# Opportunistic MSPA Demonstration #1: Final Report

**Douglas S. Abraham,* Susan G. Finley,† David P. Heckman,‡ Norman E. Lay,‡ Cindy M. Lush,▲ and Bruce E. MacNeal‡**

ABSTRACT. — The Opportunistic Multiple Spacecraft Per Antenna (OMSPA) concept seeks to provide smallsat missions with a low-attributed-aperture-fee technique for obtaining routine downlink in a manner that is very low cost to the Deep Space Network (DSN). Unlike traditional MSPA in which the number of spacecraft that can be supported is limited by the number of available receivers, OMSPA makes use of a digital recorder at each station that is capable of capturing IF signals from every spacecraft in the antenna beam within the frequency bands of interest. When smallsat missions see one or more opportunities to intercept the traditionally scheduled antenna beam of a "host" spacecraft, they can transmit open loop during those opportunities. Via a secure Internet site, the smallsat mission operators can then retrieve relevant portions of the digital recording for subsequent demodulation and decoding or subscribe to a service that does it for them. The demonstration discussed in this article was intended to provide prospective smallsat users and the DSN, as the prospective service provider, with demonstrable proof that the OMSPA concept is, in fact, an operationally viable means for obtaining routine downlink telemetry. To do this, the demonstration began by treating Mars Odyssey as a "smallsat" and Mars Reconnaissance Orbiter (MRO) as the "host" spacecraft. Using a specially created Beam Intercept Planning System (BIPS) and a DSN 7-Day Schedule Cross-Comparison (7-DSC) tool, opportunities were identified when Mars Odyssey would be transmitting while in MRO's ground antenna beam. Existing Very Long Baseline Interferometry (VLBI) Science Receivers (VSRs) were used to record the Mars Odyssey downlink telemetry during these opportunities. The recordings were played back to a secure server outside the Flight Operations Network firewall, but inside the JPL firewall. The demonstration team's signal processing personnel retrieved the recordings from this secure server and downloaded them to a workstation containing an OMSPA Software Demodulator (OSD) tool that was developed to demodulate and decode the Mars Odyssey signal. Validation of the recovered data was then accomplished by comparing the transfer frames obtained through OMSPA with those recovered via Mars Odyssey's formally scheduled downlink. The demonstration successfully achieved its

* Architecture, Strategic Planning, and System Engineering Office.
† Tracking Systems and Applications Section.
‡ Communications Architectures and Research Section.
▲ Mission Systems Engineering Section.

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

intended purpose. All of the above steps were accomplished within an operationally viable timeframe, with at least 99.95 percent of the transfer frames being successfully recovered from each demonstration recording.

## I. Opportunistic MSPA Concept Overview

The current era of highly constrained NASA budgets has generated significant interest in, and pursuit of, small satellites that can be inexpensively launched on relatively small launch vehicles or as secondary payloads on larger launch vehicles. Such smallsats typically weigh less than 500 kg. The best known of these, CubeSats, make use of multiples of $10 \times 10 \times 10$ cm units of volume (1U cubes) that are each generally capable of containing around 1.33 kg of hardware. Deep-space CubeSat missions typically make use of 6U configurations and generally cost between $5M and $20M to develop, launch, and operate. Because they are so mass, power, and volume constrained, their onboard telecommunications capability is quite limited — necessitating reliance on the Deep Space Network's (DSN's) high effective isotropic radiated power (EIRP) and gain/temperature (G/T) ground assets to close the communications link. The aperture fees allocated to these CubeSat missions while they are vying for a competitively bid opportunity, however, can amount to a substantial percentage of their total mission cost, making them potentially less attractive than their Earth-orbiting amateur-band competitors. At the same time, the significantly larger number of missions that can be launched at $5M to $20M apiece has the potential to overwhelm the traditional DSN service capacity. Opportunistic Multiple Spacecraft Per Antenna (OMSPA) is a concept for both reducing the attributed downlink cost to such missions and increasing their ability to make efficient use of the DSN's limited downlink capacity.

MSPA techniques have been used for well over a decade to increase the efficient utilization of ground network assets while decreasing the antenna fees allocated to the missions. In the DSN's traditional MSPA service, two missions that will be located within the same beam of a ground antenna (such as at Mars) can schedule to share the antenna (Figure 1). The antenna, of course, must be equipped with two separate receivers — one for each spacecraft. Applying this MSPA service to more than two spacecraft at a time requires adding receivers. But, sharing an antenna amongst many spacecraft by adding lots of receivers and associated telemetry processing chains could prove prohibitively expensive for the DSN.

Opportunistic MSPA attempts to solve this problem by employing an available recorder at each station capable of capturing IF signals from every spacecraft in the antenna beam within the frequency bands of interest (Figure 2). Smallsats can opportunistically transmit open loop when in a host spacecraft's antenna beam. Via a secure Internet site, the smallsat mission operators can then retrieve relevant portions of the digital recording for subsequent demodulation and decoding, or subscribe to a service that does it for them.

In order for a smallsat project's spacecraft to opportunistically downlink, the smallsat project must first design the spacecraft's trajectory to be within the beamwidth of a "host" spacecraft's ground receiving antenna. For properly registered smallsat missions, the DSN could make the tracking schedules and associated mission trajectory files for each station available, allowing the smallsat project to identify potential "host" spacecraft opportunities.

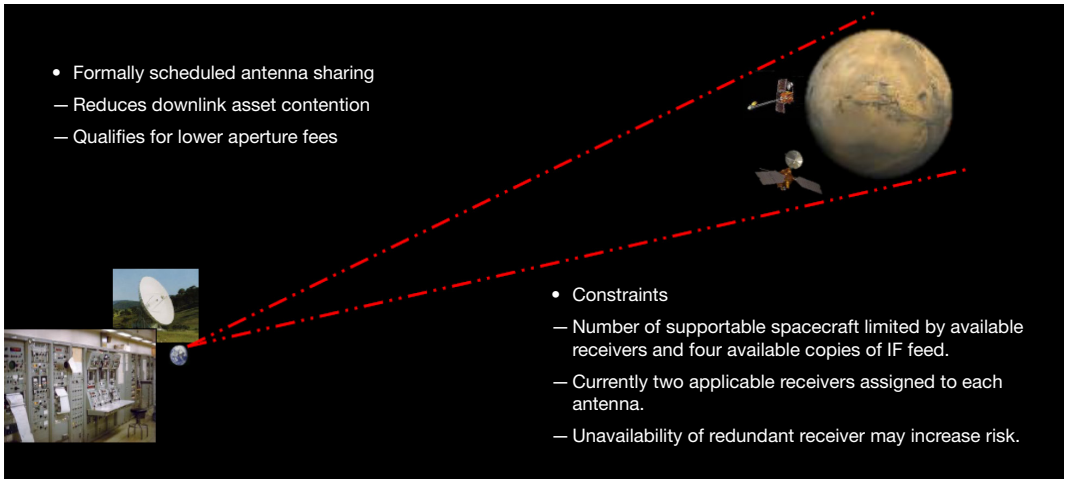


**Figure 1. Traditional MSPA.**

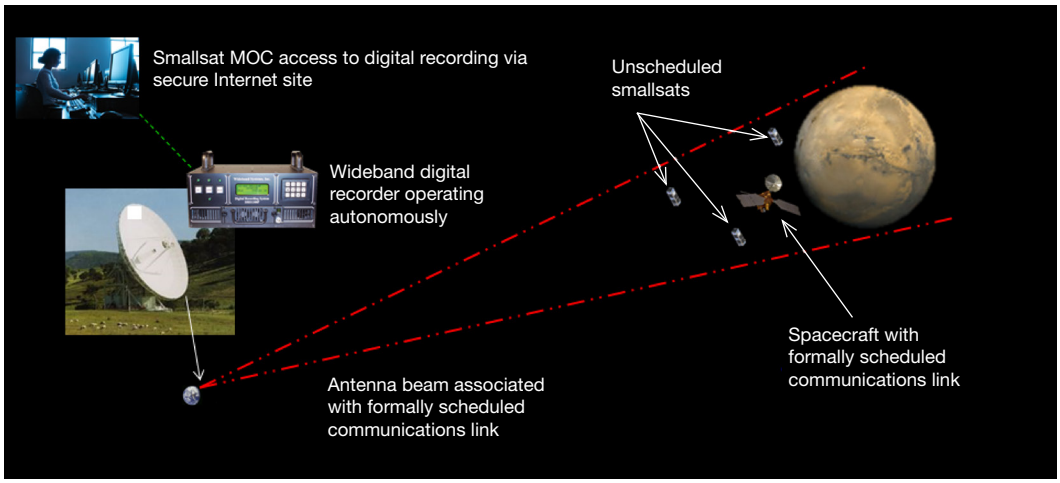


**Figure 2. Opportunistic MSPA. Everything received through the antenna beam is digitally recorded. Smallsats transmit open loop when in a host spacecraft's beam. Smallsat MOCs retrieve the relevant portion of the digital recording for subsequent demodulation and decoding.**

The smallsat projects could use that information to figure out whether or not they would be within the beamwidth of a particular "host" spacecraft's stations and for how long. The smallsat projects could then plan to downlink their data during these time periods.

Analysis shows that missions using OMSPA at the Moon, lunar Lagrange point 2, the Sun–Earth Lagrange points, Earth leading and trailing orbits, Mars, and Venus can, depending upon the spacecraft and ground receiving antenna's capabilities, support average data return rates on the order of hundreds of bps to tens of Mbps while maintaining safe separation distances from the "host" spacecraft — distances on the order of thousands to millions of kilometers [1].

OMSPA may also be applied to constellations of spacecraft operating in distant highly elliptical orbits (HEOs) and beyond. This, of course, assumes that one of the spacecraft

utilizes traditionally scheduled antenna passes and the other spacecraft at least periodically intercept its ground antenna beam.

To minimize radio frequency interference during OMSPA, the smallsat missions need to be assigned to an appropriate subband that is not in use by the surrounding "host" spacecraft. If multiple smallsats will be operating with OMSPA in the same vicinity, it may be desirable to have them share the allocated subband via code division multiple access (CDMA).

While OMSPA may require up to a couple of days for missions to access and recover their downlinked data, it allows them to avoid the level of attributed aperture fees and other charges that would otherwise be associated with scheduling formal, traditional communication links on the antennas. The smallsat mission's primary cost would be the labor to retrieve the relevant portion of the recording, pull the signal out of the "noise," and reconstruct the downlinked data. Hence, the smallsat's use of more costly, formally scheduled antenna passes could be limited to a much smaller number essential for commanding, navigation-related radiometrics, and real-time critical event telemetry. At the same time, the cost to the DSN would essentially be limited to providing and maintaining the appropriate digital recorders and the secure Internet site through which the smallsat missions can access the needed portions of the recordings.

## II. Demonstration Purpose and Objectives

Demonstration #1 was intended to provide prospective smallsat users and the DSN, as the prospective service provider, with demonstrable proof that the OMSPA concept is an operationally viable, low-cost means for obtaining routine downlink telemetry. In particular, the demonstration was designed to show the following:

(1) Given access to the DSN station schedules, associated trajectory files, and an appropriate tool, a user mission can compute the beam intercept times for its trajectory relative to that of a host spacecraft such that it has the capability to arrange its downlink operations to accord with those times.

