. **Vary constraints** to find the point where the schedule becomes infeasible.
3. **Optimize** the contact schedule over all available opportunities. See Figure 1 for reference on the optimization process.
4. **Validate** the output schedule to verify all operational requirements are met.
5. **Visualize** the performance of the system with metrics and figures.

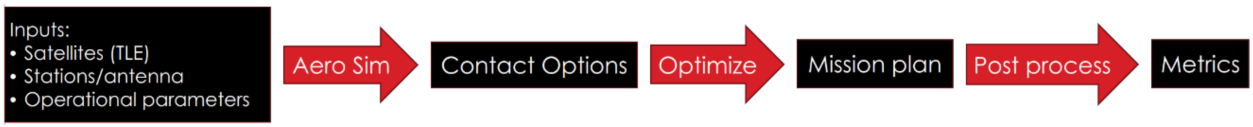


Figure 1: CNTIENT.Optimize schedule optimization process

Figure 1 shows the execution control flow for optimizing a ground contact schedule. Inputs are generated when the configuration is set and a plan is created, which are then passed to the core scheduling endpoint for optimization. The optimization process begins by finding all potential contact opportunities for each satellite, which are then optimized over to put together a viable mission schedule. Finally, these schedules are post-processed to produce the metrics and figures shown later throughout the analysis.

3 28 Satellite Results

3.1 Constellation Configuration

Table 1 specifies the missions modeled in the 28-satellite constellation with their corresponding orbits. All spacecraft are in SSOs with periods between 96 min to 102 min, specified by their Local Time of the Ascending Node (LTAN):

Satellite/Mission	**Description**
NEON	• 2 replacing JPSS (1330 LTAN) • 6 Series-1 1330 LTAN • 6 Series-2 1730 LTAN
MetOp/MetOp-SG	• 2 MetOp 2130 LTAN • 2 MetOp-SG 2130 LTAN
AWS/Sterna	• AWS-PFM 2230 LTAN • 2 Sterna 1530 LTAN • 2 Sterna 1930 LTAN • 2 Sterna 2330 LTAN
USSF	• 1 USSF 1130 LTAN • 1 USSF 1530 LTAN • 1 USSF 1930 LTAN

Table 1: 28 Satellite Constellation Configuration

Spacecraft are equally spaced within each plane, with planes randomly offset from one another. The resulting constellation is visualized in Figure 2:

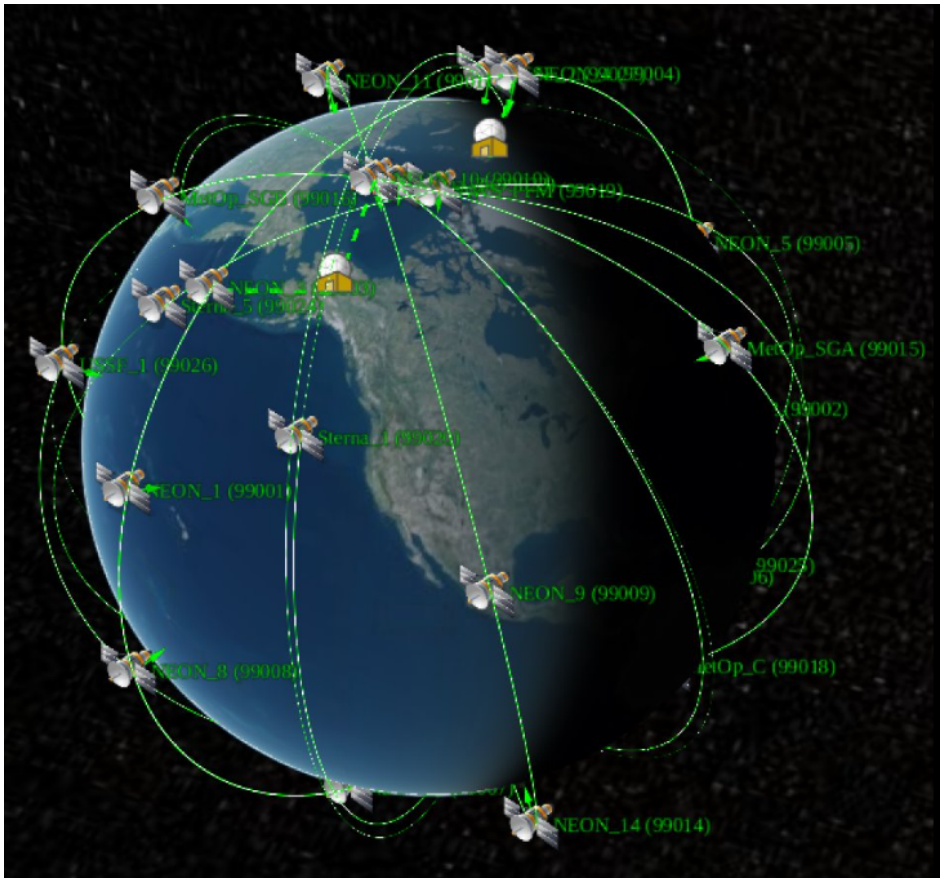


Figure 2: 28 Satellite Constellation Configuration

3.2 Scenarios

In order to answer the main questions (section 2) of our analysis, the following scenarios were run:

Base Scenario - How much capacity does the current NOAA ground system defined in subsection 2.4 provide?

KSAT Scenario - How could commercial providers, such as KSAT, be used to augment the current system to provide extra capacity? Specifically, do the current KSAT antennas being used by JPSS provide enough capacity?

X-band Conversion Scenario - How could the current ground system be updated to better support this constellation? Specifically, would converting all existing NOAA-owned antennas to support X-band be a viable alternative to buying KSAT passes?

Phased Array Scenarios - How much capacity could a phased array provide? Where should it ideally be located?

Each scenario was run varying constraints to explore the boundaries of the problem. Scenarios were evaluated on how well they supported meeting latency objectives, along with how resilient they were against scheduling conflicts.

3.2.1 Base + KSAT Results

Description. This scenario was designed to evaluate how much capacity the current NOAA-owned ground system provides. Additionally, it is intended to determine how much extra capacity the KSAT antennas available to JPSS can provide. The ground network is configured exactly as defined in subsection 2.4. The antenna reconfiguration time is slowly increased to find the point where the KSAT antennas are required to deconflict any satellite passes. For this scenario the KSAT antennas are given a negative priority to penalize the objective function for using KSAT passes. This incentivizes the optimizer to find a schedule that does not use any KSAT contacts unless it is absolutely necessary.

Results. The base scenario, consisting of solely NOAA-owned antennas, is sufficient to support the 28 satellites for each of their respective latency requirements provided the antenna reconfiguration time is kept less than 3 min. Once antenna reconfiguration time, and therefore the overhead associated with each contact, is increased past 3 min, KSAT time must be used to meet all latency requirements. In the worst-case of a 10-minute reconfiguration time, the schedule required approximately 26 to 27 KSAT passes per day to meet all mission requirements.

KSAT Allowed	Antenna Reconfiguration Time	Latency Met	Number of KSAT Contacts
No	3 min	✓	-
No	4 min	X	-
Yes	10 min	✓	426 (∼26 / day)

Table 2: Antenna Reconfiguration Time vs. Number of KSAT Contacts

Figure 3 shows the schedules for each antenna, and corresponds to the scenario represented by the first row in Table 2. Contacts are color-coded to represent sensor type, and the first day of the 16-day schedule is displayed. The top two rows in the figure represent the schdules for the KSAT antennas, which do not have any contacts scheduled because a viable schedule can be generated without them.

