GBT MEMO 181

GBT telescope and instrumentation control system hardware architecture: computers, networks, interfaces, and timing

J. M. Ford [a]

[a] National Radio Astronomy Observatory, PO Box 2, Green Bank, West Virginia, USA

ABSTRACT

The Green Bank Telescope (GBT) is designed to be a flexible instrument, accommodating many different types of front ends, back ends, and observing styles. In order to support this flexibility, a system architecture was developed according to the principle of modularity. The system is a loosely coupled group of cooperating computers, tied together to form a complete system. The hardware design of the GBT Control System will be presented. The computer architecture, network architecture, and system synchronization and timing methods will be addressed. Progress towards implementation will be described, and the lessons learned as a result of implementing the system will be covered.

Keywords: Architecture Networks Timing Synchronization

1. INTRODUCTION

The Green Bank Telescope (GBT) is a collection of mechanical and electrical subsystems, that, taken together, form a radio telescope. The coordination and control of these subsystems is accomplished using the GBT Monitor and Control (M&C) system, a hierarchical distributed system consisting of software modules running on multiple computers communicating over a local area network (LAN). The M&C system is implemented using general purpose computers as well as specially designed hardware. The system is designed to minimize real-time communications and timing dependencies between subsystems. This allows the non-deterministic Ethernet protocol to be used successfully for control purposes and for data transport. The whole system is synchronized via a combination of Network Time Protocol (NTP),[1] InterRange Instrumentation Group B (IRIG-B) and 1 Pulse Per Second (1PPS) signals distributed by the site timing center.

2. COMPUTERS AND PERIPHERAL EQUIPMENT

The top level of control is provided by UNIX workstations running the Solaris operating system. These provide the platform for the user interface programs and the daemons that collect data, as well as some soft real-time tasks. Table 1 shows the characteristics of the workstations in the GBT M&C system.

Table 1. Unix Computer Characteristics

UltraSparc Processor (167 MHz)
128 MB Memory
2.1 GB system disk
Fast/Wide/Differential SCSI adapter
Fast Ethernet

For data storage, a multiport disk array (RAID) is connected to each of the data acquisition computers via fast wide differential SCSI-II interfaces. All data is stored in Flexible Image Transport System (FITS) binary tables,[2] unless the data rates are prohibitive, e.g. the spectrometer and spectral processor have a raw mode that simply dumps the data onto the disk in an unformatted manner. Data archiving will be supported with tape and CD-ROM writers.

E-mail: jford@nrao.edu, Web: http://info.gb.nrao.edu/ jford

2.1. Data Storage

The data collected by the various instruments is written to a 50-75 GB RAID. This RAID has 4 ports on it that may be simultaneously used by 4 different hosts, which provides an aggregate data rate of 60 MB per second, enough to support all 3 existing backends simultaneously. VLBI data will be stored directly on VLBI tapes, as is currently practiced at the 140 ft. telescope.

2.2. General Purpose Computers

The telescope operators will use screens displayed on multiple monitors to control and monitor the telescope. These monitors are standard 21 inch Sun monitors, with a resolution of 1152 x 900 pixels. The operator control computers are Sun UltraSPARC workstations, with each computer driving either two or three screens. The number of screens is dependent on the Telescope Operations group's designs, which are still being finalized.

Data will be collected and written to the RAID via any one of three computers. The spectrometer and spectral processor each contains its own data collection machine, and an additional data collection computer is used to aggregate the data from the rest of the telescope systems.

A fourth computer will be connected to the RAID for read-only access to the data via the Network File System (NFS). This workstation will be located in the equipment room and will serve as an engineering workstation, to be used mainly for troubleshooting and working with the equipment in the equipment room.

The Precision Pointing System will use another workstation as a compute server, and for coordinating the tasks of the Laser Metrology system and the Active Surface system. This computer is located on the telescope alidade, approximately 2.5 km from the control room.

2.3. Single board VME computers

Closer to the machinery, VME based computers contain the interface hardware, and run most real-time tasks under the VxWorks operating system. Hardware interfaces are provided to connect the telescope servos and instrumentation to the computer systems.

The real-time needs of the telescope are met using Motorola MVME-167 68040 based SBC's. Twelve of these run M&C code, while 2 others are supplied with the contractor's servo system. See Table 2 for the features found on the MVME-167.

