IPN Progress Report 42-176 • February 15, 2009

RAIK: Redundant Array of Inexpensive Klystrons

Bruce L. Conroy* and David L. Losh*

Methods of combining output power from multiple 20-kW transmitters were examined as an alternative to developing a new 80-kW transmitter for the Deep Space Network (DSN) based on a single klystron. A promising technique using two hybrids and a waveguide phase shifter was developed and successfully tested at a power level of 300 W as a proof of concept. Implementation strategies including interface to the microwave subsystem's feed, cooling, physical arrangement of components in a beam-waveguide (BWG) antenna, and transmitter reliability were examined. The study concluded that by using five 20-kW transmitters, a capability to provide power levels from 1 to 80 kW could be developed for roughly the same cost as a single-tube 80-kW transmitter and that the multitransmitter approach would be much more reliable and probably less expensive.

I. Introduction

Deep-space missions beyond Mars depend largely upon the 20-kW transmitter on the 70-m antennas. As these antennas appear to be nearing their end of life, the Deep Space Network (DSN) has explored alternatives such as increasing the power from a 34-m antenna to compensate for the reduced gain of the smaller antenna. During a cost exercise for the development of an 80-kW transmitter for emergency commanding, the concept of combining the outputs from the 20-kW power amplifiers (PAs) — a proven technology — rather than developing a new tube and higher power filters was suggested to the authors by Mike Britcliffe and Dan Hoppe. It was also noted that having one or more extra 20-kW PAs would result in higher reliability at a modest cost.

This article describes several concepts for combining the output of 20-kW 7145- to 7190-MHz (X-band) klystron power amplifiers into a single 34-m antenna to achieve two goals:

1. To increase power for emergency commanding without the cost and risk associated with development of higher power tubes and filters.
2. To increase reliability by having redundant systems.

---

* 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.

The following section describes the various power-combining alternatives considered, the chosen approach, and some preliminary test results from low- and medium-power testing. Section III discusses implementation issues such as interface to the X-band (uplink)/X-band (downlink)/Ka-band (downlink) feed — referred to as X/X/Ka feed in this article — techniques for cooling the multitransmitter system, physical arrangements, and reliability. Section IV presents conclusions and recommendations for future development.

II. Methods Considered

In the course of the study, several methods of combining the outputs from multiple transmitters were considered. The following sections describe the methods and their strengths and weaknesses.

A. Hybrid or Magic Tee Combiner

A 3-dB hybrid or magic tee accepts two inputs. If they are of equal amplitude and properly phased, the output will their sum. To the extent that there is amplitude unbalance or improper phase, a fraction of the input power is delivered to a fourth port, typically called a "waster." The amplitude balance requirement forces the number of inputs to be a power of two. We can combine two inputs or two pairs of inputs. One way to provide flexibility is to run the inputs through a matrix of transfer switches, as shown below.

The circuit shown in Figure 1 contains nine transfer switches (a well-developed technology, which should operate easily at 100 kW) and three hybrid combiners. The hybrid combiners could be scaled from the design Dan Hoppe provided for the Goldstone Solar System Radar (GSSR), which has been operated at nearly 500 kW at a higher frequency. There are five inputs, labeled A, B, C, D, and E, and a single output, "Out."

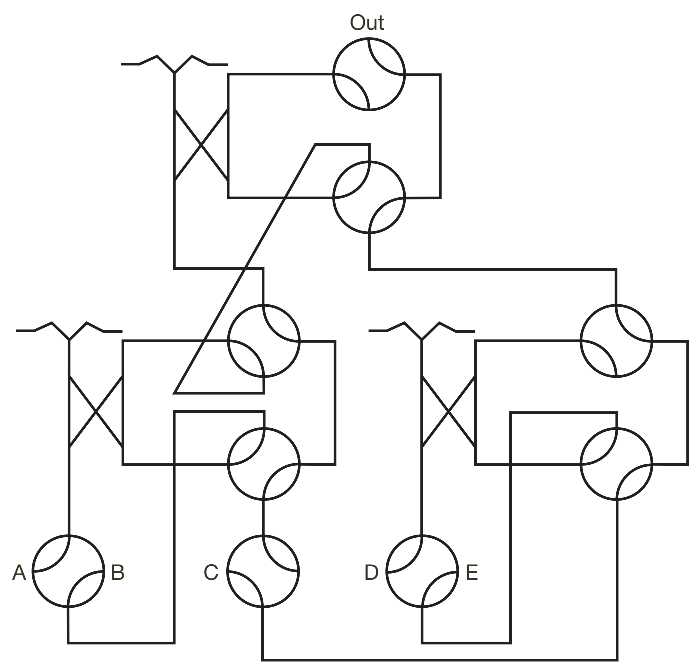

Figure 1. Flexible switch and hybrid combiner for 20 kW, 40 kW, or 80 kW.

Proper positioning of the switches allows any single input to be sent to the output to provide 20-kW operation with four redundant inputs available. The circuit also allows the sum of any two of the five to be selected to provide 40-kW operation with three standby units for redundancy, or the sum of any four of them for 80-kW operation with one redundant unit at any time.

40-kW Option

If a lower-power solution is needed, redundancy could be provided by a combiner that would allow the selection of any single transmitter to provide 20 kW or any pair of transmitters to provide 40 kW from a total of three power amplifiers. The capability can be provided with one hybrid and three transfer switches, as shown in Figure 2.

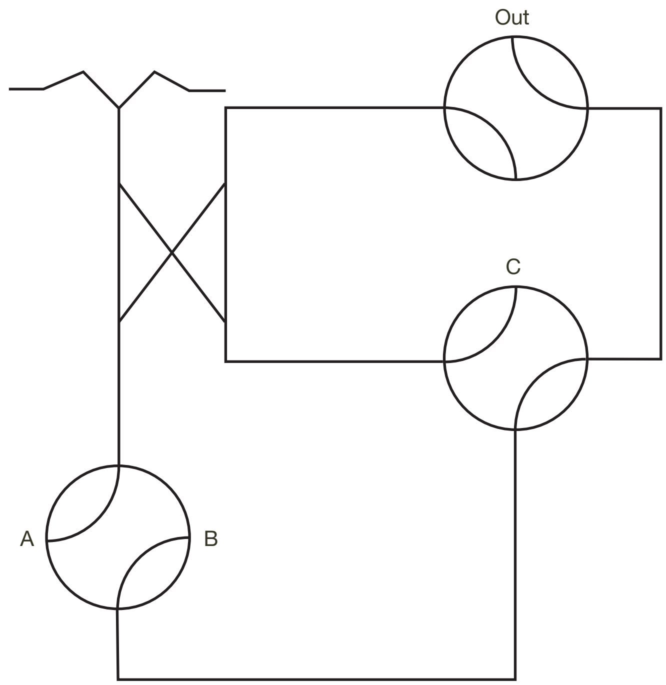

Figure 2. Simplified switch and hybrid combiner for 40 kW.

B. Common Junction

A combining method that avoids the requirement that the number of inputs be a power of two is to match the inputs to a common junction, often a cavity. A combiner of this type can be designed for any number of inputs.

The operation of this type of combiner can be complex. Consider a simple three-way combiner consisting of three vertically stacked waveguides, each of which tapers down to one-third height, and flows into a full height guide that forms the output. See Figure 3.

It is clear that as long as the three inputs are in phase and of equal amplitude, they will combine at the common junction to form a single wave towards the output port and not cause reflection. The question of what happens if one of the inputs is out of phase or of different amplitude is more complex.

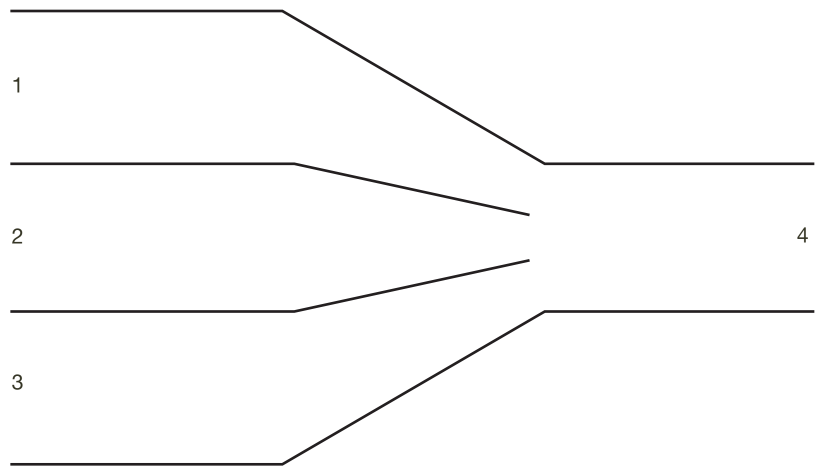

Figure 3. A three-way power combiner.

To answer these questions, a Waveguide Simulation Program for Integrated Circuits Emphasis (WGSPICE) model was created and calculated the S matrix for this junction, as shown in Table 1.

Table 1. Scattering matrix for a three-way device.

Excitation	Output	Coefficients
Node: 1 mode: TE(1,0)	Node: 1 mode: TE(1,0)	S = |0.659235| 19.9442
Node: 1 mode: TE(1,0)	Node: 2 mode: TE(1,0)	S = |0.324943| –174.313
Node: 1 mode: TE(1,0)	Node: 3 mode: TE(1,0)	S = |0.356908| –138.621
Node: 1 mode: TE(1,0)	Node: 4 mode: TE(1,0)	S = |0.576574| 179.197
Node: 2 mode: TE(1,0)	Node: 1 mode: TE(1,0)	S = |0.324943| –174.313
Node: 2 mode: TE(1,0)	Node: 2 mode: TE(1,0)	S = |0.675564| 1.79337
Node: 2 mode: TE(1,0)	Node: 3 mode: TE(1,0)	S = |0.324943| –174.313
Node: 2 mode: TE(1,0)	Node: 4 mode: TE(1,0)	S = |0.576574| 179.197
Node: 3 mode: TE(1,0)	Node: 1 mode: TE(1,0)	S = |0.356908| –138.621
Node: 3 mode: TE(1,0)	Node: 2 mode: TE(1,0)	S = |0.324943| –174.313
Node: 3 mode: TE(1,0)	Node: 3 mode: TE(1,0)	S = |0.659235| 19.9442
Node: 3 mode: TE(1,0)	Node: 4 mode: TE(1,0)	S = |0.576574| 179.197
Node: 4 mode: TE(1,0)	Node: 1 mode: TE(1,0)	S = |0.576574| 179.197
Node: 4 mode: TE(1,0)	Node: 2 mode: TE(1,0)	S = |0.576574| 179.197
Node: 4 mode: TE(1,0)	Node: 3 mode: TE(1,0)	S = |0.576574| 179.197
Node: 4 mode: TE(1,0)	Node: 4 mode: TE(1,0)	S = |0.0518326| –125.023

As expected, s_{44} is near zero, indicating that signals entering that port are well matched, but s_{11}, s_{22}, and s_{33} are huge, so that when only a single port is excited, 43 percent of the power is reflected. The factor that makes it work is that when all three ports are excited in phase, the leakage from Port 2 to Port 1 plus the leakage from Port 3 to Port 1 is just sufficient to cancel the reflection at Port 1. In other words, $s_{11}+s_{21}+s_{31} = 0$.

Because there is no interport isolation, external isolators must be provided for each port. A standard 20-kW DSN transmitter is normally provided with an isolator; however, the lack of any way to compensate for amplitude unbalance or the loss of an input creates a significant problem. When one of the three inputs is lost, the output drops from 3P0 to 1.33P0 instead of 2P0 (combining efficiency drops to 66 percent.)

To compensate for loss of one of the transmitters, Figure 4 shows how some provision can be made for redundant power amplifiers by using transfer switches.

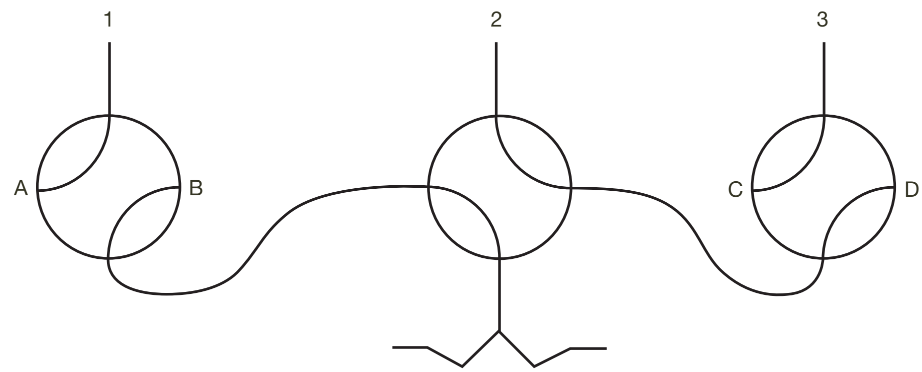

Figure 4. Redundancy using transfer switches.

In this configuration, any three of the four inputs (A, B, C, D) can be routed to the three outputs (1, 2, 3), thus giving a measure of redundancy.

C. Quasi-Optical Combining

The basis for many quasi-optical power combiners is a wire grid that will reflect an incoming wave that has the E field parallel to the wires and pass a wave where the E field is perpendicular to the wires.

By placing two sources of equal amplitude and properly phased at right angles to each other with a wire grid at the intersection, the output of the grid will contain the sum of the two input powers in a linearly polarized wave at 45 deg to either of the incident waves.

This concept still works when the two amplitudes are not equal, but the polarization angle of the resultant wave is arctan{E1/E2} instead of 45 deg. If there were an element that would accept a linearly polarized wave at an arbitrary angle and rotate it to a standard angle, it would allow combining an arbitrary number of sources, each with arbitrary power.

The Vane Polarizer

Another standard quasi-optical component is the vane, or Venetian blind, polarizer. It consists of several thin metal strips placed the path of the propagating field. When the E field is perpendicular to the strips, no current flows, so the wave continues in the transverse electromagnetic (TEM) mode. When the E field is parallel to the strips, the wave is forced into the TE1 parallel plate mode, which has a longer wavelength than the TEM mode. The thickness of the device is chosen so that the TEM mode is delayed by 90 deg more than the TE1 mode. When the device is illuminated by a wave polarized at 45 deg to the vanes, it splits the wave into two equal components and delays one of them by 90 deg. When the components recombine at the output, the result is circularly polarized. Conversely, when the vane polarizer is illuminated by a circularly polarized wave, the output is linearly polarized at 45 deg.

