



Fast Methods of Coverage Evaluation for Tradespace Analysis of Constellations

Vinay Ravindra , *Member, IEEE*, and Sreeja Nag , *Member, IEEE*

Abstract—Simulation of distributed space missions (DSMs) for the purpose of phase-A mission design studies and general tradespace analysis is computationally challenging owing to the necessity of evaluating thousands of architecture options. Machine learning and evolutionary optimization methods have enabled intelligent search of the architectural tradespace of DSMs, including spacecraft and instrument design specifications. A critical computational bottleneck in evaluating architectures is the ability to rapidly simulate orbits of many DSMs with varying parameters for global earth observation and compare their coverage-related performance. When design variables include heterogeneous payload types and characteristics, orbital characteristics, areas of interest, and user constraints, the parameter space may be in thousands. In this article, we describe the difficulty of coverage calculations for narrow field of view (FOV) and conical FOV sensors, and propose a novel algorithm, called quick search and correction (QSC), to overcome it. We also propose new temporal evaluation metrics to characterize the coverage performance of DSMs, as well as a uniform random sampling technique for fast evaluation of overall performance of DSMs. Performance of the proposed methods and metrics are verified on an example Landsat-derived DSM, showing $\sim 100x$ improvement in computational speed due to the QSC algorithm and $\sim 10\text{--}250x$ due to the sampling technique.

Index Terms—Earth observing system, numerical simulation, orbital mechanics, performance evaluation, remote sensing, statistical analysis.

I. INTRODUCTION

THE MATURITY of the small satellites and launch access to space over the last two decades has made distributed space missions (DSMs) for earth observation (EO) economically practicable [1]. This utility of and interest in DSMs, has spurred efforts to develop open access, DSM architecture design tools to aid mission designers in optimizing DSM designs at early stages of project development. NASA Goddard Spaceflight Center is leading the development of one such tool, called the Tradespace Analysis Tool for Constellations (TAT-C) [2], [3]. The overall

Manuscript received May 1, 2019; revised August 1, 2019 and September 21, 2019; accepted October 22, 2019. Date of publication November 27, 2019; date of current version February 12, 2020. This work was supported in part by the NASA Earth Science Technology Office's Advanced Information Systems Technology program, via a grant awarded to NASA Goddard Spaceflight Center, to lead the development of the Tradespace Analysis Tool for Constellations and in part by the project "Generalizing Distributed Missions Design Using the TradeSpace Analysis Tool for Constellations and Machine Learning" NASA under Grant N55617K0572. (*Corresponding author: Sreeja Nag.*)

The authors are with the NASA Ames Research Center, Bay Area Environmental Research Institute, Moffet Field, CA 94035 USA (e-mail: vrgl221@gmail.com; sreeja.nag@nasa.gov).

Digital Object Identifier 10.1109/JSTARS.2019.2952531

effort is to simulate candidate DSMs over their mission period, compare their performance, and use optimization algorithms to recommend pareto-optimal DSM architectures, within user-specified constraints and requirements. Some example optimizers include genetic algorithm (GA) to minimize revisit time for low Earth orbit (LEO) constellation [4], multiple-objective GA (MOGA) to optimize over average and maximum revisit times [5], MOGA for optimizing coverage-related metrics [6]. Reference [7] compares the performance of simulated annealing and GA performance for the problem of discontinuous coverage. The algorithms used are heuristic in nature and flexible to diverse user needs and underlying nature of problem, and thus work well for design of DSM architectures. However, convergence to optimal solutions requires many (typically 1000s) candidate DSMs (candidates are referred to as "individuals" in GA taxonomy) to be evaluated [6]. This number grows rapidly with the increase in design variables, and thus the search space.

Tradespace analysis of DSM architectures for EO needs a computational module to evaluate the observational and payload performance of each architecture. A critical component of this module is the propagation of satellite states (position and time) for every satellite in every DSM architecture, and computation of coverage (events of access) of customizable ground regions by payloads on those satellites during its mission lifetime. This component is also the most time consuming, computationally expensive part of the tradespace analysis owing to the large number of options for orbital specifications, satellite numbers, and geometry, payload characteristics, as well as large number of discrete points on the ground regions and discrete time steps in numerical propagation and coverage computation. Hence, it significantly slows down heuristic search or optimization techniques applied to DSM design in early formulation.

Rectangular narrow along-track (AT) field of view (FOV) sensors are fairly common in EO, examples being pushbroom imagers such as the Thermal InfraRed Sensor (TIRS) [8], [9] on Landsat-8, Operational Land Imager [10] on Landsat-8 (30 m along track swath), Multi-Spectral Imager for Sentinel-2 (10–60 m along track swath); panchromatic and multispectral matrix imagers such as WorldView-3 satellite sensor by DigitalGlobe, Ball Global Imaging System on Quickbird (16.5 km steerable swath) [11]; synthetic aperture radars (SARs) such as C-SAR on Sentinel-1 (80 km nadir swath) [12], PALSAR-2 of ALOS-2 [13]. Conical FOV sensors too are popular in EO, examples being the precipitation profiling radar on RainCube [14], and the MicroWave Radiometers on Sentinel 3 [15]. Both narrow (rectangular) AT FOV and Conical FOV sensors require

shorter time steps during coverage calculations, causing increased computation time. In Section II, we outline a numerical simulation algorithm [called as quick search and correction (QSC)] for coverage calculations specifically geared toward narrow AT-FOV sensors and conical sensors. In Section III, we verify the proposed coverage calculation method for various test cases and explore how to achieve minimal execution time by the right selection of one of the simulation parameters. Also, in the section, small errors (missed access events) are analyzed.

Another critical aspect of tradespace analysis of DSMs is the identification of suitable performance metrics, which provide comparative utility of EO observations across DSM architectures while being computationally efficient. For applications where the primary purpose of multiple satellites is to improve temporal resolution of the observations made over regions of interest, metrics commonly selected are temporal in nature, e.g., percentage of area of interest covered (maximize) in a given time span, and revisit time (minimize) of a region of interest. Examples of DSMs where it was important to optimize the temporal resolution of observations are the NASA-funded CYGNSS mission [16] which has demonstrated wind speed measurements taken from eight simultaneous vantage points in the same orbital plane, and thereby reduction of model uncertainties and the TROPICS constellation of microwave radiometers [17], [18], also funded by NASA, which is looking to increase science value of temperature/pressure profiles especially for tropical cyclones with six spacecraft in three orbital planes.

Calculating metrics of multiple satellites in the DSM over long mission lifetimes results in increased runtimes. In Section IV, we introduce some coverage metrics that characterizes the revisit performance of DSMs where value of making observations drops off after a certain period. We also explore the usage of a uniform, random, time sampling method for orbital propagation, and coverage calculations to estimate DSM performance over the actual mission lifetime. Typically, numerical simulations propagate orbits and compute coverage for all satellites in the DSM, evaluated using fixed time steps over the entire mission lifetime, and aggregate coverage metrics across every possible observation. Instead, we hypothesize that one could get a reasonably good estimation of the aggregated metric by simulating a finite number of time period samples (with each sample having a predefined duration, selected randomly over the mission lifetime).

Finally, in Section V, we tested the proposed metrics and method on an example (hypothetical) DSM with five satellites hosting a single sensor identical to the TIRS on Landsat-8. Our proposed and performance verified method lowers the time required to compute temporal metrics of a DSM by more than a factor of 10, and thus serves as an ideal employment alongside heuristic optimization for rapid tradespace analysis. The methods proposed in Section II (QSC method) and IV (random-sampling method) maybe used independently or together for a DSM evaluation. Section VI summarizes the key findings.




Fig. 1. Illustration of the numerically simulated footprint of a sensor with a narrow, rectangular FOV at three consecutive, predefined time steps. The area within the dotted lines is the actual area accessed, while the area shaded in green is the area (made of GPs) computed as accessed during coverage calculations at each of the three time steps. For accurate access calculations, the propagation time step should be a fraction of expected the access duration over a GP at the nadir, so that every GP within the dashed lines is computed correctly as being accessed.

II. ORBIT PROPAGATION AND COVERAGE COMPUTATION FOR NARROW FOV AND CONICAL FOV SENSORS

A. Background

Orbit propagation and coverage calculations are required to compute expected satellite states (position and velocity) and access events over regions or points of interest over the mission duration. Access events (i.e., the time interval in which a payload on a satellite accesses a point of interest) are dependent on satellite states and payload FOV. Observation metrics (e.g., coverage) can be computed from this data and characterize the performance of any DSM, (ref. [19] describes this flow of calculations).