Table 2. MV-167 Characteristics

68040 CPU, 33 MHz
8 MB RAM
SCSI Interface
10 Mbit/s Ethernet Interface
4 Serial Ports
1 Parallel Port
4 Counter/Timers
Demonstrated MTBF:
Mean 147,507 hours (16.8 years)
90% Confidence 85,522 hours(9.8 years)

2.4. Other VME hardware

Several other VME cards are used in the system including the Bancomm IRIG decoder, parallel I/O cards, IEEE-488 interface cards, and analog to digital converter cards. All the chassis used in the telescope have at least three slots open for future expansion.

The Bancomm BC-635 card is used in each chassis to receive and decode IRIG signals from the site timing center, and synchronize the chassis to the site clock.

The parallel digital interfaces are Motorola MVME-340 cards. These cards provide 48 bits of I/O and three timers. The cards have software written for them that allows the higher-level software to read and write bits on the card without having to program all the registers on the interface chips. The inputs and timers on these cards are used to generate interrupts for synchronizing the systems in the telescope; for example, the Tracking Local Oscillator (LO) can switch frequencies on the edge of a signal/reference (sig/ref) or calibrate (cal) signal.

The Mizar MZ-7500, a IEEE-488 controller/talker/listener, is used to control the IEEE-488 based instruments. This card provides 2 channels via a pair of TMS-9914 chips. These chips provide an interface that is largely software driven and is not particularly reliable in a multitasking environment. There is a plan underway to replace the 9914 chips with a newer chip, possibly the INES i9914 chips. The INES i9914 devices provide much of the IEEE-488 functions in hardware, thereby allowing better performance and reliability, and reducing the amount of code in the device drivers.

Three analog to digital converter cards from VME Microsystems International Corporation (VMIC) are used to provide monitoring of the 16 azimuth and 8 elevation drive servos. These cards have 64 channels of 16 bit differential analog input.

3. NETWORKS

The GBT control system depends on the LAN to function. The network has been designed to provide a high level of performance, yet be affordable.

The choice of a network technology was not easy. There were, as usual, several competing constraints. The networks proposed in the early stages of the project were completely different than that finally chosen for implementation. This was due to many factors. The original design included two networks, one for data transfer, and one for control information. The data load on one Ethernet network would have been too great to allow proper control of the telescope over the same LAN. Therefore, a data LAN and a M&C LAN were planned. Later, the M&C group decided[3] to use a multiport RAID to store the backend data, and so the proposed data LAN was unnecessary. At the same time, technology was advancing, bringing to market full-duplex Ethernet, switching technology, and 100 Mbit Ethernet (Fast Ethernet). The LAN technologies investigated include Fibre Channel, Asynchronous Transfer Mode (ATM), Switched Fast Ethernet, and Fiber Distributed Data Interface (FDDI). Each has its merits. FDDI is (or was) the most widely used high-speed LAN. ATM shows much promise. Fast Ethernet is a speedy version of a familiar friend. In the end, Fast Ethernet won out on cost and convenience factors. Fibre Channel is by far the highest performer, with a corresponding price tag. ATM and FDDI were also determined to be more costly, without apparent benefits. Our choice seems to have been a good one, as the Fast Ethernet market has taken off, and seems assured of a long life.

3.1. Network Architecture

The key to being able to use run-of-the-mill Ethernet to fulfill the needs of the GBT is to break the Ethernet up into segments, and use high-speed switches to connect the segments.

In a conventional Ethernet, each machine shares a portion of the 10 Mbit/s bandwidth. In a switched topology, the network is broken up into a set of smaller segments, with a much greater aggregate bandwidth. The packets from one segment are switched only to the segment where the target interface resides, and not to the entire network. If the switch does not know which segment the packet should be sent to, it sends it to all the ports, but the switch learns which devices are on each port, so this happens only when the switch is reset, or a device is moved from one port to another.

Figures 1 and 2 show the design of the telescope's networks. Several things are apparent from these figures. Many separate physical networks are defined, each with a local traffic pattern. The Antenna Control Computer talks to the servo system over a network free of other traffic. The Laser System also has a segment of the network with which to work. The Pointing UltraSPARC is connected to the network with 100 Mbit/s Fast Ethernet. In the equipment and control rooms, the computers all communicate via the Fast Ethernet. The legacy machines can still communicate with them, due to the ability of the switches to translate packets from 10 Mbit/s Ethernet to 100 Mbit/s Ethernet. As we retire the old machines, we can migrate the new ones to Fast Ethernet.

