

Nebulae: A Proposed Concept of Operation for Deep Space Computing Clouds

Joshua Vander Hook*, Julie Castillo-Rogez
Richard Doyle, Tiago Stegun Vaquero
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA, USA
hook@jpl.nasa.gov

Trent M. Hare, Randolph L. Kirk
Astrogeology Science Center
U.S. Geological Survey
Flagstaff, AZ, USA
thare@usgs.gov

Valerie Fox
University of Minnesota
Minneapolis, MN, USA
vfox@umn.edu

Dmitriy Bekker, Alice Cocoros
Johns Hopkins University
Applied Physics Laboratory
Laurel, MD, USA
Dmitriy.Bekker@jhuapl.edu
Alice.Cocoros@jhuapl.edu

Abstract—In this paper, we describe an ongoing multi-institution study in using emplaced computational resources such as high-volume storage and fast processing to enable instruments to gather and store much more data than would normally be possible, even if it cannot be downlinked to Earth in any reasonable time.

The primary focus of the study is designing science pipelines for on-site summarization, archival for future downlink, and multi-sensor fusion. A secondary focus is on providing support for increasingly autonomous systems, including mapping, planning, and multi-platform collaboration. Key to both of these concepts is treating the spacecraft not as an autonomous agent but as an interactive batch processor, which allows us to avoid “quantum leaps” in machine intelligence required to realize the concepts.

Our goal is to discuss preliminary results and technical directions for the community, and identify promising new opportunities for multi-sensor fusion with the help of planetary researchers.

success. However, a fundamental imbalance exists between our ability to gather observations from instruments, and our ability to get those observations back to Earth. Dominated by a $\frac{1}{r^2}$ relationship, downlink bandwidth decreases quickly as a spacecraft moves further out into the solar system, yet instrument capability does not. This is the second tyrannical equation for space exploration, behind the well-known rocket equation, and is the perennial issue with scientific discovery in the solar system: We are here on Earth, and we need data about what is “out there” to inform our instrument targeting, design our mission objectives, and in general make our discoveries. Other factors notwithstanding, this data-hungry process contributes to a low cadence of missions to the outer planets and ocean worlds, and is exacerbated by a lower per-year return of data than, for example an Earth, Moon, or Mars mission.

This was the starting point for a Keck Institute for Space Studies (KISS) workshop. The goal of the workshop was to increase the “science bang per downlinked bit” by looking across the spectrum of technological advancements in autonomy, onboard processing, and operations. Our findings are preliminary and necessarily high-level. Ongoing studies are preparing for a second workshop. However, in the interest of engaging the community early, we provide this summary.

The workshop focused mainly on the use of a spacecraft *not only* as a remotely operated instruments but also as remote data storage and processing asset. The key insight is that most science inquiry on Earth is already completed by interacting with data stored on a remote server (e.g., the Planetary Data System or PDS [1]), and a server placed near the planetary body of interest would have in its banks 100s or even 1000s of times more data. Furthermore, the “inquiries” themselves are algorithms, and the “answers” can potentially also be highly compressible. The challenge is to improve by orders of magnitude the available processing and data storage on spacecraft, but the benefits are potentially revolutionary, even before autonomous operations are considered. As a general-purpose, shared resource for computing, the concept resembles so-called “cloud computing” and yet is emplaced in deep space; thus, we call it a *Nebula*. We present three discussions.

First, we discuss some of the relevant state-of-the-art onboard processing techniques that operate mostly on “raw” sensor data, or at least close to the instrument. We find that these

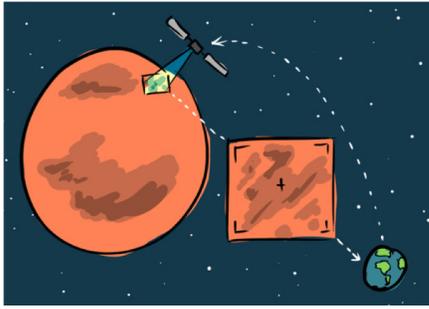
TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. ADVANCED SUMMARIZATION	2
3. GROUND SCIENCE DATA PROCESSING	3
4. NEBULAE AS A REMOTE SCIENCE DATA SERVER ..	5
5. SCIENCE IN THE AGE OF NEBULAE.....	7
6. ENGINEERING AND OPERATIONS SUPPORT	8
7. CONCLUSION	11
REFERENCES	12

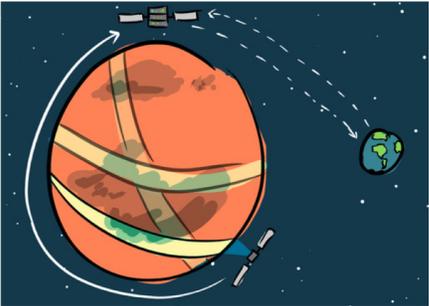
1. INTRODUCTION

In this paper, we propose a change to the conceptual design of a spacecraft from a remotely (or autonomously) operated instrument that streams data back to Earth to a remote compute and data storage node, which may also include data acquisition capabilities. Space systems consisting of instruments remotely operated through the Deep Space Network (DSN) have brought incredible discoveries. The paradigm is a great

*Corresponding Author
978-1-7281-2734-7/20/\$31.00 ©2020 IEEE. Government Sponsorship Acknowledged.



(a) Traditional spacecraft are generally treated as remote-targeted instruments, and observations are constrained to fit in downlink budgets.



(b) We propose instruments should gather much more data for aggregate analysis or prioritized downlink over time.

Figure 1: A quick conceptual drawing of the Nebula concept. Traditionally (1a), the spacecraft is commanded to take observations, and those observations are returned to the ground system for analysis (or sometimes analyzed onboard). Instead (1b), an instrument in a remote location should be allowed to gather many more observations than can be downlinked, and the *dataset* is what the ground team interacts with. The dataset is shown here on a separate spacecraft (i.e., a remote science data station) but this is not required.

technologies, while designed for highly constrained spacecraft with limited to no onboard storage, can be deployed on a Nebula to summarize a massive dataset and provide nondestructive summaries of data. Moreover, such a dataset, which is beyond reach of Earth-bound scientists because of downlink restrictions, can still be studied using any traditional or newly emerging algorithm *nondestructively*. That is, one can re-summarize the raw data if one does not value the result, unlike the highly constrained systems that cannot store data and drove the development of summarization techniques in the first place.

In contrast to the image-based (or sensor-based) data processing, ground-side analysis traditionally involves a long toolchain of algorithms which produce clean, well-rectified images for general purpose analysis. Well-understood summarization techniques are used to guide searches and produce human-digestible outputs based on this cleaned-up imagery. We discuss this in Section 3. The engineering challenges associated with porting this rectification toolchain to a spacecraft are touched upon.

The crux of the Nebula concept is the new opportunities accruing from deploying *both* advanced sensor-space preprocessing algorithms from Section 2 *alongside* traditional ground-side data processing pipelines of Section 3 *onboard*

the spacecraft. This processing is optional. A large cache of imagery can and should still be downlinked as early as possible for terrestrial analysis. However, the onboard processing is hugely enhancing because the onboard data can be used to respond quickly to downlink requests and provide context for observations, and is enabling for population-level studies. Examples of this mode of operation are discussed in Section 4.

The paper ends with a discussion of improvements to the operations of *nearby* assets, including semi-autonomous spacecraft, enabled by on-site computing infrastructure. We present case studies for the Mars 2020 rover and Dragonfly lander missions, as well as multi-agent concepts for Lunar, Mars, and small-body exploration in Section 6.

2. ADVANCED SUMMARIZATION

The major point of the KISS workshop was the increasing importance of onboard summarization techniques—a capability well-advocated and studied over the last decade [2]. Traditionally, these techniques are used to offset the problem of downlink-constrained scientific discovery. That is, these techniques are deployed when either the spacecraft cannot downlink what the scientists wish to observe or the science team cannot react quickly enough to capture and interpret transient phenomena.

Predominately this is done using onboard software to help prioritize from among available data for downlink, or even to direct the spacecraft’s sensors or actions. In this case, the downlinked data can be either a raw observation or a small thumbnail (image “chip”) that the spacecraft chose according to some human-directed ruleset. It may also be a “feature” vector, or description that is not necessarily what the sensor observed. A major common theme in this research is that *all* rely on image-space feature detection (e.g., pattern matching in non-rectified or “raw” data), and the raw data itself is *rarely* available for reanalysis or downlink to safe storage on Earth. This can be a barrier to acceptance by the scientific community.

Simple cases of summarization techniques are already in practice, as the scientific imaging cameras onboard the Mars Science Laboratory return thumbnail images prior to downlinking the full resolution images. This allows scientists to select the best focus or most scientifically useful frames for prioritized downlink, with lower-value data products are downlinked when time is available or deleted [3].

