Scheduling Downloads for Multi-Satellite, Multi-Ground Station Missions

Jeremy Castaing

Faculty Advisors: Amy Cohn and James Cutler
University of Michigan, Ann Arbor, MI

Abstract: In this paper we address the deterministic, multi-satellite, multi-ground station communication scheduling problem. As the number of small satellites in space increases, so does the demand for downloading their large quantities of acquired data. Given the capacity-constrained ground station network, efficient scheduling plays a major role in the overall performance of missions. The small-satellite dynamics of collecting, storing, using, and spilling both data and energy further complicates schedule design. In this paper we extend previous work on the single-satellite scheduling problem in order to incorporate simultaneously scheduling downloads from multiple satellites to a ground station network with the objective of maximizing the total amount of data downloaded to Earth. We assume that ground stations are restricted to receiving data from at most one satellite at a time and compare the results to those of the case where ground stations may receive data from multiple satellites concurrently in order to determine the potential download increase from such an enhanced communication capability. We create a greedy scheduling heuristic in order to compare our model's performance to a reasonable approximation of current scheduling methods. We test our model on a variety of scenarios generated from defined probability distributions to demonstrate how the model can be used to analyze the download performance of a satellite constellation. We also study the model's computational performance limits and sensitivity.

1 Introduction

Small satellites provide excellent opportunities to gather data from space and are significantly less expensive than larger, more traditional satellites. Given the recent technological advances and expansion of space applications [1], there has been a significant increase in the number of small satellites being launched into space. From 2001-2005 fewer than 15 satellites in the 1-50 kg range were launched annually. Approximately 25 satellites in the same weight range were launched annually from 2006-2012. In 2013 that number jumped to 92 satellites and is expected to rise an additional 52% in 2014 [2]. As the number of satellites in space increases, so does the total amount of acquired data. Given the finite download capacity of the ground station network, there is a growing need to efficiently schedule data downloads and make the most of the available resources.

The goal of this paper is to develop a model for scheduling data downloads from a constellation of satellites to a network of heterogeneous ground stations, so as to maximize the total amount of data downloaded to Earth while satisfying all energy and data requirements. We call this the deterministic *Multiple-Satellite, Multiple Ground Station Scheduling Problem (MMSP)*.

Our research makes multiple contributions to the small-satellite community. First, we develop a model capable of quickly solving real-world scheduling instances involving multiple satellites and multiple ground stations. Second, we identify the types of real-world instances where optimization provides the greatest gains over a simple greedy scheduling method. Third, we demonstrate the potential gains that result from ground stations that are able to simultaneously receive data from multiple satellites at once using a Multiple Spacecraft Per Aperture (MSPA) capability similar to that of the Deep Space Network (DSN) [3]. Lastly, we lay the foundation for future work that could incorporate equally distributed downloads or other factors such as prioritized data into the scheduling process.

The rest of the paper is as follows. Section 2 describes the details of our scheduling problem including important definitions and assumptions. In Section 3 we review the current literature relating to the problem of scheduling data downloads from space. In Section 4 we describe and present our model formulation. In Section 5 we present computational results from testing this model on a variety of problem instances and compare the results with both an alternative scheduling method and the case where ground stations have enhanced capabilities. We conclude by summarizing our findings and providing suggestions for future research.

2 Problem Description

The problem is to maximize the total amount of data downloaded from a constellation of satellites to a network of ground stations over a fixed planning horizon while satisfying the energy and data dynamics of the system. As opposed to the Single-Satellite, Multiple-Ground Station Scheduling Problem (SMSP) proposed by Spangelo [4], our MMSP involves multiple satellites which may simultaneously be in view of the same ground station. Throughout this paper, we define the situation where multiple satellites are simultaneously in view of a single ground station as a *conflict*. During *conflicts*, ground stations are restricted to communicating with one satellite at a time.

In order to maximize the total amount of data downloaded from space over some planning horizon $[0, T]$, we generate a schedule that determines, for specific time intervals, the optimal amount of data for each satellite to download and to which ground stations to download. We discretize the planning horizon by defining a set of time intervals I as time periods over which the view of each ground station (and thus each satellite) is constant. Whenever the view of any ground station changes, a new interval begins.

For example, in Figure 1, Ground Station 3 and Ground Station 1 can both see Satellite 1 and Satellite 3 during Interval 1. The view of all ground stations (and satellites) is constant throughout Interval 1. In Interval 2, Ground Station 3 can only see Satellite 1, while Ground Station 1 can only see Satellite 2.