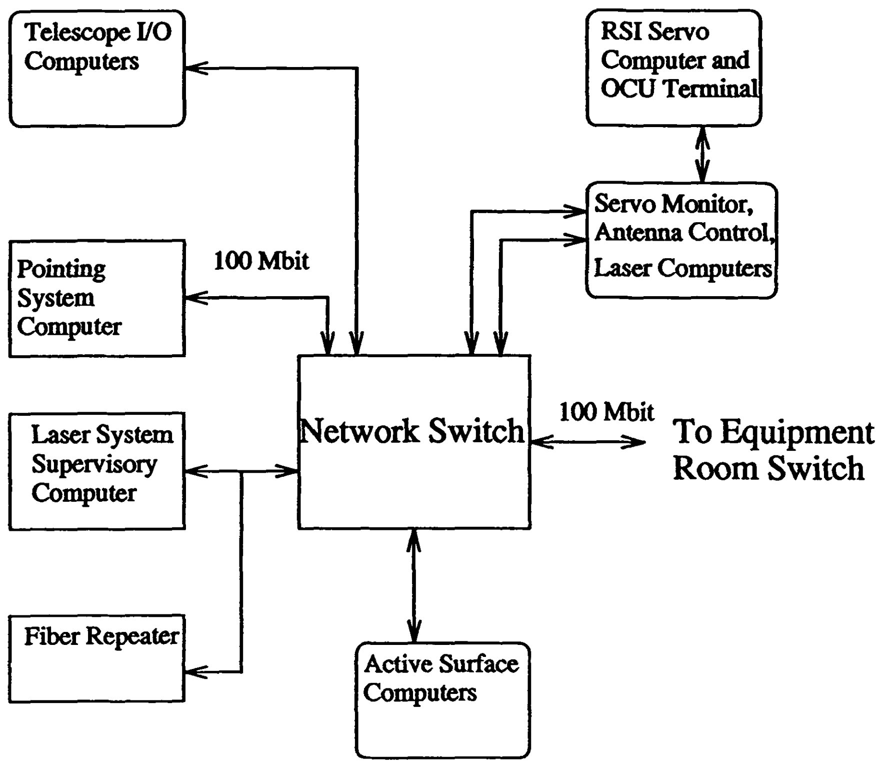

Figure 1. GBT Alidade Network Design

3.2. Network Load

Table 3 shows the estimated load on the network segments in the telescope networks shown in Figure 2 and 1. The design is based upon an Ethernet switch in the equipment room that can handle multiple Fast Ethernet segments. The data acquisition machine and the operator consoles are connected to this switch, while the 100 Mbit Fast Ethernet backbone link goes only to the alidade room switch. This topology is under consideration because it will reduce costs, while providing good performance, and allowing the network to scale even more. The maximum expected percentage utilization for any network segment on the GBT is 25%, which is the traffic from the data acquisition workstation. The maximum utilization on any network segment used for real-time control purposes is 8%. This is for the Digital Continuum Receiver, which logs its data to disk via the network. This traffic would not be active while the control commands are being issued. The traffic generated by control commands is negligible.

3.3. Network Hardware

The network uses the XYLAN OmniSwitch. This switch has been in use in our mock-up environment for about a year, and has proven to be reliable and fast. The aggregate backplane throughput for the switch is 960Mb/s, while the latency for switching packets is very low. Our configuration uses 2 switches, each with 10 Base 2 (50 Ohm Coax), 10 Base FL (full-duplex 10 Mbit/s fiber), and 100 Base Fx (100 Mbit/s single and multi-mode fiber). The switches are built with the ability to move to ATM or Gigabit Ethernet, if the need arises. Note from Sect. 3.2 that we have plenty of room to grow before we need to take that step.

4. INTERFACES

The GBT integrates many varied components. The system uses both NRAO produced devices and off-the-shelf electronic instruments, which requires us to support many different styles of hardware interfaces. The variety of interfaces that we were required to support has resulted in the ability to integrate new devices into the system with a minimum of work, both on hardware and software engineering. Integrating a "foreign" device is described in Sect. 4.5.