In this article, we consider orbital propagation as a two-body (earth and satellite) problem taking into consideration the J_2 perturbations due to nonspherical earth. This can be solved using analytical formulation of the change in Keplerian elements of the satellite for an input timestep [20, Ch. 6]. From the initial conditions of the state of satellite, the propagator steps through appropriately selected time steps, and at each time determines if a region (represented by discretized set of uniformly set of grid points) is under the FOV of the sensor. The orbit and coverage (O&C) module section in [3] and [21] describes the coverage calculation for an arbitrary closed geometrical FOV of a sensor within the TAT-C tool, as an example. This method of coverage calculation is hereby referred to as the “traditional” method and we compare our proposed coverage calculation method to this method. Running the O&C module for conical sensors (common in sampling radiometers) and sensors with narrow AT FOV (in pushbroom imagers) has proven to be an extremely time-consuming operation. Figs. 1 and 2 illustrate how selection of the numerical time step impacts the fidelity of access (coverage) calculations for differently shaped FOV

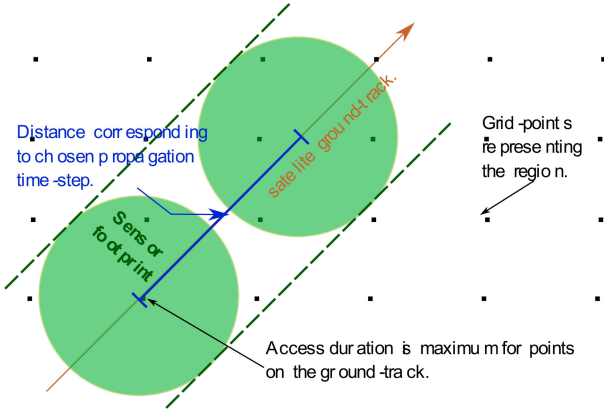


Fig. 2. Illustration of the numerically simulated footprint of a sensor with a conical FOV at two consecutive, predefined time steps. The access duration by a satellite decreases with perpendicular distance from the ground track and is nearly zero for GPs at the dotted line (flank of the FOV). Since the minimum access-duration in this case is nearly zero, the propagation time step may be set to the minimum exposure time needed by the sensor to make a valid observation.

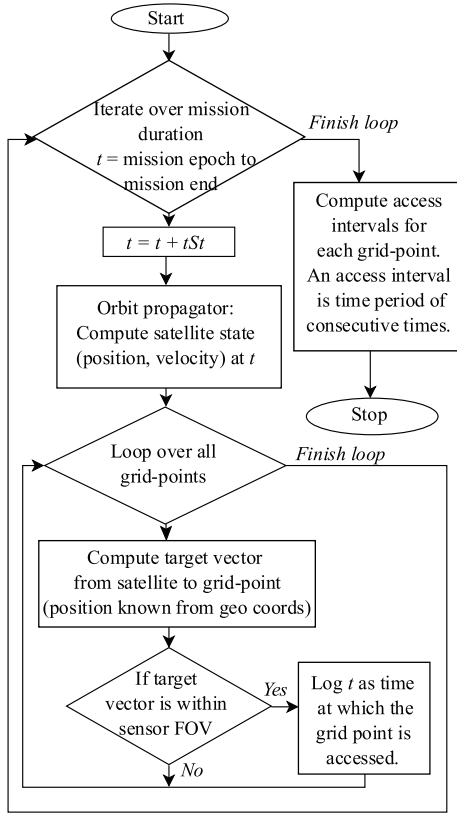


Fig. 3. Conceptual flow diagram of the numerical coverage calculations. t in the flowchart refers to the time at which the satellite state, coverage is calculated, while tSt is the propagator step size.

sensors. Fig. 3 explains the principle of the traditional O&C computations.

Access information for every grid point (GP) consists of the following fields: access event start, access event duration, and spacecraft state (position and velocity) during the access period, typically at the start and middle of access. For accurate access calculations, the propagation time step should be a fraction of

the time taken to go over the ground pixel (smaller fraction is higher fidelity but computationally more expensive). It is thus a computational burden to simulate narrow AT FOV sensors like the Landsat-8 TIRS pushbroom sensor where the AT FOV is 142 urad [8], [9] and at an altitude of 705 km, the time required to go over a ground pixel at nadir is 14.8 ms. Moreover, accurately computed access information is needed to make a fair estimate of the observational data metrics (e.g., view zenith, view azimuth, and bidirectional reflectance), which depend on the observation geometry.

B. Quick Search and Correction (QSC) Algorithm

In this subsection, we outline the proposed QSC algorithm where we divide the access calculations into a two-step calculation:

- 1) A “quick-search” step, where a proxy sensor with a larger FOV is used for orbit propagation and coverage computation, as traditionally computed within TAT-C’s O&C module for a given, fixed FOV. Since we use a large FOV, the O&C computation takes relatively less time at this step.
- 2) A “correction” step, where the actual sensor FOV is used for O&C computation over the access events recorded in the quick-search step. Here, the propagation time step is relatively small, but the computationally burden is reduced due to two factors:
 - a) The total propagation duration for each access event is small (typically in order of seconds). For example, in case of 15° AT FOV proxy-sensor at 705-km altitude used in the quick search step, the access duration at nadir is about 26 s. This becomes the total propagation duration for the “mission” corresponding to this access event in the correction step.
 - b) Each access event recorded in the quick-search step is for only one ground point, and hence the O&C computation in the correction step needs to only examine if that ground point is under the actual sensor FOV.

The selection of how large a proxy sensor FOV is to be used in the quick-search step is an important tradeoff. Selecting a too large proxy sensor FOV makes the quick-search step faster but records a larger number of access events with larger access durations. This in turns slows down the correction step. This analysis is described in Section III-A.

Fig. 4 shows the QSC algorithm for access calculations for narrow AT FOV rectangular sensors and conical sensors. It describes the selection process for the propagation time step tSt , in terms of the conditional comparison between a minimum required propagation time step (tSt_{minReq}) and a predefined minimum time step to be used in the quick-search step (tSt_{QSmin}), and thus, the proxy sensor FOV determination. The tSt_{minReq} can be determined as $f \cdot tAT_{FP}$, where $f < 1$ (f is user specified and is hereby referred to as the “overlap factor,” since it dictates the extent of overlap of consecutive simulated sensor footprints. Smaller the f , greater the overlap, and higher is the accuracy of the access information computations) and tAT_{FP} is the analytical estimated time required for the sensor footprint to go over the ground point (at nadir). tAT_{FP} can be

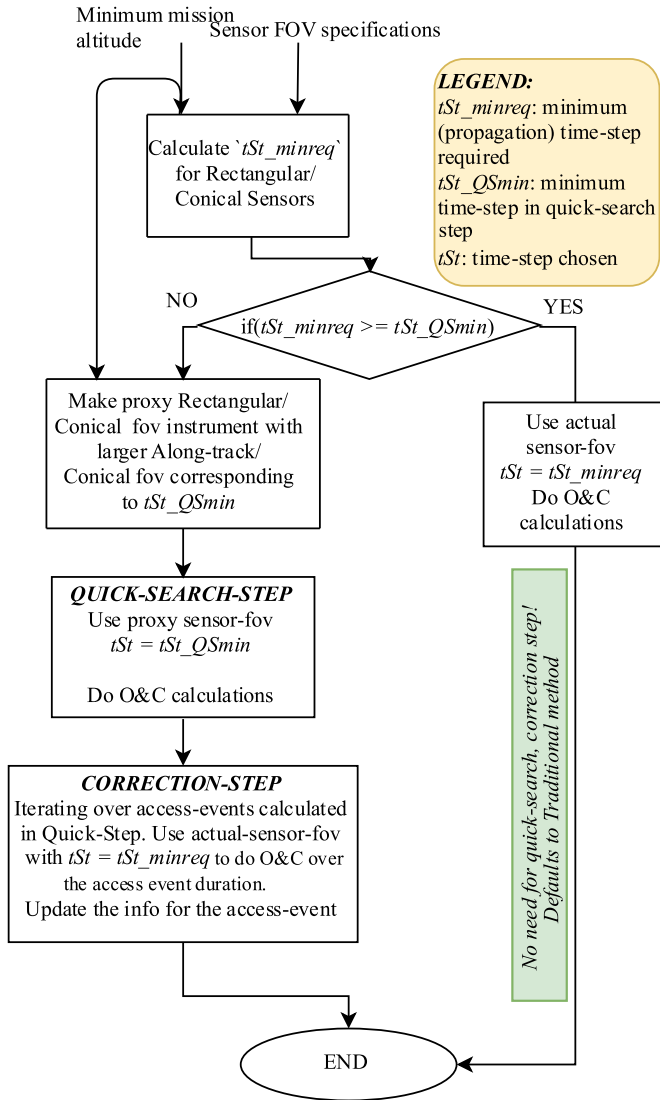


Fig. 4. Proposed quick search and correction (QSC) algorithm for access calculations for small at FOV rectangular sensors and conical sensors involving quick-search step and correction step.

estimated analytically from the earth central angle [20, Sec. 7.2]. The proxy sensor FOV is the AT FOV which gives tSt_{QSmin} as the minimum required propagation time step.

The fidelity of the coverage results produced by the QSC algorithm depends on the underlying orbit propagation model. A higher fidelity model (for example, one accounting for atmospheric drag, higher order perturbations) will improve quality of coverage data, and therefore the output of the QSC algorithm.