(2) Given the tool in item #1, a user mission can successfully retrieve the portions of the wideband recordings corresponding to its signal via a secure Internet site.

(3) Given an appropriate tool, a user mission can successfully demodulate, decode, and recover its data from the retrieved recordings within a timeframe that is operationally reasonable for the user mission.

(4) A user mission can successfully do these things irrespective of which DSN complex is in view during the pass.

## III. Methodology Overview

The demonstration was performed using the downlinks from two spacecraft already operating at Mars: Mars Odyssey and Mars Reconnaissance Orbiter (MRO). Mars Odyssey was assumed to be the "smallsat" and MRO was assumed to be the "host" spacecraft. The objec-

tive was to capture Mars Odyssey's downlink transmission with a recorder associated with MRO's downlink antenna, transfer the recording to a secure Internet site, have the smallsat team (the demonstration team in this case) retrieve this transmission from the appropriate portion of the recording from the internet site, demodulate and decode it, recover the data, and validate the data based on the traditionally downlinked data supplied by the Mars Odyssey project.

This objective was pursued for two types of passes: (1) A traditional MSPA pass between Mars Odyssey and MRO, and (2) a situation where Mars Odyssey and MRO were each scheduled to downlink on a separate antenna during the same time period.

The purpose of the OMSPA test on the traditional MSPA pass was to exercise the OMSPA hardware and software in a well-understood context to: (1) identify and correct any glitches in the tools or technique, and (2) establish an approximate "turnaround time" for signal retrieval.

The purpose of the second type of pass was to demonstrate that the smallsat's (Mars Odyssey's) open-loop transmission could be retrieved from the "host" spacecraft's (MRO's) ground antenna recording due to being in the host's antenna beam. This was important to show because it establishes that no planning has to occur between the smallsat and the host for OMSPA to work. Demonstration with the second type of pass also had to be performed once from each complex to establish that OMSPA can be successfully employed at any of the DSN complexes.

## IV. OMSPA Demonstration Tools/Systems

Figure 3 provides an overview of the key tools that were used in demonstrating OMSPA, as well as the sequence in which they were applied. Three of the tools were developed as part of the demonstration. Two of the tools/systems already existed and were shared with other users. The following paragraphs describe the origin, nature, and key features/applications of each tool in more detail.

### A. Beam Intercept Planning System

The Beam Intercept Planning System (BIPS) is a new tool referenced in the orange display of Figure 3 and given an overview in this section. The main purpose of BIPS is to compute all of the motions of the system of spacecraft ("host" and any number of "smallsats"), the communications ground stations, and possible occulting bodies (Mars, Moon, etc.), then accurately identify and report the intervals of time where opportunities for OMSPA exist. In addition to the motions of all the objects in any OMSPA scenario, OMSPA opportunities also depend on the angular beamwidth that applies to scheduled links from each ground station to the host spacecraft. Hence, ground antenna beamwidth is a required input to any scenario computed by BIPS.

In addition to the computation of the opportunity intervals, BIPS computes statistical information about the collection of opportunities that exist in any scenario. It has sev-

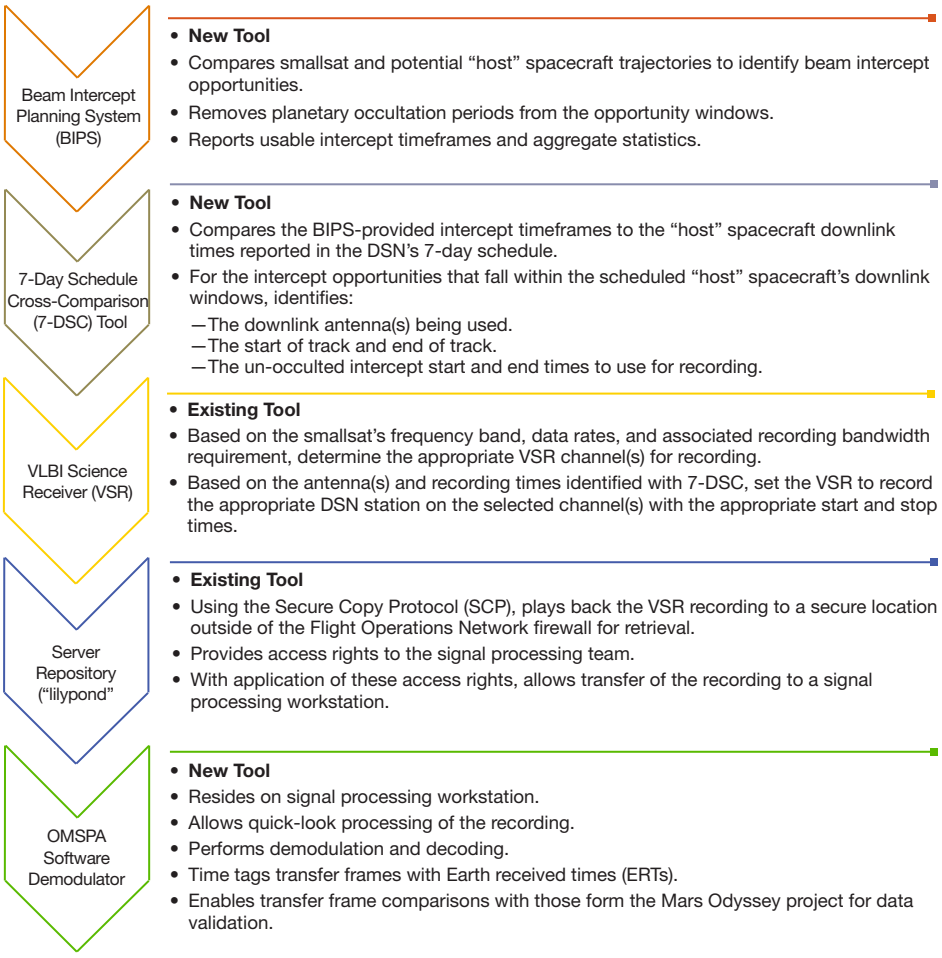


**Figure 3. Key demonstration tools.**

eral graphical and animation features that greatly help one visualize and understand the motions of the system of objects, visibility, and beam constraints, and how these factors combine to impact potential OMSPA opportunities.

The key input data required for the computation of any OMSPA scenario are a set of accurate trajectory files assigning the position evolution of the host spacecraft and all of the smallsats involved in the scenario. The file type required for any trajectory is a SPICE[1] "SPK" file, but other common file types that can be converted outside BIPS to accurate SPK representations will also serve the purpose, after processing. As the OMSPA feasibility is based on potentially fast-changing interspacecraft relationships and beam intercept intervals and not just visibility from stations on Earth, trajectory modeling accuracy is crucial for the correct identification of potential OMSPA opportunity intervals.

Figure 4 provides a view of the main graphical user interface of BIPS, where all of the trajectory files are loaded and all of the subcomponents of the system are accessed. Once a set of trajectory files has been loaded, as seen in the uppermost list, any one of them can be chosen to represent the host spacecraft. (In this example, the file representing MRO has been

[1] The SPICE (Spacecraft, Planet, Instrument, C-Matrix, Events) Toolkit is provided by JPL's Navigation and Ancillary Information Facility — http://naif.jpl.nasa.gov.

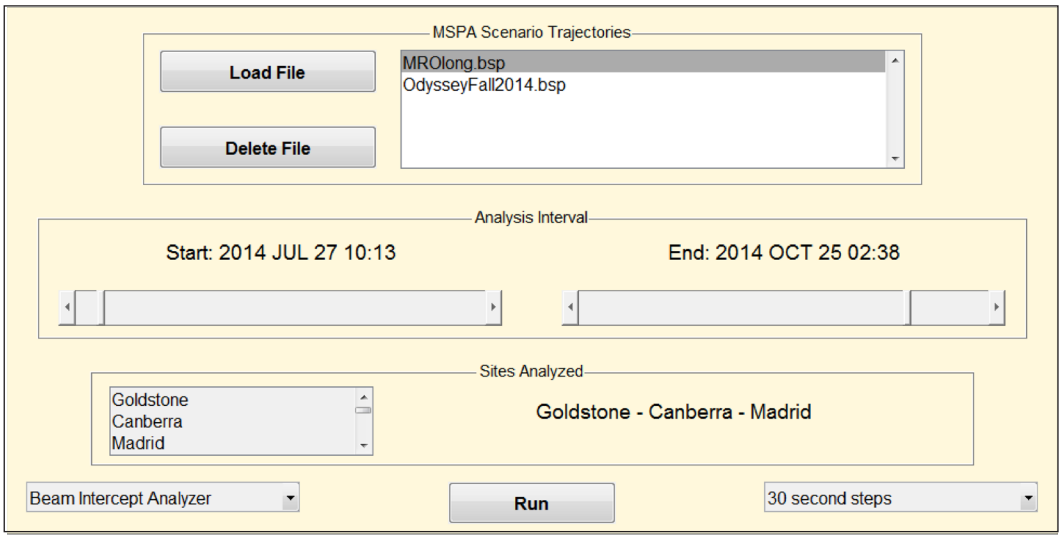


**Figure 4. Graphical user interface for the Beam Intercept Planning System (BIPS).**

selected.) The remaining files are the smallsats (or, in this example, Mars Odyssey, which was used as a representative "smallsat" in the demonstration). The tool automatically finds the largest coverage interval that is common to all of the loaded files, and the user operates the slider controls near the middle of the graphical user interface (GUI) to select a subinterval of that period for the scenario analysis. The default is the full coverage interval.

With the files loaded and a time interval selected, the user may then select a group of three Earth ground station sites for analysis, with the DSN station sites as the default group. A few different tools can be accessed using the lower-left list, and the nominal time step to be used in animation displays can be selected using the lower-right list. Clicking the "Run" button performs the scenario's motion calculation and invokes the selected subtool (for this example, the setting shown is the animated beam intercept analyzer/visualizer).

The beam intercept visualization subtool provides a three-site animated display of the relationships between the objects in the scenario and their positions with respect to the beam pattern centered on the host (Figure 5). The observations from all three selected stations are shown simultaneously, with a white background corresponding to visibility of the host spacecraft at that instant, and the colored background implying that the host is not visible from the associated station at that instant. All motions are displayed in the animation whether the host spacecraft is instantaneously visible from the station or not.

