TMO Progress Report 42-134 August 15, 1998

Remote Operations of the Deep Space Network Radio Science Subsystem

J. Caetta,[1] S. Asmar,[1] S. Abbate,[2] M. Connally,[1] and G. Goltz[1]

The capability for scientists to remotely control systems located at the Deep Space Network facilities only recently has been incorporated in the design and implementation of new equipment. However, time lines for the implementation, distribution, and operational readiness of such systems can extend much farther into the future than the users can wait. The Radio Science Systems Group was faced with just that circumstance; new hardware was not scheduled to become operational for several years, but the increasing number of experiments and configurations for Cassini, Galileo, Mars missions, and other flight projects made that time frame impractical because of the associated increasing risk of not acquiring critical data. Therefore, a method of interfacing with the current radio science subsystem has been developed and used with a high degree of success, although with occasional problems due to this capability not having been originally designed into the system. This article discusses both the method and the problems involved in integrating this new (remote) method of control with a legacy system.

I. Introduction

The first use of an open-loop receiving capability by the Deep Space Network (DSN) for science purposes was to record occultation data during the Mariner IV flyby of Mars in 1965. Although that system originally had been implemented to augment the DSN signal-tracking capability, the open-loop data soon became a primary source for significant scientific research, such as atmospheric analysis, the search for gravitational waves, and solar science. Between 1965 and 1993, significant changes to the hardware and software were made to accommodate the requirements of new spacecraft.

However, during this same period of time, a trend was noticed by the Radio Science Systems Group (RSSG) and associated investigators. Specifically, for the Pioneer Venus orbiting mission, the success rate of the radio science experiments was approximately 50 percent. When the RSSG started to monitor in real time the configuration and operation of the open-loop receiver, collectively known as the radio science subsystem (RSS), the success rate soared to around 95 percent. Over the years, real-time monitoring by the RSSG became a part of almost every radio science experiment. But this use of radio science personnel was threatened by the decreasing mission operations funding; from a team of 8 full-time equivalents for

[1] Communications Systems and Research Section.

[2] TMOD Operations Program Office, Allied Signal.

Voyager, the radio science team fell to 4 full-time equivalents for Galileo, and is currently at 1.5 full-time equivalents for Mars Global Surveyor (MGS). When these team levels are normalized against the number of experiments (one dozen for Voyager, two dozen for Galileo, and over 8,000 planned for MGS), the threat to the ability of the RSSG to monitor every experiment in real time is apparent. Although the proposed full-spectrum recorders, the hardware replacements for the RSS, would increase the reliability of the execution of radio science experiments, the recorders were not scheduled to be available until after the year 2000. Something else had to be done to ensure an acceptable success rate while attempting to stay within the funded staffing levels.

The answer came in a two-part package: to merge the responsibility of configuration and operation with the resources used to monitor experiments and to streamline and automate experiment operations as much as possible. This solution both addressed the decreasing-resources issue and allowed scientists to schedule experiments without adding to the work load of the station operations personnel needed to correctly configure and operate the RSS for extended periods of time. In some cases, scientists had been told that certain experiments could not be done due to a nonstandard configuration, as opposed to technical difficulties. Time constraints occasionally prevented responsiveness by station personnel to configuration details in the station sequence of events (SOE). Because this new solution would use radio science personnel as the RSS operator, the necessary awareness and sensitivity to changing configurations would be of a much higher level than that of personnel who had to understand every single subsystem at a DSN antenna site.

This solution commonly is called the radio science remote operations capability, or "remops." It uses a workstation-based client/transaction server interface to the radio science subsystem's front-end hardware component, the radio science control processor (RSCP), to send user configuration and operations directives from JPL. Monitor data are returned both by querying the RSCP and by inspecting the broadcast monitor data available on networks inside the JPL flight operations firewall. It has been in use, with increasing stages of capabilities, since late 1996, and is currently the primary tool for all radio science projects and experiments. During the MGS aerobraking phase experiments, only 4 of the first 125 experiments were lost due to operator error—a success rate higher than that of the previously monitored experiments, and with a decrease in the radio science personnel time invested in each experiment. Additionally, the ability to smoothly transition the hardware from one experiment to another (and one project to another) has greatly increased because of the shortened lines of communication for such a change to occur. Overall, the success of the remops project makes it a clear solution for those DSN systems that are facing many resource constraints.