After antenna reconfiguration time is increased beyond 3 min, a feasible schedule can no longer be generated without relying on KSAT passes for deconfliction. Once KSAT passes are allowed to fill in the gaps in coverage, antenna reconfiguration time is able to be increased to 10 min through buying KSAT contacts at a rate of approximately 26 passes per day. The schedule for the 10-minute antenna reconfiguration scenario with KSAT antennas is shown in Figure 4. The increase in antenna reconfiguration time can be seen in the larger gaps between each contact in the schedule compared to the schedule shown in Figure 3. The additional overhead caused by antenna reconfiguration causes the new KSAT contacts shown in the top two rows at the times where scheduling conflicts arise. Relying on KSAT is a viable strategy for buying extra resiliency against scheduling conflicts in the ground system, and can meet all requirements for the 28-satellite constellation.

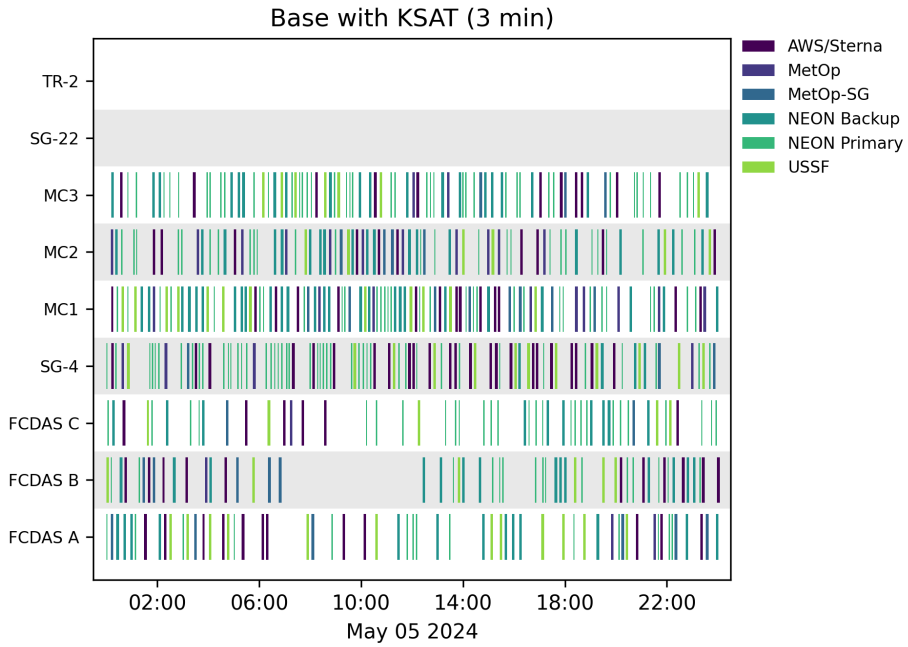


Figure 3: Antenna Schedule for Base + KSAT Scenario with 3 min Reconfiguration Time

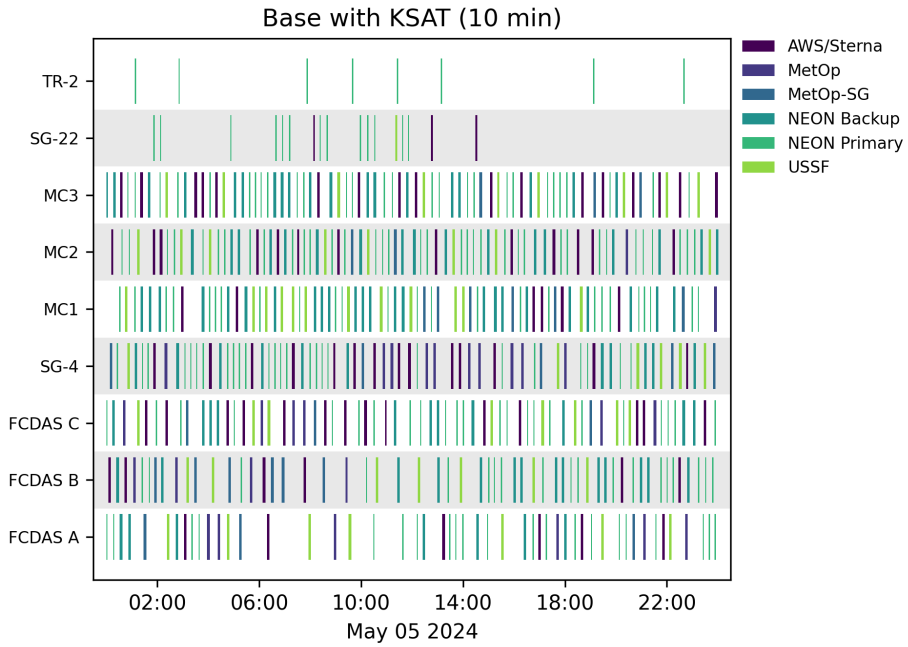


Figure 4: Antenna Schedule for Base + KSAT Scenario with 10 min Reconfiguration Time

3.2.2 X-band Conversion Results

Description. This scenario was designed to evaluate if an alternative to buying commercial passes from KSAT exists in the form of upgrading the NOAA-owned antennas to support X-band across the board. Doing this would buy two additional antennas not in the base scenario: one at Fairbanks (FCDAS 4m) and one at Svalbard (SG-60). As in the previous scenario, antenna reconfiguration time was incrementally increased to determine how much additional capacity these two antennas provide and measure system resiliency against scheduling conflicts.

Results. Upgrading the NOAA-owned antenna network to ensure X-band support across the board for the antennas available to JPSS ensures there is enough capacity to meet all mission requirements for the 28-satllite constellation. Furthermore, it allows for up to 10 min antenna reconfiguration time, therefore buying similar resiliency as the two KSAT antennas without having to purchase time from commercial vendors.

With Conversion	Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
No	3 min	✓	41% (16%)
No	4 min	X	-
Yes	3 min	✓	32% (12%)
Yes	10 min	✓	60% (41%)

Table 3: Antenna Utilization with X-band Conversion

Table 3 shows how upgrading SG-60 and FCDAS 4m to support X-band buys additional resiliency against scheduling conflicts by allowing longer for antenna reconfiguration. Figure 7 shows these two antennas now available to the satellites. Comparing this schedule to the same scenario without the X-band conversion shown in Figure 6 shows how these two antenna buy down antenna utilization by allowing the contacts to spread out more across all antennas. This is shown quantitatively in Table 3 by comparing the utilization between rows 1 and 3, which shows a 22% reduction in total antenna utilization.

Figure 5 provides a better way to visualize the antenna utilization metrics provided in Table 3. The numbers for total utilization and the portion due to antenna reconfiguration correlate to the first row in the table. Comparing rows 3 and 4 of the table shows that as antenna reconfiguration time increases, and all else stays the same, both the total utilization and the utilization due to antenna reconfiguration increase while the portion due to transmission of SMD stays the same. Furthermore, the utilization due to reconfiguration increases by about the same factor as the antenna reconfiguration time.

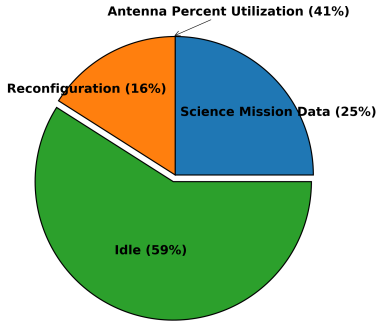


Figure 5: Breakdown of Antenna Utilization for Base Scenario with 3 min Reconfiguration Time