We now arrive at the major departing point of the study’s result from previous literature. We press the vital importance of contextual information, accurate georeferencing, and, most critically, *backup* of raw sensor data for later reanalysis. Without these factors, it has become clear that the scientific community is reluctant to “trust” onboard automated analysis. Despite a decade or more of work, this trust issue is pervasive enough that the often-quoted mantra of “Scientists want the raw data” is well-known to researchers in this area, and mission proposals seem to be predominantly scoped down to the “remote instrument” mode of spacecraft operations.

3. GROUND SCIENCE DATA PROCESSING

Where the previous section could be thought of as treating a spacecraft as a remote instrument, possibly with some advanced sensor processing algorithms embedded in it, this section discusses ground-side analysis techniques that could be deployed to a remote spacecraft.

The motivation to move to the notional Nebula spacecraft is that if a particular well-understood analysis technique, algorithm, or program can produce interesting discoveries when run on downlinked data, it may be more useful when run on significantly more data, even if that data is on the spacecraft itself. This is not straightforward since analysis of imagery and data involves a complex series of steps, and each of the steps is itself an active area of research. Merely calibrating an instrument to produce useful output data is challenging enough, let alone calibrating a series of instruments, registering the measurements, fusing sensibly, and geolocating the points of interest. Indeed, the major reason, as stated, that “scientists want the raw data” is to *enable* the research of new analysis techniques on downlinked data, particularly those techniques involving multi-sensor fusion.

Background

Any data gathered by space-borne instruments are effectively useless without adequate supporting data. Two main areas must be captured: (1) methods to characterize the instrument and (2) the spacecraft location.

First, prior to launch, methods to characterize each camera or instrument, including its internal geometrical properties (e.g., focal length) and its quantitative radiometric properties, must be known. This metadata describing the instrument is generally called the camera model or more generically the sensor model. It should be noted that this is not a static model, as an instrument sensor model should be updated throughout the life of a mission perhaps simply due to the age of an instrument or other factors impacting an instrument during flight.

Second, the relation between the primary observation and the planetary body being “imaged” depends on the spacecraft position and even more sensitively on instrument pointing (for all but radar images). The most important metadata are therefore the time histories of attitude (pointing) and trajectory (position), generally captured in an information model called “SPICE” (Spacecraft, Planet, Instrument, Camera-matrix, Events) [5]. It is “the NASA Planetary Science Divisions method of conveniently packaging, archiving, and subsequently accessing observation geometry needed to understand science data returned from robotic spacecraft.” [6].

SPICE is commonly provided in multiple versions, starting with predicted values (based on commanded pointing and forward integration in time of the trajectory) and eventually replaced with reconstructed pointing based on analyses of attitude control system data, ground tracking, and the images and altimetric data themselves. It is not uncommon when relying on these predicted positions that systematic position errors of many kilometers can occur. Therefore, it is not uncommon that reconstructed SPICE is eventually generated for the observations, but usually after the mission has completed. Lastly, it is worth noting that the eventual positional accuracy of the observation not only depends on (1) the sensor model and (2) the ability to measure the position of the spacecraft and the pointing of the instrument, but also knowledge of the shape and rotation of the target body [4].

Processing of the digital planetary image data is generally described using a multi-level process that is conducted in several stages, as follows. Note that these “levels” are as described by Batson (1990), but a variety of other systems for indicating relative levels of data processing are also in use by NASA and by specific instrument teams, and they must be distinguished.

- Level 0 generally refers to the raw data records and incorporation of spacecraft, planet, instrument, and camera positional information and navigation data (i.e., SPICE).
- Level 1 consists of radiometric correction involving correction for bad pixels, dark current, shutter effects, flat fields, and conversion of raw pixel values to intensity units (i.e., W/m²/sr or W/m²/sr/nm).
- Level 2 consists of geometric processing, generally beginning with the establishment of geometric control, hopefully tied to a single established control-base image, and calculation and updating of camera-pointing angles by bundle-block adjustment. This is followed by subpixel-level co-registration and reprojection to planimetric or panoramic views. Topographic modeling is also part of this step. At this level, the observations are spatially tied to the surface.
- Level 3 consists of first-order photometric correction, including normalization of scene brightness and removal of residual frame boundaries. Photometric corrections are best applied when the spatial location has been accurately solved.
- Level 4 consists of mosaicking of individual observations, possibly followed by cosmetic enhancement or the addition of other annotations.

Software to process digital planetary image data has been evolving for decades. Two well-known applications that proved this capability include the U.S. Geological Surveys Integrated Software for Imagers and Spectrometers (ISIS3) [4] and JPLs Video Image Communication And Retrieval (VICAR) [7]. Both ISIS3 and VICAR, as photogrammetric software, can be used to improve position, pointing, and planetary parameters based on measurements of matching ground features where images overlap. This process of geodetically controlling the observations greatly helps to reduce positional errors, but it cannot eliminate them entirely. In fact, the lack of ground control for most planets (laser altimetry data now provides an effective substitute for Mars, the Moon, and a few other bodies) leads to much greater errors than are encountered in mapping Earth. And the initial lack of topographic data for each planet or satellite introduces further positional errors. For bodies that do have topographic data, images are now routinely orthorectified. Orthorectification spatially improves, or planimetrically corrects, the observation by removing the effects of image perspective (tilt) and relief effects of the surface topography.

Moving Onboard

Given the background section above, while some challenges must be overcome, it is realistic that the existing processing pipeline could be transitioned to space-borne platforms. To realize this, several pipeline hurdles must be addressed.

First the predictive positional information (SPICE) must be captured and derived into a useful state without the need to bring it down to Earth.

The next step would be to support the photogrammetric software needed to step through the various processing levels described above. This means supporting a well vetted and tested sensor model, the ability to apply the predicted pointing (SPICE) to each observation, the ability to tie (image-

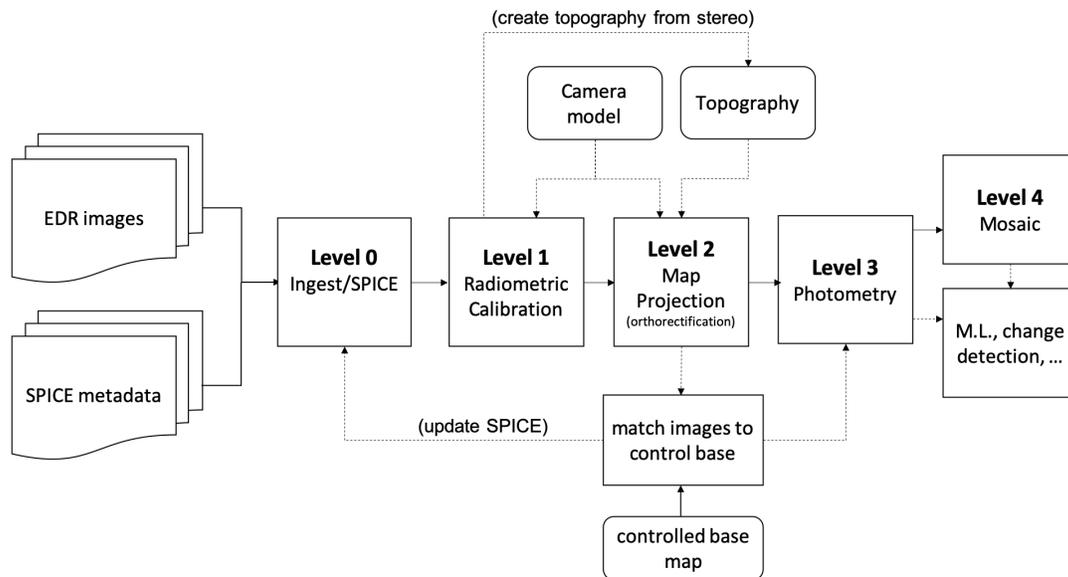


Figure 2: A generalized processing workflow for ISIS3 [4], which supports downstream analysis and is a critical component for a future Nebula-enabled remote science station.

match) the newly taken observations to an existing onboard controlled base, and lastly the ability to orthorectify the observation through an existing topographic layer or the ability for the observation to generate its own topography.

For bodies that do not have an existing geodetically controlled base map, this pipeline becomes extremely difficult. And while, in theory, creating a fully controlled and photogrammetrically bundle-adjusted global solution across images might be possible, it generally takes a lot of manual human intervention at all stages.

For bodies like Mars, the Moon, and Europa, where a controlled base map has been created, this pipeline can exist, largely intact, if the ability to automatically tie (image-match) the observed data to the control base is mostly automated. Automation of these processes has benefits even to Earth-side data stores, and is therefore a good continuing research direction.