II. The Background

During the Mars Observer (MO) experiment planning phase, the need for a capability to support large numbers of experiments with a minimum of personnel was recognized by the Radio Science Systems Group. From this, the RSSG envisioned a solution to provide both remote-control and local-control capabilities to the users of the radio science subsystem. However, the hardware and software of the RSS were not designed with the capability of remote operations in mind and only supported local control.

A. Overview

1. Hardware. Figure 1 shows the network and serial connections at a Signal Processing Center (SPC) as it applies to the radio science subsystem. The RSCP has several open serial ports and one unused ethernet port. Station personnel can connect a terminal to the console serial port in order to command the hardware locally.

2. Software. Commanding: All commands for the radio science equipment were entered at the link monitor and control (LMC) station. The commands were encapsulated into the common software

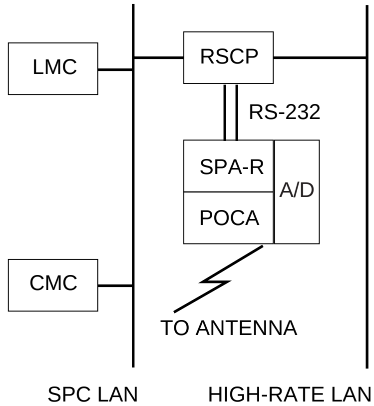

Fig. 1. SPC network and serial connections.

protocol and sent via the SPC low-rate local area network (LAN) to the RSCP, which then parsed the commands.

Monitor: Responses from the RSCP and the signal-processing assembly, radio science (SPA-R) were sent back to the LMC via the same common software protocol over the SPC LAN. When the RSCP was placed in a "link," a state where it was considered part of a virtual system, monitor data from the RSCP also would be broadcast on the high-rate LAN, where it would be routed back to both the LMC and JPL.

Support: An auxiliary file was needed in order to record spacecraft data correctly. This was the "predict" file, a set of time–frequency pairs that were transferred in a binary format to the RSCP via the command, monitor, and control (CMC) station.

B. Operational Scenarios

An end-to-end description of all actions required for each radio science experiment before implementation of remops is shown in Fig. 2. The operation of the radio science equipment was based on an SOE being generated by the project and sent to the station. During the prepass calibration period, or pre-cal, the station was to refer to the SOE to determine the radio science parameter values for the current experiment. Additionally, they were to load tapes into the three tape drives and distribute and load the indicated predict set into the RSCP. If the RSCP was configured correctly and was functioning normally, it would record data to disk (and optionally to tape) between the specified start and stop times without further interaction with the operator.

During the experiment, the operator would monitor the analog-to-digital converter (A/D) maximum and minimum values to ensure that the digitizers were not being saturated due to too little attenuation, and the A/D rms values to ensure that an unacceptably low signal level was not being recorded due to too much attenuation. Usually, these values, once they were set just after the start time for recording, could remain fixed for the remainder of the pass. Additionally, the operator would note the position of the signal within the passband using a fast Fourier transform (FFT) display. At times, the predict file would not be accurate, and frequency offsets occasionally would have to be entered in order to center the signal within the passband.

