

Uplink Array Demonstration with Ground-Based Calibration

Larry D'Addario,* Robert Proctor,* Joseph Trinh,* Elliott Sigman,*
and Clifford Yamamoto†

A set of small, separately steerable reflector antennas has been used as a transmitting phased array for the purpose of demonstrating techniques that can be used in a larger array to serve the future uplink transmission needs of NASA's Deep Space Network. We envision an operational array with 100 or more antennas that could generate the order of 1 TW of effective isotropic radiated power (EIRP). The demonstration is a small-scale version of this with five antennas and about 1 MW of EIRP. Each antenna has a 1.2-m-diameter aperture and a 2-W power amplifier, and the array operates in the 14.0- to 14.5-GHz communication satellite band. The main technical challenge for an uplink array is to ensure that the carrier phases of the signals from all antenna elements are aligned when the signals arrive at the receiver on a distant spacecraft. This requires a method of phase calibration. In the demonstration, we have shown that active receivers attached to the Earth near the array can be used as calibration targets. Measurements made at these receivers have been successfully used to calculate the phase adjustment needed at each antenna to achieve the desired alignment, even though the destination spacecraft is in a direction and at a distance very different from that of the calibrator. When the calculated adjustments are applied at the antennas, the combined power at a spacecraft has been shown to be within 1 dB of that expected for perfect alignment. Commercial satellites in geostationary Earth orbit were used for these tests. Other objectives of the demonstration, all successfully accomplished, include: (1) Show that a new and simple electronics architecture, specifically designed for phase and delay stability, can implement all functions of NASA deep space uplinks at low cost, supporting mass production for large arrays. (2) Show that phase alignment can be maintained for at least a few hours without recalibration. In fact, stability over several days has been demonstrated. (3) Show that data can be transmitted on the aligned carriers at substantial speeds with no degradation in bit error rate compared with single-antenna transmission at the same EIRP.

Note — Because of the length of this article, a list of contents with page numbers is provided to aid the reader.

I. Introduction	PAGE 2
II. Project Overview and Objectives	PAGE 4
III. System Description	PAGE 5
IV. Test Methods and Results	PAGE 16

* Tracking Systems and Applications.

† Communications Ground Systems 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. © 2009 California Institute of Technology. Government sponsorship acknowledged.

V. Subsystem Details	PAGE 44
VI. The Next Step: A Prototype Array for the DSN	PAGE 63
Appendix A: Documentation Repository	PAGE 67
Appendix B: Analysis of Costs	PAGE 69

I. Introduction

Transmission of commands and other data to distant spacecraft has traditionally relied on having a large antenna on Earth driven by a high-power amplifier, as well as a high-gain antenna (about 30 to 50 dBi) on the spacecraft. At interplanetary distances, NASA's Deep Space Network (DSN) now uses individual 34-m or 70-m-diameter ground antennas with 20-kW amplifiers operating near 2.1 GHz or 7.2 GHz. These systems are limited to data rates of a few kbps at the outer planets. To support more complex missions requiring higher data rates, to support a larger number of simultaneously flying missions, and to reduce the spacecraft cost by allowing it to carry a lower-gain antenna, it is desirable to have higher effective isotropic radiated power (EIRP) capacity on the ground. To do this, or even to build new systems with the same capacity as existing ones, we are exploring the use of phased arrays of small antennas rather than individual large antennas. However, as of this writing, no decision has been made to construct any facility based on the concepts or results that are presented here.

An accurately aligned array with identical elements produces a far-field EIRP of

$$E = N^2GP$$

where N is the number of antennas, G is the gain of each, and P is the power delivered to each. The EIRP can then be increased as much as desired by adding more elements of the same design. With a single antenna and amplifier, there are practical limits to the largest possible sizes [1] and these are already being approached by the existing systems. In addition, for a given EIRP, there is an antenna size G and amplifier size P [and thus a number of elements $N = (E/GP)^{1/2}$] that minimizes the cost of construction. We have reason to believe¹ that the minimum occurs for antennas and amplifiers much smaller than the 34-m and 20-kW sizes used in the present DSN.

To use a set of separately steerable antennas as a transmitting phased array, it is necessary that they transmit signals whose carrier phases are aligned when the signals arrive at their destination. For our applications, the destination is a distant spacecraft. Assuming that the signals are modulated to carry useful information, the maximum overall data rate (for a given error rate) is achieved if all antennas simultaneously transmit the same information, with the modulation timings adjusted so that they are also aligned when the signals arrive at the spacecraft. In general, this means that the carrier phases and modulation timings are not equal at the antennas, since some antennas are closer to the spacecraft than others. To maintain alignment as the spacecraft moves across the sky, each antenna must include

¹ L. D'Addario, "Cost Model for an Uplink Array," DSN Array Technical Note No. 37 (internal document), Jet Propulsion Laboratory, Pasadena, California, April 29, 2005. A copy may be requested from ldaddario@jpl.nasa.gov.

a mechanism for continual adjustment of the phase and timing with good accuracy. For most applications, it is sufficient that the relative phases and timings be aligned, so we can arbitrarily designate one antenna as the “reference” and measure the phases and timings of the others relative to it.²

At microwave frequencies, it is not practical to establish the relative carrier phases at the antennas by design and construction. This is because the practical tolerances on the lengths of cables and on the transfer functions of electronic components such as amplifiers and filters are greater than one wavelength or cycle. Therefore, some method of in situ post-construction calibration of carrier phases is needed. Regarding the modulation timing, it is sometimes possible (but difficult) for it to be fixed by design and construction to an accuracy $\ll 1/B$, where B is the modulation bandwidth, because usually $B \ll f_0$ for carrier frequency f_0 , but it remains desirable to have an in situ method of calibrating the modulation timing as well. Devising a practical and accurate calibration procedure is the principal challenge in constructing an array of this type.

In the absence of disturbances, the calibration need only be carried out once. But in practice, changes in temperature, mechanical stress, and component replacements cause the relative phases and timings to vary over time, necessitating recalibration. Since each calibration procedure takes time away from useful transmission to a spacecraft, the system should be as stable as possible so as to maximize the interval between calibrations.

The concept we have chosen to pursue is illustrated in Figure 1. The calibration process consists of all antennas transmitting simultaneously to a receiver at a fixed and known location on Earth, near the array. The receiver must have a small antenna so that its beam encompasses the entire array. The signal from each antenna is phase-modulated with a binary pseudorandom sequence (PRS). Unlike the situation during transmission to the target spacecraft, the modulation for each antenna is different, but each sequence is known to the receiver except for a possible timing offset. This enables the receiver to extract each antenna’s component from the composite signal and to measure its carrier phase and modulation timing. Since the measurements are relative to a selected reference antenna, the receiver’s internal timing need not be precisely synchronized to that of the array. Assuming that the positions of the antennas and of the receiver are known, the propagation delay from each antenna to the receiver can be calculated, and this can be used to calculate the carrier phase and modulation timing of the signal at each antenna (relative to the reference antenna). Optionally, this process can be repeated at additional receivers around the array, providing redundancy as well as the possibility of correcting for certain errors. Next, the phase and timing adjustments needed to produce alignment in the direction of the target spacecraft are calculated and applied, and the PRSs are replaced with bit sequences (identical at all antennas except for the timing adjustments) containing the information to be transmitted, including any necessary error detection or correction coding.

² For some navigation applications, the modulation timing at the reference antenna with respect to a clock elsewhere on the ground is also important. In addition, the carrier phase noise and phase drift over time at the reference antenna are sometimes important, even if they are common among all antennas. Such requirements can be met using the same techniques as are used with single-antenna transmission, so they will not be further discussed in this article.

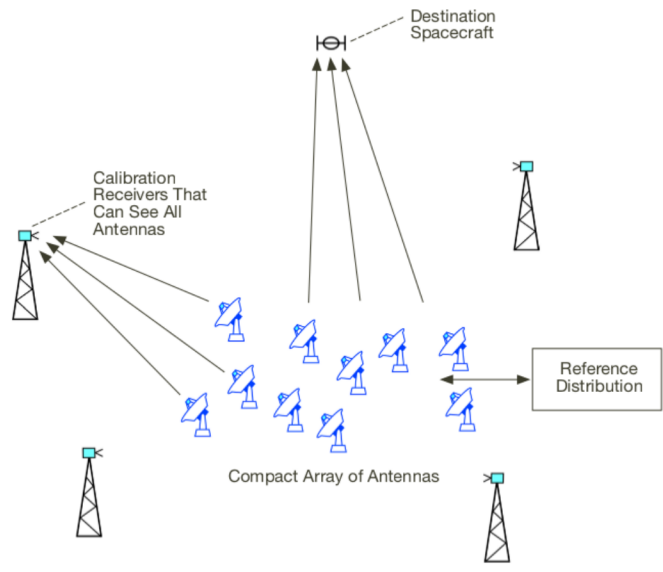


Figure 1. System concept for a future large array of which the demonstration is a scaled-down version. The antennas receive phase and timing reference signals from a central master source. Calibration is accomplished by transmitting with all antennas simultaneously to one or more calibration receivers attached to the Earth and, hence, near the array. Several calibration receivers provide redundancy and the ability to solve for additional parameters. All calibration data are then used to compute the phase and timing adjustments to each antenna's signal that will cause them to be aligned at a destination spacecraft. The adjustments are then applied and the antennas are repointed toward the spacecraft.

The above description is simplified and ignores several complications and important sources of error that will be discussed in later sections of this article.

Alternative calibration concepts include the use of natural radio sources and the addition of a receiving capability to each antenna [2]; placing the calibration receivers on aircraft or spacecraft; placing a transponder on an aircraft or spacecraft and a calibration receiver on the ground to process the transponded signal; and use of the signal reflected from a natural or man-made object (radar). In particular, the radar approach with the Moon as the target has been pursued by others at JPL [3]. Most of these concepts are far more expensive to implement than the selected approach, and each of them has serious technical difficulties or disadvantages. The selected approach has potential difficulties as well, but we show in this demonstration that they have been largely overcome.

II. Project Overview and Objectives

The demonstration project involves the design and construction of a small-scale transmitting array of the type described in Section I. Its size is limited by funding, but it is sufficient

to establish the feasibility of the calibration approach and to incorporate most of the features required in an operational array for the DSN. It allows us to meet these objectives:

- Develop an architecture for the array's electronics that provides high stability and is inexpensive to mass-produce.
- Construct and test a small array using this architecture. (Some of the devices developed can serve as prototypes for a large array.)
- Determine the practical short- and long-term phase and delay stability and thus establish the maximum practical calibration interval.
- Demonstrate a calibration technique that achieves accuracy within a predetermined error budget.

The demonstration array operates at 14.0 to 14.5 GHz and includes five antennas, each with a 1.2-m-diameter aperture and a 2-W power amplifier. This allows it to generate up to 1.0 MW of EIRP when all signals are aligned at a target. The calibration system includes four specially designed receivers mounted on towers around the array, spaced at roughly equal azimuths but at varying distances and elevations so as to explore near-field and multipath effects. The available EIRP is not sufficient to allow testing with deep-space spacecraft, but it is easily adequate for achieving high signal-to-noise ratio (SNR) at geostationary Earth orbit (GEO). Transponder channels on commercial GEO communication satellites can easily be rented over short time intervals. Therefore, such satellites were chosen as the targets for our tests, and this determined the frequency band in which the array must operate.

The antenna size and amplifier size were chosen because inexpensive units were available off-the-shelf. The number of antennas is near the minimum needed to represent a large array. Three is the nontrivial absolute minimum, and four are needed to provide enough degrees of freedom in the calibration measurements to investigate the effects of receiver position errors. With five antennas, we add redundancy and allow useful measurements to be made even if one fails. A larger number of antennas would have been difficult to support within the available budget. Finding a suitable place near JPL to install the array proved difficult, and this would have been exacerbated if we had attempted to build a larger array.

III. System Description

Figure 2 is an overall block diagram of the system. The same architecture is intended to be used if a large array is constructed in the future. It differs substantially from traditional systems in that the signals to be transmitted are synthesized at the antennas, with only reference tones and data being distributed from a central source. Each major subsystem is further described in the following paragraphs, and additional details are given in Section V.

The inexpensive reflector antennas were manufactured for fixed pointing. We added a two-axis positioner and servos so that they could be pointed anywhere in the sky. Commercial products were used and modified as necessary.

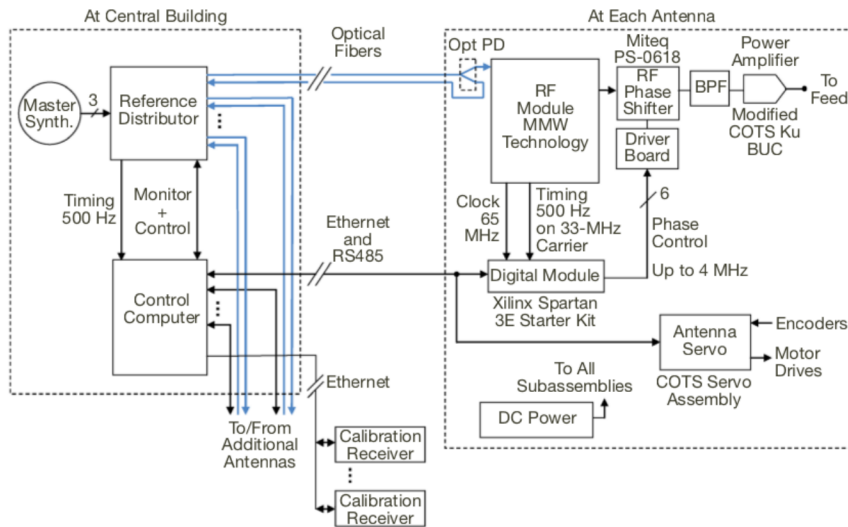


Figure 2. Block diagram of the electronics of the demonstration array, showing the major subsystems and inter-connections. The RF components and the digital module are on the feed arm of each antenna. The RF module, RF phase shifter, and optical power splitter are in a temperature-controlled box. Not shown is the fact that one antenna (#2) includes a second RF module without phase shifter to produce a frequency-offset reference tone during calibration. See Section III for a functional description of each item shown here, and Section V for more detailed discussions of their designs.

Most of the electronics were developed by our team at JPL. This includes a central reference distribution system, digital and RF electronics at each antenna, a temperature-stabilized package for the RF electronics, power supplies, and the calibration receivers. The system also includes software for control of all subsystems. The control programs allow automation of many tests via scripting and remote control over the Internet.

A. Antenna-Mounted Electronics

Central to the architecture is a digitally controlled phase shifter in the RF path at each antenna. This device, operating at the final transmitting frequency, accomplishes several functions simultaneously. It (a) subtracts the antenna-specific carrier phase offset that has been determined by the calibration process; (b) adds the phase needed to cancel the geometrical delay on the path to the target spacecraft, so as to align all carrier phases at the target; (c) shifts the carrier frequency to provide Doppler compensation at the spacecraft; and (d) applies any desired phase modulation. Items (b) and (c) involve continual adjustment as the spacecraft moves in the sky. All of this is done by providing the correct sequence of control numbers to the phase shifter at the correct times. In the demonstration array, a commercial 6-bit microwave phase shifter is used.³ This provides 5.625-deg resolution and

³ Miteq model PS-0618-360-5-5.6.

a switching speed of about 100 ns, allowing modulation rates up to a few MHz. In a large operational array, a custom design could easily provide finer resolution and higher speed.

The 6b control number sequence for the phase shifter is generated by the digital module, which consists of a circuit board containing a single field-programmable gate array (FPGA) and supporting devices. The module contains logic to vary the phase linearly at a commanded rate for tracking, and to produce binary phase-shift keying (BPSK) modulation at rates from 2 kbps to 16.384 Mbps (by factors of 2) with a modulation index from ± 2.8 deg to ± 90 deg. The bitstream for modulation can be either an internally generated PRS of length 1023 bits or taken from a double-buffered 8192-byte memory that has been filled by the array control computer. If the data buffer is refilled in a timely manner, continuous transmission of arbitrary data of any length can be supported.⁴ The PRS is used during calibration, and is known to the calibration receivers. The time at which the first bit of the buffer or the first bit of the PRS begins being transmitted can be adjusted with a resolution of 7.63 ns in order to align the bit transition times at the target. The adjustment is with respect to a 500-Hz timing reference signal that is distributed to all antennas from the central building. In addition, the starting time of the PRS can be offset by 0 to 1022 modulation clock periods. The digital module communicates with the control computer over Ethernet using a private network and Transmission Control Protocol/User Datagram Protocol (TCP/UDP). Communication is supported by software running in a microcomputer built into the FPGA. Commands can be addressed to each antenna's digital module separately, or they can be broadcast to all modules simultaneously. (In an operational system, the dominant traffic on this network would be for filling the data buffers with the information to be transmitted. Since all antennas transmit the same information, these data would be broadcast. Thus, the network data rate is independent of the number of antennas. Using 100-Mbps Ethernet, a sustained uplink rate of at least 40 Mbps could be supported for any number of antennas.)

Timing within the digital modules is controlled by two signals that are distributed from the central reference distributor to all antennas. These are at frequencies 32.768 MHz and 500 Hz, coherent with each other. The first is multiplied by 2 in the RF module and again by 2 in the FPGA to provide a 131.072-MHz clock that drives most of the FPGA logic. The second is used to select a specific clock cycle (out of each $131.072 \text{ MHz}/500 \text{ Hz} = 262,144$ cycles) to act as the timing reference for commands whose execution must be synchronized among antennas. This avoids the need to create and maintain any "real-time clock" at the antennas. The 2-ms timing ambiguity created by the 500-Hz reference is far longer than the propagation time across the array (53 ns across the 16-m demonstration array, and $3.3 \mu\text{s}$ across a 1-km array), and is long enough to allow a real-time control computer to ensure that a command is received within a specific 2-ms interval. Timing synchronization among the antennas is not perfect because the reference distribution cables are not equal in length and because of small differences in delays within the electronics. But the practical differences are small compared with the 2-ms ambiguity interval (less than 80 ns in the demonstration array and easily less than a few μs in a large array), and the actual differences are accurately measured during calibration.

⁴ The transmission of arbitrary data via the buffer was built into the design and was checked in the lab during stand-alone tests of the digital module, but the feature was never tested in the full system and was not used during operation of the demonstration array.

The RF module⁵ receives the signal sent from the reference distributor on optical fiber. The signal is intensity modulated onto a 1310-nm wavelength optical carrier and is demodulated by a photodetector. The resulting electrical signal is a composite waveform consisting of a carrier reference tone at 437.5 to 453.125 MHz added to a 32.768-MHz sinusoid that is biphasic modulated at 500 Hz. These are separated in an analog diplexing filter. The timing signal is sent to the digital module and the carrier reference is used to phase lock a voltage-controlled oscillator (VCO) at 7.0 to 7.25 GHz (16 times the reference). The VCO output is doubled in frequency and amplified to produce the 14.0- to 14.5-GHz carrier at the input of the phase shifter. (The 7.0- to 7.25-GHz signal is also available at an output not shown in Figure 2, so the same module could be used in a system designed to operate in the 7.145- to 7.235-GHz space research allocation used by many NASA spacecraft.)

After the phase shifter, the modulated signal drives a power amplifier that has about 38 dB of gain and an output power of about 2 W at -1 dB compression and about 3 W saturated. The amplifier was built by modifying a commercial block upconverter (BUC) intended for use in small satellite communication terminals.

A photograph of one antenna, showing the modules just described, is given in Figure 3.




Figure 3. A typical antenna assembly before its installation at the field site. The antenna is a 1.2-m-diameter offset-fed paraboloid. The feed, supplied with the reflector, is a corrugated horn with integrated orthomode waveguide junction for linear polarization and a bandpass filter in one branch. The main electronic modules are mounted on the feed arm. The power amplifier drives one polarization directly, and the other (filtered) polarization is connected to a commercial low-noise block converter (LNB) for receiving the transponded satellite signal. The LNB is installed on only one antenna.

⁵ This assembly was built under contract by MMW Technology Inc. using a conceptual design provided by JPL.

B. Reference Distribution

The reference distribution system consists of the master synthesizer and reference distributor assemblies (Figure 2). The master synthesizer includes an oven-stabilized crystal oscillator at 100 MHz as the frequency standard of the array, from which it synthesizes three tunable reference signals in the range 359 to 531 MHz. These frequencies are set by the control computer as functions of the carrier frequency to be transmitted, as explained below.

The tunable references are split by passive power dividers and distributed to five identical modules inside the reference distributor assembly, one for each antenna. In each of these reference generator modules (RGMs, see Figure 4), the carrier reference at $f_c/32$, where f_c is the carrier frequency, is generated and intensity-modulated onto an infrared laser coupled to a single-mode optical fiber. The output fiber from each RGM runs to an RF box on an antenna. At the RF box, the light is split in a passive optical power divider, with half of it going to the RF module and half being returned to the same RGM over another fiber in the same cable. The round-trip transmission enables the RGM to compensate for variations in the electrical length of the fiber, stabilizing the phase of the carrier reference signal at the antenna. The stabilization scheme is a crucial element of the system architecture.

A simplified block diagram of the length stabilization system is given in Figure 4. The carrier reference at $f_0 = f_c/32$ is generated by mixing an input at $f_0 + f_1$ (from the master synthesizer) with a voltage-controlled crystal oscillator (VCXO) at nominal frequency $f_1 = 77.76$ MHz. It is then modulated onto a diode laser coupled to the outgoing fiber. The signal returned from the antenna is converted back to electrical form and downconverted to $f_0 - f_1$ by mixing with the same VCXO, and the result is compared in phase against a second input from the master synthesizer at $f_0 - f_1$. The phase difference is used to drive the frequency of the VCXO via filter $H(f)$, forming a phase-locked loop. If we consider the input signals to have phase zero, it is easy to show that the loop will drive the phase θ of the VCXO so that $2\theta = (2L) \bmod 2\pi$, where L is the electrical length of the line in each direction, causing the phase of the signal at the turn-around point to remain constant as L

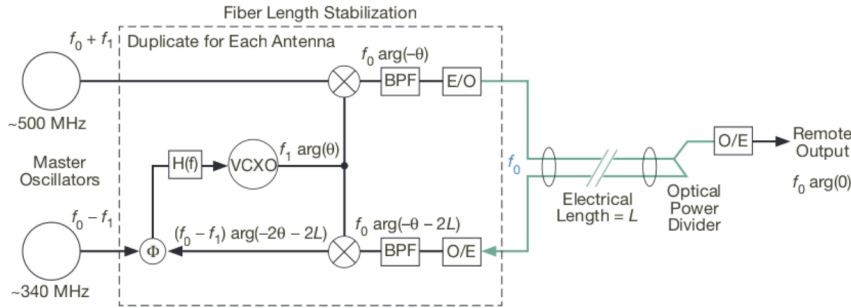


Figure 4. Simplified block diagram of the cable length stabilization, showing its principle of operation. Components in the dashed box are in the RGM. Here E/O and O/E stand for electrical-to-optical (laser) and optical-to-electrical (photodetector) conversion, respectively. In the demonstration array, the VCXO is at frequency $f_1 = 77.76$ MHz, and the output f_0 is in the range 437.5 to 453.125 MHz.

is changed. The actual phase there has a two-fold ambiguity; it could be either 0 or π . (This result, and the expressions in Figure 4, ignore the delays in the mixers, bandpass filters, and photonic components. As long as these are stable, the loop will correctly compensate for changes in L .) A system with a similar principle of operation was originally described in [4].

