Mars Coordinated Science Networking Demonstration

PI: Ross Jones
Contributors: André Girerd, Kiri Wagstaff,
Charles Baker, Rashied Amini, and Debarati Chattopadhyay
September 24, 2009

1. Introduction

Currently, spacecraft at Mars are commanded independently by their respective, separate mission operations teams. The value of coordinated observations of the same phenomenon is widely acknowledged and has led to some joint observing campaigns, such as simultaneous observations of the Martian atmosphere with the TES instrument on Mars Global Surveyor (looking down) and the Mini-TES instrument on the Mars Exploration Rovers (looking up). However, these joint observations were planned well in advance by operators on the ground and therefore were time-consuming in terms of human effort and also could not have been used to respond to a transient event. In this study, we investigated the possibility of using autonomous software onboard two remote spacecraft to enable an automated coordinated observation in response to a transient event. We focused on defining a Mars Coordinated Science Networking Demonstration, which would operate within a pre-scheduled demonstration window and then proceed without ground intervention to demonstrate autonomous communication, using Delay-Tolerant Networking (DTN) technology, and yield observations of the same surface feature by two (or more) spacecraft. We defined candidate operational scenarios, identified specific types of events or features for which fast automated coordination would be beneficial, quantified the expected latency between coordinated Mars observations, and evaluated current technology onboard existing Mars assets to determine the additional development required to achieve the coordinated science capability. We also identified several open issues that require further investigation in the next phase of this project. The next step is to scope out the amount of effort involved and to develop a cost estimate for the demonstration.

2. Operational Scenario

DTN can be used for reliable, automatic communication between spacecraft, similar to what is currently provided on Earth by the Internet. Existing Internet protocols cannot operate over the long distances and with the frequent gaps in connectivity experienced by spacecraft. This capability can be used to achieve coordinated observations of the same target. One spacecraft (the “trigger” spacecraft) detects an event or phenomenon of interest and then communicates to another spacecraft (the “followup” spacecraft) to request a second observation of the same target. We considered all of the current assets at Mars as candidates for a coordinated science demonstration. The most likely demonstration would be between two orbiters, Mars Odyssey (ODY) and Mars Reconnaissance Orbiter (MRO), and that is the scenario we focus on for the rest of the report (see Figure 1). In subsequent demonstrations, the two Mars Exploration Rovers (MER) could also participate, most likely as trigger spacecraft since their mobility is much more limited than that of the orbiters. The Mars Express spacecraft would also be a potential participant, in collaboration with the European Space Agency.

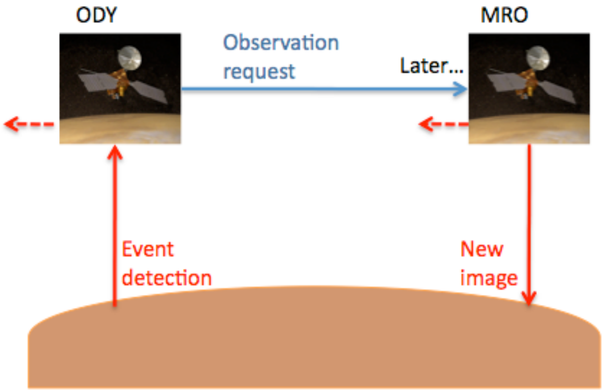

Figure 1. Concept for coordinated observations between Mars Odyssey (ODY) and Mars Reconnaissance Orbiter (MRO). ODY detects the event and sends a request to MRO. MRO images the same target at its next overflight opportunity.