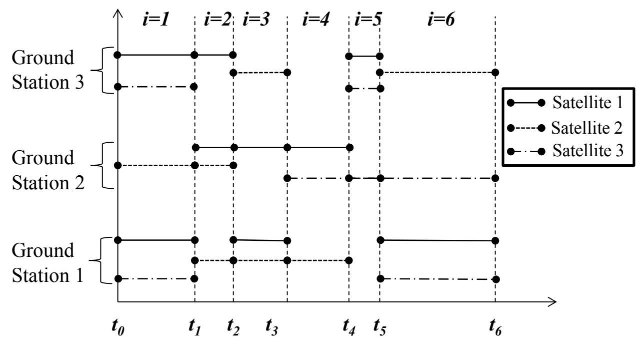

Figure 1: Defining Time Intervals

When creating a download schedule, energy and data dynamics must be considered. As satellites orbit Earth, they collect both energy and data at varying rates over time. Satellites must use their on-board energy to perform operations and to transmit their data to Earth. The ground stations to which they transmit their data vary in characteristics such as energy required per bit of downloaded data (joules), data rate of each download (bits per second), and efficiency of each download (i.e. the percentage of transmitted data that is successfully received).

2.1 Assumptions

- A constellation of satellites is orbiting the earth. Each satellite collects both energy and data at rates that vary over time depending on the view of each satellite and the line of sight of its solar panels relative to the sun. We assume the collection rates of both energy and data are linear with respect to time across any given interval.

- Each satellite has minimum and maximum levels of energy and data that can be stored in its buffers. For simplicity, we assume all satellites have the same energy and data storage capacities.

- While orbiting Earth, each satellite periodically comes in view of ground stations. Multiple ground stations may be in view at a particular time, but a satellite can only transmit data to one ground station at a time.

- The orbit of each satellite is deterministic and is known a priori, before any scheduling decisions are made. Data and energy acquisition rates for specific time intervals are known in advance. Likewise, download opportunities to specific ground stations at specific times are determined beforehand.

- The ground station network consists of multiple ground stations that each have a specific download efficiency rate, download data rate, and download energy cost associated with it. For simplicity, we assume each of these characteristics are constant over time, but this assumption can easily be relaxed.

- Each ground station can have multiple satellites in view at once, but can only receive information from one satellite at a time. In Section 5 we relax this restriction and analyze the results.

- We assume that there is no cost (time or energy) when a satellite switches between downloading from one ground station to another, or when a ground station switches between receiving from one satellite to another.

- We assume that data lost due to inefficiencies in the transmission from a satellite to a ground station is detected and is kept on the satellite for future download opportunities. For example, if a satellite attempts to download 10 MB of data but the ground station only receives 8 MB, that satellite keeps the 2 MB of data that weren't successfully transmitted and attempts to transmit them during a later opportunity.

2.2 Energy and Data Dynamics

To ensure that each satellite has enough energy and data to complete its scheduled downloads, we model and enforce the energy and data dynamics of the system. Solar energy is collected over time by each satellite's solar panels while energy is consumed in order to perform satellite operations and to download data. We can calculate $e_{(i+1)}$, the energy available to a satellite at the start of interval $i+1$, using the following recursive equation:

$$e_{(i+1)} = \min\{e_i + \delta_i^e - \alpha_i q_i,\ e_{max},\}$$

Here, δ_i^e is the net amount of energy gained during interval i, α_i is the energy consumption rate for downloading (joules/bit), and q_i is the amount of data downloaded (bits). Since each satellite has a maximum (e_{max}) amount of energy that can be stored on board, if more energy is acquired than what can be stored, the excess is lost.

Similar to the energy dynamics, data is collected over time by each satellite and is consumed when a satellite downloads its data to a ground station. We use an efficiency factor $\eta_{i,g}$ that represents the percentage of data successfully sent to ground station g during interval i. Therefore, we calculate $d_{(i+1)}$, the data stored on a satellite at the start of interval $i+1$, using the following recursive equation:

$$d_{(i+1)} = \min\{d_i + \delta_i^d - \eta_{i,g} q_i,\ d_{max}\}$$

Similar to the energy equation, δ_i^d is the amount of data acquired during interval i, q_i is the amount of data downloaded, and d_{max} is the maximum amount of data that can be stored.

We enforce these energy and data dynamics in our optimization model in order to generate schedules that are feasible in terms of both energy and data.

2.3 Download Example

To demonstrate how energy and data dynamics can create scenarios where optimization outperforms less sophisticated scheduling methods, consider the following example involving a single satellite:

Example 1: Download Scheduling

	Interval 1	Interval 2
Maximum download (bits)	10	10
Cost of download (joules/bit)	2	1

Here, the maximum download amount is based on the length of time available for the satellite to download data. Assume that the satellite has 24 joules of energy for use on downloads, over 20 bits of data, and does not gain any additional energy or data during this planning horizon. In this example, a myopic scheduling method would schedule the satellite to download the maximum of 10 bits of data during Interval 1, using 20 joules of energy. Then, the remaining 4 joules of energy would be used during Interval 2 to download 4 bits of data, for a total of 14 bits. Using optimization on this example would result in a schedule where the satellite downloads 7 bits of data during Interval 1 and 10 bits of data during Interval 2, for a total of 17 bits.