Once the data is adequately prepared, it is common to use ground-side summarization techniques to interrogate data. The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [8] is a good example where such summarization is already used on the ground and could be adapted for onboard production. CRISM is a 545 band hyperspectral visible to near infrared (VNIR) pushbroom spectrometer onboard the Mars Reconnaissance Orbiter (MRO) [8]. The instrument has two detectors; the shortwave (S) detector operates from 0.35-1.04 μm over 107 bands, and the longwave (L) detector operates from 1.01–3.91 μm over 438 bands. The spectral sampling is 6 nm, allowing for the detection of narrow spectral absorptions diagnostic of key mineralogy. CRISM takes targeted images with spatial resolutions between 9 and 18 m/pixel (4.3 Gb), as well as lower resolution mapping strips (1.8-5.7 Gb). The combination of high spectral and spatial resolution makes CRISM observations an important dataset for interpreting the surface mineralogy of Mars.

The CRISM team produces mineral "parameter maps" that are summarizations of the spectral characteristics present

within an observation [9], and reduce the 545 wavelength image cube to a 28 band browse cube that contains information about major mineral classes that may or may not be present in a given observation, and are therefore a fast way to prioritize analysis of the observation. Minerals are classified on the basis of characteristic spectral absorptions at known wavelengths, and the parameter maps calculate the strength, or band depth, of a given spectral feature for each pixel in the scene to determine the likelihood of a mineral's presence. These parameter maps are generated on radiometrically calibrated data products (usually I/F) that may or may not be map projected. Therefore, onboard processing would need to autonomously generate the level 1 data records, perform some basic de-stripping, filtering, and atmospheric correction, and calculate the parameter values to generate the summary products. The current pipeline is a mix of autonomous functions and significant human intervention, and in particular noise remediation and atmospheric corrections benefit enormously from scene-to-scene customization and expert evaluation. However, a rapidly generated browse product is useful in selecting scenes that are the most valuable for time intensive processing and scientific analysis.

Summary

Taken together, these are examples of key enabling analysis techniques that are well-understood, commonly used by the community, and could be applied to large datasets that are gathered by a spacecraft.

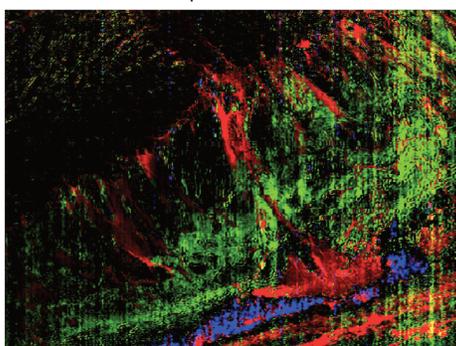
One benefit of correct orthorectification is that each image-level data product or analysis output can now be precisely geolocated on the planet. Another is the possibility of onboard multi-sensor fusion. However, orthorectification requires significant *a priori* data in the form of ground control networks of feature points. In the event of a mapping mission to a new body, this would have to be built using data the spacecraft itself gathers, rather than data available on the ground.

The benefit to scientific discovery is potentially immense, as

A. FRT0000B6F1 Near IR Albedo



B. FRT0000B6F1 Spectral Feature Parameter Map



C. Spectral Retrievals

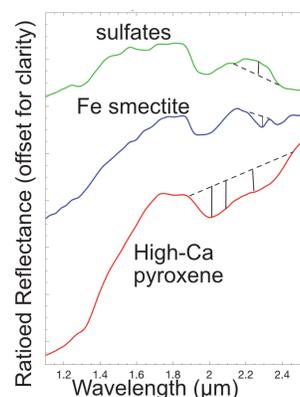


Figure 3: A. CRISM observation FRT0000B6F1 displayed in sensor space prior to georegistration.. B. Example of a set of parameter maps calculated from the sensor space imaging. Red: “HCPindex”, tuned to detect a broad absorption around 2 m characteristic of pyroxene. Green - “SIndex”, sensitive to hydrated sulfates by calculating the convexity at 2.3m caused by paired absorptions at 2.1 and 2.4m. Blue - “BD2900” - the band depth at 2.9 m, characteristic of Fe-clay minerals. C. Example spectra of each class of minerals, and sketches of the calculations made to parameterize them. The full spectral data is required for analysis, but the diversity of minerals as detected in the summary parameters provides a rapid estimate of the spectral diversity in the observation.

well vetted techniques and analysis algorithms can be run on a large, well registered and clean dataset, albeit a *remote* dataset. At Mars alone, MRO’s HiRise camera has imaged *and downlinked* less than 5% of the Martian surface, despite covering Mars every few months for the last decade. If there is one takeaway that nicely summarizes our concept it is: *Leave the camera on and keep your data onboard for later analysis.*

4. NEBULAE AS A REMOTE SCIENCE DATA SERVER

This section expounds the main concept of the paper—that of computing infrastructure. Given the myriad onboard analytics techniques emerging or already developed (Section 2), which can quickly comb data streams for interesting features, and given the robust ground-side analysis pipelines, which can clean data for general purpose use (Section 3), we propose adding sufficient computation and storage to future spacecraft to enable a remote dataset to be built up, summarized, and made available for analysis without impacting the “traditional” downlink activities.

Taking for granted the capabilities in the previous two sections, we have a spacecraft capable of rectifying and registering data. This is of marginal utility by itself if the spacecraft is still taking only targeted observations and immediately downlinking the results—that is, if we only design and operate the instrument to observe what we can downlink. Instead, we assume this is not the case, and instruments are “left on” to gather more data than what is going to be downlinked.

As a case study, consider a serendipitously-discovered Recurring Slope Lineae [10], i.e., Figure 4. Upon discovery and recognition—almost certainly well after first imaged by the spacecraft—a science team could request downlink of the *historical* observations going back years and over multiple seasons. Because (by assumption) the spacecraft has good registration capability and large onboard storage capacity, this is a trivial operation. It is mostly “infrastructure”, and is

entirely non-destructive of the “raw” data, which is kept for later downlink if required (or removed at human request to make room as required). If needed, the images could still be downlinked for reproducibility purposes or to tune or develop the registration routines.

This “query/response” mechanism is key to the concept of *Reactive Science*. The spacecraft has, and will maintain, a larger set of data than can ever be downlinked. While we can, and should, continue to downlink as much of the “raw” data as is possible, a time will come when a scientist wants to know what the rest of the data contains. This is where the discussion of lightweight analysis capabilities from Section 2 comes in. Each of these algorithms could be deployed as a targeted “query” (or, in computer science terms, a “daemon”) that can rove through the data and call out features that may have been missed. Because the raw data is preserved, re-tuning of these “science daemons” is possible, and risk of mistakes is minimized because a small set of verification data can be downlinked to verify the processing steps or the results.

Now consider the addition of onboard rectification and data management from Section 3. With well aligned historical datasets, science daemons can compare current output to past observations directly. It becomes possible to automatically build up a population-level time series of observations for the whole planet.

As a quick approximate analysis, it is possible, with onboard storage of “only” a gigabit, to track approximately a million features’ positions and identifiers across four years of annual surveys. On Venus, this could be periodic check-ins of potentially volcanically active areas. On Mars, fresh impacts or RSLs. This continues, all without loss of data and with human verification as desired, directed, or required during mission operations.

This interactivity concept is illustrated in Figure 5. Initially, the science team has only prior measurements. As the spacecraft takes and accumulates observations, they prioritize

