

# RIDESHARE-INITIATED CONSTELLATIONS: FUTURE CUBESAT ARCHITECTURES WITH THE CURRENT LAUNCH MANIFEST

**Joseph W. Gangestad**

The Aerospace Corporation, 2310 E. El Segundo Blvd., El Segundo, CA, 90245, joseph.gangestad@aero.org

**James R. Wilson, Kristin L. Gates, and John V. Langer**

The Aerospace Corporation, 2310 E. El Segundo Blvd., El Segundo, CA

## Abstract

CubeSats have achieved growing credibility among government and commercial stakeholders as a valid architecture for future space systems, and many groups have proposed fielding CubeSat constellations for applications ranging from space-weather monitoring to space-based surveillance. This paper investigates the performance of CubeSat constellation designs that could be assembled using the current rideshare manifest. CubeSat launch opportunities have been identified to provide a baseline of possible orbits going through 2017. Coverage metrics including maximum revisit time, average revisit time, and daily visibility are used to evaluate constellation architectures that make varying use of the manifest, such as flying one or two satellites on every rideshare or using only rideshares with US launch vehicles. Despite the randomness of their orbital geometry, these “rideshare-initiated” constellations are competitive against traditional, symmetric constellations such as Walker constellations on several metrics. They also offer some resilience against attrition, to which their symmetric analogues are vulnerable. An evolutionary algorithm was used to explore the tradespace of rideshare selections and combinations, and several constellation designs were found that achieved <1-hour average revisit time with only 5 satellites or <1-hour 95th percentile revisit time with only 12 satellites. In general, the most effective rideshare-initiated constellations used a combination of high-inclination and low-inclination orbits to achieve patched, uniform coverage around the globe, a result that suggests future rideshare capability should be made available to a diverse set of orbits, as opposed to favoring one regime over another.

## Introduction

Over the last decade, the CubeSat form factor—a cube-shaped satellite 10 cm on a side with a mass of 1 kg—has become popular as a means of providing relatively easy access to space for technology demonstrations, scientific research, and education.<sup>1</sup> In the year 2013 alone, nearly 100 CubeSats were launched by universities, private corporations, and governments around the world,<sup>2</sup> and more than 100 were launched in 2014. The growing demands of picosatellite applications have driven the development of increasingly sophisticated small-satellite hardware and software and the implementation of systems-engineering techniques that facilitate rapid turnaround of missions.

The popularity of CubeSats coupled with a standardized deployment mechanism has created a new phenomenon: the commoditization of access to space. Whereas primary payloads acquire a launch vehicle years in advance and customize their interface to the launch vehicle, CubeSats are distinguished from other satellite architectures by their method of deployment: the P-POD. The Poly-Picosatellite Orbital Deployer is a rectangular tube with dimensions of approximately 30×10×10 cm, which can hold three CubeSats (or, more generally, 3 units or “3U” of volume), with a spring at one end and a door at the other. The satellites are inserted into the P-POD during integration, and the P-POD is attached to the launch vehicle (e.g., on the Atlas V, with a purpose-built carrier on the Centaur upper stage); once on orbit, the door opens on command and the spring force ejects the satellites. A launch vehicle may carry many P-PODs, deploying tens of satellites in one launch. Instead of being attached to a particular launch vehicle, a CubeSat can reach space on any rocket as long as it carries a P-POD.

Several private entities<sup>3–5</sup> provide integration services or have developed P-POD carriers that fly on the Atlas V, Falcon 9, and Dnepr launch vehicles and on the International Space Station. A CubeSat program pays for a rideshare (by volume) at a market rate that varies according to demand (e.g., some rideshare orbits are more desirable than others), with lead times for delivery seldom more than 6 months.

CubeSats have achieved growing credibility among government and commercial stakeholders as a valid architecture for future space systems, and CubeSat constellations have been proposed for applications ranging from space-weather monitoring to space-based surveillance.<sup>6,7</sup> Due to their small size and mass, a large number of CubeSats can be lofted to orbit to yield resilient constellations with short refresh times and global coverage. However, the challenge of reaching the application-specific orbits necessary for some proposals is often neglected. Companies and agencies may wish to move forward with CubeSat constellation concepts, but the community has seen little practical discussion about fielding these constellations from a rideshare perspective, instead relying on the expectation of dedicated launchers in the future. Although several small-satellite launchers are in development,<sup>8,9</sup> rideshare is likely to remain the most reliable access to space for CubeSats for the foreseeable future.

This paper investigates the feasibility, design space, and optimization of “rideshare-initiated” constellations; that is, constellations that are built up piecemeal over time by taking advantage of secondary payload opportunities. A careful selection of rideshares is necessary to achieve future missions that demand distributed or disaggregated architectures. Acquisition agencies that are considering small-satellite constellations have limited guidance on how many satellites are necessary to achieve their goals and which rideshares to select, and rideshare intermediaries have limited insight into where small-satellite operators could most benefit from the implementation of rideshare capability. The upcoming rideshare manifest into 2018 is used as a baseline to estimate the performance of these rideshare-initiated constellations, and this performance is compared against example large-scale constellations on orbit or in planning, such as COSMIC-2.<sup>10,11</sup> Thereafter, an optimization algorithm is applied to the manifest to identify the rideshares with the most utility for generating effective constellations.

## The Rideshare Manifest

Many public sources are available to identify the upcoming CubeSat rideshare opportunities\*. A table of rideshares (as of March 2015) appears in Table 1, subject to two constraints: 1) only those rideshares going to low Earth orbit (LEO) were included and 2) launches already dedicated to CubeSat missions (e.g., QB50 and Super-Strypi) were excluded. Several rideshare opportunities in 2016 and 2017 plan to deploy secondary payloads to apogees greater than 20,000 km, which would be suited only for specialized (and likely experimental) CubeSats. Also, a launch already dedicated for CubeSats, such as the QB50 mission<sup>†</sup>, would disqualify the launch from being available to other CubeSats. Because most launch vehicles perform an additional maneuver after releasing the primary payload but before deployment of secondary payloads, it can be challenging to determine the exact orbit into which a rideshare payload will be deployed. The orbit altitudes in Table 1 for Atlas V launches correspond to the expected rideshare orbit, but some of the Falcon 9 and Soyuz launches—particularly those beyond 2015—use the orbit of the primary. The orbit inclination is likely to remain fixed for a particular launch, with some variability in the apogee and perigee altitudes. For launches to Sun-synchronous orbit (SSO), we have attempted to identify the most likely local time of ascending node (LTAN) based on public information. Several launches in Table 1 have launch dates marked “?”, as the date may not yet be known or has not been released. In these cases, a date

\*Including: [www.spaceflightnow.com](http://www.spaceflightnow.com), [www.spaceflightservices.com](http://www.spaceflightservices.com), and others.