The GUI of the beam intercept visualizer, displayed in Figure 5, has several controls for on-the-fly animation adjustments and can export a movie file of the animation. This subtool also has the controls necessary for accepting a beamwidth value, and that value is instantly reflected in the running animation. Once the user has an appropriate beamwidth entered into the visualization subtool, the deep analysis for identifying all of the OMSPA scenario opportunity intervals can be set up to run. The user can choose whether or not to include occultation effects due to a particular planetary body in the analysis, and he or she can choose the desired computation precision and formatting style of the output data. Clicking

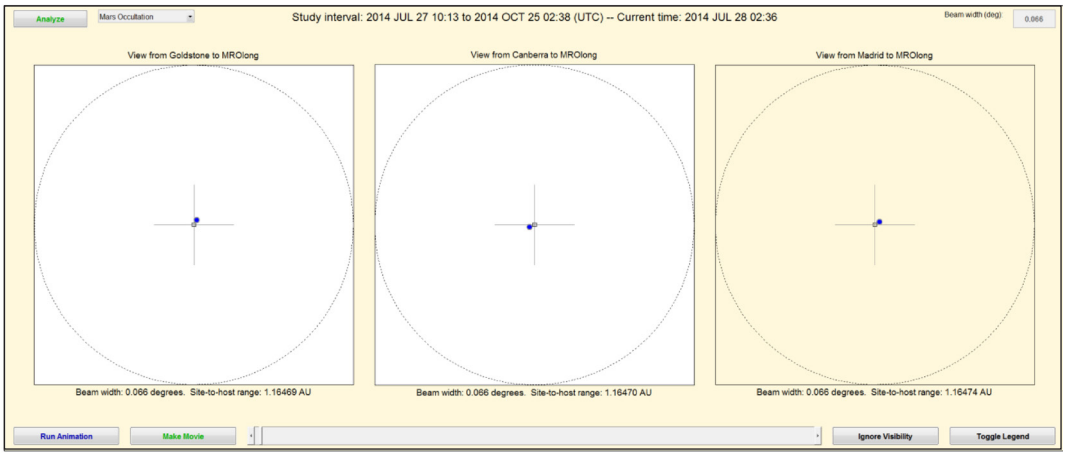


**Figure 5. Display output associated with the beam intercept visualization subtool.**

the "analyze" button then launches the computation. Upon completion of the job, usually within a couple of minutes, a Microsoft Excel workbook is exported.

The Excel files contain all the beam intercept events, both entering and leaving the host-centered beam, and the times at which those events occur, as well as a separate sheet of summary statistics in columns that match the columns of intercept data. All times are expressed in Coordinated Universal Time, with a couple of different standard time formats available. A sample of data is shown in Figure 6 that applies to a Mars Odyssey–MRO beam intercept scenario, accounting for Mars occultations. Each time in the sample columns corresponds to the "smallsat" Mars Odyssey entering or leaving feasible OMSPA geometry — with a "+" sign indicating a time of entry and a "–" sign indicating a time of departure.

In this example, to be geometrically feasible for OMSPA with a particular ground station, MRO and Mars Odyssey must both be in view from the station, with Mars Odyssey not occulted by Mars and inside the antenna beam pattern applying to MRO's link to the ground station. In all of this, the host spacecraft, MRO, maintains its antenna track even while occulted by Mars (the occulted period is too short to warrant terminating the track and going through all of the pre- and post-cals for new tracks over and over again). A small snapshot of some of the data produced by BIPS is shown in Figure 6. The first three columns give the event times delineating OMSPA feasibility for the three DSN sites, and the fourth column provides the same metrics but for the network of all three stations taken as a union — implying that any OMSPA feasibility for one or more stations applies to the network total.

BIPS was developed in MATLAB and has a built-in planetary ephemeris (provided by a file from JPL's Navigation and Ancillary Information Facility, using SPICE toolkit components). It was built to run in the Microsoft Windows environment and requires a now-common 64-bit operating system. It is entirely stable at this time. Future developments may include slight usability improvements and compiling the system into an executable deliverable package.

| Goldstone-to-OdysseyFall2014 | Canberra-to-OdysseyFall2014 | Madrid-to-OdysseyFall2014 | Network-to-OdysseyFall2014 |
|---|---|---|---|
| - 2014 AUG 01 01:34:50 | + 2014 AUG 01 02:08:10 | + 2014 AUG 01 13:25:10 | - 2014 AUG 01 01:34:50 |
| + 2014 AUG 01 02:08:10 | - 2014 AUG 01 03:33:30 | - 2014 AUG 01 13:25:50 | + 2014 AUG 01 02:08:10 |
| - 2014 AUG 01 03:33:30 | + 2014 AUG 01 04:06:50 | + 2014 AUG 01 13:59:10 | - 2014 AUG 01 03:33:30 |
| + 2014 AUG 01 04:06:50 | - 2014 AUG 01 05:32:00 | - 2014 AUG 01 15:24:30 | + 2014 AUG 01 04:06:50 |
| - 2014 AUG 01 05:32:00 | + 2014 AUG 01 06:05:20 | + 2014 AUG 01 15:57:50 | - 2014 AUG 01 05:32:00 |
| + 2014 AUG 01 20:43:00 | - 2014 AUG 01 07:30:30 | - 2014 AUG 01 17:23:00 | + 2014 AUG 01 06:05:20 |
| - 2014 AUG 01 21:20:00 | + 2014 AUG 01 08:03:50 | + 2014 AUG 01 17:56:20 | - 2014 AUG 01 07:30:30 |
| + 2014 AUG 01 21:53:20 | - 2014 AUG 01 09:28:50 | - 2014 AUG 01 19:21:30 | + 2014 AUG 01 08:03:50 |
| - 2014 AUG 01 23:18:20 | + 2014 AUG 01 10:02:10 | + 2014 AUG 01 19:54:50 | - 2014 AUG 01 09:28:50 |
| + 2014 AUG 01 23:51:50 | - 2014 AUG 01 11:27:20 | - 2014 AUG 01 21:20:00 | + 2014 AUG 01 10:02:10 |
| - 2014 AUG 02 01:17:00 | + 2014 AUG 01 12:00:40 | + 2014 AUG 01 21:53:20 | - 2014 AUG 01 11:27:20 |
| + 2014 AUG 02 01:50:20 | - 2014 AUG 01 13:16:50 | - 2014 AUG 01 21:54:40 | + 2014 AUG 01 12:00:40 |
| - 2014 AUG 02 03:15:30 | + 2014 AUG 02 01:50:20 | + 2014 AUG 02 13:41:20 | - 2014 AUG 01 13:16:50 |
| + 2014 AUG 02 03:48:50 | - 2014 AUG 02 03:15:30 | - 2014 AUG 02 15:06:40 | + 2014 AUG 01 13:25:10 |

**Figure 6. A sample of BIPS output data.**

### B. 7-Day Schedule Cross-Comparison Tool

The objective of the 7-Day Schedule Cross-Comparison Tool (7-DSC) is to enable mission planners to identify simultaneous host tracking and smallsat beam intercept opportunities in a single tool. This is accomplished by fusing data from official DSN tracking schedules with beam intercept information from the BIPS tool described above. Schedule data from the Service Preparation Subsystem (SPS) website and output from BIPS are copied to worksheets in a single Excel workbook. Visual Basic for Applications (VBA) code is then used to interpret and restructure the data so that host tracking opportunities can be identified and correlated with beam intercept information. The code also identifies cases where the host and smallsat are being tracked simultaneously by different DSN assets.[2] Results are provided in tables as well as graphically.

Official DSN tracking data are downloaded in Extensible Markup Language (XML) format from the SPS website.[3]

Several weekly schedules are downloaded in XML format and stored on a single worksheet. Since these schedules are changed frequently, it is important to use the most recent versions in order to accurately identify tracking opportunities. The data are then interpreted and stored in an object-oriented, internal data structure that includes (for each track): start and end times (to the nearest minute), setup and teardown times, and required equipment (processors, monitors, consoles, amplifiers, transmitters, etc.). MSPA and delta-differential one-way range (DDOR) tracks are included, but maintenance and calibration downtime "tracks" are ignored.

BIPS output is loaded into a separate worksheet. This tool provides, separately for each DSN complex, time intervals in which the smallsat is in the host beam. In-beam times include the effects of the local horizon mask (assumed to be 10 deg) and occultation of the smallsat by the planets (e.g., Mars) and Earth's moon. Intercept times are evaluated with a 10-s time resolution.

---

[2] This last capability is used to identify times when the smallsat is transmitting (to a different DSN antenna) and can be picked up by the antenna tracking the host. This capability is only used for demonstration purposes (i.e., when assuming Mars Odyssey to be the "smallsat").

[3] http://spsweb.fltops.jpl.nasa.gov/portalappsops/Scheduling.do?ft=mos. This is an internal JPL website.

Data summarizing host tracking opportunities are provided in tabular form, as shown in Figure 7. Any spacecraft tracked by the DSN may be used as the host spacecraft. In this particular case, MRO is used as the host. Controls are provided to initiate analysis, and adjust the amount of detail provided. The first row shows a track on 7/28/2014 using DSS-34 in Canberra lasting 6.75 h. Mars Odyssey (M01O) is also being tracked in MSPA mode for part of the track. A detailed account of the equipment used during the pass is provided as well.

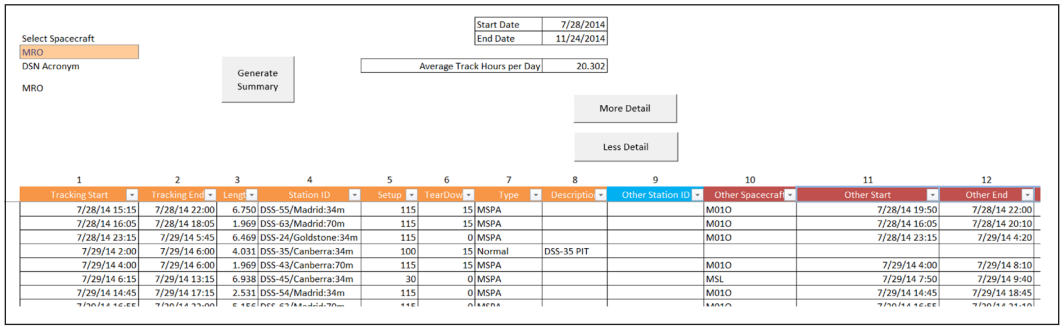


| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Tracking Start | Tracking End | Lengt | Station ID | Setup | TearDow | Type | Descriptio | Other Station ID | Other Spacecraf | Other Start | Other End |
| 7/28/14 15:15 | 7/28/14 22:00 | 6.750 | DSS-55/Madrid:34m | 115 | 15 | MSPA | | | M01O | 7/28/14 19:50 | 7/28/14 22:00 |
| 7/28/14 16:05 | 7/28/14 18:05 | 1.969 | DSS-63/Madrid:70m | 115 | 15 | MSPA | | | M01O | 7/28/14 16:05 | 7/28/14 20:10 |
| 7/28/14 23:15 | 7/29/14 5:45 | 6.469 | DSS-24/Goldstone:34m | 115 | 0 | MSPA | | | M01O | 7/28/14 23:15 | 7/29/14 4:20 |
| 7/29/14 2:00 | 7/29/14 6:00 | 4.031 | DSS-35/Canberra:34m | 100 | 15 | Normal | DSS-35 PIT | | | | |
| 7/29/14 4:00 | 7/29/14 6:00 | 1.969 | DSS-43/Canberra:70m | 115 | 15 | MSPA | | | M01O | 7/29/14 4:00 | 7/29/14 8:10 |
| 7/29/14 6:15 | 7/29/14 13:15 | 6.938 | DSS-45/Canberra:34m | 30 | 0 | MSPA | | | MSL | 7/29/14 7:50 | 7/29/14 9:40 |
| 7/29/14 14:45 | 7/29/14 17:15 | 2.531 | DSS-54/Madrid:34m | 115 | 0 | MSPA | | | M01O | 7/29/14 14:45 | 7/29/14 18:45 |

**Figure 7. Typical tabular output describing tracking opportunities for the host spacecraft.**

A similar, separate set of outputs is provided for DSN tracking of the smallsat. In actual practice, these tracks would be limited to uplink, navigation, and radio science tracking, since downlink telemetry would be obtained using OMSPA (and would not be formally tracked by the DSN). For this demonstration, however, all tracking of Mars Odyssey is provided, including telemetry downlinks (so that we know when our "smallsat" is transmitting).