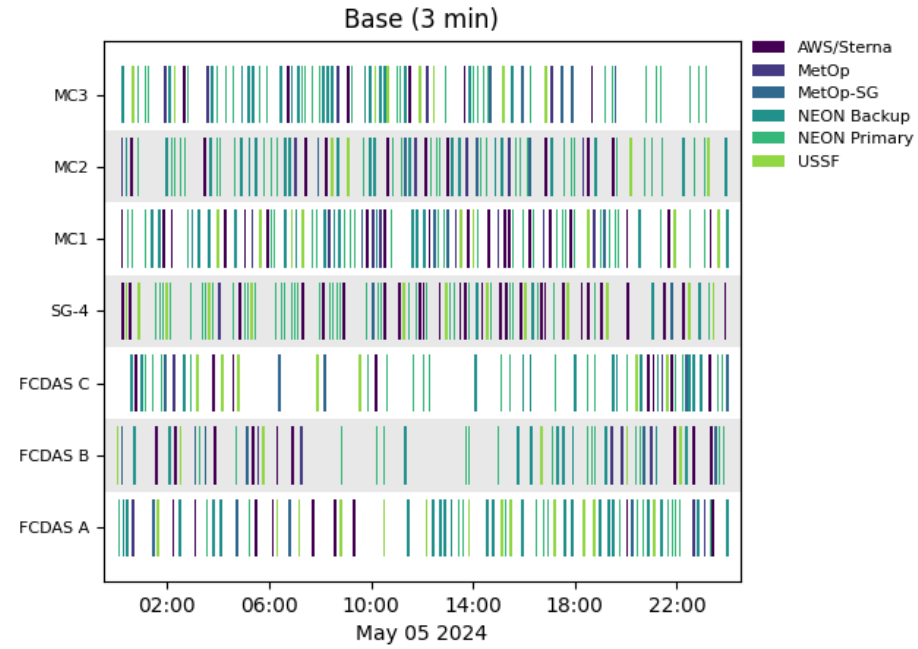


Figure 6: Antenna Schedule for Base Scenario with 3 min Reconfiguration Time

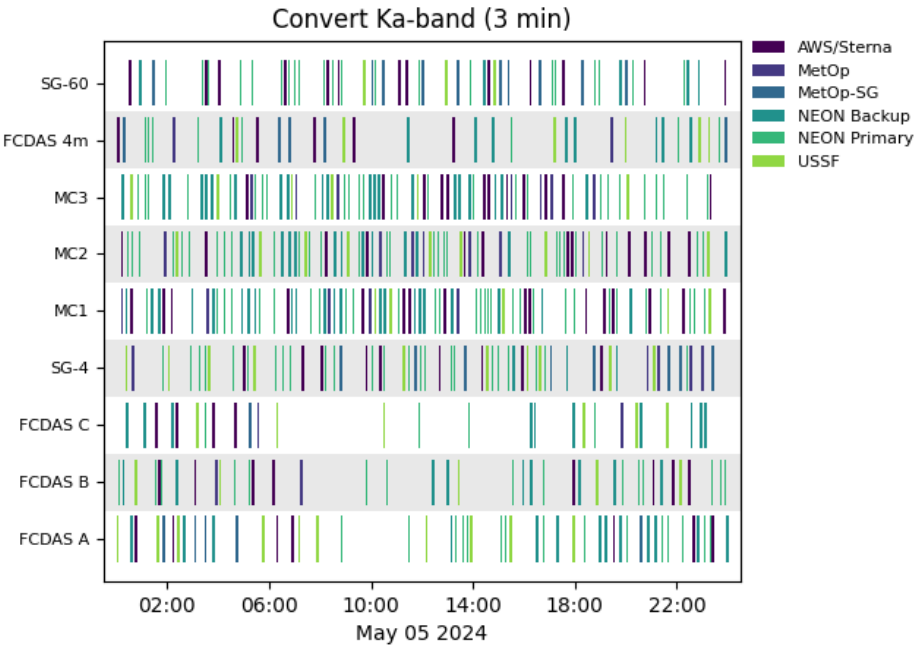


Figure 7: Antenna Schedule for X-band Conversion Scenario with 3 min Reconfiguration Time

3.2.3 Phased Array Results - Fairbanks

Description. This scenario is the first to investigate the effectiveness of phased arrays by placing one at Fairbanks. The phased array was modeled as 8 additional antennas, representing the 8 simultaneous contacts it would be capable of supporting. As in previous scenarios, antenna reconfiguration time is incrementally increased to determine the amount of additional capacity and resiliency towards scheduling conflicts this configuration provides. Antenna reconfiguration is a constraint that applies to parabolic dish antennas only and is not applied to the 8 new antennas representing the phased array.

Results. Adding a phased array at Fairbanks had a dramatic impact on reducing antenna utilization for the ground system, reducing overall utilization by 56%. However, the phased array did not have a significant impact on adding additional resiliency against scheduling conflicts, being unable to support a maximum antenna reconfiguration time beyond the 3 min shown in the base scenario.

The reason the phased array does not increase allowable reconfiguration time despite providing ample capacity is because the capacity is not being added in the ideal location. The Fairbanks Command and Data Acquisition Station is not located at a high enough latitude to provide 100% coverage to the satellites, unlike the stations at Svalbard and McMurdo. This is due to the geometry between the spacecrafts' SSOs and the station. Figure 11 can be used to get a picture of how these orbits precess along with Earth's rotation over the course of a day. As can be seen in the figure, on some revolutions the spacecraft pass too far north of the station to fall into access. These blind revolutions require any spacecraft that requires two contacts per revolution to rely on the single antenna at Svalbard (SG-4) to fulfill their contact in the Northern Hemisphere. This causes a bottleneck at this station, and as overhead due to antenna reconfiguration increases it causes unresolvable conflicts in the schedule. Because the phased array is located at Fairbanks, it is not able to address this bottleneck at Svalbard as it will also be occasionally subject to these blind revolutions.

With Fairbanks Phased Array	Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
No	3 min	✓	41% (16%)
No	4 min	X	-
Yes	3 min	✓	18% (6%)
Yes	4 min	X	-

Table 4: Antenna Utilization with Fairbanks Phased Array

Table 4 shows the dramatic reduction in mean antenna utilization, both for overall utilization and utilization from reconfiguration. The 8 simultaneous contacts being supported by this array are shown at the top of Figure 8, which qualitatively shows the huge increase in capacity now available at the station. Since the phased array does not require reconfiguration between passes it also dramatically reduces the amount of overhead required to support an equal amount of passes as the base scenario.

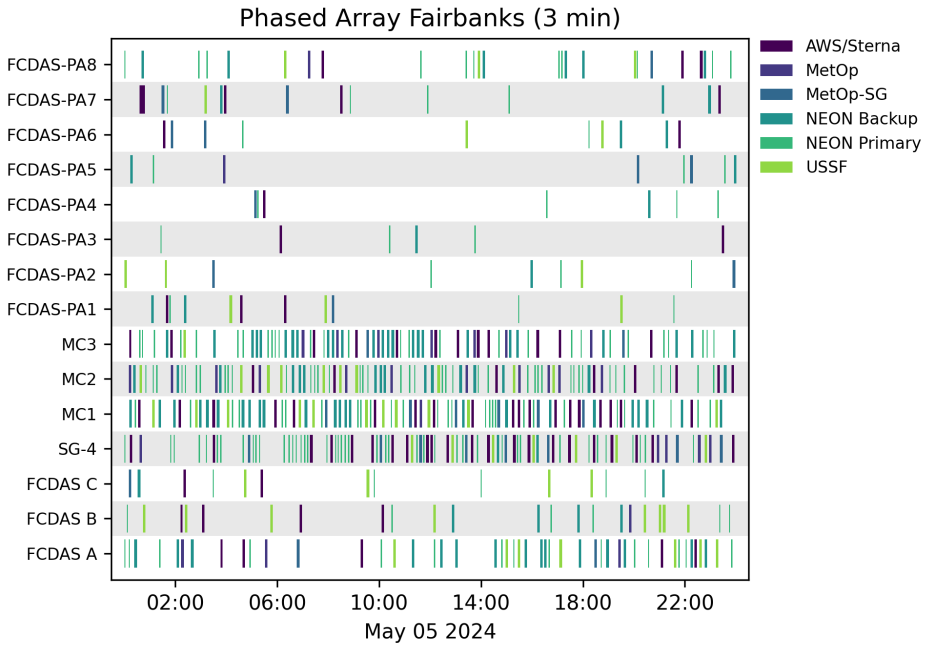


Figure 8: Antenna Schedule for Phased Array at Fairbanks with 3 min Reconfiguration Time

3.2.4 Phased Array Results - McMurdo

Description. This scenario continues the investigation into the effectiveness of phased arrays by placing one at McMurdo. The phased array was modeled as 8 additional antennas, representing the 8 simultaneous contacts it would be capable of supporting. As in previous scenarios, antenna reconfiguration time is incrementally increased to determine the amount of additional capacity and resiliency towards scheduling conflicts this configuration provides. Antenna reconfiguration is a constraint that applies to parabolic dish antennas only, and is not applied to the 8 new antennas representing the phased array.