The vane polarizer is the basis of the polarization rotator needed for the arbitrary combiner. If a vane polarizer is placed in the path of the output of the wire grid combiner and rotated so the vanes are 45 deg to the wave, the result will be circularly polarized. When this circularly polarized wave passes through a second vane polarizer, the result will be linearly polarized at the fixed angle of 45 deg to the vanes, and can be used as the input to another stage of power combining. The implementation of one such stage is shown in Figure 5.

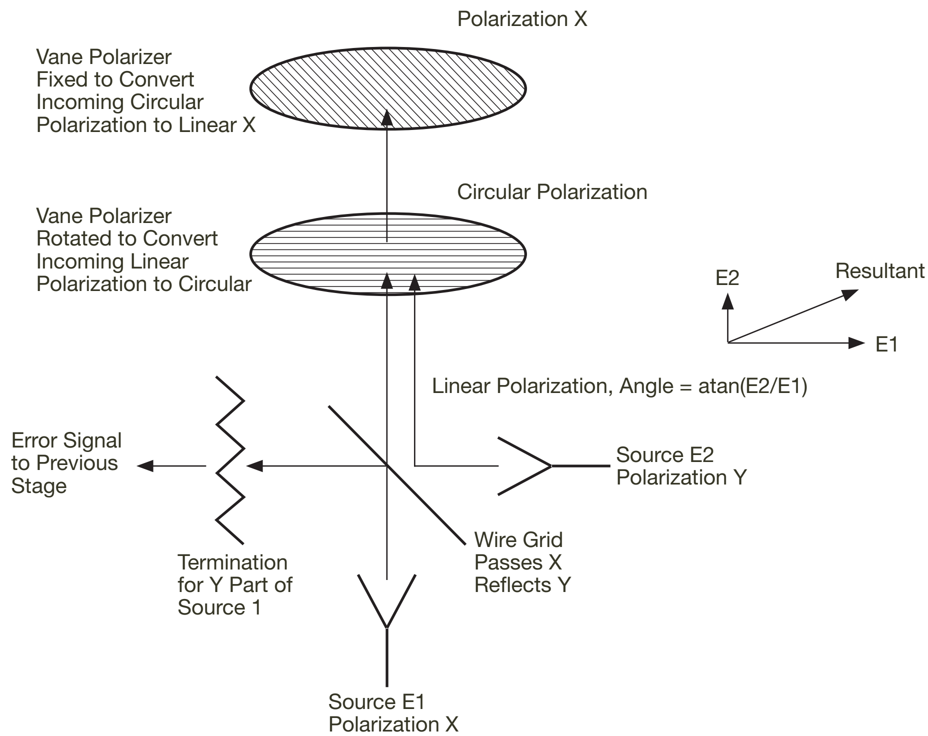

Figure 5. Possible implementation of one stage of an arbitrary power combiner.

Termination and Feedback

If the wire grid is not properly aligned with the polarization of the signal from source 1, a part of the signal will reflect off the wire grid, so a proper termination must be provided for the reflected component. The presence of the reflected component can be used as a feedback signal to correct the positioning of the wire grid.

Phase Shift

Although quasi-optical combining allows signals of arbitrary amplitude to be combined, the entire signal must be properly phased; thus, it is important that the phase shift of the polarization rotator be constant and independent of the ratio of the input powers. Unfortunately, this is not the case. Although the net effect of each polarizer is a phase shift of exactly 45 deg over the free space phase, the phase shift from the propagation in a circular mode is the length of the section (wavelengths times 2π plus the angle between the two polarizers. Thus, the rotation of the first polarizer introduces a phase into the output signal.

This is not a serious defect for two reasons. First, the phase shift is predictable from the polarizer angle, and second, the induced phase shift can easily be corrected in the next stage. By adding a phase shifter into one leg (e.g., at the input to the transmitter) and adjusting both this phase shifter and the polarizer angle to minimize the power in the waster load, the combiner becomes independent of the input phase in the other leg.

Each stage of power combining generally needs a phase shifter for fine-tuning, and considering it to be a fundamental part of the combiner provides a way of adjusting it. The only complexity at a system level is the need to coordinate the setting of the phase shifter and polarizer rotation since they are both using the same observable.

D. Waveguide Combiner Based on Orthomode Junction

Quasi-optical combining uses free-space propagation of the signals and thus requires significant physical space to reduce loss; however, there is a waveguide analog of the quasi-optical form presented above. To introduce it two other components must be considered.

Orthomode Junction

This device is an analog of the wire grid. It has three ports. Ports 1 and 2 are rectangular guides supporting a dominant mode of TE01 and Port 3 is a square guide with TE01 and TE10 as the principle modes. Ports 1 and 2 are decoupled, Port 1 couples to TE01 at Port 3 and Port 2 couples to TE10 at Port 3.

Quarterwave Polarizer

This polarizer is a waveguide analog of the vane polarizer. Typically it consists of a length of circular guide with flattened sections on opposite walls. The concept is the same as in the vane case. A wave with the E field perpendicular to the flats propagates in a different mode from one with the E field parallel to the flats. The name comes from the fact that the length of the device is chosen to have 90-deg phase shift between these two modes.

Implementation

The waveguide form of an arbitrary power combiner is shown in Figure 6. One obvious difference from the quasi-optical approach is that this form shows a second orthomode junction used in reverse to provide the termination and feedback function rather than using the wire grid from the following stage. The second orthomode junction is needed because it is a three-port device while the wire grid is natively a four-port device.

E. Waveguide Combiner Based on Hybrids and Phase Shifters

Another implementation of a combiner for signals of unequal amplitude in waveguide uses the circuit shown in Figure 7 that consists of two hybrids and two phase shifters.

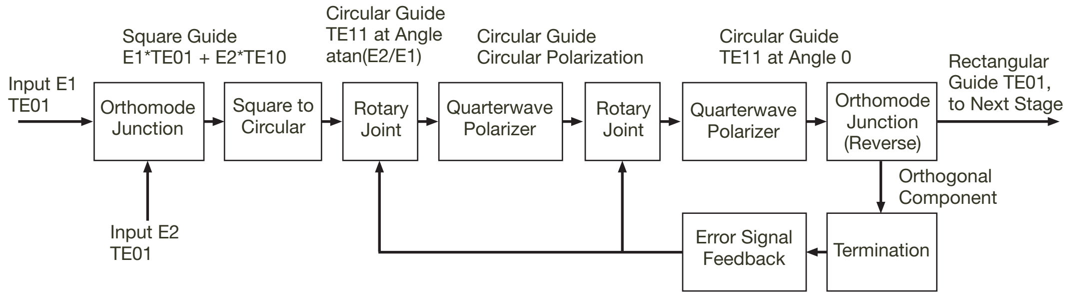

Figure 6. Waveguide implementation of arbitrary amplitude power combiner.

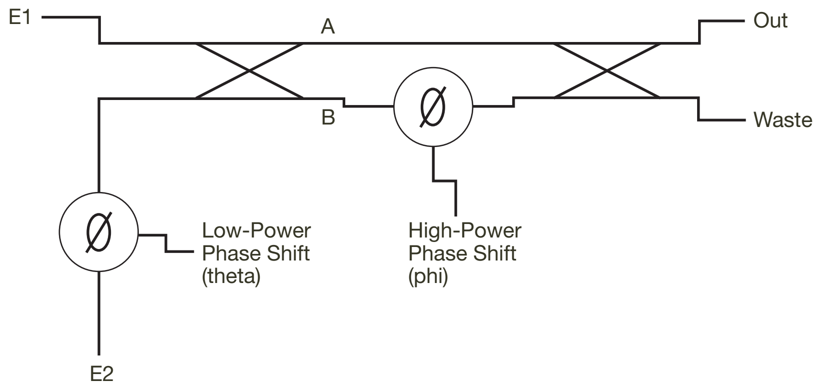

Figure 7. Another combiner architecture.

The setting of the theta phase shifter will control the ratio of the signals at A and B. No matter what the relative amplitudes of the signals at E1 and E2 are, there is always a setting for Theta such that the amplitudes at A and B are equal. It can be considered as being in quadrature, although the actual value depends on the hybrid design. Once there are equal amplitudes at A and B, there is always a setting for phi such that all the power goes to the out port, and none goes to the waste port.

This circuit is in fact a replacement for the circuit with the rotating polarizer. The orthomode junctions are replaced by 3-dB hybrids, and a single high-power phase shifter replaces the two polarizers and rotary joints. It has the following advantages:

1. Lower parts count than using orthomode junctions.
2. The hybrid is better understood than the orthomode because we have built X-band hybrids up to 500 kW.
3. The final phase shifter only sees half of the total output power.
4. It may be possible to build a ferrite phase shifter without moving parts.
5. We have some experience with high-power phase shifters from the resonant ring experiments, although that was only up to 20 kW and they did have some arcing problems.

F. Ladder Combiner

Given a method that allows combining power from two sources and does not depend on equal amplitudes, it becomes practical to add multiple power sources sequentially in a ladder. The waveguide form of the arbitrary power combiner shown in Figure 8 illustrates the process.

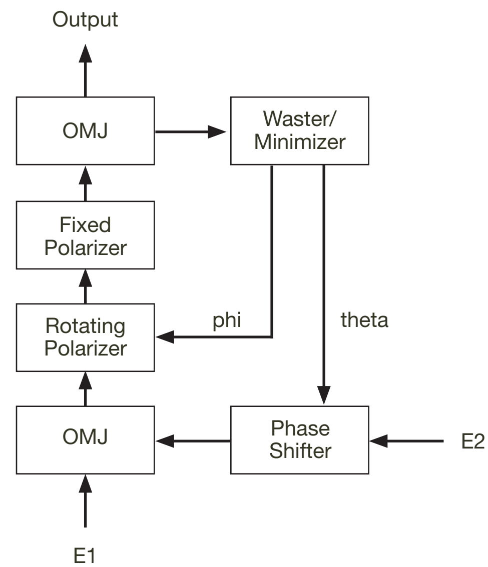

Figure 8. Simplified schematic of a waveguide power combiner.

There are two orthomode junctions, two polarizers (one rotatable), a phase shifter, and a waster load, plus some logic circuits that adjust the angle of the rotating polarizer (phi) and the angle of the phase shifter (theta) to minimize the power in the waster load. This logic (which will be covered in detail later) enables the combiner to adapt itself to the two inputs (E1 and E2), whatever the incoming amplitudes and phases.

To form the emergency command transmitter, we simply stack four of these combiners as in a ladder and move the phase shifters for each stage to the input side of the power amplifiers, as shown in Figure 9.

The ladder combiner architecture has all the advantages of the complex transfer switch matrix presented in Section III.A, plus several more:

1. Operation at 20 kW with a single power amplifier involves activating one of the five and leaving the others deactivated. The combiners will adjust themselves to produce a 20-kW output. Similarly, combiners will self-adjust for operation of any number of transmitters activated producing output power from a minimum of 1 W to 100 kW; so the capabilities far exceed those provided by the switch matrix.
2. During the transient period while the combiners are adjusting, all power is safely terminated. There is no possibility of operating a klystron into an unterminated waveguide.

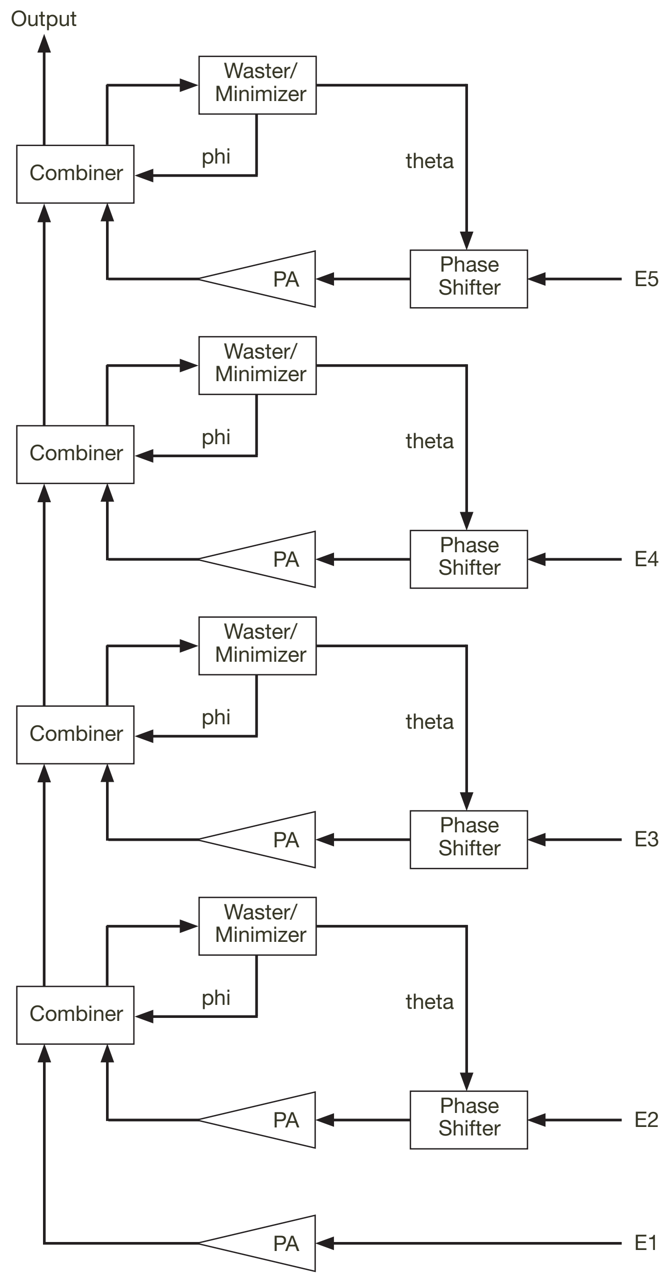

Figure 9. Architecture of emergency command transmitter.

3. Phase matching (necessary for any power combining scheme) is now included and automatic.
4. Group delay matching, a potential nightmare for the switch matrix, is easy. Assume that the electrical length of a combiner section is L. To equalize the group delay, simply place a coax with length L at the E3 input, one of 2L at the E4 input, and one of 3L at the E5 input. The group delays will be equalized for all possible combinations of inputs.
5. The architecture includes the equivalent of a transfer switch at each stage. When a combiner stage is tuned up so that all power is going to the output, a change of 90 deg in phi will switch all of the power into the waster load. This can be extremely fast, and it is safe to switch under full power if the waster load is adequate.