III. VERIFICATION OF THE QSC ALGORITHM

The QSC algorithm is demonstrated on the coverage of one satellite with Landsat-8 orbit specifications, carrying the sensor specifications of Landsat-8 TIRS over a global grid of 20000 ground points. The QSC simulation parameters used were $tSt_{QSmin} = 1$ s, $f = 0.25$ for the rectangular FOV ($15^\circ \times 142$ urad). A mission duration of 0.1 days (2.4 h) required 2853.30 s of simulation runtime when O&C was implemented traditionally

with a propagation time step of 0.0037 s, as calculated from 142 urad AT FOV of TIRS. Instead, the QSC algorithm took 29.25 s for the quick-search step, 2.73 s for the correction step, and thus a total of about 32 s of runtime. The results of the QSC algorithm matched with the results from the traditional O&C implementation. The quick-search step was implemented with a proxy sensor AT FOV 2.1960° , considering $tSt_{QSmin} = 1$ s. To summarize for this example, the QSC took \sim one-hundredth the execution time compared to a direct approach.

Several representative use cases of different sensors in different orbits over different durations were simulated to compare results from the QSC vs. the traditional method of O&C, over 2000 global GPs. Table I lists detailed parameters and execution times of the simulations. In order to succinctly show that the results from the QSC and traditional method match, the following metrics are compared in Table I: calculated number of access events and mean and variance of the calculated access durations. The use cases also verified that the access events from the QSC and the traditional method correspond to the same GPs and at the same time horizons (not shown in the table). Runtime is improved by several orders of magnitude at negligible difference of the results against the traditional method across all the metrics. As expected, the QSC algorithm causes greater improvement in computational speed over the traditional method, for smaller sensors. While the correction step in the presented cases takes < 1 s, Section III-A discusses how the overall execution time may be further minimized by selection of optimal tSt_{QSmin} parameter (and hence proxy sensor FOV for the quick-search step).

The small differences observed in the results of the traditional and the QSC method will be discussed below. The machine used for all simulations in this article was a Virtual Machine with allotted 8-GB RAM running on Intel Core i9-8950HK CPU@2.90 GHz, 32-GB RAM.

A. Performance Sensitivity to QSC Parameters

The tSt_{QSmin} parameter dictates the size of the AT FOV of the proxy sensor. While a larger tSt_{QSmin} implies larger propagation step size in the quick-search step which would reduce execution time of the quick-search step, it also forces a larger proxy sensor FOV, which covers a larger number of GPs at each propagation step (see Fig. 3). Since larger number of GPs need to be processed, this increases execution time. Fig. 5 plots the execution time of the quick search, correction steps, total execution time, and the AT FOV used in the quick-search step for different values of the user-defined parameter tSt_{QSmin} . In all cases, the orbit simulated was a 500-km SSO over a simulation period of 15 days, and the sensor considered was a rectangular FOV sensor with cross-track FOV of 45° . The QSC simulation parameter f (overlap factor) was set to 0.25. The tSt_{QSmin} parameter determines the value of the proxy-sensor FOV. While a larger tSt_{QSmin} implies larger propagation step size and lower execution time of the quick-search step, the larger proxy sensor FOV also captures a larger number of GPs, hence more processing at each propagation step and increasing execution time (see Fig. 3). In all the subfigures of Fig. 5, we see that the

TABLE I
EXAMPLE CASES TO VERIFY AND COMPARE THE TRADITIONAL AND QSC METHOD OF COVERAGE COMPUTATIONS

Method	Sub-method	FOV (deg)	tSt (s) = tSt_{minreq} (s)	# Access events	Average Access duration (s)	Standard deviation of Access duration (s)	Execution time (s)
2x6 deg FOV Rectangular Sensor in a 300km SSO circular orbit. Duration: 1 day.							
Traditional		2 x 6	0.355	82	1.04	0.101	54
QSC	Quick Search	5.63 x 6	1	83	2.88	0.422	6.5
	Correction	2 x 6	0.355	83	1.02	0.139	0.02
1x5 deg FOV Rectangular Sensor in a 500km SSO circular orbit. Duration: 15 days.							
Traditional		1 x 5	0.309	1440	0.91	0.078	1144
QSC	Quick Search	3.24 x 5	1	1455	2.91	0.403	134.2
	Correction	1 x 5	0.309	1435	0.92	0.069	0.24
15 deg Conical FOV Rectangular Sensor in a 700km SSO circular orbit. Duration: 15 days.							
Traditional		15	2.728	6847	18.2	6.2	50
QSC	Quick Search	15	2.728	6847	18.2	6.2	50
	Correction					-NA-	
3 deg Conical FOV Sensor in a 900km SSO circular orbit. Duration: 15 days.							
Traditional		3	0.727	1764	5.03	1.55	220
QSC	Quick Search	4.13	1	2359	6.94	2.16	160
	Correction	3	0.727	1762	5.05	1.56	0.3
0.001 x5 deg FOV Rectangular Sensor in a 1100km SSO circular orbit. Duration: 0.1 day.							
Traditional		0.001 x 5	7.708e-3	19	2.25e-3	~0	1887
QSC	Quick Search	1.297 x 5	1	19	3.00	~0	1.42
	Correction	0.001 x 5	7.708e-3	19	2.29e-3	~0	0.24

In all the listed cases a uniform grid of 2000 points over the entire Earth is considered. The QSC simulation parameters were: $tSt_{QScmin} = 1$ s, $f = 0.25$ for rectangular FOV, and $f = 0.1$ for Conical FOV. The average and standard deviation are calculated over all the recorded access events of the simulation run. The performance metric column headers are bolded. The row titled "correction" under QSC lists the performance metrics of the quick step *and* correction to compared to the Traditional row. However, it lists the execution time of *only* the correction step to allow readers to understand the time budget.

execution time of the quick-search step decreases rapidly with initial increases of tSt_{QScmin} , but gradually flattens out, due to the varying degree of conflict between the two factors described above, at different regimes of tSt_{QScmin} . On the other hand, execution time of the correction step has an increasing linear trend (with small local deviations that are explained below in Section III-B). This is because as the tSt_{QScmin} increases, the proxy-sensor FOV increases, and the number of accesses and the access duration of each of the access events by the proxy-sensor increases. Thus, a relatively larger number of access events needs to be processed in the correction stage. The minimum total execution time is the region where the execution times of the quick-search step and correction steps are similar.

In subgraphs (a) and (b), the number of GPs considered in coverage calculations is varied (10 and 5K, respectively) for the same sensor rectangular FOV of $1 \times 45^\circ$, to show the effect of greater number of GPs at each propagation step. The same behavior is expected by varying of the altitude of the orbit or the cross-track FOV of the sensor (larger altitude/cross-track FOV implies more GPs captured at each propagation step). Subgraph Fig. 5(b), as expected, has shorter execution times since it corresponds to smaller number of GPs. The position of the minimum execution time is roughly the same in both cases ($tSt_{QScmin} = 16 \sim 17$ s) suggesting that increased number of GPs affects (increases) the execution times of both the quick-search step and correction step to the same degree.

Comparing subgraphs (b) and (c), show the effect of the AT FOV of the sensor, 1° and 0.1° , respectively. The rate of increase of the correction step execution time is more for 0.1°

case because the correction step has a smaller propagation step size (tSt), due to the smaller associated AT FOV. The minimum execution time in this case is for $tSt_{QScmin} = 5$ s and can be said to be sensitive to AT FOV of the (actual) sensor.

An optimal (in terms of minimizing the overall execution time) tSt_{QScmin} may be chosen by the user by conducting an analysis as described above. Such analysis can be for a short mission duration, such as a few days, just to determine the optimal tSt_{QScmin} , which can then be used for calculating coverage over the entire mission duration.