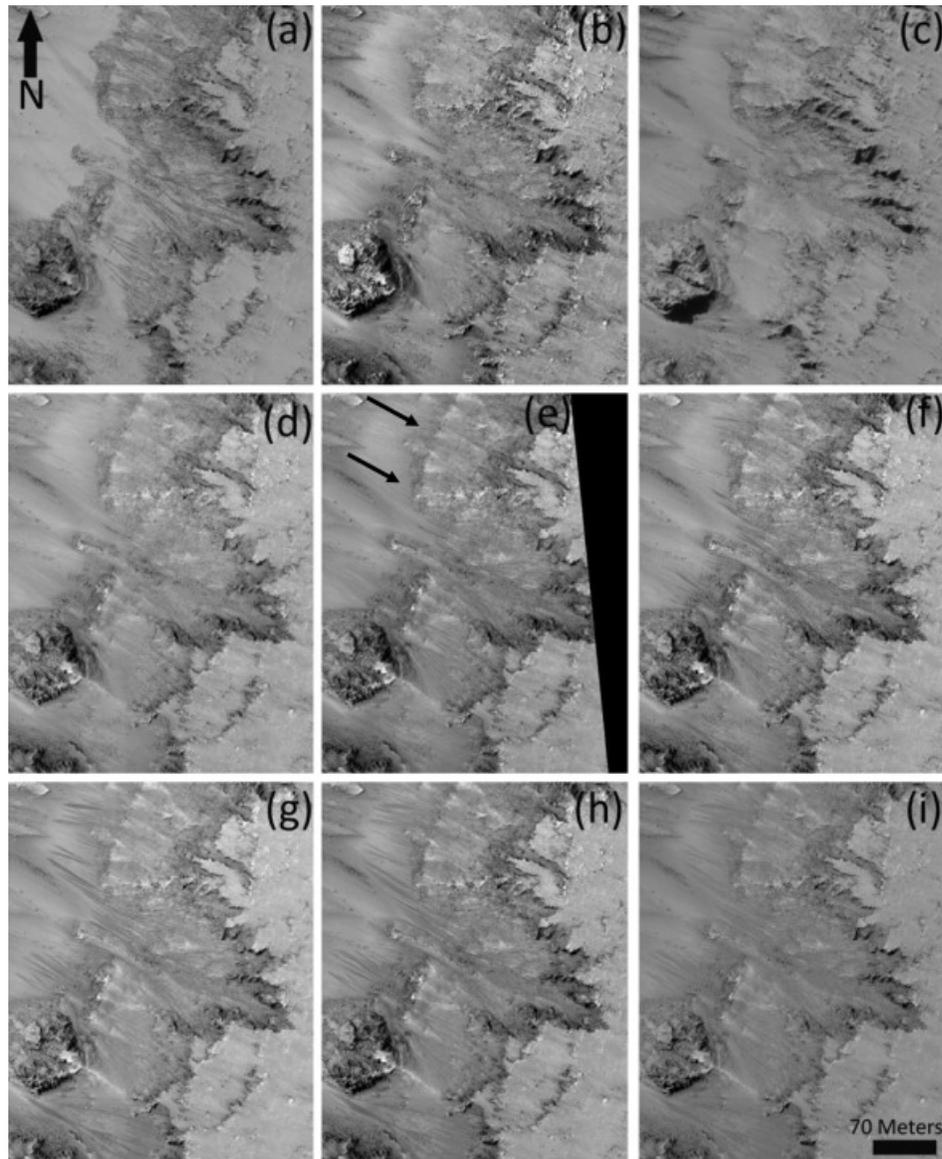


Figure 4: Historical analysis and image change tracking is enabled by precisely rectified onboard image stores. This figure, reproduced from and annotated by [11] represents three years of HiRISE targeted observations for the study of Recurring Slope Lineae. We envision capture of historical datasets for later downlink or human-cued onboard analysis after unexpected discoveries, potentially saving years of time before publications. Onboard analysis or change detection is enhancing of this concept, especially in helping scientists understand what is stored on the spacecraft. Continuous observation, large-scale storage, and some level of “uploaded” (human-tuned) registration or rectification is enabling.

downlink and ground-side processing per the usual mission operations. Each downlink may come with tiny summaries of the rest of the data collected by the prototype onboard analysis. The downlinked data can be compared to the summarization to validate the output of onboard analysis as part of normal operations. As time goes on, sufficient data is accumulated on the ground to enable development of more trusted summarization techniques and the mix of raw vs. pre-processed data will shift.

The capabilities suggested do require technological leaps from current spaceflight computing. But, we are optimistic this will come, since a modern commercial solid state disk, for example, has been released at the time of this writing that can store 100 Terabytes. This would store a 50 meter survey

of Mars across 32 spectral bands annually for 16 years.

A final stage of the evolution of this concept is the *Remote Science Data Station*. A years-long observation campaign can produce prodigious amounts of data. By storing as much as possible, a revisiting of priorities (i.e., new “science daemons”) is as simple as unlinking a new piece of software. A global study of weather patterns, distribution of geysers and frequency of eruptions, or even a global survey of landing sites can be accomplished by refining a procedure on the ground using whatever data is available, and uploading that procedure to the spacecraft for use on the significantly larger amount of data available.

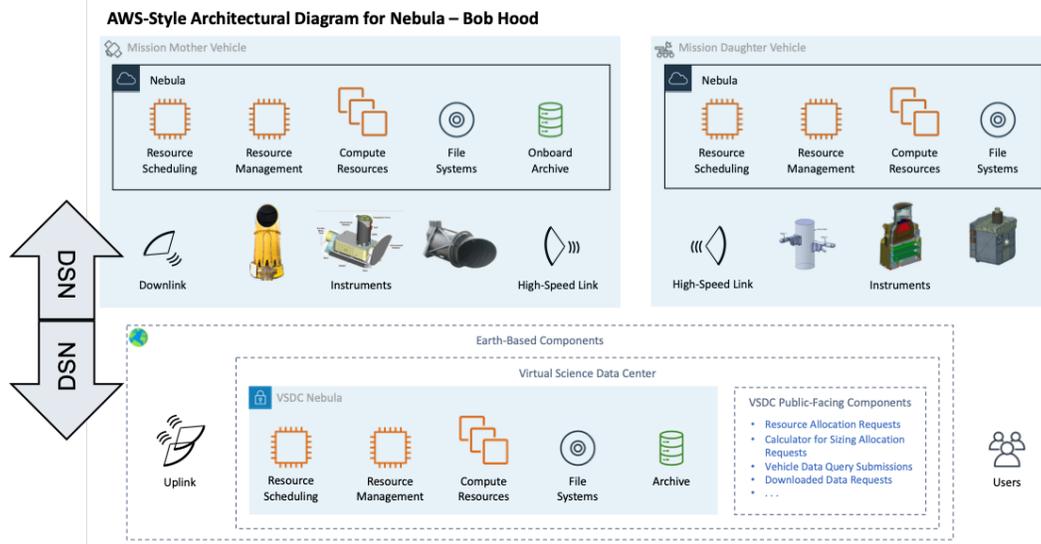


Figure 5: An example interactivity pipeline. The top layer represents the major technological pieces of a server and an instrument client. The bottom represents the ground-side toolkit, including tools for analysis of downlinked imagery and also testbed to validate and verify the spacecraft’s ongoing analysis “daemons”.

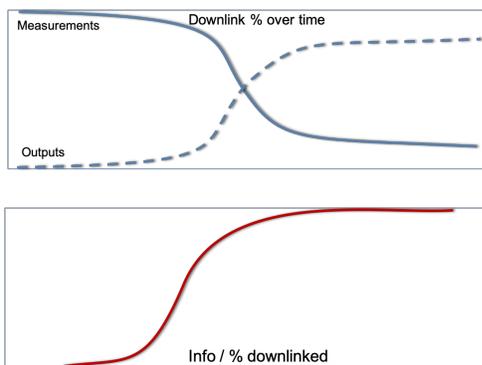


Figure 6: A notional mission downlink prioritization over time. On a long-running Nebula-like mission, the amount of data gathered far exceeds that which can be downlinked. Initially, only the raw or mildly processed sensor data will be downlinked, similar to today’s operation. However, onboard processing enables downlinking of “syntheses” or summaries of the data, and even allows science inquiries and change tracking over time. Since these outputs from onboard processed inquiries can be more informative, the outputs from onboard experiments are hypothesized to dominate as time goes on.

5. SCIENCE IN THE AGE OF NEBULAE

An exciting potential advantage of the Nebulae concept applies in the outer solar system, where the cost and frequency of access are highly constrained. These realities can lead to the frustrating circumstance whereby a mapping mission to say, the Jupiter or Saturn system, reveals new science questions based on inevitable discoveries, but the opportunity to design and deploy new observing and measurement capability to pursue those science questions may be decades away.

Under the Nebulae concept, a mapping mission would arrive with deep resources for computing, data storage and net-

worked communications. If also augmented with observing capability held in reserve, e.g., deployable daughter instrument payloads, this can lead to an unprecedented scenario in which a mission campaign can be accommodated within a single deployment opportunity at the remote planetary system. Expanding this concept further, in the first, mapping phase, the mission would collect a broad range of new observing data of the target system and, with the help of onboard analytics, generate initial results and insights. While the full set of collected data remains stored robustly in the outer solar system, selected intermediate results are downlinked for deeper interpretation by the scientific community. The results of that analysis can lead to specific measurement objectives and concepts, for a more focused, second phase of the mission.

The advantage is that the already-deployed capability can be commanded, perhaps reprogrammed to execute what amounts to a second mission at the same target, without the usual significant delays and new costs, not to mention the value accruing to scientists careers under such enabled efficiency of discovery and response.

The exploration of far outer solar system bodies, such as Centaurs, Kuiper Belt objects (KBOs), dwarf planets (e.g., Eris), long-period comets, and the moons of icy giant planets (e.g., Triton), offers compelling Nebula applications. First, these missions are in most cases expected to be fast flybys. Orbiting these bodies is theoretically possible but requires resources that can put these missions in highly competed cost classes. Instead, enabling the exploration of these bodies under low- or medium-cost (\leq \$1B) class programs could lead to very attractive missions at increasing pace. Centaurs and KBOs may also represent targets of opportunities for missions to the outer solar system, such as the icy giant flagship mission under consideration for the 2023-2032 planetary science decadal survey.

Encounters beyond Neptunes orbit set stringent telecommunication requirements with data rates on the order of 1 kbps. Here, we take a KBO flyby tour to illustrate the various

benefits offered by large storage capacity and computing capability.