The dynamics of the combiner and the tuning algorithm were analyzed using a suite of MATLAB scripts. For the tuning algorithm, each iteration consists of:

1. Start with combiner 1, make small (0.02 rad) step in phi.
2. If the waster power increases, undo the step, and make a small step in the other direction.
3. Keep making steps in the same direction as long as the waster power keeps decreasing.
4. Make a small step in theta.
5. If the waster power increases, undo the step, and make a small step in the other direction.
6. Keep making steps in the same direction as long as the waster power keeps decreasing.
7. Do the same at combiner 2.
8. Do the same at combiner 3.
9. Do the same at combiner 4.

The results seem quite impressive. There doesn't seem to be a set of initial conditions that takes more than two iterations.

Figure 10 plots the output power and waster power from the four combiners for the above algorithm. For purposes of the simulation, the two combiner options are similar. Each has two inputs, phi and theta. In both cases, theta is a low-power phase shifter at the input of a PA, and phi is a mechanical device run by a motor. The differences in the MATLAB model are small, and both produced identical simulation results with the same tuning algorithm. In the simulation, all the initial angles are 45 deg, and each transmitter is at 1 W. After the second iteration, every stage is tuned correctly.

Changes to the phase shifter, theta, can be made at electronic speed, but changes in the polarizer rotation or high-power phase shifter, phi, will be much slower because it in-

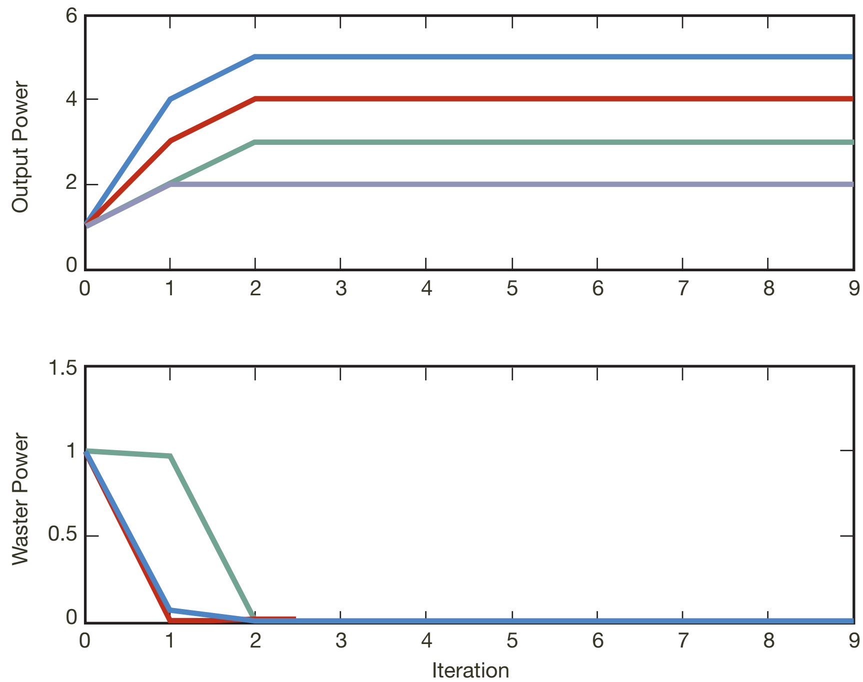

Figure 10. Transient at initial turn-on of five transmitters.

volves turning on a motor to physically move something. An initial guess would be several seconds per iteration. It will be important to watch the time required for digitization of the waster power. Each iteration may require several step/read/decide what to do stages.

Figure 11 shows output power when one of the five transmitters fails. In this case, transmitter 1 fails at time zero. As expected, the initial output of the first combiner drops from 2 W

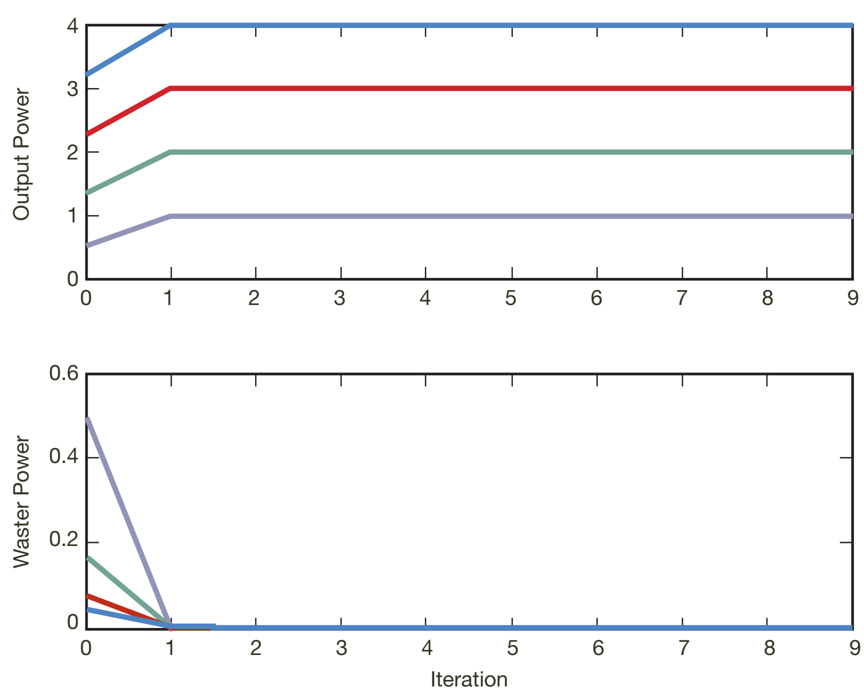

Figure 11. Transient at failure of one transmitter.

to 0.5 W (6 dB), and overall output drops from 5 W to 3.2 W (1.94 dB.) The system recovers, however, and the full 4 W available is all going to the output after just one iteration.

G. Combiner Development

To further verify that the ladder combiner can be implemented, a model was made with commercial off-the-shelf (COTS) parts, as shown in Figure 12.

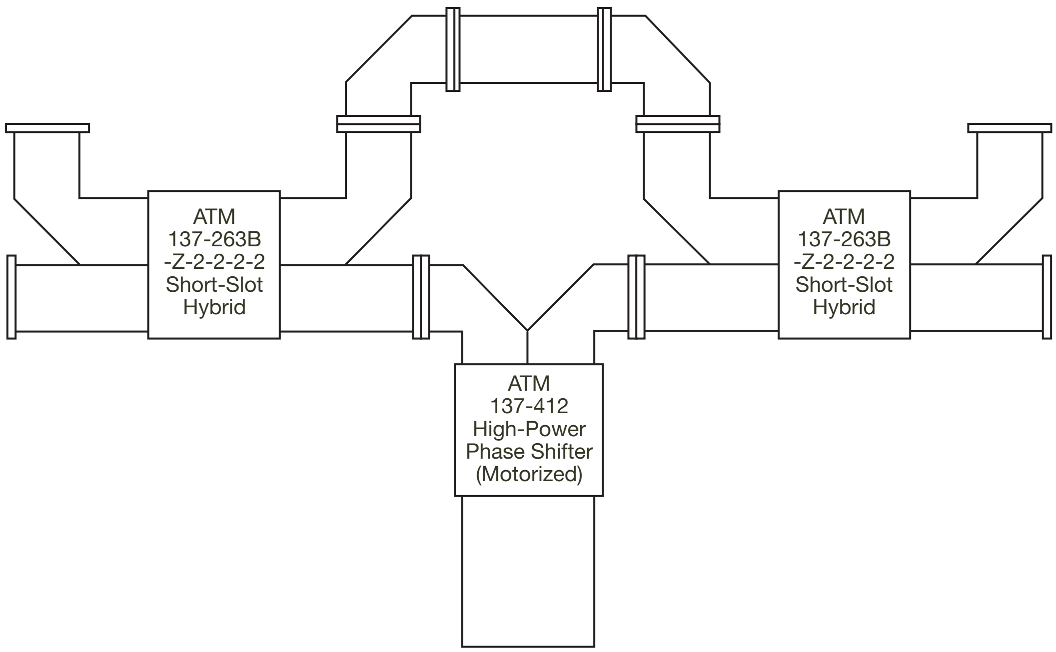

Figure 12. Basic waveguide combiner section.

The limiting component in the model is the motorized phase shifter, which has an average power rating of 3.5 kW. Its peak power is rated at 515 kW, so breakdown is not expected to be a problem at any power level needed for a 100-kW transmitter. The chief objective of the test was to develop and test algorithms to implement the combiner controller. The experiment setup is shown in Figure 13.

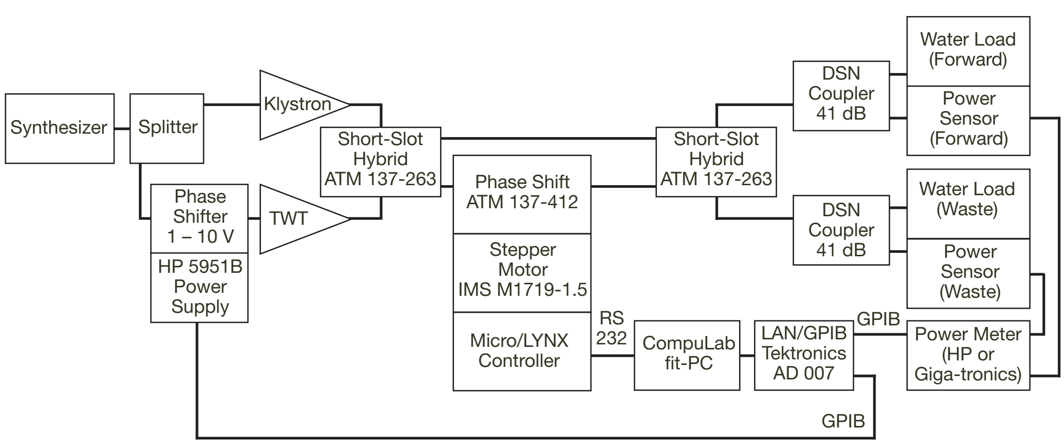

Figure 13. Block diagram of experimental system.

Hardware Setup

Power from a synthesizer is split into two paths. One path goes directly from the splitter to a standard DSN 20-kW klystron power amplifier, in this case the prototype for the 70-m antennas, and the other goes through a low-power phase shifter to 350-W traveling-wave tube (TWT) amplifier from Amplifier Research.

The low-power phase shifter is model PQ60-LC from KDI/Triangle. This unit is linearized for operation from 1 to 10 VDC, and provides almost 400 deg of shift over the DSN uplink band (7145 to 7190 MHz.) It is driven by an HP 59501B programmable power supply, which can be set in 0.01 V increments through general-purpose instrument bus (GPIB) control, giving a total of 1000 steps of phase shift.

The directional couplers are standard forward/reverse power couplers used by the 20-kW power amplifiers. In the experiment, the directional couplers are arranged to use the reverse power arm, which has a coupling of about 41 dB.

The high-power phase shifter is a model 137-412 from Advanced Technical Materials (ATM) and consists of a short-slot hybrid with choke-jointed sliding shorts on two of the arms. The input is split into the two shorted arms. The two shorts are moved by a common block (so they stay in sync) and when the two reflected waves pass back through the hybrid, the phase is such that they combine at the output port but have a phase shift proportional to the distance to the sliding shorts. The full travel provides almost 400 deg in the DSN band.

The block moving the shorts is driven by a screw with about 19 turns from stop to stop, and has backlash of about 1/8 of a turn. There are no limit switches at the end of travel, only hard stops. The screw, in turn, is driven by an M1719-1.5 stepper motor from Intelligent Motion Systems (IMS). The shaft passes through the motor to a knob to allow manual setting when the stepper motor is unpowered.

The stepper motor is controlled by a MicroLYNX controller, also provided by IMS. This is a very versatile controller that accepts high-level commands over an RS232 link (among others).

Commands are available to:

1. Define user units rather than having to work in absolute step numbers. The stop-to-stop range was chosen to be 1015 user units: 1000 units for the useful range so each step was 0.1 percent of the travel, with 15 extra units for the backlash.
2. Set the current position to zero.
3. Move forward or back a commanded number of units.
4. Move to an absolute position in user units.
5. Set rates for travel and acceleration.
6. Slew forward or backward at a specific rate.

As long as the controller is powered on, it remembers the current position of the stepper motor. There is also a problem that any time the acceleration rate is too high, the motor may fall behind, and the controller doesn't know it. This involved some experimentation to find what rates were acceptable.

The lack of limit switches meant the only way to get initial position calibration was to slew the motor in one direction, at a slow enough rate that hitting the stop wouldn't damage anything and for a long enough time to be sure it was at a stop, and call that zero.

At various phases of the experiment, the power measurement was done by:

1. An HP 438, where the time required to switch between channel A and channel B was too high.
2. Two HP 438s, so no switching was required.
3. A model 8542C from Giga-tronics, which did not have a delay switching between channels.

GPIB control for the power meters and the low-power phase shifter was provided by a Tektronix AD0007 GPIB/local area network (LAN) unit. The experiment controller was a Fit-PC, which is a 5-W, paperback-book-sized computer from CompuLab. The operating system is Xubuntu 6.06 from Canonical Software.

Control Software

The major components of the control system are patterned after a typical transmitter controller. They consist of:

1. DLL1, a shared memory technique that allows the modules to communicate.
2. RAIK_Mon, the main control program.
3. TXR_Data, a program that listens on the Internet, sends measured data from the Dynamically Linked Library (DLL), accepts and parses commands, and if the commands are valid and authenticated, passes them to the DLL for RAIK_Mon to execute. This is the same as used on the beam-waveguide (BWG) transmitters, but listens on port 9995.
4. A calibration file (r_cal.txt) that contains items for:

   a. Forward Power Coupling

   b. Reverse Power Coupling

   c. Phi 1 (the lowest setting of the high-power shifter required, in user units)

   d. Phi 2 (the highest setting of the high-power shifter required)

   e. Theta Factor (used to convert degrees into user units for the low-power phase shifter). There is a command to RAIK_Mon that causes it to go to a MySQL server and regenerate the cal file.

The control program has a "Standby" mode, where it monitors the DLL for an operate command, but does no further monitoring. This is useful for freeing the power meters from remote control for calibration and zeroing.

Upon receiving a command to go into the "Operate" mode, it reads the cal file and establishes communication with the power meters, power supply, and the MicroLynx stepper motor control.

Since the two output ports are the same at the hardware level, calling one the forward power and one the waste power is arbitrary. The control program has a mode to maximize the forward power (minimize waste) and one to maximize the waste power.

Tuning Techniques

Fine Tune. This technique is what was simulated in MATLAB. Using a step size of 0.1 percent of full scale, it first adjusts the high-power shifter (phi) then the low-power shifter (theta). For each iteration, it takes a step and sees if the power in the port to be minimized went up or down. If up, it undoes the step, then takes a step in the other direction. It continues in the same direction as long as the power keeps going down. When an increase is noted, it undoes the last step. The two drawbacks to this are 1) it is the slowest technique, and 2) because the step is so small, it often ends too early due to noise. It is most useful when the start point is close to the real optimum.

