FPGA Cluster Computing in the ETA Radio Telescope

C. Patterson, B. Martin, S. Ellingson
Dept. of Electrical & Computer Engineering
Virginia Tech
Blacksburg VA 24061
{cdp, bm92, ellingson}@vt.edu

J. Simonetti, S. Cutchin
Dept. of Physics
Virginia Tech
Blacksburg VA 24061
{jhs, scutchin}@vt.edu

Abstract

Array-based, direct-sampling radio telescopes have computational and communication requirements unsuited to conventional computer and cluster architectures. Synchronization must be strictly maintained across a large number of parallel data streams, from A/D conversion, through operations such as beamforming, to dataset recording. FPGAs supporting multi-gigabit serial I/O are ideally suited to this application. We describe a recently-constructed radio telescope called ETA having all-sky observing capability for detecting low frequency pulses from transient events such as gamma ray bursts and exploding primordial black holes. Signals from 24 dipole antennas are processed by a tiered arrangement of 28 commercial FPGA boards and 4 PCs with FPGA-based data acquisition cards, connected with custom I/O adapter boards supporting InfiniBand and LVDS physical links. ETA is designed for unattended operation, allowing configuration and recording to be controlled remotely.

1 Introduction

The Eight-meter-wavelength Transient Array (ETA) is a new radio telescope designed to observe a variety of postulated but as-yet undetected astrophysical phenomena which are suspected to produce single pulses detectable at relatively long wavelengths. Potential sources for these pulses include the self-annihilation of primordial black holes (PBHs) and prompt emission associated with gamma ray bursts (GRBs). PBHs are postulated to arise from density fluctuations in the early universe rather than the gravitational collapse of a star, and may be of any size. If black holes evaporate as suggested by specific combinations of general relativity and quantum mechanics, those with masses below about 10^{14} g should be approaching a state of runaway evaporation, terminating in a massive burst of radiation [18]. A number of models have been proposed to explain short-duration GRBs, including the merger of closely-separated compact objects such as neutron stars and/or black holes. The formation of a single black hole would release an immense amount of energy over a few seconds [14]. Prompt radio emission from GRBs would be very useful in pinning down the physics of the bursts, the nature of the progenitor object, and possibly the medium in which it occurs. Recently, a dispersed pulse of duration < 5 ms was detected during the analysis of archival pulsar survey data [5]. The brightness (30-Jy) and singularity of the burst suggest the source was not a rotating neutron star. Although pulses can be quite strong by astronomical standards, they are difficult to detect due to their transient and unpredictable nature, and the narrow field of view provided by existing telescopes.

ETA, in contrast, is designed to provide roughly uniform (albeit very modest) sensitivity over most of the visible sky, all the time. The complete array consists of 12 dual-polarized dipole-like elements (i.e., 24 radio frequency inputs) at the Pisgah Astronomical Research Institute (PARI), located in a rural mountainous region of Western North Carolina ($35^\circ\ 11.98'$ N, $82^\circ\ 52.28'$ W). Each dipole is individually instrumented and digitized. The digital signals are combined to form fixed "patrol beams" which cover the sky, and the output of each beam is searched for the unique time-frequency signature expected from short pulses which have been dispersed by the ionized interstellar medium. ETA has the computational resources and communication bandwidth to implement up to eight beamforming datapaths across all antennas, permitting it to operate as eight independent telescopes simultaneously recording data from different regions of the sky. A patrol beam can be quickly reoriented by updating its beamforming coefficients. Additional information about ETA and its science objectives are available at the project web site [10].

This paper is organized as follows. Section 2 describes the ETA architecture, with particular emphasis on the digital back-end. Section 3 discusses system validation issues such as determining bit error rates and ensuring channel synchro-

1-4244-1472-5/07/$25.00 © 2007 IEEE

Figure 1. ETA's ten-element core array. One antenna stand is in the center, and the remaining nine stands form a circle of radius 8 m.

Figure 2. A close-up view of an antenna stand. Each of the two dipoles is connected to an active balun (dipole-to-coaxial converter and preamplifier) located inside the PVC mast.

nization. Current status and conclusions are summarized in Sections 4 and 5.

2 System Design

ETA is designed to operate in the range 29–47 MHz, which is a response to a number of factors. First, some astrophysical theories suggest the possibility of strong emission by the sources of interest in the HF and lower VHF bands, limited at the low end by the increasing opacity of the ionosphere to wavelengths longer than about 20 m (15 MHz). Useable spectrum is further limited by the presence of strong interfering man-made signals below about 30 MHz (e.g., international shortwave broadcasting and Citizens Band (CB) radio) and above about 50 MHz (e.g., broadcast television), which make it difficult to observe productively outside this range. At these frequencies, however, the ubiquitous Galactic synchrotron emission is extraordinarily strong and can be the dominant source of noise in the observation [17].

The ETA system receives RF input from 24 dipole-like antennas mounted on 12 stands. Figure 1 shows the ten-stand core array; not shown are a pair of outrigger stands located several hundred feet away to the north and east. Each stand supports two orthogonally-polarized, dipole-like elements shown in Figure 2. Figure 3 is a snapshot of typical spectrum seen by the antenna. Buried coaxial cable connects antenna outputs to analog receivers that amplify and filter the signals. The digital signal processing and data recording are the focus of this paper, and are shown in Figure 4.

The architecture consists of three tiers: receiver nodes, FPGA cluster nodes, and data acquisition nodes. At each level, data must be synchronized and processed in parallel. The system receives antenna signals through the receiver nodes consisting of twelve Altera Stratix DSP development boards, referred to as S25s [7]. Each S25 board performs A/D conversion and digital filtering for two dipole inputs, and provides a source-synchronous stream of data to a cluster node. The cluster nodes are Xilinx ML310 FPGA boards, referred to as ML310s [12], and form twelve outer and four inner nodes networked together. Cluster nodes process, combine, and prepare data for recording on the four Dell SC430 server-class PCs. Stored data may be analyzed while the system is not acquiring data, or archived to tape. A fifth PC controls and monitors system functions. The following sections describe each level of the system in more detail.