A separate module in the reference distributor produces the 32.768-MHz timing signal with 500-Hz coherent biphase modulation for the digital modules. This signal is passively split and sent to each of the RGMs, where it is added to the carrier reference as part of the signal used to modulate the laser (not shown in Figure 4).

The reference distributor also contains a sixth RGM. It is similar to the others but produces an output at $f_0 - 125$ kHz, and its fiber pair is connected to a second RF module on antenna #2 (only). After multiplication by 32, that RF module produces an unmodulated output at $f_c - 4$ MHz, which is transmitted by antenna #2 during calibration (along with its normal, modulated signal). This is further discussed in Section III.D. This RGM uses a VCXO at 77.885 MHz.

The master synthesizer must produce three signals for distribution, at $f_0 - f_1$ and $f_0 + f_1$ for the five main RGMs, and at $f_0 - f_1 - 250$ kHz for the offset RGM. All of these must be retuned to change the transmitted carrier frequency $f_c = 32 f_0$. This is accomplished with several direct digital synthesizers.

The decision to distribute the carrier reference at $f_c/32$, rather than at a higher frequency, was dictated by cost. Our objective was to demonstrate that a large array can be built with low-cost electronics. For an array of many antennas, the construction cost of the reference distribution system is dominated by that of the photonic components (lasers and photodetectors). At the time of our design choices, the prices of these devices increased sharply above about 800 MHz. While 10-GHz laser transmitters and receivers are readily available, they typically cost many thousands of dollars, compared with the US\$262 cost of the laser assemblies used here.⁶ The complete cost of parts to construct an RGM was US\$586, and with assembly labor the reproduction cost was under US\$1000.

A more detailed description of the reference distribution system is given in Section V.

C. Antenna Mechanical Subsystem

The antennas use 1.2-m offset paraboloids fed at their prime focus, as shown in Figure 3. The reflectors, feed support structures, and waveguide feed horns were purchased off-the-shelf as assemblies⁷ for US\$205 each. They are manufactured for use in small satellite terminals. As supplied, the antenna is intended for mounting in a fixed position to point at a single GEO satellite, but our application requires that it be pointed in a wide range of directions under computer control. We therefore added an elevation-over-azimuth motor-driven positioner using a commercial product⁸ modified by the manufacturer to our specifications. These de-

⁶ The costs quoted in this article are the actual amounts paid to the vendor at the time of purchase. They do not include shipping, cost of procurement borne by JPL, acceptance testing, or other overhead.

⁷ Prodelin (division of General Dynamics) Model 1132. http://www.gdsatcom.com/Antennas/Data_Sheets/1132.pdf

⁸ Quickset International Inc. Model QPT-130. <http://www.quickset.com/filebin/PDF/QPT130.pdf>

vices include a DC motor, gearing, and an incremental encoder for each axis. We connected each motor-encoder pair to a commercial motor driver and position servo.

Figure 5 shows photographs of the rear of the antenna. The positioner is supported on a mounting plate welded to a steel pipe, and the bottom of the pipe is welded to another plate that is bolted to a large concrete slab. (The concrete slab was preexisting at the installation site.) The support structure is extremely stiff, and is predicted to have negligible deflection in 100 km/h wind. The antenna is roughly balanced about the elevation axis by counterweights in order to minimize the torque required from the positioner. The left-side counterweight is a box containing DC power supplies for the electronics, and the right side consists of steel slabs. Each motor driver/servo is a small electronic box that we mounted inside the rear cover of the positioner. Each axis is controlled independently. All 10 servos of the array communicate with the control computer over a single RS485 serial network.

An important consideration in the mechanical design is that the two rotation axes and the antenna's boresight (beam) axis should all intersect. Intersection of the rotation axes provides an effective position for the antenna that is independent of pointing direction and




Figure 5. (a) A rear view of an antenna, showing the counterweights, positioner, support post, and base electrical box. The left side counterweight carries DC power supplies for the electronics. The base box contains a 24-VDC unregulated supply for the motors. (b) A close-up of the positioner, showing the aluminum brackets that connect it to the reflector and to the counterweights. (c) The interior of the positioner, with the two motor driver/servo boxes mounted to the inside of the cover at the bottom.

is fixed to the Earth. These axes intersect in the positioner with an accuracy determined by the machining tolerances of its 6-mm-thick steel case and the runout of its bearings. Although the intersection mismatch is not part of the positioner's specifications, examination of the design suggests that the error is not more than about 0.1 mm. If the boresight axis does not intersect the rotation axes, then small pointing errors can lead to significant phase changes (or changes in the antenna's effective location). Keeping the phase error less than $1/50$ cycle with a pointing error of $1/2$ beamwidth implies an intersection miss distance less than $D/25$, where D is the aperture diameter (48 mm for our $D = 1.2$ m). We achieved this by carefully selecting the location of the reflector with respect to the positioner in the design of the aluminum mounting brackets that form the interface between them.

D. Calibration Receivers

A conceptual block diagram of a calibration receiver is shown in Figure 6. The receiver's antenna must be small so as to produce a broad beam that can see all antennas of the array. In the demonstration, the receivers are 20 m to 900 m from the array antennas. The antenna for the most distant receiver is a 20-dB gain horn, and for the others it is an open-end WR75 waveguide (theoretical gain 5 dB). The signal level at the antenna's waveguide port is 10 to 200 μ W per array antenna, which is strong enough that no low-noise amplifier is needed. In fact, except at the most distant receiver, attenuators of 9 to 12 dB are used just after the antenna to avoid saturation.

During calibration, all array antennas transmit to the receiver simultaneously at the same carrier frequency but modulated with different pseudorandom sequences. In addition, one antenna transmits an unmodulated tone offset in frequency by 4.0 MHz from the modulated carriers. This allows the receiver to be very simple: It uses the offset tone to downconvert the set of modulated signals to a 4-MHz intermediate frequency (IF), and the IF signal is immediately digitized. The down-conversion is accomplished in a square-law detector, which is the only microwave component other than the antenna and attenuator.

Although we have said that each antenna uses a different PRS, we actually use the same sequence at every antenna but each begins at a different bit in the sequence. If the timing differences at the antennas are large compared with the largest difference in propagation times to a receiver, then the sequences are nearly uncorrelated at the receiver. Specifically, we have chosen to use a length-1023 PRS of the class generated by a linear feedback shift register (an LFSR sequence) and to modulate at a clock rate of 1024 kHz. Such a sequence has an autocorrelation function of 1023 at zero offset and -1 at all other integer-clock offsets. One clock period is just under 1.0 μ s, which corresponds to 300 m in free space, whereas the demonstration array is only 16 m in extent. Even for a large array, this sequence would remain nearly uncorrelated at a receiver if the smallest time shift among antennas is at least a few clock periods. In the demonstration, we have used a minimum time shift of 21 clocks.

The digitized IF is distributed to five identical signal processors, one for each antenna, which operate in parallel. These are implemented in a small FPGA. Processors for several more antennas could have been included, and we see no obstacle to implementing hun-

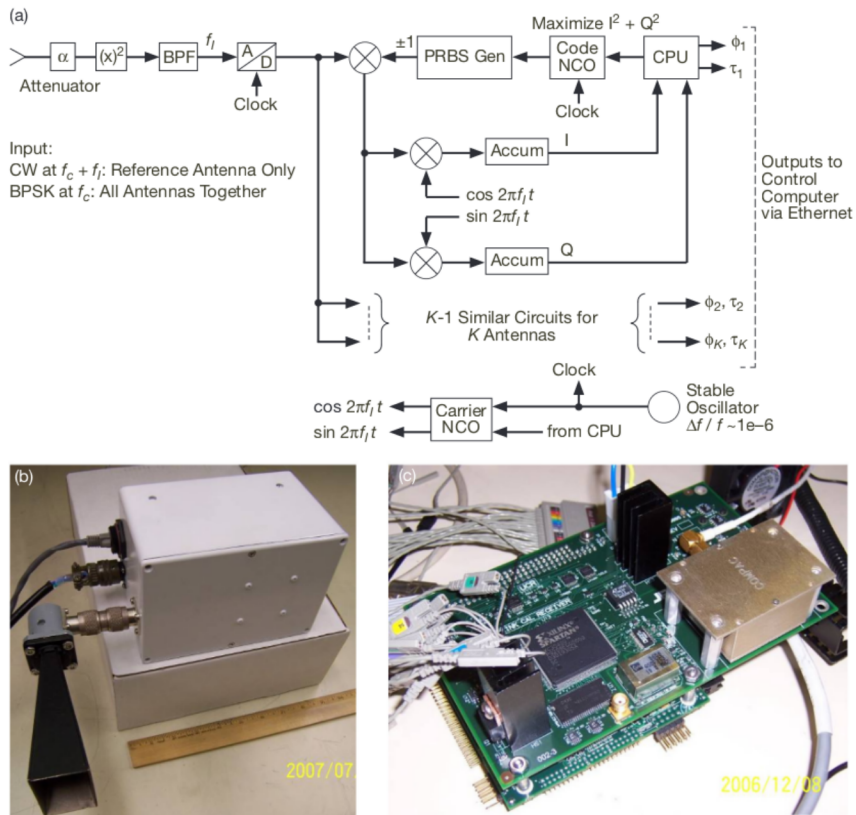


Figure 6. (a) Simplified conceptual block diagram of a calibration receiver. The only microwave components are an attenuator and a square-law detector. **(b)** An assembled receiver in the lab. A horn antenna is attached at one end, and the only other connections are for AC power and Ethernet communication. **(c)** The main interior components during bench testing. The small aluminum box contains an input IF amplifier and filter (0.1 to 5 MHz), and the largest IC on the main circuit board is the FPGA. Attached below the main board is a commercial PC104 single-board computer. Signal processing is shared between the FPGA and this computer.

dreds of processors in a larger FPGA. As shown in the block diagram, each processor multiplies the signal by the known PRS. If the local and received PRSs are aligned in time, then this multiplication removes the modulation and leaves just the carrier. The resulting signal is multiplied by in-phase and quadrature versions of a local estimate of the carrier, and the results are accumulated for one cycle of the PRS. If the PRSs are aligned and if the estimated carrier is sufficiently close to the correct frequency, then the accumulator contents represent the real and imaginary parts of the complex amplitude of that antenna's carrier. Meanwhile, the signals from the other antennas, which are also present at this processor's input but are not aligned with the local PRS, appear as noise and contribute at most $(N-1)/L$ as much as the aligned antenna's signal to the accumulators, where $N = 5$ is the number of antennas

and $L = 1023$ is the PRS length, assuming that the signals from all antennas have the same amplitude at the receiver. (If N were much larger, it would be necessary to make L proportionally longer to maintain high signal-to-noise ratio, but this could easily be done.)

The receiver includes an embedded single-board computer running a real-time operating system. This computer controls the signal processing, records the resulting measurements, and communicates with the array control computer. The timing of the local PRS sequence of each processor can be separately adjusted by the computer via the numerically controlled oscillator (NCO) that generates the local PRS clock [“Code NCO” in Figure 6(a)]. In this way, software in the computer enables the FPGA-based processors to acquire and maintain timing alignment with the signals from all antennas. Details of how this is done are given in Section V.

The receiver includes an accurate and stable crystal oscillator as its frequency reference, and the master oscillator of the array is also an accurate and stable crystal oscillator, but the two are independent. (We decided to avoid the complexity and expense of distributing the array’s master reference to all receivers.) Each has short-term fractional frequency variation near 10^{-10} but the absolute fractional frequency errors may be as much as 10^{-5} . This is close enough so that over one cycle of the PRS ($1023/1024$ kHz = 0.99902 ms) the difference between the local and received carrier phases at the 4-MHz IF changes by less than 0.04 cycle. The receiver makes a new measurement every PRS cycle, so meaningful measurements can be made. Nevertheless, during signal acquisition the receiver measures the apparent carrier frequency and adjusts its carrier NCO to match. This enables the local carrier estimate to be within 0.0023 Hz (resolution of the NCO) of the received signal initially, and we find that in actual operation it remains within about 0.05 Hz even after several hours of signal tracking.

After each PRS cycle, the receiver can record the I and Q accumulator contents, providing a measurement of the carrier phase $\phi = \arctan(Q/I)$. It can also record the time offset τ of the received PRS relative to the nominal starting time of the receiver’s PRS, providing a measurement of the modulation timing offsets of the antennas. Again due to the difference in oscillator frequencies, τ will slowly change. Both ϕ and τ are measured for all antennas, and in later analysis the values for the selected reference antenna are subtracted from the others, cancelling the effect of the residual difference in frequencies. For τ , the large timing differences among the antennas that were deliberately inserted in order to make their PRSs distinguishable at the receiver are also subtracted, leaving the antenna-specific timing differences whose calibration is needed.

E. Control Computer and Software

The array control computer is a generic rack-mounted PC running Microsoft’s Windows XP. It includes I/O cards for general-purpose analog/digital data acquisition, general purpose interface bus (GPIB) instrument control, and a second Ethernet port (in addition to one on the motherboard). The data acquisition card is used to monitor and control the reference

distributor, and the GPIB card provides an interface to test instruments such as a spectrum analyzer. One Ethernet port connects to a local TCP/IP network that includes the calibration receivers and a router to the public Internet, and the other Ethernet port connects to a private network that includes the digital modules at the array antennas. A photograph of the rack containing this computer and the reference distributor is shown in Figure 7.

Software with the following features has been built to run in the control computer:

- A single program integrates control and monitoring of the antenna motion servos, the antenna electronics, the master synthesizer and reference distributor, all calibration receivers, and any GPIB device, even though communication with each of these is via a different interface.
- A graphical user interface (GUI) is provided with buttons and data entry boxes for frequently used commands, along with numerical and colored-light status displays. Separate sections handle the antenna motion, antenna electronics, and calibration receivers. The hardware is periodically polled to refresh the status displays. There is also a command line on which command text with parameters can be typed, and a syntax allowing this to support all subsystems, along with a scrolling text window to show responses to commands and error messages.
- A scripting facility allows a sequence of commands in a text file to be called for execution. The script-calling command can be included in such a file, so that execution of the files can be nested. The script language includes “pause” and “resume” commands for




Figure 7. Indoor equipment of the demonstration array. The master synthesizer and reference distributor each occupy a 2U-high chassis. The control computer is a rack-mounted PC occupying another 6U of space. An Ethernet switch supports a local network that includes a connection to the Internet. Other space in the 1-m-high rack is empty or unused. A keyboard and monitor allow local interaction with the control computer, but most operation of the array has been done remotely over the Internet.

manual pacing of execution, as well as a “pause <*n*>” command to wait *n* seconds and then resume, for automatic pacing. Although this facility is elementary compared with other scripting languages, it is very powerful and has enabled us to automate most of our tests.

- A logging facility automatically writes status information to a file. It supports a user-selected sampling period of 2 s to 60 min. Examination of these log files has been very useful for troubleshooting.

Other commercial or public software installed on the control computer is also important in operating the array. This includes Windows Remote Desktop (part of Windows XP) for general remote control; a File Transfer Protocol (FTP) server for remote transfer of files; and MATLAB⁹ for data analysis.

F. Installation and Geometrical Configuration

The array has been installed at a site called Loop Canyon, located in the Angeles National Forest about 17 km NNW of the Burbank, California, airport and 1144 m above sea level. The site is leased from the U.S. Forest Service for use as an antenna testing facility.¹⁰ Figure 8 gives a satellite view of the area and a map of the principal facilities.

The array is in a linear configuration oriented 4.6 deg east of north with a nearly uniform spacing of about 4.2 m. We identify the antennas as ANT1 through ANT5, numbered from south to north. The calibration receivers are identified as CAL1 through CAL4 in Figure 8 and are at distances of 20 m to 904 m from array antennas, at heights of 7.7 m to 55.8 m above the array, and appear at elevation angles from 3.5 deg (CAL2) to 24.0 deg (CAL1 from ANT5). The precise locations of the antennas and receivers were determined by a survey, discussed in Section IV.B.

IV. Test Methods and Results

A. Calibration Procedures

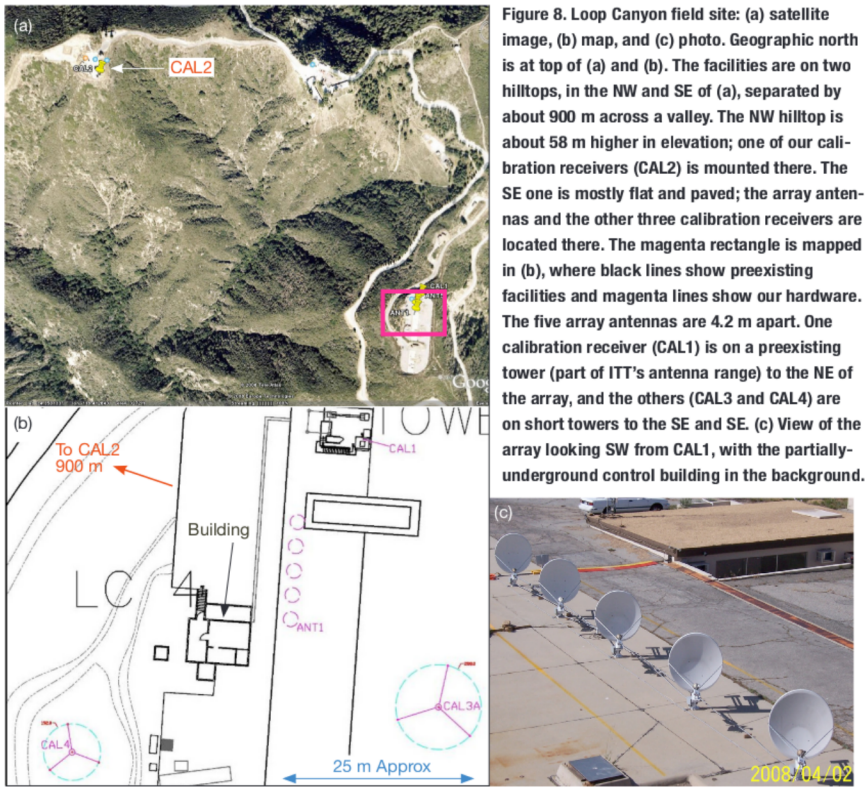
We desire to compute the carrier phase adjustment required at each antenna (from its nominal or “zero” setting) to produce in-phase signals at a distant spacecraft. Imagine a hypothetical signal radiated from a point at an antenna’s effective position (intersection of its rotation axes¹¹), with the (electric or magnetic) field magnitude at that point given by

$$a_i \cos[\omega t + \Phi_i + m_i(t)] \quad (1)$$

⁹ A product of The MathWorks, Inc.

¹⁰ The lessee is ITT Radar Systems–Gilfillan, a division of ITT Corporation, and our use of the site is under a contract with that company. Besides use of indoor and outdoor space, they provide Internet service via a radio link to a router at their offices in Van Nuys, California, as well as electrical power.

¹¹ There is no need to determine the location of the antenna’s phase center. The signal suffers a fixed delay between any point in the electronics (e.g., the waveguide flange of the feed) and the phase center or the axis intersection point or any other point of interest in the actual antenna’s optics. In the absence of severe mechanical distortions, that delay is independent of pointing direction. (For large antennas, it could depend slightly on temperature.) Selection of the point defined as the antenna’s position is therefore arbitrary, but selecting the intersection of axes allows it to be constant in the Earth-fixed coordinates in which the target position and the (fixed) positions of the calibration receivers are expressed.



where ω is the angular carrier frequency, Φ_i is the carrier phase, $m_i(t)$ is the phase modulation, and subscript i is the antenna number. Let $\Phi_i = \theta_i + \phi_i$, where θ_i is defined as the antenna-specific instrumental phase and ϕ_i is a controllable phase adjustment. These expressions establish the sign convention for our calculations. At distance r , the field magnitude becomes

$$\alpha a_i \cos[\omega(t - r/c) + \Phi_i + m_i(t - r/c)] = \alpha a_i \cos[\omega t - 2\pi r_i/\lambda_0 + \theta_i + \phi_i + m_i(t - r/c)] \quad (2)$$

where α accounts for the space attenuation, c is the propagation speed, and λ_0 is the wavelength.

All signals received at the target will be phase aligned if the adjustments are set to

$$\phi_i = 2\pi(r_i - r_0)/\lambda_0 - (\theta_i - \theta_0) \quad (3a)$$

where r_i is the propagation distance from that antenna to the target and r_0 and θ_0 are the distance and instrumental phase for the reference antenna, respectively. Here we have chosen to make no adjustment at the reference antenna ($\phi_0 = 0$). Since only the differences are important, let $\delta\theta_i = \theta_i - \theta_0$ and $d_i = r_i - r_0$, so

$$\phi_i = 2\pi d_i/\lambda_0 - \delta\theta_i \quad (3b)$$

and our task reduces to finding d_i and $\delta\theta_i$ for each antenna.

We assume that the target is in the far field of the array (not merely of the individual antennas) so that its direction from all antennas is the same, and that it is outside the atmosphere and at fairly high elevation angle (in practice, above about 10 deg). This ensures that the portions of the paths within the atmosphere are the same, so the path difference d_i can be taken to be in vacuum and λ_0 is the vacuum wavelength.

The two terms in Equation (3b) are determined separately. The first, involving d_i , depends only on the geometry of the signal paths. The second, $\delta\theta_i$, is derived from measurements made at the calibration receivers.