Medium tune. This mitigates some of the drawbacks of the fine-tune method. It first iterates through phi and theta as described above with a step size of 0.4 percent, then iterates through them again with a step size of 0.1 percent.

Coarse tune. This option adds a third iteration. The first time through, the step size is 1.6 percent, then 0.4 percent, and finally 0.1 percent. This technique is quite fast and robust. The only remaining problem was hitting a stop in the middle of a tune.

Reset coarse tune. This technique was introduced after some attempts to subtract 360 deg from the setting when hitting or nearing a stop (since both shifters have adequate range) but the step would have to be exact and that required calibration. It involves just setting each shifter to the middle of its range, then going into the coarse tune described above. It was always successful.

Other techniques. Although the theta value is unconstrained, the phi value depends only on the ratio of the two input powers. One technique investigated was to calibrate the high-power shifter to find the minimum and maximum values required for all power ratios and over the frequency band. A tuning request would then start at the minimum value and always search upward.

The result of calibration of the phi (high-power phase shifter) setting vs. frequency (7145 to 7190 MHz) on the RAIK for the cases (1) where all the power is coming in the klystron port, and (2) all the power is from the other port is shown in Table 2.

Table 2. Calibration of high-power phase shifter vs. frequency.

	User Units for All Power From:	
Frequency	Klystron Port	TWT Port
7145	293	730
7167	293	727
7190	293	725

Phi setting vs. input port and frequency. The actual setting for any distribution of input power will be between these limits, thus reducing the search area for tuning.

With power from the klystron, there was no change of setting vs. frequency; with the other port, the change was about 0.5 of the full scale. This means that the two paths are the same length in the first case, and almost the same in the second. The practical effect is that a search algorithm will not need to consider frequency.

This approach may have merit for further investigation, but there was no noticeable speed improvement, and the additional calibration adds complexity and possible failure modes.

Tuning Speed

The high-power phase shifter lacked both limit switches and an activity indicator. It also had a tendency to overshoot the requested position if allowed to go too fast. These deficiencies limited the speed of tuning. It was necessary to limit the slewing speed to avoid the overshoot, and wait after a move command to insure getting a valid stable reading.

The maximum speed used was about one quarter of that available. This was determined experimentally and may have been slower than necessary.

Another problem was the hysteresis. It was about 1 percent of the full travel, which is about 10 times the step size needed for tuning. This required setting the zero position away from the hard stop, so there would be room to go past it and always approach from the same direction, and always going past a down step, and coming back. This also limited the speed of tuning.

Another factor that limited the tuning speed was the decision to use a power meter to monitor the outputs' power rather than detectors and DC voltages. Although the DC readings could be faster, it seemed that actual power readings would be more valuable at this stage of development. The very first experiments used a Hewlett-Packard Model 438 dual-channel power meter, which took a full second to switch from one channel to the other. The next phase used separate 438s for the two channels, which was much faster. The final configuration used a Giga-tronics 8542C dual-channel meter, which was as fast as the two HP 438s.

Commands

Commands are passed through the DLL as a pair of integers. The first number is an index into the DLL and the value in that location is a parameter to the command. The commands defined are presented in Table 3.

Table 3. Combiner commands.

Number	Value	Description
7	1	Dump all values (see below)
9	1=Operate, 2=Standby	Set mode
10	1=Waste, 2=Forward	Set channel to maximize
19	Any	Read cal file
20	User units	Increase phi
21	User units	Decrease phi
22	User units	Set phi
23	User units	Tune phi only (see below)
24	User units	Increase theta
25	User units	Decrease theta
26	User units	Set theta
33	Any	Abort tuning
40	1=Fine tune; 2=Medium tune; 3=Coarse tune; 4=Reset and coarse tune	Tune
42	Any	Zero phi
44	User units	Tune theta only (see below)

Dump all values. Command 7 with a value of 1 will cause the program to send all DLL values, whether changed or not, on the next 1-s boundary. The command is used to re-synchronize a remote system's copy of the DLL if there is any doubt about the validity.

Tune phi or theta only. For these commands, the program adjusts only a single phase shifter, using the command value as a step size. It makes an initial step. If the result is in the desired direction, it keeps stepping in that direction until the outcome is not in the desired direction. If the initial step is worse, it tries the other direction. It always undoes the final step where the result got worse.

User Interface

The user interface is provided by a TCL script that connects to the controller through the internet. A screen capture of the user interface is shown in Figure 14.

The input fields are:

- Server: for the IP address or alias of the combiner
- Password: for the control password
- Step/Set: used in manual control of the phase shifters

The left column of indicators blinks green to show that the Ethernet connection is active and various tasks in the control program are running. The middle column has an indicator that shows yellow when the tuning algorithm is busy, and one that shows red if the password is rejected. The right column indicators show red if there are errors in the GPIB links to power sensors or power supply that controls the low-power phase shifter.

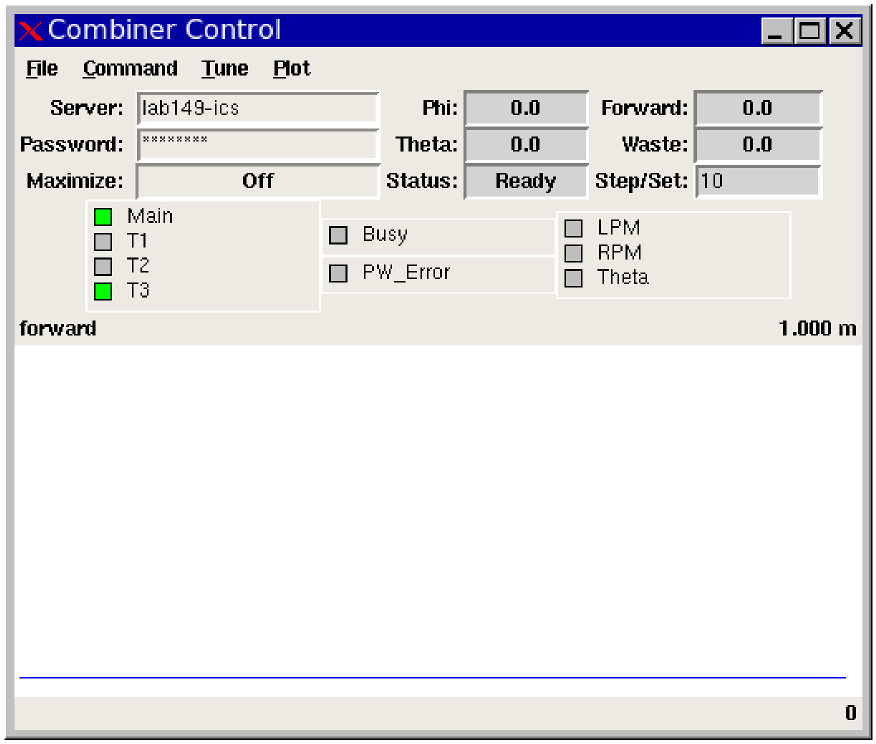

Figure 14. Screen capture of the interface program.

Control over the combiner and the plotting is through drop-down menus.

The "File" menu includes:

- Start — Opens the internet connection to the combiner, and starts filling a 1000-s memory buffer with data.
- Log start(stop) — Starts (or stops) recording all incoming data to a local file.
- saVe — Dumps the 1000-s buffer to a local file.
- Read cal — Sends a command to the combiner to refresh its calibration data from a MySQL server.
- eXit — Exits the interface program (it has no effect on the combiner).

The "Command" menu has

- Standby — Puts the combiner into a quiescent mode.
- Operate — Wakes up the combiner, establishes the GPIB and RS232 links, and starts reading the power meter channels.
- max Forward — Tells the combiner that forward power is desired.
- max Waste — Tells the combiner that maximum power at the waster port is desired.

- Zero phi — Starts the routine to establish a zero reference. Slews the phase shifter for enough time to hit the stop, takes out the backlash, and calls that zero.
- Phi -> Increase — Increases the phi setting by the number shown in the Step/Set window.
- Phi -> Decrease — Decreases the phi setting.
- Phi -> Set — Sets phi to the number in the Step/Set window.
- Theta -> Increase, Decrease, or Set — Same for the theta setting.

Under the "Tune" menu are

- Coarse
- Medium
- Fine
- Reset
- Phi
- Theta
- Stop

The "Plot" menu lets one choose which parameter (Phi, Theta, Forward, Waste, or Status) from the 1000-s buffer to plot.

Low-Power Testing

Most of the original testing and software development was done in a low-power configuration, without the power amplifiers and directional couplers, as shown in Figure 15. The gray waveguide components are the two short-slot hybrids and motorized phase shifter (at lower right, labeled "ATM") plus the copper waveguide sections that make up the basic combiner and correspond to the schematic in Figure 12. (Note that the boxes labeled "Altivar 16" are not part of the test setup.)

An initial test of the concept is to find the value of the low-power shifter (theta) that produces a null (in either channel), then run the high-power shifter (phi) through the full range. Figure 16 shows that we can get a deep null in either channel with the high-power shifter, and that the full range of the shifter is greater than 360 deg because there are two nulls in one channel.

The performance of the convergence algorithm is shown in Figure 17.

Medium-Power Testing

The test setup for medium-power testing using the 20-kW transmitter and 350-W TWTA is shown in Figure 18. In the figure, the input from the 20-kW transmitter is through the silver-colored waveguide on the right, while the input from the TWTA is through the aqua-colored coax. Connected to the output ports are two (black-colored) dual-directional couplers that are terminated in water loads, one of which can be observed at the upper left of the figure.

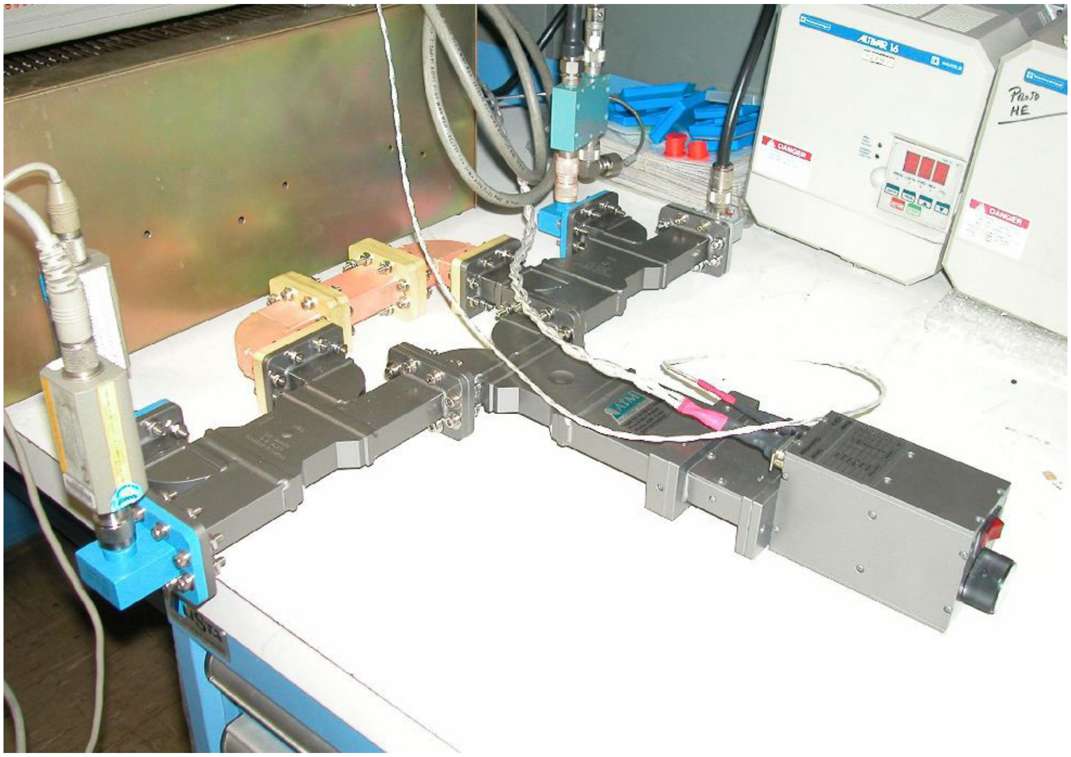

Figure 15. Photograph of the low-power test setup.

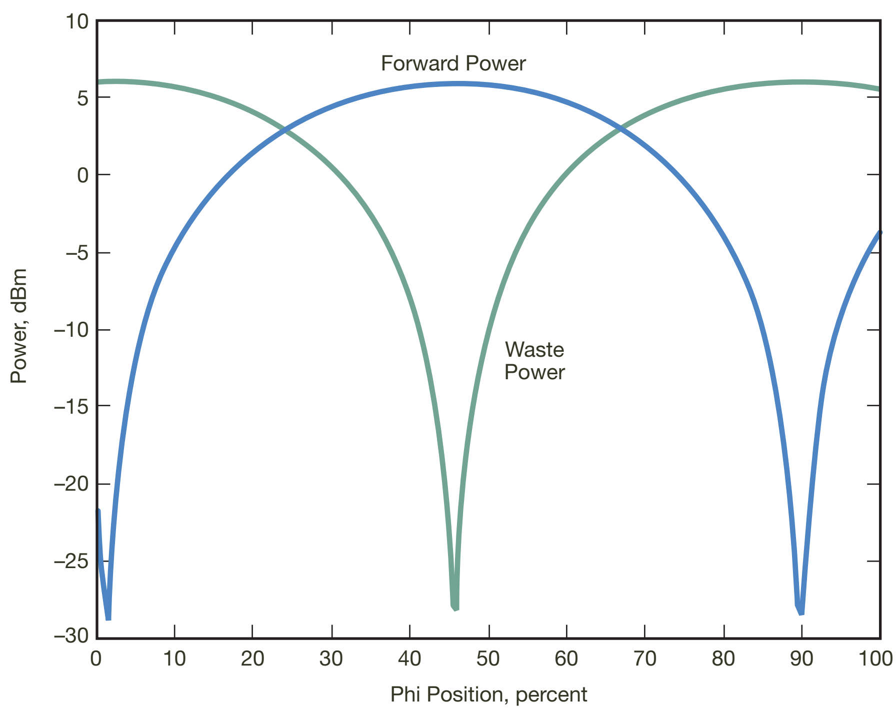

Figure 16. Forward and waster power vs. high-power phase shifter position.