We first consider the scenario in which ODY detects an event of interest and requests a follow-up observation by MRO. In this scenario, ODY uses onboard data analysis algorithms to evaluate data collected by the THEMIS instrument (a visible and near-infrared imager). When ODY detects an event of interest, it formats a message to be sent from its radio to MRO's radio at the next line-of-sight communication opportunity. This message contains information about the detected event, including its location from ODY's perspective. MRO, which briefly checks for just such a message at each line-of-sight opportunity, receives the request. MRO converts the target detection into its own coordinate system and adds a new item to its upcoming list of targets to be imaged by HiRISE or CRISM (depending on the type of the event). The next time MRO flies over the location where ODY detected the event, the appropriate instrument autonomously collects a new observation to complement the THEMIS data. Both images of the event of interest are downlinked to Earth at the next Deep Space Network (DSN) communication opportunity for each spacecraft. The two observations complement each other since THEMIS is able to continuously observe the planet, at 100 meters per pixel, while HiRISE and CRISM can provide higher-resolution (30 cm per pixel and 20 meters per pixel respectively), targeted follow-ups of specific locations. Figure 2 summarizes the major components needed to support this capability.

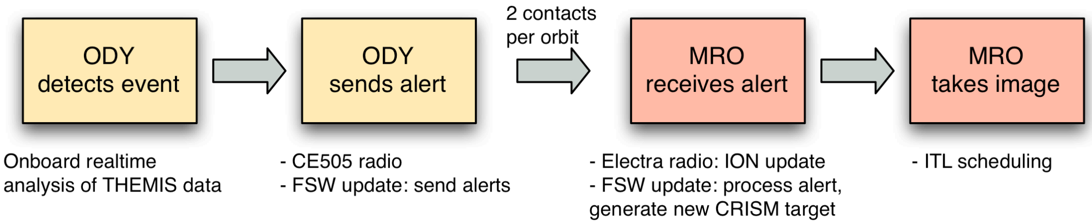

Figure 2. Major components needed to enable coordinated ODY->MRO observations.

In a similar scenario, MRO acts as the trigger spacecraft, e.g., in data from its CTX (Context Camera) instrument. MRO then sends a message to ODY, which responds by taking a follow-up image with THEMIS. CTX provides grayscale images at 6 meters per pixel; THEMIS provides visible and IR information at 100 meters per pixel. Information at a variety of different wavelengths greatly augments the scientific value of the original observation.

Science Phenomena of Interest

There are many transient phenomena of interest on Mars that would benefit from coordinated science using multiple Mars assets. Examples include:

- <u>Thermal anomaly (hotspot):</u> These could indicate subsurface hydrothermal activity. These features have been predicted, but never yet observed on Mars, so any such detection would have tremendous scientific significance. Since they have never been observed, the event's longevity is unknown, so a fast follow-up would also be important to obtain confirmation before it disappeared.
- <u>Dust devils (and tracks):</u> These have been observed to occur at all latitudes (but most likely between 45 and 75 S), and occurrence rates vary according to season and time of day (Whelley and Greeley, 2008). Longevity for dust devils that have been repeatedly observed is on the order of minutes, but there has been a strong sampling bias towards short-lived dust devils that means this is likely to be an under-estimate. In addition, there is value in imaging the area even after the dust devil is no longer active, to get observations of fresh dust devil tracks (e.g., to test the hypothesis that in some cases they may be triggering the formation of new dark slope streaks).
- <u>Meteoroid strikes:</u> A study by Byrne et al. (2009) found that "Twenty impacts created craters 2 to 150 meters in diameter within an area of 21.5 x 106 square kilometers between May 1999 and March 2006." The craters themselves persist for long periods, but Byrne et al. also found that these craters can sometimes excavate buried water or ice, especially at high latitudes (five examples are reported in the cited paper). Since water (and water ice) is not stable at the latitudes where these craters were found (45-55 N), fast repeat imaging is the only way to observe them before they disappear.
- <u>Polar geysers:</u> These have been hypothesized, due to curious "spider" deposits on the polar caps, but never yet observed in action. They are likely caused by CO2 sublimating under transparent ice until the increasing pressure bursts through the overlying ice layer (Kieffer et al., 2006).
- <u>Dark slope streaks:</u> These have long persistence (order of years to decades) so do not require immediate follow-up to obtain multiple images. However, the mechanism by which they form is still not well understood, and early follow-up could help determine which of the formation theories are most likely. They primarily occur on steep slopes between 30 S and 30 N (Schorghofer et al., 2007).
- <u>Gullies:</u> These appear without warning on slopes, from 30-70 N and 30-70 S. Although they persist for a long time, their formation mechanism is a mystery (Malin et al., 2006). Early repeat imaging could capture evidence of volatiles or other materials involved in their formation.