To calculate the path differences d_i , assume that an accurate ephemeris of the target is available and that it can be transformed (if necessary) into Earth-centered-fixed (ECF) coordinates. This system rotates with the Earth, has its origin at the Earth's center of mass, Z axis toward the geographical North Pole, and X axis toward zero geographical longitude. (This can be done for any spacecraft because the transformations from ephemerides expressed in other coordinates are straightforward and well known.) For the demonstration, where the targets are limited to GEO satellites, the ephemeris is often given directly in ECF coordinates by the satellite's operator. For those operated by Intelsat, a predicted time series is published on the Internet [5] and updated weekly.

Next, assume that the position of each array antenna is known in a local coordinate system. For convenience, we choose local topocentric horizon (LTH) [6] coordinates with the origin at the position of the reference antenna. Also assume that the LTH origin is known in ECF coordinates (since the ECF representation of any point is a fixed function of its geographical longitude, latitude, and height [7]). As described in Section IV.B, we have used a combination of local laser metrology and a prior GPS survey of the site to obtain these coordinates.

With this information, the desired path differences are given by

$$\begin{aligned} \mathbf{s} &= (\mathbf{y} - \mathbf{x}_0)/|\mathbf{y} - \mathbf{x}_0| \\ d_i &= (\mathbf{x}_i - \mathbf{x}_0) \cdot \mathbf{s} \end{aligned} \quad (4)$$

where \mathbf{y} is the (vector) position of the target, \mathbf{x}_i is the position of an antenna with $i = 0$ denoting the reference antenna, and \mathbf{s} is a unit vector in the direction of the target. Numerical evaluation of Equation (4) requires that all positions be in the same coordinate system,

so we first transform \mathbf{y} from ECF to LTH, and note that $\mathbf{x}_0 = (0, 0, 0)$ in the selected LTH system.

To determine the instrumental phase differences $\delta\theta_i$, we transmit to one or more of the calibration receivers with all of the adjustments ϕ_i set to zero. The carrier phases measured at receiver j are then

$$\rho_{ij} = \theta_i - 2\pi d_{ij}/\lambda + \gamma_j \quad (5)$$

where d_{ij} is the effective distance from antenna i to receiver j , λ is the wavelength in the local medium, and γ_j is the phase of the receiver's estimate of the carrier at the time of the measurement. (γ_j varies slowly over time because the receiver's reference oscillator is independent of that of the array.) Subtracting the measurement of the reference antenna from each of the others removes the receiver's phase:

$$\delta\rho_{ij} \equiv \rho_{ij} - \rho_{0j} = \delta\theta_i - 2\pi(d_{ij} - d_{0j})/\lambda \quad (6)$$

Whereas the receivers are inside the atmosphere (unlike a target in space), we must use $\lambda = \lambda_0/n$, where n is the average index of refraction along the path. Since $n - 1$ is small ($< 3 \times 10^{-4}$) and since the antennas and receivers are in a relatively small region, we assume that n is nearly constant throughout that region, and that it can be accurately estimated by local meteorological measurements. An accurate empirical formula is [8]

$$10^6(n - 1) = (0.77689 \text{ K/mbar})(p - p_w)/T + (0.712952 \text{ K/mbar})p_w/T + (3753.63 \text{ K}^2/\text{mbar})p_w/T^2 \quad (7a)$$

with

$$p_w = (611 \text{ mbar})(T_d/273.16 \text{ K})^{-5.3} \exp[25.2(T_d - 273.16 \text{ K})/T_d] \quad (7b)$$

where p is the total pressure, p_w is the partial pressure of water vapor, T is the (absolute) temperature, and T_d is the (absolute) dew point temperature. For the demonstration, we used hourly readings of p , T , and T_d obtained from a government weather station located about 1 km from the array and about 200 m west of the CAL2 receiver [9].

Near-field correction. The physical distance from each antenna to each receiver was determined by surveying (see Section IV.B). However, if the receiver is not in the far field of the antenna (roughly, $d_{ij} > D^2/\lambda$ for antenna diameter D), then the effective distance is slightly larger because rays from the edge of the aperture must travel farther than those from near the center. This is illustrated in Figure 9. If the wavefront in the aperture has zero phase and constant amplitude, then using geometrical optics [10], we find that the field at the receiving point is

$$E_r = E_a \int_0^{2\pi} \int_0^{D/2} \frac{r}{j\lambda d} e^{-j2\pi d/\lambda} dr d\theta \quad (8)$$




Figure 9. Development of model for the effective distance to a point not in the far field of an antenna.

where E_a is the field in the aperture and $d = \sqrt{r^2 + R^2}$ is the distance from a point in the aperture to the receiving point. Thus, there is an excess phase at the receiving point of $\arg(E_r) - 2\pi R/\lambda$ compared with transmission from a point source over the same distance R . Equivalently, the effective distance is $R + (\lambda/2\pi)\arg(E_r)$. It is this effective distance that must be used in Equation (6).

In practice, the apertures of reflector antennas are not uniformly illuminated, but rather the amplitude of the field at the edge is typically 20 to 30 percent of that at the center. We nevertheless use the above model as a reasonable approximation, but with an illumination diameter smaller than the physical diameter so as to account for the illumination. For the $D = 1.2$ m antennas of the demonstration array, it was empirically determined that using an illumination diameter of $D_{\text{eff}} = 0.94$ m gave the most consistent results among antennas at varying distances from the calibration receivers. This value was used in all the results reported here. The largest effect occurs for the shortest path, which in the demonstration array is 20.33 m from ANT5 to CAL1. Here, the calculated excess phase at 14.25 GHz is 0.77 rad, making the effective distance about 2.6 mm (0.013 percent) larger than the physical distance.

This approach allows placement of calibration receivers closer than the far-field distance, but they must not be too close. At slightly closer than $R_{\text{min}} = D_{\text{eff}}^2/4\lambda$, the fields from different parts of the aperture cancel to produce a null in amplitude, and many such nulls occur at closer distances [11]. This is accompanied by rapid changes in excess phase with axial distance and with radial offset from the beam center. Such rapid changes preclude accurate modeling for a practical antenna. For $D_{\text{eff}} = 0.94$ m at 14.25 GHz, $R_{\text{min}} = 10.5$ m.

Modulation timing. Not only the carrier phase, but also the modulation timing, must be aligned at the target. Returning to Equation (1), let $m_i(t) = m_0(t + \tau_i + \beta_i)$, where τ_i is an antenna-specific instrumental delay (analogous to the antenna-specific instrumental phase θ_i) and β_i is a controllable delay adjustment (analogous to the controllable phase adjustment ϕ_i). Then the timings will be aligned at the target if the adjustments are set to

$$\beta_i = d_i/c - \delta\tau_i \quad (9)$$

[cf. Equation (3b)], where $\delta\tau_i = \tau_i - \tau_0$ is the instrumental delay relative to the reference antenna and we choose to apply no adjustment at the reference antenna ($\beta_0 = 0$). d_i was determined in Equation (4), and $\delta\tau_i$ is obtained from the modulation timing delay dif-

ference measured at a calibration receiver, reduced to that at the antennas by subtracting $(d_{ij} - d_{0j})/(c/n)$. Here, c is the speed of light in vacuum and c/n is its value in the local medium. For the small demonstration array, the measured values of instrumental delay have been small (20 to 80 ns) and the baselines are short enough that the geometrical delays d_i/c are also small (<40 ns) compared with the minimum modulation bit period used in our satellite tests (7.8 μ s at 128 kbps). Therefore, no attempt was made to apply delay adjustments during those tests.

B. Survey of Geometry

To determine the antenna positions and the calibration receiver positions, a physical survey of the installed equipment was carried out in December 2007 and April 2008. A laser-based total station and a laser tracker¹² were placed about 10 m to the east of the array. Measurements of three-dimensional positions relative to the surveying instruments were made for all of the antennas and for receivers CAL1, CAL3, and CAL4. CAL2 was too far away (900 m) for distance measurement, but its angular position was measured. Also measured were various temporary reference targets on the ground and a nearby monument that had been included in a GPS survey commissioned by ITT in 2006.¹³ By combining these data with the GPS position of the latter monument and that of another monument measured in the same 2006 survey, we determined the antenna and receiver positions listed in Table 1.¹⁴ The details of the survey measurements and their analysis are somewhat complicated and are described in an internal report,¹⁵ so they will not be covered here.

From the surveying instruments' data sheets, the measurements of the antenna positions (relative to the origin at ANT1) are believed accurate to 0.2 mm, and those of the three nearby receivers to <0.7 mm. For CAL2, the position is much less critical. The angular measurement has an accuracy of about 1 arcsec (4.85×10^{-6} rad), corresponding to a transverse position error of 4.4 mm. The distance (determined by a complicated procedure¹⁶) may be in error by up to 0.3 m. For the receivers, errors in the horizontal transverse positions (i.e., in directions normal to the direction from the center of the array) have the largest effect on the calibrations. This is shown in Table 2, which gives the ratio of the change in $d_{5j} - d_{1j}$ to the change in coordinate x_j , y_j , or z_j for receiver $j = 1, \dots, 4$.

The tests described below were spread over a six-month period following the survey. It is possible that some position changes occurred during this time due to temperature variations, seismic activity, wind, and other stresses. From the consistency and stability of the calibration measurements it appears that any such changes have been small (see Section IV.F). The antennas, especially, are unlikely to have moved more than a small fraction of a mm because they are supported on stiff steel pipes bolted to a single, large, thick concrete slab. CAL1 is on a substantial steel structure close to the array and on a connected

¹² Leica TDM5005 and LTD500 series, respectively. The measurements were taken and initially analyzed by Timothy Sink of JPL's Communications and Ground Systems Section.

¹³ The 2006 survey was carried out by Stanton Land Surveys under contract to ITT.

¹⁴ The antenna coordinates used for the tests reported below were adjusted by 0 to 2.6 mm from those in Table 1 in order to make the results more consistent among the four calibration receivers. This is discussed in Section IV.H.

¹⁵ L. D'Addario, "Analysis of the Loop Canyon Survey Results," version 2.0 (internal document), Jet Propulsion Laboratory, Pasadena, California, June 16, 2008. See Appendix A.

¹⁶ Ibid.

Table 1. Coordinates from surveys.

Local Topocentric Coordinates, m			
	X (South)	Y (East)	Z (Up)
CAL1	-30.92259	12.80672	9.03611
CAL2	-541.17442	-721.73830	55.77300
CAL3	15.13267	25.06134	7.73003
CAL4	20.69329	-38.63926	8.51589
ANT1	0.00000	0.00000	0.00000
ANT2	-4.19189	0.33646	0.00285
ANT3	-8.34033	0.66544	-0.01358
ANT4	-12.57758	1.01452	-0.01007
ANT5	-16.76270	1.35682	-0.00756
Geographical position of origin (ANT1)			
Longitude, deg E			-118.4089533
Latitude, deg N			34.347804
Height above ellipsoid, m			1144.113

Table 2. Cal error sensitivity.

Change in ANT5-ANT1 Range Difference per Unit Coordinate Change			
	X	Y	Z
CAL1	0.19377	-0.19557	0.18421
CAL2	-0.00940	-0.01282	0.00062
CAL3	-0.24221	-0.28808	-0.06419
CAL4	0.14415	-0.21201	-0.03703

concrete foundation, so no secular movement should have occurred. CAL2 is also on a substantial structure, but it is across a deep valley from the array so seismic or other Earth movements might be possible; however, repeated GPS measurements of monuments in the area show no movement at the 1-cm level over several years. CAL3 and CAL4 are on guyed aluminum towers, but these are relatively short (9.1 and 6.1 m, respectively) and very stiff. Thermal distortions of each are possible; a 10 deg C change will affect the heights by 2.1 and 1.4 mm. CAL3 is on a concrete slab that is connected to that of the antennas, but CAL4 is on a nearby hill and is anchored to the soil by concrete-filled piers. At the end of the measurement period, the survey was repeated,¹⁷ showing small differences in coordinates but no significant change in calibration consistency.

C. Phase Calibration Transfer Among Receivers

Since we have four calibration receivers, we can use each separately to determine the antenna-specific carrier phases and modulation timings. In the absence of errors, the four sets of results should agree. Any disagreement provides an indication of possible errors. We can go further by using the measurements at one receiver to align the phases and timings at

¹⁷ L. D'Addario, "Surveys of Uplink Array at Loop Canyon" (internal document), Jet Propulsion Laboratory, Pasadena, California, November 25, 2008. See Appendix A.

another receiver, then using the target receiver to measure the alignment errors. That is, we use a calibration receiver to substitute for a spacecraft as the signal destination. This enables us to check the alignment performance without the complications of tracking a satellite. Results from such tests are shown in Table 3. A similar set of tests on each of four different dates, and at several different frequencies, is shown. In each set, calibration measurements were first made sequentially at all four receivers, then each receiver's data were analyzed to determine the phase adjustments that would cause carrier phase alignment at each of the receivers. This resulted in 16 separate sets of phase adjustments. We then pointed the array again at each receiver in turn, but this time applied the phase adjustments computed from each receiver's data. Two sets of measurements were made for each set of adjustments: first with the PRS modulation timing set differently for each antenna, allowing the receiver to separate the signals as is done during calibration; and second with all of the modulation timings aligned. The results of these measurements are given in the table. (Although Table 3 contains many numbers, the main results are given in the columns labeled "Phase/rms" and "Efficiency/Predict," which are in bold font. Each row reports measurements for one of the 16 combinations of calibration receiver and target receiver.) Ideally, the first measurements should show that all carrier phases are equal; any deviation represents an error. The measured phases are given relative to ANT1, so the ideal value in all cases is zero. The receiver also measures the signal amplitude from each antenna separately. From these, we can predict the amplitude that should be obtained when the modulation timings are aligned:

$$A_{predict} = \left| a_1 + \sum_{i=2}^5 a_i e^{j\phi_i} \right| \quad (10)$$

where a_i and ϕ_i are the measured amplitude and phase from antenna i . The ideal amplitude for aligned signals is the sum of the separate amplitudes, i.e., Equation (10) with all phases set to zero. These computed values are given in the columns labeled "Amplitudes/Predict" and "Amplitudes/Ideal," respectively. The ratio of these is the predicted combining efficiency, given in the last column. The measured amplitude with modulation timings aligned is also given; usually this does not agree with the prediction, as discussed below.

The "Phases/rms" column, which is the standard deviation of the five measured phases (taking ANT1 as zero), is a good measure of our success in achieving phase alignment. In the first complete set of measurements (from 9/25/2008, at 14.24425 GHz), this ranges from 0.097 to 0.389 rad. Not surprisingly, the best alignment for each target receiver is obtained when the calibration data from that same receiver are used. Even then, the alignment is not perfect; this reflects errors in implementing the computed adjustments, mainly due to the phase shifters. When calibration data from a different receiver are used, calibration errors (due mainly to antenna and receiver position errors, multipath interference, imperfect near-field correction, and refractive index errors) are added to the adjustment errors.

The "Efficiency/Predict" column shows the combining efficiency predicted from the measured phases and amplitudes. In the first two data sets of Table 3, these are always better than -0.43 dB and usually better than -0.2 dB, in spite of some instances of large phase errors (phases that depart from zero by more than 0.5 rad are shown in red). These numbers are believed to be representative of the efficiency achievable with a spaceborne target.

However, the measured amplitudes with modulation aligned (“Amplitudes/Meas.” column, with the corresponding combining efficiency shown in the “Efficiency/Meas.” column) are always less than predicted, and sometimes much less. One reason for this is that the RF detector in the calibration receiver is not truly square-law (linear in power); when the modulations are aligned, the average power increases by about 3 dB, leading to partial saturation of the detector and an apparent amplitude smaller than the actual amplitude. This effect should be smallest at CAL2, where the signals are weakest, and indeed the CAL2 measurements show the least deficit in measured efficiency. Another reason depends on details of the calibration receiver’s signal tracking and amplitude measurement, which assumes that the cross-correlation function of the received signal with the known PRS has a triangular shape extending ± 1 clock from a peak; in the presence of noise and non-ideal modulation, this is not quite true, leading to errors in the amplitude (but not the phase) measurements. It is also possible that occasionally one antenna’s modulation timing did not become synchronized as commanded. These effects explain most of the deficits in measured efficiency seen in Table 3.

The third and fourth data sets in the table (8/25/2008 and 8/21/2008) are different from the others in two ways. First, each set used only one receiver as the target, rather than all four, but it included measurements at four different frequencies covering most of the commercial uplink band, using calibration measurements at the same frequencies from all four receivers. Second, in implementing the calculated phase adjustments for each case, no correction for errors in the phase shifters was applied.

Each phase shifter is controlled by a 6-bit number, so that its nominal phase is $2\pi k/64$ rad for control value k (0 to 63). This gives a resolution of 5.625 deg (0.098 rad), and the manufacturer specifies a maximum error of 6 deg (0.105 rad). We have found that most of them have much larger errors, and that the error is somewhat frequency dependent. (See Section V for further discussion.) In our earlier tests, we simply rounded the desired phase value to the nearest nominal phase shifter setting, then applied the corresponding control code, ignoring any error. In our most recent tests, we made use of prior measurements of the actual phase shift at each control value for several specific frequencies. We then used the control value that produces an actual phase shift closest to the desired value. This accounts for the fact that the phase errors are generally larger for the lower two (earlier) data sets in Table 3.

D. Phase and Timing Alignment at Geostationary Satellites

Between May 16, 2008, and September 26, 2008, the performance of the array was verified by conducting test transmissions to GEO satellites. Four different satellites were used, ranging in longitude from 74 deg W to 123 deg W. All are operated by Intelsat. Figure 10 shows their positions in the sky as seen from the array, along with the positions of the calibration receivers.

The satellite amplifies the signal it receives, downconverts it in frequency by 2300 MHz, and transmits it back to Earth. We receive the downlink by means of a commercial low-

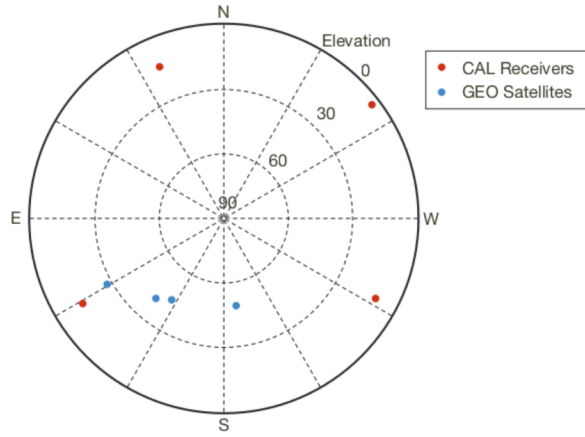


Figure 10. Azimuth and elevation of the calibration receivers and geostationary satellites, with zenith at the center, as seen from antenna #1. Proceeding east from north, the receivers are CAL1, CAL3, CAL4, and CAL2. The satellites are Horizons 2 (longitude 74 deg W), Galaxy 26 (93 deg), Galaxy 16 (99 deg) and Galaxy 18 (123 deg). Our site is at longitude 118.4 deg W.

noise block converter (LNB) mounted on one of our antennas (see Figure 3), and we observe its 950- to 1450-MHz IF on a spectrum analyzer. Whereas the satellite’s transponder is linear, the power that we receive on the downlink is proportional to that received by the satellite. To determine the combining efficiency of the array, we first transmitted with all antennas and then with each antenna separately, measuring the downlink power in each case. The ideal maximum combined power was then computed as

$$P_{\text{ideal}} = \left(\sum_i \sqrt{P_i} \right)^2 \quad (11)$$

where P_i is the power measured with antenna i alone, and the efficiency is $\eta = P_{\text{all}}/P_{\text{ideal}}$, where P_{all} is the power measured with all antennas transmitting. Measurements have been made with unmodulated carriers and with BPSK modulation at rates from 16 kbps to 128 kbps. Results of these tests are listed in Table 4. Each line of the table represents a separate test, and most were conducted within a 15-min interval whose starting time is given. The settings of the phase shifters needed to produce alignment at the satellite were computed as described in Section IV.A, using calibration measurements that were made 20 to 60 min prior to the test, and using the predicted satellite ephemeris for the time of the test as published by Intelsat [5].

Over the five months during which these tests were carried out, various refinements were made such that the latest results are believed to be the most accurate. For example, prior to July 16, the power measurements were obtained by manual measurements of high-resolution spectrum analyzer traces, limiting the accuracy of each value to ± 0.3 dB. Later, the analyzer’s bandwidths were optimized and its internal measuring facilities were used to achieve 0.01-dB resolution and accuracy better than ± 0.1 dB.

Table 4. Combining tests.