Despite this being a simple example with only one satellite and a few bits of data, it's clear that optimization can add value to the scheduling process. As more satellites, ground stations, and intervals are considered in the planning horizon, the added value of optimization becomes significant. With optimization, all future download opportunities in the planning horizon are evaluated and the best possible schedule is generated.

3 Literature review

Numerous articles studying satellite operations (imaging, acquiring or downloading data, resource management, etc.) can be found in both Aerospace and Operations Research journals. However, very few articles address the problem of conflicts in a multi-satellite network. This brief literature review aims to present some of these articles and show that our study makes an important contribution to this field.

Small satellites provide timely and often more affordable opportunities to gather data from space than larger, more traditional satellites. As a result, many organizations including universities are now building and using their own satellites. A description of additional advantages of small satellites can be found in [1] and [5]. Miniaturized satellites, often called Cubesats, are a popular choice for university research and are used for a variety of missions such as: [6],[7],[8],[9],[10],[11] and [12].

Since resources such as power and contact time are limited when using satellites, optimizing the operations schedule is crucial for acquiring data and communicating with ground stations. The extensive literature on scheduling for imaging satellites has a similar objective to our research: scheduling a sequence of typical tasks such as taking pictures to maximize a certain objective associated with those tasks while satisfying constraints related to the cost of operating a satellite. [13] compares different repair strategies for imaging satellites, while scheduling under uncertainty is studied in [14] and [15].

However, most studies neglect data and energy dynamics. These dynamics are described in a problem closely related to our research, the *Single satellite Multi-ground Station Problem* (SMSP), which addresses the download from a single spacecraft to a network of ground stations, such as the *Deep Space Network* [16], in a deterministic environment: [4] and [17].

Following the recent launch of numerous satellites [2], communication from a constellation of satellites to a network of ground stations has been studied in several articles: the satellites might be collaborating on a single mission or have different tasks [18]. A simulation approach to assess network capacity is introduced in [19], while a heuristic to resolve conflict amongst satellites trying to download to the same ground station is presented in [20].

For our research, we address the problem of scheduling downloads from a multi-satellite constellation to a multi-ground station network while enforcing energy and data dynamics. Although each of these components have been studied independently, we are unaware of any studies that have considered all of the components together.

4 MMSP Model Formulation

In this section we present a mathematical programming formulation of the MMSP model. This formulation enforces the data and energy dynamics described in Section 2.2 and schedules satellite downloads to ground stations over the planning horizon. Given the linear energy and data dynamics, this formulation can be solved using well studied and efficient linear programming methods. The objective is to maximize the total amount of data received by the ground station network.

We first describe the parameters and decision variables of the model:

Sets and Subsets

- S is the set of satellites.
- G is the set of ground stations.
- I is the set of time intervals.

Parameters

- $\gamma_{sig} = 1$ if satellite s is in view of ground station g during interval i. $\gamma_{sig} = 0$ otherwise.
- η_{ig} is the efficiency (fraction of downloaded data successfully received by the ground station) during interval i when downloading to ground station g.
- t_i is the duration of interval i, measured in seconds.
- ϕ_{ig} is the data rate associated with downloading data to ground station g during interval i, measured in bits/second.
- α_{ig} is the energy cost associated with downloading to ground station g during interval i, measured in joules/bit.
- e_{min}, e_{max}, and d_{max} are the minimum and maximum allowable amounts of energy and data to be stored in the buffer, measured in joules and bits.
- e_{start} and d_{start} are the amounts of energy and data stored in the buffers at the beginning of the planning horizon, measured in joules and bits.
- δ^e_{si} and δ^d_{si} are the total net amounts of energy and data that are acquired by satellite s during interval i, measured in joules and bits.

