MAXIMIZING DUST DEVIL FOLLOW-UP OBSERVATIONS ON MARS USING CUBESATS AND ON-BOARD SCHEDULING

Robyn Woollands*, Federico Rossi ; Tiago Stegun Vaquero, Marc Sanchez Net, S. Sandra Bae ** Valentin Bickel, Joshua Vander Hook

Several million dust devil events occur on Mars every day. These events last, on average, about 30 minutes and range in size from meters to hundreds of meters in diameter. Designing low-cost missions that will improve our knowledge of dust devil formation and evolution, and their connection to atmospheric dynamics and the dust cycle, is fundamental to informing future crewed Mars lander missions about surface conditions. In this paper we present a mission for a constellation of low orbiting Mars cubesats, each carrying imagers with agile pointing capabilities. The goal is to maximize the number of dust devil follow-up observations through real-time, on-board scheduling. We study scenarios where cubesats are equipped with a 2.5 degree boresight angle camera that accommodates five slew positions (including nadir). We assume a concept of operations where the cubesats autonomously survey the surface of Mars and can autonomously detect dust devils from their surface imagery. When a dust devil is detected, the constellation is autonomously re-tasked through an on-board distributed scheduler to capture as many follow-on images of the event as possible, so as to study its evolution. The cubesat orbits are propagated assuming two-body dynamics and the ground tracks and camera field of view are computed assuming a spherical Mars. Realistic inter-agent communication link opportunities are computed and included in our optimization, which allow for real-time event detection information to be shared within the constellation. We compare against a powerful "omniscient" oracle which has a priori knowledge of all dust devil activity to show the gap between predicted performance and the best possible outcome. In particular, we show that the communications are especially important for acquiring follow-up observations, and that a realistic distributed scheduling mechanism is is able to capture a large fraction of all dust devil observations that are possible for a given orbit configuration, significantly outperforming a nadir-pointing heuristic.

INTRODUCTION

Martian dust storms have had a major impact on Mars exploration efforts. Planet-wide storms can bring exploration operations to a halt by reducing available power or obscuring view. This has serious implications for both the explore-ability and habitability of Mars. An understanding of the

^{*}Mission Design Engineer, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.

[†]Assistant Professor, Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, IL 61801.

[‡]Robotics Technologist, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.

[§]Data Scientist, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.

[¶]Telecommunications Engineer, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.

^{II} Summer Intern, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.

^{**}PhD Student, University of Colorado Boulder, CO 80309.

^{††}Max Planck Institute for Solar System Research

^{‡‡}Group Supervisor, Jet Propulsion Laboratory, California Institute of Technology, CA 91109.



Figure 1. A CTX image and HiRISE zoom-in (inset) of a large dust devil in Amazonis Planitia. A framing version of CTX, likely well within today's photoimaging capabilities, may be able to record near-real-time imagery and track dust devil size, distribution, and shape over time when combined with on board processing. Credit: NASA-JPL, University of Arizona

dust cycle, with its connections between regional and global storms, thermal exchange between the surface and atmosphere, and the circulation that can eventually lead to atmospheric escape, remains a top scientific priority.¹

Yet, as dynamic, widespread phenomena, these processes are challenging to capture with global context from the lower orbits that have dominated Mars orbiter missions, and nearly impossible to capture in the detail required if higher orbits are assumed (e.g., Areostationary). Our goal in this paper is to find a middle ground using a constellation of small orbiters collaborating at 200-400 km altitude. In this paper, we present initial feasibility studies of a constellation of small orbiters using medium-resolution cameras to track dust devils on the surface of Mars in real-time.

By using a combination of medium-resolution imagers, on board processing, and cross-platform cueing, the constellation will enable short and long time-scale studies of dust devil formation, propagation, and global distribution over seasonal weather variations. We take, as inspiration, the recent NASA study on a multi-satellite network (MOSAIC)¹ for Mars.

We proceed as follows. We evaluate with a five satellite network. We assume that the five satellites were sufficiently small to be delivered as a secondary payload, perhaps on an ESPA ring accompanying a main mission.² The satellites are assumed to have basic telemetry and commanding interfaces direct to Earth, but otherwise relay their data through the main mission platform. We do not model the main satellite or nearby telecommunications relays, and simply assume there is "sufficient" capacity for this role. Note that MOSAIC¹ does include such relays for their mission concept, and significant literature points to the need and high likelihood of such relays being deployed after the retirement of the Mars Reconnaissance Orbiter.^{3,4}

However, we do assume that the five satellites are equipped with communications arrays to communicate between members of the network. The main use of this limited communication ability is to exchange information about the surface conditions between each observer and "cue" other members. There is precedent for this concept of operation. For instance, it resembles the "A-train" configura-



Figure 2. Performance illustration of an exemplary, CNN-based dust devil detector, operating on rectangular tiles of various sizes cropped from full-frame CTX images. A) & B) correct detections of small and large dust devils (white, TP = true positive); C) a dust devil missed by the detector (yellow, FN = false negative); D) false detections (red, FP = false positive). The network confidence score is displayed with its respective detection. The smallest (and faint) dust devil in this example is ~100 meters across.

tion near Earth, but includes significantly more autonomy to cope with the long propagation delays between Earth and Mars.⁵ Similarly, it is also very close to the "sensor network" configuration, in that cues are sent between members of the network to schedule follow-up observations.⁶ The total data volume exchanged over the network will be kept small so that communications is not the limiting factor of the design.