C. Operational Problems

The RSS commanding process, while seemingly simple, was complicated by the need for simultaneous support of different projects/experiments with different requirements/considerations. For example, Galileo requests that the A/D rms values be maintained at 1-V peak to peak, whereas MGS requests that the A/D rms values be as high as possible without saturating the A/D converters. Narrow passbands

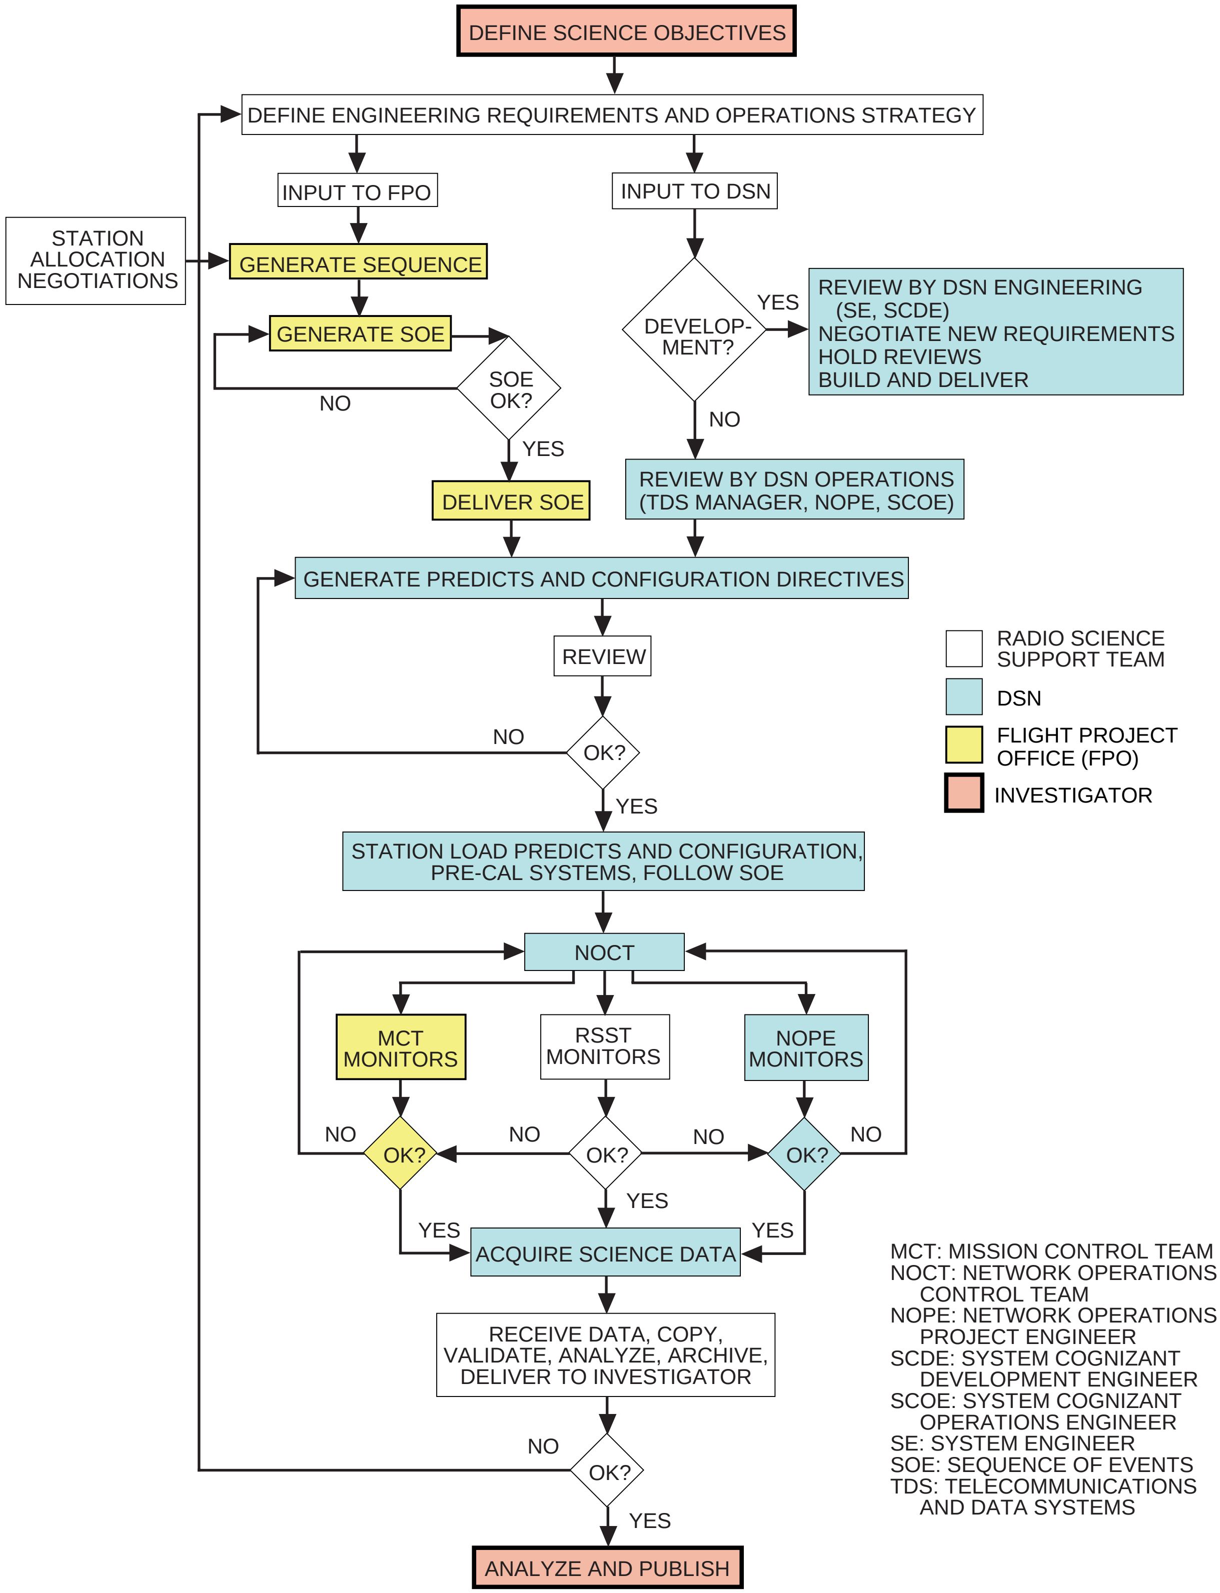

Fig. 2. End-to-end radio science experiment process prior to remote operations.

would be more difficult to contain a moving signal, yet they are required for certain experiments. Additionally, both attenuation and frequency adjustments are disruptive to the data and cause great difficulty for analysis; occultation data are extremely sensitive to frequency shifts and attenuation changes because it is difficult to separate effects from an atmosphere and effects from operator interaction.

Because of the to-be-expected unfamiliarity of the DSN operators with the science behind the experiment, and the lack of a method of communicating such information in an operations scenario, the operators would concentrate on meeting the project requirements. For example, operators would maintain a 1-V peak-to-peak A/D rms value via changing attenuations regardless of the type or timing of the experiment; although a gravitational wave experiment would not be greatly impacted, an occultation experiment could be decimated. The only two methods of communication available were (1) the SOE, which would not suffice because it was generated via a "keyword" lookup file and (2) the prepass briefing, which was not considered a viable option because it contained a lot of other important information and directives for the station personnel.

III. The Ideal Solution

The two main causes of "operator error" at the station, then, were oversubscribed station personnel and the lack of a clear method of communicating a scientific understanding of each radio science experiment. The solution was obvious: let those who already understood the science and did not have dozens of other systems to operate run the radio science subsystem.

A. Overview

This goal—to simplify the path between the users and the data—is shown clearly in the design of the remops operational scenario, as shown in Fig. 3. The level of assistance needed for RSSG personnel to effect a change in the configuration of the RSS went from three intermediary persons (the mission controller, the track controller, and the LMC operator) to none.

1. Hardware. Figure 4 shows the adaptation of the network and serial connections to include the radio science remote operations workstation (RSRW), the host that handles the transaction processing between a remote site (such as JPL) and the RSCP. The RSRW uses one of the open serial ports on the RSCP and connects to the network via the SPC LAN.