<sup>†</sup><https://www.qb50.eu/index.php/project-description-obj/launch-vehicle>, accessed 3 March 2015

**Table 1:** Upcoming CubeSat Rideshare Opportunities

#	Date	Vehicle	Apogee [km]	Perigee [km]	Inclination [deg]	LTAN	Comment
1	Apr 2015	Falcon 9	423	418	51.6		ISS resupply
2	May 2015	Soyuz	830	670	98.4	12:00?	SSO
3	May 2015	Atlas V	700	350	55		
4	Jun 2015	Falcon 9	423	418	51.6		ISS resupply
5	Jul 2015	Falcon 9	720	720	86		
6	Aug 2015	Atlas V	720	500	65		
7	Sep 2015	Falcon 9	423	418	51.6		ISS resupply
8	Nov 2015	H-2A	550	550	31		
9	Nov 2015	Atlas V	423	418	51.6		ISS resupply
10	Dec 2015	Minotaur	560	560	97.6	10:00?	SSO
11	Dec 2015	Falcon 9	423	418	51.6		ISS resupply
12	Dec 2015	Soyuz	836	836	98.8	9:30?	SSO
13*	? 2015	Minotaur	500	500	45		
14*	? 2015	? (Russian)	390	390	63		
15	Feb 2016	Falcon 9	421	418	51.6		ISS resupply
16	Mar 2016	Falcon Heavy	720	720	24.5		
17	May 2016	Falcon Heavy	520	520	24		
18	May 2016	Falcon 9	421	418	51.6		ISS resupply
19	Jun 2016	Falcon 9	620	620	98.9	6:00?	SSO
20	Jun 2016	Atlas V	680	680	98.1	22:30?	SSO
21	Dec 2016	Soyuz	836	836	98.8	9:30?	SSO
22*	? 2016	? (US)	600	500	63.4		
23	Nov 2017	Atlas V	550	450	94		
24	? 2018	Atlas V	700	350	55		

\*These launches and their orbits are listed by Spaceflight Inc. without information regarding precise dates or launch vehicles.

of 1 June was used.

This manifest is in constant flux due to launch slips, decisions by primary payloads to remove or add rideshare, and refinement of orbits as a launch date nears. While containing the most up-to-date information available to the authors, Table 1 cannot serve as a final or definitive list for future planning. Thus far, only one confirmed rideshare opportunity each in 2017 and 2018 has been identified, although those years are almost certain to be as plentiful as 2015 and 2016. Rather, Table 1 provides a representative distribution of orbits to which a CubeSat mission planner can expect to have access, and subsequent sections of this paper identify how the capabilities and coverage of a CubeSat constellation coverage build up over time with such a manifest and how rideshares can be selected judiciously to yield desired performance.


### ***Initial Conditions***

Evaluating the performance of a constellation, as through revisit and coverage metrics, requires the propagation of satellites in time, which in turn requires knowledge of initial conditions. The manifest in Table 1 does not provide a complete set of initial conditions. Six orbital elements and an epoch are necessary to initialize a satellite, and the table provides only three (apogee altitude, perigee altitude, and inclination). For satellites in SSO, the LTAN provides a fourth. The following assumptions were made regarding the CubeSat epochs and initial conditions:

- For resupply missions to the International Space Station (ISS), although several months can elapse between launch and deployment from ISS, the date of launch was used for initialization.
- For all SSOs, the right ascension of ascending node (RAAN) was calculated on the date of launch using the estimated LTAN. For all other orbits, the RAAN was initialized to 0 deg.
- The initial argument of perigee and mean anomaly at epoch were set to 0 deg. In cases where more than one CubeSat is launched on the same rideshare, the satellites are spaced evenly in mean anomaly. That is, for two satellites in a plane, the mean anomalies would be initialized to 0 and 180 deg. Although a wide separation in mean anomaly, such as that made possible via differential drag,<sup>12</sup> might take weeks or months to achieve, modeling that deployment process was considered beyond the scope of this work in light of the many other modeling simplifications made throughout.

Propagations were carried out using a force model that included Earth point-mass gravity and the first term of Earth oblateness (“ $J_2$ ”).<sup>13</sup> No atmospheric drag was included; the potential impact of drag on CubeSat orbit lifetime—and on how one uses the manifest—is addressed later.

Although all satellites were initialized with the same RAAN (excepting SSOs), argument of perigee, and mean anomaly, the orbit planes are not necessarily aligned, because each satellite has its own epoch. When calculating coverage statistics at a particular starting date, each satellite is first propagated from its initialization epoch to that date, a length of time which is different for each satellite. The perturbations from Earth’s oblateness, which cause secular precession of the RAAN and argument of perigee, move those orbital elements by different amounts, ensuring that the starting configuration of a constellation is spread out. For example, Fig. 1 depicts all of the orbits in Table 1 propagated to 1 Jan 2019, where the differences in altitude and inclination have led to spreading of the orbit planes.



**Figure 1:** The orbits of the entire rideshare manifest in Table 1 propagated to 1 Jan 2019, showing the natural separation of orbit planes due to differences in altitude and inclination.

## Coverage with the Current Manifest


### *Definitions*

The following sections present constellation coverage metrics of several different varieties. There are two broad categories:

- **Revisit Time:** the time interval between two periods of visibility to a ground point from one or more satellites. During a propagation, many such intervals of varying duration accumulate, and revisit time can be expressed as a statistic of these intervals. For example, “maximum revisit” would be the longest period of time over a propagation that a ground point does not see a satellite.
- **Daily Visibility:** the total amount of time each day that one or more satellites is visible to a ground point. When considering propagations over many days, this metric could be expressed as an average, minimum, maximum, or other statistic.

The surface of the Earth is broken up into a grid of latitude-longitude points, and for a given constellation and starting date, visibility intervals are calculated during the course of a propagation at each point. The visibility intervals are then distilled into visibility or revisit statistics.

Figure 2 shows an example world map of maximum revisit time for the 7-day period starting 2 Nov 2015. Each point on the Earth is color-coded according to the maximum revisit time: regions with short revisit are in blue, and regions with longer revisit are in orange and red. The constellation used to generate Fig. 2 included every rideshare orbit from Table 1 up to 2 Nov 2015, using 1 satellite per rideshare. For all visibility metrics in this analysis, a minimum elevation angle of 10 deg was used; that is, a satellite was considered visible only if it was greater than 10 deg above the horizon.



**Figure 2:** Example plot of maximum revisit time for the 7-day period starting 2 Nov 2015 with the current rideshare manifest, using 1 satellite per rideshare.

### ***Satellites on Each Rideshare***

The simplest approach to building up a rideshare-initiated constellation is to take advantage of every rideshare opportunity by acquiring one or more P-PODs on each launch. The cases of using 1 or 2 satellites on each launch are considered here. Figure 3 shows several plots of coverage metrics as a function of time as the constellation grows with each launch. On each plot, the horizontal axis is time, starting in June 2015 and ending in January 2018. The two colored lines on each plot correspond to their respective vertical axes: the blue line (left axis) shows the coverage metric (global maximum revisit, global 95th percentile revisit, and minimum daily visibility) and the green line (right axis) shows the accumulated number of satellites as the launch manifest progresses. By the start of 2018, 23 or 46 satellites have accumulated (the last launch, on 1 Jun 2018, does not appear on the plots) depending on the number of satellites per launch. The coverage metrics were calculated over consecutive 7-day intervals for the 30-month period.