Date, 2008	Time, UTC	Cal File	Satellite/ W Long.	Freq., GHz	Mod, kHz	Power, nW		Ratio η ,		Notes
						Ideal	Meas.	dB	Δ	
5/16	22:30	may16c	G16 261°	14.39125	0	0.280	0.282	0.025		Slope 0.8 cycles
6/13	0:00	jun13a	G26 267°	14.38825	0	4.012	3.162	-1.033	[1]	Slope 0.595 cycles
6/13	13:00	jun13b	G26 267°	14.38825	0	7.121	5.012	-1.525		Slope 0.595 cycles
6/25	2:45	jun25b	G26 267°	14.37925	0	4.229	3.548	-0.763		Slope 0.595 cycles
7/3	2:30	jul3b	G18 118°	14.46500	0	n/a	0.794	n/a	[1]	1st of 12-hr spacing
7/3	14:30	jul3e	G18 118°	14.46500	0	0.544	0.562	0.144		2nd of 12-hr spacing
7/10	14:00	jul10b	G18 118°	14.46500	0	0.845	0.794	-0.268		Adjusted power balance
7/10	22:00	jul10e	G18 118°	14.46500	0	0.526	0.562	0.290		
7/11	2:00	jul11a	G18 118°	14.46500	0	0.641	0.711	0.455		
					16	0.394	0.316	-0.957		
7/16	0:00	jul15b	G18 118°	14.46500	0	0.345	0.356	0.128	[2]	
7/17	16:15	jul17a	G18 118°	14.46500	16	0.457	0.224	-3.095		
					128	0.110	0.066	-2.205		
7/17	23:30	jul17b	G18 118°	14.46500	0	0.476	0.376	-1.024	[3]	1st of 12-hr spacing
					16	0.334	0.224	-1.738		
					128	0.095	0.066	-1.560		
7/18	0:30	jul18a	G18 118°	14.46500	0	0.480	0.421	-0.571		2nd of 12-hr spacing
					0	0.536	0.455	-0.713		
					16	0.355	0.290	-0.877		
7/28	15:30	jul28b	G18 118°	14.46500	0	0.615	0.601	-0.097		
					16	0.276	0.267	-0.149		
7/31	22:15	jul31e	G18 118°	14.46500	0	0.500	0.390	-1.081		Software revisions
					0	0.514	0.405	-1.037		
					0	0.514	0.431	-0.767		
8/19	22:00	aug19b	G18 118°	14.46500	0	0.478	0.445	-0.316	[4]	
					16	0.204	0.188	-0.350		
8/20	17:45	aug20b	G18 118°	14.46500	0*	0.738	0.640	-0.619		Power settings -6
	19:15				0*	0.730	0.642	-0.556		Power settings -6
8/27	23:45	aug27d	G18 118°	14.47100	0	0.358	0.324	-0.437		1st of 12-hr spacing
8/28	11:30	aug28a	G18 118°	14.47100	0	0.774	0.622	-0.947		2nd of 12-hr spacing
					16	0.319	0.267	-0.769		
8/31	3:00	aug31a	G18 118°	14.47100	0	0.700	0.600	-0.673		Phase shifters uncal
					0	0.700	0.612	-0.583		phShCal jun24b 14.38825
					16	0.336	0.269	-0.980		phShCal jun24b 14.38825
9/9	20:15	sep9a	G18 118°	14.46500	0*	0.505	0.449	-0.510	[5]	phShCal sep5b 14.465
					64*	0.061	0.051	-0.755		
	22:00	sep9b	G18 118°	14.46500	0*	0.495	0.407	-0.849		
					16*	0.222	0.169	-1.180		
					32*	0.123	0.100	-0.914		
					64*	0.061	0.043	-1.496		
9/13	1:00	sep13a	G18 118°	14.46500	0	0.379	0.395	0.182		
9/15	16:00	sep15b	G18 118°	14.46500	0	0.164	0.112	-1.641		
9/16	0:00	sep15e	H2 74°	14.24425	0	2.033	1.614	-1.001	[1]	phShCalSep5b 14.465
					16	0.859	0.752	-0.578		
					128	0.125	0.092	-1.300		
	22:30	sep16b	H2 74°	14.24425	16	0.666	0.555	-0.795		Power balance adjusted
9/17	23:00	sep17c	H2 74°	14.24425	0	2.155	1.807	-0.765		phShCalSep16b 14.24425.
					16	0.937	0.883	-0.259		
9/26	18:00	sep26c	H2 74°	14.26675	0	1.495	1.208	-0.927		phStCalSep22a 14.26675.
					16	0.649	0.528	-0.895		
					128	0.071	0.058	-0.909		

When modulate rate is non-zero, "power" is spectral density in nW/10 kHz.

*Measurements at 55.2 MHz 5th IF; power in μ W or psd in μ W/10 kHz.

Significant changes (Δ column):

[1] New satellite.

[2] Improved spectrum analyzer setup.

[3] Begin offset pointing of "off" antennas.

[4] Fix error in refractive index calculation. Adjust pos. of ANT4, 5 from fits.

[5] Begin using phase shifter calibrations.

The earliest measurements, using the Galaxy 16 and Galaxy 26 satellites, achieved high combining efficiency only after modifying the calculated settings by adding a phase slope across the array, shifting the beam on the sky by about one array beamwidth. This was due to a software error that was corrected after July 1, 2008. Retrospective calculations show that in the absence of this error, the same results would have been obtained with zero added slope. Subsequent measurements have used the Galaxy 18 and Horizons 2 satellites. The early measurements of individual antenna power were made with the power from the other antennas reduced by 20 to 30 dB but not turned off completely. This was corrected during the 7/17/2008 23:30 test and thereafter.

Looking only at the measurements on and after 7/17/2008 at 23:30 UTC, the carrier-only efficiency results ranged from -1.08 dB (9/15/2008) to -0.10 dB (9/13/2008) with an average of -0.68 dB. We attribute about 0.5 dB of the variation to measurement error and the rest to actual differences in efficiency. Modulated signals usually (but not always) gave slightly worse efficiency during the same session. For these, the antenna power measurements are less accurate because the spectrum analyzer is looking at only a portion of the signal and the spectra are not ideal. Nevertheless, the 16 kbps results cover a similar range to the unmodulated results: -1.3 dB to -0.15 dB, except for one early result at -1.7 dB, with an average of -0.74 dB. There were only a few measurements at 128 kbps, and they showed slightly worse efficiency, -1.6 to -0.9 dB, average -1.25 dB.

Measurements prior to 8/31/2008 at 03:00 were made without attempting to correct for phase shifter errors, while later measurements attempted to use control codes based on prior calibration, as discussed at the end of Section IV.C. This had no obvious effect on efficiency.

On two occasions (July 17–18 and August 27–28), measuring sessions were conducted 12 hr apart and timed to coincide with the extremes of diurnal motion of the satellite as seen from the array (see Figure 11). Efficiencies of -1.0 dB or better were observed each time, confirming accurate tracking of the satellite motion. (A similar test on July 3 also showed accurate tracking by a different method.) Calculations show that if the phase settings that are optimum for one extreme were used at the other, the efficiency would have been less than -9 dB, even though the angular movement was only about 0.1 deg. In addition, the fact that accurate phase alignment was obtained multiple times for four different satellites at substantially different positions (Figure 10) using exactly the same procedures and software provides strong evidence that the system can achieve alignment at any position in the sky.

E. Error Budget

The results in Table 4 can be understood with the aid of the error budget given in Table 5. This gives the estimated value of each of the main contributors to the error in phase alignment. It includes separate estimates for the present demonstration and for a hypothetical operational array at 7.19 GHz. For the demonstration, the largest contributions are dis-

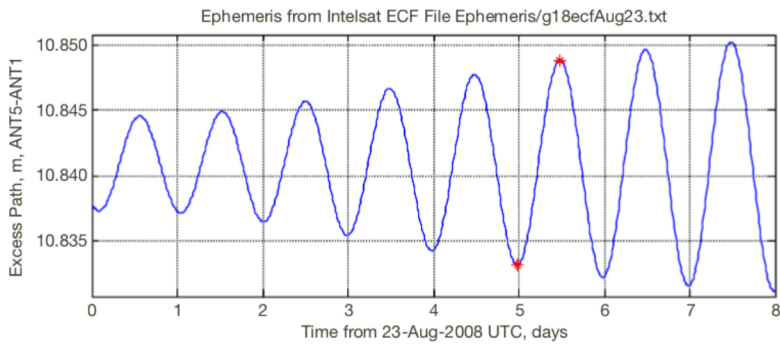


Figure 11. Path difference between ANT5 and ANT1 in the direction of Galaxy 18 for eight days beginning 8/23/2008. The times of two of the combining tests (8/27/2008 23:45 and 8/28/2008 11:30) are marked with red stars. The path difference changed by 1.5 cm, or 0.72 wavelength, between the tests. For the five-element array, this would cause a combining efficiency change of -9.0 dB if the phase settings were left unchanged.

Table 5. Phase alignment error budget.

Item Description	Small-Scale Demo at 14.5 GHz			Operational Array at 7.19 GHz		
	Value	Units	Effective Val., rad	Value	Units	Effective Val. rad
Noise in calibration meas. [1]	0.02	rad	0.02	0.02	rad	0.02
Knowledge of cal rcvr positions [2]	1.4	mm	0.16	0.25	mm	0.04
Knowledge of antenna positions [2]	0.5	mm	0.15	0.25	mm	0.04
Pointing errors during cal.	0.25	beam	0.05	0.1	beam	0.01
Freq. dependent error (incl. multipath)			0.30			0.10
Errors in near-field correction [4]			0.05			0.05
Refractive index error [3]	25	ppm	0.11	11	ppm	0.17
Electronics drift after calibration [5]			0.10			0.10
Mechanical changes after cal. [6]	0.1	mm	0.03	0.1	mm	0.02
Phase shifter inaccuracy [8]			0.20			0.05
Pointing errors on target	0.25	beam	0.05	0.1	beam	0.01
Atmosphere to target [7]	0.08	mm	0.02	2.1	mm	0.32
RSS, rad			0.46			0.40
EIRP loss, dB			-0.91			-0.68

Notes

“Operational” array is hypothetical. No decision to construct it has been made.

- [1] Includes effects of other antenna signals, receiver noise, and atmospheric turbulence.
- [2] Operational array will have improved initial survey for antennas and receivers, and sensors to measure receiver movement.
- [3] Local measurements of air pressure, temperature, and dew point are used. Accuracies 5 mbar and 2 deg C for demo, 5 mbar and 1.5 deg C for operational.
- [4] Near-field correction error assumed to be 10 percent of actual.
- [5] Electronics drift: new architecture for both arrays, 8 hr after cal, based on experience with demo.
- [6] Mechanical: Guess based on experience. Mainly wind and thermal. Demo worse for wind due to weak mounts, better for thermal due to smaller dishes.
- [7] Atmosphere: Operational array is worse because of its larger extent. Value is for 98th percentile weather at 20-deg elevation angle, based on Very Large Array data.
- [8] Phase shifters: Operational array will use custom-designed phase shifters integrated with RF electronics.

cussed later in this section. For the hypothetical array (which includes 156 antennas of 6-m diameter located on concentric rings with an outside radius of 319 m), various improvements in the instrumental errors are expected as described in the notes to Table 5 (partly because of operating at a lower frequency), but the effect of atmospheric turbulence is expected to be worse because the array must be spread over a larger area.

The error budget is intended to be conservative, i.e., biased toward overestimating the errors or using worst-case values. It takes account of two observations, each of which is discussed more thoroughly in later subsections of this paper. First, the antenna phases derived from measurements at the four calibration receivers have peak-to-peak discrepancies of 0.2 to 0.5 rad, implying that the calibrated carrier phases at the antennas may be in error by up to 0.25 rad. Second, studies of the variation of calibration results with frequency show deviations from linear up to 0.75 rad peak-to-peak. Some of this may be due to differences in transfer functions of the electronics (which would not imply any error since calibration and target transmission are done at the same frequency), but some of it may be due to multipath interference (which would cause errors).

For each contributor, Table 5 gives the estimated value in its natural units followed by its effect on the phase error. If the errors are independent, zero mean, and normally distributed, then the standard deviation σ of each antenna's total error is the root-sum-of-squares of the contributors. If the total errors are independent among antennas, then it is easily shown that the combining efficiency is $\eta = \exp(-\sigma^2)$. The resulting value in Table 5 (for the small scale demonstration) is $\eta = -0.91$ dB. This can be compared with the average measured losses from Section IV.D of -0.69 dB (carrier only) and -0.78 dB (16 kbps), confirming that the budget is conservative.

Knowledge of calibration receiver positions. Positions used for the calibration receivers were determined by the survey (Section IV.B). Based on the specifications of the survey instruments,¹⁸ the estimated error in the survey results is 0.4 mm at CAL1 and CAL2, and 0.7 mm at CAL4 in each coordinate. At CAL2 (900 m from the array), errors could be 5 mm, 30 mm, and 0.3 m in the transverse vertical, transverse horizontal, and range coordinates, respectively. To these errors we add the possible movement since the survey. CAL1 was removed and reinstalled once, and although this was done with some care an offset of up to 1 mm is possible. It is on a steel tower 9.1 m high, so a 10 deg C temperature change will cause a height change of 1.2 mm, but the structure is substantial so we expect no horizontal displacement. CAL3 and CAL4 are on aluminum towers 9.0 m and 6.1 m high, respectively, so 10 deg C change will change the height by 2.1 mm and 1.4 mm, respectively, and it is estimated that wind can cause up to 1 mm of horizontal deflection.

Combining these position error estimates with the sensitivities in Table 2, and taking the horizontal errors to be in the worst direction, we find that the largest error in path difference occurs for a 1.4-mm horizontal error at CAL3. This is 0.53 mm, or 0.16 rad at 14.5 GHz. These values are used in the budget.

¹⁸ The Leica laser tracker has specified 3-D errors of $\pm 15 \mu\text{m} + (6 \mu\text{m}/\text{m}) \cdot \text{range}$. It was used to measure the antennas and CAL1, CAL3, and CAL4. The Leica total station has angular accuracy of 0.5 arcsec, and was used to find the angular position of CAL2. Additional errors are expected for the antennas due to motion during the measurements, and for CAL2 due to reconciling of the coordinate system (see footnote 15: L. D'Addario, "Analysis of the Loop Canyon Survey Results"). These are included in the values given in the text.

Knowledge of antenna positions. Like the calibration receiver positions, the antenna positions rely primarily on the survey. In this case, the specified accuracy of the measuring instruments¹⁹ corresponds to 0.1 mm in all coordinates. Since the antennas were rotated during the survey measurements, additional errors are possible from bearing runout and from structural distortions due to wind and gravity. In some cases, we have repeated measurements on the same day (2007 Dec 13), and these show a maximum discrepancy of 0.49 mm. We therefore take the survey error to be 0.5 mm (and about 0.3 mm typical per coordinate). In view of the stiffness of the antenna support and underlying concrete slab, movement since the survey is believed to be negligible.

Some of the horizontal coordinates were adjusted from the survey values by solving for the offsets that minimize the measurement discrepancy among calibration receivers. This is described in Section IV.H. The final survey (after all tests were completed)²⁰ showed good agreement with the adjusted horizontal coordinates (0.25 to 0.7 mm), but improved techniques imply that the vertical coordinates may have been in error by up to 2 mm.

Multipath interference and other frequency-dependent calibration errors. This item in the budget attempts to account for the deviation from linear phase-vs.-frequency seen in swept-frequency measurements (Section IV.I). It is very difficult to see how multipath effects can cause errors as large as 0.3 rad rms, so some of the observed variation with frequency is likely due to other causes.

Phase shifter accuracy. The commercial RF phase shifters that are used to implement the phase adjustment and modulation at each antenna are not as accurate as expected from their data sheets. They have 6 bits of resolution (5.625 deg), and should have 6 deg maximum error, but tests show that errors as large as 30 deg (0.5 rad) occur. Typical errors are closer to 12 deg (0.2 rad), so that value was used in the error budget. The largest errors are associated with particular bits of the phase shifter, so they may be repeatable and nearly independent of frequency. The errors can be reduced by using the measurements to select, for any desired phase, the control code that produces the nearest phase (usually different from the nominal code). If this is done, the peak errors are predicted to be about 12 deg and the typical errors around 4 deg.

Most of the measurements in Table 4 were done using the nominal phase codes, but those made on 8/31/2008 and later attempted to use the calibrated codes. The effectiveness of the calibration depends on whether the errors are repeatable, and we have insufficient data on this. Therefore, the typical uncalibrated errors are used in the budget.

For a future operational array, much better phase shifter accuracy will be possible using a custom design rather than commercial off-the-shelf devices.

Atmosphere to target. Turbulence in the atmosphere when transmitting to the target can be a significant error source for a large array such as the hypothetical operational array.

¹⁹ Ibid. The survey instruments were closer to the antennas than to the calibration receivers, making errors smaller.

²⁰ See footnote 17.

However, the largest effects are caused by variations in air density over hundreds of meters, so the demonstration array (with 16 m extent) is expected to be affected very little. The value given in Table 5 for the operational array is based on results in [12], and it is scaled as $S^{5/3}$ where S is the largest antenna spacing in the array (following Kolmogorov turbulence theory) to obtain an estimate for the demonstration array.

F. Data Transfer

The combining efficiencies reported in Table 4 for the cases with modulated signals are based on power spectral density measurements, so they do not verify that no degradation of the modulation has occurred. In the absence of such degradation, it should be possible to demodulate the combined signal with a bit error rate (BER) that depends only on the SNR E_b/N_0 , where E_b is the energy received per bit and N_0 is the noise power spectral density, both of which can be determined from the power measurements. The only known array-specific effect that might cause degradation is a misalignment of the modulation timing among the antennas. As discussed in Section IV.A, the worst-case misalignment expected for the small demonstration array is 120 ns, which is 1.5 percent of the bit period at 128 kbps, the highest rate we have used. Therefore, we expect to see no measurable array-specific degradation of the modulation.

To verify this, we conducted a series of tests in which the downlink signal from a satellite was demodulated and its BER was measured. This was accomplished with the equipment setup shown in Figure 12, which made use of commercial hardware borrowed from other groups at JPL.²¹ The signal was received via the same LNB on one antenna (ANT4) as was used for the combining efficiency measurements. It was then converted in frequency from the LNB's IF to the input range of a satellite modem. The modem is capable of demodulating the BPSK signal, locking to its carrier and data clock, and producing a synchronous serial bit stream output (data and clock). (The modem can perform more complex processing, including descrambling and decoding, and it also includes a transmitting channel. We did not use any of these.) To examine the synchronous serial bitstream at the modem's output, we used the parallel port of a single-board computer running simple software to implement a data recorder. After each clock transition, a data bit was captured and saved in memory; when a requested number of bits had been captured, all were written to a file for off-line analysis.

The data transmitted by the array was the same repeating 1023-bit PRS as is used for calibration, but with the timing aligned at the antennas. The analysis software cross-correlated the recorded bit stream against the known sequence, reporting the number of errors found. The error rate can appear to be high if a bit slip occurs, but this did not happen for any of our data sets.

Results are listed in Table 6, which includes multiple sessions on each of three dates, and two different satellites. N is the number of array antennas used, which was always either

²¹ Norman Lay of the Communication Architectures and Research section provided the Icom communications receiver, and Thomas Jedrey of the Flight Communications Section provided the CommQuest satellite modem.

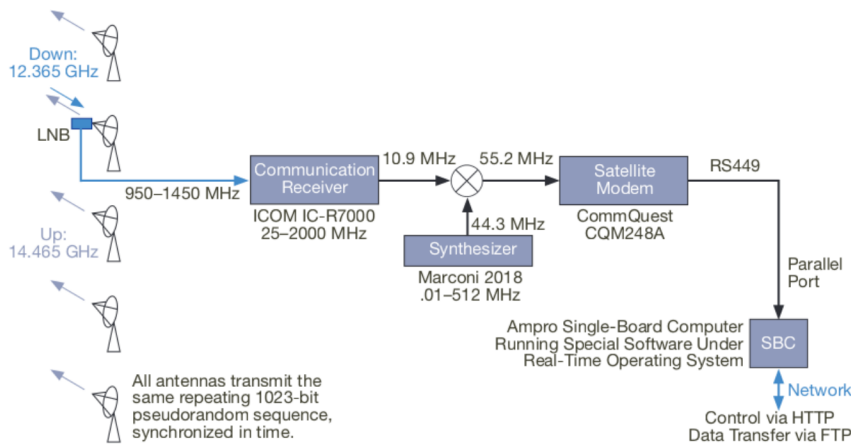


Figure 12. Block diagram of signal demodulation setup. One antenna has an LNB attached to its feed. A communication receiver is tuned to the LNB IF corresponding to our downlink carrier and the receiver's IF output is upconverted to the input band of a satellite modem. The modem locks to the signal and provides the demodulated bitstream and clock. We used a single-board computer and real-time software to capture data blocks up to several Mb long.

all ($N = 5$) or just ANT1 ($N = 1$). Before each data recording session, the total carrier power P_{carr} without modulation and the nearby noise spectral density N_0 were measured with the spectrum analyzer; these are believed to be accurate to ± 0.3 dB and ± 0.1 dB, respectively. From these and the data rate, the bit SNR E_b/N_0 was calculated (± 0.4 dB). The power being transmitted, and hence the SNR, was adjusted via the step attenuators in the RF modules. A normal (high-power) setting was previously determined for each antenna so that the radiated power was close to the same for all antennas (typically 3 dB peak-to-peak); to achieve lower levels, the attenuator settings at all antennas were increased by the same number of steps, where each step is approximately 1 dB. The level in steps relative to normal is given in Table 6, and P_{carr} was measured separately for each setting. Normally 1 Mb of data was recorded in each test, but occasionally 8 Mb were recorded. Data rates from 16 kbps to 128 kbps were used.

As expected, the tests with high SNR produced error-free recordings, giving only an upper limit on the BER. The only exceptions are some early tests at 16 kbps. It was later realized that this rate is below the minimum in the modem's specification (19.2 kbps); although the modem accepted commands for 16 kbps and appeared to lock normally, its performance was poor. By operating the full array (with carrier phases aligned) at levels of 12 to 14 dB below normal, we achieved nearly the same EIRP and SNR at the satellite as with antenna #1 alone at normal level.

The results highlighted in green in Table 6 are plotted in Figure 13 as a function of E_b/N_0 , along with the theoretical performance of an ideal BPSK channel with additive white

Table 6. Data demodulation tests.