Variables

- $x_{sig} \in [0, 1]$ is a continuous variable representing the percentage of interval i during which satellite s downloads to ground station g.
- q_{sig} is the amount of data downloaded by satellite s during interval i to ground station g, measured in bits.
- e_{si} and d_{si} are the amounts of energy and data available for satellite s at the beginning of interval i, measured in joules and bits.
- h^e_{si} and h^d_{si} are the amounts of excess energy and data spilled by satellite s throughout interval i, measured in joules and bits.

$$\max \sum_{s\in S}\sum_{i\in I}\sum_{g\in G} \eta_{ig} q_{sig} \tag{1}$$

Subject to:

$$\begin{aligned}
& x_{sig} \le \gamma_{sig} && \forall s \in S, i \in I, \quad g \in G && (2)\\
& \sum_{s\in S} x_{sig} \le 1 && \forall i \in I, g \in G && (3)\\
& \sum_{g\in G} x_{sig} \le 1 && \forall s \in S, i \in I && (4)\\
& q_{sig} \le t_i \phi_{ig} x_{sig} && \forall s \in S, i \in I, \quad g \in G && (5)\\
& e_{s0} = e_{start} && \forall s \in S && (6)\\
& e_{min} \le e_{si} \le e_{max} && \forall s \in S, i \in I && (7)\\
& e_{s,i+1} = e_{si} + \delta^e_{si} && && \\
& \qquad - \sum_{g\in G} \alpha_{ig} q_{sig} - h^e_{si} && \forall s \in S, i \in I && (8)\\
& d_{s0} = d_{start} && \forall s \in S && (9)\\
& 0 \le d_{si} \le d_{max} && \forall s \in S, i \in I && (10)\\
& d_{s,i+1} = d_{si} + \delta^d_{si} && && \\
& \qquad - \sum_{g\in G} \eta_{i,g} q_{sig} - h^d_{si} && \forall s \in S, i \in I && (11)\\
& 0 \le x_{sig} \le 1 && \forall s \in S, i \in I, \quad g \in G && (12)\\
& q_{sig}, e_{si}, d_{si}, h^e_{si}, h^d_{si} \in \mathbb{R}^+ && \forall s \in S, i \in I, \quad g \in G && (13)
\end{aligned}$$

Description of each constraint:

- (1) The objective maximizes the total amount of data that is successfully downloaded from each satellite during the planning horizon.
- (2) Downloads are only allowed if the satellite is in range of the ground station.
- (3) Each ground station cannot receive data for more than 100% of each time interval.

- (4) Each satellite cannot transmit data for more than 100% of each time interval.

  Note: (3) and (4) ensure that we generate a schedule where satellites only download to one ground station at a time and ground stations only receive data from one satellite at a time.

- (5) The amount of data downloaded from a satellite to a ground station is limited by the time length of the download and the download rate of the ground station.

- (6-11) Data and energy dynamics

- (12-13) Definition of variables

5 Computational Results

We present computational results to address the following questions:

1. What types of scenarios benefit most from our optimization model?

2. How do the results of our model compare to those from a more common scheduling method?

3. How large of a problem instance can the model solve?

4. How sensitive is our model's solve time to the problem's parameters?

5.1 Data Generation

For computational tests, we generated data sets for a variety of problem instances. It should be noted that these problem instances are for testing purposes only and are not intended to replicate a real-world network of satellites and ground stations. However, given any real-world data set, we could conduct similar analysis. For parameters where real-world variation is expected, such as the energy gained during each time interval, we randomly generated the values from defined probability distributions. For each set of parameters tested in the experiments, we generated and sequentially solved 50 random problem instances. For each of the computational experiments, we started with the base case of parameters and distributions described below, and modified them as indicated in each specific experiment.

To generate the orbit information, we first generated a set of time intervals with uniformly random lengths between one and thirty seconds. Using the probabilities listed in Table 1, we randomly generated the number of ground stations each satellite would see during each time interval. We then randomly selected that number of ground stations from the set of ground stations using a uniform distribution.

Table 1: Base Case Model Parameters

Description	Default Value
Number of satellites	20
Number of ground stations	15
Number of time intervals	100
Time Interval Length	Uniform (1,30)
Prob (see 0 ground stations)	25%
Prob (see 1 ground station)	25%
Prob (see 2 ground stations)	25%
Prob (see 3 ground stations)	25%
Prob (see 4 ground stations)	0%
Prob (see 5 ground stations)	0%

Additional base case parameters indicated in Table 1 include the number of satellites, ground stations, and time intervals in the planning horizon.

Our base case satellite parameters are included in Table 2. We assume that satellites are storing their maximum amount of energy and data at the start of our planning horizon. As indicated, the amount of energy and data gained during each interval follow normal distributions that are restricted to non-negative values.

Table 2: Satellite Parameters

Satellite Data Descriptions	Default Value
Minimum energy level (e_{min}) (Joules)	0
Maximum energy level (e_{max}) (Joules)	100
Starting energy level (e_0) (Joules)	e_max
Maximum data level (d_{max}) (bits)	100
Starting data level (d_0) (bits)	d_max
Energy Gain (δ^e_{si}) (Joules per interval)	Normal (30,15)
Data Gain (δ^d_{si}) (bits per interval)	Normal (10,5)

Our base case ground station parameters are in Table 3. We truncate the efficiency parameter at 100% and round up negative realizations of the other ground station parameters to zero.