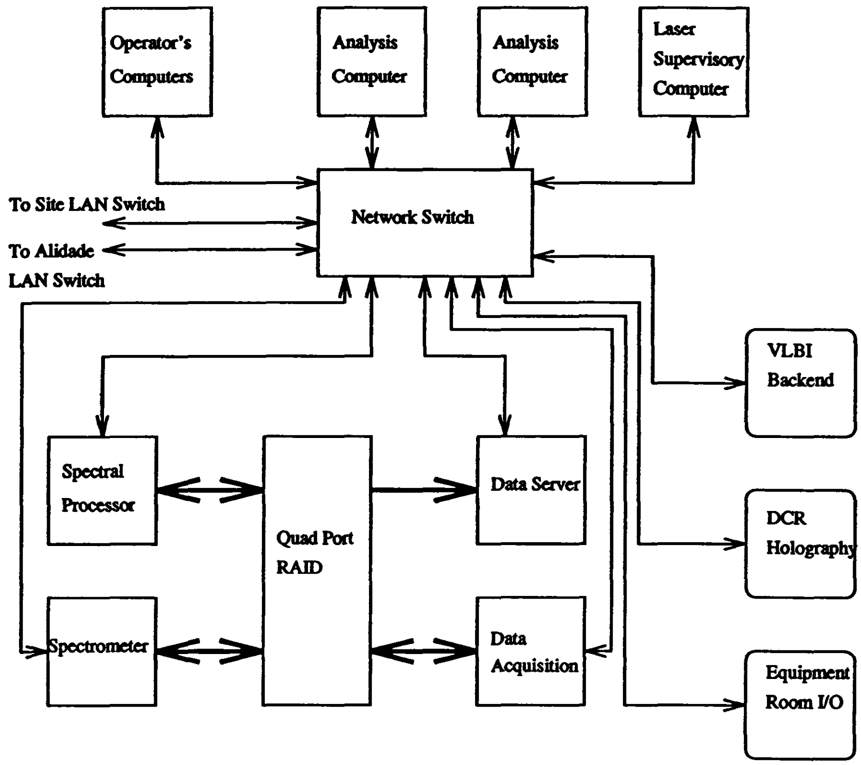

Figure 2. GBT Equipment and Control Room Network Design

4.1. VLBA Monitor and Control Bus

The VLBA Monitor and Control Bus (MCB) interface is an RS-485 serial interface used for connecting most NRAO built equipment to the GBT computers. The device has the ability to respond to monitor or control commands about every 3 ms using the MCB protocol. Specifications for this protocol are available from NRAO.[4,5]

4.2. IEEE-488

The IEEE-488 interfaces are used for communication and control between the GBT real-time computers and commercially procured instruments. In our system, these devices are synthesizers, spectrum analyzers, and counters. A fiber-optic extender is used to pipe the IEEE-488 bus up the the receiver room, where it is used to control the tracking LO equipment.

4.3. Direct Memory Access (DMA)

The two fast backends use DMA techniques to transport the data from the instrument onto a computer where it can be formatted and written to disk. The DMA interface on the Spectral Processor can transfer data at up to 10 MB/Sec. The GBT Spectrometer has a VME64 interface that can transfer data at the full design rate of the machine, approximately 13 MB/Sec.

4.4. Miscellaneous

Several miscellaneous interfaces have been implemented into the system, or are planned. Parallel interfaces are used in several places, such as to communicate with the 140 ft. telescope for servo control system testing. A raw encoder parallel readout will be implemented to allow a redundant check of the contractor supplied servo system outputs. Analog I/O points are provided to complete the interfacing of the computers.