Date, 2008	Time	Sat.	Freq., GHz	Pwr Steps	N	Rate, kbps	P_{carr} dBm	E_b/N_0 dB	Length, bits	Errors, bits	BER
9/9	J32	G18		0	5	32	-34.90	—	999471	0	<1E-6
9/9	M64	G18		0	5	64	-34.90	24.0	998448	0	<1E-6
9/9	N128	G18		0	5	128	-34.90	21.0	999471	0	<1E-6
9/9	P128	G18		0	5	128	-34.90	21.0	999471	0	<1E-6
9/9	Q128	G18		-9	5	128	-43.30	12.6	998448	10	1.00E-05
9/9	AH16	G18		0	5	16	-33.48	31.5	998448	6723	6.73E-03
9/9	Al64	G18		0	5	64	-33.48	25.5	998448	0	<1E-6
9/18	1	H2	14.24425				noise	-104.2	dBm/Hz		
9/18	1	H2	14.24425	0	5	16	-28.04	34.1	1021815	37175	3.64E-02
9/18	1	H2	14.24425	-6	5	16	-34.26	27.9	1021092	37384	3.66E-02
9/18	1	H2	14.24425	-12	5	16	-41.57	20.6	1021487	70495	6.90E-02
9/18	1	H2	14.24425	-12	5	64	-41.57	14.6	1021529	5	4.89E-06
9/18	1	H2	14.24425	-12	5	128	-41.57	11.6	1021254	46	4.50E-05
9/18	1	H2	14.24425	-12	5	128	-41.57	11.6	8189124	426	5.20E-05
9/18	1	H2	14.24425	0	5	128	-28.04	25.1	8189426	0	<1.2E-7
9/18	1	H2	14.24425	0	1	128	-28.04	25.1	1021233	18	1.76E-05
9/18	1	H2	14.24425	0	1	128	-28.04	25.1	5963232	131	2.20E-05
9/18	1	H2	14.24425				noise	-103.8	dBm/Hz		
9/18	1	H2	14.24425	0	5	32	-28.01	30.7	1021574	0	<1.0E-6
9/18	1	H2	14.24425	0	5	64	-28.01	27.7	1020996	0	<1.0E-6
9/18	1	H2	14.24425	-6	1	128	-48.51	4.2	1021713	22318	2.18E-02
9/18	1	H2	14.24425	0	1	128	-41.95	10.7	1021406	31	3.04E-05
9/18	1	H2	14.24425	-3	1	128	-44.51	8.2	1021052	2256	2.21E-03
9/18	1	H2	14.24425				noise	-102.7	dBm/Hz		
9/18	1	H2	14.24425	0	5	128	-26.86	24.7	1021890	0	<1.0E-6
9/18	1	H2	14.24425	0	5	128	-26.86	24.7	1021504	0	<1.0E-6
9/18	1	H2	14.24425	-6	5	128	-33.50	18.1	1021490	0	<1.0E-6
9/18	1	H2	14.24425	-12	5	128	-40.26	11.3	1020942	131	1.28E-04
9/18	1	H2	14.24425	0	5	128	-26.86	24.7	1021402	0	<1.0E-6
9/18	1	H2	14.24425	-6	1	128	-46.96	4.6	1021233	19674	1.93E-02
9/18	1	H2	14.24425	0	1	128	-40.11	11.5	1021054	27	2.64E-05
9/18	1	H2	14.24425	-3	1	128	-43.16	8.4	1021295	1013	9.92E-04
9/23	1	H2	14.26675				noise	-103.4	dBm/Hz		
9/23	1	H2	14.26675	0	5	128	-26.25	26.1	1023000	0	<1.0E-6
9/23	1	H2	14.26675	-14	5	128	-42.26	10.1	1023000	312	3.05E-04
9/23	1	H2	14.26675	0	1	128	-39.90	12.4	1023000	10	9.78E-06
9/23	1	H2	14.26675				noise	-104.5	dBm/Hz		
9/23	1	H2	14.26675	0	5	128	-27.57	25.8	8191161	0	<1.2E-7
9/23	1	H2	14.26675	-10	5	128	-37.84	15.5	1023000	0	<1.0E-6
9/23	1	H2	14.26675	-12	5	128	-40.11	13.3	1023000	9	8.80E-06
9/23	1	H2	14.26675	-14	5	128	-41.98	11.4	1023000	110	1.08E-04
9/23	1	H2	14.26675	0	1	128	-40.26	13.1	1023000	9	8.80E-06
9/23	1	H2	14.26675	-2	1	128	-42.19	11.2	1023000	91	8.90E-05

Gaussian noise. Within the measurement accuracy, the array and the single antenna produce the same error rate at the same SNR, with both being 2.5 to 3 dB below ideal. Many things can contribute to non-ideal performance. The array's modulators have response times of about 0.1 μ s, as well as errors in modulation index up to 10 deg. The timing alignment among the array antennas was not adjusted for antenna-specific offsets that are measured during calibration to be 20 to 80 ns, nor for geometrical delay differences to the satellite, which are up to 36 ns. These known effects within the array are theoretically negligible (<<1 dB from ideal). Other effects are common to the array and single-antenna tests. The modem's specification claims that its demodulator contributes <0.5 dB degradation from ideal. The receiving chain contains a total of five frequency conversions ahead of the modem, each of which contributes phase noise. (The LNB has a 10.75-GHz free-running dielec-

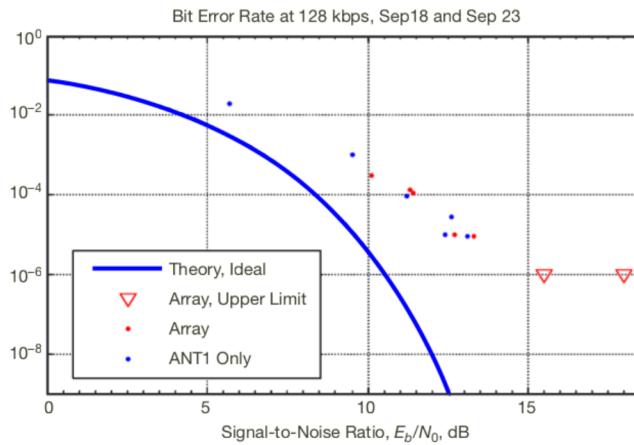


Figure 13. Measured and ideal bit error rate vs. signal to noise ratio, showing the results for antenna #1 (ANT1) alone separately from those for the phase-aligned five-antenna array. The measured SNR values are estimated to be accurate within ± 0.4 dB.

tric resonator oscillator; the Icom receiver has three stages of crystal oscillator-controlled frequency conversions; and our final upconverter from 10.9 to 55.2 MHz uses a commercial synthesizer.) There is also the frequency translation oscillator in the satellite's transponder (2.300 GHz). The phase noise from all of these oscillators could be a significant contributor to non-ideal performance.

The purpose of these tests was to look for unknown array-specific degradation, and none was found.

G. Calibration Stability

As mentioned in Section I, it is important that the carrier phases be as stable as possible so that the interval between recalibrations can be maximized. Design features intended to achieve this in the demonstration include the line-length stabilization servo of the reference distribution system (Section III.B) and temperature stabilization of the RF and microwave circuitry at each antenna.

An example of the stability achieved is shown in Figure 14. This shows the results of repeated calibrations at the same frequency over an interval of more than three weeks. Over a five-day gap in measurements near the beginning of the interval, the antenna carrier phase of one antenna (#5) had an unexplained change of 3.2 rad while the others changed about 0.2 rad, but over the remaining 16 days most antennas' phases remained stable within 0.3 rad. For unknown reasons, antenna #5 drifted more than the others, varying 1.24 rad peak-to-peak in the 16 days. For measurements made within any 24-hour period, no antenna's phase varied more than 0.3 rad peak-to-peak. All of these results are relative to the




Figure 14. Results of repeated calibration measurements at the same frequency over 21.3 days.

phase of antenna #1, which was taken as the reference. Each result plotted in Figure 14 is the average of the calculated phases from all but one of the available calibration receivers, where the receiver whose result was farthest from the average was excluded.²²

The phase changes in Figure 14 are unrelated to the ambient air temperature. The extremes of temperature among the measurement times occurred on July 25 (31.7 deg C) and July 28 (12.2 deg C), between which the phases of all antennas changed very little (+.08 to +0.17 rad).

Note that the phase error budget (Table 5) allows 0.1 rad due to electronics drift between calibration and target transmission, and an additional 0.03 rad for mechanical changes. These limits can almost always be met by making the interval between calibration and target transmissions sufficiently short, but our goal is to make that interval as long as possible. Apparently the architecture used here can reliably achieve this level of stability for at least 24 hr.

H. Calibration Consistency and Position Error Solutions

Figure 15 gives more detailed results from repeated calibrations at the same frequency, and it covers an even longer time interval (81 days) than Figure 14. It shows that the relative phases of ANT2 (used as reference), ANT3, and ANT4 remained very constant (less than 0.5 rad peak-to-peak) for the entire time, but those of ANT1 and ANT5 had larger variations. The large changes at ANT1 between 9/2 and 9/5/2008, and at ANT5 between 7/11 and 7/15/2008 and also between 8/19 and 9/2/2008 may have been due to some maintenance work or other deliberate change, but none has been identified. Note that there were many intervals of several days, and one of two weeks, when no measurements were made at this

²² See Section IV.H for a discussion of discrepancies among receivers. Excluding the outlier appeared to produce a slight improvement in phase estimates at one time, but it had negligible effect on the stability results reported here.




Figure 15. Calibration data at 14.465 GHz over an 81-day interval. Here, ANT2 is used as the reference rather than ANT1, because the latter had a large phase change (for unknown reasons) between 9/2 and 9/5/2008. For each antenna, the results computed from all calibration receivers are plotted separately, along with the average after excluding the outlier. All four plots have the same scale sizes.

frequency. It remains true that measurements made within 1 to 2 days of each other give very nearly the same results.

Figure 15 shows the results from each calibration receiver separately. We followed a procedure in which measurements were almost always made at all four receivers in immediate succession. A complete sequence was normally completed in 7 min. Occasionally difficulties caused it to take up to 30 min. Disagreements among results from different receivers in the same sequence can be caused by errors in the antenna positions, receiver positions, refractive index, or near-field correction; by multipath interference; or by variation of the effective antenna positions with pointing direction. During the first half of the period shown in Figure 15, the peak-to-peak discrepancy among receivers was 0.2 to 0.6 rad. During the second half of the period, and especially after 9/9/2008, the discrepancies tended to get larger for most antennas. This is not understood.

A complete sequence of measurements provides 16 separate antenna phase results (four antennas relative to the reference, using each of four receivers). Our main objective is to estimate the four actual antenna phases, but this leaves 12 additional degrees of freedom. We can use this redundancy to solve for some of the possible errors. It is not possible to solve for all identifiable parameters, since these include three coordinates for each antenna and each receiver, the refractive index n [(Equation 7)], and the effective antenna diameter D_{eff} [(see discussion after Equation (8)) — a total of 26 parameters. Furthermore, phase errors from non-parameterized effects like multipath interference can confuse the solutions. Nevertheless, attempts were made to solve for various subsets of the parameters over many data sets, looking for trends and consistency. In all cases, the solution is taken to be the one that minimizes

$$J = \sum_{i \neq 0} \sum_{j=1}^4 (\delta\theta_{ij} - \overline{\delta\theta}_i)^2 \quad (12)$$

where $\delta\theta_{ij}$ is the carrier phase difference at antenna i (from the reference antenna, $i = 0$) determined by receiver j , and $\overline{\delta\theta}_i$ is the average for antenna i among all receivers.

First, consider solving only for errors in the antenna positions. This neglects other causes of discrepancy such as receiver position errors and multipath, or assumes that their effects are separable. The four measurements for each antenna contain three separate discrepancy values, but we solve for only the X and Y (horizontal) components of position error because the geometry makes the measurements insensitive to errors in the vertical component (although CAL1 has significant sensitivity to the Z component, especially for ANTS; see Table 2). In an attempt to average out the effects of other errors, we have done this for many separate measurement sets.

In particular, multipath interference should be frequency sensitive, so the measurements include frequencies spread across the 500-MHz band. The solutions for 17 measurement sets taken at four frequencies between 7/16/2008 and 8/17/2008 are plotted in Figure 16(a). There is considerable scatter. In Table 7, the average at each frequency as well as the overall

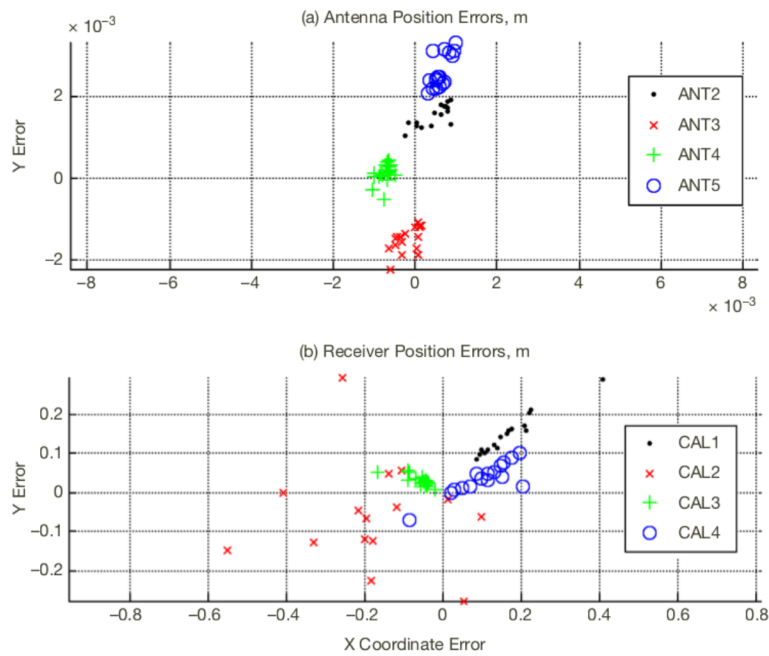


Figure 16. Solutions for (a) antenna position errors and (b) calibration receiver position errors from 17 calibration data sets, each including measurements from all four receivers. See also Table 7.

Table 7. Antenna position error solutions, mm.

Freq., GHz	No. Data Sets	A2 dX	A2dY	A3 dX	A3 dY	A4 dX	A4 dY	A5 dX	A5 dY
14.15	6	-0.031	1.261	-0.497	-1.645	-0.885	-0.025	0.497	2.551
14.23	1	0.872	1.331	-0.312	-1.557	-1.033	-0.296	0.589	2.243
14.44	2	0.548	1.688	-0.277	-1.608	-0.700	0.216	0.743	2.739
14.465	8	0.713	1.679	0.072	-1.351	-0.621	0.233	0.687	2.663
Overall Averages	17	0.441	1.512	-0.192	-1.497	-0.748	0.109	0.621	2.608
Standard Dev.	17	0.387	0.257	0.276	0.307	0.160	0.241	0.209	0.411
Pk-Pk by Freq.		0.903	0.427	0.569	0.294	0.411	0.528	0.247	0.496
6/10/2008 Adjust		0.875	1.400	-0.532	-1.400				
8/17/2008 Adjust						-0.822		0.623	2.562

average is given for each coordinate. To decide whether the solution for a particular coordinate is significant, we compare it with the standard deviation among all 17 solutions and with the peak-to-peak variation among averages at the four frequencies. In five of the eight coordinates considered, the error exceeds both by a factor of 3; for these, the average result was used to adjust the coordinate in Table 1. Two of the other coordinates (ANT2 X and ANT3 X) were also adjusted on the basis of earlier but less extensive data (6/10/2008 adjustments in Table 7). The present data do not confirm the latter adjustments, but they were nevertheless retained. For two others (ANT2 Y and ANT3 Y), the adjustment values determined earlier are confirmed by the present data. The remaining three adjustments (ANT4 X , ANT5 X and Y) are based on the present data. The adjustment values are not exactly the same as the average solutions because a different weighting was used, but the differences are not significant.

The adjusted antenna coordinates were used for all the ground-only tests reported here, including Table 3, Figure 14, and Figure 15. (Tests conducted before the adjustments were determined were reanalyzed using the adjusted coordinates.) For tests involving satellites, the signals were aligned by analyzing the calibration data using the best coordinate values available at the time of the test. Thus, almost all of the results in Table 4 (combining efficiency) included the 6/10/2008 adjustments and those on and after 8/19/2008 also included the 8/17/2008 adjustments. All results in Table 6 (data demodulation) included both sets of adjustments.

Similarly, we can now consider solving for only the calibration receiver position errors, keeping the antenna positions constant and ignoring all other sources of phase error. Again we consider only the horizontal coordinates (X and Y). Those solutions, for the same 17 data sets, are plotted in Figure 16(b). Here, there are several reasons to believe that the solutions do not represent actual position errors. First, they are mostly too large. Both the individual solutions and the averages are typically larger than 0.1 m, and it is hard to see how the survey errors could be that large. Second, the standard deviations among the solutions, as well as their variation with frequency, are large. After accounting for the sensitivities in Table 2, most of the standard deviations correspond to path differences greater than one wavelength, and sometimes many wavelengths. Third, the solutions do not reduce the phase discrepancy among the receivers by a large factor; in many cases, the peak-to-peak discrepancy remains more than 0.5 rad. In Fig. 16(b), it can be seen that the CAL2 solutions are widely scattered, as expected since its distance makes the results insensitive to its position (see Table 2). It can also be seen that the other solutions tend to fall on a straight line through the origin. These lines all correspond to displacement toward or away from the center of the array, which is the direction of minimum sensitivity, so the scatter along them is attributable to the other sources of error (not receiver position) and measurement noise. For these reasons, the solutions do not justify any adjustment to the receiver positions, and none were made. This does not mean that the receiver positions have no significant errors, but merely that we are not able to make corrections from the available measurements.

I. Frequency Dependence and Multipath Interference

The instrumental phase at each antenna θ_i [Equation (3)] is in general frequency dependent, for two reasons. First, the effective delay T_i between the master oscillators and the reference point of an antenna causes the carrier phase at the reference point to be $2\pi f_c T_i$, where the phases of the master oscillators are taken to be zero. The length compensation servo in the RGM (Figure 4) makes T_i small by cancelling the delay through the optical fiber, but T_i is not zero because of delays in the cables from the master oscillators to the RGM, from the optical splitter at the antenna to the RF module, through the RF electronics and power amplifier, and through the antenna optics and free space to the reference point. Each of these noncompensated delays is small, and their total corresponds to a physical signal path of less than 2 m. Furthermore, care has been taken in construction to make them nearly equal among the antennas, so that the residual delay differences $T_i - T_0$ are considerably smaller, where T_0 is the effective delay for the reference antenna. Second, carrier phase dependence that is not proportional to frequency occurs because the transfer functions of various components in the carrier generation signal path are frequency dependent. These include filters, mixers, and amplifiers, as well as imperfectly terminated transmission lines.

Such frequency dependence has no effect on the accuracy of phase alignment at a target spacecraft provided that calibration and target transmission are done at the same frequency. In an operational system, it would sometimes be convenient to avoid this constraint. This is feasible if the phase variation is smoothly varying over the frequency range of interest, so that it can be determined by calibration measurements at a few frequencies. The residual delays $T_i - T_0$ can easily be determined by multifrequency calibration. The effect of very small frequency changes, such as Doppler shift variation over a tracking pass (even for low-Earth-orbit satellites), is already negligible in the demonstration system.

However, frequency dependence that arises in the signal paths via multipath interference can degrade the phase alignment. This is a concern mainly for transmissions at low elevation angles, where reflections from the ground and objects near the ground, including the other antennas of the array, could be significant. The geometry of the demonstration array has been chosen to provide wide separation (>3.5 m) between the beam of each antenna and all other antennas in the direction of each calibration receiver and for all GEO satellites of interest. An operational array can be designed to provide any desired beam separations s for targets above a selected minimum elevation e_{min} by spacing the antennas at least $s/\sin e_{min}$ apart. If calibration receivers must be placed at lower elevations, the geometry can usually be designed to provide similar separation.

If the signal amplitude at the target (or calibration receiver) via the direct path is A and the net amplitude from all (undesired) indirect paths is A_u , then the maximum phase error caused by the undesired signal is $\tan^{-1}(A_u/A)$, and this occurs when the direct and indirect signals are in phase quadrature. Since the paths are of different lengths, their phase relationship varies with frequency and the peak-to-peak phase variation of the net signal is $2\tan^{-1}(A_u/A)$.

For the demonstration array, we have made swept-frequency measurements of the carrier phases at each calibration receiver over most of the 14.0- to 14.5-GHz commercial uplink band. A typical result (for CAL3) is shown in Figure 17(a). The plot shows each antenna's carrier phase relative to ANT1, reduced to the antenna's reference point. The average slopes are due to the residual delays $T_i - T_0$, which vary from -0.99 ns (ANT2) to $+2.27$ ns (ANT3). The best match to ANT1 is ANT4, whose residual is $+0.12$ ns. These results are the same within ± 0.08 ns when determined from measurements at the other receivers. Whereas all the cable lengths are within a few cm of identical among antennas and all electronic devices are nominally the same, it is believed that the residual delays are dominated by a VHF high-pass filter in the RF module (designed to separate the 437.5- to 453-MHz carrier reference from the 32.768-MHz clock reference). In any case, the residual delays are small and repeatable.

We are more concerned about the departures from constant slope. To the extent that they appear the same at all receivers for a given antenna, they could be attributed to the phase part of the transfer function of the antenna's electronics. Figure 17(b) shows the results from all receivers for one antenna (ANT3) after removing the common residual delay; other antennas are similar. While some features are reproduced at all receivers, there are substantial differences, including some large, narrow-band deviations at CAL1 and CAL4. These phase variations are hard to explain.

One explanation is multipath interference. Any indirect path involves reflections, and calculations indicate that it is very difficult to produce a net reflected signal larger than -30 dB from the direct one, and such a interfering signal would produce a maximum phase change of only 0.03 rad. Spillover past the dish goes mostly to the sky (due to the offset feed), and the small amount that hits the flat ground behind the dish is reflected away from the receiver. Even in the near field of the forward beam, less than 10 percent of the power falls outside a cylinder twice the dish diameter. For CAL2, which appears at elevation 3.5 deg, as much as 1 percent (-20 dB) of the radiated power might strike the flat ground in front of the antenna, but only a small fraction of this can be reflected into the receiver, and only a small fraction of that will be in phase quadrature with the direct signal. The other receivers are at much higher elevations (8.7 deg to 24.0 deg), so reflections from near the antennas should be insignificant. CAL1 is on a complicated steel structure, so reflections from near the receiver are likely. But all others are well isolated. (CAL2 is also on a large steel structure, but it has a 20-dB-gain horn with a clear field within -15 dB from boresight.) For CAL3, especially, there are no objects within a 20-m radius of its 9-m-high tower, and the ground between it and all array antennas is level pavement. The towers for CAL3 and CAL4 are stiffened by guy wires, but we also did swept-frequency measurements with the guy wires removed and could see no change in the phase-frequency response.