B. Error Analysis

Minor differences in coverage performance were seen in the QSC results compared to results using the traditional method. The calculated access duration and the access event start date by the QSC method can deviate upto the calculated minimum required propagation step size used in the correction-step ($\pm tSt_{minreq}$). This error can be seen in Table I, where the average and the standard deviation of the access durations of the traditional method and the QSC method are slightly different, since the resolution of the orbit propagation is tSt_{minreq} . The user can control this parameter by specifying a different value of the overlap factor f . A lower overlap factor f corresponds to smaller propagation step size used in the correction step of the QSC algorithm and hence enhanced time resolution. Limited time resolution may also cause access events to be captured in one method (QSC or the traditional) but not in the other method. These events correspond to those with access durations lesser

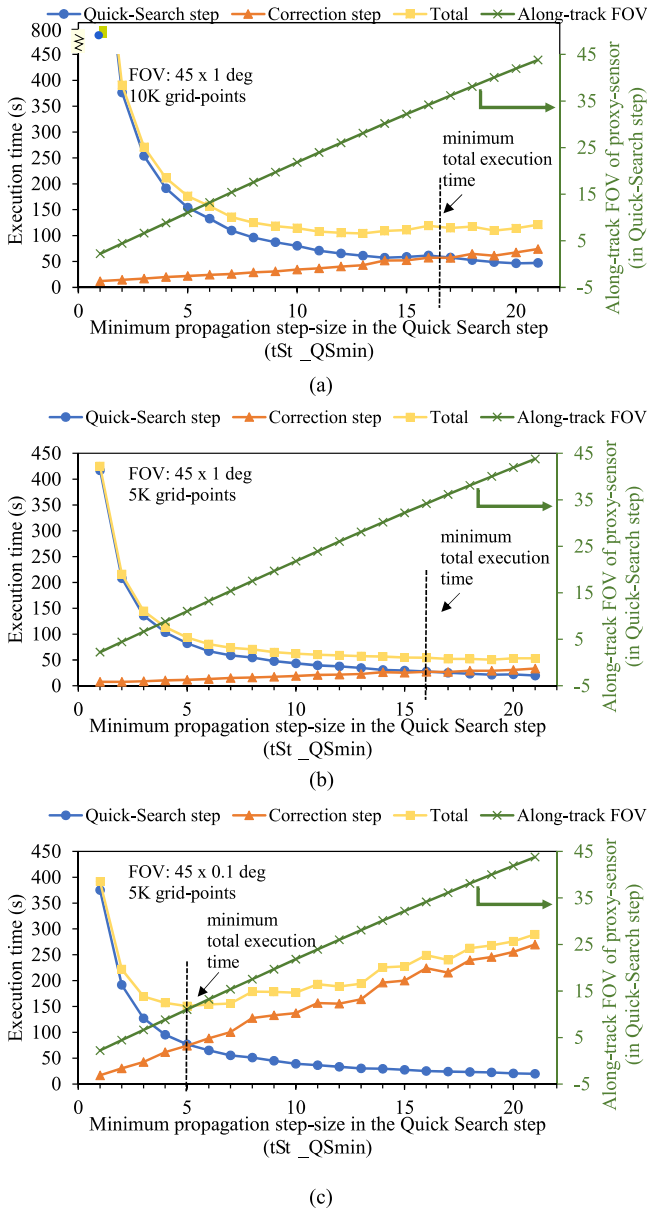


Fig. 5. Sensitivity of the execution time to the tSt_QSmin parameter for a fixed factor $f = 0.25$. In all cases, the orbit simulated was a 500-km SSO with a rectangular FOV sensor of 45° cross-track FOV, over a period of 15 days. Comparing (a) and (b), varies grid point number. Comparing (b) and (c), varies the along-track FOV of the sensor. The tSt_QSmin corresponding to minimum total execution time is seen to be most sensitive to the along-track FOV of the sensor. (a) FOV: $45 \times 1^\circ$, 10 K grid points globally. (b) FOV: $45 \times 1^\circ$, 5 K grid points globally. (c) FOV: $45 \times 0.1^\circ$, 5 K grid points globally.

than one tSt_minreq . Capturing of such events is dependent on the epoch at which the numerical orbit propagation started, and not based on the method used.

In some cases of Table I, the number of accesses (see “correction” row) is slightly different as calculated by the calculated by the QSC vs. traditional method. While some of the disparity corresponds to access events with duration less than the time resolution of orbit propagation, there are missing access events (in the QSC method) with duration longer than the propagation time step. Further, Fig. 5(c) shows small visible deviations (at the

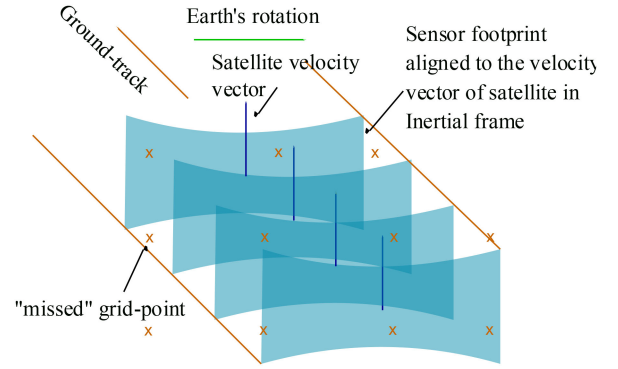


Fig. 6. Exaggerated illustration of coverage error due to use of large propagation step sizes. An orbit at 90° inclination is shown, with the satellite moving from bottom to top of the page. In the coverage calculations, the sensor is aligned in a Nadir-pointing frame, with the y-axis along the negative orbit normal, defined by the satellite position vector and the satellite velocity vector in inertial frame. Due to Earth’s rotation, there are regions which the numerical simulated sensor footprints do not cover, even though they may have a reasonably large longitudinal overlap.

larger proxy sensor FOV regime) from an otherwise monotonic increase in correction step execution times. An increasing proxy sensor FOV implies a greater number of access events collected during the quick-search step to be processed in the correction step. However, the number of access events (of access durations larger than the minimum propagation step size) calculated by the correction step *decreases* with increasing proxy sensor FOV used in the quick-search step, leading to dips in the otherwise monotonic increase.

We hypothesize that the decrease number of access events, is due to discretization errors in the coverage area calculations resulting from the use of a fixed step size by the numerical orbit propagator. In the quick-search step, larger the proxy sensor AT FOV, bigger is the corresponding propagation time step used. While this serves to decrease the execution time, it introduces discretization errors as illustrated in Fig. 6. The coverage is typically calculated with the sensor aligned to the nadir-pointing frame, and we assume one of the sensor axes to be aligned to the orbit normal, which is derived from the satellite velocity vector in inertial frame. There is a nonzero angle between the ground track and the satellite velocity vector in inertial frame due to the rotation of the earth, which manifests into missed access events (missed GPs).

We test this hypothesis in two ways as follows:

- 1) Evaluate the errors (difference in number of accesses compared to the traditional method) for orbits of different inclinations and different proxy sensor FOVs used in quick-search step [see Fig. 7(a)]. The error is zero for equatorial orbits, which is consistent with our hypothesis, because the satellite velocity vector in inertial frame is aligned with the ground track in an equatorial orbit (the satellite revolves around the same axis as the earth’s rotational axis). Moreover, the errors for nonzero inclination orbits increase with the increase in the proxy sensor FOV used in the quick-search step.
- 2) Decrease the QSC overlap factor f to 0.20 (and have corresponding changes in the tSt_QSmin to maintain the

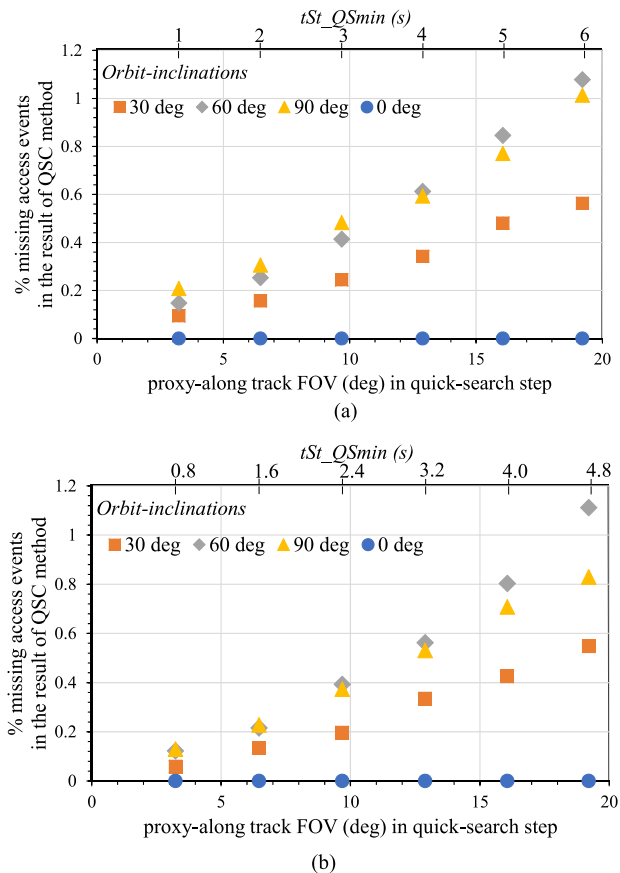


Fig. 7. Percentage of missing access events recorded by the QSC method (baselined to the number of access events calculated by the traditional method) for cases of $f = 0.25$ and $f = 0.20$. The satellite orbit simulated was 500-km SSO with $20 \times 20^\circ$ rectangular FOV sensor, with 8000 global grid points. The error increases with increase in the proxy-sensor FOV used in the quick-search step and decreases with larger overlap between the sensor footprints [Case (b)]. It is zero for the case of equatorial orbits. (a) Case: $f = 0.25$, $tSt_QSmin = 1x$. (b) Case: $f = 0.20$, $tSt_QSmin = 0.8x$.

same set of proxy sensor FOVs) and reevaluate the errors [compare Fig. 7(a) and (b)]. With decreasing f and thus increasing sensor footprint overlap for the given set of proxy sensor FOVs, the propagation time steps used in the quick-search steps is smaller and should result in smaller discretization error to be consistent with our hypothesis. There is indeed an observed decrease in the error relative to the case of Fig. 7(a), where $f = 0.25$ was used.