- Cave skylights: These are collapsed pits that can only be distinguished from impact craters by combining visible and near-IR observations (Cushing et al, 2007).

Odyssey Onboard Data Analysis

Algorithms currently exist for analyzing THEMIS data onboard ODY to detect certain events of interest, including thermal anomalies, the polar cap edge, and atmospheric concentrations of dust and water ice (Castano et al., 2007). These algorithms were ported to VxWorks for use by ODY but would require integration into the ODY flight software by Lockheed Martin. However, they have already demonstrated the ability to keep up with realtime data acquisition. Similar algorithms for identifying other features of interest could be added to the same software package.

Communication between Odyssey and MRO

ODY uses a CE505 radio, while MRO uses an Electra radio. These radios are compatible, and currently both are used to communicate with the Mars Exploration Rovers on the surface. Basic one-way communication between ODY and MRO has already been demonstrated, for the purposes of a radio science experiment. However, Electra is a software-defined radio, while the CE505 is not. Updating MRO's radio to support ION (the JPL implementation of DTN) requires only a standard Electra software update. Updating ODY's radio would require a patch to the main ODY flight software, and therefore would require more extensive review and testing. Incorporating ION into MRO's flight software (FSW) is also an option if it cannot be added to Electra.

Electra on MRO provides 2 MB of SRAM, of which 1 MB is already consumed by the Proximity-1 protocol used for communication with the MER rovers. ION required about 900 kB for the DINET experiment and could be compressed further if it excluded the Licklider Transport Protocol (LTP) and instead used the existing Proximity-1 implementation (or no protocol at all) to perform data transfer. The Proximity-1 consumption could also be pared down; e.g., it currently includes a very generously sized data buffer. ION also requires some buffer space to hold waiting (delayed) bundles. 13 MB of SDRAM (slower than the SRAM) and 1 MB of EEPROM (non-volatile) are also available. Figure 3 summarizes, for MRO, the architecture of the CPU (a RAD750) and the Electra radio as well as how they communicate.

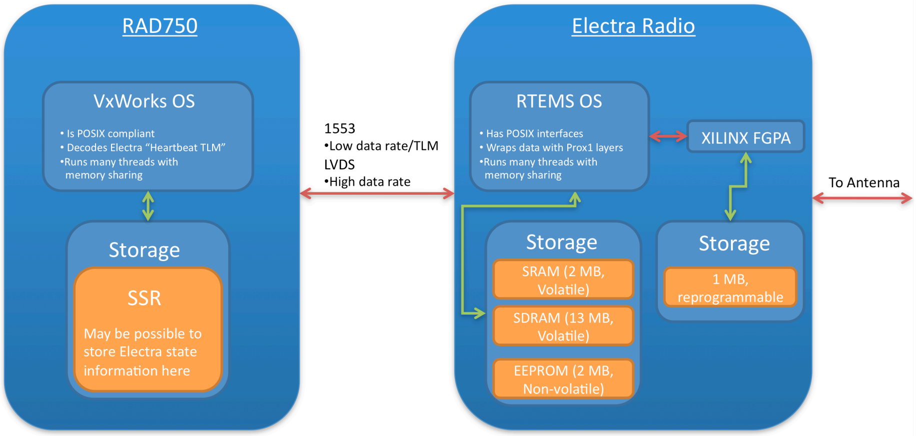

Figure 3. Architecture of CPU (RAD750) and Electra radio on MRO.

Integrating ION with Electra would have benefits beyond just this demonstration. Software-defined radios are the communication device of choice for future missions; e.g., the Mars Science Laboratory will also use an Electra radio. Each of these future missions could make use of the ION integration, or an adapted version of it, for their own radios.