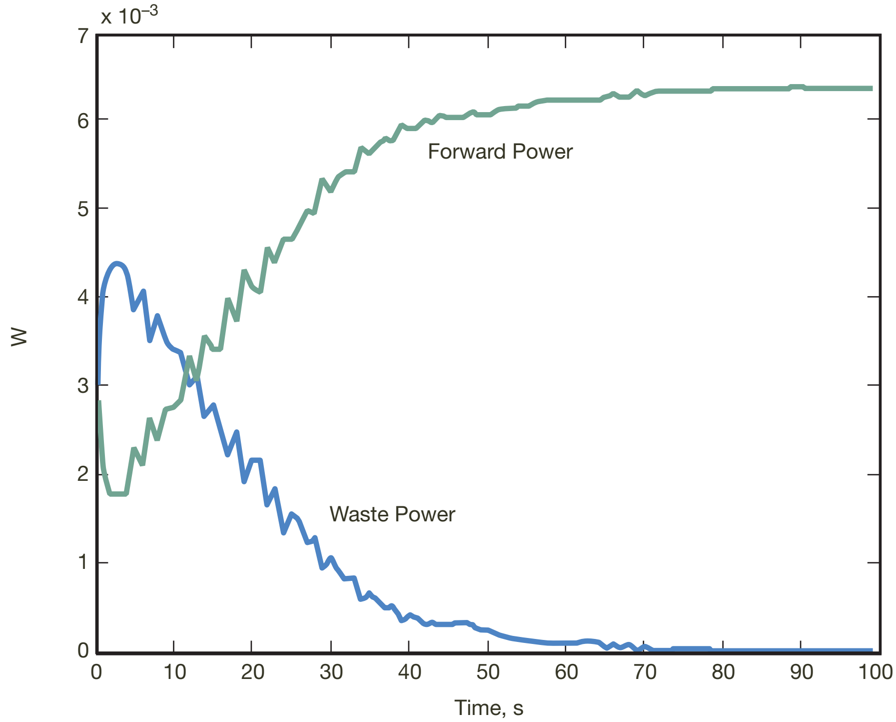

Figure 17. Forward and waster power vs. time during a "Tune" command.

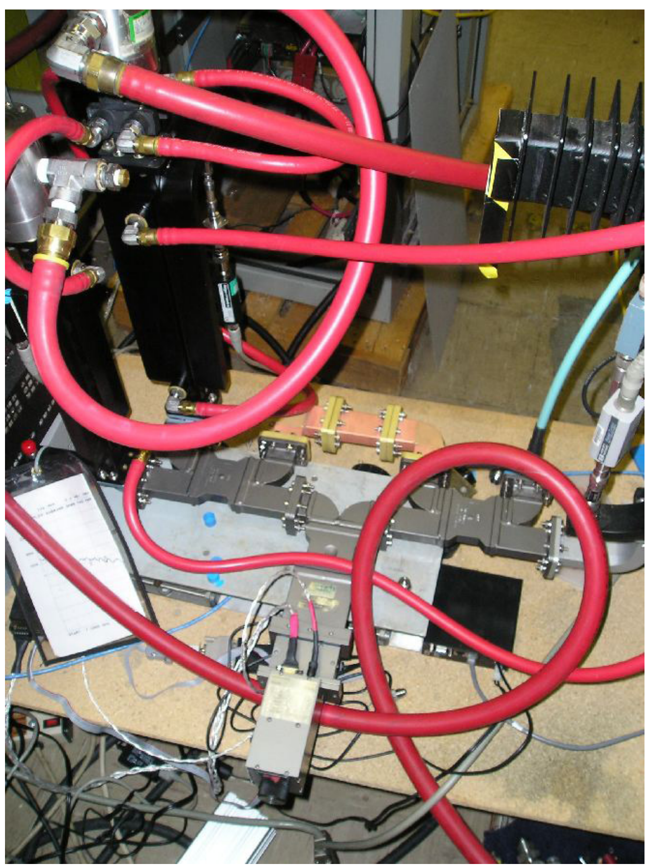

Figure 18. Photograph of medium-power test setup.

Figure 19 shows a plot of forward and waster power during a coarse tune.

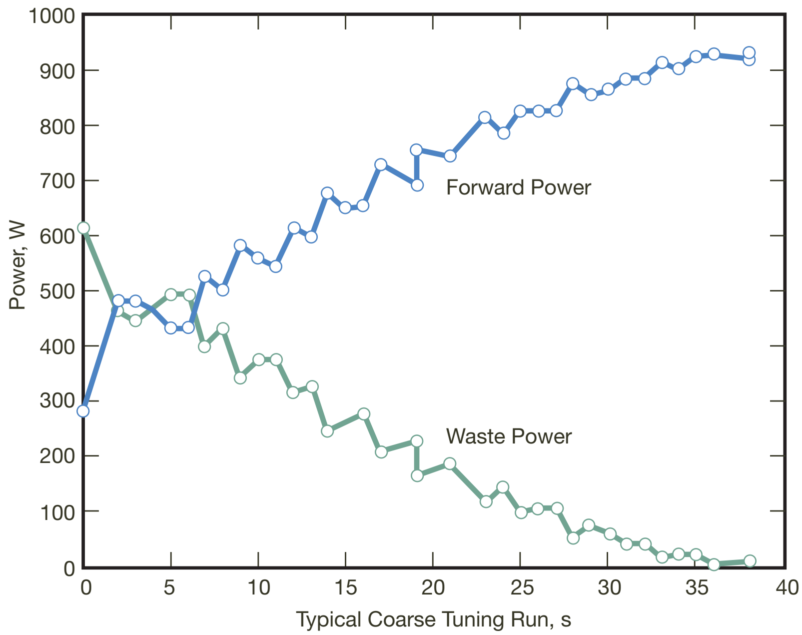

Figure 19. Medium-power convergence test results.

To help understand the long-term stability of the combiner, the setup was run for a period of 6 hr. At the beginning of the test, the klystron output was 159 W and the TWTA output was 250 W, for a total of 409 W input. Output from the combiner was 333 W, or a loss of about 0.9 dB. Figure 20 shows the plot of forward power and waste power vs. time.

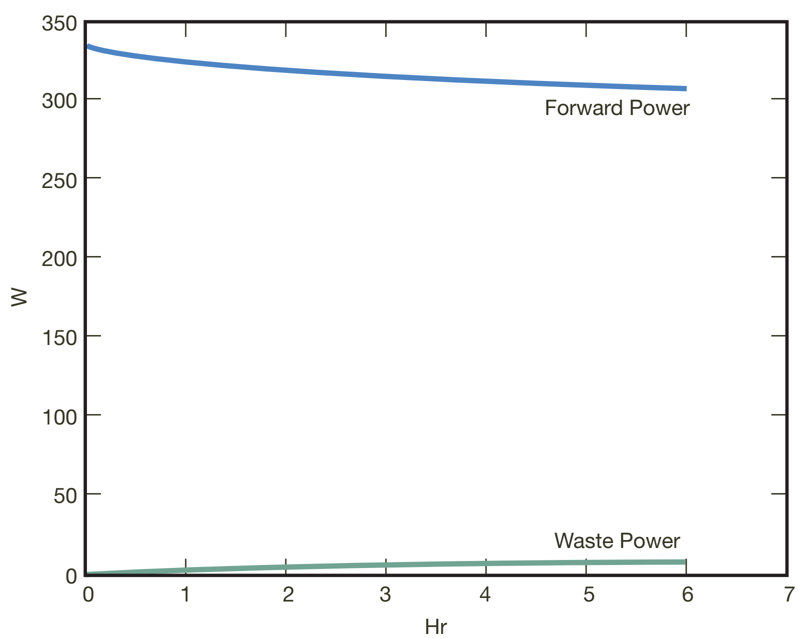

Figure 20. Stability test results.

At the end, the klystron was 151 W, and the TWT was at 232 W (383 W total). Output was 306 W, and the waster was 6 W. The unbalance in the inputs caused the rise in waste power, and the loss stayed constant.

Calculation from several of the tests gave a total insertion loss of a little under 1.0 dB. To get a better figure, which wasn't dependent on so many calibrations (i.e., directional couplers and power meters), we did a full measurement of the four port S parameters. Using this more accurate method, the calculated loss was 0.11 dB or about 2.5 percent.

Summary

The effort has developed a workable method for combining outputs from two transmitters having arbitrary phase and amplitude differences. The method provides a relatively compact package that could be physically integrated with multiple 20-kW transmitters to produce an 80- to 100-kW uplink. The actual RF testing showed that this form of combining works as predicted. The power level for actual testing was limited due to a lack of water cooling on the combiner parts.

III. Possible Implementation

Given that a feasible means of combining the outputs from two or more transmitters is available, the effort examined a variety of implementation issues, including the interface of the transmitter to the X/X/Ka feed used in the BWG stations, the development required of microwave components to support operation at 100 kW, cooling of the transmitters and waveguide components, physical arrangement of the hardware to minimize the footprint used in the BWG pedestal, and operational concerns and scenarios. The following sections describe the results of the studies.

A. Interface to the Feedhorn

By design, the JPL X/X/Ka feedhorn must be driven by four separate inputs, with equal amplitude and phase. The genesis of this effort was to devise a means of attaching a 20-kW transmitter to each of these inputs, so a question is: If there are four PAs and the feed requires four inputs, is there a better way than power combining the PAs into a transmission line, and then splitting it back to four lines for interface with the feed?

The answer is "yes and no." The simple way to meet the requirement that the four inputs be of equal amplitude and phase is to combine and split in such a way that each output contains a fraction of each input. There is a possibility of cross-strapped amplitude and phase loops around the PAs, but such an approach can't handle a single failure, so it's not robust. Therefore it is not acceptable, nor feasible, from a technical or operational perspective.

The schematic shown in Figure 21 uses 16 hybrids to accomplish the task, and has the added benefit that the power level never exceeds the individual input power.

A downside to this schematic is that trying to merge the concept into the goal of being able to pick any four out of five inputs — one of the features required for operational availability — would be very "messy."

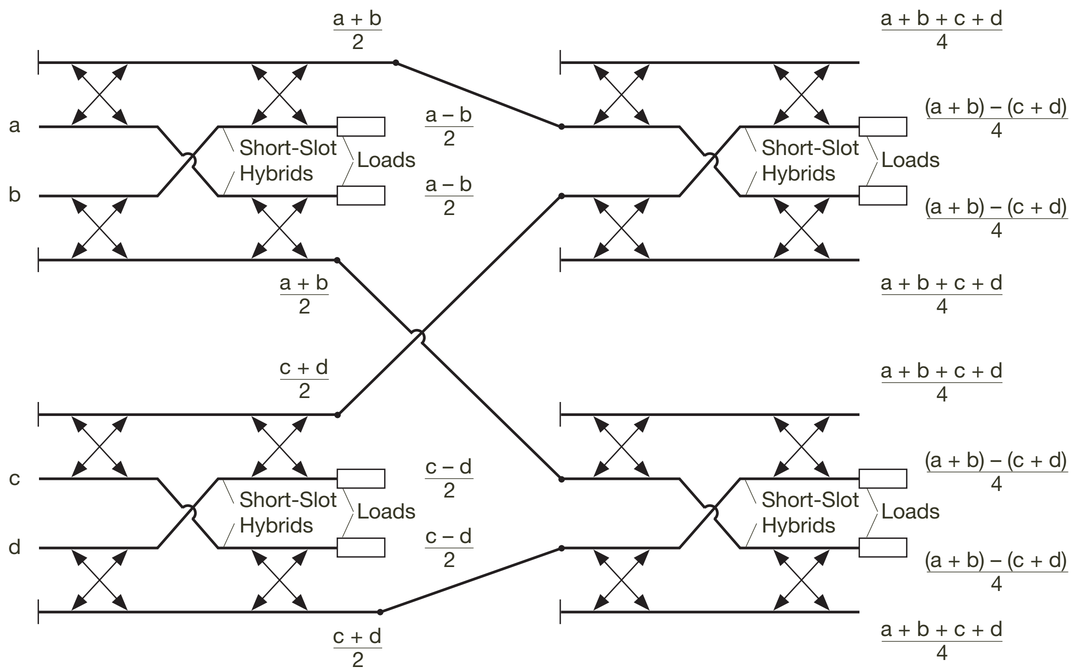

Figure 21. A 16-hybrid approach for interfacing four transmitters to the feed.

An improvement to the 16-hybrid approach is a similar system that started by combining the inputs in pairs (resulting in a doubling of the power in a single waveguide), then splitting, which requires only eight hybrids.

The obvious solution of combining all four inputs (three hybrids) to a single waveguide (four times the power) and splitting (three more hybrids) has the following features:

1. Lowest hybrid count (good).
2. Compatible with the switch matrix (Figure 1) for operation with one, two, or four inputs (good).
3. A single point for the interface between transmitter and microwave (good).
4. All of the power in a single waveguide (bad).

There is another approach when using a ladder combiner based on hybrids and high-power phase shifters. The output of each stage in the combiner consists of two signals of equal amplitude and (potentially) equal phase no matter what has happened in the previous stages. To get the four equal signals required for the feedhorn, we omit the last hybrid in the fourth combiner, and split the two equal amplitude signals that we were going to combine, as shown in Figure 22.

It is unfortunate, but the fourth high-power phase shifter is still required. Although the amplitudes will be equal, the relative phase of the signals depends on the ratio of the input powers. There is also the problem that we lose the power in the fourth waster load to use in setting the phase shifters. The obvious solution is to provide a "simulated waster" by

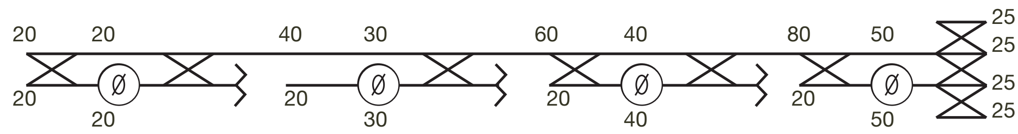

Figure 22. Ladder combiner producing four equal outputs.

combining samples from the two signals to be split in a low-power hybrid, as shown in Figure 23. Although the directional couplers and cables must be carefully calibrated, it should be stable when completed, and logic and control system for all four stages may be kept the same.

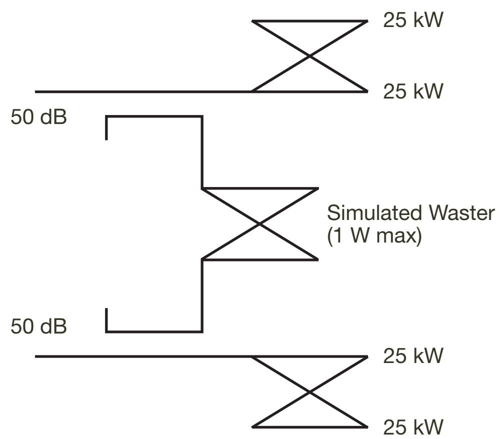

Figure 23. Four-way splitter with simulated waste load.