The global maximum revisit time, in the first row of Fig. 3, was calculated by taking the largest maximum revisit time over the global grid for each 7-day run. This metric is highly variable throughout 2015–17. The addition of several satellites in 2015 drops the maximum revisit time from 600+ minutes into the range of 100–200 minutes, but the variability remains high through 2017, including a large spike in mid-2016. Because rideshare-initiated constellations by definition lack the symmetry and phasing of a planned constellation such as a Walker constellation,<sup>14</sup> there are occasional times that the haphazard phasing among the satellites conspires to yield a particularly poor (i.e. high) revisit time, which is captured in the maximum revisit plot.

The maximum revisit time may be a common metric used to evaluate the suitability of a constellation for a particular mission, but with a rideshare-initiated constellation the mission designer is hampered by the occasional poor performance in maximum revisit, even though the maximum case may only occur once in a year. If CubeSats are in general not considered platforms intended for near-100% reliability, it may behoove the mission designer to also consider relaxing the relevant revisit metric. Instead of imposing a requirement on maximum revisit time, which captures the worst performance over all time, one may consider the second row of Fig. 3, the 95th percentile revisit time. Whereas the maximum revisit time varies  $\pm 100$  min over the course of two years, the 95th percentile revisit is more stable *and* lower in absolute value. For 1 satellite per rideshare, the 95th percentile revisit stabilizes around 80 min; for 2 satellites per rideshare, around 40 min.

The third row of Fig. 3 shows the global minimum daily visibility of the accumulating constellation as a



**Figure 3:** Coverage metrics vs. time, building up a rideshare-initiated constellation using the rideshare manifest in Table 1. The left column shows metrics using 1 satellite per rideshare, and the right column shows metrics using 2 satellites per rideshare.

function of time; that is, the minimum total amount of time in one day that any point on Earth sees one or more satellites. Many locations on Earth have *more* daily visibility. At the start of the propagation, each point on Earth is covered at least  $\sim 2$  hours per day, but that amount rapidly increases with the number of satellites, such that by the end, each point is covered at least 7 hours a day for 1 satellite per launch and at least 11 hours a day for 2 satellites per launch. In early stages of build-up with few satellites, the amount of coverage increases rapidly with the number of satellites; because the satellites are randomly phased across planes, there is little overlap of coverage between two or more satellites. However, when the total number of satellites is high, later in the scenario, the benefit of going from 1 to 2 satellites per launch decreases: by 2017, doubling the number of satellites (by going to 2 satellites per launch) increases the visibility by 50%, whereas in 2015 the payoff was largely 2 to 1. This decreasing marginal return occurs because when 40+ satellites are distributed randomly in orbits around the globe, overlap in coverage becomes more common. In contrast, a deliberately constructed 40-satellite Walker constellation would be designed to avoid such overlap as much as possible.

### ***US Rideshares with Decaying Orbits***

The effect of drag on orbit lifetime has not yet been accounted for, although it is significant for rideshares that use ISS resupply missions. The lifetime for CubeSats deployed from the ISS is approximately one year. To capture the effect of ISS-deployed CubeSats leaving the constellation, a new scenario was evaluated where any ISS rideshare was removed after one year on orbit. Furthermore, to reflect the constraints for agencies that are enjoined from using non-US resources for access to space, the scenario was also limited to rideshares on US launch vehicles in Table 1. Figure 4 shows coverage statistics for this scenario versus time.

In this scenario, the number of satellites increases early on in the same fashion as in Fig. 3. However, by the end of 2016, the series of CubeSats from ISS deployments leave the constellation as they notionally reenter Earth's atmosphere. Despite the reduction in satellite population by about one third through late 2016 and 2017, the revisit statistics do not change appreciably, indicating that the ISS orbit does not contribute strongly to the tail of the distribution of the revisit times (i.e., the maximum and 95th percentiles). The global minimum daily visibility does track with the number of satellites in the constellations, peaking when all of the ISS-deployed CubeSats are on orbit and dipping as they leave the scenario.

### ***Example Baseline: COSMIC-2***

The previous sections considered the performance of rideshare-initiated constellations in isolation. Comparison to other constellations—particularly those designed for desirable revisit or coverage—provides insight into how many CubeSats might be necessary in a rideshare-initiated constellation to match or exceed its performance.

The example baseline considered here is based on the COSMIC-2 constellation, consisting of two sub-constellations of six satellites each in Walker constellations. One sub-constellation, planned for launch in 2016, is in a Walker 6/6/4 configuration with a low altitude (520 km) and inclination (24 deg); the second sub-constellation, which has not yet been manifested, is in a Walker 6/6/2 configuration with a higher altitude (720 km) and inclination (72 deg). The two sub-constellations are sometimes referred to as COSMIC-2 “Equatorial” and “Polar,” respectively.

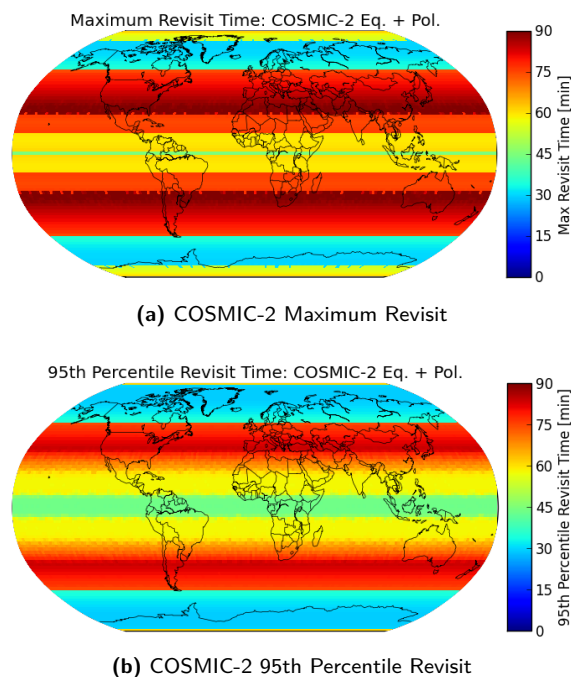
Global maps of the COSMIC-2 maximum and 95th percentile revisit times appear in Fig. 5. The mid-latitudes suffer from revisit times as high as 90 min, but the poles and equatorial regions are better served by the two sub-constellations, with revisit times not exceeding  $\sim 45$  min.

The primary payload on COSMIC-2 will collect weather data through the occultation of GPS radio signals by





**Figure 4:** Coverage metrics versus time for a rideshare-initiated constellation that uses only US launch vehicles *and* assumes that ISS-deployed CubeSats leave the constellation after 1 year due to atmospheric reentry. The left column shows metrics using 1 satellite per rideshare, and the right column shows metrics using 2 satellites per rideshare.