Table 3: Ground Station Parameters

Ground Station Descriptions	Default Value
Efficiency Percentage (η)	Normal (1,0.2)
Data Rate (ϕ) (bits/sec)	Normal (4,2)
Energy Cost (α) (Joules/bit)	Normal (5,2.5)

5.2 Unrestricted Ground Station Model

In addition to solving the MMSP problem as stated in Section 4, we also consider a relaxed version that allows ground stations to simultaneously receive downloads from multiple satellites. This relaxed model provides an upper bound on our model's objective function and provides insights on the system's total download potential given such a ground station capability [21]. The formulation of this model is obtained directly from the MMSP model by simply removing constraint (3): $\sum_{s \in S} x_{sig} \leq 1 \; \forall i \in I, g \in G$.

5.3 Greedy heuristic

One reasonable approach to scheduling downloads from a single satellite to a network of ground stations is to download as much data as possible during every download opportunity. However, as noted in Example 1, this greedy method may not generate an optimal solution. As a benchmark and alternative to our scheduling model, we propose the following scheduling heuristic for the MMSP:

1. Divide each interval of the planning horizon into k pieces (we use $k = 100$).
2. Start with the first piece of the first interval in the planning horizon.
3. Generate the set D of all possible download amounts from satellites to ground stations.
4. Select the maximum value in D, schedule the corresponding download from satellite s to ground station g, and remove any download involving satellite s or ground station g from D. Return to step 3 as long as the maximum value in D is positive and additional satellites and ground stations are available.
5. Update the energy and data available on each satellite for piece $i+1$. Move to piece $i+1$ and return to step 3.
6. End once the last piece of the planning horizon is scheduled.

This greedy heuristic creates a download schedule where satellites download to no more than one ground station during each piece of the interval and ground stations receive data from no more than one satellite during each piece of the interval. Using this method does not necessarily result in the maximum amount of data being downloaded for each individual time interval. Consider Example 2 below which involves two satellites and two ground stations. Since the best option is for Satellite 1 to download to Ground Station 1, this download will be scheduled and the total amount of data download will be 20 MB. However, scheduling Satellite 1 for Ground Station 2 and Satellite 2 for Ground Station 1 results in a total download of 25 MB, the best solution.

Example 2: Greedy Download Amounts

	Satellite 1	Satellite 2
Ground Station 1	20	10
Ground Station 2	15	0

5.4 Objective Value Analysis

In this section we explore objective value differences between the MMSP model and both the greedy scheduling heuristic and unrestricted ground station problem under a variety of parameter scenarios. Since the unrestricted model relaxes the restriction that ground stations can only communicate with one satellite at a time, the optimal objective value of the MMSP can be no better than that of the unrestricted problem. The unrestricted problem also provides insight into the potential gains from an enhanced ability that allows simultaneous communication with multiple satellites.

In general, we find that for problem instances where there are a large number of conflicts, where ground station utilization is low, or where energy acquisition is low, our model performs significantly better than the greedy heuristic and generates schedules with optimal values close to those of the unrestricted ground station problem.

5.4.1 Number of Satellites

In Figure 2 we fix the number of ground stations at 20 and increase the number of satellites from 10 to 130. We plot the objective values as a percentage of the unrestricted model's objective value to highlight the relative performance of each model. As the number of satellites increases and the constellation becomes more congested, the objective value of the greedy heuristic approaches that of the MMSP. This is due to the fact that for congested constellations, ground station utilization is close to 100% and it is therefore rarely beneficial for satellites to wait for future download opportunities. In instances where congestion is lower, MMSP's objective is very close to that of the unrestricted problem. Intuitively this makes sense since the higher the ground station utilization, the more beneficial it is to increase the download capacity of the ground station network, which is what the unrestricted problem does. The optimal objective value of the unrestricted model continues to increase with the addition of satellites since every satellite can download data at the same time. This increases the total capacity of the ground station network. In contrast, the ground station network of the MMSP and the greedy heuristic have a fixed download capacity and therefore there is an upper bound on the total amount of data downloaded. Since we plot the objective values as a percentage of the unrestricted model's, we observe the download trends of MMSP and Greedy in Figure 2.