Timing relationships between tracks of different spacecraft during MSPA can also be viewed graphically (Figure 8). Single and MSPA[4] tracks are correlated with beam intercept opportunities during the same time interval to give planners a complete picture of potential OMSPA opportunities. While detailed output of tracks correlated with beam intercepts is provided in tabular form (Figure 7), the graphical output provides a clear view of the timing relationships. The example in Figure 8 shows a traditional MSPA track of MRO and Mars Express (MEX) on 8/03/2014 using DSS-25 at Goldstone. The MRO track lasts 8.0 h, while the MEX track lasts 2.92 h, toward the end of the MRO track. The DSN is also tracking M01O in the same time frame, first on DSS-65 (Madrid) for 6.75 h. As this first track is ending, the traditional MSPA track of MRO, and eventually MEX, is just beginning. The DSN then switches to DSS-24 at Goldstone to continue tracking of M01O for another 7.67 h. As M01O orbits Mars, it is in the host beam (tracking MRO and MEX) for brief periods of 1.42 h, separated by short periods during which the "smallsat" is blocked by Mars. These beam intercept periods are shown as short bars at the bottom of Figure 8.

The situation shown in Figure 8 portrays a good opportunity to demonstrate OMSPA tracking of M01O in the host beam of MRO. MRO is being tracked by DSS-25 while, at the same time M01O is transmitting to DSS-24. Because the beam of DSS-25 covers more than the

[4] DDOR tracks are ignored because they are not an opportunity for OMSPA. DDOR tracks of host spacecraft orbiting Mars are difficult to make use of because they alternate between the spacecraft and a quasar.

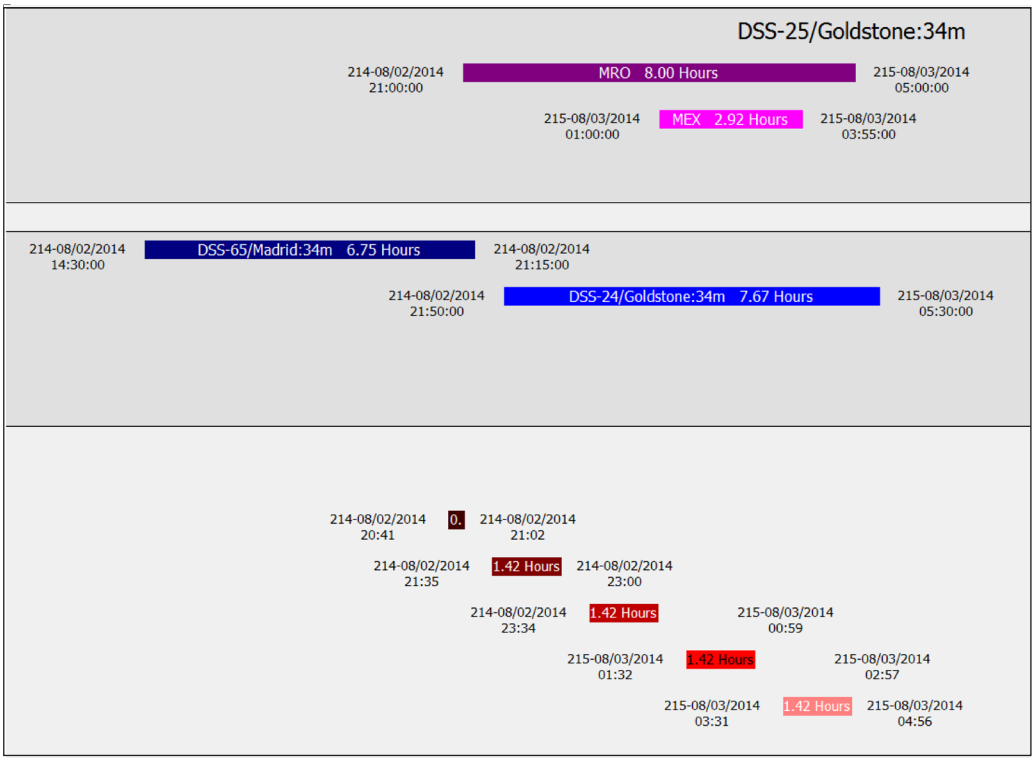


**Figure 8. Timing of host tracking of MRO and Mars Explorer (in traditional MSPA mode), along with tracking for M01O and the corresponding beam intercept times with the host tracking.**

planet Mars, the signal from M01O will be detectable in the signal received at DSS-25 using the recording and recovery techniques of OMSPA.

### C. Very Long Baseline Interferometry Science Receiver

For radio science and navigation (i.e., DDOR purposes, the DSN maintains a variety of open-loop receivers at the Signal Processing Center (SPC) for each complex. The analog IF output from each antenna can be connected with one or more of these open-loop receivers, depending upon the band and bandwidth requirements involved. Only the 26-GHz band is currently unsupported. The open-loop receivers then convert the analog IF signals into a digital format and, among other things, provide a recording capability. The various types of open-loop receivers available at each complex include two to three Radio Science Receivers (RSRs), one Very Long Baseline Interferometry (VLBI) Science Receiver (VSR), and two Wideband VLBI Science Receivers (WVSRs). Each one of these can process two IF inputs independently. An overview of the interfaces associated with the open-loop receivers appears in Figure 9. The RSR recorders make use of three disks with a capacity of 71.6 GB, 8.9 GB, and 4.4 GB, respectively. The VSR recorders also have three disks, each with a capacity of approximately 8.9 GB. And, the WVSRs each have a single disk with a capacity of approximately 1.7 TB. Specific details regarding each type of receiver, such as its bandwidth and channel characteristics, are available in [2]. Because the VSRs are older and less capable than the WVSRs, they tend to be used less frequently. Hence, in the interests of having open-loop receivers readily available for the demonstration, the VSRs were selected.

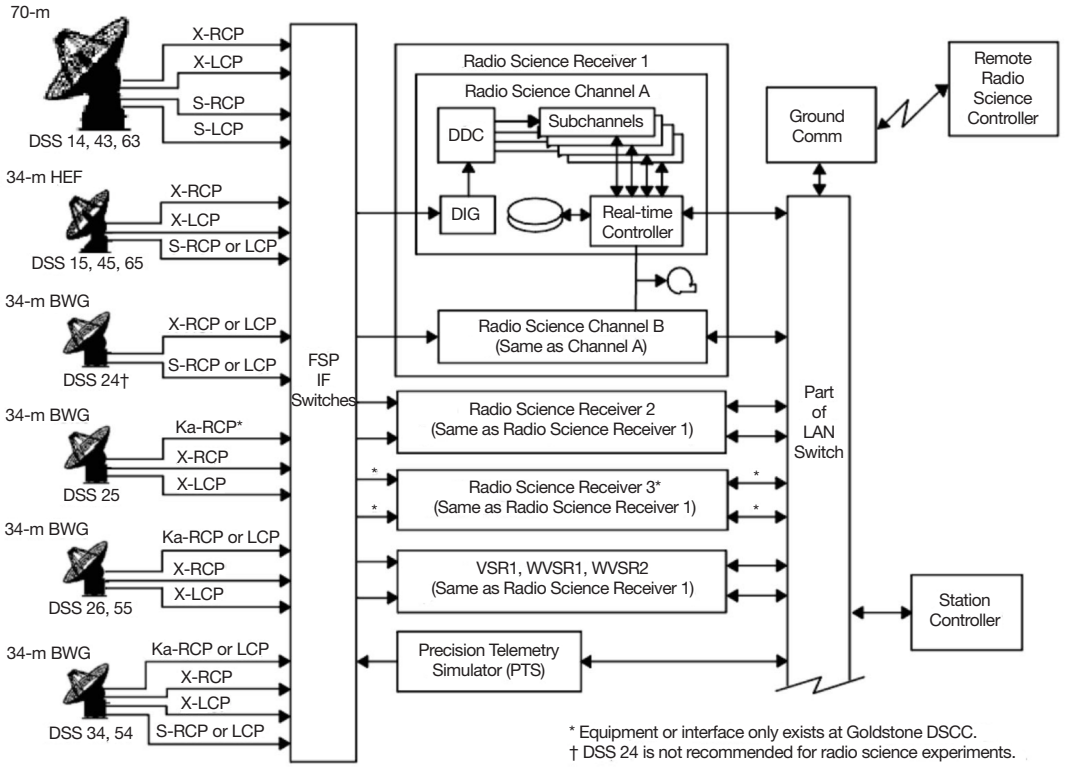


**Figure 9. Radio science receiving equipment configuration per [2].**

In the demonstration, the 7-DSC tool and the Mars Odyssey sequence of events were used to figure out the specific times when the spacecraft was in view of Earth and transmitting in a manner that could be captured on the "host" spacecraft's ground antenna. The VSR's record times on the "host" spacecraft's ground antenna were scripted accordingly. The center frequency was chosen to be in the middle of the frequency span by examining Mars Odyssey's downlink frequency predicts. The recordings were then executed according to the script.

### D. Server Repository ("lilypond")

When convenient, the data files were "copied" to a virtual server named "lilypond," which is administered by the DSN contractor Exelis at its Monrovia, California, facility. Whereas the VSR resides on the DSN's restricted Flight Operations network, "lilypond" resides on the JPL network — making it easier to retrieve the appropriate data file for demodulation and decoding. The directory, UNIX permissions, file sizes (in bytes), file dates and time, and recording name (by recorder type and time in year-day-hour-minutes-seconds) of the four files are shown in Figure 10.

### E. OMSPA Software Demodulator (OSD)

The software receiver used for the OMSPA demonstration is based on prior and ongoing investment by the DSN Advanced Engineering Office towards the development of an analysis and decoding tool for low-to-medium rate deep space telemetry downlink waveforms.

```
/dat/ftp/pub/lilypad/dsn/wvsr/omspa
total 56286701
-rw-rw-rw-+1  14400547200 Aug 11 14:39 VSR1A.1W1.14-220-222145
 -rw-rw-r--+1 14400547200 Aug 25 19:34 VSR1A.1W1.14-236-170128
 -rw-rw-r--+1 14400547200 Aug 29 00:21 VSR1A.1W1.14-240-080132
-rw-rw-r--+1  14404547352 Sep  8 11:02 26VSR1A.1W1.14-249-212127
```

**Figure 10. VSR file directory on lilypond.**

An early description of the background, conceptual architecture and functional design for the receiver can be found in [3]. The receiver is implemented using a set of custom developed signal processing functions written in MATLAB. The receiver typically operates from a recorded file interface based on the signal recording format of many of the DSN's open-loop recorders.[5,6] For this demonstration effort, the receiver software was augmented by integrating forward error correction (FEC) code decoders. The waveform type used in all the demonstration runs consisted of Mars Odyssey's residual carrier binary phase-shift keying (BPSK) modulation with a concatenated FEC code composed of an inner (15,1/6) convolutional and outer (252,220) Reed–Solomon code. Integration of these decoders required converting custom C-software (developed by the Communications Architectures and Research Section's Information Processing Group) into MATLAB functions.