**Figure 5:** Global maps of maximum and 95th percentile revisit time for the complete, 12-satellite COSMIC-2 constellation.

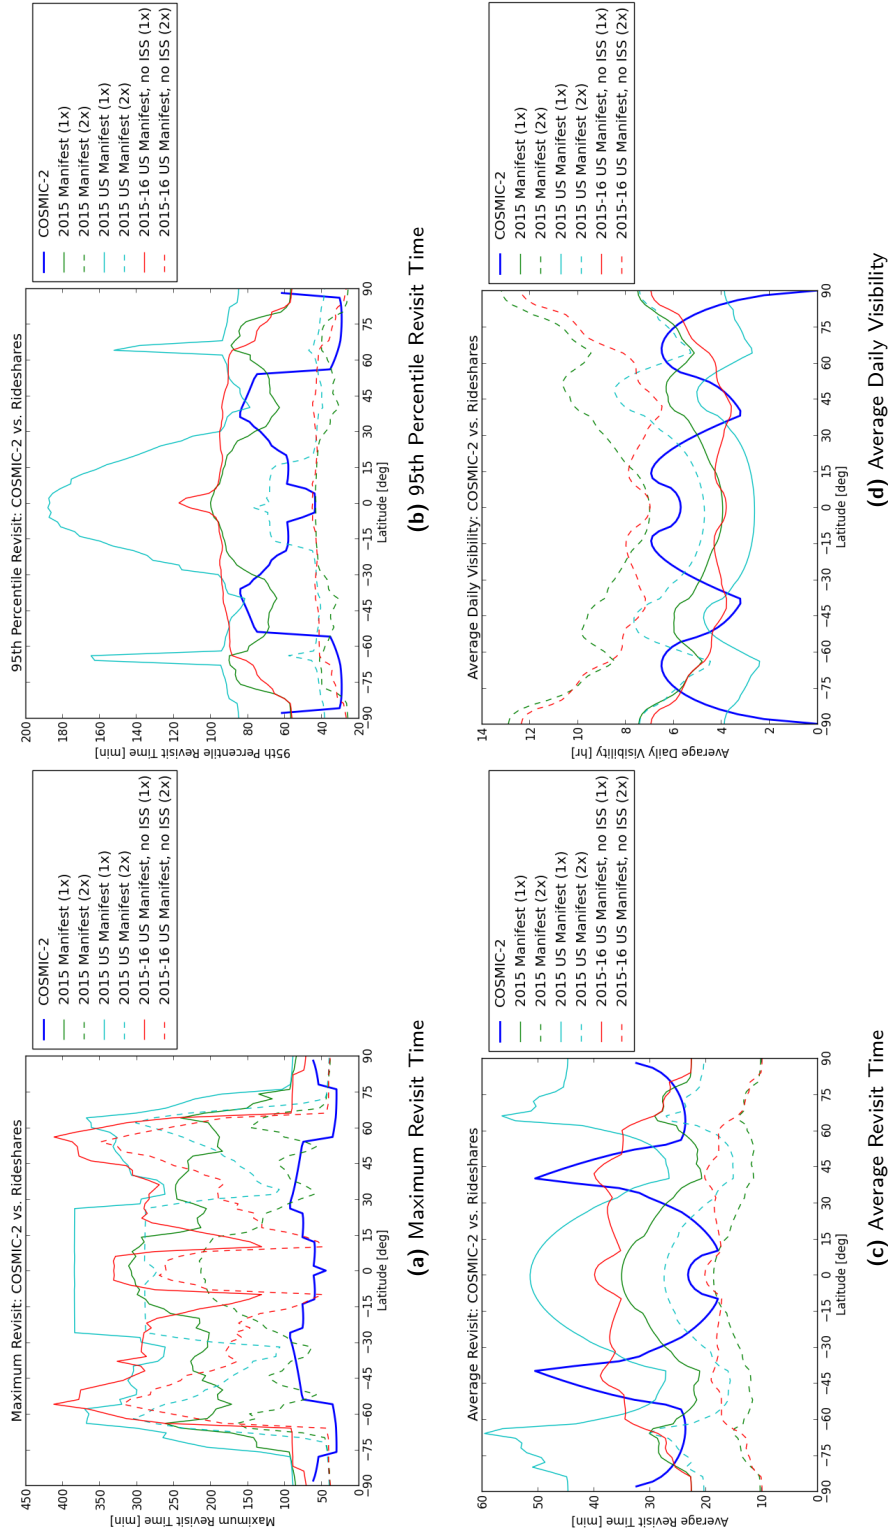
the Earth's limb, a mission that is not driven directly by coverage requirements over the Earth's surface. However, COSMIC-2 does serve as a valuable baseline for comparison to rideshare-initiated constellations for three reasons: 1) COSMIC-2 is one of the only large-scale constellations planned for deployment in the near future, providing a challenging test case to match with a rideshare-initiated constellation, 2) weather has been identified as a near-term candidate mission for small, micro-, and nanosatellites, and 3) the global distribution of radio occultation measurements is affected, in part, by the global coverage pattern of the constellation.

The plots in Fig. 6 compare four key coverage statistics between COSMIC-2 and several subsets of the rideshare manifest as a function of latitude. For both the maximum and 95th percentile statistics, the maximum value along a line of latitude for that statistic was taken to represent that latitude; for the average revisit and average visibility, the average value along the line of latitude was used.