Results. Similar to the scenario with the Fairbanks phased array, the phased array at McMurdo dramatically reduced antenna utilization. Additionally, even though McMurdo station provides full coverage for all revolutions, it is still not adding capacity in the right spot to address the bottleneck at the Svalbard SG-4 antenna. Therefore, even though it significantly builds out capacity in the network, it does not provide any additional resiliency to scheduling conflicts over the base scenario configuration.

With McMurdo Phased Array	Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
No	3 min	✓	41% (16%)
No	4 min	X	-
Yes	3 min	✓	17% (5%)
Yes	4 min	X	-

Table 5: Antenna Utilization with McMurdo Phased Array

Table 5 shows the dramatic reduction in mean antenna utilization, both for overall utilization and utilization from reconfiguration. The 8 simultaneous contacts being supported by this array are shown at the top of Figure 9, which qualitatively shows the huge increase in capacity now available at the station. Since the phased array does not require reconfiguration between passes it also dramatically reduces the amount of overhead required to support an equal amount of passes as the base scenario.

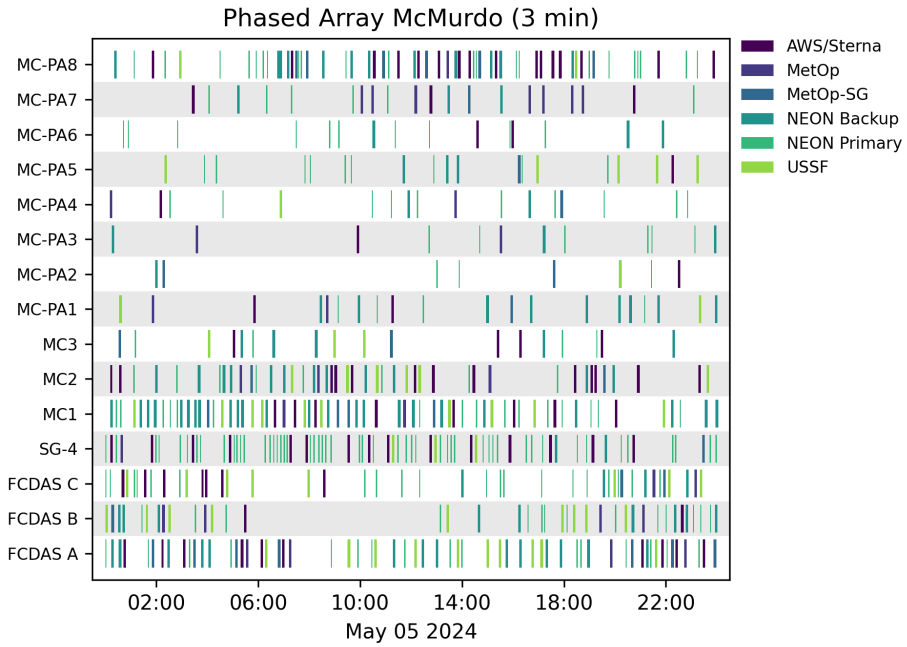


Figure 9: Antenna Schedule for Phased Array at McMurdo with 3 min Reconfiguration Time

3.2.5 Phased Array Results - Svalbard

Description. This scenario is the final one in the investigation into the effectiveness of phased arrays by placing one at Svalbard. The phased array was modeled as 8 additional antennas, representing the 8 simultaneous contacts it would be capable of supporting. As in previous scenarios, antenna reconfiguration time is incrementally increased to determine the amount of additional capacity and resiliency towards scheduling conflicts this configuration provides. Antenna reconfiguration is a constraint that applies to parabolic dish antennas only, and is not applied to the 8 new antennas representing the phased array.

Results. Adding a phased array at Svalbard finally addresses the bottleneck in capacity at that station, and therefore provides the necessary means to prevent scheduling conflicts for the 28-satellite constellation. Doing so allows for the full 10 min antenna reconfiguration time, and actually serves to move the bottleneck to the South at McMurdo station, which has plenty of capacity with its 3 antennas to meet all mission requirements. As in the previous two scenarios, the phased array served to significantly reduce antenna utilization across the network by adding much more capacity to the system, as well as reduce overhead associated with reconfiguration.

With Svalbard Phased Array	Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
No	3 min	✓	41% (16%)
No	4 min	X	-
Yes	3 min	✓	17% (6%)
Yes	10 min	✓	27% (16%)

Table 6: Antenna Utilization with Svalbard Phased Array

Table 6 shows how the maximum allowable antenna reconfiguration time is increased with this configuration, as well as the dramatic reduction in mean antenna utilization seen in the previous two scenarios. The 8 simultaneous contacts being supported by this array are shown at the top of Figure 10, which qualitatively shows the huge increase in capacity now available at the station. Since the phased array does not require reconfiguration between passes it also dramatically reduces the amount of overhead required to support an equal amount of passes as the base scenario. The application of the antenna reconfiguration constraint is evident from the figure by observing the difference in minimum spacing between contacts for the phased array antennas and the parabolic dish antennas. Furthermore, the larger supported antenna reconfiguration time is shown qualitatively in a similar manner when comparing against the other antenna schedules in Figure 8 and Figure 9.

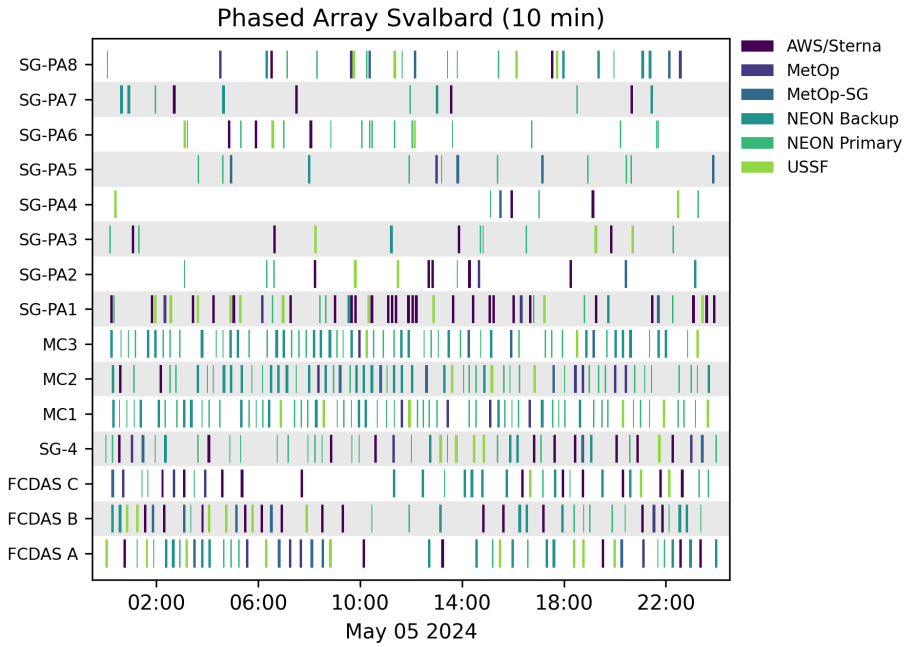


Figure 10: Antenna Schedule for Phased Array at Svalbard with 10 min Reconfiguration Time

3.3 Summary

Ultimately, the current NOAA ground system as available to JPSS, outlined in subsection 2.4, provides sufficient capacity to support the 28-satellite constellation provided antenna reconfiguration time does not exceed 3 min. There are several viable methods to provide more resiliency against scheduling conflicts, allow for more overhead, and increase capacity:

Commercial Antenna Vendors. Relying on the two KSAT antennas available to JPSS to resolve the occasional conflict provides enough resiliency to meet all mission requirements without risk. The worst-case antenna reconfiguration time examined in this analysis was 10 min, which required up to 27 KSAT contacts per day.