Another possible explanation involves the RF circuitry of the receivers. CAL1, CAL3, and CAL4 use short WR75 cutoff waveguides followed by waveguide-to-coax adapters as antennas. Reflections at the aperture and the adapter (as well as at a Mylar window between the waveguide flanges) can create a resonance, causing both large phase shift and loss. Indeed,

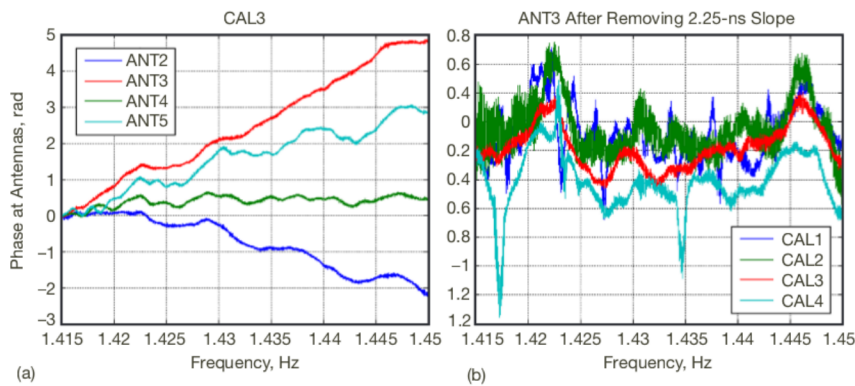


Figure 17. (a) Carrier phase measured at CAL3 for each antenna relative to ANT1 as frequency was varied from 14.150 to 14.500 GHz. The sweep was done at an accurately controlled rate of 3.8889 kHz/ms, which is slow enough to allow all elements of the reference distribution system and the RF modules to remain phase-locked. The receiver recorded a 1 ms integration every 10 ms. A calculated phase slope due to the geometrical delay from each antenna to the receiver was subtracted from the raw measurements, giving the estimated phase at the antenna reference point. (b) Similar measurements at all receivers for ANT3, after subtracting a common residual delay of 2.25 ns. (Note difference in vertical scales.)

the two large narrow-bandwidth features in the CAL4 results and many of those in the CAL1 results — Figure 17(b) — correspond to drops of 2 to 6 dB in signal amplitude from all antennas.

Further investigations of these effects were not possible before the array was disassembled on 9/30/2008. However, the calibration receivers are available for laboratory tests, so measurements of their frequency responses are planned. Any problems found should be correctable by minor modifications.

In view of the uncertainty here, the error budget (Table 5) includes a conservatively large item called “freq. dependency error” as a placeholder for all effects not otherwise listed that might be different during calibration and during transmission to a distant target. In spite of this, both the budget and the experimental results (Tables 3 and 4) show good combining efficiency.

J. Lessons Learned

We had some difficulties as well as some positive experiences from which to draw lessons.

Logistical planning should start early. The project was delayed because a suitable field installation site on the JPL campus could not be arranged as originally assumed. The delay might have been avoided if site selection had started earlier. Similarly, installation of the new towers proved more time consuming than expected, but it could have started earlier.

Funding profile affects efficiency and cost. The project eventually received all of the funding initially requested, but most of the time it was working with a smaller budget. This forced a deliberate slowdown for several months and increased the total cost.

Reliability of electronics could have been better. Most problems were caused by outdoor connectors that were inadequately protected from weather or whose weather protection was compromised, leading to corrosion. In some places, we used silicone sealant that was later found to be corrosive. More care in selection of components and materials might have reduced or eliminated these problems.

Protection circuitry is a good investment. Several Xilinx logic boards were destroyed when external power transistors connected to FPGA I/O pins failed during power transients. Circuitry to isolate and protect the more expensive parts should have been included from the beginning.

Calibration receiver development proved much more difficult than expected. Although there were no problems with the basic concept, implementation was intrinsically complex and led to a lengthy debugging process. This might have been anticipated.

Adapting and modifying commercial-off-the-shelf (COTS) products is effective. We used readily available assemblies and components whenever possible, achieving substantial savings of time and money over custom designs. Examples include modification of commercial BUCs to make the power amplifiers; choice of crystal oscillator and VCXO frequencies so that the parts were available off the shelf; and use of commercial “evaluation” boards to implement the digital module and the direct digital synthesizers (DDSs) in the master synthesizer.

Software requirements can be kept simple. By limiting the features of the control program to those actually needed, it was possible for one person to develop a rather powerful system in a short time.

V. Subsystem Details

A. Reference Distribution

The reference distribution subsystem consists of the master synthesizer and reference distributor assemblies, which reside in the control rack (Figure 7), and the network of optical fibers that connect the reference signals to the antennas. A functional block diagram is shown in Figure 18.

The master synthesizer contains an oven-stabilized crystal oscillator (Wenzel 501-04516), which acts as the master carrier reference for the array. The assembly synthesizes three UHF tones whose frequencies must be accurately set according to the final carrier frequency to be radiated. This is achieved using three DDSs, all clocked at 100 MHz, along with three submodules in which a UHF VCO is locked to the desired output frequency. The DDSs are on a single quad-DDS chip (Analog Devices AD9959) capable of operating at clock speeds

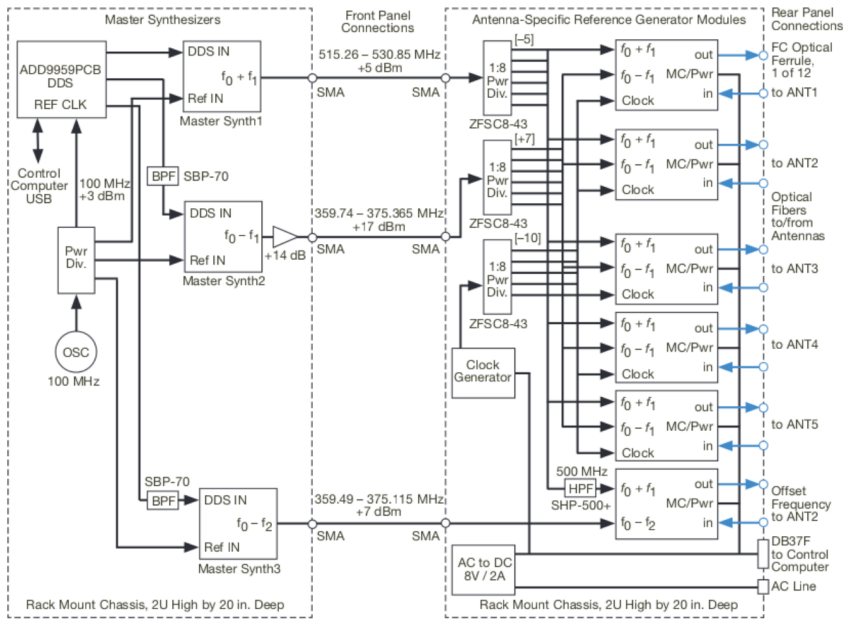


Figure 18. Reference distribution subsystem, consisting of master synthesizer (left) and reference distributor (right) assemblies. The master synthesizer generates the three tones that must be adjusted in frequency as a function of the carrier frequency to be transmitted. The reference distributor contains six RGMs (see Figure 20), all connected to the antennas via two-way optical fibers, as well as a clock generator module. Five of the RGMs provide a carrier reference tone (at $f_0 = f_c/32$) and clock signal to the antennas, and the sixth provides an offset carrier reference [at $(f_c - 4 \text{ MHz})/32$] to ANT2 for use during calibration.

to 500 MHz. Each of the three outputs has a different frequency range and different power level requirement, so each submodule works somewhat differently, but all have the same topology and share a common board design. The three frequency ranges are

$$\begin{aligned} f_0 + f_1 &= f_c/32 + f_1 = 515.260 \text{ to } 530.885 \text{ MHz} \\ f_0 - f_1 &= 359.740 \text{ to } 375.365 \\ f_0 - f_2 &= f_0 - f_1 - 2\delta f = 359.490 \text{ to } 375.115 \end{aligned}$$

where we have selected $f_1 = 77.760 \text{ MHz}$ as the frequency of the reference generator's line length correction VCXO, and where $f_0 - \delta f$ is the output frequency of the offset RGM.

Figure 19 is a block diagram of a master synthesizer submodule. The 100-MHz input drives a mixer at high level (more than +10 dBm) so as to generate harmonics. In the first synthesizer ($f_0 + f_1 = 515$ to 531 MHz), the fifth harmonic is mixed with a sample of the output, producing a 15- to 31-MHz IF for comparison with the signal from the first DDS. A double-balanced mixer acts as a phase detector, after which an active filter drives the VCO to form

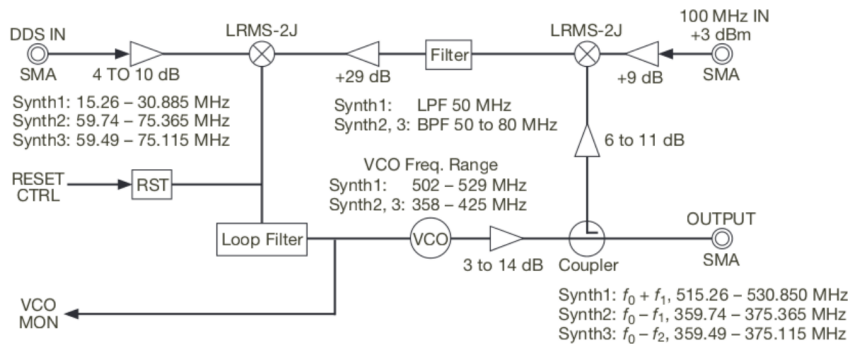


Figure 19. Master synthesizer submodule. Three of these are needed in the master synthesizer assembly. All have the same block diagram and use the same board layout, but some components vary in order to support different frequency ranges and output levels.

a phase-locked loop. In the other two synthesizers, the third harmonic of the 100-MHz reference is mixed with a sample of the output to produce a 59- to 75-MHz IF, which is compared in the phased detector with the second and third DDS signals, respectively. Note that the second and third DDSs operate in the second Nyquist band (50 to 100 MHz for 100-MHz clock), controlled by 55- to 85-MHz coaxial bandpass filters, whereas the first DDS operates in the first Nyquist band (0 to 50 MHz), controlled by a discrete low-pass filter on the DDS board. The gains of several amplifiers must be different among the three synthesizers because they have different output power requirements (see Figure 18) and because various components have different gains and available powers. Each of these amplifiers consists of a gain block integrated circuit (IC) cascaded with an appropriate attenuator.

The first two master synthesizer outputs, $f_0 + f_1$ and $f_0 - f_1$, are split by passive power dividers in the reference distributor to provide identical inputs to each of five RGMs. Care was taken to make the five cables from each divider the same length. The third drives only the offset RGM. All splitters are eight-way, and the reference distributor contains space for additional RGMs, so it could easily be expanded to support more than five antennas.

The reference distributor includes the clock generator module. This contains a fixed 32.768-MHz crystal oscillator and logic to convert its output to a square wave whose phase is reversed every 1.000 ms (500-Hz biphasic modulation). The resulting composite signal is passively split and delivered to the five RGMs. The 500-Hz square wave is recovered at each antenna and used to synchronize the modulation timing. The clock generator also has a 500-Hz square-wave logic output, intended to provide an interrupt to the control computer so that it too could be synchronized with the modulation timing. This was not needed during the demonstration. It would be needed for transmitting continuous data streams, so that the computer would know when to refill the data buffers at the antennas. The clock generator's oscillator is free running and thus not coherent with the master oscillator from which the carrier is generated. Making them coherent would have some small advantages, including slight simplification of the calibration receivers, but it was not necessary for the demonstration.

The RGM concept was discussed in Section III.B and illustrated in Figure 4. A more detailed block diagram is shown in Figure 20. The optical output of the diode laser (upper right) varies nearly linearly with its current, which is the sum of three signals: DC bias from a laser driver IC, the UHF carrier reference tone at f_0 , and the composite clock signal centered at 32.768 MHz. Clock harmonics are limited by a low-pass filter. The clock frequency was chosen to ensure that no harmonics fall in the carrier reference range. The DC bias is adjusted to provide 1.0 mW of average optical power. The carrier reference level is then set to produce an optical modulation index of approximately 0.25 (0.25 mW peak to peak), and the clock level is set about 6 dB lower. The laser driver measures the optical power by means of a photodiode that is built into the laser's package and servos the bias current to maintain constant average power. Monitor voltages proportional to laser current and optical power are provided to the control computer, and a control bit can disable the laser.

The bandpass filter in the carrier reference path suppresses undesired mixing products. An identical filter (and amplifier) is used in the return signal path to keep the two directions as similar as possible, even though the phase-lock loop (PLL) would be sufficient to suppress any undesired signals. Any change in phase delay through the electronics will be cancelled by the servo in the same way as for changes in the fiber length provided that both directions change by the same amount.

An analog phase detector (double-balanced mixer) is used in the PLL. A digital phase-frequency detector could not be used here because the undesired mixing products in the signal combine nonlinearly in the detector's comparator and cannot be adequately filtered by the loop. Lock acquisition (sweep) circuitry is not needed because the loop bandwidth

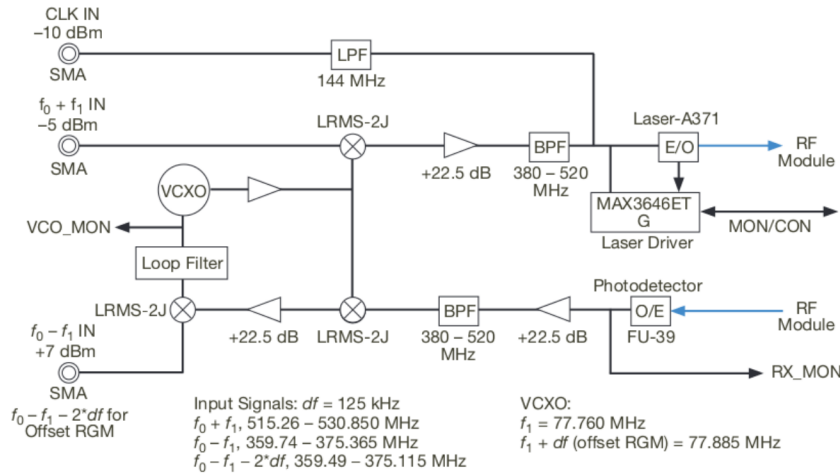


Figure 20. Reference generator module block diagram. Signals common to all antennas are input on the left, and optical signals to and from the antenna-mounted RF module are connected on the right. The clock signal (32.768-MHz carrier with 500-MHz biphasic modulation) is added to the carrier reference (437.5 to 453.125 MHz) and the combined signal is used to modulate a diode laser. See text for further explanation.

is comparable to the VCXO tuning range and the signal level is large enough to overcome detector DC offset. Similar considerations apply to the design of the master synthesizer PLLs, but there the VCO ranges are much larger. To enable lock acquisition, an intentional DC offset is included to drive the integrator slowly toward either minimum or maximum VCO tuning voltage, according to the state of a “reset” bit from the control computer. The computer toggles this bit after each frequency change to ensure locking.

The overall system for carrier generation at any one antenna is relatively complex, involving a master oscillator (100 MHz), two DDSs, three locked oscillators common to all antennas (two VCOs in the master synthesizer and one VCXO in the RGM), and one microwave-locked oscillator at the antenna. Each DDS and each oscillator contributes phase noise to the final carrier, as does the additive noise of the optical link. There are also numerous opportunities for generation of spurious signals. Careful choice of frequencies, selection of components (especially the oscillators), and design of the PLLs is necessary to minimize the phase noise and avoid significant spurious signals. Various architectures were considered during development before the one described here was selected. In the final design, all PLLs are type 2 and second order, with noise bandwidths of approximately 20 kHz for the master synthesizers, 5 kHz for the line length correction (VCXO), and 200 kHz for the RF module (7 to 7.25 GHz VCO). A model of the overall phase noise has been computed using information from the manufacturers’ data sheets of the various components. The result is plotted in Figure 21, which shows the total noise power spectrum referred to the 14-GHz output along with each part’s contribution. The integrated total from 10 Hz to 1 MHz is 0.43 deg rms.

In the original concept for the master synthesizer, it consisted of three off-the-shelf laboratory synthesizers locked to a common reference. Unfortunately, we found that the available units had unacceptable phase noise. While the newest and highest-quality instruments may have been adequate, buying new ones was not feasible within our budget. Units available for rent (commercially and internally to JPL) were found not to meet their original phase noise specifications. This led to the in-house design described here, where much better noise was achieved by tailoring the synthesizers to the limited tuning range required for our system.

B. RF Assembly and Power Amplifier

The main RF components at each antenna are contained in a temperature-controlled and weatherproof assembly called the RF box, mounted near the feed as shown in Figure 3. The interior of an RF box is shown in Figure 22.

The critical components, including the RF module, phase shifter, and optical power splitter (see Figure 2), are mounted on an aluminum plate supported on a 24-W thermoelectric cooler (TEC) whose external heat sink and fan can be seen at the bottom of the box in Figure 3. The cooler is driven by a proportional-integral servo connected to a thermistor on the plate, resulting in temperature stability of ± 0.3 deg C within the RF module during diurnal

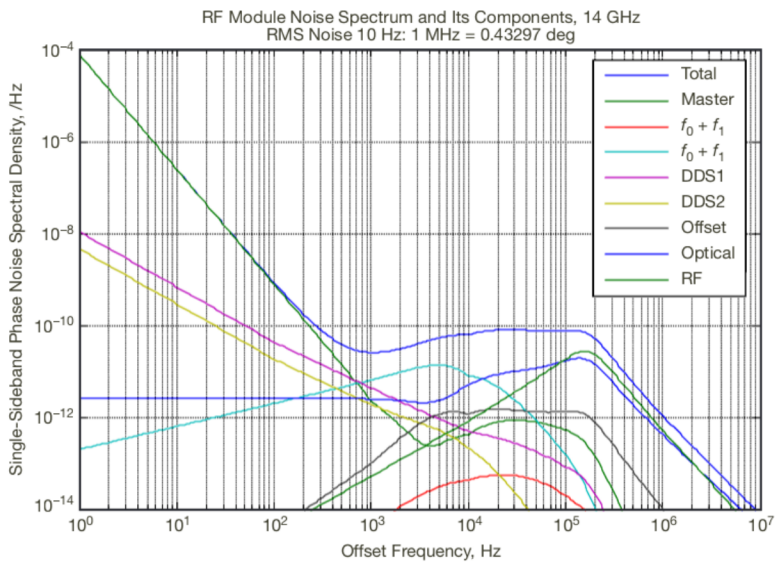


Figure 21. Model of carrier phase noise vs. offset frequency, computed for the final architecture using specifications from data sheets of the components. Contributions from all oscillators, from the DDSs, and from the optical link (laser noise) are included and plotted separately, along with the total.




Figure 22. RF box with cover removed. The critical components are mounted to an aluminum plate supported by a thermoelectric cooler. The optical power splitter lies along the lower edge of the RF module, with excess fiber loops taped to the temperature-controlled plate. The interior is insulated with Styrofoam blocks, some of which were removed for this photograph. Noncritical components are thermally connected to the outside wall.

changes of outside temperature exceeding 20 deg C. The TEC could also provide heating at low outside temperatures by reversing its current, but the available commercial servo provides only unidirectional control. Therefore, a 9-W (fixed) resistive heater is also mounted to the plate, and it is turned on by a thermostat when the outside (box wall) temperature falls below 15 deg C. When the heater is off, the total dissipation of components on the controlled plate is 7.4 W. Given the dissipation, the specified performance of the TEC, and measured thermal resistances, accurate temperature control is expected over the outside temperature range -1 deg C to +40 deg C. The TEC, its controller, and the heater all operate from a dedicated 115-VAC to 15-VDC switching power supply. No other devices are connected to this supply, none share a return connection with it, and neither of its terminals is grounded to the structure. The TEC controller uses pulse width modulation, but we have seen no evidence of the switching frequency on any of our signals.

Also included in the box is a small board carrying an eight-channel analog to digital converter (ADC) for monitoring temperatures and signal levels. The ADC has a serial data connection to the digital module.

A block diagram of the main internal assembly, the RF module, is shown in Figure 23. The optical input is demodulated by a photodetector and amplified, then the carrier reference (near 450 MHz) and the timing reference (32.768 MHz with 500-Hz modulation) are separated in a passive diplexing filter. The timing reference is doubled to 65.536 MHz, removing the modulation; this is the clock for the digital module's logic. The modulated signal is also delivered to the digital module, where the 500-Hz timing signal is recovered. Meanwhile, the carrier reference is delivered to a phase-frequency detector to lock a microwave VCO at 7 to 7.25 GHz, 16 times the reference. The power level from the locked oscillator is adjusted by a step attenuator with 0.5-dB resolution, controlled by the digital module. The signal is then doubled in frequency and amplified to produce the 14- to 14.5-GHz output. The doubler acts roughly as a square-law device, causing the step attenuator to have about 1 dB resolution at the doubled output. The signal at 7 to 7.25 GHz is also delivered to an output, allowing the same module to be used in the 7.145- to 7.235-GHz space research band of many NASA spacecraft.

As mentioned in Section IV.C, the phase shifters have proved to be less accurate than originally expected. The nominal phase shift is $2\pi k/64$, where $k = 0$ to 63 is the 6-bit control number, but tests show that the actual phase shift is often 0.25 to 0.5 rad from nominal. The error pattern is systematic and appears to be repeatable, so in many of our later tests we attempted to compensate for the error by using the previously measured phase at each value of k to compute the setting that produces the phase closest to the desired value. An example is shown in Figure 24.

Each power amplifier was built by modifying a commercial block upconverter. The BUC expects an input at about 1 GHz for conversion to 14 to 14.5 GHz using an internal local oscillator (LO) phase-locked to a 10-MHz reference. We bypassed the upconverting mixer, disabled the LO, and added the ability to turn the power supply on and off under control of a logic signal from the digital module. The BUC assembly is rugged and weatherproof,

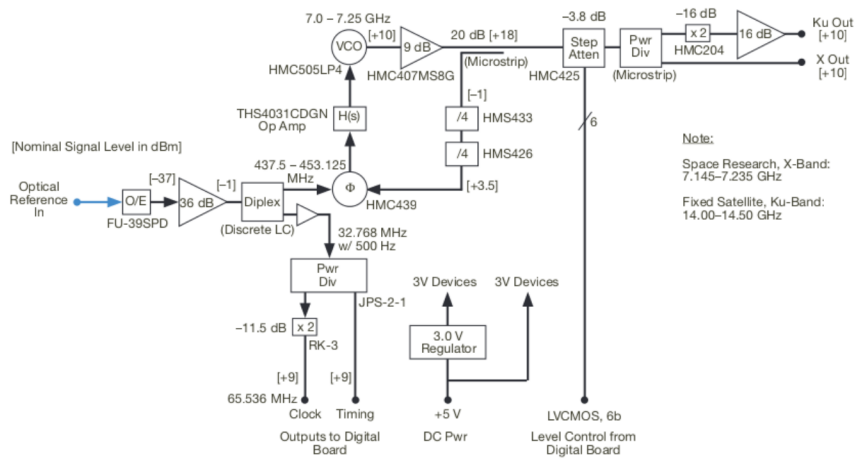


Figure 23. RF module block diagram. The UHF carrier reference is separated from the modulated 32.768-MHz clock/timing reference in a passive diplexer. The carrier reference locks a microwave VCO at 7 to 7.25 GHz. Power is adjusted by a 0 to 31.5 dB step attenuator with 0.5-dB resolution. The 14- to 14.5-GHz output is produced via a frequency doubler and amplifier.