2. Software. The remops software is divided into two parts, a client run at a remote site (such as JPL) and a transaction server that runs on the RSRW at each SPC site.

The remops client: The client is a graphical user interface written in Tcl/Tk (a toolkit based on Xlib) that allows users to select actions from a specific set of allowable commands and parameters. Once the user has entered a command, the client delivers the command to the specified server (e.g., the remops server at SPC 10, SPC 40, SPC 60, or DSN Test Facility (DTF) 21). Responses from the server are delivered to a message window in the client, allowing the user to see error or completion messages as they happen.

The remops server: The remops server is more correctly called a transaction server because, when it receives action requests from the client, it actually reformats those requests into commands recognized by the RSCP and passes them through. The server side, however, is prominent because it handles the serial connection between the RSRW and the RSCP, accepts incoming requests from several clients, and outputs monitor data to all clients. It also allows the RSRW to act as a staging area for script files (for unattended operations) and RSSG-generated predict files. This staging area is necessary in case the network connection between an SPC site and a remote facility is broken after a script has been configured.

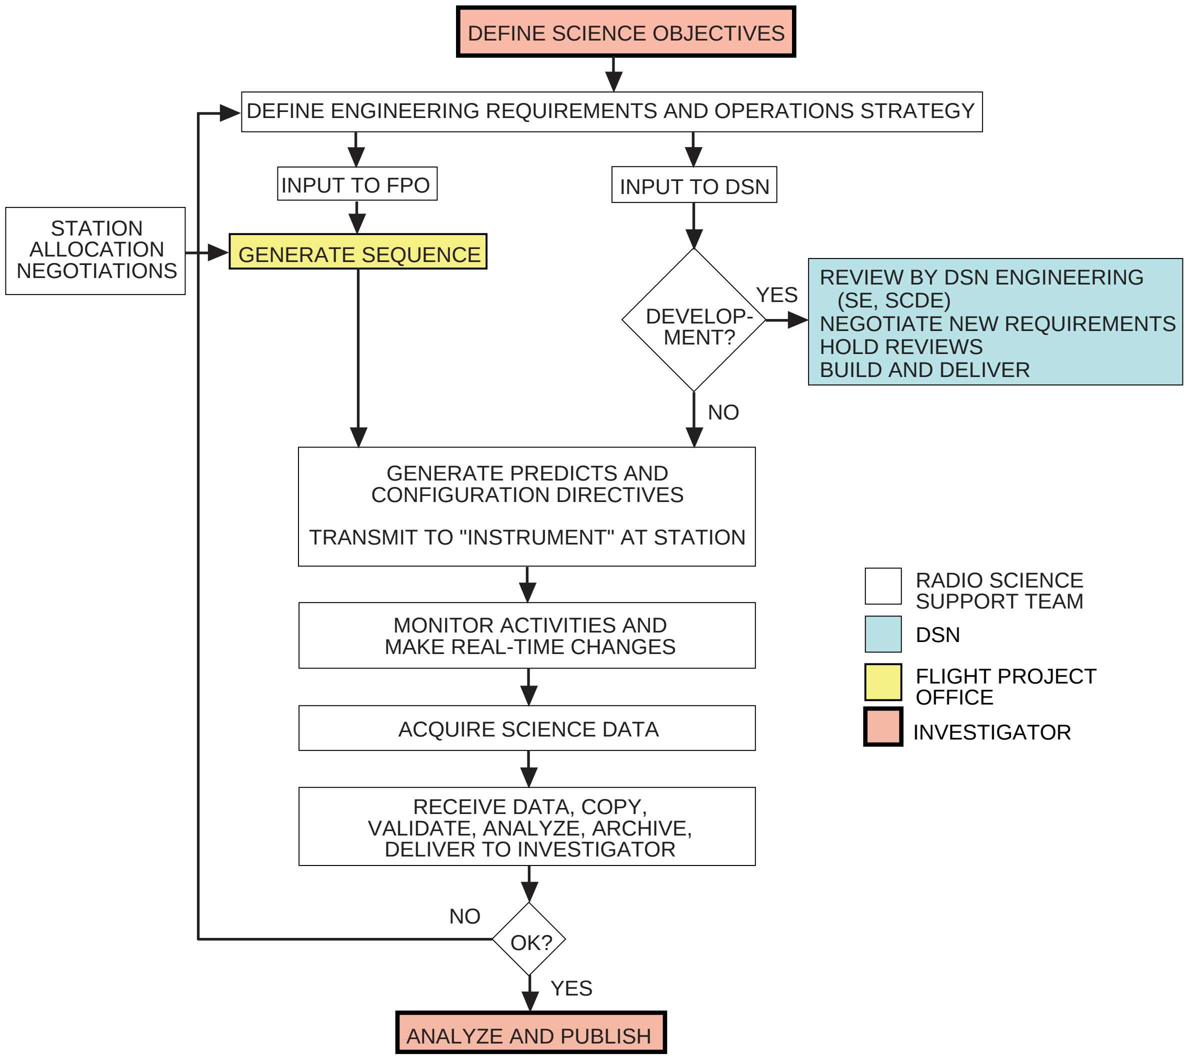

Fig. 3. End-to-end radio science experiment process after remote operations.

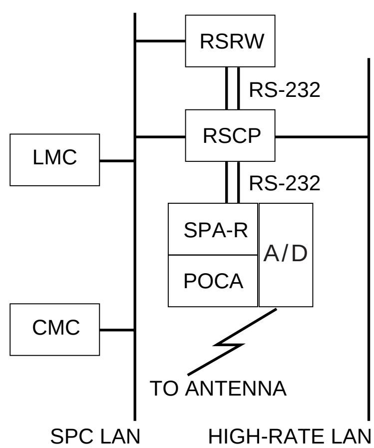

Fig. 4. SPC network and serial connections: remote operations configuration.

B. The Reality

While the client/transaction server base of the remops implementation was a valid solution, the difficulties involved in interfacing with a system that had no hooks for remote operations proved, at times, quite frustrating. Initially, the transaction server could not even log into the RSCP because the extra serial ports were not being initialized upon RSCP boot-up. The software on the RSS had more responses than documented (causing some synchronization problems), and the hardware itself did not always let the software