ODY and MRO are in sun-synchronous orbits of opposite direction (ODY is descending; MRO is ascending), which affords at least two opportunities for communication per orbit. However, their relative geometry does not support a high data rate; both antennas point towards Mars and slews are limited to 30 degrees (even if slews were to be permitted for the communication link). However, communication should still be possible for a short period at a low rate such as 8 kbps. ODY and MRO both use the Proximity-1 protocol to communicate with the MER rovers, but they cannot use it to communicate with each other because it requires a duplex link, and their transmit and receive frequencies are not compatible. ODY transmits at 437.1 MHz and receives at 401.585625 MHz. MRO in full-duplex mode can transmit on any frequency between 435 and 450 MHz, and it can receive on frequencies between 390 and 405 MHz, but this does not match with ODY's frequencies. MRO in half-duplex mode can use any frequency between 390 and 450 MHz, so half-duplex link is theoretically possible. However, since ODY's radio does not have a half-duplex mode, in reality a simplex link would be used between these two spacecraft. ODY would employ its radio's "raw" data mode, which simply sends or receives bits as instructed by the flight computer.

There are two MRO-specific considerations relevant to the demonstration. First, Electra currently is not always on, primarily for power reasons. The radio is powered down when not being actively used to communicate with the MER rovers. Anything (such as the ION software) being stored in SRAM or SDRAM is lost and must be re-loaded. One option would be to request that Electra stay on continuously, at a low power level (not transmitting) for the duration of the demonstration. Alternatively, the ION software (and any waiting/delayed bundles) can be stored in EEPROM and reloaded each time Electra is turned on. Second, the Electra radio interferes with other instruments on MRO and therefore minimizing its use is desired, to enable more data collection opportunities for those instruments. The instruments affected are CRISM, SHARAD, and MCS, and they must be turned off when Electra is operating for thermal, radio-frequency interference, and pointing reasons, respectively.

MRO Targeting

When a new target request from ODY is received by MRO's Electra radio, that information must be communicated to the spacecraft flight software. There is a regular Electra "heartbeat" packet sent to the main flight CPU that could include this supplemental information. The MRO flight software, with an appropriate upgrade, would convert ODY's target into MRO's Mars coordinates, determine when the next overflight would be, and schedule the follow up observation. A coordinate conversion is necessary because ODY does not have any onboard knowledge of its current position, nor the position of any target it detects on Mars. All it can report is the time-on-orbit when the data was collected that triggered the detection. MRO, however, does have onboard ephemeris for itself, as well as a more capable CPU. It could store ephemeris for ODY as well to enable the conversion from ODY time-on-orbit to Mars Mean Equator coordinates, which are used onboard MRO.

There are two ways to schedule the new observation using existing MRO operational modes. The first is to add the new target to the Integrated Target List (ITL) at the appropriate position. When MRO reaches that item in the list, it will be executed just like any other item. It is possible that the new target would conflict with other pre-scheduled target(s). Assigning a priority to each targeted observation would enable MRO to decide which target to retain. The second option is to disable the ITL and use a purely direct-commanded mode, as is currently done about every six months for calibration purposes to permit the instruments to observe the moon.

Latency Analysis

For a follow-up observation to be collected by the second spacecraft, it must have an opportunity to fly over the same area that was imaged by the first spacecraft, in which the event or target of interest was detected. We investigated the latency (elapsed time between detection by one spacecraft and the next overflight by the other spacecraft) for observations by ODY and MRO at a variety of latitudes on Mars. To provide quantitative summaries of the latency experienced at each latitude, we used the Satellite ToolKit (STK) to obtain a list of each ODY and MRO overflight from 75 S to 75 N, at latitude increments of 15 degrees (22 such events per orbit, per spacecraft). This list was tallied over the six-month period from January to July, 2009, yielding over 4000 overflight events for each spacecraft. For each such overflight by one spacecraft, we identified the next time at which the other spacecraft crossed the same latitude and was sufficiently close in longitude to be able to image the same original location. We assumed that the first spacecraft was nadir-pointed, and the second spacecraft was permitted to slew up to 30 degrees off-nadir. We did not model the field of view of either spacecraft's instruments, instead assuming a point observation in both cases.