B. Components

One of the advantages of using multiple 20-kW transmitters is that there is little development required of the microwave components; however, some development must occur for those components and some care must be taken in distributing the inputs to the transmitters. The following subsections discuss those issues.

Filters

One of the most severe requirements imposed on the DSN X-band transmitters is rejection of the second, third, and fourth harmonics. Meeting this requirement with minimum loss requires a reflective "waffle iron" filter that has small gaps and is a difficult component to produce at 20 kW because arcing can occur between the gaps. In the RAIK design, the harmonic can be filtered out at the 20-kW level, so existing filters will suffice. However, there are no companies that produce the filters that we currently use for the 20-kW transmitters. New sources need to be developed and qualified.

Loads

Water loads would need to be qualified at 80 kW or more. There is a reasonable chance that the existing DSN loads would suffice if the water flow were increased to 30 gpm. There is also the GSSR load that is not yet satisfactory at 300 kW, but is close, to scale to the DSN frequency.

Couplers

Although not shown on the simplified diagrams (Figures 8 and 9), there is a need for forward/reverse power couplers at 40 and 80 kW. Currently these do not exist; however, their development should be low risk.

Excitation

Because the geometry of the combiner array is fixed, unlike the uplink array concept, all units can be driven by the same modulated signal with, at most, a static phase offset so that the exciter modifications can be as simple as adding an amplifier and five-way splitter as an adjunct to an existing DSN exciter.

In providing inputs to multiple transmitters, any variation in the group delay between various paths limits the dynamic bandwidth. A 10-ns variation will still allow nearly 100-MHz bandwidth, but care must to taken to:

1. Match the physical line lengths in the various paths both in the exciter splitter and in the waveguide combiners.
2. Verify that the klystron tuning results in similar group delays.

This concern will require that there be constant temperature controls among the various systems.

C. 80-kW Transmitter Cooling Architecture

It would be desirable to provide dual cooling manifolds to allow for independent maintenance of one transmitter while others are in operation. Figure 24 shows one potential water distribution arrangement that allows one or more of the PAs to be operated in maintenance mode while others are supporting operations. Not shown in the figure is a waveguide switch between the transmitter and combiner that would allow an individual PA to be operated into the water load for maintenance purposes.

A key to understanding the cooling system is that the maximum power of the entire system is no more than the sum of all five transmitters working at the same time. Because the sum of power through the loads is always less than 100 kW, 30 gpm is adequate for cooling all loads. The configuration shown in Figure 24 allows for operation by any combination of one to five transmitters.

The water cooling assembly design needs additional thought. It could use six BWG-style 30-gpm units that are already designed, or a new scaled-down design based upon the 400-kW GSSR transmitter. One of the key elements in any design is that the failure of one part of it should jeopardize no more than the operation of a single PA.

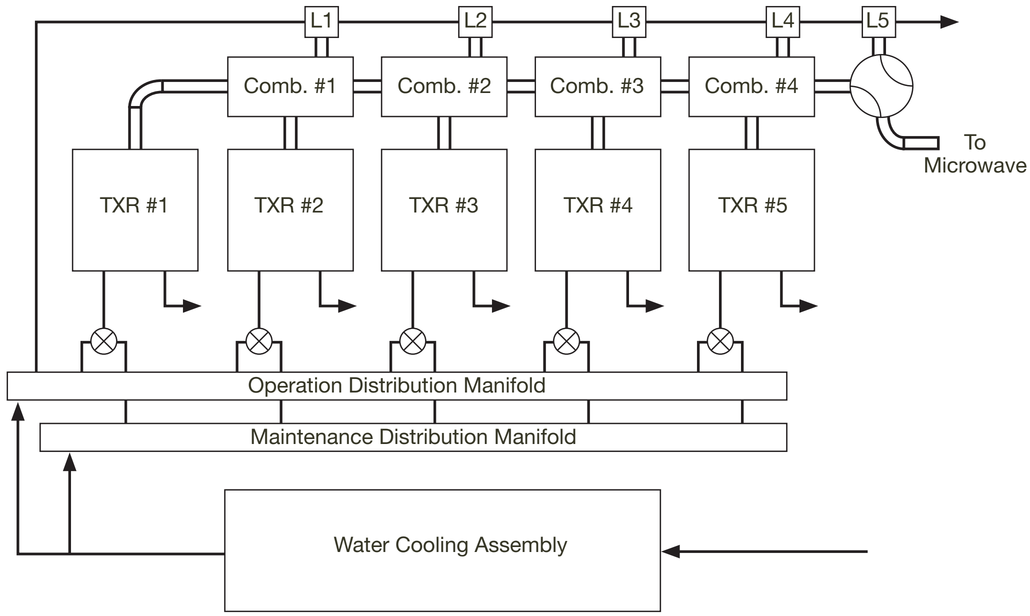

Figure 24. Candidate water distribution architecture.

D. Physical Arrangement Options

This section considers several possible ways of laying out the multiple transmitters. The factors are floor space near the X-band feed and adequate access for maintenance on at least one PA while others are supporting operation. In these figures, we have used an approximation of the BWG 20-kW X-band PA that, at the top level, consists of a half-rack for electronics, a custom enclosure for the klystron and high-voltage area, and a waveguide assembly that contains the arc elbow, couplers, filters, isolators, and waveguide switch. The space required is significantly deeper than a rack because space must be provided for the waveguide components extending from the klystron.

Figure 25 shows a front view of five PAs in line. This arrangement requires approximately 108 ft^2 of floor space in an area 12.6 by 8.5 ft. The overall height is estimated to be about 9.3 ft.

The disadvantages are that the arrangement leaves no access to the sides of the PA cabinets. Such access is needed when working on the klystrons, because the manifolds are separated from the PAs and require some redesign of the existing 20-kW transmitter for a stacked arrangement. Figure 26 shows the top view of the same configuration.

Another possible arrangement that uses only 90 ft^2 of floor space is shown in Figure 27. This concept stacks the electronics sections of two PAs or combines them into a single full-sized rack and separately stacks the klystron and high voltage (HV) areas for the same two, so that a PA has its electronics and klystron side by side. The arrangement improves the access to the klystron area somewhat for maintenance but would require some redesign of the PA to accommodate a different physical arrangement of the key components. Access to the waveguide components of the stacked PAs would be difficult. The fifth PA would retain the current BWG design.

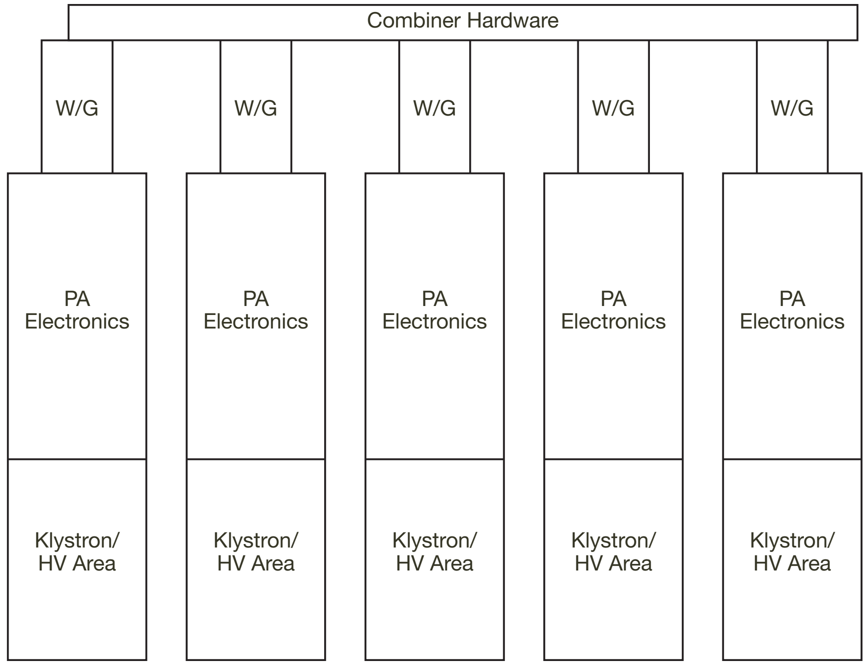

Figure 25. Five-in-line front view.

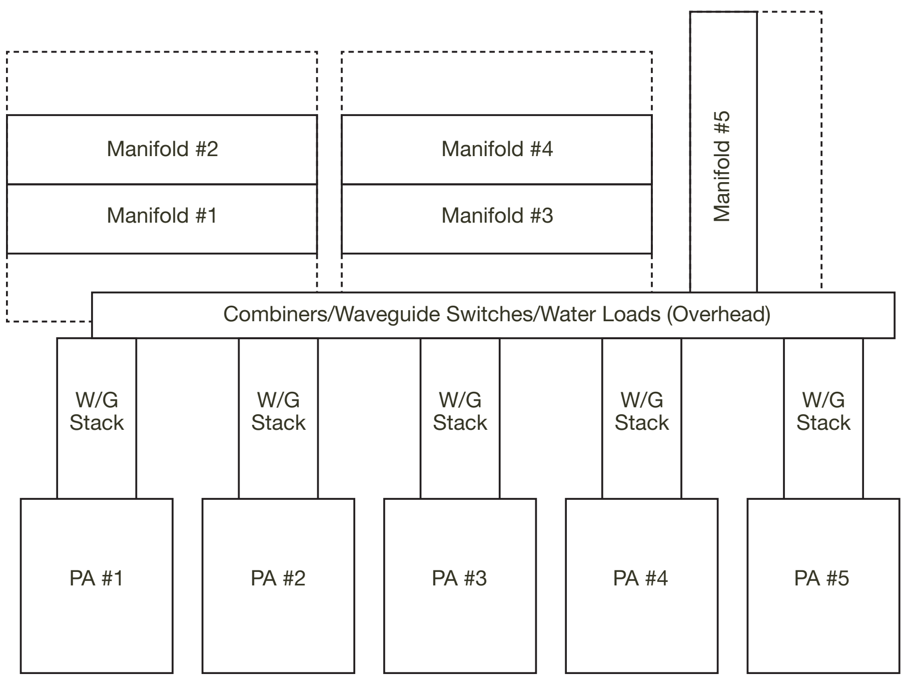

Figure 26. Five-in-line top view.

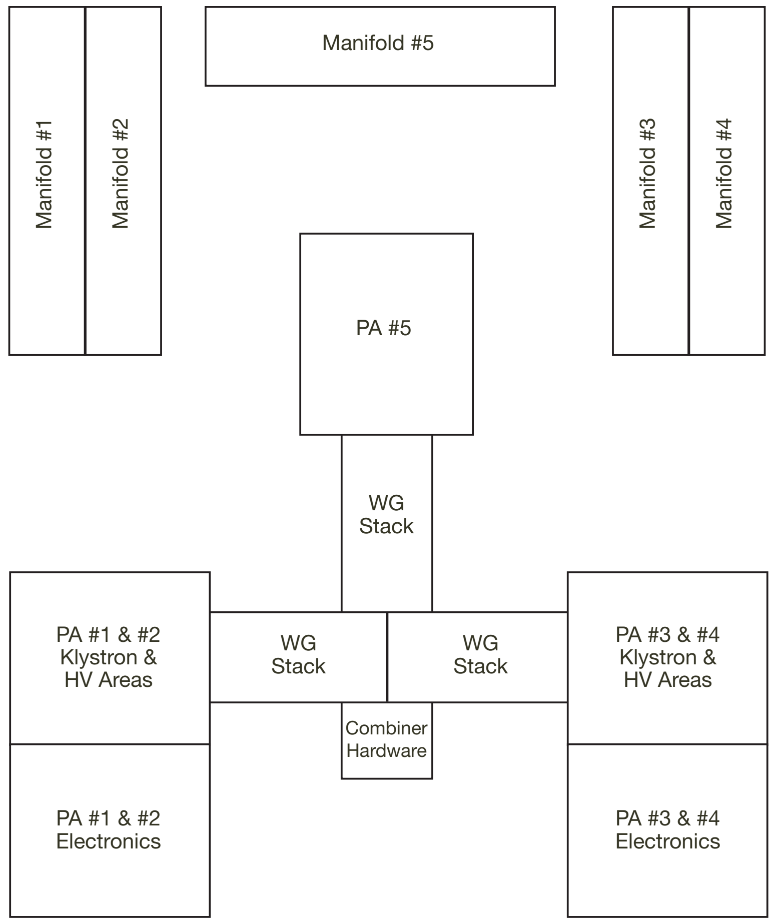

Figure 27. Compact physical arrangement.

Disadvantages are that the waveguide area would be difficult to access. Some redesign of existing 20kW transmitter would be needed.

A third option, a pentagonal arrangement as shown in Figure 28, uses about 170 ft^2 of floor area but permits about the same access to the transmitter as in the current BWG X-band. The combiner hardware would be in the middle and potentially be difficult to access for maintenance. Also, the waveguide would require custom 72-deg bends to align properly with the combiner.

E. Operational Concerns

The sheer number of components and the number of ways things can fail will require a more complex control system, but the basic architecture envisioned of having a collection of tasks communicating with each other would appear to handle it, and would be tractable.

Control System

A simplifying constraint to control would be to limit operation to a small set of fixed power levels, rather than requiring each transmitter to operate at any level from 200 W to 20 kW

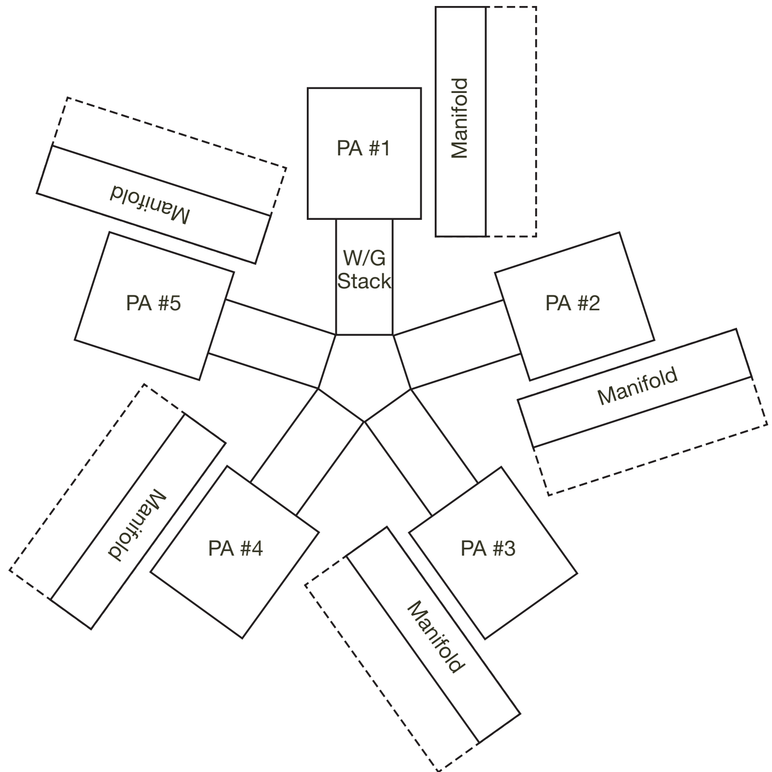

Figure 28. Pentagonal physical arrangement.

and achieve saturation at any level from 2 kW to 20 kW. Having fixed power levels would allow precalibration of the voltage and drive required within a small window to allow for tube aging rather than having to search for any arbitrary operating point.