In contrast to real-time receiver designs that are governed by fairly strict latency constraints, the software receiver is able to reprocess the recorded data multiple times and is not limited to operating on an incoming signal flow in a continuous, pipelined fashion. This flexibility is useful in deriving and refining various signal parameter estimates to improve the overall detection performance of the receiver. Another distinct aspect of this receiver is the independent, block-oriented nature of its processing. This simply refers to being able to fully process arbitrary time segments of the recorded data (i.e., waveform samples in, time-tagged transfer frames out). This characteristic provides scalability for increasing the aggregate processing speed by allowing concurrent processing of different data blocks. Typical desktop computers provide multicore processors that can ideally reduce processing time by up to the number of cores if the total processing memory requirements are satisfied by the amount of the computer's physical memory.

In Figures 11 and 12, both the receiver's functional processing steps and the corresponding data block segmentation are shown. For every data block, the first step performed is an open-loop frequency estimate of the residual carrier averaged over the block duration within an uncertainty range of ±100 kHz. The entire sample block is frequency compensated to allow approximate centering of the residual carrier tracking loop. The receiver next

---

[5] *Deep Space Mission System (DSMS) Subsystem Software Interface Specification*, "0367-DOR_WVSR Full Spectrum Processing Subsystem Wideband VLBI Science Receiver (WVSR) and Differential One-Way Range Subsystem (DOR) Interface," JPL 820-016, Rev. A (internal document), Jet Propulsion Laboratory, Pasadena, California, April 15, 2009.

[6] *Deep Space Network (DSN) External Interface Specification*, "0159-Science Radio Science Receiver Standard Formatted Data Unit (SFDU)," JPL D-16765, Rev. C (internal document), Jet Propulsion Laboratory, Pasadena, California, February 24, 2010.

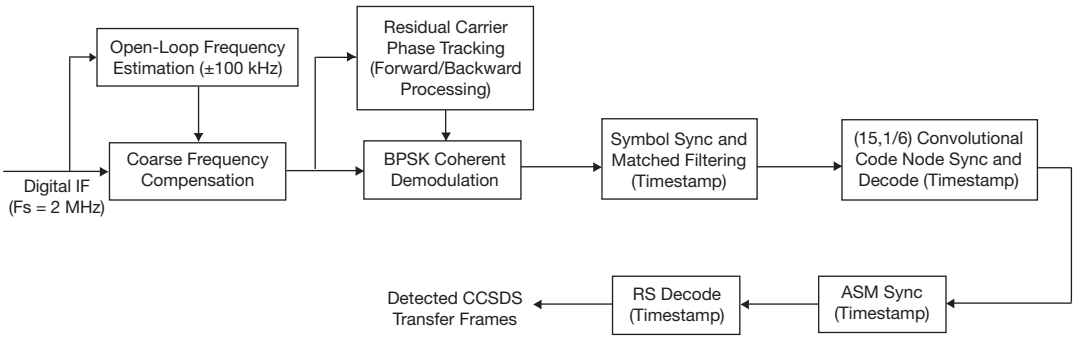


**Figure 11. Receiver/decoder functional block diagram.**

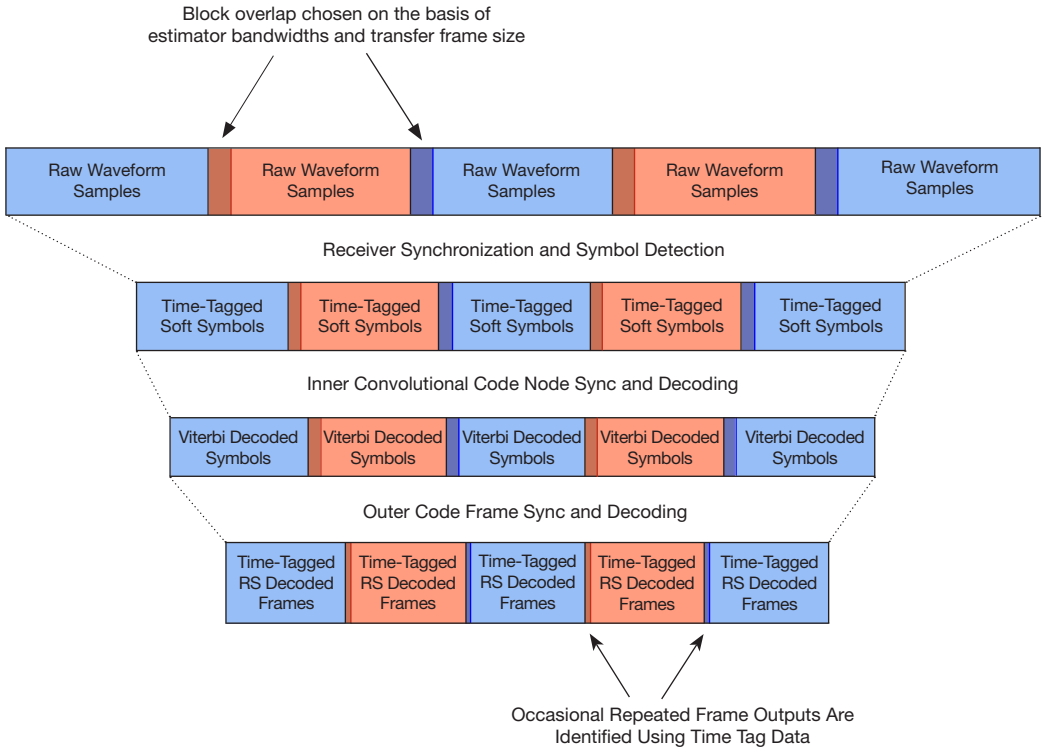


**Figure 12. Segmented processing.**

uses the phase-locked carrier to coherently demodulate the BPSK signal. Symbol synchronization is performed using a modulation independent delay-and-multiply technique. This processing also utilizes a two-pass approach, wherein the average symbol rate is estimated for the entire block and then used to re-center the tracking filter with a narrowband setting for improved noise rejection. Matched filter outputs are generated along with accompanying timestamps. The generation of soft decision symbols is followed by node synchronization for the inner convolutional code and Viterbi decoding. The output bit sequence is searched for a standard Consultative Committee on Space Data Systems (CCSDS) attached sync marker (ASM) and the timestamp associated with the leading bit of the ASM marker is saved as the transfer frame time tag. After Reed–Solomon decoding, both the systematic (information) bits of the codeword and the associated Earth received time (ERT) are written to file for eventual comparison to operational data provided by the Mars Odyssey project.

The segmented data view shown in Figure 12 shows how data blocks are organized at the input to the receiver and the subsequent flow as more levels of processing are executed. As noted above, the receiver processing flow concept is to have each input data block operate completely independently of the others to allow asynchronous, parallel processing for computing speed scalability. The block sizes could then be selected to optimize concurrent performance. For example, overall computer memory limitations might set limits on individual block sizes, as would the number of available computing cores. This might generally lead to fairly small block sizes. However, the nature of the receiver processing and signal memory overlap due to parameter estimator performance and FEC constructs suggest having a certain amount of overlap between data blocks to avoid lowered performance or missing data due to "edge effects." For this demonstration, the amount of overlap between blocks is related to two key parameters — the length of the Reed–Solomon code block and the integration time of the symbol synchronizer. Clearly, any data overlap between blocks results in some reduced processing efficiency that can be minimized by making the data segments larger. Consequently, there exists some tension between optimizing for communications performance and optimizing based upon computing constraints. For much of the processing performed for the OMSPA tests, the individual data block sizes ranged from 3 to 10 s. Another effect of the data overlap is the possibility of generating repeated transfer frames at the output of the block processing. These repeated frames are easily identified on the basis of time tags that are measured to fractional symbol times (submicroseconds).

One of the drawbacks of independent block processing is accommodating narrowband tracking loops and/or parameter estimators with long integration times, which may limit the minimum block size. To help address this challenge, the receiver applies forward/backward processing in the residual carrier tracking loop. An example of estimating the carrier phase trajectory using this process is shown in Figure 13. The real component of the narrowband filtered residual carrier is shown along with the cosine arms of the forward phase-locked loop (PLL) estimate (in red) and the backward PLL estimate (in green). Figure 14 provides a close-up view of the PLL processing at the beginning of the data segment and at the end. Note that the poor phase estimate generated during the forward loop acquisition period is substantially improved by running the loop backwards from the end point of the data block. In the receiver processing for the OMSPA demonstration, the individual data block sizes and loop bandwidths were chosen such that the forward/backward cycle was exercised only once. In future work, we intend to investigate running this cycle multiple times for very short block lengths.

## V. Procedure

Sections III and IV provide a general sense of the nature of the demonstration and the tools needed to identify, initiate, record, and post-process each pass. This section enumerates the specific steps that the investigators followed during the execution of each demonstration pass.

(1) *Identification* of potential intercept opportunities using BIPS and 7-DSC.