Key science questions at KBOs include understanding the origins of the various sub-populations (e.g., scattered, cold classical) and their compositions as constraints on the temperature, fugacities, and redox conditions in the protoplanetary disk. In turn, that information can be extended to exoplanetary systems to map the distribution of volatiles and their relationship to habitable zones. Addressing these objectives requires searching for mineralogical and isotopic markers on the surfaces and potential outgassing activity at these bodies. This involves imaging in the ultraviolet and infrared, which are data intensive techniques. Visible imaging in stereo, color, and under multiple phase angles is also expected to yield gigabytes of data. As these bodies may in general be in thermal equilibrium, outgassing might be absent or limited. In order to collect volatiles for isotopic characterization, outgassing could be triggered by an impactor, similar to the Deep Impact mission. These observations are all challenging due to the relative velocity (10s km/s) between the spacecraft and target, the target low albedo and irregular shape (e.g., 2014 MU69), combined with the low luminosity at 20+ AU, and potentially unfavorable illumination angles at the time of the encounter.

Because of the difficulty in finding targets with ground-based telescopes, a spacecraft in the Kuiper-Belt may first be used as a local observatory. Regular survey of the sky to search for potentially accessible targets, including lightcurve observations, can be accomplished onboard. Potential targets accessible to the spacecraft may then be sent to the ground for assessment and follow-up with Earth-bound telescopes, along with key supporting data which remains onboard for re-analysis if required. On the way to a selected target, a Nebula-enabled spacecraft would use autonomous navigation, starting a few days inbound when the target shape is barely resolved. This would involve optimizing the flyby distance based on the reconnaissance of potential hazards, i.e., rings or companions, optimizing instrument settings, and prioritizing observations based on the target properties. Performing these activities with the ground in the loop is prohibited by the relatively short lead between the time at which the target is resolved and closest approach. For the same reason, sequencing of the impact deployment and observations of the resulting dust and volatile cloud would have to be planned onboard the mothership. During these critical operations, the spacecraft needs to be capable of self-health characterization and healing in order to avoid or get out of safing. During the flyby, data are acquired at high rates and directly sent to storage. Following the encounter, the Nebula interrogates the archive to extract and return representative data products to the ground. Scientists can query the archive during the cruise between targets. The archive data can then serve as a basis for informing follow on flybys.

6. ENGINEERING AND OPERATIONS SUPPORT

Given their infrequency, high cost, and stringent payload requirements, missions to the outer solar system could greatly benefit from nearby networked compute and storage elements that would enable data curating and intermediate data product generation. While direct-to-Earth communications from these distant missions is severely data rate limited, we envision a much higher bandwidth possible to exchange data with regional server-like resources.

For instance, a surface exploration spacecraft could transfer navigational and observational data from instruments such as optical cameras, radar, or LiDAR directly to the supporting orbiting compute resource, where the data could be analyzed to produce intermediate data products and more quickly inform scientists of follow-on observations. Such data products could also be used to support autonomous surface exploration by providing detailed maps of the planetary body to aid in daily exploration tasks.

For example, NASA's latest New Frontiers mission, Dragonfly [12], will be sending an octocopter to explore the surface of Titan. The mobility subsystem aboard the rotorcraft consists of flash LiDAR, radar, cameras, pressure altimeters, and IMU sensors, in addition to actuators, structure, electronics, and processing to support Terrain Relative Navigation (TRN) and Guidance, Navigation, and Control (GNC) [13]. Due to bandwidth, processing, and storage limitations, only some of the sensor data collected in flight is retained onboard or transmitted to Earth. If a nearby compute and storage system were available, it would enable the constant collection of sensor data, which could be analyzed and archived on this system.

In particular, Titan presently has very limited high-resolution maps, images, or digital elevation models (see Figure 7 for a current map of Titan, comprised of data taken during the Cassini mission). If this external system existed, optical and LiDAR data taken during Dragonfly's flights could be continuously saved to the system, allowing for the generation of a map that stretches over the full duration of the flight. These data products are costly, many Gigabytes of data could be generated per daily² 30 minute flights, but with the external system nearby, memory and data rates would place much more malleable limits on what can be saved. In addition, with the compute power of the external system, these maps could be further analyzed onboard the system and summarization products could be sent down to scientists on Earth to better inform future flights and selection of raw data for downlink to Earth.

Also, we see some examples of small, co-dependent systems included in flagship missions (MarCO, PUFFER, and MarsHeli). This is not surprising given current planetary surface exploration is limited to benign operating areas due to the inability to land in or traverse challenging terrain. Missions may leave behind the most interesting science opportunities which are often in hard-to-reach areas. Secondary mobile sensor platforms, in particular, are beneficial because they can investigate transient targets without endangering or delaying the primary mission timeline e.g., left behind to investigate the transient methane detection from MSL [14]. Similarly, cubesat-borne instruments could, for example, be deployed separate from a "mothership", which would allow them to independently target observations or optimize their orbits for particular science questions, without constraining the other instruments or platforms. In both cases, using a centralized data store and processing station (on ground or in orbit) allows these secondary instruments to "stream" data back to a server which can collate and register with other data, cache for later use, or downlink as required.

Constellation Mothership

Understanding the formation of small Solar System bodies holds promise to unlock the answers to a number of hitherto-unanswered questions concerning the formation of our Solar

²A Titan solar day is approximately 16 Earth days.

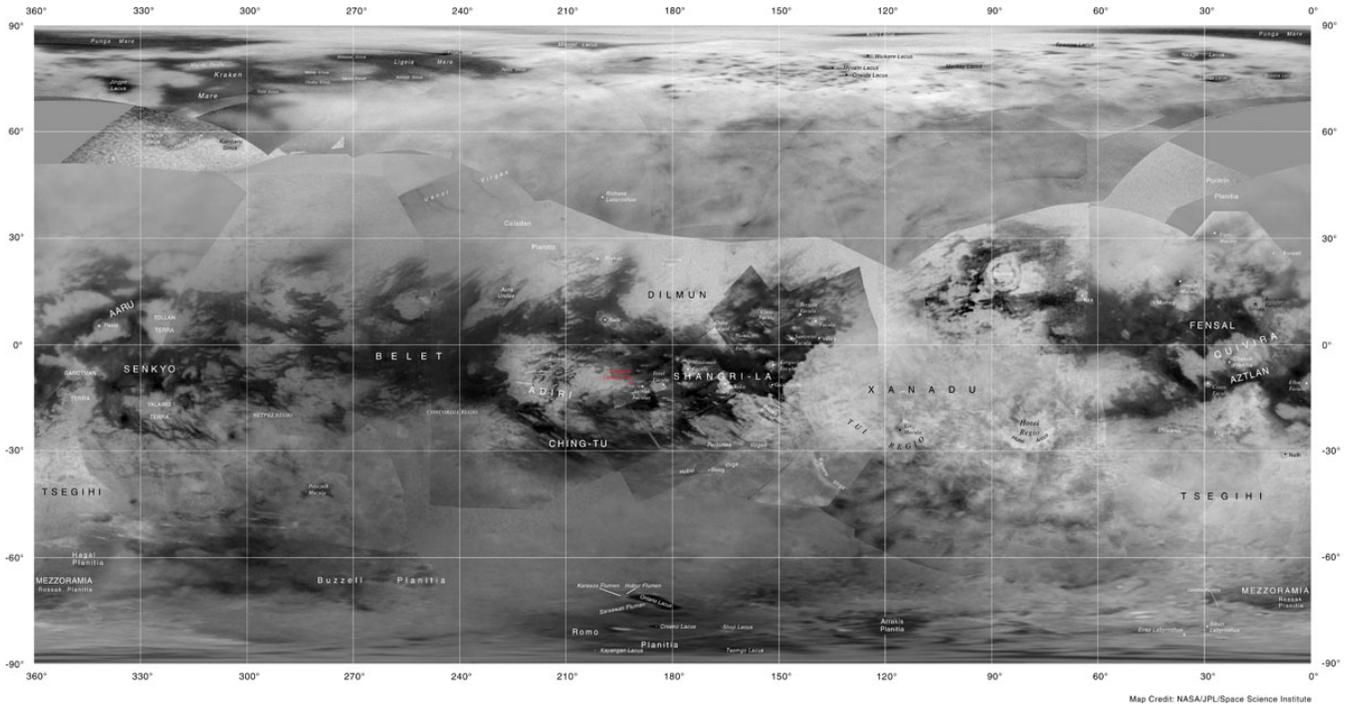


Figure 7: Cassini imaging coverage of Titan as of June 2015

System. Exploration of such bodies typically requires a multi-instrument approach, including cameras, imaging spectrometers, X-ray and gamma-ray spectrometers, laser and radar altimeters, and possibly sampling probes [15, 16]. The operations of such exploration missions are challenging due to the incompatible orbit requirements of certain instruments; for instance, a common requirement for an imaging instrument is to overfly its target at a sun angle of 45 degrees to achieve best illumination, whereas an X-Ray/gamma-ray spectrometer achieves best performance at a zero-degree sun angle [17]. This results in complex sets of orbits and long periods of idle time for individual instruments.

