EXPLORING MARS VIA AUTONOMOUSLY NETWORKED SPACECRAFT.
E.J. Wyatt (e.jay.wyatt@jpl.nasa.gov), S.C. Burleigh, L.P. Clare, J.L. Torgerson, K.L. Wagstaff, Jet Propulsion Laboratory, California Institute of Technology

Introduction: Autonomous operations are transforming how spacecraft plan activities, analyze collected data, and select appropriate responses. The next step in capability is to enable multiple spacecraft to coordinate their activities autonomously by relying on new space networking protocol techniques [1] that allow autonomous management of communication links and data transmission. Each spacecraft can use any communication opportunity that presents itself, rather than requiring pre-planned communication sessions. This will significantly improve NASA's ability to perform ambitious new missions, such as Mars sample return and human exploration of Mars, by providing major leaps in capability [2]:

• *Collaborative observations* - one rover signals another rover to help with an observation, or one orbiter notifies another orbiter (or rover) about a short-lived event, such as a dust devil or seismic event so followup observations can be made quickly.

• *Improved reliability for human support on planetary surfaces* – for improved effectiveness of routine communications and to enable exploration of deep craters or other scientifically interesting locations where direct line-of-sight may be limited or non-existent.

• *Automated resource sharing among space assets to improve science return* – similar to the DARPA "fractionated spacecraft" concept, this approach enables entirely new design approaches for future missions. A rover mission could be designed with limited onboard resources (storage, fault diagnosis, science processing, etc.) if nearby or on-orbit vehicles are networked with it. More missions could be accomplished for the same cost or have access to more capability by relying on resources on other vehicles, including already emplaced assets.

• *Fully automated relay operations* – lander-orbiter relay communications are commonly used but are not fully automated. Further automation would improve throughput and reduce operations cost.

• *Greater Communications redundancy* – autonomous ability to coordinate communications means assets can take advantage of any nearby spacecraft, not just their primary relay. This reduces risk and enables new mission concepts.

These capabilities in turn enable the following major mission scenarios:

Surface-Orbit Distributed Processing: Autonomy often imposes a computational burden on spacecraft, since additional calculations must be performed to analyze incoming data and determine how to respond. For example, autonomous navigation on the MER rovers reduced their driving speed by a factor of 10 (5 cm/s to 0.6 cm/s) [3]. Imagine a rover that can request that an orbiter with a faster or more available CPU perform the Hazcam image analysis and obstacle avoidance processes for it, greatly accelerating the process. Even if outsourcing navigation is considered too risky, distributed processing could greatly increase the image analysis and data prioritization capabilities of a rover or lander.

Orbiters Relay through Landers: Autonomy provides the ability to detect and respond to dynamic or short-lived events (e.g., dust devils or polar avalanches on the surface of Mars). For circular, polar orbits (the type most typically used for mapping and monitoring), the time between repeat overflights of the same location on Mars is about two hours at the poles and increases to days or even months at lower latitudes. In some cases, a different spacecraft would view a given location before the first spacecraft returned to it. Autonomous networking would allow the first spacecraft to request that the other spacecraft collect observations of the site. If two orbiters do not immediately have a cross-link opportunity, they could communicate via relay through other spacecraft, such as landers or rovers acting as ground stations.

Onboard Anytime Data Classification: Communication relay windows between a lander/rover and an orbiter, whether pre-planned or opportunistic, may not always arise when the lander has completed the initial analysis steps allowing it to make fully informed decisions about link management. These decisions include: what kind of data compression to use, how to prioritize data for downlink (e.g., determining whether a dust devil is present in a given image), etc. If full processing is forced to halt with, e.g., only 80% of the data classified, then 20% of the data will never be examined at all. In contrast, *progressive refinement* [4] devotes computational resources to providing an anytime result: no matter when it is interrupted, it always has at least an approximate result for every item which improves over time. Tests have shown that not only does this increase coverage significantly, but the *total* time required to process all items is less than analyzing each one from scratch.