The maximum disparity of recorded events is $\sim 1\%$. Their occurrence depends on the density of the grid, the placement of the grid relative to the orbit path, the satellite state at which the simulation is started, orbit inclination, and sensor alignment for coverage calculations (e.g., some coverage calculations may be done with frames constructed using satellite ground-velocity vector, geodetic frames, etc.).

IV. PROPOSED METRICS AND RANDOM SAMPLING FOR RAPID DSM EVALUATION

One of the key advantages that DSMs have over single satellites is a better temporal resolution of observations made over regions [1]. Ref. [20, Sec. 7.2.3] describes standard figures

of merit for coverage, to quantify the temporal resolution of observations possible by DSMs. Much of the objectives used in DSM optimization are revisit time or coverage gaps [4]–[7]. These metrics are often averaged to give a “mean” quantity over several events recorded over the entire mission duration of the DSM (e.g., mean coverage gap). Further, since these averaged metrics are computed for each ground point, they may be further aggregated over several ground points in a region.

This section proposes an updated coverage metric (useful revisit time) for quantifying a DSM’s temporal characteristics, and briefly outlines some instantaneous observation metrics important to earth science products. The section also describes a novel sampling method for rapid computation of metrics over space and time. The example simulations in Section V show the applicability and computational efficiency of the proposed sampling method on the proposed metrics.

A. Useful Revisit Time

For DSMs in the LEO, continuous coverage globally is not possible without a large number of satellites in the DSM (e.g., Iridium commercial constellation of 66 communication satellites). The quantification of “revisit” of any global point therefore becomes important for DSMs which cannot afford a large number of satellites. Our proposition of a new metric called “useful revisit time” is motivated by effectively quantifying only those revisits who observational data is useful to the user, especially in applications of rapid response (e.g., wild fires), events or phenomena being of finite time duration (e.g., transient precipitation and flash floods), low latency satellite data supplementing slower ground based instruments (e.g., aircraft or ship tracking). In such applications, a DSM with revisit rate longer than the phenomena duration or data expiry does not serve any use.

A useful-revisit event is defined as an event when a DSM revisits and *can* make a successful observation of a region within a user-specified maximum useful revisit period. For example, for an application involving urban floods due to heavy precipitation on small streams in the Atlanta area, the user may specify the maximum useful revisit period as 1 day. Thus, events over the region of interest with revisit < 1 day qualify as a useful revisit event, and the period of the revisit (also called the coverage gap) is the useful revisit period. This framework also allows us to define a “useful visit” (what qualifies as a successful observation?) by setting a threshold level on the expected instantaneous data metrics, some of which are defined in Section IV-B. For example, we may define a minimum signal-to-noise ratio (SNR) of an image taken by an optical sensor, such that visit of the satellite at local dawn, dusk or night would not qualify as a useful visit result. By setting the maximum useful revisit period to the entire mission simulation period and clearing all the thresholds on data metrics will make useful revisits equivalent to the traditionally calculated revisit period.

This event can be quantified by the following aggregating metrics.

1) *Aggregate Value of Useful Revisit Periods*: This metric aggregates the coverage gaps between useful revisit events. The

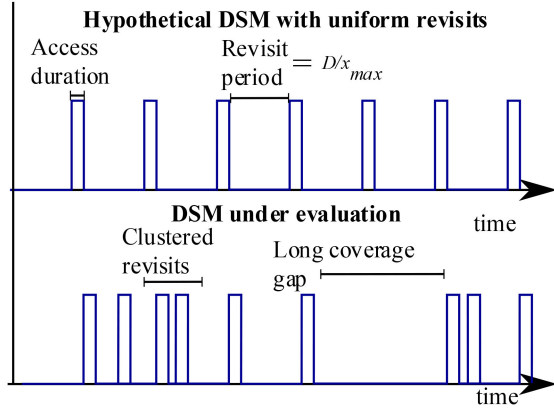


Fig. 8. Comparison of the occurrence of observation events by a hypothetical ideal DSM, defined by a user-derived minimum number of observations, against an example DSM whose performance is being quantified. Performance is evaluated not only by number of observations, but also its spread.

statistical aggregator can be the median, a percentile (e.g., 90%-ile or upper quartile), or a standard mean period over all the useful revisit events during the mission, represented as

$$\bar{x} = \sum_{i=1}^N x_i / N \quad (1)$$

where

- \bar{x} is the mean of the useful revisit periods;
- x_i is the period of the i th useful revisit;
- N is the total number of useful revisits.

2) *Normalized Number of Useful Revisit Events*: We define a baseline hypothetical DSM which has the following number of useful revisits, equally spread over mission life. This represents the minimum number of observations that a DSM would need to make to meet the user's useful revisit criteria.

$$M = D / x_{\max} \quad (2)$$

where

- M is the number of useful revisits of any region or point of interest by the hypothetical DSM;
- D is the total mission duration;
- x_{\max} is the user specified maximum useful revisit period.

The normalized number of useful revisits is defined as

$$r = N / M \quad (3)$$

where

N is the number of useful revisits of any region or point of interest by the DSM whose performance is being quantified.

While it is trivial to see that $0 \leq r < \infty$, $r > 1$ does not necessarily meet the revisit requirements of the user because the spread of observations is as important as the number. Fig. 8 illustrates a situation where a DSM has $r > 1$, but since some revisits clustered around a certain period (which may occur in a DSM with "string-of-pearls" architecture), others are further apart than x_{\max} , thus useful revisit is not met for all times of the mission.

Ref. [18] introduces a similar aggregate metric called Continuous High Revisit Coverage which is the percentage of time where a GP is either in an access, or in a gap shorter than a threshold gap duration. The normalized number of useful revisits metric, on the other hand, considers the number of imaging opportunities, and not the time available for any opportunity. The access duration for a GP (which limits the observation time, or exposure time available to the sensor) can be reported as an instantaneous observation metric in our proposed framework.

3) *Variance of Useful Revisit Periods*: Ideally, a mission designer may like the revisits by the DSM to be spread uniformly over the entire mission duration. The uniform distribution of revisits in the hypothetical baseline DSM is motivated by this ideal. The variance of the revisit period for the ideal DSM is zero. We can get a sense of the distribution of the revisits, for any DSM being evaluated, by calculating the variance of the useful revisit period as follows:

$$v = \sum_{i=1}^N (x_i - \bar{x})^2 / N \quad (4)$$

If a user wants to avoid clustered useful revisits, DSMs with lower variance of the useful revisit period may be selected.

B. Instantaneous Observation Metrics

While coverage and revisit quantify the overall temporal performance of a DSM, instantaneous metrics are very useful for evaluating sensor specific performance as time series. For tradespace analysis purposes, we propose two standard observation metrics: SNR ratio and the Noise-Equivalent Delta Temperature (NEDT) for radiometric performance determination in optical/near-optical sensors. Metric such as the noise-equivalent sigma zero can be used for quantifying the radiometric performance of SARs. Irrespective of the payload sensor, observation geometry of the satellite and Sun with respect to the observed ground point are instantaneous parameters that serve as critical inputs to the above metrics. Observational geometry parameters [22] are not just important for standard metrics like SNR, but also for more specific data product driven metrics that are dependent on the spectral characteristics and type of sensor, e.g., bidirectional reflectance distribution functions (BRDF), leaf area index (LAI), normalized difference vegetation index (NDVI). The ability to rapidly compute observation geometry time series enables dependent data products to be computed, and DSM trades can be analyzed based on higher fidelity, science product objectives.

In Section V, where we have conducted a simulation case study, we use the following instantaneous observation metrics: range, observation-zenith angle, SNR, and NEDT. A brief description of the evaluation of these standard metrics is as follows. Range is defined as the distance from the satellite to the target ground point at the middle of any observation period. The observation-zenith angle is defined as the angle between the vector from satellite to ground point and the nadir vector. The SNR is calculated from the framework of passive-optical-sensors given in [20, Chap. 9]. The framework is described briefly here: earth is modeled as a blackbody radiator at a temperature of 290 K,

while the sun is modeled with temperature 6000 K. The angle of the sun to the local frame at the pixel (centered at the ground point of dimensions equal to the spatial resolutions in the AT, cross-track directions) and angle to the satellite is calculated from the computed access data. Radiance at the sensor aperture is taken as the sum of the radiance radiated from the earth and radiance of sun reflected of earth integrated over the imaging band. A unity surface albedo is assumed. Integration time at the sensor aperture is set to be the minimum of the sensor hardware specification or the access duration. Efficiency of the optical transmission system and detector efficiency of converting the incident photons to circuit electrons is considered and the final signal electrons at the sensor electronics is estimated. Shot noise model is considered to calculate the number of noise electrons, and finally the ratio of the signal electrons to the noise electrons is presented as the SNR of any observation. The NEDT metric (used for thermal sensors such as Landsat-8's TIRS) is calculated as the ratio of number of noise electrons to the change in number of signal electrons for 1K raise in scene temperature. A lower NEDT corresponds to a higher quality observation. Depending on the user application, process-driven computation for BRDF, LAI, NDVI, or any other observation metric can be implemented similarly.