A promising concept of operations to overcome this difficulty is to deploy instruments with incompatible requirements on small, Cubesat-sized buses from a carrier spacecraft. In such a concept (shown in Figure 8)³, each instrument can be placed in an orbit suitable for its own science objectives, enabling faster science returns and removing the need to trade-off between different instruments' requirements.

However, such a concept of operations leaves open the problem of relaying data from individual instruments to scientists on Earth. Direct-to-Earth transmission generally requires large antennas - and carrying one antenna per instrument is mass-prohibitive and wasteful. In addition, the proposed concept of operations produces significantly more data per unit time compared to a monolithic spacecraft (by enabling multiple instruments to operate at the same time), whereas the downlink capacity of the Deep Space Network remains unchanged. Accordingly, a multi-spacecraft small body exploration mission would almost certainly rely on *local data aggregation* and *carrier-based relaying*, whereby each instrument relays data to the carrier spacecraft which, in

turns, acts as a local archive, and then transmits the data to Earth as bandwidth becomes available. In-space computing is an enabling technology to realize this vision. Namely, the carrier should:

- Orchestrate the collection of data and the relaying from individual sensors to the carrier spacecraft through delay-tolerant communication protocols such as DTN [18];
- Archive data from the science instruments and downlink it to Earth as bandwidth is available;
- Produce higher-level data products for scientists to downlink by merging observations from multiple sensors (while storing all data for future downlink), enabling scientists to quickly assess collected data and prioritize future observations;
- Select *communication-aware* orbits for the science instruments and the carrier spacecraft that maximize the amount of relevant scientific data collected and, critically, *relayed* to the carrier and to Earth, incorporating newly-collected information about the small body's shape and gravity field.

The mothership concept applies beyond multi-orbiter scenarios. Another concept in development involves networks of small rovers, one of which is known as PUFFER [19]. This conceptual Lunar or Mars mission (Figure 9) is based on a PUFFER network combined with a parent platform (e.g., base station or flagship rover) to accurately place a PUFFER's instrument (microimager) on a terrain feature. This operation occurs within the parent platform's direct communication and sensing line-of-sight (LoS). PUFFER must be capable of autonomously navigating the environment, homing in on the feature. It may leverage the better computation capabilities of the parent platform, as well as its sensors that offer a more advantageous perspective of the drive to improve its placement accuracy. Analysis completed in [20] showed that leveraging nearby computing, when available, could enable 3x more science processing (for capabilities like AEGIS [21]), by offloading engineering processes.

³Reviewer note: At the time of this writing a C.30 PMCS proposal has been accepted to study a similar concept, but this conops has not been cleared for publication at this time.

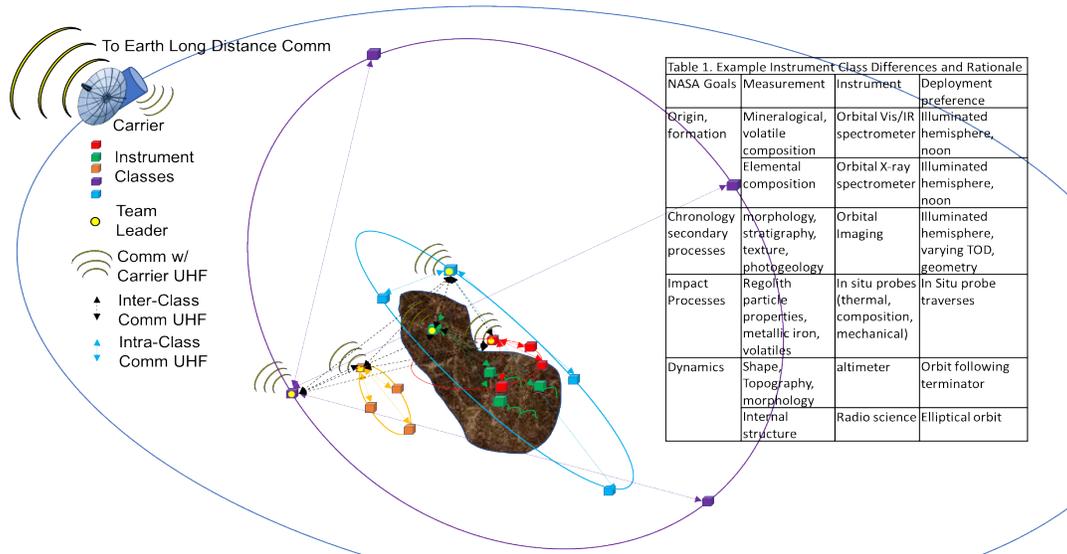


Figure 8: Design Reference Mission and proposed architecture for a distributed network of small spacecraft monitoring an asteroid.

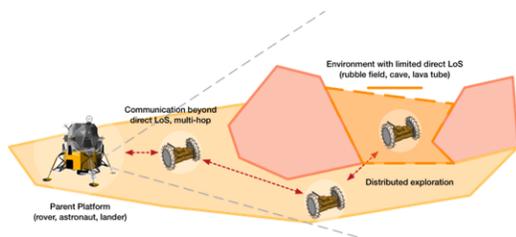


Figure 9: A lander is essentially a local server that can support small rovers.

Similarly, we analyzed the Mars 2020 mission. One defining feature of Mars Sample Return mission concepts is the likelihood of revisiting the same area with subsequent launches to fetch, retrieve, and launch the samples [22]. If an on-site computing asset were available to multiple rovers in the area, they could make use of it for off-loading their required engineering tasks, in order to take advantage of opportunistic science processing and sensing. Thus, the assisting asset(s) could provide an “infrastructure upgrade” and could remain on-site, providing communication, computation, and data analysis services for all subsequent phases of the campaign.

To explore any potential benefit, we next consider a strategic drive campaign by a Mars 2020 rover. The Mars 2020 conceptual path-planning pipeline, presented in [23], is simplified for our use in Figure 10a. The randomized time associated with Select Path is understandable given the mission analysis from [24]. This data used in simulation is adapted from [25].

We created the model software network for Mars 2020 illustrated in Figure 10b. The required tasks are constructed to model the timings given in Figure 10a. From [25], we also included the ability for the rover to use imagers to classify the terrain, but only as an optional algorithm, since the current Mars 2020 pipeline does not include it.

To model the terrain in our simulations, we use terrain data classified from HiRISE imagery from [25]. Multiple terrain

types are grouped into different classes or as obstacles (terrain that cannot be traversed). We do not currently take slope into account; therefore we model the velocity of a rover in a given terrain class based on the average speed over multiple slopes for that classification.

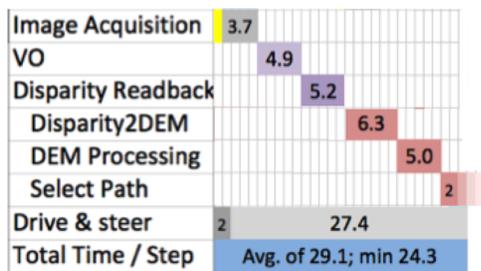
In order to model the different fidelity of data obtained in orbit and on the ground by the rover, we assume certain terrain types as unknown. When a rover is in an unknown terrain type, it will move at the velocity of the real terrain class; however, it will plan a path assuming a terrain with the fastest traverse velocity. Nevertheless, if a rover is able to perform terrain classification, we assume it will be able to correctly classify the terrain within a given radius.

Repeating the analysis of the software network produced the data shown in Table 11a. From this analysis, we isolated four operating regimes for the rover. In the first regime, the rover has no access to the assisting resource (regime 0). Regimes 1-4 represent increasing bandwidth, and therefore increasing savings from assisted computation. To reveal the *strategic* benefits of these computational regimes, we simulate the four rover regimes across a Mars-like strategic drive.

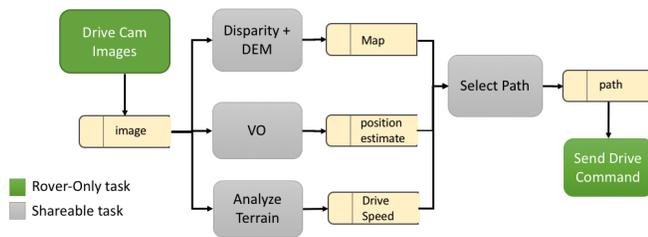
To test the different communication and computation regimes, simulations for four different regimes were run on three different terrain subsections 10 times each (resulting in 30 total runs) using stochastic durations for the path planning and terrain analysis activities.

The baseline regime is Regime 1, where the rover performs all path planning onboard and does not perform any terrain analysis. In Regime 2, the rover sends data to an orbiter where the path planning algorithm is performed and the results sent back to the rover. Regime 3 is the same as Regime 2, except that with the extra time, the rover performs terrain analysis onboard, which can be used for the next planning cycle. In Regime 4, terrain analysis is also performed on the orbiter and the results communicated back to the rover.

Figure 11b shows an example of the different paths that are taken for the different regimes when some of the terrain



(a) Simplified model of Mars 2020 path planning.