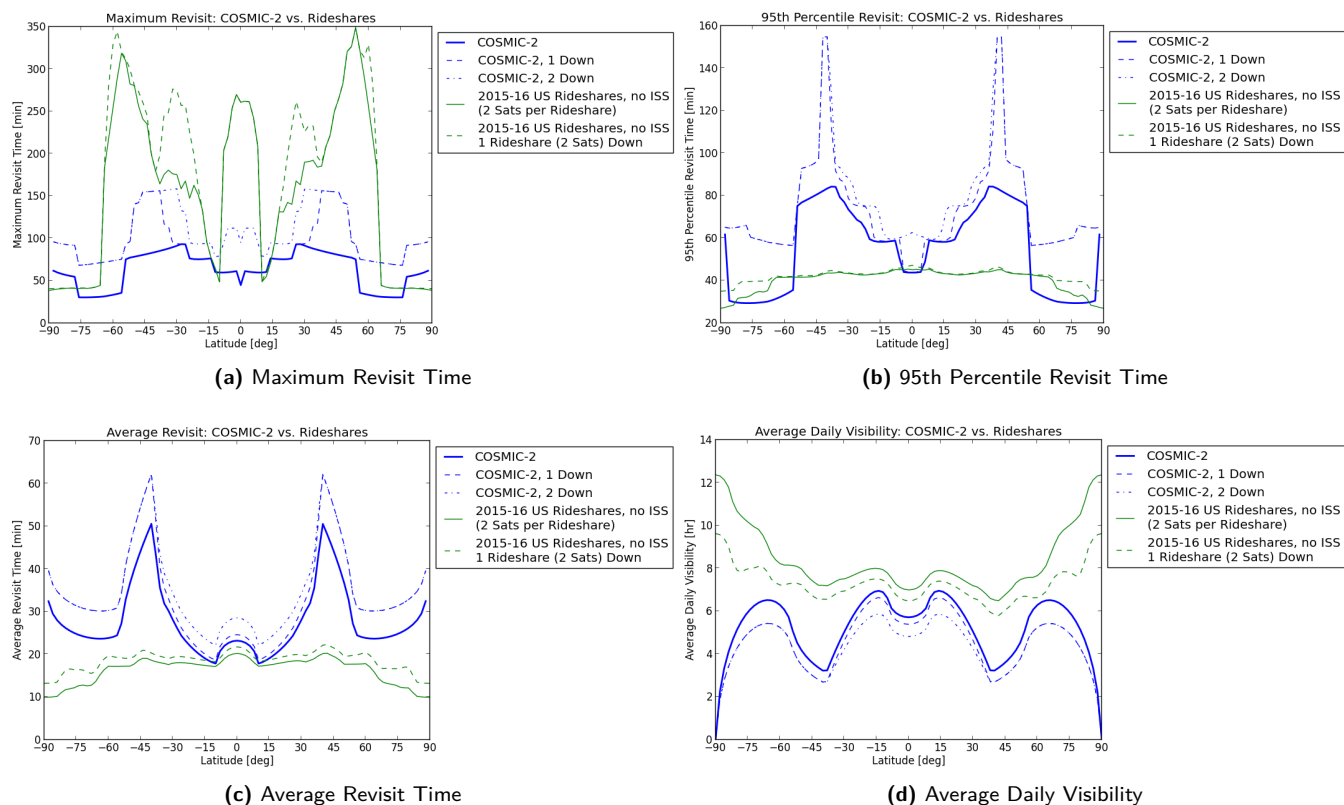
Each coverage statistic for COSMIC-2 is plotted in Fig. 6 with a thick blue line. Three subsets of the rideshare manifest were used: 1) the entire manifest in 2015, 2) the US-only manifest in 2015, and 3) the US-only manifest, excluding ISS deployments, in 2015 and 2016. Each subset was run with both 1 and 2 satellites per launch. The three cases are plotted in green, cyan, and red, respectively, and in solid and dashed style for the 1- and 2-satellite-per-launch cases.

In the case of maximum revisit, the symmetrical phasing of COSMIC-2's Walker constellations ensures excellent performance compared to the rideshare-initiated cases. None of the rideshare opportunities in 2015 go to inclinations as low as the COSMIC-2 Equatorial orbit, making it challenging to match maximum revisit within 30 deg of the equator. Maximum revisit for COSMIC-2 does not exceed 100 min, whereas most of the rideshare constellations are a factor of two or more higher. Even with 2016 rideshares that do have low inclinations, the maximum revisit remains largely higher than COSMIC-2's.

The 95th percentile revisit, on the other hand, shows promising results compared to COSMIC-2. For all three subsets of the manifest, using only 1 satellite per launch is insufficient to match COSMIC-2, but 2 satellites per



**Figure 6:** Comparison between COSMIC-2 and several subsets of the rideshare manifest. The “1x” and “2x” in the legend designate 1 or 2 satellites per rideshare, respectively.



**Figure 7:** Comparison between COSMIC-2 and the US-only subset of the rideshare manifest (*sans* ISS) when one or more satellites are down. This rideshare-initiated constellation uses 2 satellites per launch. In the case of COSMIC-2, the “1 down” case corresponds to the loss of a polar satellite, and “2 down” corresponds to the loss of one polar and one equatorial satellite. For the rideshare-initiated constellation, an entire rideshare is treated as down, removing 2 satellites from the constellation.

launch exceed COSMIC-2’s performance at nearly all latitudes. The same result holds for average revisit and average daily visibility. The dashed red lines in Fig. 6, which correspond to the US-only manifest without ISS deployments, uses 20 satellites (2 per launch) to achieve comparable revisit and daily visibility statistics as the 12-satellite COSMIC-2 constellation.

Another aspect of constellation operations to consider is the impact of a satellite failure on orbit. For a Walker constellation, the symmetry of which has a strong influence on its desirable coverage statistics, the loss of one or more satellites could seriously degrade the constellation’s performance. Figure 7 shows plots of the COSMIC-2 and rideshare constellations including the loss of one or more satellites (or rideshares). The solid lines correspond to the baseline constellations, and the dashed lines to correspond to scenarios where satellites have been lost. For COSMIC-2 (blue), the “1 down” case corresponds to the loss of a polar satellite, and “2 down” corresponds to the loss of one polar and one equatorial satellite. No re-phasing of the COSMIC-2 constellation has been assumed; the plots in Fig. 7 show the impact to coverage immediately after a satellite loss. The rideshare-initiated constellation (green) uses 2 satellites per rideshare with only US launches and excludes ISS deployments. Figure 7 includes the impact of losing one rideshare (i.e., two satellites in the same plane), specifically launch #5 in Table 1, which is

an 86-deg inclined orbit at 720 km altitude.

For both COSMIC-2 and the rideshare-initiated constellation, the loss of one or two satellites has a strong impact on the maximum revisit, roughly doubling the maximum revisit time at the mid-latitudes. However, the difference in impact diverges between the constellations in the case of 95th percentile revisit. With one or two satellites down, the 95th percentile revisit for COSMIC-2 doubles in the mid and high latitudes. This impact is due to the loss of symmetry of COSMIC-2's Walker constellations. In the baseline constellation, each ground point is visited regularly by each evenly phased satellite in succession. When a satellite is lost, a significant gap in coverage opens that the symmetry of the surviving satellites ensures will never be filled. In contrast, the loss of a 2-satellite orbit plane in the rideshare constellation has almost no effect on the 95th percentile revisit, except for a  $\sim 25\%$  increase at latitudes above 60 deg. In this case, the random distribution of orbit planes in the rideshare constellation works to the advantage of the mission designer, who can count on a satellite randomly filling the gap at least 95% of the time. Average revisit time is impacted roughly the same for both constellations, with a  $\sim 50\%$  increase with the loss of two satellites, as is the average daily visibility, which is driven more by the number of satellites on orbit than by their relative phasing.