The members of the constellation are assumed to have medium-resolution framing imagers. The ground sample distance is assumed to be similar to the Mars Context Camera⁷ (CTX). This setup (approximately 6 meters / pixel) allows the largest dust devils to be resolved.⁸ For example, Figure 1 shows a famous, large dust devil at CTX resolution.

Global coverage is achieved by "being vigilant" that is, taking significantly more imagery than is usual. Yet, detail is preserved since the images are taken at lower orbit and with higher temporal frequency than is usual. This introduces a data bottleneck: the constellation is gathering significantly more imagery than can be downlinked, and so significant on board processing is required to prepare downlink reports of surface activity, and on board autonomy is required to schedule follow-up observations of the most promising dust devil candidates, as discussed in.⁹

One way of automating the detection of dust devils is by using convolutional neural networks (CNNs). Testing of a basic, ResNet 50-based CNN trained with a few hundred labels shows that one can achieve Average Precision (AP) values of around 0.7 in CTX data (a detailed description

is out of scope but forthcoming). The AP could be further optimized by applying a stricter score threshold value (optimize precision), which would reduce the recall, however. Initial testing with CTX data indicates that a CNN-based detector is fairly robust against varying geomorphic backgrounds (dunes, craters, etc.), image quality, and illumination conditions, while being able to detect small, large, faint, and distinct dust devils alike (see Figure 2). Experience has shown that dust devils as small as ~ 60 meters across can be detected and mapped in CTX images. In this work, we assume that a well trained classifier is available for use onboard.

With on board processing and modest cross-platform communication, we now have our core problem statement: How to position the constellation and schedule the observations to achieve maximum science return. The focus of the remainder of the paper is on designing an autonomous on board scheduling system that trades exploration (look in areas that have not been photographed recently) vs. exploitation (take another image of a known dust devil).

STATE OF THE ART

Automated tasking as part of a ground-orbiting system has a long history in space missions,^{10,11} including successful demonstrations for earth observing spacecraft.^{12,13} A more detailed review of autonomous agents for space exploration can be found in¹⁴ and.¹⁵ The state of the art is timelinebased, meaning exact subsystem sequences are derived for a single agent using constraints and resource-based cost functions¹⁶ in a framework known as ASPEN. ASPEN is based on the previous generation, Continuous Activity Scheduling Planning Execution and Replanning (CASPER), and along with Eagle Eye, a specialized version of the software, this scheduler generates a baseline mission operations plan from observation requests. The scheduler "greedily" schedules commands in a priority-first manner. In addition to generating an initial plan, ASPEN/Eagle Eye, steps through a *plan execution* phase, an *onboard image analysis* phase, an *onboard replanning* phase, and a *target reimaging* phase. When rescheduling, the algorithm again follows a greedy approach and searches the current plan within the desired time window and replaces the earliest available observation that can be replaced with an updated observation request. Recent work (demonstrations of which are two decades in the making) has moved toward expressing mission goals and decomposing them into activities onboard the spacecraft.¹⁷

There is little precedent for multi-agent variants of these planning and scheduling technologies, and fewer examples of demonstrations. See¹⁸ for an example of multi-rover autonomous operations and¹⁹ for multi-aircraft / multi-spacecraft networks for earth observations, and²⁰ for recent work on constellation communication networks and tasking. A particularly close example is²¹ which discusses the science benefit of a constellation for multi-angle simultaneous observations of the Earth atmosphere. In that work, Nag et. al optimize the constellation's configuration to maximize science gain. However, onboard autonomy was not considered. In²² an optimization framework is presented that can schedule multi-satelite constellations. However, the framework in² ((like all related work mentioned) is optimized for quick convergence; in contrast, the work in this paper produces *provably optimal* schedules. In this work we focus on analyzing the benefit of communications and onboard planning. We take as input models of dust devil activity and a simplified model of the spacecraft dynamics and controls to allow a comprehensive, long-term study of the efficacy of the network.

It is worth noting that automated detection of dust devils on Mars is not a new concept, although it was previously done from the surface.²³ The automated collection of scientifically relevant imagery from rovers has been well vetted.²⁴

Given the decades-long history of automated planning and scheduling for space operations, our study takes as granted a framework for turning high level commands into low level control sequences, such as ASPEN,²⁵ and the network design considerations discussed in.²⁰

PROBLEM STATEMENT

We formulate the constellation follow-up observation scheduling as follows.

We assume that dust devils occur across the Martian surface according to a stochastic process with known spatial distribution. Once a dust devil forms, its temporal duration and spatial evolution are also stochastic, with known distribution.

A constellation of spacecraft observes the Martian surface from orbit with cameras. Each spacecraft is equipped with a dust devil detector (with imperfect precision and recall) that can detect dust devils present in images and infer their location on the surface. The spacecraft's orbits are fixed; however, the spacecraft can decide where to point their cameras, subject to slewing constraints that enforce a maximum slewing rate.

Spacecraft can share their observations (specifically, the location of detected dust devils) through communication links when they are in line-of-sight contact of each other.

The goal of the problem is to maximize the number of follow-on observations of dust devils; specifically, we assume that it is more valuable to collect multiple observations of the same dust devil than to gather single observations of distinct dust devils, up to a maximum number of observations per event.