Handling of interlocks will require special care. There are numerous switches, many of which may be either a “don’t care” or part of a critical path at a given time.

Clearly, there must be an elaborate control system governing the operation of the ladder combiner. Following are some initial thoughts on some of its desired properties.

Assumptions

Following are some basic assumptions regarding the requirements for the control system and some of its basic features:

1. Each power amplifier will have its own control system very similar to the BWG transmitters but simplified. Each will be self-protecting and maintain its own calibration.
2. Each power combiner will have its own control system, with a very simple interface, with inputs like “tune for output,” “tune for waster,” or “do nothing,” and outputs like “busy,” “ready,” or “fault.”
3. Although the control system will be distributed, the transmitter will present a single high-level interface to the uplink controller. This implies a top-level controller that coordinates the individual PAs and combiners and provides the external

interface. Possible extensions to the current interface include an indication of operational priority (e.g., routine or critical to influence the number of transmitters brought on-line to support a given uplink) and a means of “lock out” of a particular PA so maintenance can be performed while other PAs are operating.

4. Saturated mode, if supported at all, will be even more relaxed than in the BWG transmitters.
5. There will be no closed-loop control over either phase or amplitude after initial power setting.
6. It is desirable to minimize the calibration required on klystron change and to try to maintain a database of performance to allow the system to learn its settings during operation.
7. It is likely that most of the uses for this transmitter will be in the 1 W to 20 kW range (just like the current BWG PAs).
8. It is desirable to avoid unnecessary filament hours on the individual klystrons.

Although the following are not necessary, they are processes that will simplify operations:

1. A unified database for all performance and calibration data.
2. Separate private LANs for each PA and combiner. Because there is so much equipment, a private LAN will allow, e.g., each data collection unit to have its own Internet Protocol (IP) address, independent of where it is being used. Setting IP addresses on spare instruments is a problem in the field.

Scenarios

This subsection describes a selection of operational scenarios to illustrate some of the control system issues.

Cold start. Transmitter is off. A command is received to configure for *P* kW. (Of course, the transmitter is probably not totally off — the heat exchanger will be operating to keep the temperature of all PAs at about 40 deg C to allow the system to achieve stable operation within about 15 min. There may also be continuous logging of environmental parameters to assist in troubleshooting and continuing engineering support.

The top-level controller (TLC) calculates the number of 20-kW PAs required; then, if *N* is less than the number available, it adds one (in case of turn-on problems). It checks the database of filament hours. It chooses the units with the lowest hours and sends a “warm up” command to those PAs. If the support is identified as critical, the TLC would use all available PAs for the track to minimize the impact should one PA fail during the track.

Warm start. The TLC has *N* PAs warmed up. A command is received to configure for *P* kW. If *P* is less than 20 *N*, no additional PAs are needed. If *P* is greater than 20 *N*, the TLC sends warm up to enough PAs to satisfy *P*, with one spare if available.

This implies that the number of warm PAs will only increase (up to the number available.) The only time warm PAs are shut down is on an explicit “Shut Down” command.

Set power or calibrate. A command to calibrate at P kW is received. There are M PAs with warm filaments.

The TLC calculates N = floor([P/20]). If $N > M$, the command is rejected. If not, the TLC somehow picks N PAs. It calculates and sets a beam voltage for (P/N) kW. It sends beam-on commands to the N PAs, followed by a command to set power to (P/N). It waits up to some time limit for calibration-valid flags from the PAs. It then cycles through the combiners, i.e., tells combiner #1 to tune for output and waits for completion. Then combiners #2, #3, and #4 in turn are given similar commands. It goes through this at most twice. The TLC reports success if all the waster powers are below some threshold AND output is within a window, perhaps 0.2 dB of P. Otherwise, the TLC reports a failure. The TLC does not automatically go to standby, because an explicit command is required for that.

A single PA fails during operation. The TLC does nothing without operator command. The output drop will be as shown in Table 4. For cases where there were four or five PAs on before the failure, a drop of less than 3 dB is acceptable to operations.

If the combiner is retuned after the failure, the power loss (compared to before the failure) is shown in Table 5.

Table 4. Power output drop for single PA failure with no recovery.

1 of	dB
5	1.93
4	2.50
3	3.53
2	6.02

Table 5. Power output drop for single PA failure with recovery.

1 of	dB
5	0.97
4	1.25
3	1.76
2	3.01

Maintenance mode. The TLC will provide high-level commands for most normal cases. For example, a "Zero meters" command to the TLC will relay the command to each of the subsystems. Individual transmitters will be able to be made unavailable for operations so that maintenance can perform any required corrective or preventative maintenance function. During operation, it will be possible (through port mapping on the individual routers) to access the maintenance mode of the individual subsystems.

F. Reliability Considerations

Uplink reliability has increased in importance over the past five years and is likely to remain a significant factor in transmitter design. In the past, the DSN has achieved uplink redundancy by having multiple stations at a complex scheduled to support a critical spacecraft operation, and very critical operations are scheduled so that the spacecraft is in view from two complexes during the critical operation. The motivation for having an 80-kW capability is to provide a replacement for the 20-kW transmitters on a 70-m antenna, of which there is only one per complex. It is unlikely, therefore, that there will be more than one 80-kW transmitter implemented per complex. Consequently, the uplink redundancy available for 20-kW uplinks through multiple stations will be unavailable, resulting in the need to provide a very reliable 80-kW transmitter.

Types of Redundancy

Hot spare. One way the redundant PAs may be used is as "hot spares." Having a hot spare means that there is at least one unit online that can replace a failed unit in a matter of seconds (if the filaments are warm) or minutes (if the filaments need to be heated). Having a hot spare can result in a dramatic reduction in the mean time to restore service (MTTRS); however, loss of the redundant PA will likely cause temporary loss of the link, so a means to provide graceful degradation is needed.

Graceful degradation. Rather than having a hot spare available, another way of using multiple PAs to increase availability is by having several units contributing to a link so that an individual failure can occur without losing the link. The operating assumption here is that the DSN assigns a minimum of 3 dB margin to the uplink transmitter power.

For the case of combining four inputs with hybrids, either the obvious three-hybrid case or the balancing network discussed in Section IV.A, the loss of a single input results in a drop of 2.5 dB in the output, which is within the DSN margin. Also, as shown in Tables 4 and 5, the loss of a single PA when using the ladder combiner results in a drop of less than 2 dB in output power with no retuning and less than 1 dB output if retuning is possible.

To quantify, assume there is a probability p of a single PA failure (and assume for the moment that failures are independent.) If a link is using four PAs, then the probability that one (or more) will fail is $4p$. Even though this is worse than the single-transmitter case, the link is still good unless two units fail, and that probability is $6p^2$. (For completeness, the probability of three failures is $4p^3$ and p^4 for four failures.)

Thus, if p is small, there is a dramatic increase in the odds of having a good link. For example, if the probability of a single-transmitter failure during an 8-hr link is 0.1 percent, the probability that the link power drops more than 3 dB below its starting level is (6×10^{-6} + 4×10^{-9} + 1×10^{-12}) for an estimated reliability during the link of 0.999994.

The assumption that the failures are independent is very optimistic. There is a lot of common infrastructure, so a failure in one is likely to affect all the units. Because the transmitters are in close proximity, high-power arcs create enough electromagnetic interference or perturbation to the supply voltage that adjacent transmitters could easily be shut down as well. This analysis is not rigorous, but serves to point out the relative merits of graceful degradation.

Redundancy using only two transmitters. When combining two PAs with a hybrid, the loss of either one results in a drop at the output of 6 dB, which exceeds the DSN margin. There are two possible approaches to alleviating problems:

1. Convince the spacecraft operations and the DSN to assign a 6 dB margin to the transmitter. While the probability is unlikely, it might be acceptable in some critical situations.
2. Phase the two inputs in quadrature, so that half of the total power goes to the output, and half to the dummy load. In this case, the loss of one input will result in only a 3 dB drop in output (along with a 45 deg phase jump.)

Reliability Models

To assess the reliability of the various configurations, we used the Item Toolkit by Item Software and generated models for four architectures.

Functional Blocks are taken from the existing 20-kW BWG transmitters, and are common for all the models. Reliability estimates are based on those presented at the BWG 20-kW transmitter Preliminary Design Review (PDR).

The module definitions are:

- Monitor and control (M&C) is the transmitter control (TXC) portion of the high-voltage power supply (HVPS).
- Power is the HVPS (less the TXC) plus the motor-generator (MG) and MG controller.
- Cooling is the heat exchanger and manifold for one PA.
- PA and waveguide (PA & WG) is the single PA and waveguide stack.

Figure 29 shows the model for a single 20-kW transmitter. The predicted MTBF is 1815 hr, which is consistent with the value presented at the BWG 20-kW transmitter PDR.

The single-string architecture can be transformed into a multiple-transmitter model by replicating the power, cooling, and PA & WG module, using a common M&C module and adding a combiner module. The result is shown in Figure 30. Because all four strings are

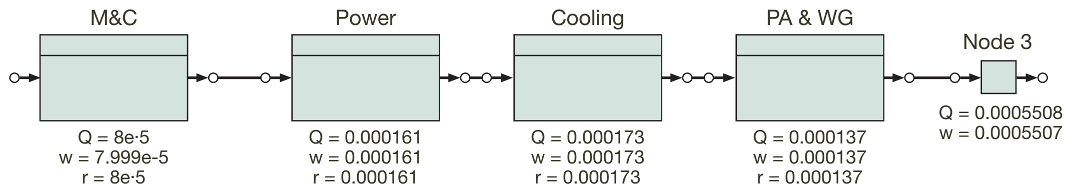

Figure 29. Single-string 20-kW architecture.

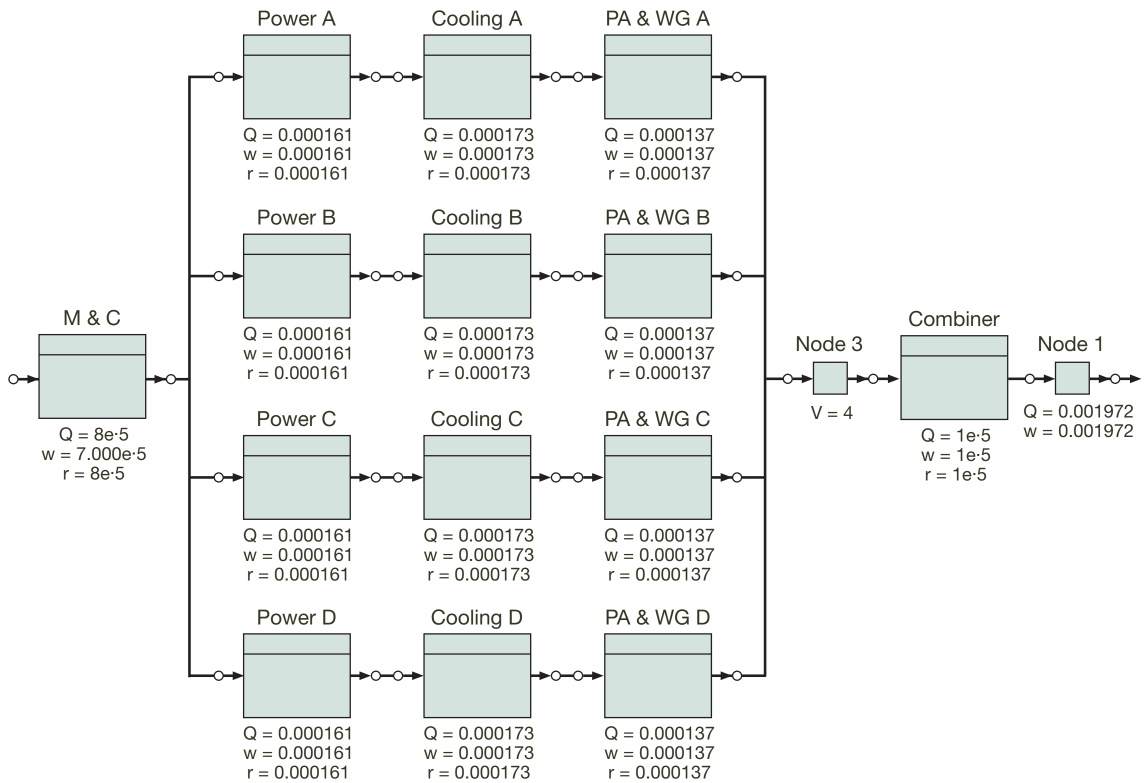

Figure 30. Four-element RAIK architecture.

required to achieve full output, the reliability of this configuration at 80 kW output is only about 500 hr. There is essentially no redundancy at all in this configuration.

Adding a fifth transmitter string, as shown in Figure 31, markedly improves the overall reliability to an MTBF of over 10,000 hr.

Because the M&C and the combiner functions are not redundant in this model, they are drivers in system reliability. If we were to lower the failure rate on M&C and combiner to 1e–6 each, the MTBF would be over 200,000 hr. The model shown here is not consistent with a scheme that uses independent combiner modules and distributes the M&C among the various transmitters, but it does provide a worst-case estimate of overall transmitter reliability.