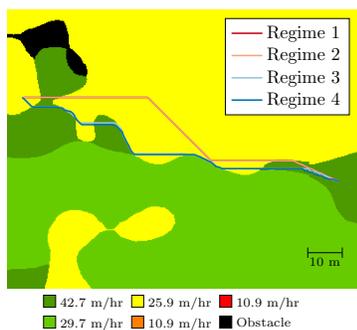


(b) A corresponding software network.

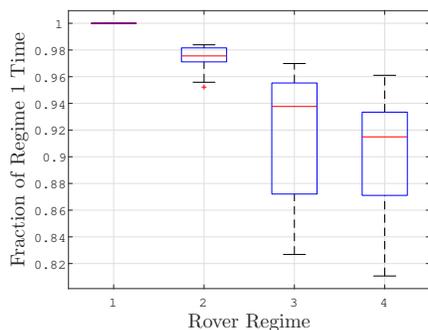
Figure 10: A model for the timing of Mars 2020 as discussed. The Select-path task is modelled as a random process taking a minimum of two seconds, but widely varying. The over-runs associated with any runtime longer than 30 seconds is the primary contributor to lost drive distance. The secondary contributor was a lack of terrain awareness, caused by insufficient processing power to run onboard terrain analysis.

b/w (Mbps)	Time	Image	Mapping	Extra Observations	Plan Path	Confirm / Drive	SPOC-lite
(0 - .1]	27	Rover	Rover	N/A	Rover	Rover	N/A
(0.1 - .3]	29.3	Rover	Rover	N/A	Assist	Rover	Rover
(0.3 - 1]	(29.7 - 28.2]	Rover	Rover	Rover	Assist	Rover	Rover
(1, 100]	(27.3 - 15.3]	Rover	Assist	Assist	Assist	Rover	Rover

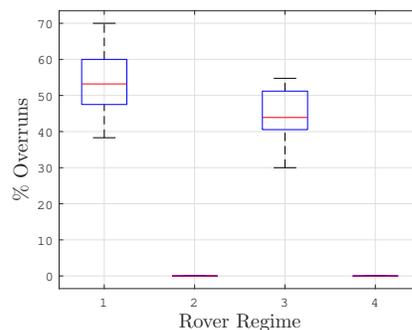
(a) A Mars 2020 rover adaptation of assisted drive. The adaptation was made using the pipeline information given in Figure 10a.



(b) Example paths



(c) Time to traverse waypoints per regime



(d) Overruns per traverse for each regime

Figure 11: Effect of computing regimes on a Mars 2020-like mission. 11b shows an example of different path choices. The main effects of additional computational assistance is reduced planner overrun and better terrain classification, resulting in more efficient paths, as shown in 11c and 11d. Terrain types are designated as different colors and the darker terrain (darkest except for black) can only be identified using terrain analysis.

is unknown without terrain analysis. The yellow terrain requires terrain analysis to be identified and is also slower to traverse. From this example, it is shown that with the terrain identification knowledge, Regime 3 and Regime 4 are able to come up with more efficient paths.

Since it is assumed that the rover must operate on a fixed 30 second cycle, if the path planning and/or terrain analysis are not completed within the allotted 8 seconds, an overrun will occur, causing the rover to stop until computation is completed. The distribution of percentage of overruns are shown as box plots in Figure 11d. As expected, Regime 2 and Regime 4 result in no overruns, whereas Regime 1 and Regime 3 have overruns around 50% of the time.

Another metric for the improvement of the rover performance is the time it takes to traverse a terrain. Figure 11c shows the time to traverse a terrain for each regime compared to the baseline (Regime 1). From these results, it is shown that being able to perform terrain analysis, and therefore being able to plan a path with better terrain knowledge, improves

the time to travel between two points.

We note a measurable increase in strategic drive efficiency using this limited study technique. Future work can focus on a more realistic terrain model, including that of the intended landing site. In addition, we can more realistically model the communication network. Intermittent loss of connectivity and varying data rates are significant impediments to this approach over long drives. Finally, modelling multiple assets would involve not only competing for the computational resources, but forwarding terrain classifications and drive rates between rovers.

7. CONCLUSION

We have presented a wide ranging set of results summarizing a subset of the Keck Institute for Space Studies workshop called Nebula. The workshop focused on the use of a future spacecraft *not only* as a remote-operated instrument, but also as remote data storage and processing assets.

Our hypothesis is that most science inquiry on Earth is already completed by interacting with data stored on a remote server (e.g., the Planetary Data System or PDS [1]), and a server placed near the planetary body of interest would have in its banks hundreds or even thousands of times more data and therefore be incredibly useful for scientific discovery.

Furthermore, the spacecraft itself has a reconfigurable set of algorithms to manage and summarize the data, and the downlinked information can be a varying mix of “raw” data or the synthesized data products from onboard analysis. The cost to realize this vision is to improve by orders of magnitude the available processing and data storage on spacecraft, but the benefits are potentially revolutionary, even before autonomous operations are considered. As a general-purpose, shared resource for computing, the concept resembles so-called “cloud computing”, yet emplaced in deep space, thus, we call it a *Nebula*.

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. The authors would like to thank the participants of the workshop who contributed ideas and discussions to this paper: Leon Alkalai (JPL), Morgan Cable (JPL), Les Deutsch (JPL), James Dickson (Caltech), Andrew Dittrich (USAF), David English (Lockheed Martin Space), Eric Frew (University of Colorado), Joseph Goldfrank (Stanford University), Shayn Hawthorne (Amazon Web Service), Jason Hofgartner (JPL), Robert Hood (ASRC Federal/NASA Ames), Ashish Mahabal (Caltech), Lukas Mandrake (JPL), Sreeja Nag (NASA Ames / BAERI), Mario Parente (University of Massachusetts), Raphael Some (JPL), David R Thompson (JPL), Jason Tichy (NVIDIA)

REFERENCES

- [1] S. K. McMahon, “Overview of the planetary data system,” *Planetary and Space Science*, vol. 44, no. 1, pp. 3–12, 1996.
- [2] D. R. Thompson, R. C. Anderson, B. Bornstein, N. A. Cabrol, S. Chien, T. Estlin, T. Fong, R. Hogan, R. Lorenz, D. Gaines, M. S. Gilmore, M. Parente, L. Pedersen, T. L. Roush, G. Marzo, , and D. Wettergreen, “Onboard science data analysis: Implications for future missions,” *Whitepaper submitted to the National Academy of Science’s Planetary Science Decadal Survey*, 2009.
- [3] A. Vasavada, J. Grotzinger, R. Arvidson, F. Calef, J. Crisp, S. Gupta, J. Hurowitz, N. Mangold, S. Maurice, M. Schmidt *et al.*, “Overview of the mars science laboratory mission: Bradbury landing to yellowknife bay and beyond,” *Journal of Geophysical Research: Planets*, vol. 119, no. 6, pp. 1134–1161, 2014.
- [4] K. L. Edmundson, D. Cook, O. Thomas, B. Archinal, and R. Kirk, “Jigsaw: The isis3 bundle adjustment for extraterrestrial photogrammetry,” *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, vol. 1, no. 4, pp. 203–208, 2012.
- [5] C. H. Acton Jr, “Ancillary data services of nasa’s navigation and ancillary information facility,” *Planetary and Space Science*, vol. 44, no. 1, pp. 65–70, 1996.
- [6] C. Acton, N. Bachman, J. Diaz Del Rio, B. Semenov, E. Wright, and Y. Yamamoto, “Spice: A means for determining observation geometry,” in *EPSC–DPS Joint Meeting*, 2011.
- [7] T. Duxbury, R. L. Kirk, B. A. Archinal, and G. Neumann, “Mars geodesy/cartography working group recommendations on mars cartographic constants and coordinate systems,” *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences*, vol. 34, no. 4, pp. 743–748, 2002.
- [8] S. Murchie, R. Arvidson, P. Bedini, K. Beisser, J.-P. Bibring, J. Bishop, J. Boldt, P. Cavender, T. Choo, R. Clancy *et al.*, “Compact reconnaissance imaging spectrometer for mars (crism) on mars reconnaissance orbiter (mro),” *Journal of Geophysical Research: Planets*, vol. 112, no. E5, 2007.
- [9] C. E. Viviano-Beck, F. P. Seelos, S. L. Murchie, E. G. Kahn, K. D. Seelos, H. W. Taylor, K. Taylor, B. L. Ehlmann, S. M. Wiseman, J. F. Mustard *et al.*, “Revised crism spectral parameters and summary products based on the currently detected mineral diversity on mars,” *Journal of Geophysical Research: Planets*, vol. 119, no. 6, pp. 1403–1431, 2014.
- [10] K. Runyon and L. Ojha, *Recurring Slope Lineae*. New York, NY: Springer New York, 2014, pp. 1–6.
- [11] L. Ojha, A. McEwen, C. Dundas, S. Byrne, S. Mattson, J. Wray, M. Masse, and E. Schaefer, “Hirise observations of recurring slope lineae (rsl) during southern summer on mars,” *Icarus*, vol. 231, pp. 365 – 376, 2014.
- [12] R. D. Lorenz, E. P. Turtle, J. W. Barnes, M. G. Trainer, D. S. Adams, K. E. Hibbard, C. Z. Sheldon, K. Zacny, P. N. Peplowski, D. J. Lawrence *et al.*, “Dragonfly: a rotorcraft lander concept for scientific exploration at titan,” *Johns Hopkins APL Technical Digest*, vol. 34, no. 3, p. 14, 2018.
- [13] K. E. Hibbard, D. S. Adams, A. A. Cocoros, N. L. Mehta, C. A. Sawyer, and B. F. Villac, “Dragonfly: In situ terrain-relative navigation for titan surface exploration,” in *16th International Planetary Probe Workshop (IPPW-16)*, 2019.
- [14] C. R. Webster, P. R. Mahaffy, S. K. Atreya, G. J. Flesch, M. A. Mischna, P.-Y. Meslin, K. A. Farley, P. G. Conrad, L. E. Christensen, A. A. Pavlov *et al.*, “Mars methane detection and variability at gale crater,” *Science*, vol. 347, no. 6220, pp. 415–417, 2015.
- [15] C. T. Russell and C. A. Raymond, “The dawn mission to vesta and ceres,” *Space Science Reviews*, vol. 163, no. 1, pp. 3–23, Dec 2011.
- [16] L. Prockter, S. Murchie, A. Cheng, S. Krimigis, R. Farquhar, A. Santo, and J. Trombka, “The near shoemaker mission to asteroid 433 eros,” *Acta Astronautica*, vol. 51, no. 1, pp. 491 – 500, 2002.
- [17] J. I. Trombka, L. R. Nittler, R. D. Starr, L. G. Evans, T. J. McCoy, W. V. Boynton, T. H. Burbine, J. Brckner, P. Gorenstein, S. W. Squyres, R. C. Reedy, J. O. Goldsten, L. Lim, K. Hurley, P. E. Clark, S. R. Floyd, T. P. Mcclanahan, E. McCartney, J. Branscomb, J. S. Bhangoo, I. Mikheeva, and M. E. Murphy, “The near-shoemaker x-ray/gamma-ray spectrometer experiment: Overview and lessons learned,” *Meteoritics & Planetary Science*, vol. 36, no. 12, pp. 1605–1616, 2001.
- [18] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, “Delay-tolerant net-