We formalize the scheduling problem as follows. We define:

- \mathcal{R} , a set of regions of interest $\mathcal{R} = \{r_1, ..., r_R\}$ on the surface of Mars where dust devils may occur; spacecraft observe these regions to look for previously-undetected dust devils. Each region should be imaged at most once over a given time window.
- \mathcal{E} , a set of observed transient events $\mathcal{E} = e_1, \ldots, e_E$ on the surface of Mars, modeling dust devils that have been detected and for which follow-up observations are desirable. Each event is associated with a given time window, which captures the expected duration of the dust devil. Within that window, the event should be observed as many times as possible.
- $H = (H_{start}, H_{end})$, a planning time horizon with start and end time points.
- S, a set of spacecraft agents $S = \{a_1, ..., a_N\}$ orbiting Mars.
- \mathcal{A} , a set of attitudes that each spacecraft $a \in \mathcal{S}$ can assume at every time step.
- O, a set of observation opportunities {o_{i,a,t}} for agent i ∈ S at time t in attitude a within the time (H_{start}, H_{end}). For each observation opportunity, an observation opportunity function (o_{i,a,t}) → (O ∪ E ∪ Ø) reports what regions and what transient events can be observed by agent a ∈ S in attitude a ∈ A at time t ∈ (H_{start}, H_{end}).
- U(r, o): O ∪ E, N → R, a scoring function that maps regions of interest r ∈ R and transient events e ∈ E, and the number o of previous observations of the region or event of interest, to a score value. For regions R, the priority function encodes the expected number of new dust devils that will be observed there, multiplied by the value of a new observation; for events E,

the scoring function denotes the scientific interest of collecting the *o*-th follow-on observation of the event.

Our objective is to produce an observation schedule by selecting $\hat{O} \subseteq O$ to maximize $\sum_{r \in R} U(r)$ subject to agents' attitude constraints.

We are now in a position for formalize the constellation follow-up observation scheduling problem.

Problem 1 (Constellation follow-up observation scheduling problem). Find the set of observations $\hat{O} \subseteq O$ that maximize $\sum_{r \in B} U(r, o)$ subject to instrument constraints.

AUTONOMOUS POINTING FOR SCIENCE PLANNING



Figure 3. Architecture of the proposed constellation scheduler. (A) Each spacecraft captures images of the Martian surface. A dust devil detector searches for dust devils in captured imagery. (B) When a dust devil is detected, its location is added to a map database which tracks possible active dust devils that should be targeted for revisit. Each spacecraft executes an instance of the constellation planner (described in the Autonomous Pointing for Science Planning section), which uses the locations of the detected dust devils to prescribe pointing locations for every spacecraft. (C) The spacecraft then executes the pointing and image capture commands assigned to itself by the constellation planner, and discards the assignments for other spacecraft. The map database of detected dust devils is opportunistically synchronized across spacecraft when a communication link is available.

A centralized ILP algorithm

We start by presenting a centralized integer programming formulation to solve Problem 1.

We discretize the time horizon of the problem in equally-spaced intervals $\mathcal{T} = [t_1, \ldots, t_T]$. We also discretize the set of attitudes that each spacecraft can point to into a set $\mathcal{A} = [k_1, \ldots, k_K]$ of discrete attitudes.

Pointing function We assume that a pointing function $p : (S, A, T) \mapsto \{R \cup E \cup \emptyset\}$ is known. The set p(i, a, t) represents the regions $r \in R$ and the known, ongoing transient events $e \in E$ that can be observed by spacecraft $i \in S$ in attitude $a \in A$ at time $t \in T$.

Slewing constraints We also assume that the spacecraft cannot transition between arbitrary pairs of attitudes in one time step due to slewing constraints. For a given attitude $k \in A$, we denote the set of feasible prior attitudes (i.e., the set of attitudes from which it is possible to transition to k in one time step) as P(k).

Variables We define the following variables:

• `

- $X(s, a, t), i \in \mathcal{S}, a \in \mathcal{A}, t \in \mathcal{T}$, a set of Boolean variables. X(s, a, t) is 1 iff spacecraft s assumes attitude a at time t.
- $Y(s, a, t), i \in \mathcal{S}, a \in \mathcal{A}, t \in \mathcal{T}$, a set of Boolean variables. X(s, a, t) is 1 iff spacecraft i captures data in attitude a at time t.
- $O(e, o), e \in \mathcal{R} \cup \mathcal{E}, o \in [1, \dots, \overline{O}(e)]$, where $\overline{O}(e)$ is the maximum number of observations of region or event e. A set of Boolean variables. O(e, o) is 1 iff region or event e is observed at least o times.