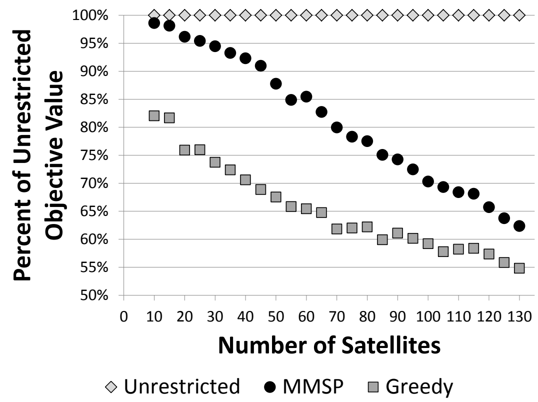

Figure 2: Objective Value Differences With 20 Ground Stations

5.4.2 Number of Ground Stations in View

In Figure 3 we fix the number of satellites and ground stations and increase the expected number of ground stations in view of each satellite during each interval. When the expected number is low, the objective value of the greedy heuristic is closer to the two optimizations. In these instances, satellites see ground stations less frequently and therefore are more likely to have enough data and energy for all of their download opportunities. Additionally, it is more likely that only one satellite is in view of a ground station at a time. Therefore, it often makes sense to download data whenever there is an opportunity. However, as the number of ground stations in view increases, more *conflicts* occur, and energy and/or data levels become more important, the value of optimization becomes more proclaimed. When each satellite expects to see 2 ground stations during each interval, the system appears to approach a steady state. This indicates that the satellites have enough opportunities to download their data, but do not have enough energy and/or data to use each opportunity. Therefore having additional opportunities for the satellites to download data is not beneficial, but having more energy would be beneficial.

5.4.3 Energy Acquisition

To assess the potential benefit from an improved ability to acquire energy, such as improved solar panels, we varied the average amount of energy acquired by each satellite during each interval. Results are included in Figure 4. Again, we are able to identify the point where the system reaches its maximum capacity and the marginal benefit of acquiring additional energy approaches zero. For instances where the average amount of energy acquired is large, the greedy heuristic performs almost as well as the optimizations because there is

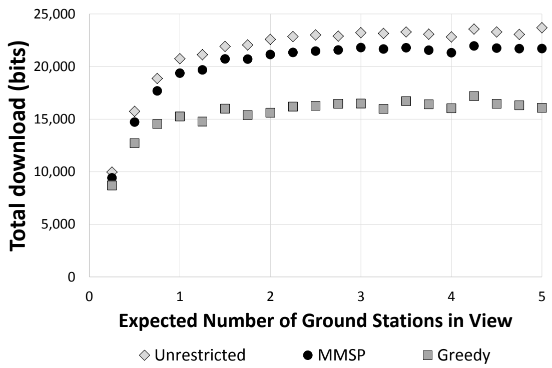

Figure 3: Objective Values for Expected Number of Ground Stations in View

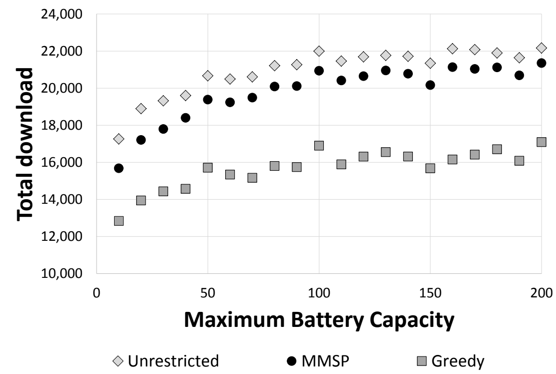

Figure 5: Objective Values for Energy Capacity

an abundance of energy and it is rarely beneficial for satellites to save energy for future download opportunities. However, in instances where less energy is acquired and satellites may become starved of energy, there is a clear and significant benefit to using optimization for scheduling.

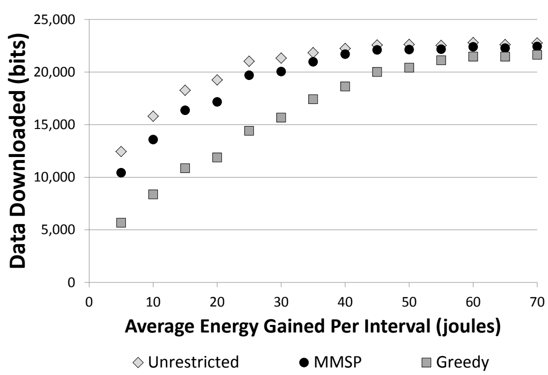

Figure 4: Objective Values for Energy Acquisition

5.4.4 Battery Size

To evaluate the potential benefit from an increase in the energy capacity of satellites, we varied the maximum amount of energy that can be stored on each satellite. Results are included in Figure 5. Given our base case of parameters, we can identify the point where the system reaches a steady state and there is no additional benefit from improved energy storage capabilities.

5.5 Solve Time Analysis

In this section we study the effects of various scenario parameters on the model's solve time. For each experiment we solved the model 50 times using an Intel Xeon E3-1230 quad-core running at 3.20 GHz with hyper-threading and 32 GB of RAM. We used IBM ILOG Optimization Studio (*CPLEX*) 12.6 C++ API software package. In general, we found that although the solve times increase for certain problem instances, the model solves quickly with a wide variety of parameter settings.