## Optimized Rideshare Selections

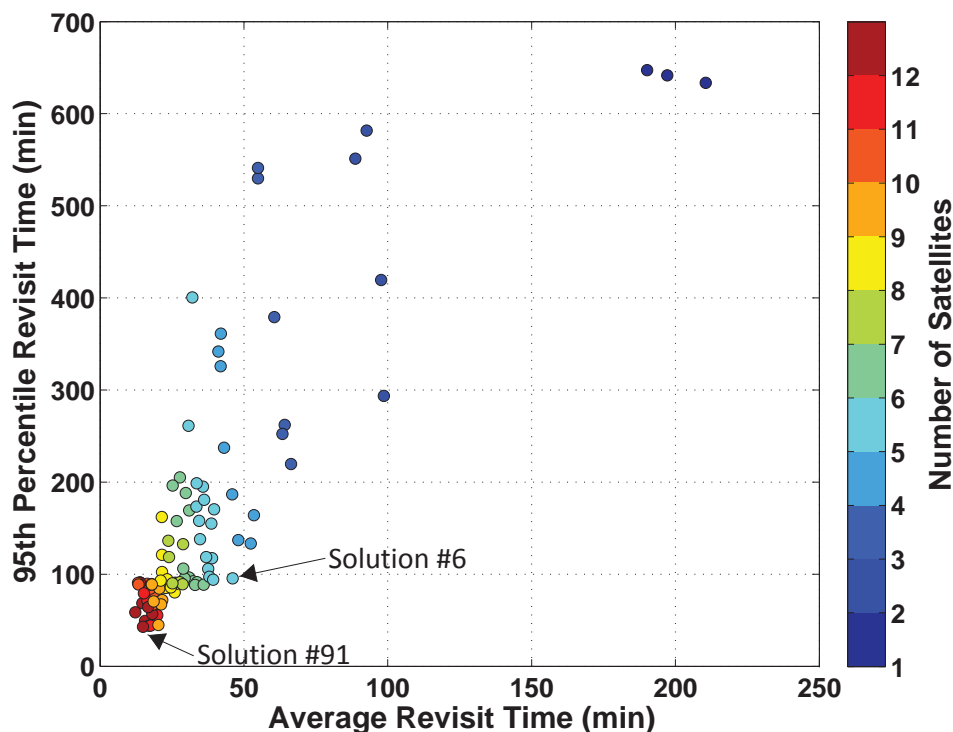
Adding CubeSats incrementally to each rideshare opportunity in the manifest is effective at generating a distributed constellation that on some metrics competes well with large-scale, pre-planned constellations such as COSMIC-2. However, the number of satellites required for such a scheme may be challenging for some programs to acquire (up to 40 or more CubeSats), and it is not clear which rideshares have the greatest effect on coverage statistics. A program with limited resources—even at the CubeSat scale—may be limited to perhaps 10 or 12 CubeSats, in which case more insight is necessary to select the most effective rideshares that achieve the desired coverage.

To explore the tradespace of rideshare selections, we have utilized a tool called GRIPS, Genetic Resources for Innovation and Problem Solving, developed by The Aerospace Corporation.<sup>15–17</sup> GRIPS uses an evolutionary algorithm to find non-dominated solutions to a problem subject to several different, often conflicting, objectives. In the context of this paper, we seek to find a series of rideshare selections (the solutions) that provide the best performance on several different coverage metrics (the objectives), such as maximum revisit time and average revisit time. The non-dominated solutions identified by GRIPS are equally optimal across all objectives simultaneously. That is, each solution cannot be changed to improve one metric without suffering on one or more other metrics. For example, if one non-dominated solution (a series of rideshare selections) provides the lowest maximum revisit time but a high average revisit time, another solution with a lower average revisit time must necessarily have higher maximum revisit time.

For exploring the tradespace of rideshare-initiated constellations, the following objectives were used in GRIPS:

1. Minimize global maximum revisit time,
2. Minimize global average revisit time,
3. Minimize global 95th percentile revisit time, and
4. Minimize total number of satellites.

A maximum of 12 total satellites was permitted. The number of launches was unrestricted: because access to space with CubeSats has been commoditized with the P-POD, each satellite is agnostic to the selection of launch vehicle. The cost of launching 12 satellites on 1 launch vehicle is the same as launching 2 satellites each



**Figure 8:** GRIPS output of rideshare-initiated constellations that tradeoff 95th percentile revisit time, average revisit time, and the total number of satellites. Each point on the plot corresponds to an entire constellation with a particular selection of rideshares. Two solutions addressed in more detail are marked.

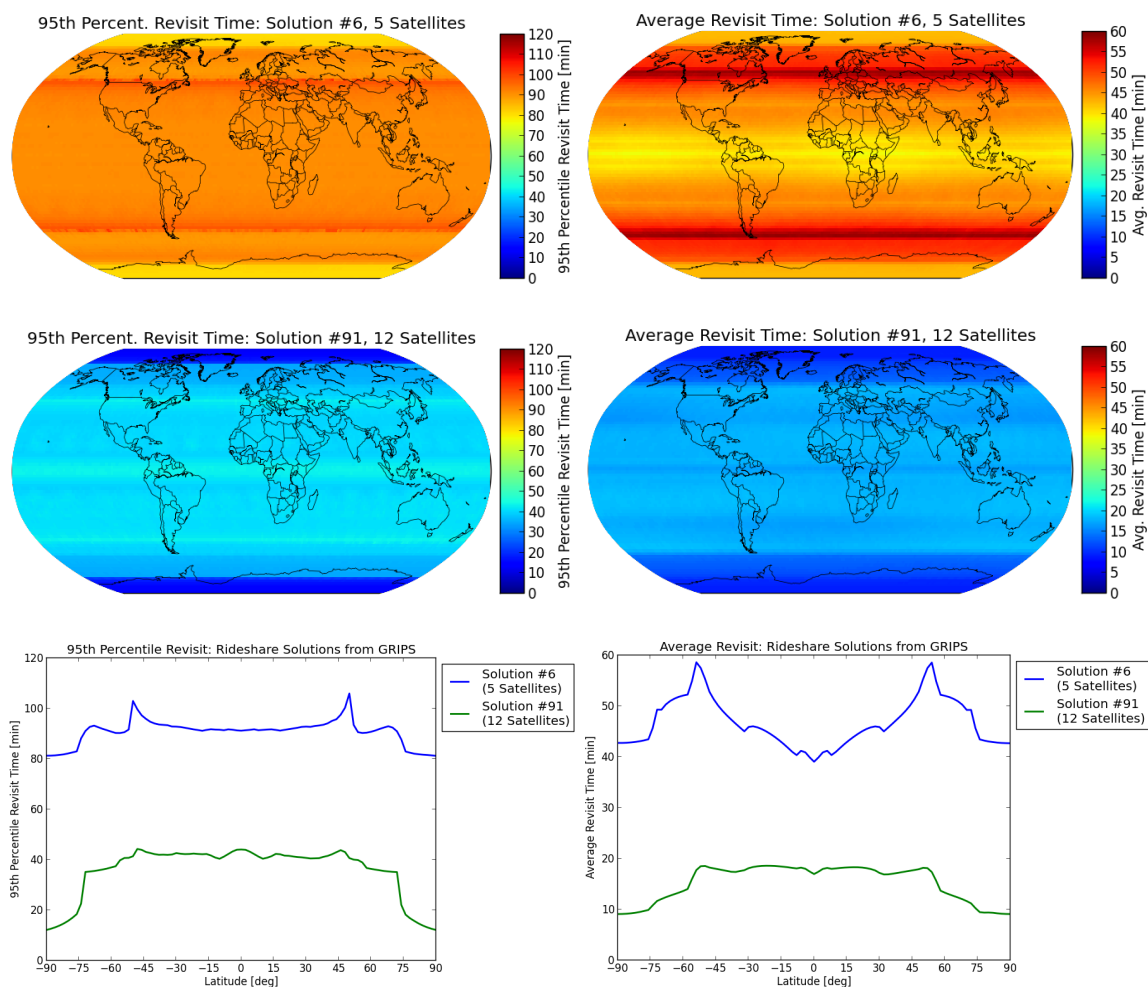
on 6 launch vehicles. The input for optimization was the rideshare manifest in Table 1. GRIPS was free to select any combination of rideshares from the manifest and the number of satellites on each rideshare, subject to the maximum constraint of 12 total satellites. A GRIPS run produces many candidate solutions, each of which is a series of rideshare selections, and these solutions form a Pareto front of non-dominated rideshare-initiated constellations.