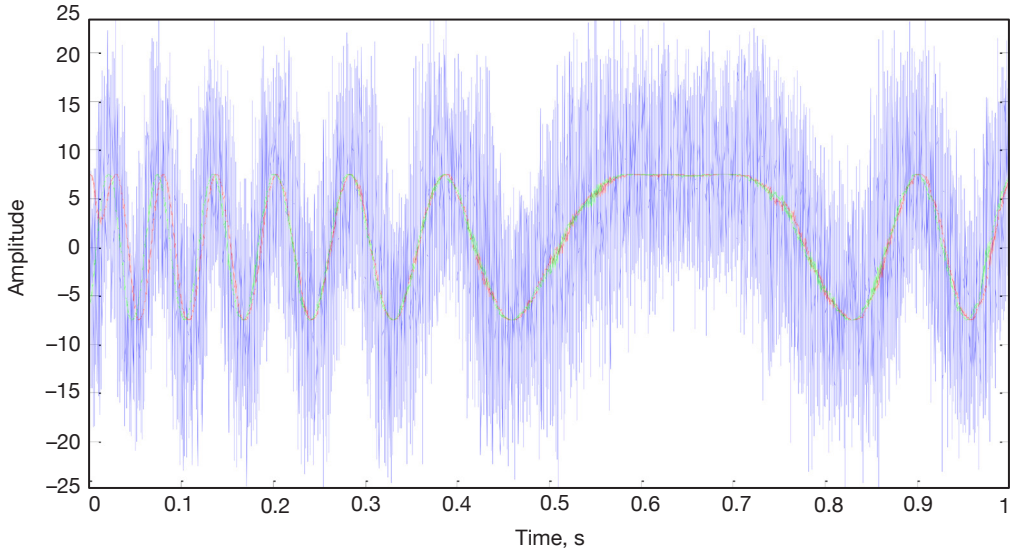


**Figure 13. Forward/backward phase-locked loop tracking.**

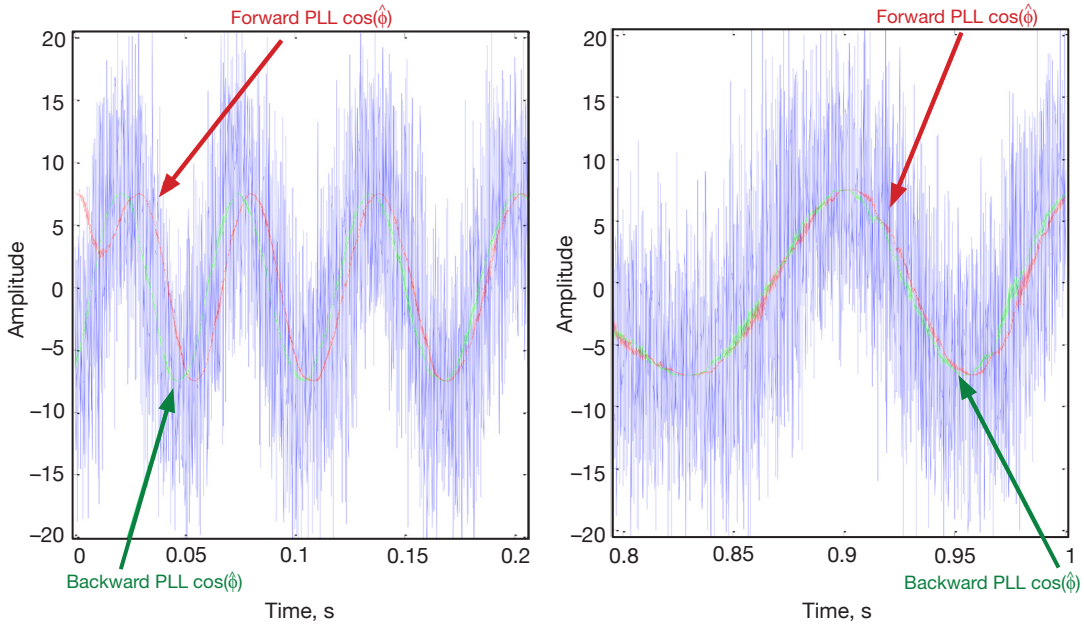


**Figure 14. Forward/backward phase-locked loop tracking (beginning and end of block details).**

(2) *Selection* of the best candidate for the desired demonstration pass. Different passes can involve different data rates, coding schemes, etc. Hence, the investigators discussed the relative merits of some of these pass aspects before selecting a particular pass intercept opportunity. They also checked the 7-day schedule for any possible changes to it.

(3) *Scripting* the recording interval on the VSR. The script contains the IF input, the channel width, the center frequency, and record times. It is scheduled for execution just prior to the record time. While one might expect the recording interval to span the entire intercept opportunity, predemonstration test recordings showed

that the "smallsat" (Mars Odyssey) transmission did not always terminate when BIPS and 7-DSC projected. This was generally due to factors pertaining to Mars Odyssey operations that the investigators could neither know nor control — such as the onboard spacecraft antenna reaching its cable-wrap limit prior to the end of the scheduled transmission time. Presumably a real smallsat and its operations team would know such things and could build it into their planned recording times. For the purposes of the demonstration, however, the recordings were scheduled for an hour of downlink that examination of Mars Odyssey's (our "smallsat's") sequence of events indicated would involve nominal transmission.

(4) *Playback* of the recording to the server repository ("lilypond"). Since the VSR is located at each complex's SPC, the recording was played back across the Flight Operations Network firewall to a secure server within the JPL firewall, "lilypond." The playback times were carefully documented since they constituted one of the largest components of the total turnaround time, from open-loop transmission to completion of demodulation and decoding. Documentation of the associated file size was also integral to turnaround time assessment.

(5) *Retrieval* of the recording for signal processing. This step pertained to the process of actually retrieving the recording from the secure server ("lilypond") and moving it to a workstation for signal processing. As with the playback time, retrieval times were documented, given their contribution to the total turnaround time.

(6) *Processing* of the recording to achieve demodulation and decoding. This step actually consisted of three substeps. First, investigators performed a "quick-look" file assessment to verify that the desired portion of the recording displayed a "normal" signal profile. To the extent that it did, the investigators then ran the OMSPA software demodulator base code on the recording input. Finally, the demodulated and decoded transfer frame output was carefully compared to Mars Odyssey's traditionally downlinked transfer frames to validate that the signal processing had been correctly performed. The time required for the first two substeps was documented and factored into the computation of the total turnaround time. The third substep was not included in the time computation since the signal processing validation was unique to the demonstration and should not be necessary for subsequent use of this version of the OMSPA software demodulator.

These six steps were repeated four times — once for the recording corresponding to the traditional MSPA pass involving Mars Odyssey and MRO, and three times for the case where Mars Odyssey's transmissions to its own ground antenna were being intercepted and recorded via MRO's antenna. The latter three occurrences corresponded to one pass for each of the three complexes: Madrid, Canberra, and Goldstone.

## VI. Results

### A. Playback

Table 1 shows the time required to move each recording from its respective VSR to "lilypond." Because the recordings of the host spacecraft's antenna were always set to record

**Table 1. Demonstration recording dates, times, and durations.**

| DOY | Date | DSS | Record Start | End | Playback Start | End2 | Duration | Notes |
|---|---|---|---|---|---|---|---|---|
| 221 | 8/9/14 | 26 | 355 UTC | 455 UTC | 223/19:17 | 21:39 | 2:22 | MRO harmonic |
| 236 | 8/24/14 | 55 | 1720 UTC | 1820 UTC | 238/16:53 | 239/02:34 | 9:41 | 2 way |
| 240 | 8/28/14 | 45 | 820 UTC | 920 UTC | 241/00:40 | 7:20 | 6:40 | |
| 249 | 9/6/14 | 26 | 2130 UTC | 2230 UTC | 251/16:39 | 18:02 | 1:23 | |

for 1 h from the desired start time and did so at 2 MHz with 8-bit sampling, the file sizes were always around 14.4 GB. Hence, the average effective transfer rate from Goldstone was roughly 2.13 MB/s. The effective transfer rate from Madrid was on the order of 413 kB/s. The effective transfer rate from Canberra was on the order of 600 kB/s.

### B. Retrieval

Data retrieval from the "lilypond" server was relatively consistent for each waveform recording from DOYs-221, -236, -240, and -249. The time to transfer each data file to a local processing machine was about 53 min. The average transfer rate across the network from "lilypond" was slightly less than 5 MB/s. In combination with the playback times, the total time up to processing ranged between 2.3 h and 10.6 h.

### C. Processing Overview

As noted in Section V, each recording was "quick look" processed for a nominal quality assessment. The processing consisted of spectral analysis and total input power reports for the duration of the pass. Estimates of the residual carrier frequency, power, and total input signal power were made in 1-s increments using a resolution bandwidth of 10 Hz. Results for the DOY-221 pass are shown in Figures 15, 16, and 17. Since the IF waveform recordings were captured in a "frequency predict-less" manner, the residual carrier varied over a peak-to-peak uncertainty range of 140 kHz, which was accommodated within the software demod processing using the coarse frequency compensation previously described. In Figure 16, the large reduction observed in the residual carrier power during the middle of the pass was really an artifact due to the narrow-resolution bandwidth of the spectral analysis processing and the integration time. During the middle of the pass, the rate of change of the downlink frequency was greatest and, consequently, the carrier power was additionally distributed across adjacent frequency bins, thereby lowering the observed power in any given bin. As shown in Figure 17, the total input power varied by less than 2 dB for the entirety of the pass and, in general, exhibited much smaller variation. Overall, the DOY-221 quick-look assessment did not reveal any artifacts, signal outages, distortions, or other obvious concern within the recorded data set.

After the quick-look assessment, the data were processed in whole or in part on two different computing stations. The general capabilities of the systems used for the demodulator processing are described below.

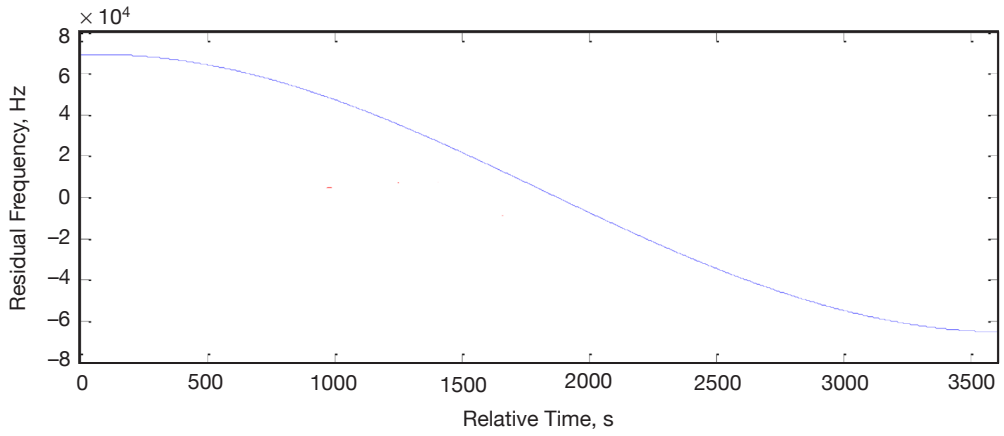


**Figure 15. DOY-221 residual carrier frequency estimate (10 Hz resolution, 1 s average). The data record was captured without frequency predict compensation.**

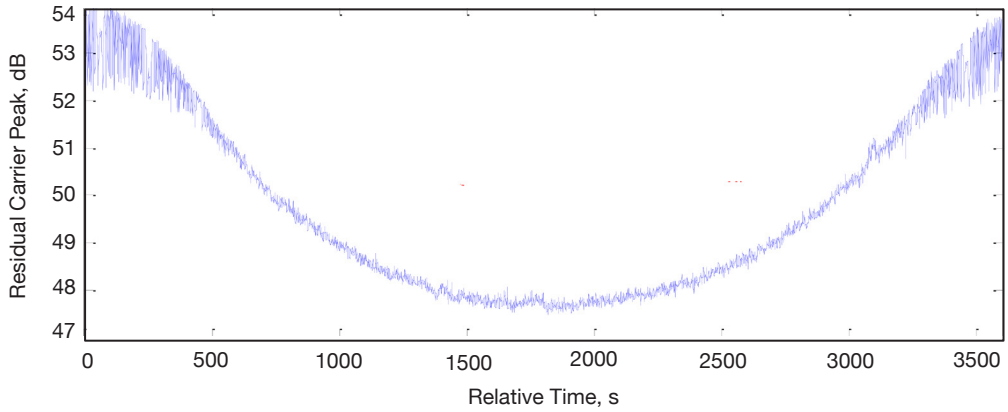


**Figure 16. DOY-221 residual carrier frequency estimate (10 Hz resolution, 1 s average). The dip in peak power is likely due to narrow bin width and high rate of change of carrier frequency near the center of the observation interval.**

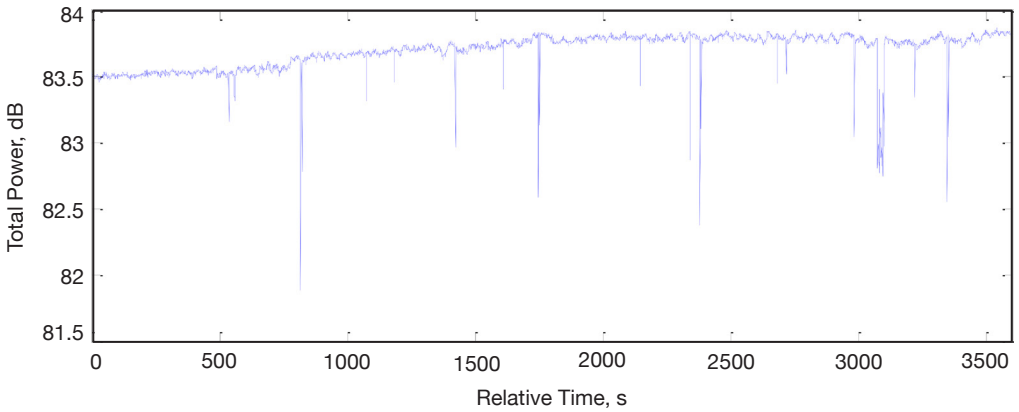


**Figure 17. DOY-221 total received signal power (10 Hz resolution, 1 s average).**

*Desktop computer #1*

| | |
|---|---|
| • Number of processors | 1 |
| • Processor type | Intel Xeon CPU E3-1245 v3 |
| • Clock speed | 3.4 GHz |
| • Number of cores per processor | Four |
| • Total physical memory | 16 GB |