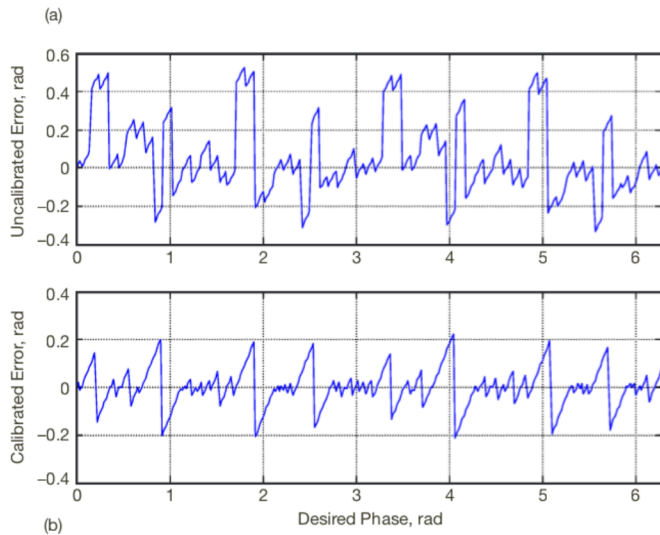


Figure 24. Phase shifter error for antenna #5, showing (a) measured error vs. desired phase using the nearest nominal 6-bit control number, and (b) using the control number producing phase closest to desired.

being designed for outdoor use over a wide temperature range. We nevertheless added an internal temperature sensor, an external fan, and software in the digital module to control the fan thermostatically.

C. Digital Assembly

The digital module consists mainly of a commercial board (Xilinx Spartan 3E Starter Kit) containing a small FPGA (XC3S500E), several memories, and input/output devices, including an Ethernet transceiver and connectors wired to many of the FPGA's I/O pins. It is housed in a weatherproof enclosure called the digital box, which is mounted on the feed support arm of each antenna, as seen in Figure 3.

The FPGA logic runs on a 131.072-MHz clock, obtained by doubling the 65.536-MHz signal received from the RF module. A 2-ms-period "tick" is generated by demodulating the 500-Hz reference from its 32.768-MHz carrier, and this is used to synchronize timing among the antennas.

In our application, the digital module's primary purpose is to generate the sequence of 6b numbers needed to drive the phase shifter in the RF box. It also controls the 6b step attenuator in the RF module, turns power on or off to the RF module and the power amplifier, and monitors various voltages and currents via the ADC in the RF box. At any moment, the phase shifter control output consists of the sum of three terms:

$$k = k_0 + nr + a_m M(t)$$

where k_0 , r , and a_m are fixed-point numbers that can be set by command from the control computer; n is the number of 20-ms ticks since the last phase initialization command; and $M(t)$ is the binary modulation function, a stream of bits clocked at the modulation frequency. The phase rate r is intended to facilitate tracking of a rapidly moving spacecraft; it was not needed during this demonstration, so we always set $r = 0$. Setting $a_m = 32$ produces full-index BPSK (since the phase difference between $M = 0$ and $M = 1$ is then 32/64 cycle, or 180 deg), and $a_m = 0$ produces an unmodulated carrier. The modulation bit stream can be taken from a buffer memory filled by the control computer, or from a 1023b PRS generator built into the FPGA.

The modulation clock is generated in the logic by dividing down the main clock, and it can be set by command to frequencies from 1 kHz to 16.384 MHz by factors of 2. The timing of the modulation clock can be adjusted relative to the 2-ms ticks in order to control the synchronization of the modulation among antennas. The logic causes the first rising edge of the modulation clock to occur a specified number of 131-MHz clocks from the next 2-ms tick after a "trigger" command. Thus, the time-setting resolution is $1/(131 \text{ MHz}) = 7.63 \text{ ns}$. Due to the small extent of the demonstration array and the fact that the natural synchronization error (with all timing offsets at zero) turned out to be less than 80 ns, this feature was not needed during the demonstration. When the 1023b PRS is being transmit-

ted, the first bit transmitted after a trigger command is bit S of the sequence, where $S = 0$ to 255 can be set by command. This provides the offset that allows signals from different antennas to be distinguished at the calibration receivers.

The digital module communicates with the control computer via its Ethernet port using TCP/UDP and a simple set of binary commands. Each module has a unique IP address in the same subnet. It also accepts packets sent to the subnet's broadcast address so as to allow a single command to be sent simultaneously to all antennas (especially trigger commands). The command syntax provides for a response to each received command within 1 ms, so the connectionless UDP protocol is adequate. Every command and response fits within a single 1500-byte Ethernet packet. Command processing in the digital module is handled by a "soft" microprocessor constructed within the FPGA. It is a 32b RISC processor whose FPGA implementation is generated by tools from Xilinx. We have programmed it in the C language using a development system and libraries also provided by Xilinx. The processor is driven by a free-running 50-MHz clock from an onboard crystal, rather than from the 131-MHz system-synchronous clock used by the rest of the logic. This allows communication to be maintained even when the external clock is absent. The board also has modes in which the logic clock and tick signals are simulated from the 50-MHz clock to facilitate laboratory testing of an isolated module when the rest of the system is not available.

The FPGA logic definition and the executable software are kept in nonvolatile memories on the board, and both are loaded automatically when power is applied.

D. Calibration Receivers

The calibration receiver concept was discussed in Section III.D, and a more detailed block diagram is given in Figure 25. The required functions are divided among a few analog components, an FPGA, and a microprocessor.

Input circuitry. The estimated signal power at the receiver's antenna terminals ranges from $5.8 \mu\text{W}$ (20-dB-gain horn at 900 m distance with one array antenna transmitting 1.0 W) to 4.2 mW (cutoff WR75 waveguide at 20 m distance with all antennas transmitting). Attenuators are used at the nearby receivers to reduce the maximum to about $263 \mu\text{W}$, producing a 16-dB dynamic range at the square-law detector. The detector mixes the combined modulated signals from all antennas with the reference tone from ANT2 to produce an IF centered at 4 MHz. This signal is filtered and amplified so that its minimum peak-to-peak amplitude is 10 times the quantization interval of the ADC ($1.0\text{V}/2^{12} = 0.244 \text{ mV}$). Analysis shows that the system noise at the ADC input (including detector shot noise, preamp noise, resistor thermal noise, quantization noise, and additive ADC noise) is then 27 dB below the signal. The saturation level of the 12-bit ADC is 52 dB higher. Thus, both the sensitivity and the dynamic range are more than adequate.

Timing. The 10-MHz crystal oscillator is the receiver's timing reference. It provides the ADC's sampling clock and the clock for all FPGA logic. Within the FPGA, NCOs synthesize

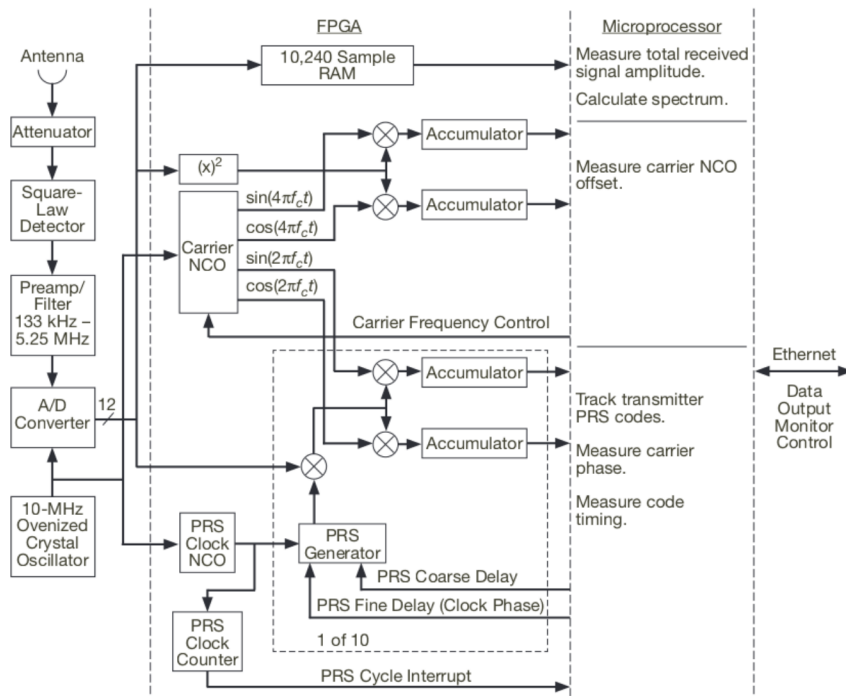


Figure 25. Functional block diagram of a calibration receiver. Modulated signals from the antennas are received along with an unmodulated tone 4 MHz lower in frequency. These are mixed in a microwave detector to form a 4-MHz IF, which is digitized and then processed by logic implemented in an FPGA. The logic has three main functions. A RAM captures short blocks of samples for diagnostic purposes; a carrier detector allows measuring the carrier frequency; and 10 copies of a code-synchronous detector allow tracking the signals from all antennas. The microprocessor implements a delay tracking loop by controlling the timing of the local PRS in each detector, while simultaneously recording the timing information and the complex amplitude of each carrier. These measurements are then available for transfer to the array control computer.

the frequencies of the carrier (nominally $f_c = 4.0$ MHz)²³ and the modulation clock (nominally 1.024 MHz). These are adjustable under software control to match the actual signals, whose frequencies depend on the independent oscillators of the array. For each complete cycle of the 1023-bit PRS (1023 cycles of the synthesized modulation clock), an interrupt to the microprocessor is generated; this synchronizes the hardware- and software-based parts of the processing.

Signal processing logic. The 12b, 10-MHz sample stream from the ADC is delivered simultaneously to multiple processing functions in the FPGA.

²³ In this subsection, f_c refers to the carrier IF in the receiver, whereas elsewhere in this article it refers to the carrier RF at 14 to 14.5 GHz.

First, a memory allows capturing (upon command) 10,240 successive samples. These can then be read by the microprocessor, where statistical and spectral analysis functions are implemented in software. This feature has been valuable for testing during development and for diagnosis of problems during operation.

Second, a carrier detector is implemented by squaring each sample to remove the modulation from any one antenna's signal,²⁴ producing a sinusoidal signal component at twice the frequency of the unmodulated carrier. This is then correlated against the $2f_c$ outputs of the carrier NCO for 1.0 ms, allowing a measurement of the phase difference between the received carrier and its local estimate. By making several successive measurements, the frequency difference can be calculated by the microprocessor. This works reliably provided that the frequency difference is less than about 50 Hz, or 25 ppm with respect to the nominal 8-MHz frequency, so that the phase change in 1 ms is less than 0.1 cycle, but in practice the oscillators are never more than 1 or 2 ppm apart.

Finally, the sample stream is delivered to 10 identical code-synchronous detectors whose operation is further discussed below. Each multiplies the samples by a delayed version of the PRS, and then correlates the result against the f_c outputs of the carrier NCO. The timing of the PRS sequence within the 1023-clock period can be adjusted separately for each detector to a resolution of 100 ns, or about 1/10 of a clock period. Once signals from the antennas have been detected, two of these processors are used to track each antenna's signal by adjusting the PRS timing of one to be slightly earlier and the other slightly later than that which produces peak correlation amplitude.

Hardware configuration. The hardware arrangement is shown in Figure 26. The detector (HP33330B) is a zero-bias Schottky diode with flat response from 0.1 to 18 GHz. The pre-amp/filter uses an operational amplifier to provide an appropriate load impedance to the detector and differential drive to the ADC; it is constructed in a small aluminum box that plugs into the main board (see Figure 6). The FPGA (Xilinx XC3S400 5PQ208), clock oscillator, ADC, and most other components are on the main printed circuit board. The microprocessor is part of a commercial single-board computer (Ampro model CM2-420) on a PC104 board that connects as a daughter board to the main board and communicates with the FPGA over an ISA bus. Also packaged in the same box are a switching power supply (115 VAC to 5 VDC) and a small fan. The box is weatherproof and the main board temperature is kept between 30 deg C and 50 deg C by software control of the fan and some resistive heaters. Communication with the control computer is via 100 Mbps Ethernet. (The 256 Mb SDRAM in Figure 26 was included for possible future expansion and is not used in the current design.)

Algorithms. The receiver operates as a slave to the control computer by responding to commands that it receives. Many commands are supported by the software, including the retrieval of status and diagnostic information, file-handling utilities, and detailed control

²⁴ With signals from multiple antennas, squaring does not remove the modulation completely unless the modulation clocks of all antennas are aligned at the receiver, but it does produce a component at twice the carrier frequency. Besides, for the small demonstration array, the clocks are usually aligned within 100 ns.

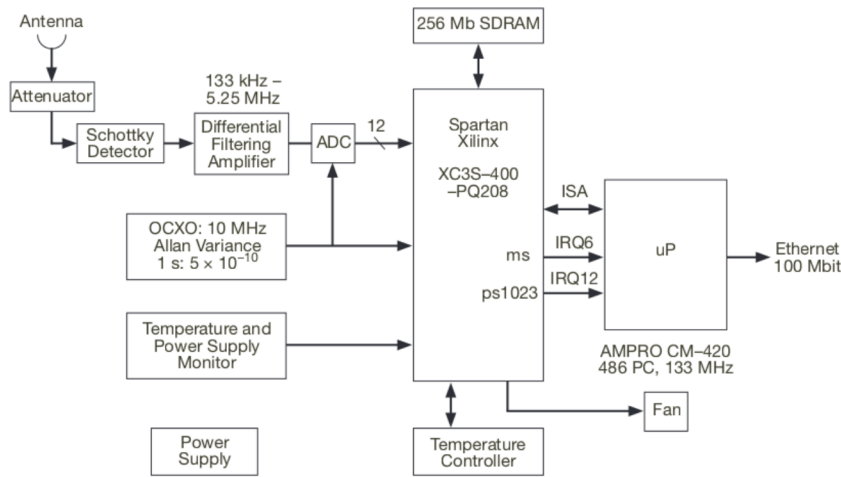


Figure 26. Hardware configuration of a calibration receiver. The antenna, attenuator, and detector are the only microwave components. The microprocessor is a PC104 single-board computer. The remaining components are on a custom PC board. All components except the antenna are mounted in a 17.3 × 11.9 × 10.7 cm box, along with a 5-V power supply and fan.

of the signal processing to facilitate testing. Only the most important operating commands will be discussed here. The main algorithms are illustrated by the flow chart in Figure 27.

There are two main modes of operation, IDLE and TRACK, controlled by a logical variable *track_flag* that can be set or cleared by commands. When an “acquire” command is issued, the software leaves the IDLE mode and proceeds through a sequence of steps to detect signals from as many of the five antennas as possible, and if any are detected it enters a continuous loop to track them. The software is organized into threads, which run as separate tasks under a real-time operating system. Two of the threads, called “command” and “code tracking,” are shown in Figure 27. The first is driven by the receipt of a command from the control computer, where communication is implemented as a Telnet server over TCP/IP. The second is driven by interrupts from the FPGA, which occur every cycle of the PRS code (every 1023 cycles of 1.024 MHz, or 0.9990234 ms). The procedure for signal acquisition (response to an acquire command) can be traced in Figure 27 as follows.

- Within the command thread, a block of captured samples is retrieved from the FPGA RAM and the total power is calculated. If the power is below a preset threshold (set about 10 dB lower than the smallest expected signal), the receiver remains in IDLE mode.
- Otherwise, the software repeatedly retrieves the accumulator contents of the carrier detector and uses them to calculate the frequency difference between the carrier NCO and the received carrier. It then resets the carrier NCO to make the difference zero, and sets *track_flag* to true.

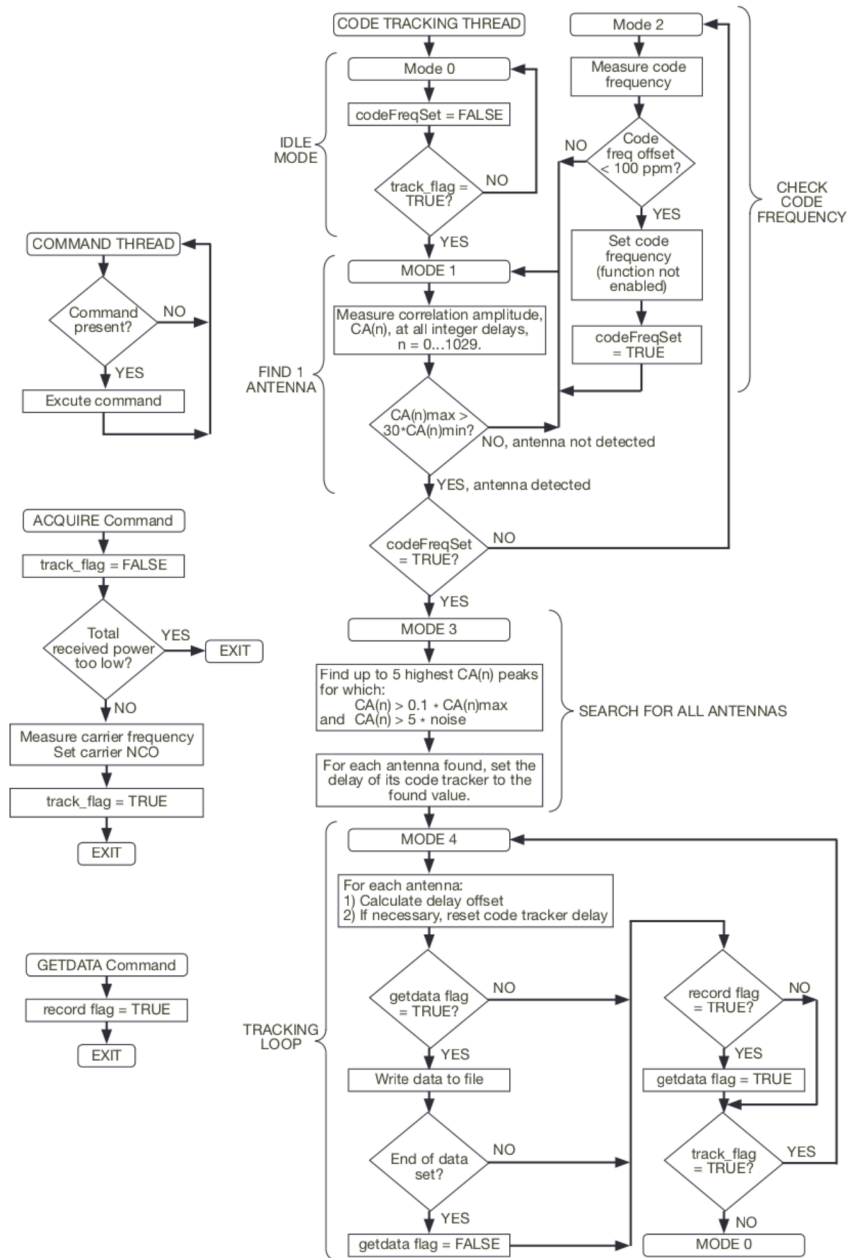


Figure 27. Flow chart for main algorithms of calibration receiver software. There are two main threads: command processing, driven by receipt of commands from the array control computer; and code tracking, driven by an interrupt from the hardware every code cycle. Many commands are supported, but only two are shown here. The ACQUIRE command triggers a search for signals from up to five antennas. If one or more antennas are detected, the signals are then continuously tracked. During such tracking, the GETDATA command can trigger the recording of a specified set of samples to a file.

- Next, within the code tracking thread, the software uses the code-synchronous detectors to measure the cross-correlation function of the received signal with the receiver's PRS (Mode 1 in Figure 27). This is done over the full range of time offsets from 0 to 1022 clock periods, in steps of 1 clock period. Since there are 10 detectors, they are configured to make 10 measurements at different offsets simultaneously, thus requiring 103 cycles of the PRS to step through all offsets. (The last set includes offsets 1023 through 1029 clocks, which are the same as 0 through 6.) For each offset n , the amplitude $C_a(n) = \sqrt{I^2 + Q^2}$ is calculated, where I and Q are the contents of the cosine and sine accumulators, respectively. It then looks for the peak value $\max(C_a)$, which should occur when the offset produces near alignment of the local PRS with that of an antenna. If $\max(C_a)/\min(C_a) > 30$, an antenna is tentatively detected. Otherwise, the cross-correlation measurement is repeated.
- When one antenna has been tentatively detected, an attempt is made to track its code timing (Mode 2). One of the code-synchronous detectors is set to a delay -0.5 clock from the estimated peak, and another is set to $+0.5$ clock from the same point. The ratio of amplitudes from these two detectors provides a measure of the tracking error, i.e., the difference in delay between the actual peak and the estimate. Over the next 10 code periods, this "fine delay" offset is repeatedly measured. Its rate of change determines the frequency error of the clock NCO relative to the actual clock frequency of the received signal. If the clock frequency error is too large (>100 ppm), then this is considered a false antenna detection and we return to Mode 1 to repeat the cross-correlation measurement. (The design includes a provision for readjusting the clock NCO to remove the frequency error, but this was found unnecessary since the oscillator frequencies are very stable.)
- If the clock frequency is sufficiently accurate, the antenna detection is confirmed. We nevertheless return to Mode 1 to make a new cross-correlation measurement, but when it is finished we proceed to Mode 3. Here the software searches for all peaks in $C_a(n)$ that meet certain criteria, in an attempt to find signals from all antennas. Up to 10 values that exceed $\max(C_a)/10$ and also exceed 5 times the noise (represented by the mean of the remaining $C_a(n)$ measurements) are examined, working from strongest to weakest. If two of the 10 are at adjacent values of $n \bmod 1023$, they are assumed to bracket the timing offset of one antenna. Otherwise, the stronger of the two neighbors of a peak is used as the bracketing value. This process finishes when all five antennas have been detected or no more peaks above the thresholds remain.
- For each detected antenna, a pair of detectors is set to -0.5 clock and $+0.5$ clock from the estimated peak delay, and the software enters the tracking loop (Mode 4 in Figure 27). Thereafter, following each code cycle interrupt, the accumulator contents for all detected antennas are retrieved and the fine delay offsets are calculated as before. The detector delays are then adjusted so as to keep the delay error as close as possible to zero. Recall that the delay setting resolution is approximately 0.1 clock. A small difference between the frequencies of the array clock and the clock NCO will eventually cause the delay adjustment to require changing both the coarse (integer clock) and fine settings, and may even cause the coarse setting to wrap around from 1022 to 0 or 0 to 1022. All such cases are handled by the software, allowing continuous tracking of unlimited duration as long as the signal-to-noise ratio remains sufficiently large.