NOAA Network Update. As an alternative to relying on KSAT to fill in gaps in coverage, ensuring all NOAA-owned antennas available to JPSS can support X-band would also buy enough capacity to fulfill mission requirements. This makes the FCDAS 4m and SG-60 antennas available to the spacecraft. This serves not only to deconflict contacts when needed, but also lowers overall antenna utilization on the network, thereby reducing stress on the system.

Phased Arrays. Phased arrays significantly increase capacity and lower stress on the network no matter where they are placed, and make a much larger individual difference than any parabolic dish antenna. However, to maximize utilization and provide resiliency against scheduling conflicts the phased array should ideally be placed at Svalbard to relieve the bottleneck there at SG-4 caused by the occasional blind revolution at Fairbanks.

4 100 Satellite Results

4.1 Constellation Configuration

The 100-satellite constellation is a homogeneous fleet composed of spacecraft modeled after the NEON satellites. The 100 satellites were divided equally into 20 planes of 5 satellites per plane each:

- Planes were equally spaced every 30 min between 0600 and 1800 LTAN
- No orbits were placed between 1100 and 1300 LTAN to avoid sun glint
- Planes alternate between ascending and descending

Spacecraft are equally spaced within each plane, with planes evenly offset from one another. The resulting constellation is visualized in Figure 11:

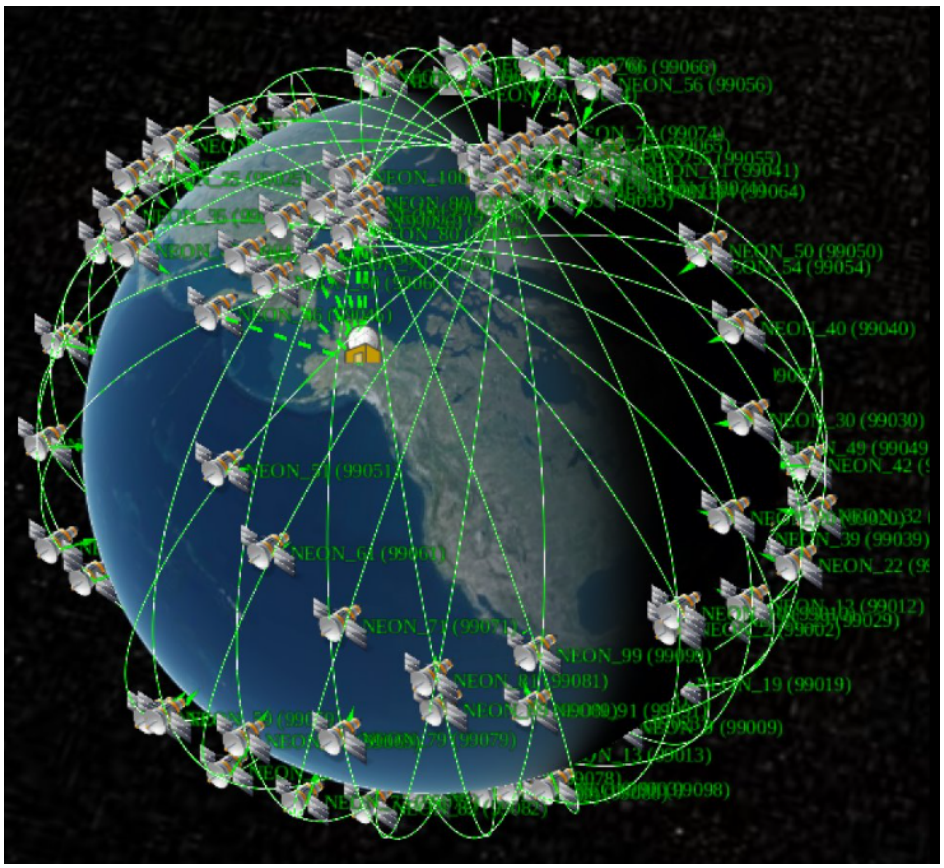


Figure 11: 100 Satellite Constellation Configuration

4.2 Scenarios - Supporting One Contact per Revolution

In order to identify potential solutions that would satisfy both the 100-satellite and the 28-satellite constellations, all previous scenarios were re-run with the new satellite configuration. Additionally, a deeper dive into how a phased array at Fairbanks could be supported was conducted, since explicit interest was shown in the feasibility of a phased array at that station. The following scenarios to augment the FCDAS phased array were run:

FCDAS Phased Array + X-band Conversion - Could a phased array at Fairbanks provide enough capacity if augmented by updating all NOAA-owned antennas to support X-band?

FCDAS Phased Array + Additional Antennas - If additional parabolic dish antennas were constructed to supplement the phased array at Fairbanks, how many are needed and where should they be built?

Each scenario was run varying constraints to explore the boundaries of the problem. Scenarios were evaluated on how well they supported meeting latency objectives, along with how resilient they were against scheduling conflicts.

4.2.1 Scenario Overlap

Description. Each of the following scenarios were re-run with this new satellite configuration to identify potential overlap in solutions that would serve to support the 100 satellites at a single contact per orbit:

Base Scenario - Reflects the current ground network available to JPSS.

KSAT Scenario - Supplements the current ground system with the two KSAT antennas available to JPSS.

X-band Conversion - Updates all JPSS antennas to support X-band, making two additional antennas available.

Phased Array - Investigates adding a phased array at each of the three NOAA stations used by JPSS.

What No Longer Works.

Base Scenario - The current NOAA-owned ground network is not capable of supporting the 100-satellite constellation.

X-band Conversion - Even with the additional capacity added by converting all NOAA antennas to support X-band, there is not enough capacity to provide full coverage for all 100 satellites. New antennas are needed to guarantee mission latency requirements.

Phased Array at Fairbanks - Due to the lower latitude position of the station causing occasional blind revolutions, the phased array at Fairbanks was unable to support the 100-satellite constellation.

KSAT Results. The two antennas supplied to JPSS by KSAT are just barely able to supply enough capacity to support all 100 satellites at one contact per revolution. However, they can only do so if antenna reconfiguration is kept at an absolute minimum. The maximum allowable antenna reconfiguration time for this scenario was 1 min, and once this was exceeded no feasible schedule could be found. As shown in Table 7 this comes with a very steep antenna utilization across the board.

Antenna Reconfiguration Time	Latency Met	Number of KSAT Contacts	Mean Antenna Utilization (% from Reconfiguration)
1 min	✓	4445 (∼277 / day)	69% (10%)
2 min	X	-	-

Table 7: Antenna Utilization with KSAT

The corresponding schedule is shown in Figure 12, where it is apparent how closely packed together all contacts are. Even though this schedule technically satisfies all requirements, due to its extremely high utilization and bare minimum allowable antenna reconfiguration time it is not a very feasible solution for meeting mission requirements. It does not provide much resiliency against conflicts between passes, and if a constellation was in a less than ideal configuration more conflicts could easily arise. Furthermore, this requires very high usage of third party antennas, which is not ideal. For this scenario to feasibly work, most likely additional KSAT antennas would need to be contracted to support these missions.