5.5.1 Number of time intervals

Using the base case of parameters, we varied the number of time intervals in our mission planning horizon. Figure 6 shows that although the time required to solve increases with the number of intervals, with 5000 intervals the average solve time is only slightly greater than one minute. For interval lengths of 5 minutes, 5000 intervals represents a planning horizon of 17 days. The vertical bars in Figure 6 represent the maximum and minimum solve times recorded for each scenario.

5.5.2 Number of Satellites

In addition to recording the objective values, we recorded the solve times for the experiment described in section 5.4.1. The experiment used the base case parameter values with 20 ground stations, varying the number of satellites from 10 to 130 in increments of 5. In this range, the average solve time increased approximately linearly from 0.14 seconds to 8.04 seconds. Although this is a relatively large

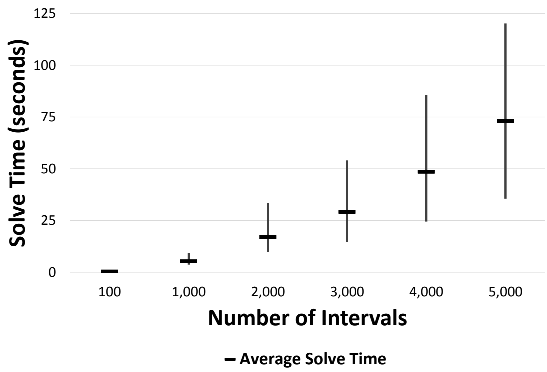

Figure 6: Minimum, Maximum, and Average Solve Times for Number of Intervals

increase, the solve time remains practical and does not increase at an exponential rate with an increasing number of satellites.

5.5.3 Number of Ground Stations in View

To test how solve times are affected by the number of *conflicts*, we varied the expected number of ground stations in view of each satellite during each interval. Increasing the number of conflicts increases both the number of download opportunities and the number of download options to consider for each opportunity. Starting with an expected value of 0.25 ground stations, we increased this value to 5 ground stations in increments of 0.25. Over this range the model's average solve time increased from 0.17 to 1.02 seconds. Although this is a significant increase, solve times remain fast.

6 Conclusion and Future Research

With the rapidly growing number of small satellites being launched, there has been a significant increase in the amount of data being collected from space. Given the energy and data constrained satellites, finite number of download opportunities, and limited capacity of the ground station network, it is essential to use the available resources as efficiently as possible. In this paper, we defined the *Multiple-Satellite, Multiple Ground Station Scheduling Problem* and presented an optimization formulation for solving it.

Through computational testing, we showed that our model quickly solves a wide variety of problem instances. In all instances, our model generates schedules that increase the total amount of data downloaded from space as compared to a traditional greedy scheduling method. Using a variety of computational experiments, we showed how our model can be used to identify the effects of satellite characteristics such as energy acquisition capabilities and data/energy storage capacities. By also solving the problem where ground stations have the ability to simultaneously communicate with multiple satellites, we are able to access the potential satellite download gains from such an enhanced communication capability. The ability to identify the bottlenecks of the system provides useful information about where improvements can be made to create the greatest impact.

Lastly, we identified the characteristics of problem instances where optimization provides the most benefit over a simpler scheduling method. Specifically, we found optimization most helpful for scheduling instances with less congested satellite constellations, where satellites see more ground stations during each time interval, and when satellites have low energy acquisition rates. In general, we conclude that optimization adds the most value when data and/or energy is a limiting factor for the satellites.

Although not considered in this paper, our model can easily incorporate multiple download options per ground station. Including additional options provides opportunities to optimize the tradeoffs between such things as data rate and energy consumption and further improve the overall performance of the system. We could also use our model to analyze the amount of data that is held on satellites over time due to limited capacities of the satellites or the ground station network.

For future research, the system's inherent stochastic nature could be modeled in order to solve problem instances where the specifics of download opportunities are not known in advance. Addressing prioritized downloads where specific downloads are more important than others, and possibly have expiration times, provides another area for future work. Furthermore, incorporating fairness into the download schedule is an important issue, especially for satellite constellations involving multiple, independent users. This issue could be addressed by guaranteeing each satellite some minimum amount of downloaded data based on its total download capacity. Lastly, we hope to apply our model to real-world data sets in order to help design the next generation of small satellites and create download schedules that optimally use the available communication resources.

7 Acknowledgments

We gratefully acknowledge the financial support of the National Science Foundation (CNS-1035236), the Air Force Research Laboratory (FA9550-12-1-0401), and the Bonder Foundation. In addition, we recognize the significant computational contributions of Mr. Brian Lemay.