working: an approach to interplanetary internet,” *IEEE Communications Magazine*, vol. 41, no. 6, pp. 128–136, 2003.

- [19] J. T. Karras, C. L. Fuller, K. C. Carpenter, A. Buscicchio, D. McKeeby, C. J. Norman, C. E. Parcheta, I. Davydychev, and R. S. Fearing, “Pop-up mars rover with textile-enhanced rigid-flex pcb body,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5459–5466.
- [20] J. Vander Hook, T. Vaquero, F. Rossi, M. Troesch, M. S. Net, J. Schoolcraft, J.-P. de la Croix, and S. Chien, “Mars on-site shared analytics information and computing,” in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 29, no. 1, 2019, pp. 707–715.
- [21] T. Estlin, B. Bornstein, D. Gaines, R. Anderson, D. R. Thompson, M. Burl, R. Castano, and M. Judd, “Aegis automated science targeting for the mer opportunity rover,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2012.
- [22] R. Mattingly and L. May, “Mars sample return as a campaign,” in *2011 Aerospace Conference*, March 2011, pp. 1–13.
- [23] R. R. Rieber, “Designing for a martian road trip: The mobility system for mars-2020,” 2017, keynote Talk: Mars Forum (URS: URS270204, CL17-5707).
- [24] M. Ono, T. J. Fuchs, A. Steffy, M. Maimone, and J. Yen, “Risk-aware planetary rover operation: Autonomous terrain classification and path planning,” in *Aerospace Conference, 2015 IEEE*. IEEE, 2015, pp. 1–10.
- [25] M. Ono, B. Rothrock, E. Almeida, A. Ansar, R. Otero, A. Huertas, and M. Heverly, “Data-driven surface traversability analysis for mars 2020 landing site selection,” in *Aerospace Conference, 2016 IEEE*. IEEE, 2016, pp. 1–12.



the 2019 Keck Institute-funded study on in-space cloud computing.

Joshua Vander Hook received a Ph.D. in Computer Science in 2015 from the University of Minnesota. He is currently a distributed systems technologist in the Autonomous Systems Division of NASA Jet Propulsion Laboratory. He is most interested in navigation, experiment design, and activity planning problems for both space flight and military applications. He was principal investigator on



Scout CubeSat mission, and a co-investigator on various instruments and technologies intended to enable planetary exploration with small spacecraft.

Julie Castillo-Rogez received a Ph.D. in Planetary Geophysics in 2001 from the University of Nantes, France. She is currently a planetary scientist at the Jet Propulsion Laboratory, specializing in the exploration of small bodies, icy moons, dwarf planets, and the Moon. She is the Project Scientist for the Dawn Extended Mission, the science Principal Investigator for the Near Earth Asteroid



ing, space-based computing, autonomous systems, and data science.

Dr. Doyle is an Associate Fellow of the American Institute of Aeronautics and Astronautics. He holds the Ph.D. in Computer Science / Artificial Intelligence from MIT. He is past Executive Council member of the Association for the Advancement of Artificial Intelligence.

Richard J. Doyle is the Program Manager for Information and Data Science at the Jet Propulsion Laboratory (JPL) in Pasadena, California. His interests span data science, autonomous systems, computing, software engineering, space asset protection and other topics applying computer science principles space missions. He has authored articles on machine learning, model-based reason-



the University of Sao Paulo, Brazil. Tiago previously held a MIT research scientist position and a joint Caltech/MIT research position where he worked on Resilient Spacecraft Systems and Risk Sensitive Planning/Scheduling algorithms. At MIT, Tiago also worked on Risk Sensitive Planners and Executives for Autonomous Underwater Vehicles and Autonomous Cars. Tiago also held a research position at the University of Toronto where he worked on Robotics for Healthcare and Multi-Robot Coordination

Tiago Stegun Vaquero is a Technical Group Leader in the Artificial Intelligence Group, Planning and Execution Section, of the Jet Propulsion Laboratory, California Institute of Technology, where he works on multi-agent coordination methods for multi-rover cave exploration and surface site characterization. Tiago holds a B.Sc., M.Sc., and Ph.D. in Mechatronics Engineering from



engineering with a focus on spacecraft on-board processing, embedded machine learning, FPGAs, digital signal processing, and computer architecture. He received his M.S. and B.S. degrees in Computer Engineering from Rochester Institute of Technology in 2007. Prior to joining APL, Dmitriy worked at NRL, JPL, and Draper.

Dmitriy Bekker is Chief Technologist of the Space Systems Implementation Branch at Johns Hopkins Applied Physics Lab and supervisor of the Processing Architectures and Algorithms Section. He is the Image Processing FPGA Lead on DART and the Computer Vision Architecture Lead for NASA’s Dragonfly mission. Dmitriy’s area of expertise is hardware/software



Alice Cocoros received a Ph.D. in Physics and Astronomy in 2018 from Johns Hopkins University. She is a physicist in the Space Exploration Sector at Johns Hopkins University Applied Physics Laboratory. She is interested in autonomous navigation and onboard data processing for space flight applications, and more generally modeling, simulation, and data analysis for a variety

of engineering and science applications.



Trent Hare is a cartographer, project lead and supervisor at the USGS Astrogeology Science Center supporting planetary scientists and engineers since 1989. His work includes advancing Geographic Information Systems and photogrammetric software to not only facilitate planetary research and analysis, but to support various missions for Mars, the Moon, Titan and other bodies. He

is an activate member on several working groups to facilitate standards and interoperability across mapping software, formats and geospatial services. He received his M.S. in Computer Science from Northern Arizona University in 2011.



Randolph Kirk received his Ph.D. in Planetary Science from Caltech in 1987. He joined the USGS Astrogeology Science Center shortly thereafter as a geophysicist. He has been a member or guest investigator on 9 NASA/ESA/ISRO imaging teams and an associate of 4 others, as well as a Principal Investigator in the NASA Planetary Cartography program from 1988 to 2016. He devel-

oped unique software for photogrammetry and pioneered the application of digital photogrammetry to planetary optical and radar images, leading to the production of controlled image mosaics and digital topographic models of bodies from Venus to Triton. Since 2014 he has been a Scientist Emeritus at the USGS.



Valerie Fox received her Ph.D. in Planetary Science from Washington University in St Louis in 2017, and is now a postdoctoral researcher at the University of Minnesota. She studies alteration mineralogy on Mars using space-based Near Infrared spectrometers, and is a collaborating scientist with the Mars Science Laboratory rover Curiosity.