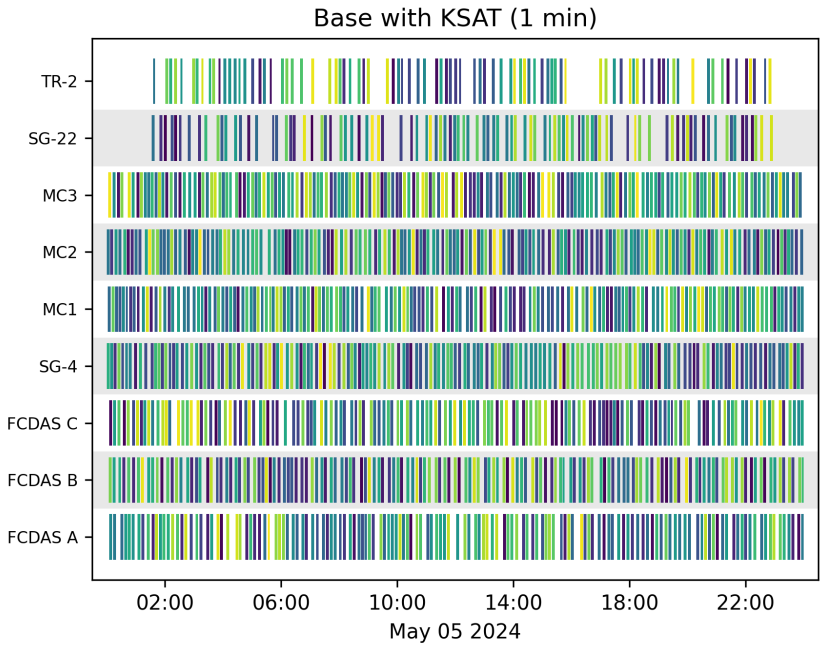


Figure 12: Antenna Schedule Using KSAT with 1 min Reconfiguration Time

Phased Array at Svalbard/McMurdo. Adding a phased at either Svalbard or McMurdo provides enough capacity and coverage to guarantee the 100 satellites their latency requirements. Since the satellites only require one contact per revolution, it does not matter whether this capacity is added in the North or the South, so long as it is added in a location that has 100% coverage of every orbit.

Location of Phased Array	Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
Svalbard	10 min	✓	59% (19%)
McMurdo	10 min	✓	60% (20%)

Table 8: Antenna Utilization with Svalbard/McMurdo Phased Array

The schedules for the Svalbard and McMurdo phased arrays are shown below in Figure 13 and Figure 14, respectively. It is apparent that, although the schedules are closely packed in, there is still plenty of capacity to support all 100 satellites in meeting latency requirements and sorting out any conflicts. The phased array at Svalbard in particular represents an overlap in solutions between the two constellation configurations, which makes it particularly well-suited to support future operations.

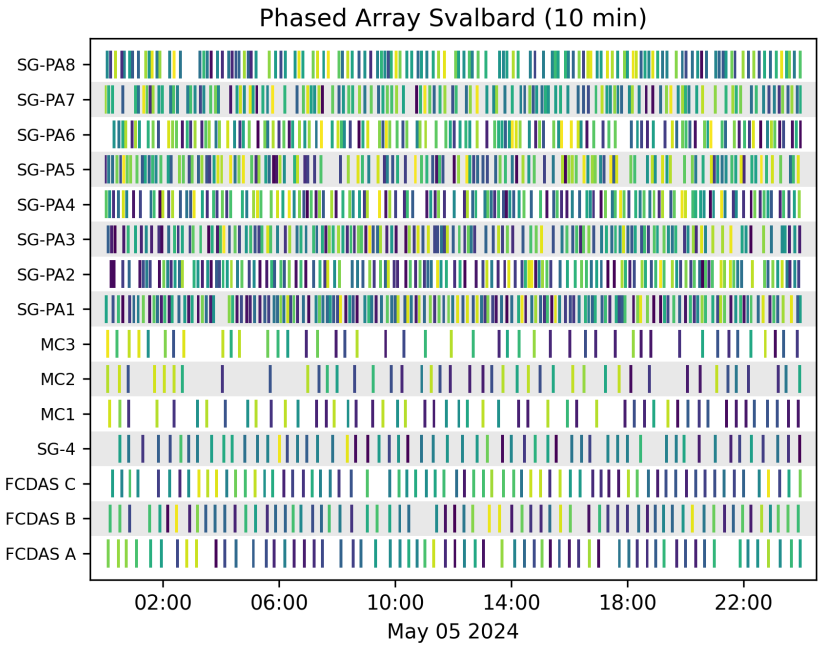


Figure 13: Antenna Schedule for Svalbard Phased Array with 10 min Reconfiguration Time

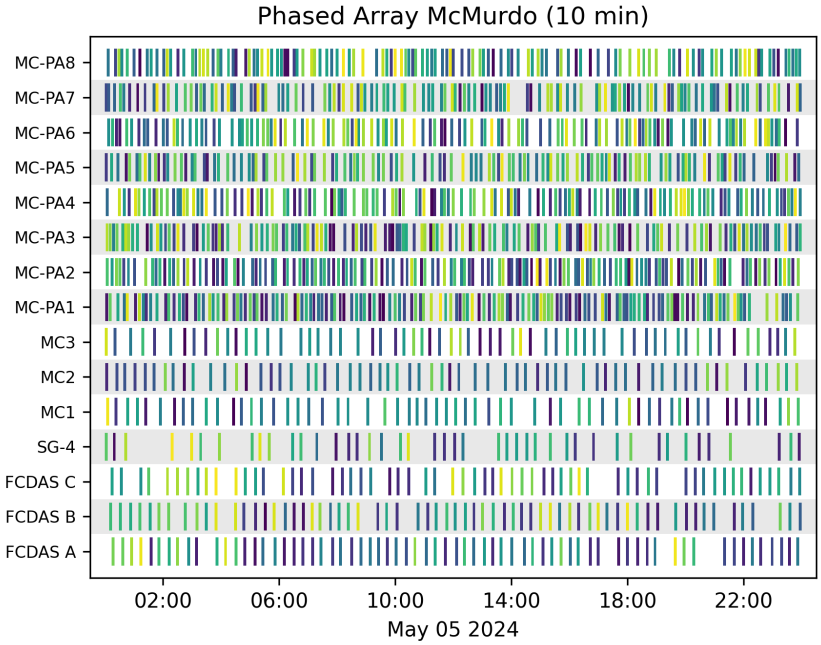


Figure 14: Antenna Schedule for McMurdo Phased Array with 10 min Reconfiguration Time

4.2.2 Supplementing a Phased Array at FCDAS

Description. The next two scenarios take a deeper dive into the feasibility of a phased array at Fairbanks since a phased array at this location was specifically mentioned in conversations with NOAA. While the Fairbanks phased array cannot provide 100% coverage of every revolution, it does cover the vast majority of revolutions. Since phased arrays at Svalbard and McMurdo were shown to provide ample coverage to support the constellation, the vast majority of passes are conceivably supported by the Fairbanks phased array. Therefore, the problem is finding the most efficient way of adding additional capacity to fill in for FCDAS on the occasional blind revolutions. To accomplish this, two main methods were investigated:

X-band Conversion - Ensure all NOAA-owned antennas can support X-band so they can all be made available.

New Antennas - Investigate adding new parabolic dish antennas in the network to fill in the gaps in coverage at Fairbanks.

X-band Conversion. As discussed in subsubsection 3.2.2, updating all NOAA-owned antennas makes two new antennas available (FCDAS 4m and SG-60). As shown in Table 9, doing so provides enough capacity to fill in for the outages at Fairbanks, provided the antenna reconfiguration doesn't exceed 3 min. It does not provide much resiliency against scheduling conflicts, but it does help relieve the bottleneck in capacity at SG-4, as shown in Figure 15.