*Desktop computer #2 parameters*

| | |
|---|---|
| • Number of processors | 2 |
| • Processor type | Intel Xeon CPU E5-2650 0 |
| • Clock speed | 2.0 GHz |
| • Number of cores per processor | Eight |
| • Total physical memory | 64 GB |

In general, the bulk of the processing was performed using a single core on computer #1. This was chosen mainly from a convenience standpoint as this machine was the primary development platform for the processing and validation code and was also used as the main repository of the recorded data files transferred from "lilypond." The base processing time using a single core for each of the recorded waveform data sets was in the 6× real-time range (i.e., 5 to 6 h elapsed time for a 1-h recording). Key factors that govern the processing rate include the sample rate, data rate, and FEC code type. In particular, for the Mars Odyssey signaling format, decoding the (15,1/6) convolutional inner code was a significant contributing factor to the overall processing time. Note that this code is approximately 256 times more complex than a standard (7,1/2) convolutional code. On the demodulator side, the software algorithms utilize multirate signal processing to improve computational efficiency. Ratios for up and down sampling at different stages of the receiver depend upon the recorded waveform sample rate, uncertainty in the signal frequency, and the data rate. As a general proposition, we expect to be able to accelerate the effective processing speed in a proportional manner as multiple processor cores are employed.

A number of benchmarking tests were run with computer #2 to test out this concept. In these tests, from 3 to 5 cores were utilized to process different portions of a waveform recording. These were set up manually and the elapsed time recorded and compared to a single processor measurement. For this platform, the single processor processing rate ranged from 8 to 9× real-time. When subdividing the processing into multiple cores, a linear improvement was obtained as long as the overall memory utilization did not exceed the available physical memory (RAM) of the platform. For these tests, the amount of RAM was, in fact, the limiting factor in the number of cores employed. In determining the software's memory utilization, we identified the (15,1/6) Viterbi decoder as the component needing the majority of the allocated RAM. It was further determined that the memory utilization by the decoder could be drastically scaled back with some minor modifications to its processing algorithm.

### D. Processing Observations

The outputs of the baseline processing for each data set were two files consisting of the data payload of the demodulated and decoded transfer frames and a corresponding set of ERTs

tagged to the first bit of each frame's attached sync marker. Failures to decode were also recorded based upon the status provided by the Reed–Solomon outer code decoder. The ERTs were then analyzed to determine if there were any gaps in the time tags indicating frames missed by the software. Analysis of the results from early processing trials revealed three different types of errors:

- "Overlap" errors resulting in an occasional missed frame due to insufficient overlap of data segments (as depicted in Figure 12).
- Open-loop frequency estimation errors due to the presence of unexpected tones in the signal at signal levels higher than the residual carrier and within the large frequency uncertainty range (±70 kHz).
- Demodulation errors due to nonoptimized carrier and symbol synchronizers in the presence of signal fluctuations.

Of these, the "overlap" errors were straightforward to fix by reprocessing the segments where the missing frames occurred. Such errors can be eliminated by extending the overlap region between data segments.

Some examples of the frequency estimation errors can be seen in the quick-look processing of the DOY-249 pass data in Figure 18, which is basically replicated on a per-data segment basis during demodulator processing. The excursions from the typical nonpredict-compensated frequency trajectory were measured at ±120 kHz offsets. The origin of these estimation errors was traced to the presence of tones in the Mars Odyssey downlink signal. A typical downlink telemetry spectrum is shown in Figure 19, and the signal spectrum with these additional tones is shown in Figure 20. The tones appear as odd-order harmonics and, upon closer measurement, exhibit spectral content consistent with a square wave appearing at half the channel symbol rate of 238,898.5 bps (i.e., consistent with an alternating 1/0 data pattern of coded symbols). Recalling Figure 11, the consequence of the open-loop frequency estimation errors was such that the gross data segment frequency compensation would be incorrectly applied, resulting in the loss of that segment. An algorithm was developed to extract accurate frequency compensation estimates, even in the presence of data tones, and the affected data segments were reprocessed.

As noted earlier, the third observed error type was analyzed and appeared to correspond to a short-duration drop in signal level. An example of the power in the demodulated data channel (i.e., a measurement of the matched filter output) over a 10-s processed segment is shown in Figure 21. In some of the investigated cases, this signal condition resulted in a bit slip that would throw off the subsequent node synchronization of the Viterbi decoder. Other times, there were no detected slips, just a short-duration reduction in effective signal-to-noise ratio (SNR). To compensate, a limited amount of parameter tuning of the demodulator was performed. Table 2 provides details of synchronizer bandwidth settings for both the carrier tracking loop and the symbol synchronizer, both for initial and tuned settings. Owing to limited analysis time, the demodulator was not exhaustively tuned and a small number of residual frame errors were observed with the modified tracking band-

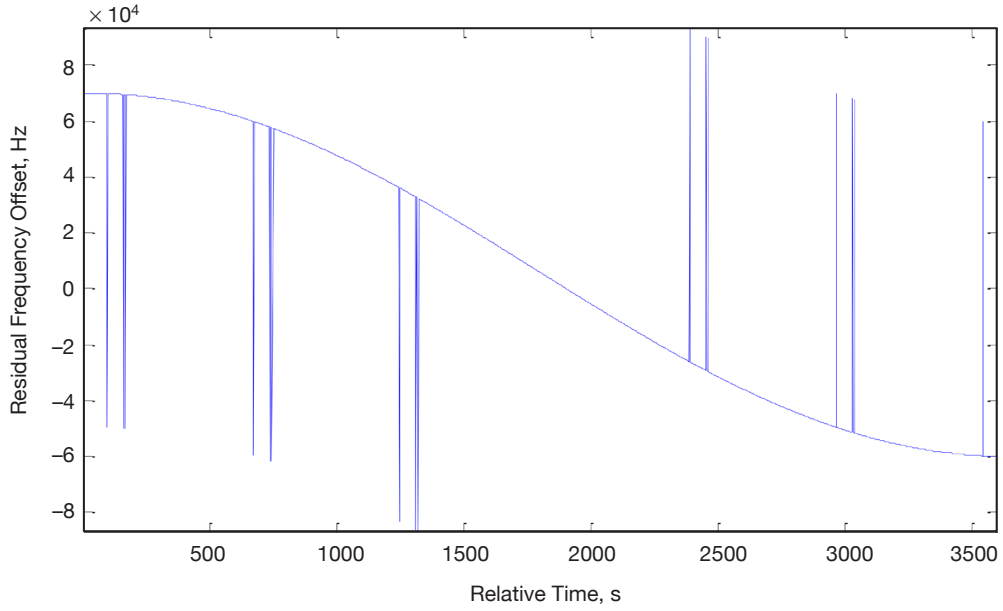


**Figure 18. DOY-249 residual carrier frequency estimate (10 Hz resolution, 1 s average).**

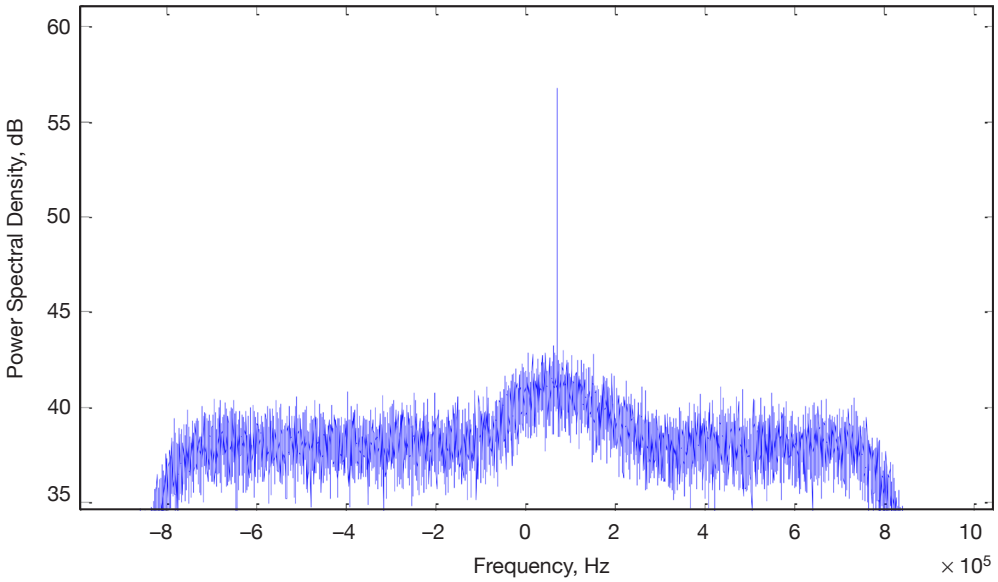


**Figure 19. DOY-249 downlink spectrum (expected telemetry and residual carrier only).**

widths. Note that the tuning also involves selection of the data segment size, since the coarse frequency compensation applies only a single correction across the entire segment. Consequently, the initial frequency offset conditions for the residual carrier PLL have a dependency on the length of the data block.