C. Uniform Random Sampling for Rapid DSM Evaluation

The traditional way of quantifying a DSM's performance has been to numerically simulate orbits of all the satellites in each DSM architecture, compute all the access events, compute coverage and instantaneous observation metrics from the access data per ground point (GP)—hereon referred to as the level-0 metrics. After that, level-0 metrics of all the events at any ground point are aggregated over time using a statistical measure like mean, median, or variance—hereon referred to as level-1 metrics. These level-1 metrics at each GP can again be aggregated over all GPs to get a performance metric for a region or globe—hereon referred to as level-2 metrics. The evaluated level-2 metrics of any DSM can be compared with that of other DSMs to determine tradeoffs in performance, as compared to tradeoffs in other evaluation criteria such as cost and risk. Note that the above introduced level-0/1/2 metrics are not to be confused with the level-1/2/3 terms used in taxonomy of remote-sensing data products at different stages of their processing pipeline. Our proposed level-0 simply indicates spatio-temporally varying metrics, which when aggregated over time produces level-1 metrics, and when further aggregated over space produces level-2 metrics.

In this article, we explore the possibility of quantifying the performance of a DSM by conducting numerical simulations of the O&C at randomly chosen small intervals within the mission duration, instead of over the entire mission duration. The randomly chosen small-intervals are hereon referred to as “samples” and the duration of each of these temporal samples are referred to as the “sample duration”. During each sample computation, we compute one or several level-0 measures of the coverage, data metrics over all the GPs representing the region of interest. We hypothesize that a small number of samples is sufficient to determine the level-1 mean metric at a GP, rather

than computing the mean over all the observations at that GP during the entire mission period. We can thereby estimate the performance of a DSM by making the O&C calculations for a fraction of the mission duration. A uniform random sampling method is chosen so that we can capture the potential changes in the level-0 metrics at different periods of the mission. The nature of orbital dynamics involves several periodic phenomena such as the orbit of satellite, rotation of Earth, etc. and a random sampling strategy is required to avoid bias.

There are two parameters to be decided upon for the random sampling strategy.

1) *Sample Duration*: The sample duration should be such that all the level-0 metrics are measurable within the duration. Increasing the sample duration guarantees sufficiency, but at the cost of increased computational load. Therefore, a threshold maximum level must be decided depending on the metrics of interest to the user in a given application, i.e., maximum across all requirements. Level-0 coverage metrics are likely to require the largest sample duration, since the revisit time/coverage gap is expected to be large. If the application is interested in a level-0 metric such as useful revisit period, the sample duration can be set to the user specified maximum useful revisit period (x_{\max}). Instantaneous metrics typically require smaller sample duration. For example, if the expected access duration is 25 s, the sample duration must be greater than 25 s so that the level-0 observation metrics can be computed at-least once during the sample.

2) *Number of Samples*: Increasing the number of samples randomly selected from the mission duration increases the computation load. On the other hand, selecting too small number of samples can lead to collection of less than adequate number of level-0 metrics at a GP, and hence wrong estimation of the level-1 metric. Therefore, a threshold level must be decided. Note that the number of samples is not necessarily the same as number of level-0 metrics collected at a GP. It may happen that during one sample run, multiple, or zero number of measures of the level-0 metric are made at a GP.

V. VERIFICATION OF THE RANDOM SAMPLING METHOD

We apply our proposed sampling method with our proposed metrics to the simulated spacecraft and sensor described in Section III's first paragraph, to demonstrate comparable fidelity of results at orders of magnitude better computational efficiency. While in our example the QSC algorithm is applied for O&C calculations, the random sampling method can also be applied with any other coverage calculation method. The example DSM simulated in this section comprises of five satellites in a uniform Walker constellation [20, Sec. 7.6] with the specifications of the Landsat-8 TIRS pushbroom sensor. One of the satellites in the DSM is simulated to be in the same orbit as Landsat-8. The mission duration is assumed at 180 days and the user-specified useful revisit period is 7.5 h. The DSM performance metrics computed in this example are range, observation-zenith angle, SNR and useful revisits.

A. Baseline Simulation for the Control Experiment

The example DSM was simulated, and metrics computed in the traditional manner over the entire mission duration of

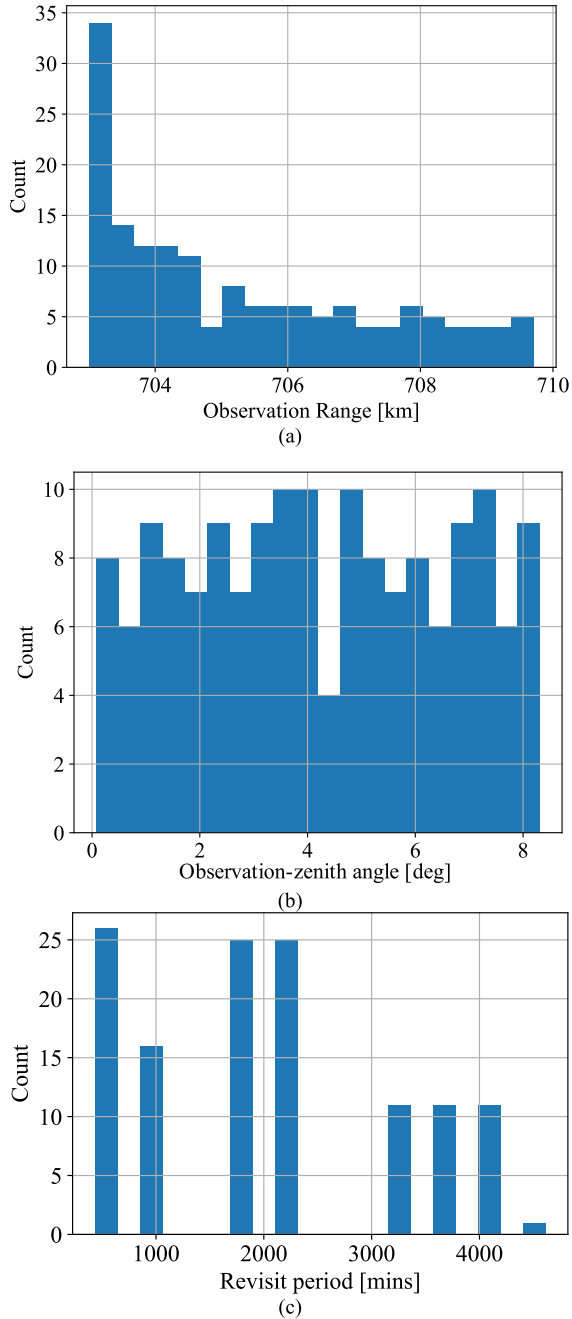


Fig. 9. Frequency plots showing some level-0 metrics at GP (40.00° N, 98.01° W): observation range, observation-zenith angle and revisit period computed for the example five satellite Walker DSM with the Landsat-8 TIRS sensor, over 180 days of mission duration (baseline simulation without subsampling the time horizon described in Section V.A). In (c), all the revisit periods are shown.

180 days, to represent the baseline dataset. Fig. 9 shows the frequency plots of level-0 metrics range, observation-zenith angle, and revisit period. All the revisits are shown in the plot, including those longer than the useful revisit duration. The key aspect to note is that all frequency distributions have finite variance, and from the Central Limit Theorem, we can obtain a measure of the error expected when we calculate the level-1 mean metrics by taking uniform random samples from these distributions. The standard error (s_E) associated with taking k samples from a

TABLE II
 s_E FOR LEVEL-0 METRICS FOR VARIOUS NUMBER OF SAMPLES AT GRID POINT (40.00° N, 98.01° W)

# Samples	s_E of range (km ²)	s_E of observation-zenith angle [deg ²]	s_E of useful-revisit period [min ²]
1	4.1	5.66	14.9
5	1.8	2.53	6.7
10	1.3	1.79	4.7
20	0.9	1.27	3.3

population with variance σ^2 is given by [23]

$$s_E = \sigma^2 / \sqrt{k}. \quad (5)$$

Table II lists the standard error s_E expected for the level-0 metrics, as calculated from the numerical variance of the full data set of distributions, as a function of different number of temporal samples drawn. All distributions are obtained at GP 40.00° N, 98.01° W and all samples are 10 orbital periods long.

B. Uniform Random Sampling Method

The uniform random sampling method described in the previous section is applied to the example case of the five satellite Walker DSM with Landsat-8 TIRS sensor. The sample duration is chosen as 10 orbital periods ~ 16.5 h. This is twice greater than the user-specified maximum useful revisit of 7.5 h, and hence capable of potentially capturing two revisits within the sample. A threshold NEDT = 0.4 K (corresponding to the TIRS requirements [8]) is defined as the maximum NEDT for an observation to be deemed useful during a visit.