Antenna Reconfiguration Time	Latency Met	Mean Antenna Utilization (% from Reconfiguration)
3 min	✓	47% (11%)
4 min	X	-

Table 9: Antenna Utilization with Fairbanks Phased Array and X-band Conversion

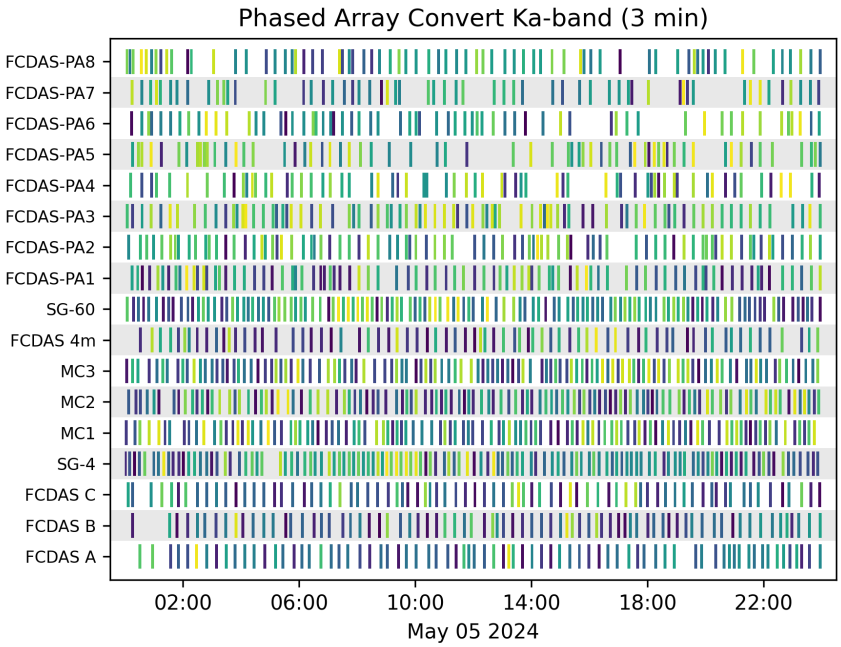


Figure 15: Antenna Schedule for Fairbanks Phased Array and X-band Conversion with 10 min Reconfiguration Time

New Antennas. The aim of this scenario is to identify the different methods of adding new antennas that would allow for the full 10 min antenna reconfiguration time. Only adding new capacity at Svalbard and McMurdo is explored since the purpose of the new antennas is to provide coverage during the blind revolutions at Fairbanks. New antennas are incrementally added to each station, and each time the maximum allowable antenna reconfiguration time is determined by incrementally increasing the antenna reconfiguration time constraint, as done in the previous scenarios.

Number of New Antennas at Svalbard	Maximum Allowable Antenna Reconfiguration Time	Mean Antenna Utilization (% from Reconfiguration)
+1	3 min	49% (11%)
+2	4 min	49% (14%)
+3	10 min	60% (26%)

Table 10: Antenna Utilization with Fairbanks Phased Array and Extra Capacity at Svalbard

Number of New Antennas at McMurdo	Maximum Allowable Antenna Reconfiguration Time	Mean Antenna Utilization (% from Reconfiguration)
+1	2 min	47% (8%)
+3	3 min	46% (11%)
+4	4 min	47% (14%)
+5	5 min	47% (16%)
+6	10 min	56% (27%)

Table 11: Antenna Utilization with Fairbanks Phased Array and Extra Capacity at McMurdo

Table 10 and Table 11 show the number of antennas needed to support the phased array at Fairbanks by adding new antennas at Svalbard and McMurdo, respectively. Adding new capacity to Svalbard is nearly twice as efficient as adding capacity to McMurdo, the former requiring a max of 3 antennas compared to the 6 required by the latter. This is likely due to the added capacity at Svalbard providing a better balance of capacity between the two hemispheres. This allows some satellites to take their pass in the North and others to take their pass in the South, thereby minimizing the total number of satellites competing at a single station. The antenna schedules for these two scenarios are shown below in Figure 16 and Figure 17.

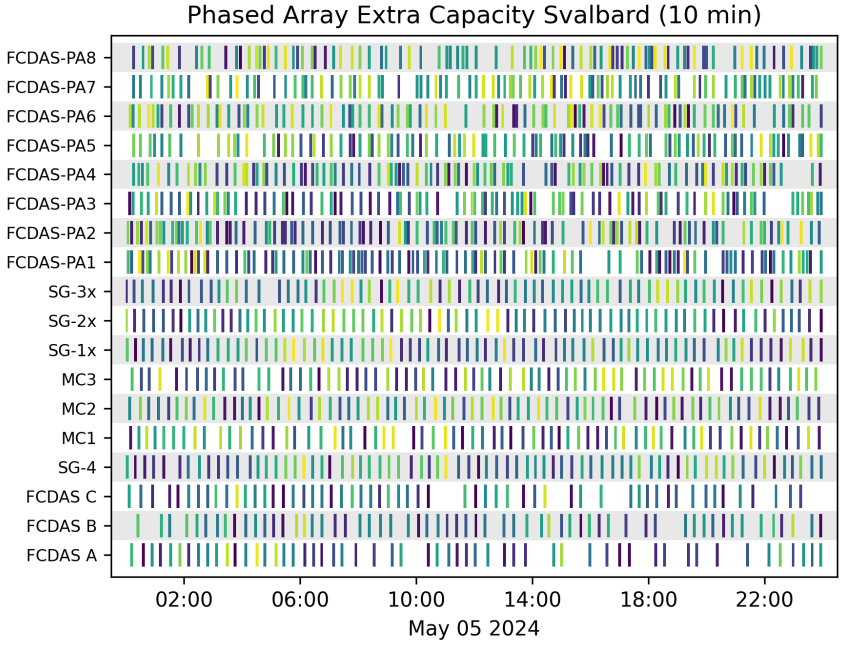


Figure 16: Antenna Schedule with Fairbanks Phased Array and Extra Antennas at Svalbard (10 min Reconfiguration Time)

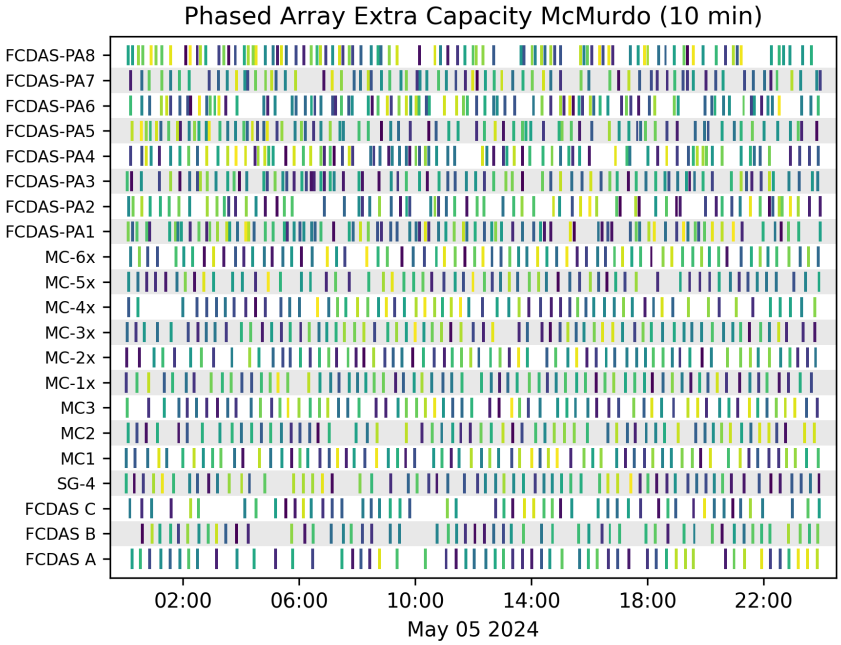


Figure 17: Antenna Schedule with Fairbanks Phased Array and Extra Antennas at McMurdo (10 min Reconfiguration Time)

Once the amount of new antennas required for each station was determined, all possible combinations of adding capacity to both antennas to support the full 10 min reconfiguration time were found to complete the trade space. These results are reflected in Table 12.

Number of New Antennas at Svalbard	**Number of New Antennas at McMurdo**	**Mean Antenna Utilization (% from Reconfiguration)**
+1	+5	58% (29%)
+2	+3	58% (28%)

Table 12: Antenna Utilization with Fairbanks Phased Array and Extra Capacity at Svalbard

4.3 Scenarios - Supporting Two Contacts per Revolution

Finally, all the previous scenarios shown thus far were attempted to support two contacts per orbit for each of the 100 satellites to identify any further overlap in solutions. In addition, one more scenario was investigated:

Dual Phased Array Scenario - Would building two phased arrays provide enough capacity to support all 100 satellites at two contacts per orbit? If so, where should they be built?