Onboard Anytime Data Prioritization: One of the key decisions that a remote mission must make is how to prioritize the data it collects for transmission back to Earth, especially when (as is typical) it has the ability to collect more data than it can transmit. Further, as science priorities become more sophisticated, the process of batch prioritization may take a significant amount of time. To make use of opportunistic links, the rover/lander must either suddenly prioritize all data it has collected, or rely on hard-coded pre-assigned priorities anticipated by science teams before the data is collected. Instead, we propose an onboard, anytime solution that continually updates the current prioritization of the data that has been collected, each time new data is obtained. No matter when the link is available, a prioritization is already available for use. This approach to prioritization is inherently dynamic, adapting the assigned priorities to fit the ongoing context in which they are collected.

Link-Aware Adaptive Data Collection Rate: Given limited onboard storage, critical decisions must be made about how much data to collect before the next download opportunity. When all links are entirely predictable, a data collection schedule follows from the available storage space and how much time will elapse before the next link opportunity. When communication links may be used opportunistically, the remote spacecraft should determine an appropriate data collection rate in response to current conditions. This capability is also needed for mobile rovers that move out of range or behind an obstacle and for any asset exposed to potential communications dropouts.

Link-Aware Adaptive Mobility: Rovers can also make driving decisions based on their observed link history. For example, if a rover has not received a communications link for an unusually long period, it might conclude that it had driven into an area that, for some reason, was blocking or restricting its communication pathway. The rover could decide to suspend current data collection activities and backtrack or otherwise move to try to re-establish communication capability. The robust ability to detect and recover from communications dropouts will permit rovers to explore a wider range of terrain.

Multi-Asset Coordination: Mayday Alerts: If a robot or human agent experiences a loss of connectivity, they must find alternate ways to transmit situational information to mission control for aid in troubleshooting and recovery. For multi-rover missions, a second rover could function as a relay assist for the compromised rover or human, maintaining an open path for communication with the stuck agent and with orbiters for relay to Earth.

Multi-Asset Coordination: Divide up Territory: A collection of rovers tasked with the exploration of a surface region could communicate to divide up the terrain for efficient exploration. While a rough division of territory may be part of the initial mission plan, autonomy enables updates based on observations made on the ground. In addition, any significant discoveries made by one rover can be shared with others, enabling them to autonomously alter their plans (or adjust their science priorities) given the new information about the terrain content.

Multi-Asset Coordination: Rovers and Scouts: Spacecraft with different capabilities can provide complementary data to inform each others' activities. An example would be a mission composed of a single MER-like rover and several smaller, faster, rover scouts. While the MER-like rover would be fully instrumented with cameras, spectrometers, weather sensors, etc., the scouts might have only a single sensor such as a camera. Their greater mobility and higher speed would allow them to quickly survey a large area and report back on potentially high science value targets. The rover would coordinate the information from all of the scouts, rank the candidate targets, and decide which one most merits the investment of time and power to investigate, and even direct the scouts to the next destination.

Summary: Autonomy and networking technology developments have been ongoing for several years. Combining the two provides the potential to a) enable autonomous networking among most or all Mars assets, b) enable entirely new ways to explore Mars through collaborative human-machine observations, and c) provides capabilities to enable NASA to meet potential exploration goals of returning a sample from Mars and sending humans to Mars.

Acknowledgments: The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References: [1] Barkley, E. et.al. Moving Toward Space Internetworking via DTN: Its Operational Challenges, Benefits, and Management, *SpaceOps 2010,* **2010.** [2] Wyatt, J. et.al. "Enabling autonomous exploration via the Solar System Internet, *IEEE Intelligent Systems,* 2010. [3] Maimone, M. et.al. "Autonomous navigation results from the Mars Exploration Rover (MER) mission," *9th International Symposium on Experimental Robotics*, 2004. [4] Wagstaff, K.L et.al. "Progressive refinement for Support Vector Machines," *Data Mining and Knowledge Discovery*, 20(1), p. 53-69, 2010.