do the same thing every time. Software limitations had been hard coded into the RSCP due to several assumptions about users entering commands at a local terminal, and the response of the RSCP to these limitations being reached had been neither tested nor documented. Moreover, the network configuration at the testing facility (DTF 21) was being changed to a 10baseT network, while the stations all had 10base2 networks. This change at DTF 21 allowed access to the high-rate LAN from the RSRW—something that was not allowed at the stations—but led remops development to assume that that connectivity was a valid option for using the internet protocol (IP)-activated RSCP port on the high-rate LAN.

Also, the ability to receive monitor data into the RSRW for automatic value-checking was shaky—the technical capability was there, but the RSS software would have required more time and money than was available to make such modifications. Because the monitor data path had been firmly coded into the RSCP and the SPA-R, changing that path would have been much more difficult than if the potential for remote operations had been considered in the initial implementation. Additionally, while querying the RSCP was workable, no data about important SPA-R parameters could be determined using this method; only variables in the RSCP shared memory were available.

C. Compromises

The final configuration of the remops solution was to use the Network Operations Control Center (NOCC) workstation displays for near-real-time information on the status and configuration of the RSCP and SPA-R. As of 1993, this workstation could capture broadcast monitor data from the DSN stations and display parameters grouped by DSN subsystem. While this compromise does not allow for the automatic handling of out-of-range parameter values, because the broadcast data cannot be sent back to the remops server, it does allow users to get a lot of important information quickly. This compromise does not mean that remops will not work at all without an NOCC workstation, but it does work more efficiently with one. The NOCC workstations normally are positioned near the workstations running the remops client.