The tracking loop attempts to continue even if all signals disappear (e.g., if the array is pointed away from this receiver). It exits and returns to the IDLE mode when a new acquire command is issued or when certain other commands that set *track_flag* to false are issued.

If the receiver has acquired and is tracking at least one antenna, it will accept a “getdata $k L$ ” command (see Figure 27). This tells the receiver to record the I and Q accumulator contents and the current code timing for all tracked antennas every k code cycles for a total of L measurements. For example, with $k = 10$ and $L = 3000$, the recording contains approximately 30 s of measurements at 10-ms intervals. These are the parameters used for most of the calibration measurements reported in Section IV. The recorded data are written to a set of files on a software-simulated disk drive in the microprocessor’s memory. When the requested measurements are completed, the data can be retrieved by commands over the Telnet interface, or the files can be transferred by FTP (since the embedded software also implements an FTP server).

The antennas are distinguished from each other because their code sequences start at substantially different times, so the receiver sees them as separate peaks in the cross-correlation function. The receiver does not know the relative times, so it attempts to find the five strongest peaks regardless of their delays. Determining which peak corresponds to which antenna is left to off-line analysis software, but to make this possible the timing offsets must be carefully chosen. For all calibration measurements, we set the antenna PRSs to begin at a synchronization tick from bit $S_i = 0, 21, 63, 147,$ or 252 for antenna $i = 1, 2, 3, 4,$ or 5 , respectively. In this way, the time differences are always fairly large (at least 21 clocks) and unique for every pair of antennas. The measured delays are then $(S_i + t_r) \bmod 1023$, where t_r depends on the receiver’s clock. From the measured values, an algorithm in the off-line software can unambiguously associate a specific antenna with each signal, even if one or more antennas are not detected.

E. Real-Time Software

The Uplink Array Control program runs in the array control computer and is used for real-time operation of the array. It allows centralized control and monitoring of the various subsystems. It has been tailored specifically to the requirements of this demonstration and is not intended to meet the demands of an operational system. As such, many of the commands are terse and error handling is minimal. The user is assumed to have high-level knowledge of the how the uplink array system operates, what its limitations are, and how to go about diagnosing problems.

The program supports five antennas and four calibration receivers, but its design allows these numbers to be easily changed, and it smoothly handles cases where one or more elements of the system are absent or not functioning. Nevertheless, the design is limited to relatively small numbers; it would not be easily expandable to an array of 100 antennas.

Development environment. The program was developed under Microsoft Visual Basic 6.0 (known as VB6). There were no requirements that forced this choice. Any number of other

integrated development environments that provide a GUI for the operator would have sufficed. The programmer in this case had prior familiarity with VB6 and was aware of some difficulties other developers in the programming community were having in dealing with network socket connections under the newer Visual Studio 2006.

There was also no requirement that the operating system be Microsoft Windows, but some of the data-collection instrumentation only came with drivers that worked under Microsoft products such as VB6, Visual Studio 2006, or DOS. The configuration program used by the motor servo manufacturer operates only under Microsoft Windows. Hence, the lack of drivers and support software made it difficult to consider other operating systems such as Linux.

Timing. The use of a non-real-time operating system for essentially real-time functions is possible here because the real-time requirements are fairly weak. Most time intervals, such as polling rates for status information, are long (seconds), and most event-driven processing, such as command executions, can be delayed without serious consequences. Periodic logging and the “PAUSE n ” command in scripts are intended to provide accurate timing, but resolution is coarse (1 s) and occasional errors are acceptable.

Some commands would have tighter real-time requirements if we were tracking rapidly moving objects, such as low-Earth-orbit spacecraft. It would then be necessary to ensure that updates to the antenna pointing, carrier phases, and modulation timing be transmitted within a few seconds of the nominal times. In a larger array, the timing accuracy requirement would become proportionally tighter.

There is one desirable feature, not implemented for the demonstration array, that produces a hard and tight real-time requirement. Continuous transmission of arbitrary data (e.g., from a file) is supported by the hardware provided that a buffer in the digital module at each antenna is updated with the next data block in a timely manner. The hardware is double-buffered, so the software has available all of the transmission time of one buffer to fill the other buffer. That time varies from 16.384 s at 4 kbps to 64 ms at 1024 kbps. Transmission of a new buffer always starts on the rising edge of the 500-Hz timing reference (2-ms tick), so the design assumes that this signal provides an interrupt to the control computer. Even at low data rates, it would not be possible under Windows to guarantee that such interrupts are always serviced and that data transmission is always completed when required.

There is one time-critical command to the digital modules, called “trigger,” whose purpose is to synchronize the modulation timing among the antennas. If buffered data are being transmitted, then this command should be received by all antennas within the 2-ms interval preceding the timing reference tick, after which a new buffer should begin being transmitted. In the simpler case of transmitting pseudorandom sequences, the trigger command must still be received within the same 2-ms interval at all antennas, but it need not be a particular 2-ms interval. This is accomplished by sending the trigger command as a broadcast to all digital modules.

If true real-time control became necessary, our plan was to off-load communication with the digital modules from the main control computer to a slave single-board computer running a single task. This was not done for the demonstration.

An important design principle was that the control computer should be the master of all timing, so that none of the devices within the array needs to know the time of day. To the extent that any event must occur at a particular time or any record must be marked with a time stamp, this is the control computer's responsibility.

Software timing in our VB6/Windows system was handled by several background timer tasks. There is no guaranteed system timer available to provide predictive and preemptive time scheduling and prioritization of tasks, so these timers run asynchronously and autonomously from other event driven threads such as those initiated by a mouse click or input from the keyboard. To avoid deadlock or throughput issues between tasks, all activities that depend on a single resource (such as the availability of serial or network data) are entirely contained in their own timer task. The following timer tasks are started during program initialization:

- Timer1: handles all serial traffic from the RS485 network that communicates to the antenna motion servos and updates the corresponding status displays (period = 50 ms).
- Timer2: handles all incoming UDP traffic from the digital module boards and updates the corresponding status displays (20 ms).
- Timer5: handles all incoming TCP traffic from the calibration receiver subsystem and updates the corresponding status displays (100 ms).

Two other timers are used only when script file operation is invoked. These are:

- Timer3: reads the script file line by line and then hands off that line to the command processor (100 ms).
- Timer4: used by the script processor to support the "PAUSE n " command. When that command is invoked, Timer4 is programmed for n seconds duration and enabled while Timer3 is disabled. When n seconds have elapsed, Timer4 executes, re-enabling Timer3 and disabling itself.

Communication. The control program must communicate with at least 21 separate devices (five digital modules, 10 motor servos, four calibration receivers, and one or more GPIB devices) over four physical interfaces (two Ethernets, asynchronous serial RS485, and GPIB) using multiple protocols (binary UDP/IP for the digital modules, proprietary binary serial for the motor servos, Telnet over TCP/IP for the calibration receivers, and ASCII GPIB) within the demonstration array, as well as with the public Internet.

For the digital modules, UDP over a private Ethernet was chosen for the digital modules partly to simplify their embedded software, but mainly to ensure precise timing of broadcast commands. For the calibration receivers, Telnet via TCP over a separate Ethernet was

chosen because it was supported by preexisting software that could be reused in its embedded processor, and because no time-critical commands were contemplated.

Files. Three types of files are handled by the control program: initialization, logs, and scripts. All are ASCII text files.

An initialization file called "UplinkAry.INI" is read during program startup. It contains configuration parameters that should be changeable without recompiling, such as IP addresses, antenna-specific pointing parameters, and default settings for the digital modules.

Log files are written when enabled by the operator. They record all commands executed and their responses, whether the commands are entered via the command line, a button push, or a script. They also periodically record monitor data from the digital modules and the reference distributor. Each log record contains a date/time stamp.

Scripts. A script file is read upon entry of "SCR RUN <filename>" on the command line. Each line in the file is a command to either a digital module, an antenna motion servo pair, a calibration receiver, a GPIB device, or the script processor. Another "SCR RUN <filename>" command can occur within a script file, so file execution can be nested. Normally, the commands are executed sequentially every 100 ms (using timer3). Operation can be paced to match external delays (like antenna motion) by including "SCR PAUSE *n*" commands, or made to wait for operator intervention by including "SCR PAUSE" commands. If an error is encountered, the default action is to ignore it and continue, but the operator can select one of two other options: pause, where an error causes execution to halt until the operator types "SCR RESUME"; or stop, where an error causes all script processing to end. "SCR PAUSE" can also be executed from the command line while a script is running. "SCR RESUME" and the initial "SCR RUN <filename>" command must always be entered via the command line.

The scripting facility is thus very simple. It includes no conditional or looping commands, nor any clock-paced commands. It nevertheless proved very powerful, and allowed us to automate much of the demonstration array's operation. We eventually created about 280 script files, not including those that were automatically generated by off-line software as part of the calibration process.

F. Analysis Software

All of the calibration measurements and other test data have been processed and analyzed using MATLAB, with some further analysis and plotting performed in Excel spreadsheets. A total of 136 MATLAB functions and scripts was written (not counting those used to assist in system design), including functions for reading the data files recorded by the calibration receivers, identifying the antennas, computing the carrier phase differences at the receiver, reducing those phase differences to those at the antennas via the known geometry, and computing the phase adjustments needed to produce alignment at a satellite or other tar-

get. (These calculations were described in Section IV.A.) Some special routines were written to analyze the survey data (Section IV.B), to analyze system tests like antenna pointing, and to calculate bit error rates from the demodulated signals (Section IV.F). Utilities to plot data from the log files and from the satellite ephemeris files [5] were also written.

VI. The Next Step: A Prototype Array for the DSN

Whereas the techniques needed to implement a large uplink array have now been demonstrated, it is appropriate to consider the next step toward creation of an operational uplink array for the DSN [13].

As explained in Section I, the long-term needs of the DSN for uplink generation would be well served by a large array of relatively small antennas, especially if EIRPs exceeding the present capability are needed [14]. For now, however, we consider only a limited-size prototype, constructed in a way that allows it to be substituted for one of the large single-dish systems currently used in deep space operations. The prototype array should have these characteristics:

- Size of each element (antenna effective area and power amplifier output) is selected to be nearly optimum for use in a large array capable of transmitting an EIRP of 1 TW or more in the 7.19-GHz space research uplink band.²⁵
- Number of elements is selected to produce EIRP sufficient to transmit useful data to a spacecraft in deep space, say, at least 4 kbps to Mars. With typical current spacecraft electronics, this requires about 16 GW.
- Antennas, signal processing electronics, control software, and support systems are designed for high stability and inexpensive mass production.
- Calibration system is arranged to support an eventual large array.
- Site is selected to allow expansion of capacity by simply adding more elements of the same size and design.

All of these are achieved by an array of 16 antennas, each having a 6-m-diameter reflector and a 470-W power amplifier. Such an array can be inexpensively deployed at the Goldstone Deep Space Communications Complex. The signal processing electronics and ground-based calibration system would use the low-cost architecture developed in current demonstration. The prototype can be built so that all components have the same design as a larger, fully operational array. It would then meet these objectives:

- Validate the design and methodology by using the new array to support those operating spacecraft for which it satisfies mission requirements.
- Demonstrate operational feasibility and validate operational concepts.

²⁵ A compact array for the 34-GHz band would also be possible if the troposphere is sufficiently stable, but whereas statistics of troposphere turbulence are still being studied, we suggest building the prototype for the 7.19-GHz band.

- Validate the low-cost assumption.
- Prove that an array is flexible, has graceful degradation, and is expandable.

Figure 28 shows one concept for the geometry of a large array of which the prototype is a subset. This configuration places the antennas in concentric rings, and the prototype antennas occupy 16 of the 39 positions of the innermost ring. A calibration receiver is placed on a 28-m-high tower at the center. This puts it at an elevation of 5 deg for antennas in the outermost ring — high enough to avoid ground reflections. Once the prototype is successfully tested, more antennas can be added according to requirements and budget. The first step is to fill the 39 positions of the inner circle; this produces 95 GW of EIRP (37 antennas would produce 85.5 GW, about the same as a 34-m antenna with 20-kW transmitter). The next three steps involve adding more rings of 39 antennas each, leading to EIRPs of 379 GW (more than a 70-m antenna with 20 kW), 855 GW, and 1519 GW, respectively. Each ring is rotated slightly from the previous one so as to provide a clear path from every antenna to the calibration receiver. The rings are separated by 34 m so as to guarantee that no shadowing occurs when tracking down to 10-deg elevation. Expansion beyond four rings is not practical for this geometry because the path to the calibration receiver could not be kept clear for all antennas.




Figure 28. Concept for the geometry of an expandable array. A 28-m-high calibration tower (black triangle) is at the center. The innermost ring has radius 217 m, and the 16 blue circles on that ring represent the prototype array. The red circles represent additional antennas that could be added later. The additional rings are spaced 34 m apart, and each ring has places for up to 39 antennas, for a maximum of 156 antennas.

Acknowledgments

This work was funded by NASA through the DSN Array Project (FY2006 and part of 2008) and the DSN Technology Program at JPL (FY2007–2009). Support and encouragement from JPL managers Faramaz Davarian, Andre Jongeling, and Joseph Statman, and from NASA Program Manager Barry Geldzahler, were much appreciated. ITT Radar Systems – Gilfillan was especially accommodating in supporting our installation and operations at their Loop Canyon test facility. Wayne Rothermich and Steve Smith of MMW Technology, Inc., did the detailed design and construction of the RF Module. Peter Kinman of California State University at Fresno reviewed much of the analysis software and many interim reports, as well as this report, providing valuable suggestions. Many people at JPL made important contributions at particular times. Summer student Christopher Gonzales designed the logic of the digital modules. Timothy Sink was responsible for the laser-based surveys. Leslie White worked on the power amplifiers, packaging, and parts lists. Norman Lay and Thomas Jedrey loaned us the equipment that made the data transfer test (Section IV.F) possible.

References

- [1] S. von Hoerner, "Design of Large Steerable Antennas," *Astronomical J.*, vol. 72, pp. 35–47, 1967.
- [2] D. S. Bagri, "Phasing an Uplink Array Using Radio Sources." *The Interplanetary Network Progress Report*, vol. 42-172, Jet Propulsion Laboratory, Pasadena, California, pp. 1–6, February 15, 2008. http://ipnpr.jpl.nasa.gov/progress_report/42-172/172C.pdf
- [3] V. Vilnrotter, R. Mukai, and D. Lee, "Uplink Array Calibration via Far-Field Power Maximization," *The Interplanetary Network Progress Report*, vol. 42-164, Jet Propulsion Laboratory, Pasadena, California, pp. 1–16, February 15, 2006. http://ipnpr.jpl.nasa.gov/progress_report/42-164/164D.pdf
- [4] L. Primas, G. Lutes, and R. Sydnor, "Fiber Optic Frequency Transfer Link," *Proc. of the 42nd Annual IEEE Frequency Control Symposium*, Baltimore, Maryland, pp. 478–484, June 1988.
- [5] Intelsat Corp. ephemeris data. <http://www.intelsat.com/resources/satellitedata/ephemeris.asp>
- [6] R. R. Bate, D. D. Mueller, and J. E. White, *Fundamentals of Astrodynamics*, p. 84, Dover Publications, 1971.
- [7] T. S. Kelso, "Orbital Coordinate Systems, Part II." <http://celestrak.com/columns/v02n02>
- [8] E. K. Smith, Jr. and S. Weintraub, "The Constants in the Equation for Atmospheric Refractive Index at Radio Frequencies." *Proc. IRE*, vol. 41, pp. 1035–1037, 1953. Cited and explained in A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*, chapter 13, Wiley, New York, 1986 (and later editions).

- [9] National Weather Service Forecast Office, Los Angeles; CAMP9 automated weather station, identifier CNIC1.
<http://www.wrh.noaa.gov/mesowest/getobext.php?wfo=lox&sid=CNIC1&num=48>
- [10] J. W. Goodman, *Introduction to Fourier Optics*, McGraw-Hill, 1968.
- [11] V. Jamnejad, "A Study of Near to Far Fields of JPL Deep Space Network (DSN) Antennas for RFI Analysis," *2004 IEEE Aerospace Conference*, paper #1546, vol. 2, pp. 975–986, March 2004.
- [12] L. D'Addario, "Estimates of Atmosphere-Induced Gain Loss for the Deep Space Network Array," *The Interplanetary Network Progress Report*, vol. 42-160, Jet Propulsion Laboratory, Pasadena, California, pp. 1–7, February 15, 2005.
http://ipnpr.jpl.nasa.gov/progress_report/42-160/160E.pdf
- [13] F. Davarian, "Uplink Arraying Next Steps." *The Interplanetary Network Progress Report*, vol. 42-175, Jet Propulsion Laboratory, Pasadena, California, pp. 1–18, November 15, 2008. http://ipnpr.jpl.nasa.gov/progress_report/42-175/175C.pdf
- [14] L. D'Addario, "Large Transmitting Arrays for Deep Space Uplinks, Solar System Radar, and Related Applications," URSI General Assembly, New Delhi, India, October 25, 2005. <http://hdl.handle.net/2014/37533>

Appendix A

Documentation Repository

I. Electronic Documentation

A collection of electronic documentation for this project has been assembled. It includes

- Internal reports generated during the project;
- Software source code for control computer, calibration receiver embedded processor, digital module embedded processor, and offline analysis (MATLAB);
- Configuration source code for programmable logic (FPGAs in calibration receiver and digital module, complex programmable logic devices [CPLDs] in digital module and clock generator);
- Schematics and PC board layout files (mostly Mentor Graphics files);
- Detailed parts lists (as an Excel spreadsheet); and
- Photographs.

Table A-1 lists the top part of the directory tree, including sizes in kB. The collection of files is available on a set of compact disks, copies of which may be requested²⁶ from one of the authors or from

Tracking and Applications Section (335), M/S 238-638
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

A list of selected internal reports is given below. All of these are contained in the "reportsAndPublications" directory. All of the MATLAB source code used for analyzing the measurements is contained in the "matlab" directory.

A separate collection of files contains data obtained during the tests, including both raw data and files processed in various ways. Most of these are ASCII text or Excel spreadsheets. A copy can be requested from the first author via e-mail at ldaddario@jpl.nasa.gov.

II. Selected Internal Reports

1. "New Operating Paradigms for the Deep Space Network," presentation to IND System Engineering Seminar, February 23, 2005. File: newDSNparadigms.ppt
2. "Uplink Array Calibration Receiver: Preliminary Design Document." DSN Array Tech Note No. 39, 2005 Dec 27. File: techNote39-uplinkCalibrationReceiver.pdf

²⁶ Requests from outside JPL are subject to management review. Some documents might not be released.

Table A-1. Documentation directory tree.

Size, kB	Name
14,292	./subsystems/antennaMechanical
171,292	./subsystems/calibrationReceiver
333,260	./subsystems/controlComputer
464	./subsystems/powerAmp
1,184,456	./subsystems/digitalBox
64	./subsystems/interconnections
121,348	./subsystems/referenceDistributionSubsystem
36,044	./subsystems/rfBox
1,861,224	./subsystems
62,768	./reportsAndPublications
612	./matlab
1,432	./procuments
2,104	./partsList
449,236	./controlComputerFilesBackup
43,572	./reviews
324,968	./pictures
37,116	./datasheets
1,528	./fieldInstallation
2,784,628	

3. "Nearly-Orthogonal Sequences for Uplink Array Calibration," DSN Array Tech Note No. 40, December 27, 2005.
File: techNote40-orthogonalSequences.pdf
4. "A New Method of Transmission Line Length Stabilization for Reference Tone Distribution," draft report, November 27, 2006.
File: roundTripCorrectionReport.pdf
5. "Analysis of the Loop Canyon Survey Results," ver. 2.0, June 16, 2008.
File: surveyAnalysisR2.pdf
6. "Further Analysis of the Loop Canyon Survey Results," June 18, 2008.
File: surveyAnalysisNew.pdf
7. "Surveys of Uplink array at Loop Canyon," November 25, 2008.
File: surveySummaries.pdf
8. "Phase Tracking of Satellites with Uplink Array Demonstration," June 18, 2008.
File: satellitePhaseTracking.pdf
9. "Tests with Galaxy 16 Satellite during May 2008," May 19, 2008.
File: g16testsMay2008.pdf
10. "Uplink Array Demonstration: Tests With Galaxy 26 Satellite during June 2008," June 13, 2008. File: g26testsJun2008.pdf
11. "Uplink Array Demonstration: Satellite Tests, Part III." July 3, 2008.
File: satelliteTestsJuly3.pdf
12. "Satellite Tests, Part IV." September 3, 2008. File: satelliteTestsPartIV.pdf
13. "Uplink Array Demonstration: Calibration Stability and Consistency," August 3, 2008.
File: calibrationStability.pdf

Appendix B

Analysis of Costs

For a few specific items, the prices paid to vendors have been given in the main text. The direct cost of all materials needed to construct the array, based on the project's parts lists and actual prices paid to vendors, was US\$86,763. It is possible that the actual cost was higher by at most a few percent because of omitted and unpriced items. Other items were obtained at no cost because they were available as surplus from earlier projects in our lab, or because they were borrowed and later returned. It is estimated that the total direct cost of materials would have been about US\$95,000 if everything were purchased. This does not include the costs of shipping, acceptance testing, extra parts procured as spares, or procurement overhead.

Based on JPL's institutional accounting procedures, the total cost of this effort was US\$1997k, with the following breakdown:

Direct labor (including benefits), 7.1 person-years	1014.8 k\$
JPL services (computers, network, instrument rental, etc.)	69.8
Purchase orders	166.9
Contracts	37.9
Travel	5.9
Other direct costs	7.5
JPL indirect costs	696.0

When the project was proposed in September 2005, the requested budget was US\$1958k and it was expected to be completed in 18 months. The actual duration was 35 months, partly because some tasks proved more difficult than anticipated (see Section IV.J) and partly because of the funding profile.