Figure 8 shows results from GRIPS, plotting 95th percentile revisit time versus average revisit time and using color to represent the number of satellites. The fourth objective, maximum revisit time, was not included in this plot. For this GRIPS run, the minimum elevation angle at the ground was reduced to 0 deg, meaning that the satellite was considered visible if it was above the horizon. In Figure 8 the most desirable direction is towards the lower left, reducing both 95th percentile and average revisit times. However, finding solutions in that direction involves the tradeoff of an increased number of satellites. The majority of solutions appears in the lower left, with 95th percentile revisit times between 0 and 100 min, average revisit less than 50 min, and using as few as 5 satellites.

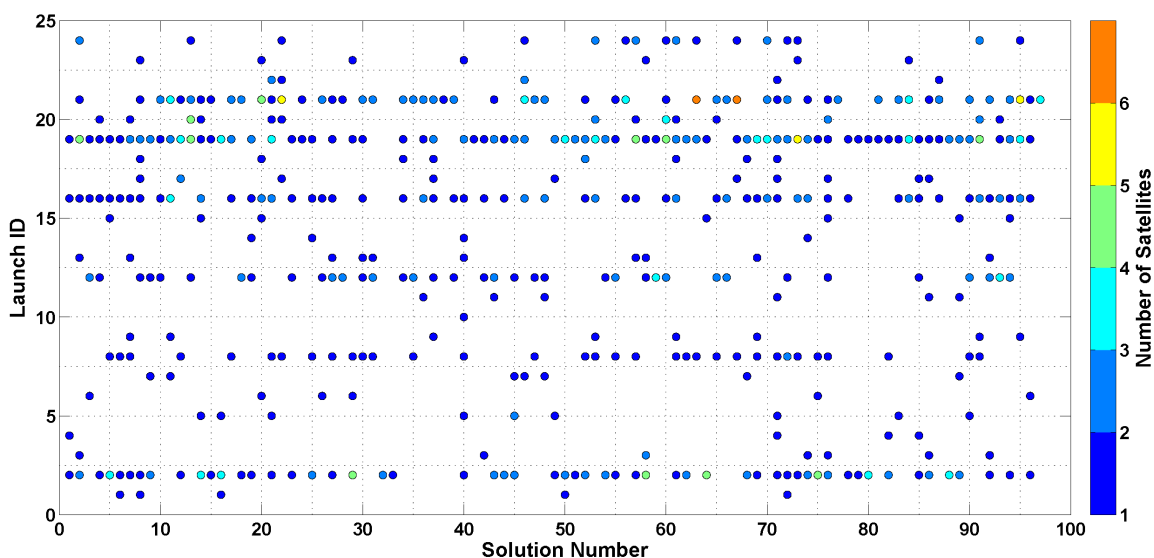
From the set of solutions in Fig. 8, the coverage from two examples was plotted in Fig. 9. The two solutions are labeled “Solution #6” and “Solution #91,” using indices that correspond to the solution space presented in Fig. 10, which is discussed in the next section. The launches used by the two solutions appear in Table 2. The left column of Fig. 9 shows the 95th percentile revisit time and the right column the average revisit time for the two constellations. The top row, Solution #6, uses 5 satellites to achieve 95th percentile revisit times around 90

**Table 2:** Launches from Example Solutions

	Launch ID (see Table 1)
Solution #6	1, 2, 8, 16, 19
Solution #91	8, 9, 16 (x2), 19 (x4), 20 (x2), 24 (x2)



**Figure 9:** The 95th percentile revisit time (left column) and average revisit time (right column) for two rideshare-initiated constellations found by GRIPS. The bottom row plots the two metrics as a function of latitude.



**Figure 10:** The launch selections for each candidate constellation generated by GRIPS. The launch ID on the vertical axis corresponds to the first column in Table 1. The color of each point denotes the number of satellites assigned to that particular launch.

min and average revisit times ranging from 40 min near the equator up to ~55 min in the mid-latitudes. The middle row, Solution #91, uses 12 satellites to reduce 95th percentile revisit time to ~40 min and average revisit time to <20 min. Referring back to Figs. 3 and 4, which showed the gradual build up a constellation by flying CubeSats on each rideshare, revisit times of this kind were not achieved until well more than 20 CubeSats were on orbit. The solutions from GRIPS identified the most effective selections and combinations of rideshares to achieve desirable coverage statistics, eliminating extraneous launches that do not benefit metrics like revisit time.

### ***Most Effective Rideshares***

One goal of exploring the rideshare tradespace with GRIPS was to identify those rideshare opportunities that were most effective at yielding desirable coverage statistics. The tradespace in Figure 8 shows 96 solutions (constellations) that were retained by GRIPS, and each solution used a different combination of rideshares to achieve its respective coverage. Figure 10, in turn, breaks down each solution according to its use of rideshares. The horizontal axis denotes each solution generated by GRIPS, and the vertical axis shows the launch ID of each satellite in the constellation. The launch ID numbers match the first column of the rideshare manifest in Table 1. The color of each point in Fig. 10 denotes the number of satellites assigned to that particular launch.

Six launches are clearly favored: 2, 8, 12, 16, 19, and 21. Rideshares #8 and #16 are launches to low inclinations, 31 and 24.5 deg, respectively and the other four are to Sun-synchronous orbits. Rideshares #19 and #21 receive the lion's share of satellites, with many solutions having 3–6 satellites in those orbits. Rideshare #21 in particular is likely preferred over other Sun-synchronous orbits because of its high altitude (836 km), which provides an edge for satellite coverage. In practice this orbit altitude is probably reserved for the primary payload; the orbital lifetime for CubeSats at 836 km is too long to satisfy most debris-mitigation requirements. Without any additional information about the rideshare payloads, the published orbit altitude was used in this analysis.