Problem Formulation We are now in a position to formalize Problem 1 as a ILP.

$$\max_{X,Y,O} \sum_{e \in \mathcal{R} \cup \mathcal{E}} \sum_{o \in [1,\dots,\overline{O}(e)]} U(e,o)O(e,o)$$
(1a)

subject to

 $a \in A$

$$O(e, o) \le O(e, o - 1) \qquad \forall e \in \mathcal{R} \cup \mathcal{E}, o \in [2, \dots, \overline{O}(e)] \qquad (1b)$$
$$\overline{O}(e)$$

$$\sum_{o=1}^{O(e)} O(e,o) = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} Y(s,a,t) \cdot \mathbf{1}_{e \in p(s,a,t)} \qquad \forall e \in (\mathcal{R} \cup \mathcal{E})$$
(1c)

$$Y(i, a, t) \le X(i, a, t) \qquad \forall i \in S, a \in \mathcal{A}, t \in \mathcal{T}$$

$$\sum X(i, a, t) = 1 \qquad \forall i \in S, t \in \mathcal{T}$$
(1d)
(1e)

$$X(i, a, t+1) \le \sum_{l \in P(a)} X(i, l, t) \qquad \forall i \in \mathcal{S}, \forall t \in \mathcal{T} \setminus t_T, \forall a \in \mathcal{A}$$
(1f)

Equation (1a) captures the reward for observing a given event or region for the o-th time. Equation (1b) ensures that the o-th observation of a region or event is only performed after the (o-1)-th observation. Equation (1c) enforces that the overall number of observations of a region matches the number of times a spacecraft has captured an image pointing at the region. Equation (1d) guarantees that a spacecraft only captures an observation in a given attitude if it is indeed pointed in that attitude. Equation (1e) ensures that each spacecraft assumes only one attitude at every time step. Finally, Equation (1f) ensures that the sequence of scheduled observations is compatible with the spacecrafts' attitude slewing constraints.

Distributed implementation

Solving Equations (1a)-(1f) requires knowledge of all observed events, and the solution prescribes a set of attitudes for all spacecraft - therefore, the problem cannot directly be used to control a constellation where spacecraft are not in constant communication with each other and must make independent decisions based on their own observations. To overcome this, we propose a sharedworld implementation of the ILP, shown in Figure 3. Every spacecraft independently detects dust devils in its own observations and adds detections to a local copy of a global event map. When

two spacecraft are in contact, they reconcile their copies of the global event map by sharing their observations; if a conflict arises (e.g., if two spacecraft have observed the same region but only one has detected a dust devil), we optimistically assume that the detection was correct. Each spacecraft individually solves Equations (1a)-(1f) based on its local copy of the global event map, obtaining a set of prescribed attitudes for the entire constellation. Each spacecraft then executes the pointing commands prescribed to itself by the ILP, and disregards the pointing commands for other spacecraft.

With this approach, if the spacecrafts' global events map are perfectly synchronized, all spacecraft act in accordance with the same solution to Equations (1a)-(1f), If the global events map are not synchronized (because, e.g., information about a detected event has not yet propagated to all spacecraft), each spacecraft is guaranteed to have access to a solution, even though spacecraft may act in an inconsistent manner (e.g., a region may be imaged twice, or it may not be imaged because each spacecraft thinks that another one will capture it). In the Numerical Results, we show that, for the orbits and network topologies considered in this paper, information propagates readily through the constellation, and the impact of imperfect synchronization between the spacecrafts' global event maps is negligible.

Receding-horizon approach

Obtaining a solution to Problem 1 requires knowledge of observed dust devils: therefore, the problem should be re-solved whenever a new dust devil is observed. To address this, we implement the algorithm in a receding-horizon fashion. Each spacecraft solves Equations (1a)-(1f) with a fixed time horizon - in the simulation results, we use a fifteen-minute time horizon. As soon as new information is received (either through a direct observation, or through an update to the global event map provided by another spacecraft), the spacecraft re-solves the problem with the same horizon. Also, before the horizon is reached, the problem is solved again, irrespective of whether new information is available. The receding horizon approach ensures that the computation cost is highly contained (as shown in the Numerical Simulation & Results section), while incorporating new information in the solution as soon as it becomes available.

NUMERICAL SIMULATION & RESULTS

Dust Devil Model

Dust devils form around low-pressure air pockets, where the air is drawn into a narrow rising column through the surrounding cooler air. A study that combined data from the High Resolution Stereo Camera (HRSC) and the Mars Orbiter Camera (MOC) found that some large dust devils on Mars have a diameter of 700 m and last for at least 26 minutes.⁸ The largest dust devils can reach a height of 8 km. A recent study found that on average, during any given day, one 13 m wide dust devil appears in every square kilometer on Mars.²⁶

As a dust devil moves across the surface of Mars, it lifts the top layer of dust into the atmosphere and exposes the dark underlying surface. These dark tracks can last a few weeks, after which they are either covered up as a result of wind action or they become oxidized from exposure to sunlight and the Martian atmosphere, and thus return to same red color as the rest of the planet. Dust devil tracks can be more than 30 m wide and extend for more than 4 km.

In our simulations, we use a simple *random model* for dust devil generation. We assume that no dust devils exist at latitudes greater than 70° and no dust devils exist before 11 am or after 4 pm

local solar time. We center the active dust devil region at 1:30 pm local solar time. Over this region, we simulate dust devils with an average density of one per ten square kilometers. This is a more conservative estimate than those found in the literature,²⁶ but it is adequate to test our on-board scheduler. In our simulations, the mean dust devil lifetime is taken to be 26 minutes and start and end times are computed for each dust devil. Furthermore, we have not taken into account the motion of the dust devils across the surface and their associated ground track. This will be considered in future work, along with a more sophisticated dust devil model.