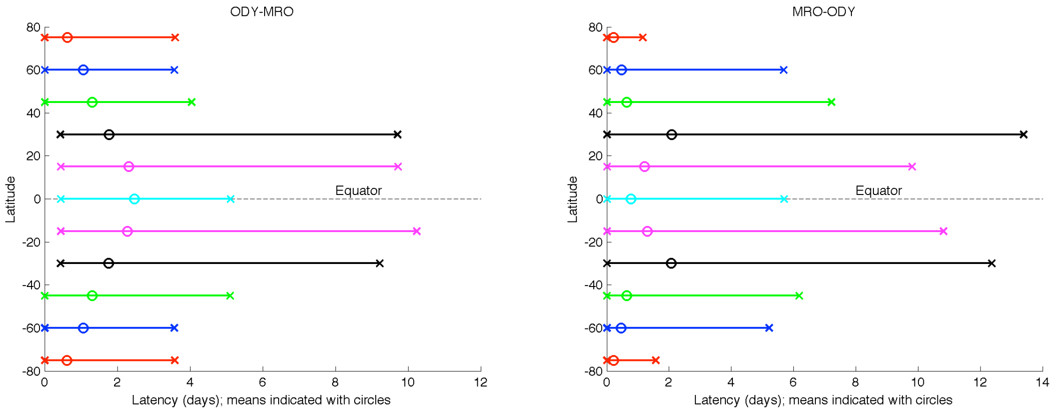

Figure 4. Latency results for coordinated observations from ODY to MRO (left) and MRO to ODY (right), as a function of latitude, tallied across six months of observations. Minimum and maximum latencies are indicated with x's, and the mean latency at each latitude is indicated with an o.

Figure 4 shows the latency results we obtained from this simulation. With ODY as the trigger spacecraft, the trend was as expected; the lowest latencies occurred near the poles, where coverage is higher, since both ODY and MRO are in polar orbits. The mean latencies ranged from 15 hours (at 75 N or S) to 2.5 days (at the equator). These times are short enough that onboard analysis and autonomous communication and retargeting could provide significant benefits over current operations requiring human involvement from the Earth. Although the round-trip communications time is only 10-40 minutes, there are other significant delays involved, such as the trigger spacecraft waiting for its next scheduled Deep Space Network (DSN) link, ground operators processing the data, conversations between the mission planners for each spacecraft, ground planning to add the new observation, and uplink of the new target to the follow-up spacecraft at its next DSN link. Current operations are not set up to provide a fast turnaround for this kind of operation (MRO generally schedules observations two weeks in advance), and it would impose a significant additional burden to the operators. The worst cases, at 15 and 30 N or S, took up to 10 days.

With MRO as the trigger spacecraft, the results were different; the lowest latencies were again observed near the poles (5 hours), but the highest mean latencies were not at the equator. Instead, they occurred for 30 N or S (2 days); mean latency at the equator was only 18 hours. This curious pattern is a result of the relative geometry of the two spacecrafts' orbits. The worst cases for MRO took up to 13 days for a repeat observation opportunity by ODY.

Note that these results include some impossibly short latencies (e.g., 4 seconds) that occasionally happen when one spacecraft reaches the same location right after the other one. While these are coordinated observation opportunities, they could never be utilized since it would take longer than that to detect the event, communicate to the other spacecraft, compute the new target location, and schedule the new observation. With a good estimate for (or via modeling of) the time required to achieve these necessary steps, we could refine these latency results. Most latencies would increase slightly. Over a six-month period, the maximum elapsed time between ODY/MRO communications opportunities was 17.6 minutes, with an average of 15.4 minutes.

3. Additional Work Needed

Mars Odyssey Orbiter

- Store list of possible communications times for MRO (can be pre-calculated)
- Integrate THEMIS analysis software with ODY FSW (Lockheed)
- Update FSW to control radio and send/receive trigger events as packets to MRO (Lockheed?)