Fig. 10 shows the computed level-1 metrics: mean of observation range, mean of observation-zenith angle, and mean of useful revisits for different number of samples and at 10 randomly chosen GPs within the global grid set (20 000 points). As the sample size S_N increases, the estimated level-1 metrics for any GP gets closer to the level-1 metric, as evaluated from the complete simulation dataset (control experiment). The variance (Section IV.A-4) of the useful revisit periods is nearly 0 for all cases due to the uniform nature of the chosen example DSM.

The aggregated error of a level-1 metric is defined as

$$\sum_{i=1}^P \left| |m_{l1,i} - m_{l1,i}^B| \right| / P \quad (6)$$

where

$m_{l1,i}$ is a level-1 metric (e.g., mean observation range) at GP i , calculated using the proposed sampling method;

$m_{l1,i}^B$ is the level-1 metric at GP i , calculated from the baseline simulation, serving as the control experiment;

P is the total number of GPs observed using the corresponding uniform random sampling method. Note that the error calculation in (5) does not consider GPs which are not observed by the random sampling technique, although they may have been observed sometime during the mission.

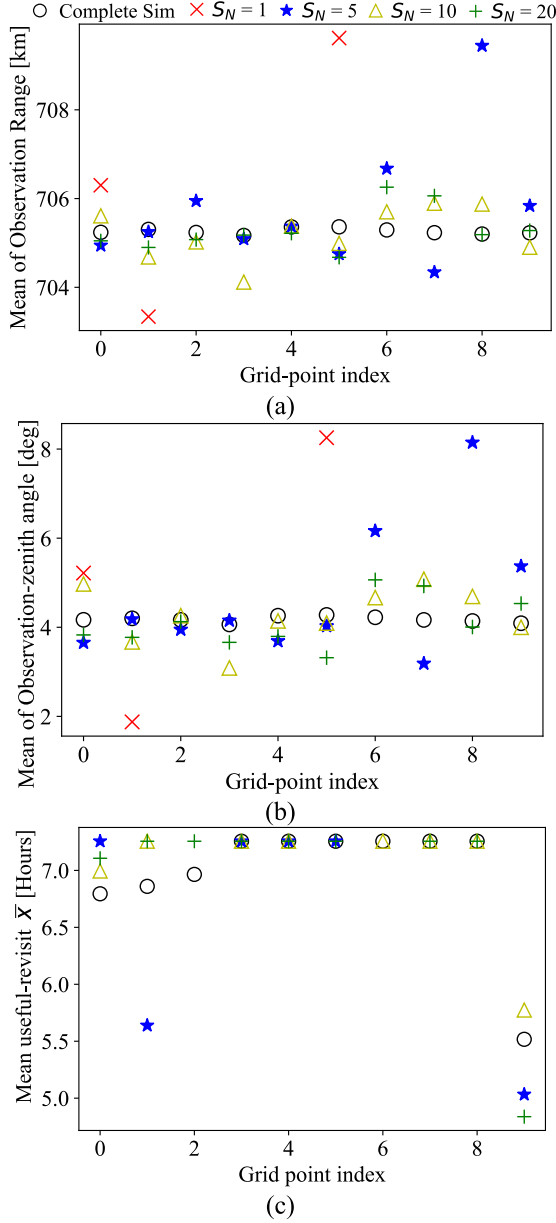


Fig. 10. Level-1 metrics of the five satellite Walker DSM with Landsat-8 TIRS at ten randomly chosen GPs, among the 20000 GPs populating the globe. S_N is the number of temporal samples used for the random uniform sampling strategy. The level-1 metrics calculated the complete simulation of the DSM is plotted as circles to provide a baseline dataset to validate the random sampling strategy. Higher number of samples increases the accuracy of the evaluated level-1 metric as expected. Too few samples (such as $S_N = 5$) may lead to missing metrics at some GPs.

Fig. 11 shows these error terms in absolute units for four metrics (normalized number of useful revisit events, range, observation zenith angle, useful revisit period), as aggregated over all the GPs over the globe, for the uniform random sampling method as a function of sample sizes. Increasing the sample size reduces the aggregate error in all metrics, albeit nonlinearly and to differing extents. The errors are calculated using the evaluated metric value of the baseline simulation as the reference.

The complete, baseline simulation of 180 mission days took a runtime of 28 h for the O&C employing the QSC algorithm to complete on the previously described computer and OS




Fig. 11. Aggregated level-1 metric errors over all 20000 GPs as a function of the number of temporal samples in the uniform random sampling technique, as baselined against metrics calculated from the complete numerical simulation dataset. There are four different y-axes for the four plot lines, as matched by the color of the axis label, tick marks, and plot line.

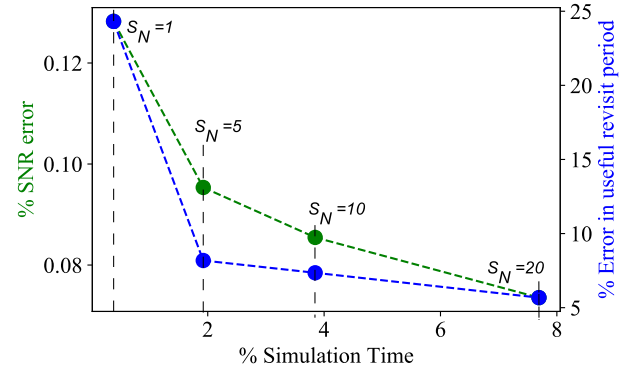


Fig. 12. Percentage SNR error, percentage useful revisit period error versus percentage of the simulation time (referenced to the baseline, complete sim). The % SNR error is observed to be low even in case of $S_N = 1$ because the variation of range and observation-zenith angle across all observations made by the DSM (Fig. 9) is small. The useful revisit error is comparatively higher.

(Sections II and III). The random sampling technique successfully *reduced* computational load in inverse proportion to the number of samples. The total simulation length is the product of the number of samples and sample duration, therefore reducing either term reduces the runtime proportionately. In the verification case, the reduction in simulation time is a factor of 262, 52, 26, and 13 for randomly picked sample size of $S_N = 1, 5, 10, 20$ cases, respectively. If runtime is assumed to be proportionate to simulation time, the expected, sequential runtime (for the O&C employing the QSC algorithm) for the four #samples buckets in Fig. 11 are 6.4, 32.3, 64.6, and 129.2 min, respectively, while the actual runtimes recorded were: 6.7, 32.7, 66.5, and 143.1 min, respectively (running on machine with specifications as listed in Section II).

The accuracy of results from the proposed methods is sensitive to the metrics. Fig. 12 shows the relative SNR error and relative useful revisit period error (baselined against the metric values computed in the complete simulation) vs. the simulation time required, as a percentage of the time taken by the complete simulation. The percentage SNR error is low even in the case

of $S_N = 1$, owing to the small variance in range at any given ground point [see Fig. 9(a)]. SNR primarily varies with range, and thus the variance in SNR too is low. The small variance in range is further attributed to the relatively small off-nadir ($\pm 7.5^\circ$ cross-track FOV [8], [9]) allowed for imaging. The trend in both Figs. 11 and 12 shows decreasing error with increased sample size (and therefore, simulation time) as expected.

C. Practicable Implementation of the Proposed Method

The proposed random sampling method was demonstrated for the case of DSM with five satellites in uniform Walker constellation. DSMs can vary significantly in time-varying topology and therefore metrics as a function of their type (constellations, formations, etc.), number of satellites, orbits, geometry, etc. The random sampling method has not been verified for all possible use cases. We anticipate the following issues (and suggest workarounds) which may arise during practical implementation of the method:

1) *Selection of Number of Samples S_N* : During any sample simulation, only a subset of all GPs of interest maybe “seen” by any satellite in the DSM. It is difficult to predict when a GP will be accessible in a sample. In an extreme case, there may be some GPs in the user’s region of interest that are never “seen” by the DSM. To circumvent this unpredictability, we propose an algorithm that adaptively sets the number of samples. First, it can process an initial fixed number of samples (say S_{N0}), and later, it forms a set of all the GPs seen at least once during the S_{N0} sample collection. The algorithm iteratively continues to sample until the required number of metrics per GP and over all of the GPs is acquired.

2) *Orbit Propagator*: In our presented demonstration, we considered a simple orbit-propagation model with perturbations due to J2 effect on the argument of perigee and right ascension of ascending node [20, Sec. 6.2.2]. Both these perturbations remain constant throughout the mission; the altitude, inclination, and eccentricity do not perturb. This allows for propagator to make large “time-jumps” between periods of propagation without any propagation errors. This may not be true for a more sophisticated propagation model and errors due to large time-jumps may need to be considered. On other hand, if the satellites in DSM are equipped with station keeping abilities, we can assume that the orbits are corrected frequently, and it may be reasonable to propagate with a simple model due to near equivalence of results.