Orbits

In this paper we propagate two-body orbits for a train of five cubesats in a low Mars (200 km altitude), Sun synchronous orbit. Two different train constellations are considered. The first has satellites in the train separated by one minute time intervals, and second uses six minute intervals. We have selected a low Mars orbit such that the cubesat camera (for example, HiREV²⁷) is capable of imaging the surface with a resolution of less than 10 m. A Sun synchronous orbit is chosen to ensure that the lighting conditions (shadow angles and shadow lengths) are the same for every pass. Although we do not simulate dust devil event detection through image processing in this paper, we assume that using the same lighting conditions for each pass will aid the eventual dust devil event detection algorithm, thus reducing the potential for false positives. Each cubesat camera has a field of view with 2.5° boresight angle.²⁷ For these simulations, we have selected five camera positions: nadir, $\pm 5^{\circ}$ and $\pm 10^{\circ}$ with respect to nadir. We assume that the slew/settle time to move between camera positions is one minute.²⁸

Figure 4 shows a snapshot taken during the simulation where the cubesats in the train are separated by one minute intervals. The top panel shows a spherical Mars that is experiencing a southern hemisphere summer. Each yellow dot on the surface represents an active dust devil. The low-Mars orbit is also shown in cyan/red, along with the ground track. Note that the five satellites are colored red, blue, green, black and cyan respectively. The magenta circles represent currently detected dust devils. The bottom panel shows a 2D projection of Mars, along with the same information plotted as in the top panel.

Figure 5 shows an enlarged view for six snapshots taken at one minute time intervals. In these figures, the ground coverage bounding box is shown for all five possible camera positions simultaneously. Of course, in reality, the camera can only image one of these regions at a time. Looking at the top left panel, if the lead satellite (red) was pointing 5° ahead of nadir (i.e. *camera position 2*), it would detect one dust devil event (yellow dot with magenta circle), and if it was pointing 5° behind nadir (i.e. *camera position 4*) it would detect two dust devil events. The next snapshot (top right) shows that one minute later the second satellite (blue) is capable of imaging the dust devils that were previously detected by the lead satellite a minute earlier. The remaining panels show the satellite train moving further along in its orbit as each minute passes.

Field of View

The camera field of view is modelled as a square and the corresponding observable surface of Mars is computed as the region enclosed within the four corners of this square projected onto the surface of a spherical Mars. A dust devil is detected if it falls within this visible surface element on Mars. In our simulations, all the orbits are sampled at one minute time intervals. Thus if the lead spacecraft detects an event in, say, *camera position 1*, it can communicate to the second spacecraft to slew from its current camera position to *camera position 1* in order to reimage the dust devil event

one minute later, thus enabling the dust devil time history to be recorded. Similarly, the second spacecraft can communicate to the following one to slew to the same position, and so on. If the lead spacecraft detects a dust devil event and the second spacecraft also detects an event, then the on-board scheduler, running in real-time, will determine which positions the following spacecraft should slew to in order to maximize the total number of follow-on dust devil observations.

Communication

In our simulations, we assume that each cubesat is equipped with a patch antenna. For this orbit geometry (200 km altitude), with the one minute separation between satellites, all elements of the chain are constantly within line-of-sight communication, thus enabling information to be transferred between the satellites at any time. For the case where satellites are separated by six minutes, the line-of-sight, and hence the communication link, only extends to the satellite immediately in front and behind it. We assess the performance of the proposed approach in both scenarios.

Results

Comparison with upper and lower bounds First, we compare the performance of the proposed constellation scheduling approach with an "oracle" upper bound and a naive lower bound.

The oracle solves Equations (1a)-(1f) once with full knowledge of the location of all present and future dust devils. The problem is solved in a centralized manner. The solution provided by the oracle represents an (unattainable) upper bound on the performance of the constellation.

The lower bound is obtained by simply setting the attitude of all spacecraft to be in a nadirpointing direction at all times.

The two scenarios are compared with the outcome of the simulation where each spacecraft solves Equations (1a)-(1f) in a receding-horizon fashion, as described in the Autonomous Pointing section, and spacecraft exchange observations along available communication links as discussed in the Communication section. To ensure a fair comparison, all spacecraft are assumed to have access to an ideal event detector with perfect precision and recall (that, is, it never produces false positives or negatives); the effect on performance of a non-ideal event detector is discussed in the next section.

The results are reported in Table 1 and Figure 6 for a set of orbits with one minute spacing between spacecraft, and in Table 2 and Figure 7 for orbits with six minute spacing.

Remarkably, the proposed approach results in a more than two-fold increase in the number of events that receive at least one follow-up observation compared to the nadir-pointing approach for both sets of orbits considered. The average number of follow-up observations per event increases from 0.13 to 0.65, an almost five-fold increase, for widely spaced satellites; for closely-spaced satellites, a three-fold increase is obtained. In both cases, the proposed approach collects 7% more observations compared to the nadir-pointing case. The number of unique events detected is 15% smaller in the closely-spaced case and 25% smaller in the widely-spaced case: as expected, the autonomous pointing approach privileges follow-on observations as opposed to new event detections, in line with the proposed reward function, which privileges revisits over new detections.

Effect of inter-spacecraft communication As discussed in the Distributed implementation section, the distributed implementation relies on inter-spacecraft communication to maintain a common picture of observed events; this can result in suboptimal, and potentially inconsistent, behavior at





Figure 4. A snapshot taken during the simulation that shows the five sun synchronous, low-Mars orbiting satellites, flying over a region of dust devil activity (yellow dots) during a southern hemisphere summer. The top panel shows a spherical Mars and the bottom panel shows a two-dimensional projection. The respective orbits and ground tracks are also shown.