Additionally, funding to upgrade the RSCP software had been obtained after discovering some problems in the 1993 RSCP upgrade task, which had added a larger internal disk and moved the output from 9-track tapes to 8-mm tapes. This new task eventually was augmented to include the addition of several commands that would allow RSSG-generated predicts to be loaded into the system; two modification kits were sent out to change the operating system (OS) configuration files in order to initialize the necessary serial port and enable IP on the SPC LAN.

IV. Results

Use of the remops system already has had a profound, positive impact on the ability of the RSSG to efficiently handle a high-level workload with a decreasing number of available support personnel, and this trend is expected to continue during the MGS mapping phase and the initial Cassini mission. The data in Table 1 support this claim since, when station support (i.e., an LMC operator) was not used, the RSSG support personnel configured and commanded the RSS, showing that fewer people could produce equivalent results with the remote operations resource.

The ability of the RSSG to remotely configure the RSS in parallel with the station's pre-cal period allowed the pre-cal time to be reduced—a budget savings. Additionally, special station-wide training

Table 1. Success rates versus personnel requirements.

Experiment	Success rate	Total experiments	RSSG personnel	Station support
MGS Aerobraking	97.6	166	1.5	No
GLL Occultations	97.8	47	4.0	No
GLL Solar Wind 1994	98.1	222	5.5	Yes
Gravity Wave 1994	97.1	138	4.6	Yes
GLL Solar Wind 1993	96.5	146	4.6	Yes

sessions, called operational readiness tests (ORTs), no longer were required beyond the basic checkout of the hardware. This was because the RSSG personnel could be trained by experienced users during radio science experiments, thereby not requiring any additional station time. When the cost of station time is taken into account, the cost savings due to the use of remops (which reduced the required station time by an average of 45 minutes per pass) is significant.

The flexibility and reliability of gathering useful radio science data using remops has allowed the RSSG to deal with the growing ratio of project experiments to project personnel levels. In addition, out-of-the-ordinary hardware configurations are no longer outside the realm of possibility, due to the higher level of specialized knowledge about each individual experiment that the RSSG operators can draw upon when conducting remote operations. The time required to set up radio science experiments has dropped, and the station time and cost due to radio science alone has dropped to a level that requires only hardware support.

This was achieved by using a new type of control on an old system, which was not designed for such control. The benefits, therefore, would be even greater if the station hardware and software had been designed with remote operations in mind. This use of the radio science remops software has effectively demonstrated that, in cases where similar problematic situations exist or in new system developments, the ability to remotely control a system can greatly improve performance. Even with the unforeseen problems encountered while interfacing with the legacy radio science subsystem, the benefits to the RSSG from the remote operations method were worth the effort.

Appendix
Acronyms

A/D = analog-to-digital (converter)
CMC = command, monitor, and control (station)
DSCC = Deep Space Communications Complex (a DSN antenna site)
DSN = Deep Space Network
DTF = DSN Test Facility
FFT = fast Fourier transform (of the incoming signal)
FPO = Flight Project Office
GLL = Galileo
IP = internet protocol
JPL = Jet Propulsion Laboratory
LAN = local area network
LMC = link monitor and control (station)
MCT = Mission Control Team
MGS = Mars Global Surveyor
MO = Mars Observer
NOCC = Network Operations Control Center
NOCT = Network Operations Control Team
NOPE = Network Operations Project Engineer
ORT = operational readiness test
OS = operating system
POCA = programmable oscillator control assembly
rms = root mean square (of the incoming signal)
RSCP = radio science control processor
RSRW = radio science remote operations workstation
RSS = radio science subsystem
RSSG = Radio Science Systems Group
RSST = Radio Science Support Team
SCDE = System Cognizant Development Engineer
SCOE = System Cognizant Operations Engineer
SE = System Engineer

SOE = sequence of events

SPA-R = signal processing assembly, radio science

SPC = Signal Processing Center (the open-loop receiver site)

TDS = telecommunications and data systems