References

[1] Daniel N Baker and S Pete Worden. The large benefits of small-satellite missions. *EOS, Transactions American Geophysical Union*, 89(33):301, 2008.

[2] Dominic DePasquale, AC Charania, Hideki Kanayama, and Seiji Matsuda. Analysis of the earth-to-orbit launch market for nano and microsatellites. *AIAA SPACE 2010 Conference and Exposition*, 2010.

[3] Mark D Johnston and Daniel Tran. Automated scheduling for nasa's deep space network. In *7th International Workshop on Planning and Scheduling for Space (IWPSS 2011), Darmstaft, Germany, June 8-10, 2011.* Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2011., 2011.

[4] Sara C Spangelo. *Modeling and Optimizing Space Networks for Improved Communication Capacity*. PhD thesis, University of Michigan, 2013.

[5] Kirk Woellert, Pascale Ehrenfreund, Antonio J Ricco, and Henry Hertzfeld. Cubesats: Cost-effective science and technology platforms for emerging and developing nations. *Advances in Space Research*, 47(4):663–684, 2011.

[6] Robert M Robinson and Therese Moretto. Small satellites for space weather research. *Space Weather*, 6(5), 2008.

[7] Geoff Crowley, Chad Fish, Charles Swenson, Robert Burt, Eric Stromberg, Tim Neilsen, Steve Burr, Aroh Barjatya, Gary Bust, and Miguel Larsen. Dynamic ionosphere cubesat experiment (dice). 2011.

[8] DE Rowland, AT Weatherwax, JH Klenzing, and J Hill. The nsf firefly cubesat: Progress and status. In *AGU Fall Meeting Abstracts*, volume 1, page 07, 2009.

[9] RP Lin, GK Parks, JS Halekas, DE Larson, JP Eastwood, L Wang, JG Sample, TS Horbury, EC Roelof, D Lee, et al. Cinema (cubesat for ion, neutral, electron, magnetic fields). In *AGU Fall Meeting Abstracts*, volume 1, page 09, 2009.

[10] X Li, SE Palo, DL Turner, D Gerhardt, T Redick, and J Tao. Cubesat: Colorado student space weather experiment. In *AGU Fall Meeting Abstracts*, volume 1, page 1585, 2009.

[11] DM Klumpar, HE Spence, BA Larsen, JB Blake, L Springer, AB Crew, E Mosleh, and KW Mashburn. Firebird: A dual satellite mission to examine the spatial and energy coherence scales of radiation belt electron microbursts. In *AGU Fall Meeting Abstracts*, volume 1, page 08, 2009.

[12] John Springmann, Benjamin Kempke, James Cutler, and Hasan Bahcivan. Initial flight results of the rax-2 satellite. 2012.

[13] Guo Yu-hua, Jing Ning, Li Jun, and Wang Jun. A comparison of iterative repair strategies for earth observing satellites imaging scheduling. In *Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, 2008. SNPD'08. Ninth ACIS International Conference on*, pages 93–98. IEEE, 2008.

[14] Da-Yin Liao and Yu-Tsung Yang. Satellite imaging order scheduling with stochastic weather condition forecast. In *Systems, Man and Cybernetics, 2005 IEEE International Conference on*, volume 3, pages 2524–2529. IEEE, 2005.

[15] JC Pemberton and LG Greenwald. On the need for dynamic scheduling of imaging satellites. *INTERNATIONAL ARCHIVES OF PHOTOGRAMMETRY REMOTE SENSING AND SPATIAL INFORMATION SCIENCES*, 34(1):165–171, 2002.

[16] William A Imbriale. *Large antennas of the deep space network*, volume 1. John Wiley & Sons, 2005.

[17] Laura Barbulescu, Jean-Paul Watson, L Darrell Whitley, and Adele E Howe. Scheduling space–ground communications for the air force satellite control network. *Journal of Scheduling*, 7(1):7–34, 2004.

[18] HJ Li, Y Lu, FH Dong, and R Liu. Communications satellite multi-satellite multi-task scheduling. *Procedia Engineering*, 29:3143–3148, 2012.

[19] SC Spangelo, JW Cutler, AT Klesh, and DR Boone. Models and tools to evaluate space communication network capacity. *Aerospace and Electronic Systems, IEEE Transactions on*, 48(3):2387–2404, 2012.

[20] Xiao Dong Ling, Wei Kang Zhu, Jin Mei Wu, and Xiao Yue Wu. Research of multi-satellite tt&c scheduling problem. *Applied Mechanics and Materials*, 263:476–484, 2013.

[21] Sara Spangelo and James Cutler. Analytic modeling framework and simulation toolkit for space network communication capacity assessment. *IEEE Transactions on Aerospace and Electronic Systems (In Progress)*, 2011.