Each scenario was run varying constraints to explore the boundaries of the problem. Scenarios were evaluated on how well they supported meeting latency objectives, along with how resilient they were against scheduling conflicts.

4.3.1 Scenario Overlap.

There is no overlap in solutions, none of the previous scenarios in this analysis worked to provide coverage for two contacts per orbit for the 100-satellite constellation. Additional capacity is needed in both hemispheres to fulfill latency requirements, and needs to be enough to provide contacts for all 100 satellites every revolution.

4.3.2 Dual Phased Array Results

In this scenario two phased arrays were added to investigate whether they would each provide enough capacity to take a contact from all 100 satellites every orbit. Since capacity is needed in both the Northern and Southern hemispheres, one phased array was added at McMurdo and the other at Svalbard. Fairbanks was not chosen because it does not provide 100% coverage for the satellites. In order to test how much capacity this configuration could provide on its own all other NOAA antennas were left out of this scenario.

Table 13 shows that this configuration is indeed capable of supporting all 100 satellites at two contacts per orbit. Additionally, Figure 18 shows additional capacity is available in the system to support an even larger, or much less ideally spaced constellation. These findings are also quantitatively supported by the relatively low antenna utilization in Table 13. The large difference between the minimum and maximum antenna utilizations indicates that some antennas are very underutilized. The dual phased array configuration is ideal for supporting this constellation because each of these arrays provides support for more simultaneous contacts than provided by the entire network in the base scenario, and does so with no overhead required for antenna reconfiguration.

Antenna Reconfiguration Time	-
Mean Latency	47:44 min
Maximum Latency	56:42 min
Latency Met	✓
Minimum Antenna Utilization	6%
Mean Antenna Utilization	37%
Maximum Antenna Utilization	65%

Table 13: Antenna Utilization for Dual Phased Array

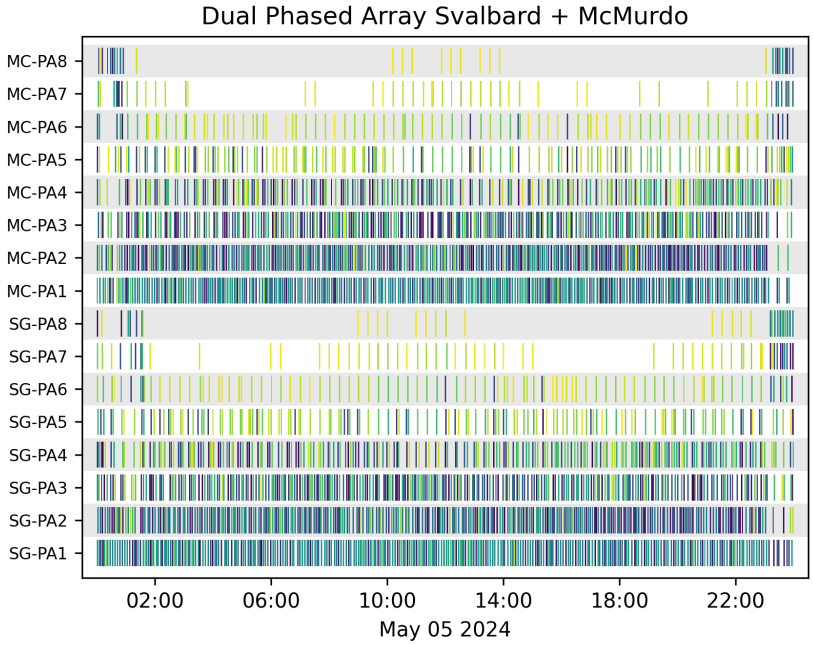


Figure 18: Antenna Schedule for Dual Phased Array

4.4 Summary

Overall, there is very little overlap in solutions between the 28-satellite solutions and the 100-satellite solutions. The current ground network as available to JPSS, defined in subsection 2.4, does not provide sufficient capacity to meet latency requirements while allowing for sufficient overhead.

KSAT. While a feasible schedule was found for the KSAT scenario, it is heavily constrained and over-utilized. Because it is not very resilient to scheduling conflicts and heavily relies on third-party providers, it is not an ideal solution for real-life operations, especially given the ideal spacing of satellites in this mock constellation configuration.

Phased Array at Svalbard/McMurdo. A phased array in either location provided enough capacity to support the full 100 satellites at one pass per orbit. Adding a phased array at Svalbard represents the only overlap in solutions between the two constellation configurations, making it an ideal candidate for building future capacity.

Supporting a Phased Array at Fairbanks. A phased array at Fairbanks can be supported as long as additional capacity is added to Svalbard or McMurdo. Although enough capacity at either location could support this constellation, adding new antennas at Svalbard is by far the most efficient way to do so. Adding capacity to Svalbard in any form is the best bet at improving ground system performance for future missions.

Dual Phased Array. The only viable way to support such a large constellation at two passes per orbit is to build a phased array at both Svalbard and McMurdo. Doing so would provide even more capacity than required for these 100 satellites, and provides plenty of resiliency against scheduling conflicts.

5 Discussion

5.1 Summary of Results

The main obstacle to contact scheduling stems from insufficient capacity in the Northern Hemisphere. Since FCDAS is not at high enough of a latitude to provide 100% coverage to every revolution, it creates a bottleneck in the single antenna that supports X-band at Svalbard (SG-4). This bottleneck causes unresolvable conflicts if either too many satellites require two contacts per orbit, the constellation is too large, or too much overhead (such as antenna reconfiguration) is associated with each pass.

In order to address this bottleneck, additional capacity must be added to Svalbard. This can come in the form of buying KSAT passes, converting SG-60 to support X-band, or building new antennas/phased arrays. Phased arrays are the most efficient way to add capacity to the ground system by supporting many simultaneous passes with no overhead required for reconfiguration between them. In particular, adding a phased array at Svalbard should both address the bottleneck at that location and buy plenty of capacity for future missions, while ensuring maximum resiliency against scheduling conflicts.

5.2 CNTIENT.Optimize

This analysis was performed to demonstrate the scalability and performance of CNTIENT.Optimize's ground contact scheduling optimization software. With simple configuration changes the product can easily adapt to changes in the ground network or satellite constellation, scaling easily to support larger sized scenarios such as the 100-satellite ones. Although none of the scenarios in this analysis represent current missions, they were designed to show how the product could be used in real-time operations.

CNTIENT.Optimize is primarily meant to be used in a real-time operational environment to generate mission schedules for satellites on-orbit. It is capable of supporting any mission, only requiring an instance to be configured to reflect all relevant constraints and infrastructure. It can be directly integrated with C2 systems to dynamically plan constellations, supporting many scenarios encountered in day-to-day operations such as replanning around satellite/antenna maintenance windows all while reporting relevant metrics on the schedules being generated.

5.2.1 Performance

Schedule Optimization. As shown in this demonstration, CNTIENT.Optimize can find truly optimal schedules for exceedingly large and complex problems. For example, it was able to produce the contact schedule shown in Figure 12, which was a heavily constrained problem requiring over a thousand contacts per day.

Solve Time. Typical run time for schedule optimization varies with the complexity and size of the configuration, but scales linearly with the number of days. For the 28-satellite constellation solve times were typically between 1 min to 2 min per day, requiring up to 5 min per day in the worst case for a heavily-constrained problem. For the 100-satellite constellation solve times were typically between 5 min to 10 min per day, requiring up to 20 min per day in the worst case for a heavily constrained problem.

Graph Neural Network. Research into using Graph Neural Networks (GNNs) to pre-prune the graph before solving the constraint programming problem shows a 90% reduction in solve times. Work is currently underway to bring this feature into the CNTIENT.Optimize product in support of other programs, and promises to bring down the solve times shown in this analysis even further.