Table 3. Network utilization

Machine	Destination Equipment	Link Speed, Mb/s	KB/s	Percent utilization
Equipment/Control Rooms				@80%
IF SBC, Weather SBC, SCRB	Equipment Room Lan Switch	10	1.5	0.1875
DCR	Equipment Room Lan Switch	10	64	8
Spectrometer	Equipment Room Lan Switch	10	0.4	0.05
Spectral Processor	Equipment Room Lan Switch	10	0.4	0.05
VLBA Station Computer	Equipment Room Lan Switch	10	0	0
Operator Consoles	Equipment Room Lan Switch	100	2000	25
Data W.S.	Equipment Room Lan Switch	100	2000	25
Equipment Room Switch	Backbone	100	150	1.875
Engineering/Server W.S.	Control Room LAN Switch	100	1000	12.5
Observer W.S. 1	Control Room LAN Switch	100	1000	12.5
Observer W.S. 2	Control Room LAN Switch	100	1000	12.5
Alidade Room				
Laser SBC	Backplane Network	10	8	1
Antenna SBC	Backplane Network	10	8	1
Servo Monitor SBC	Backplane Network	10	0	0
Antenna SBC	Alidade Room LAN Switch	10	15	1.875
Servo Monitor SBC	Alidade Room LAN Switch	10	20	2.5
Laser SBC	Alidade Room LAN Switch	10	40	5
Telescope I/O SBC 1&2	Alidade Room LAN Switch	10	2	0.25
Laser LAN Repeater	Alidade Room LAN Switch	10	50	6.25
Receiver Room LAN,SCRB	Alidade Room LAN Switch	10	0	0
Actuator Control Room LAN	Alidade Room LAN Switch	10	3	0.375
Pointing W.S.	Alidade Room LAN Switch	100	110	1.375
Alidade Room LAN Switch	Backbone	100	150	1.875
Total Backbone Traffic		100	150	1.875
Total Control Room Switch Traffic		1000	4000	5
Total Equipment Room Switch Traffic		1000	4066.3	5.082875
Total Alidade Room Switch Traffic		1000	150	0.1875

4.5. Requirements for "foreign" instrumentation

In order to connect and integrate a device into the GBT control system, there are only a few requirements that must be met. The device must have a way for the GBT control system to tell it when to start, and it must be able to report to the control system the earliest time it can start. In the native GBT systems, this intelligence is built into the device managers. In a device that is to be integrated, a simple device manager would be written that can translate the M&C commands into those suitable for the device, and the status of the device can be passed along to the control system. The interface to the device could be virtually anything; for example, Ethernet, serial ports, parallel ports, TTL signals, etc.

5. TIMING

5.1. Site timing center

The site timing center contains the hydrogen maser frequency standards used throughout the site, as well as the IRIG source and 1PPS sources.

The site timing center is also the source of all the phase-locked LO reference signals used by the telescope. A 100 MHz signal generated by the maser is multiplied by 5, and combined with the 10 MHz signal. It is then split into 4 outputs, each of which may be transmitted to a different location. The combined signal is then put on an optical fiber and transmitted to a user station. The Green Bank site has user stations at the 140 ft telescope, the GBT equipment room (at the Jansky Lab), and on the GBT. A fourth station will be installed in the Jansky lab for other telescope usage. At the user stations, the 500 and 10 MHz signals are recovered, resynchronized, and made available to the users. The 1PPS is also resynchronized to the 500 MHz and the 10 MHz. A Round Trip Phase Measurement Module (RTPM)[6,7] as described in Sect. 5.4 is provided to measure the changes in phase of the signal for each user station.

5.2. IRIG and 1PPS synchronization

The IRIG and 1PPS signals are generated in the timing center by a Trak Systems clock driven by the maser. The IRIG time code is synchronized to the 1PPS by the Trak clock. The clock is synchronized to GPS when the system is first powered up, and then it is left to free-run from the maser signal. These signals are distributed to the user instruments via fiber optics or coaxial cable, depending on the location of the user instruments.

The 1PPS signal generated is buffered and transmitted to various locations around the site. Because the 1PPS signal is used by the system backends and the tracking LO system for high precision timing, it is necessary to know the delays introduced into the signal by the delays in the transmission equipment and the transmission medium itself. This is accomplished by using a return cable to loop the signal from each location back to a multiplexer at the site timing center. The originating signal is also supplied to the multiplexer. Using this multiplexer, each signal can be applied to the gates of a counter, and the delay between the pulses may be measured and logged by the M&C system.

Each VME chassis in the system is fitted with a BC635 (See Sect. 2.4) card that processes the IRIG signal and extracts the timing information. The card uses interrupts to signal the computers in the VME chassis each 20 ms. The system controller in the chassis latches the current time into the memory. This time is then read out by the other processors on the bus. This timing has 100ns resolution. Accuracy is better than 1 microsecond.

The software driving this card is slated for upgrades; with the current software, if the IRIG signal goes berserk (which happens if the signal from the timing center jumps, for whatever reason), the system time happily follows along. A more graceful degradation of synchronization is desired. This is important for many reasons, not the least of which is that the conversion to Azimuth-Elevation servo coordinates from commonly used astronomical coordinate systems depends on knowing the correct time.