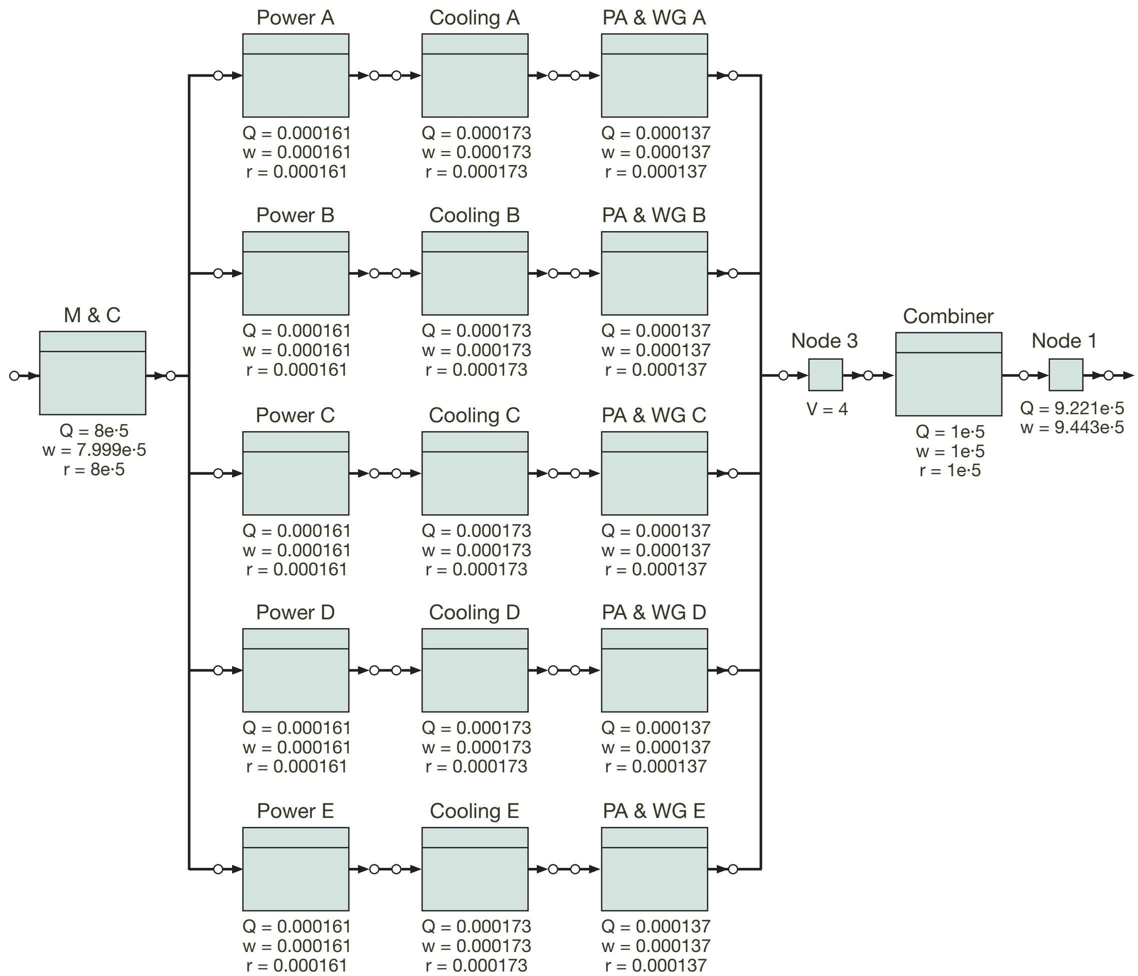

Figure 31. Four of five RAIK architecture.

Another approach to redundancy is to use multiple functional modules so that a failure in Power A and another in Cooling B will remove only one transmitter from operation rather than two. Multiple functional modules are modeled in Figure 32. The difficulty with this approach is that the "bussing" of power, cooling water, or RF may, in itself, be unreliable. In particular, combining high voltage from several sources can lead to common-mode failures that would shut down the entire system. Reliability is still driven by M&C, and RF combiner failure rates, but adding failure rates for the power buss and cooling combiner actually lowers the MTBF to less than 10,000 hr.

If combining could be done reliably, using multiple functional modules would yield the best reliability, but the simpler technique of combining single-string transmitters meets or betters existing reliability requirements. The results of reliability modeling are summarized in Table 6.

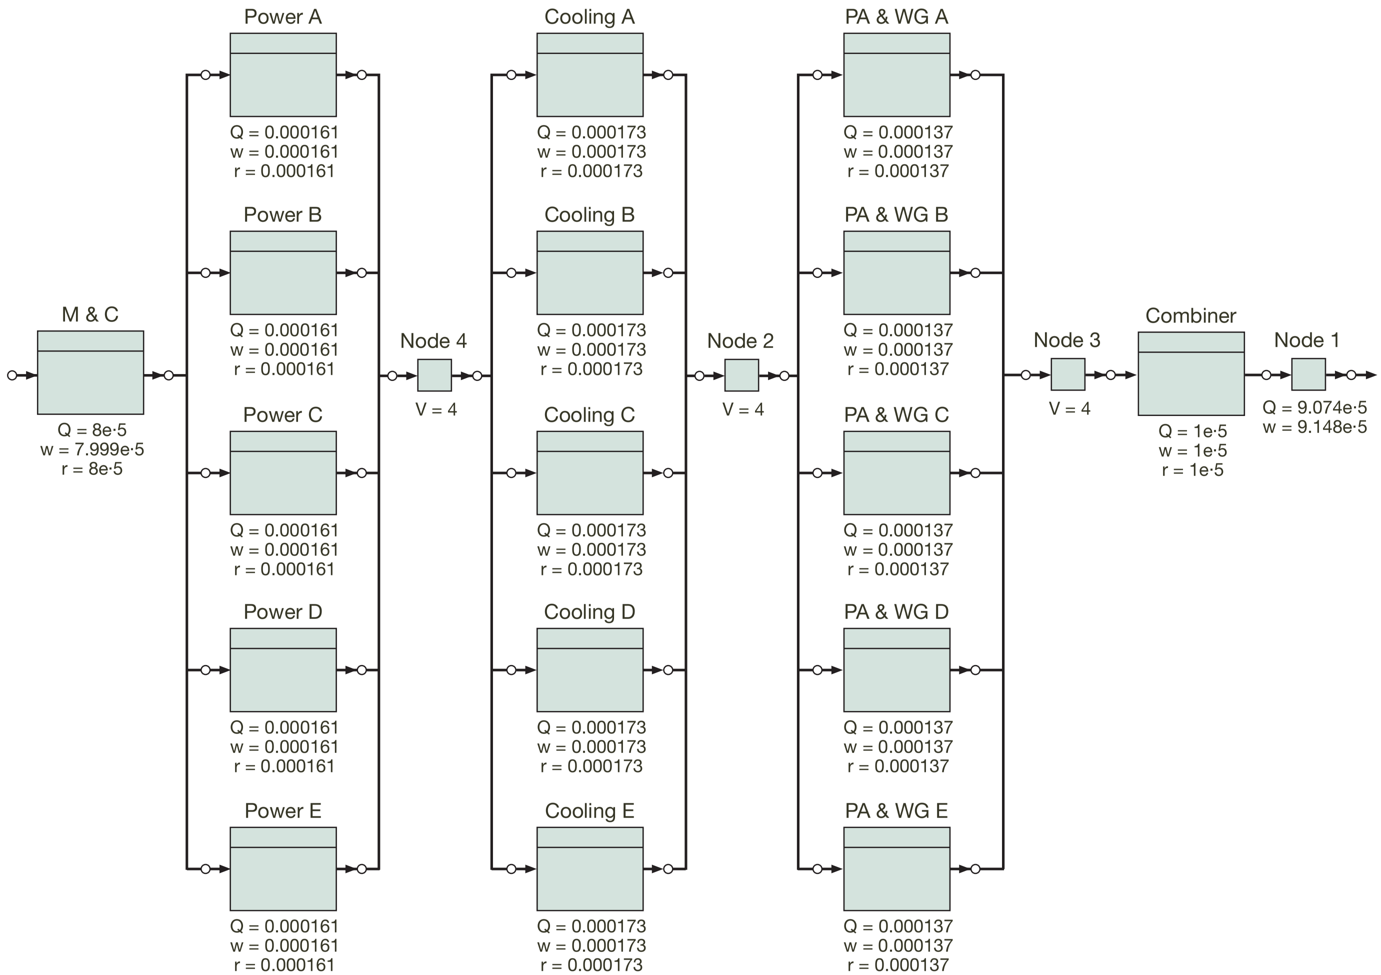

Figure 32. Four of five RAIK with multiple functional modules.

Table 6. Reliability summary.

Configuration	MTBF, hr
Single string	1,815
Four element	500
Four of five	>10,000
Multiple functional modules	<10,000

IV. Comparison of RAIK and a Single 100-kW Transmitter

The objective of the study was to determine if using a multiple-klystron transmitter would be a viable option to development of a 100-kW transmitter. The advantages, disadvantages, and associated issues are summarized in Table 7.

A detailed cost comparison was performed early in the effort; however, there have been enough changes over the past two years that the information used to develop the estimate is outdated and the estimate is now of questionable value. The initial cost comparison shows that the two approaches had approximately equal cost. Clearly, the single-tube approach will require significantly more nonrecurring engineering (NRE) than the approach of using multiple 20-kW transmitters. The tube doesn't exist and, although built from existing designs, will be built in very small numbers in comparison to the current 20-kW tubes.

Similarly, the development of harmonic filters at the 80- to 100-kW power levels will require significant NRE. It is unlikely that the reflective "waffle-iron" filter used in the 20-kW transmitter will be usable at 100 kW unless pressurized nitrogen or SF_6 is used in the waveguide. Using a multistage absorptive filter to achieve the required harmonic suppression may result unacceptable loss at the primary frequency.

The HVPS will require a complete redesign for a new tube. The existing HVPS has a transformer/rectifier section that works satisfactorily without requiring an oil tank to prevent arcing in air. Implementing a design that requires an oil tank would require significant effort and impose additional requirements on the installation for oil containment and oil quality management.

The transmitter group has proposed for many years an effort to purchase and evaluate a solid-state AC-to-DC inverter-style power supply that might better handle the higher voltages needed for the 100-kW tube and also be a more compact replacement for the current 20-kW HVPS and its associated motor-generator hardware.

In producing the BWG 20-kW transmitters, we discovered that the cost reduction from the "learning curve" of producing more than just three transmitters (as in the 70-m 20-kW transmitter task) was significant. It is likely that producing five, 10, or 15 20-kW transmitters will yield significant unit cost reductions. It is probably worth investing some NRE to improve the reliability and maintainability of the 20-kW assemblies prior to "large-scale" production to reduce the maintenance burden during operation. Prime candidates for improvement are the arc detector and data acquisition assemblies.

Table 7. Comparison of RAIK and single-tube approaches.

Adv/Disad	RAIK Approach	Single-Tube Approach
Advantages	• Development cost for a single PA in a BWG is low because few changes are needed to the existing BWG PAs to build more. • Replacement klystrons are relatively inexpensive (i.e., about $80K, in quantity). • High reliability is possible by using five 20-kW transmitters to provide the 80-kW capability.	• Requires about the same space in the BWG pedestal as the existing 20-kW transmitter. • Less total hardware than the RAIK approach.
Disadvantages	• Physical space in the BWG pedestal is large. • More transmitter hardware is needed so more maintenance will be required.	• Significant NRE costs for the klystron, HVPS, PA, and waveguide components. • There are many single points of failure. The overall reliability is likely to be somewhat worse than a single 20-kW transmitter due to the higher voltages and RF power levels. • Klystron cost is estimated to be about $200K each.
Development Needed	• Combiner needs development for operation at 100 kW. • Monitor and control software for overall transmitter control. • Filters for 20 kW (need to find a new source). • Couplers for 100 kW (need to check current design and modify if needed).	• Klystron must be developed. • HVPS must be developed. Due to higher voltage levels, the transformer/rectifier section may need to be in an oil tank. • A heat exchanger for ~350 kW needs to be developed. • The power amplifier assembly can use some of the 20-kW subassemblies, but significant changes are required to accommodate the larger tube/magnet. The cathode may require an oil tank. Some existing assemblies may require change to accommodate tube requirements. • Waveguide components: —All waveguides need to be qualified for 100-kW power levels. —Arc elbow. —Couplers. —Harmonic filters. Note: filter technology may not be able to achieve the same harmonic suppression with acceptable loss. May have to use gas such as SF_6 to prevent arcing. —Circulator.
Other Issues	• The 20-kW PAs can be assembled and tested at JPL. A pair of transmitters can be tested with a combiner with little or no modification to the existing facility. Testing of a full-sized transmitter would require facilities modifications or would have to be done at DSS-13.	• Testing at JPL would require facilities modifications. Performing the testing at DSS-13 would be more expensive than final testing of the RAIK design.

V. Conclusion

After developing and testing a working waveguide combiner, the use of multiple 20-kW transmitters to provide a higher-power uplink appears to be quite feasible. Although the parts tested are not suitable for deployment due to a lack of cooling, the concept is practical and should be considered for further development should there be a desire to go to higher uplink power in the DSN. The main disadvantage of this approach is the physical space required for the hardware and the total amount of hardware itself. The reliability of the individual transmitter will not be significantly enhanced, but the ensemble will have much higher availability than any current 20-kW transmitter and will have significantly better availability than a single-tube transmitter at the 80-kW output level. The design concepts presented here should permit maintenance to be done on individual transmitters while other transmitters are still providing uplink to spacecraft. Because of the reduced NRE requirements for the multiple 20-kW transmitter approach and reductions in unit cost from large-scale production, it is likely that the RAIK approach will be less expensive and lower risk than a single PA based on a 100-kW klystron.

The technology still needs additional development in the following areas:

- Combiner power handling at 100 kW average power.
- Combiner loss needs to be reduced to less than 0.5 dB.
- The overall control of the multiple 20-kW transmitters needs to be developed. The development could probably be done effectively by using transmitter simulators in software rather than requiring that all hardware be present.

Concerns that should be addressed before committing to implementation include:

- Startup dynamics and individual calibration: A scheme is needed to ensure an orderly startup. One approach might be for each PA's power to be set combining by having a separate switch and water load.
- Dynamic phase control to minimize waster power: Should there be an active phase-control loop? What is the effect on the uplink phase to the spacecraft? Is waster load power an adequate measurement to determine appropriate phase settings?
- Does there need to be an amplitude-control loop? It is conceivable that the loss of power from a single transmitter could be offset by increasing power output from the remaining transmitters.
- Cross-strapping of heat exchangers seems feasible, but cross-strapping of high-voltage supplies presents many problems.
- How are the water loads cooled? Is a separate heat exchanger needed?
- Will the system need to support "saturated" operation? If it is more linear, variations in exciter drive are more likely to cancel, but saturation might improve amplitude stability.

Acknowledgment

This activity was funded under the NASA–JPL Interplanetary Network Directorate (IND) Technology Office in 2007 and 2008.