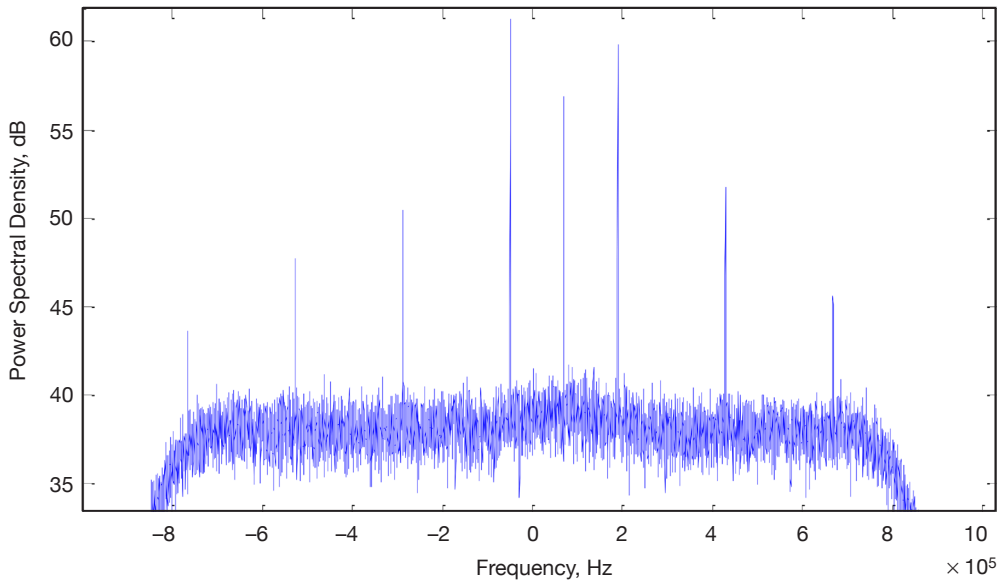


**Figure 20. DOY-249 downlink spectrum (unexpected tones).**

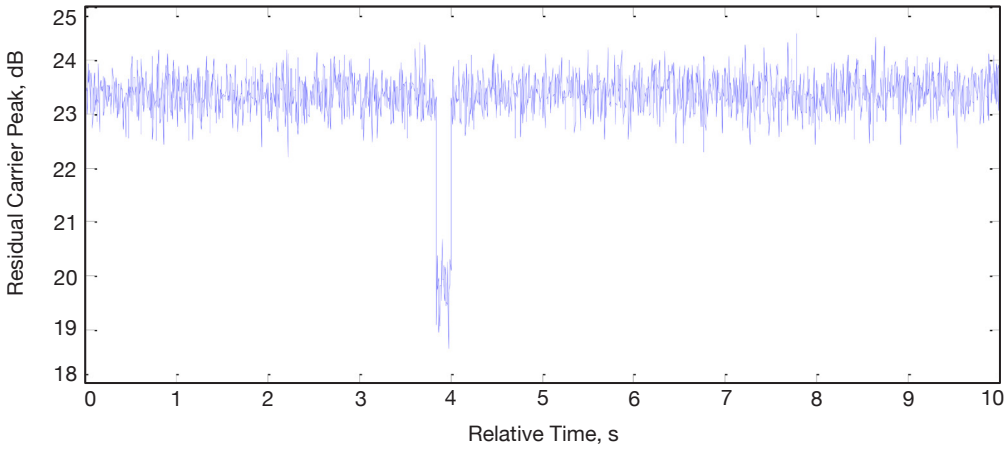


**Figure 21. DOY-221 example of intermittent signal outage observed in data channel power.**

**Table 2. Software demodulator settings.**

| Configuration | Carrier Tracking Loop Bandwidth | Symbol Synchronizer Bandwidth |
|---|---|---|
| Initial Settings | 100 Hz | 100 Hz |
| Partially Tuned Settings | 20 Hz | 5 Hz |

### E. Data Validation

Final validation of the results produced from the software demodulator/decoder processing of the OMSPA waveform recordings was performed by comparing the transfer frame results to the equivalent data (i.e., transfer frames and ERTs) generated by the DSN operational ground telemetry processing systems and provided to this effort by the Mars Odyssey project. A summary of these comparisons is shown in Table 3. We conjecture that the frame errors present in the DOY-236 and DOY-240 data sets may be eliminated if we apply a finer-grained frequency compensation approach to the recorded data sets. That is, measure frequency offsets at a fine time resolution and compensate for these offsets continuously across a data segment (rather than once) to allow use of much narrow tracking bandwidths.

**Table 3. Data validation summary.**

| 2014 DOY (Antenna) | Nominal Symbol Rate | Waveform Recording Interval | Number of Operational Transfer Frames | Number of Validated Transfer Frames |
|---|---|---|---|---|
| DOY-221 (DSS-26) | 398,164 bps | 03:54:59 – 04:54:58 | 23624 | 23624 |
| DOY-236 (DSS-55) | 398,164 bps | 17:20:00 – 18:19:59 | 23624 | 23619 (5 frame errors) |
| DOY-240 (DSS-45) | 398,164 bps | 08:19:59 – 09:19:58 | 23624 | 23613 (11 frame errors) |
| DOY-249 (DSS-26) | 238,898.5 bps | 21:29:59 – 22:29:59 | 14174 | 14174 |

## VII. Conclusions and Recommendations

As noted in Section II, the demonstration was designed to show four things:

(1) Given access to the DSN station schedules, associated trajectory files, and an appropriate tool, a user mission can compute the beam intercept times for its trajectory relative to that of a host spacecraft such that it has the capability to arrange its downlink operations to accord with those times.

(2) Given the tool in item #1, a user mission can successfully retrieve the portions of the wideband recordings corresponding to its signal via a secure Internet site.

(3) Given an appropriate tool, a user mission can successfully demodulate, decode, and recover its data from the retrieved recordings within a timeframe that is operationally reasonable for the user mission.

(4) A user mission can successfully do these things irrespective of which DSN complex is in view during the pass.

With respect to objective (1), by the time of the formal demonstration passes, BIPS and 7-DSC were able to compute beam intercept and associated record times to within a minute or so of the observed transmission times. Earlier practice passes provided an opportunity to stabilize the programming in 7-DSC and to add the treatment of light-time delay effects in BIPS. With both tools stable and meeting all immediate needs, the only difficulties that arose had to do with the predicted opportunity end times. Since Mars Odyssey was not actually our own smallsat, there were operational features such as the cable wrap limits on

its high-gain antenna (HGA) of which we were unaware and unable to factor into our end time estimates. To the extent that a real smallsat project would know and understand these operational factors, it would be in a position to adjust the BIPS and 7-DSC estimates accordingly. Hence, the demonstration clearly succeeded with respect to objective (1).

With respect to objective (2), the arrangement with "lilypond" successfully provided a secure site to which the VSR demonstration recordings could be transferred and from which they could be retrieved for subsequent demodulation and decoding on the OSD workstation. However, the time needed to transfer the recordings from the VSRs to "lilypond" varied dramatically, depending upon the location of the VSR and the nature of the competing traffic on the wide-area network (WAN). As noted in Table 1 in Section IV.A, transfer times ranged from 1 h 23 min from Goldstone to 9 h 41 min from Madrid, with the total file size in both cases being about 14.4 GB. This translates into effective average transfer rates of 2.13 MB/s from Goldstone and 413 kB/s from Madrid. While the Goldstone rate seemed reasonable, neither the Madrid nor the Canberra rates did. The demonstration team engineer responsible for the recordings believes the problem is associated with the relative bandwidth priority allocated to the VSR data on the WAN. Typically, VSR data does not receive the same priority as spacecraft telemetry. And, during test recordings, DDOR also appeared to receive higher priority than the VSR data (and resulted in even longer playback times). The DSN systems engineer has been provided with the data and is looking into the matter.

On top of the VSR-to-"lilypond" transfer rates, another key aspect of objective (2) was the time needed to retrieve the recording from "lilypond" for demodulation and decoding on the OSD workstation. This duration was consistently around 53 min, which translates into about 5 MB/s on JPL's local-area network (LAN). Combined with the VSR-to-"lilypond" transfer rates, the total time up to demodulation ranged from 2 h 16 min to 10 h 34 min for 1 h worth of Mars Odyssey data (~14.4 GB). Hence, a full 8-h pass could be ready for demodulation within a 24-h period at the Goldstone rate, but not at the Canberra or Madrid rates. However, one should note that real smallsat data rates tend to be at least 15 times lower than Mars Odyssey's, with only a rate 1/2 coding overhead instead of a rate 1/6 coding overhead. Hence, the required bandwidth for the recording diminishes somewhat proportionally. So, even the demonstration's Canberra and Madrid transfer rates become operationally tenable under those circumstances. Nonetheless, solving the WAN traffic prioritization and associated architecture issues should be a key component of any follow-on to this demonstration.

With respect to objective (3), the OSD enabled successful demodulation, decoding, and recovery of the Mars Odyssey data in a timeframe that was roughly 5 to 6 times longer than the recorded duration. As discussed in Section VI.C, this processing time can be accelerated by employing multiple processor cores. And, real smallsat missions tend to use data rates that are at least 15 times lower than Mars Odyssey's, as well as FEC codes that are 256 times less complex, enabling much faster processing even without multiple processor cores. Of all the signal processing challenges that were encountered during the test and demonstration passes, the only residual errors that could not be entirely eliminated through algorithm modifications were those due to signal fluctuations. While a limited amount of parameter tuning of the demodulator was attempted, analysis time limitations precluded exhaustive

tuning, which would have involved revisiting the fundamental data segment size used in the processing. Nonetheless, the worst-case demonstration pass had only 11 unresolved frame errors out of 23,624 — equivalent to 99.95 percent data return.

With respect to objective (4), Mars Odyssey data was successfully recovered irrespective of which DSN complex was in view during the pass. Total data recovery times, from VSR transfer all the way through data demodulation and decoding, ranged from about 8 h to 19 h, largely depending on whether the recording had to be transferred from Goldstone, Canberra, or Madrid. While this is well under the "day or two" of data latency initially advertised with the OMSPA concept, any follow-on activities to this demonstration should strive to further reduce the turnaround time by:

(1) Solving the WAN traffic prioritization and associated architecture issues that are leading to asymmetric data return rates from the Deep Space Communications Complexes. (This is not only an issue for OMSPA, but also a potential issue for future DSN forays into optical communications.)

(2) Developing a recommended set of standard hardware (e.g., specific multiple core processors) and software (e.g., specific MATLAB toolsets or conversion to executable C code) to employ during signal processing in order to achieve various predefined categories of operational performance.

Beyond these recommendations, future efforts should probably focus on refining the choice of fundamental block segment size to be used in the OSD, expanding the software's applicable demodulation and decoding capabilities to such things as quadrature phase-shift keying (QPSK) and low-density parity-check (LDPC) codes, and, in so doing, expanding OSD's utility to higher data rate spacecraft. While OMSPA was originally conceived as a means for providing a simultaneous low-cost downlink service to multiple smallsats that happen to be within the beam of a single 34-m antenna, OMSPA has the potential, in the case of routine downlink, to also enable larger, high-data-rate Mars spacecraft to make more efficient use of DSN assets than the current 2-MSPA capability allows. As other solar system destinations become more heavily trafficked, OMSPA may have similar applicability for those regions as well.

## Acknowledgments

We gratefully acknowledge predemonstration advice from Faramaz Davarian, Dayton L. Jones, Charles J. Naudet, and Stephen P. Rogstad. The demonstration was funded by the DSN through the DSN Technology Office. Data validation support was funded by the Mars Odyssey project. Additional software receiver and decoder development support was provided by Kenneth Andrews, Meera Srinivasan, and Albert Zhu. The demonstration sponsor was E. Jay Wyatt, Manager, Space Networking and Mission Automation Technology Program Office (973).

## References


[1] D. S. Abraham and B. E. MacNeal, "Low-Cost Communications Concept for Smallsats: Opportunistic MSPA," *NASA Tech Briefs*, vol. 38, no. 5, pp. 56 and 58, May 2014.

[2] A. Kwok, "Open Loop Radio Science," *DSN Telecommunications Link Design Handbook*, DSN No. 810-005, Radio Science, Module 209, Rev. A, Jet Propulsion Laboratory, Pasadena, California, June 1, 2010.

[3] N. Lay, M. Lyubarev, A. Tkacenko, M. Srinivasan, K. Andrews, et al., "Software Receiver Processing for Deep Space Telemetry Applications," *The Interplanetary Network Progress Report*, vol. 42-180, Jet Propulsion Laboratory, Pasadena, California, pp. 1–8, February 15, 2010.
http://ipnpr.jpl.nasa.gov/progress_report/42-180/180E.pdf