The current plan is to integrate NTP into the software driven by this card. This will allow a certain amount of filtering and validation to be performed on the time returned from the card, and will prevent the system time from taking a large step.

5.3. Unix workstation synchronization

Network Time Protocol is used to synchronize the Unix hosts with the IRIG time. A VME computer that is integrated with the site time equipment is used as a time server for the site. This time server is locked to the IRIG signal, and so is a stable time source for the site computers.

If the time source is disabled, the internal clocks of the computers will take over and continue to provide the time to the computer. Once the NTP establishes that the site time server is back in operation, it will sync back up to it automatically.

5.4. Round Trip Phase Monitors (RTPM)

The original design of the RTPM was developed as part of the VLBA project. This description is just a summary of the information available about the design and performance of the RTPM. See Weimer[6] for more information.

The purpose of the RTPM is to measure the changes in the delay of the signals transmitted from the timing center to the telescope and the equipment room. The delay must be known so that the phase differences in the signals can be accounted for. The measurement is made by reflecting some of the 500 MHz signal back to the timing center with sidebands 1.953 KHz away, which is mixed with a 500MHz + 2.083 KHz signal at the timing center. This results in a 130.2 Hz signal that is compared with a locally generated reference, and exhibits the same phase change as the original 500.001953 MHz signal. The measurement is repeated and logged by the M&C system every 10 seconds.

6. PROGRESS TOWARDS IMPLEMENTATION

The GBT is still under construction. It was fortunate that the engineers resisted the temptation to buy the control system computers early in the life of the construction project, as an entire generation of computers has come and gone since then.

All the VME based real-time hardware is in place, some of which have unfortunately been placed on end-of-life status by the vendor already. The general purpose computers have not been purchased, as we are waiting until we need to start integrating them into the final system.

This summer should see the completion of all the computer acquisition and integration, as the telescope should be turned over to NRAO for out-fitting by the beginning of 1999. We procured our network switches last year, and we have been getting experience on them. The site LAN has also undergone some improvement since last year. The link to the Internet has been upgraded from a 56K baud link to a 768Kb Frame Relay link. The internal network has had Ethernet switches added, and 100Mbit Ethernet has been deployed to link the switches.

7. CONCLUSION

We believe that we have designed the control system to allow the telescope to live up to its full potential. We have been conservative in our application of computer power to each part of the telescope.

There is substantial room for growth in most of the subsystems. Each VME chassis has been specified with many spare slots, and each SBC is running at only a fraction of its capacity. Each of the components is more or less independent of the others. The real-time systems can be upgraded one by one, if the need arises.

The telescope network is key to the control system performance. The switched Ethernet allows us to keep the utilization, and hence the latency, under control, and thereby allow more data to flow over the network. Using the Ethernet protocol keeps things simple. We can still cosider any of the higher-speed options, because our network switches are ready to accept higher-performance network links, like Gigabit Ethernet and ATM.

ACKNOWLEDGMENTS

This GBT M&C hardware design is the product of a very good group of people. Several that have had a particularly big hand in the hardware design are Joe Brandt, Ed Childers, and Tim Weadon, all of NRAO Green Bank.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

REFERENCES

1. D. Mills, "Internet time synchronization: the network time protocol.," *IEEE Trans. Communications COM-39* **10 (October)**, pp. 1482–1493, 1991.
2. W. D. Cotton, D. Tody, and W. D. Pence, "Binary table extensions to FITS.," *Astronomy and Astrophysics Supplement Series* **113**, pp. 159–166, 1995.
3. A. Benett, J. Brandt, M. Clark, and J. Ford, "Monitor/control hardware design review," Tech. Rep. GBT Archive MC022, National Radio Astronomy Observatory, 1996.
4. B. Clark, "Monitor and control bus at VLBA stations," Tech. Rep. A550001N001, National Radio Astronomy Observatory, 1984.

5. D. Weber, B. Clark, and L. R. D'Addario, "Monitor and control standard interface," Tech. Rep. A550001N002-A, National Radio Astronomy Observatory, 1985.
6. R. Weimer, "Round trip phase measurement module," Tech. Rep. 23, National Radio Astronomy Observatory, 1991.
7. L. R. D'Addario and M. Stennes, "Transmission of timing reference to sub-picosecond precision over optical fiber," in *Advanced Technology MMW, Radio, and Terahertz Telescopes*, T. G. Phillips, ed., *Proc. SPIE* **3357**, 1998.