Figure 5. A series of six snapshots taken at one minute intervals that show the instantaneous ground coverage for each of the five cubesats in the chain. Note that all five camera slew positions are shown simultaneously in this figure. Yellow dots represent active dust devils and magenta circles represent possible detections.



Figure 6. Closely-spaced orbits: distribution of the number of follow-on observations for the proposed approach, compared to an upper and lower bound. The effect of a non-ideal event detector are also shown.

the constellation level if inter-spacecraft links are sporadic or if communication between spacecraft requires multiple hops.

To explore the impact of inter-spacecraft communications on system performance, we compare the proposed approach, where spacecraft communicate with their neighbors when a line-of-sight link is available, with two scenarios: an *idealized* scenario where all spacecraft can exchange information instantly, and a *no-communications* scenario where each spacecraft only relies on information it has collected itself. We focus our attention on the widely-spaced scenario, where each spacecraft can only communicate with its immediate neighbors in the constellation.

Results are shown in Table 2 and Figure 8.

The difference between the proposed approach and the upper bound where all spacecraft can communicate with each other is negligible, showing that the proposed approach is able to effectively spread information about event detections across the constellation. In contrast, the case where satellites do not communicate with each other only presents a small improvement with respect to the nadir-pointing case, with a 38% improvement in the average number of follow-up observations (vs. the almost five-fold improvement obtained with communications). Collectively, these results show that the proposed distributed optimization approach is able to approach the performance of a centralized controller, and that constellation-level optimization provides an order-of-magnitude improvement in performance compared to uncoordinated planning by individual satellites.

Non-ideal event detector We also assess the impact of using a non-ideal event detector with non-perfect precision and recall. Specifically, we simulate the performance of a detector with 70%

Table 1. Performance of the proposed distributed pointing approach compared to an "oracle" upper bound and a fixed-attitude lower bound. Closely spaced orbits (one minute gap between spacecraft) are considered.

	Events observed				Follow-ups	
	Total	Unique	2+	3+	Mean	Max
Oracle (upper bound)	1607	560	431	354	1.87	8
Nadir-pointing (lower bound), ideal detector	743	646	95	0	0.15	2
This paper, ideal detector	794	548	198	38	0.45	4
Nadir-pointing (lower bound), non-ideal detector	507	463	44	0	0.095	2
This paper, non-ideal detector	527	423	92	11	0.25	3

Table 2. Performance of the proposed distributed pointing approach compared to an "oracle" upper bound and a fixed-attitude lower bound. More widely spaced orbits (six minute gap between space-craft) are considered.

	Events observed				Follow-ups	
	Total	Unique	2+	3+	Mean	Max
Oracle (upper bound)	1753	914	507	227	0.92	8
Nadir-pointing (lower bound), ideal detector	832	734	93	4	0.13	4
This paper, ideal detector	897	552	235	85	0.63	5
Nadir-pointing (lower bound), non-ideal detector	576	531	43	2	0.08	3
This paper, non-ideal detector	648	445	157	39	0.45	5
This paper, ideal detector, continuous comms	896	557	234	83	0.608	5
This paper, ideal detector, no comms	838	713	115	6	0.18	5

recall, i.e., a 30% false negative rate. Due to limitations of the simulation environment, the study of the effect of non-ideal precision is left as future work. Results are shown in Tables 1 and 2 and in Figures 6 and 9.

Predictably, the imperfect recall results in a roughly 30% reduction in the overall number of observations collected. Specifically, we observe 23% and 19% reductions in the number of unique detected events in the closely-spaced and widely-spaced cases respectively; the reduction in the average number of follow-up images collected for each event is 53% in the closely-spaced case, and 29% in the widely-spaced case. Compared to a nadir-pointing case with imperfect recall, the proposed approach performs similarly to the perfect case, with a twofold to fourfold increase in the number of events with at least two observations, and a threefold to fivefold increase in the average number of follow-up observations per event.

Computational performance The ILP was solved with the CPLEX solver. The time required to formulate the problem (i.e., encode the cost function and constraints) and solve it on a commodity desktop workstation is shown in Figure 10. The time required to formulate the problem is consistently below 40ms, and the time required to solve it is almost always below 20ms - results which suggest that the proposed approach could be readily implemented on more modest CPUs well-suited for spaceflight.



Figure 7. Widely-spaced orbits: Distribution of the number of follow-on observations for the proposed approach, compared to an upper and lower bound.



Figure 8. Widely-spaced orbits: Effect of communications on distribution of the number of follow-on observations.



Figure 9. Widely-spaced orbits: Effect of non-ideal detector on distribution of the number of follow-on observations.



Figure 10. Computation times to solve the proposed ILP on a Xeon E5-2687W.

VISUALIZATION SYSTEM



Figure 11. A visualization system (a) Side-bar menu with options to change the map projection and orientation (b) Coverage map depicts the ground track of the constellation (c) Overview provides the state of the constellation and dust devils (d) Counters for the total number of dust devil events and dust devil observations (e) Time scrubber depicts the current time position within the simulation.

To visualize and interact with the simulation, we designed a visualization system that adheres to the same mission design methodology as discussed in the previous sections. The system is composed of three interactive views: overview, coverage map, and time scrubber (Figure 11). All views are interlinked: the overview and coverage map visualizes data corresponding to the particular timestamp from the time scrubber.