3) *Parallel Processing*: The simulation of the example 5 satellite DSM took 2.15 h using uniform sampling technique with $S_N = 20$. This can be reduced further by means of processing all (or part of the) 20 samples in parallel. Of course, the same maybe applied to the case of complete simulation, i.e., dividing the entire mission period into contiguous chunks of say 10 orbits duration each and propagating these chunks in parallel. Nevertheless, the number of parallel threads required for the complete simulation case far outnumbers that in the uniform random sampling case, and hence is relatively inefficient.

VI. CONCLUSION

The time-consuming propagation of orbits and performance quantification of DSM architectures has been a longstanding problem to apply heuristic optimization on DSM design and tradespace analysis. In this article, we propose and demonstrate two methods, the QSC algorithm and the random sampling method (with associated metrics of useful revisit events) which can be used to either independently or together in DSM evaluations. The fidelity of the results of the QSC algorithm depends on the underlying orbit-propagation model chosen by the user. The random sampling method on other hand approximates aggregated performance characteristics (level 1 or level 2 metrics) of the DSM by sampling the metrics (level 0) over the mission lifetime.

The QSC algorithm was demonstrated to accurately process narrow AT-FOV and conical FOV sensors, using the Landsat-8 TIRS pushbroom sensor as an example. Runtimes are shown to be two orders of magnitude faster than the traditional O&C calculations. We present results from use cases comparing the traditional O&C and the QSC method. The sensitivity of the execution time to the selection of the minimum propagation step-size (and hence the proxy-sensor FOV) to be used in the quick-search step was also studied. The QSC algorithm may be improved by considering a more general N-step search and correction process. While the QSC demonstrated in this article has just one “search” step, it may be replaced with N-search steps, and further optimize the execution time for coverage calculations.

Novel metrics for useful revisit events, such as mean/variance of useful revisit period and normalized number of useful revisits, instead of the traditional aggregation of all revisits, when quantifying the overall temporal response of a DSM, was proposed. We demonstrated the uniform random sampling method for quantifying DSM performance over mission lifetime for an example simulation of five satellite Walker DSM with Landsat-8 TIRS sensors. Simulation times and run times were shown to decrease by a factor of 13 upon using the random sampling technique with 20 samples (each of whose durations was 10 orbits), as compared to a baseline case when the DSM was simulated over the entire mission period of 180 days. This improvement is over and above that shown by the QSC algorithm.

REFERENCES

- [1] M. N. Sweeting, “Modern small satellites-changing the economics of space,” *Proc. IEEE*, vol. 106, no. 3, pp. 343–361, Mar. 2018, doi: [10.1109/JPROC.2018.2806218](https://doi.org/10.1109/JPROC.2018.2806218).
- [2] J. Le Moigne *et al.*, “Tradespace analysis tool for designing constellations (TAT-C),” in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Fort Worth, TX, USA, 2017, pp. 1181–1184, doi: [10.1109/IGARSS.2017.8127168](https://doi.org/10.1109/IGARSS.2017.8127168).
- [3] S. Nag, S. P. Hughes, and J. J. LeMoigne, “Navigating the deployment and downlink tradespace for earth imaging constellations,” in *Proc. 68th Int. Astronaut. Congr.*, Sep. 25–29, 2017, pp. 1–17.
- [4] T. J. Lang, “A parametric examination of satellite constellations to minimize revisit time for low earth orbits using a genetic algorithm,” *Adv. Astronaut. Sci.*, vol. 109, pp. 625–640, 2001.
- [5] E. A. Williams, W. A. Crossley, and T. J. Lang, “Average and maximum revisit time trade studies for satellite constellations using a multiobjective Genetic Algorithm,” *J. Astronaut. Sci.*, vol. 49, pp. 385–400, 2001.

- [6] N. Hitomi and D. Selva, "Constellation optimization using an evolutionary algorithm with a variable-length chromosome," in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, 2018, pp. 1–12, doi: [10.1109/AERO.2018.8396743](https://doi.org/10.1109/AERO.2018.8396743).
- [7] W. A. Crossley and E. A. Williams, "Simulated annealing and genetic algorithm approaches for discontinuous coverage satellite constellation design," *Eng. Optim.*, vol. 32, no. 3, pp. 353–371, 2000.
- [8] D. Reuter *et al.*, "The thermal infrared sensor (TIRS) on landsat 8: Design overview and pre-launch characterization," *Remote Sens.*, vol. 7, no. 1, pp. 1135–1153, 2015.
- [9] T. Arvidson *et al.*, "Landsat and thermal infrared imaging," in *Thermal Infrared Remote Sensing*. Dordrecht, the Netherlands: Springer, 2013, pp. 177–196.
- [10] R. Morfitt *et al.*, "Landsat-8 operational land imager (OLI) radiometric performance on-orbit," *Remote Sens.*, vol. 7, no. 2, pp. 2208–2237, Feb. 2015, doi: [10.3390/rs70202208](https://doi.org/10.3390/rs70202208).
- [11] N. Longbotham, F. Pacifici, S. Malitz, W. Baugh, and G. Camps-Valls, "Measuring the spatial and spectral performance of World View-3," in *Fourier Transform Spectroscopy and Hyperspectral Imaging and Sounding of the Environment*, Paper HW3B.2. 2015.
- [12] F. Rostan, S. Riegger, M. Huchler, R. Croci, and R. Torres, "The Sentinel-1 C-SAR instrument," in *Proc. 9th Eur. Conf. Synthetic Aperture Radar*, Nuremberg, Germany, 2012, pp. 166–169.
- [13] Y. Kankaku, M. Sagisaka, and S. Suzuki, "PALSAR-2 launch and early orbit status," in *Proc. IEEE Geosci. Remote Sens. Symp.*, Montreal, QC, Canada, 2014, pp. 3410–3412, doi: [10.1109/IGARSS.2014.6947213](https://doi.org/10.1109/IGARSS.2014.6947213).
- [14] E. Peral *et al.*, "The Radar-in-a-Cubesat (RAINCUBE) and measurement results," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Valencia, Spain, 2018, pp. 6297–6300, doi: [10.1109/IGARSS.2018.8519194](https://doi.org/10.1109/IGARSS.2018.8519194).
- [15] M. Á. Palacios Lázaro, M. Bergadà Pujades, R. González Sola, A. Gamonal Coto, J. L. García Fernández, and Y. Camacho, "Design, development and calibration of the MWR microwave radiometer on board Sentinel-3," in *Proc. 44th Eur. Microw. Conf.*, Rome, 2014, pp. 1671–1674, doi: [10.1109/EuMC.2014.6986776](https://doi.org/10.1109/EuMC.2014.6986776).
- [16] C. S. Ruf, S. Gleason, and D. S. McKague, "Assessment of CYGNSS wind speed retrieval uncertainty," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 12, no. 1, pp. 87–97, Jan. 2019.
- [17] W. J. Blackwell *et al.*, "An overview of the TROPICS NASA earth venture mission," *Quart. J. Roy. Meteorol. Soc.*, vol. 144, pp. 16–26, 2018.
- [18] P. G. Buzzi, D. Selva, N. Hitomi, and W. J. Blackwell, "Assessment of constellation designs for earth observation: Application to the TROPICS mission," *Acta Astronautica*, vol. 161, pp. 166–182, Aug. 2019.
- [19] S. Nag, V. Ravindra, and J. L. Moigne, "Instrument modeling concepts for tradespace analysis of satellite constellations," in *Proc. IEEE Sens.*, New Delhi, India, 2018, pp. 1–4, doi: [10.1109/ICSENS.2018.8589764](https://doi.org/10.1109/ICSENS.2018.8589764).
- [20] J. R. Wertz and W. J. Larson, *Space Mission Analysis and Design* (Space Technology Library), 3rd ed. Portland, OR, USA: Microcosm Press, 1999.
- [21] S. Nag, S. P. Hughes, and J. L. Moigne, "Streamlining the design tradespace for earth imaging constellations," in *Proc. AIAA Space Conf.*, Long Beach, CA, USA, 2016, pp. 1–17.
- [22] D. C. Morton *et al.*, "Amazon forests maintain consistent canopy structure and greenness during the dry season," *Nature*, vol. 506, pp. 221–224, 2014.
- [23] E. W. Weisstein "Standard error," 2003. [Online]. Available: <http://mathworld.wolfram.com/StandardError.html>



Vinay Ravindra (Member, IEEE) received the Ph.D. degree in electrical engineering and information systems from the University of Tokyo, Tokyo, Japan, in 2017.

He is a Research Scientist with NASA Ames Research Center, Mountain View, CA, USA, contracted by Bay Area Environmental Research Institute. His research interests include satellite mission design and analysis, synthetic aperture radar development, instrument modeling, and attitude determination and control of satellites.



Sreeja Nag (Member, IEEE) received the Ph.D. degree in space systems engineering from Massachusetts Institute of Technology, Cambridge, MA, USA.

She is a Senior Research Scientist with NASA Ames Research Center, Mountain View, CA, USA, contracted by Bay Area Environmental Research Institute. Her research interests include distributed space systems, automated planning and scheduling of constellations, swarm decision making in space, and space traffic management.