Mars Reconnaissance Orbiter

- Store list of times to listen for a signal from ODY (can be pre-calculated)
- Store ODY ephemeris
- Integrate ION into Electra
- Update Electra to accept messages from ODY, decode, and pass to FSW
- Update FSW to convert trigger event from ODY coordinates to Mars coordinates, create a new target observation, and add it to the ITL (or use non-ITL mode)

- Add priorities for observations, so that the trigger events only replace existing observations if they are at a higher priority (priorities are currently possible but not used)

4. Outstanding Questions and Issues

1. What is MRO's pointing precision? Given target coordinates, what is the probability of hitting the same location previously observed by Odyssey? Is this within the swath of the instrument being used to follow up?
2. How much of the new MRO capabilities should be done via a flight software update for the main processor, and how much should be done by instrument software updates?
3. How to ensure that ODY communication to MRO does not interfere with other instruments onboard ODY? Does Electra conflict with any of them? Electra itself may already be in use to talk to MER. Same for MRO (we already know that Electra interferes with CRISM, SHARAD, and MCS).
4. Can MRO's ITL be dynamically modified? (i.e., does it open a file and progressively read one target at a time from it, or does it read it all into memory (thus modifiable) and then go through it?) If not, the non-ITL mode will need to be used.
5. Can MRO figure out when it will be able to see the new target (and how much to slew), to determine when the new image should be taken? Can MRO figure out where in the ITL to insert the new target? How much additional onboard planning or modeling capability will it need?

Acknowledgements

The information contained in this report was compiled from discussions with a large number of generous individuals, including David Beaty, David Bell, Scott Burleigh, John Callas, Steve Chien, Stuart Demcak, Chad Edwards, Kurt Eggemeyer, Pasquale Esposito, Samad Hayati, Peter Illot, Tom Jedry, Ricardo Mendoza, Ken Peters, Jeff Plaut, Joseph F. Smith, Dan Wenkert, and Rich Zurek.

References

Byrne, S., Dundas, C.M., Kennedy, M.R., Mellon, M., Shean, D., Daubar, I., Cull, S., Seelos, K.D., Murchie, S., Cantor, B., Arvidson, R.E., Edgett, K., McEwen, A., Harrison, T., Posiolova, L., Seelos, F.P., and the HIRISE, CTX, and CRISM teams, "Excavation of Subsurface Ice on Mars by New Impact Craters," *Proceedings of the 40th Lunar and Planetary Science Conference,* paper 1831, 2009.

Castaño, R., Wagstaff, K.L., Chien, S., Stough, T.M., and Tang, B. "On-board Analysis of Uncalibrated Data for a Spacecraft at Mars," *Proceedings of the Thirteenth International Conference on Knowledge Discovery and Data Mining,* p. 922-930, 2007.

Cushing, G.E., Titus, T.N., Wynne, J.J., and Christensen, P.R., "THEMIS Observes Possible Cave Skylights on Mars," *Geophysical Research Letters*, 34, L17201, doi:10.1029/2007GL030709, 2007.

Kieffer, H.H., Christensen, P.R., and Titus, T.N., “CO_2 Jets Formed by Sublimation Beneath Translucent Slab Ice in Mars’ Seasonal South Polar Ice Cap,” *Nature*, 442, doi:10.1038/nature04945, 2006.

Malin, M.C., Edgett, K.S., Posiolova, L.V., McColley, S.M., and Dobrea, E.Z., “Present-Day Impact Cratering Rate and Contemporary Gully Activity on Mars,” *Science*, 314, doi:10.1126/science.1135156, 2006.

Schorghofer, N., Aharonson, O., Gerstell, M.F., and Tatsumi, L., “Three Decades of Slope Streak Activity on Mars,” *Icarus*, 191:132-140, 2007.

Whelley, P.L. and Greeley, R., “The Distribution of Dust Devil Activity on Mars,” *Journal of Geophysical Research*, 113, E07002, doi:10.1029/2007JE002966, 2008.