Overview. The overview (Figure 11b) provides a quick assessment of the constellation performance in terms of dust devil observations. Dust devils are represented as brown triangles, which appear and disappear throughout the simulation run. A counter (Figure 11d) helps count the total number of dust devil events. In contrast, spacecraft are represented by their bounding boxes and spacecraft number. As mentioned in the Field of View subsection, these bounding boxes are determined from the five slew positions each spacecraft is capable of achieving. All bounding boxes are mapped with two visual encodings. First, for each spacecraft, a bounding box has solid edges if it represents the spacecraft's current camera position, while the rest are represented with thin dashed edges. Second, a bounding box is colored red if a spacecraft 4 has observed one or more dust devil in its field of view at a particular time point (e.g., Spacecraft 4 has observed a dust devil at t = 32580). Otherwise, the color is transparent. The visual encoding for the dust devil follows a similar logic, where the dust devil is colored red if it is observed and remains brown otherwise. To help users keep track, a second counter (Figure 11d) counts the total number of dust devil events the constellation has observed.

The combination of these visual encodings allows the user to not only visually see when and where spacecraft have observed dust devils, but also highlights where the vehicles *could* have observed a dust devil event (i.e., a dust devil is within the field of view of the four bounding boxes that are not in use).

The overview also visualizes the terminator to showcase the daylit side (white illuminated region) and the night side (solid dark grey filled-in region) of Mars throughout the simulation run. This corresponds to the timer (unit in hours) at the top of the overview. The default map projection is Kavrayskiy VII, but users can change to different map projections as well as zoom-in, rotate, and pan using the side-bar menu (Figure 11a). Users can also see details on-demand (i.e., via an interactive tooltip).

Coverage Map. The coverage map (Figure 11b) dynamically visualizes the ground track of the constellation throughout the simulation (t = 0 to t = 88320). The ground track is colored dark steelblue, and the path corresponds to the movement of the spacecrafts in the overview (Figure 11b). The coverage map provides a visual insight as to which regions the constellation has frequently visited and vice versa. This can help analysts compare different constellation formations, and in particular different inter-spacecraft spacing.

Time Scrubber. The time scrubber (Figure 11e) is a horizontal slider control that lets the user dynamically seek any time position within the currently running simulation. The scrubber 'knob' (circle) moves to indicate the current time position. The position of the scrubber knob determines the information being displayed in the overview and coverage map. Users can also pause and play when analyzing the simulation.

CONCLUSIONS

Several million dust devil events occur on Mars every day. These events last, on average, about 30 minutes and range in size from meters to hundreds of meters in diameter. Designing low-cost missions that will improve our knowledge of dust devil formation and evolution, and their connection to atmospheric dynamics and the dust cycle, is fundamental to informing future crewed Mars lander missions about surface conditions. In this paper we presented a mission for a constellation of low orbiting Mars cubesats, each carrying imagers with agile pointing capabilities. The goal was to maximize the number of dust devil follow-up observations through real-time, on-board scheduling. Realistic inter-agent communication link opportunities were computed and included in our simulation, which allowed for real-time event detection information to be shared within the constellation. In particular, we found that the communications are especially important for acquiring follow-up observations, and that the proposed distributed scheduling mechanism is able to capture a large fraction of all dust devil observations that are possible for a given orbit configuration, significantly outperforming a nadir-pointing heuristic. This is a significant result that will aid in the study of dust devils science on the surface of Mars, in preparation for future crewed missions.

REFERENCES

- R. J. Lillis, D. Mitchell, L. Montabone, N. G. Heavens, T. Harrison, C. M. Stuurman, S. D. Guzewich, S. England, P. Withers, M. Chaffin, *et al.*, "Mars Orbiters for Surface-Atmosphere-Ionosphere Connections (MOSAIC)," *AGU Fall Meeting 2020*, AGU, 2020.
- [2] M. Inc., "ESPA: The Evolved Secondary Payload Adapter," MOOG Data Sheet, 2020.
- [3] J. C. Breidenthal, C. D. Edwards, E. Greenberg, G. J. Kazz, and G. K. Noreen, "End-to-end information system concept for the Mars Telecommunications Orbiter," 2006 IEEE Aerospace Conference, IEEE, 2006, pp. 13–pp.
- [4] J. Breidenthal, H. Xie, C.-W. Lau, and B. MacNeal, "Space and earth terminal sizing for future mars missions," 2018 SpaceOps Conference, 2018, p. 2426.
- [5] T. S. L'Ecuyer and J. H. Jiang, "Touring the atmosphere aboard the A-Train," *Phys. Today*, Vol. 63, No. 7, 2010, pp. 36–41.