The selection of Sun-synchronous orbits appears intuitive—the nearly polar orbits provide coverage over all latitudes—but the simultaneous preference for the low-inclination Rideshares #8 and #16 may not be. The three



polar orbits provide excellent coverage at high latitudes, which are inaccessible to the vast majority of the 45–65 deg inclined rideshares in the manifest, but the high inclination comes with the tradeoff that coverage is reduced near the equator, where the polar-orbit ground tracks are at their most separated. The evolutionary algorithm in GRIPS solves this problem for many of the rideshare constellations by adding one or more satellites to a near-equatorial orbit. The combination of near-polar and near-equatorial orbits patches the coverage together to yield desirable metrics, an architecture not dissimilar from the pair of COSMIC-2 sub-constellations in low and high inclinations.

## Conclusions

As CubeSats grow in popularity and rideshares become more plentiful, the deliberate planning for rideshare selections is necessary to ensure an effective allocation of limited resources while trying to achieve desirable coverage with a nanosatellite constellation. The evaluation of these constellations may require mission designers to revisit their traditional metrics. The inherent randomness in a rideshare-initiated constellation's orbit planes makes maximum revisit a challenging statistic to match against more traditional, symmetric constellations like Walker constellations. However, rideshare-initiated constellations can begin to match traditional constellations with 95th percentile revisit time, average revisit time, daily visibility, and other metrics. In an environment where nanosatellite concepts typically achieve lower cost by accepting lower reliability or availability, shifting performance requirements from maximum revisit time to a more relaxed metric, such as 95th percentile revisit, may make for a valuable compromise. Furthermore, the randomness inherent to the rideshare-initiated constellation brings with it some benefit of resilience against attrition: whereas the performance of a symmetric Walker constellation can be severely compromised by the loss of one or two satellites, coverage in a rideshare-initiated constellation does not always depend on a specific satellite or plane. In one example, it was shown that the loss of two satellites in one plane of such a constellation had almost no effect on the global 95th percentile revisit time. This advantage of randomness may also have implications for large space programs and constellations that seek resiliency and also have the advantage of selecting their own orbits.

With the power of evolutionary algorithms, it is possible to explore the large tradespace of rideshare selections and combinations to identify scenarios that achieve the best performance with the fewest number of satellites. With the current rideshare manifest, improvements in coverage level off roughly when the 6th satellite is added to the constellation; the addition of more satellites largely benefits the 95th percentile revisit time. By carefully selecting the rideshares for these 6-12 satellite constellations, it is possible to achieve global revisit statistics that required 20 or more satellites when each available rideshare in succession. Specific launches from the manifest stand out as the most effective for global revisit and coverage statistics: those in Sun-synchronous and in low-inclination orbits. The combination of these two regimes patches coverage in the equatorial and polar regions to achieve favorable, and roughly uniform, coverage across the globe. Although large demand exists for rideshare opportunities to Sun-synchronous orbits, a rideshare manifest dominated by them runs the risk of neglecting other orbits necessary to construct constellations that provide truly global coverage. The long-term feasibility of rideshare-initiated constellations to provide flexible, resilient, and global solutions depends on the availability of a set of diverse orbits.

## References

- <sup>1</sup> Nugent, R., Munakata, R., Chin, A., Coelho, R., and Puig-Suari, J., “The CubeSat: The Picosatellite Standard for Research and Education,” in *AIAA SPACE 2008 Conference and Exposition*, San Diego, CA, 2008, AIAA Paper 2008-7734.
- <sup>2</sup> Swartwout, M., “The Long-Threatened Flood of University-Class Spacecraft (and CubeSats) Has Come: Analyzing the Numbers,” in *AIAA/USU Conference on Small Satellites*, Logan, UT, 2013, SSC13-IX-1.
- <sup>3</sup> Spaceflight, Inc., <http://www.spaceflightservices.com>, accessed 5 March 2015.
- <sup>4</sup> Pournelle, R., and Johnson, M., “Nanoracks CubeSat Deployment Services,” in *CubeSat Developers’ Workshop*, Logan, UT, 2014, SSC14-WK-25.
- <sup>5</sup> Tyvak Nano-Satellite Systems LLC, <http://www.tyvak.com>, accessed 5 March 2015.
- <sup>6</sup> Blackwell, W., et al., “MicroMAS: A First Step Towards a Nanosatellite Constellation for Global Storm Observation,” in *AIAA/USU Conference on Small Satellites*, Logan, UT, 2013, SSC13-XI-1.
- <sup>7</sup> Wertz, J., Van Allen, R., and Barclay, T., “NanoEye—Military Relevant Surveillance for Less Than \$5 Million Total Recurring Mission Cost,” in *8th Responsive Space Conference*, Los Angeles, CA, 2010, RS8-2010-1008.
- <sup>8</sup> Sorensen, T., Nunes, M., Chan, J., Pilger, E., and Yoneshige, L., “New Gateway to Space for Small Satellites by the Hawaii Space Flight Laboratory,” in *The 4S Symposium 2014*, 2014.
- <sup>9</sup> Foust, J., “Virgin Galactic Opens LauncherOne Facility in Long Beach,” *Space News*, 12 February 2015.
- <sup>10</sup> Juang, J.-C., Tsai, Y.-F., and Chu, C.-H., “On Constellation Design of Multi-GNSS Radio Occultation Mission,” *Acta Astronautica*, Vol. 82, Issue 1, 2013, pp. 88–94.
- <sup>11</sup> Silva Curiel, A., Lambert, M., Liddle, D., Sweeting, M., Chu, C.-H., Fong, C.-J., and Chang, G.-S., “Introduction to FORMOSAT-7/COSMIC-2 Mission,” in *AIAA/USU Conference on Small Satellites*, Logan, UT, 2013, SSC13-VI-1.
- <sup>12</sup> Gangestad, J.W., Hardy, B.S., and Hinkley, D.A., “Operations, Orbit Determination, and Formation Control of the AeroCube-4 CubeSats,” in *AIAA/USU Conference on Small Satellites*, Logan, UT, 2013, SSC13-X-4.
- <sup>13</sup> Vallado, D., “Fundamentals of Astrodynamics and Applications,” 2nd Edition, Microcosm Press, El Segundo, CA, 2001.
- <sup>14</sup> Wertz, J., “Orbit and Constellation Design and Management,” Renewed Edition, Microcosm Press, El Segundo, CA, 2009.
- <sup>15</sup> Ferringer, M., and Spencer, D., “Satellite Constellation Design Tradeoffs Using Multiple-Objective Evolutionary Computation,” *Journal of Spacecraft and Rockets*, Vol. 43, No. 6, 2006, pp. 1404–1411.
- <sup>16</sup> Ferringer, M., Clifton, R., and Thompson, T., “Efficient and Accurate Evolutionary Multi-Objective Optimization Paradigms for Satellite Constellation Design,” *Journal of Spacecraft and Rockets*, Vol. 44, No. 3, 2007, pp. 682–691.
- <sup>17</sup> Ferringer, M., DiPrinzio, M., Thompson, T., Hanifen, K., Whittecar, W., and Reed, P., “A Framework for the Discovery of Passive-Control, Minimum Energy Satellite Constellations,” in *AIAA SPACE 2014 Conference and Exposition*, San Diego, CA, 2014, AIAA Paper 2014-4158.