- [6] S. Chien, B. Cichy, A. Davies, D. Tran, G. Rabideau, R. Castano, R. Sherwood, D. Mandl, S. Frye, S. Shulman, et al., "An autonomous earth-observing sensorweb," *IEEE Intelligent Systems*, Vol. 20, No. 3, 2005, pp. 16–24.
- [7] M. C. Malin, J. F. Bell, B. A. Cantor, M. A. Caplinger, W. M. Calvin, R. T. Clancy, K. S. Edgett, L. Edwards, R. M. Haberle, P. B. James, *et al.*, "Context camera investigation on board the Mars Reconnaissance Orbiter," *Journal of Geophysical Research: Planets*, Vol. 112, No. E5, 2007.
- [8] H. Kahanaa, M. T. Lemmon, D. Reiss, J. Raack, E. Mason, and M. Battalio, "Martian Dust Devils Observed Simultaneously by Imaging and by Meteorological Measurements," 49th Lunar and Planetary Science Conference, No. 2083, 2018.
- [9] J. Vander Hook, J. Castillo-Rogez, R. Doyle, T. Stegun-Vaquero, T. M. Hare, R. L. Kirk, D. Bekker, A. Cocoros, and V. Fox, "Nebulae: A Proposed Concept of Operation for Deep Space Computing Clouds," *Proceedings of the IEEE Aerospace Conference*, 2020. Accepted.
- [10] S. Chien, A. Jonsson, and R. Knight, "Automated planning & scheduling for space mission operations," 2005.
- [11] S. Chien, G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, *et al.*, "ASPEN-Automating space mission operations using automated planning and scheduling," *SpaceOps 2000*, 2000.
- [12] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davis, D. Mandl, S. Frye, B. Trout, et al., "Using autonomy flight software to improve science return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication*, Vol. 2, No. 4, 2005, pp. 196–216.
- [13] S. Chien, B. Cichy, A. Davies, D. Tran, G. Rabideau, R. Castano, R. Sherwood, D. Mandl, S. Frye, S. Shulman, *et al.*, "An autonomous earth-observing sensorweb," *IEEE Intelligent Systems*, Vol. 20, No. 3, 2005, pp. 16–24.
- [14] Y. Gao and S. Chien, "Review on space robotics: Toward top-level science through space exploration," *Science Robotics*, Vol. 2, No. 7, 2017.
- [15] R. Amini, S. Chien, L. Fesq, J. Frank, K. Kolcio, B. Mennsesson, S. Seager, and R. Street, "Enabling and Enhancing Astrophysical Observations with Autonomous Systems," arXiv preprint arXiv:2009.07361, 2020.
- [16] A. Fukunaga, G. Rabideau, S. Chien, and D. Yan, "Aspen: A framework for automated planning and scheduling of spacecraft control and operations," *Proc. International Symposium on AI, Robotics and Automation in Space*, 1997.
- [17] S. Chien, B. Smith, G. Rabideau, N. Muscettola, and K. Rajan, "Automated planning and scheduling for goal-based autonomous spacecraft," *IEEE Intelligent Systems and their applications*, Vol. 13, No. 5, 1998, pp. 50–55.
- [18] T. Estlin, A. Gray, T. Mann, G. Rabideau, R. Castano, S. A. Chien, and E. Mjolsness, "An integrated system for multi-rover scientific exploration," *AAAI/IAAI*, 1999, pp. 613–620.
- [19] S. Nag, "Sensor Webs of Agile, Small Satellite Constellations and Unmanned Aerial Vehicles with Satellite-to-Air Communication Links," *1st IAA Latin American Symposium on Small Satellites*, 2017.
- [20] S. Nag, M. Sanchez Net, A. Li, and V. Ravindra, "Designing a Disruption Tolerant Network for Reactive Spacecraft Constellations," ASCEND 2020, p. 4009, 2020.
- [21] S. Nag, C. K. Gatebe, and O. d. Weck, "Observing system simulations for small satellite formations estimating bidirectional reflectance," *International Journal of Applied Earth Observation and Geoinformation*, Vol. 43, 2015, pp. 102–118.
- [22] S. Nag, A. Aguilar, R. Akbar, A. Azemati, J. Frank, R. Levinson, A. Li, M. Moghaddam, V. Ravindra, and D. Selva, "D-SHIELD: Distributed Spacecraft with Heuristic Intelligence to Enable Logistical Decisions," *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*, 2020.
- [23] A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, and S. Chien, "Automatic detection of dust devils and clouds on Mars," *Machine Vision and Applications*, Vol. 19, No. 5-6, 2008, pp. 467–482.
- [24] T. A. Estlin, B. J. Bornstein, D. M. Gaines, R. C. Anderson, D. R. Thompson, M. Burl, R. Castaño, and M. Judd, "Aegis automated science targeting for the mer opportunity rover," ACM Transactions on Intelligent Systems and Technology (TIST), Vol. 3, No. 3, 2012, pp. 1–19.
- [25] S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Integrated planning and execution for autonomous spacecraft," *1999 IEEE Aerospace Conference. Proceedings (Cat. No. 99TH8403)*, Vol. 1, IEEE, 1999, pp. 263–271.
- [26] B. Jackson, R. Lorenz, and K. Davis, "A Frame Work for Relating the Structures and Recovery Statistics in Pressure Time-Series Surveys for Dust Devils," *Icarus*, Vol. 299, 2018, pp. 166–174.

- [27] D.-H. Cho, W.-S. Choi, M.-K. Kim, J.-H. Kim, E. Sim, and H.-D. Kim, "High-Resolution Image Video CueSat (HiREV): Development of Space Telescope Test Platform Using a Low-Cost CubbeSat Platform," *International Journal of Aerospace Engineering*, 2019.
- [28] M. L. Tibbs, "Design and Test of an Attitude Determination and Control System for a 6U Cubesat using AFIT's Cubesat Testbed," *Department of the Air Force Air University, Air Force Institute of Technology,* Wright-Patterson Air Force Base, Ohio, 2015.

ACKNOWLEDGEMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).