2.1 Receiver Nodes

The main functions of the receiver nodes are to digitize, downconvert and channelize (filter to a narrower bandwidth) the analog antenna input. Since designing custom hardware to perform this function would be costly and time consuming, this functionality was implemented on commercially available Altera S25 boards. This board was chosen because it contains the necessary hardware elements for the design, and was familiar to the design team. Each S25 board contains two 12-bit 125 MSPS A/Ds, allowing all 24 analog input signals to be digitized with 12 boards.

To maintain synchronization across all receiver nodes, one board generates the clock and reset signals for the other S25 boards. A counter, synchronous across all S25 boards, is encoded with each time sample. This counter is used in the cluster nodes to align data streams from different S25s. Two antenna streams are combined into a single stream and transmitted to a cluster outer node at 30 MB/s. The combined stream consists of a 60 MHz clock, four bits of source synchronous data, and a counter bit. These LVDS signals are transmitted over Precision Interconnect's medical grade "Blue Ribbon" brand coaxial cable [3] with MICTOR connectors to help maximize noise immunity and minimize radio frequency interference (RFI) generation.

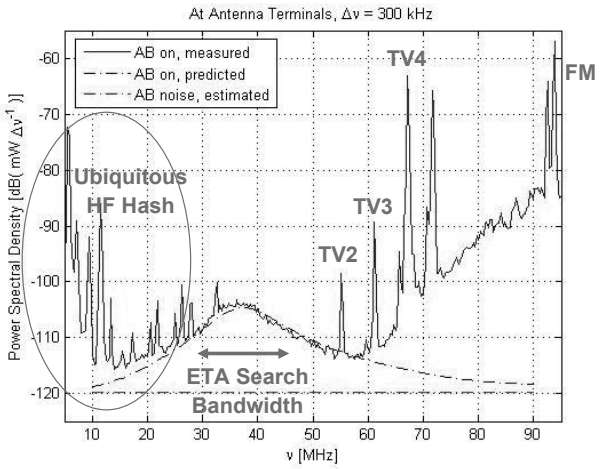

Figure 3. A typical power spectrum measurement for an ETA dipole's active balun (AB). Over the 29–47 MHz range, sensitivity is limited only by Galactic noise (the upper dashed line) and not AB noise (the lower dashed line).

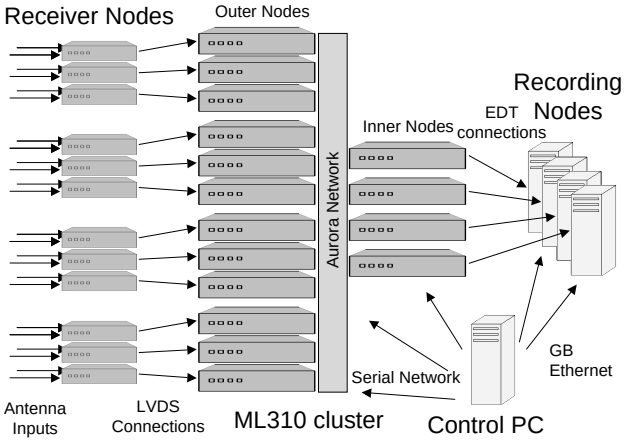

Figure 4. ETA's digital front-end and back-end.

2.2 FPGA Cluster

The FPGA computing cluster consists of 16 ML310 boards interconnected with 1X InfiniBand cable assemblies supporting a 2.5 GHz signaling speed. Like the S25 board, the ML310 was chosen because of it was familiar to the design team and provides sufficient hardware resources and connectivity. Cluster nodes are divided between 12 outer nodes, which connect directly to the receiver nodes, and 4 inner nodes, which connect directly to the PCs. The FPGA cluster provides a versatile, synchronous, high speed network for merging antenna streams to support operations such as FFTs, beamforming and RFI mitigation, and reformatting either raw or reduced data for transfer to the PC recording nodes. In fact the FPGA cluster is as much a router as a computing engine. The cluster rack is shown in Figure 5.

The ML310 board shown in Figure 6 has a Xilinx 2VP30 Virtex-II Pro FPGA, 256 MB DDR DIMM, 512 MB CompactFlash card, standard PC ports, and personality modules interfaces for RocketIO and LVDS access. Custom adapter boards were designed for the personality module interfaces. A four-layer adapter board provides InfiniBand HSSDC2 connector access to the FPGA's eight RocketIO transceivers. These connections implement the lightweight Aurora link-layer, point-to-point protocol [1]. A six-layer adapter board provides a MICTOR connector interface to an S25 receiver node, and a source synchronous, 16-bit parallel LVDS cable interface to a PC node's EDT PCI CDa LVDS/600E data acquisition card [9]. Each ML310 also communicates with the control PC through a UART, allowing all cluster nodes to be configured and queried remotely.

Routing in this architecture is simplified since the data flows only from the outer to the inner nodes. Outer nodes are responsible for a large portion of the beamforming computation, while the inner nodes synchronize and combine the signals produced from the outer nodes. As illustrated in Figure 7, each inner node receives an input stream from up to six outer nodes, combines them, and outputs the result to a PC. The upper and lower halves of the FPGA and PC cluster are independent and process antenna input from the north-south and east-west polarizations respectively.

ETA uses all three forms of clocking methodologies: system-synchronous within receiver, inner and outer nodes; source-synchronous for data transmission between S25s and ML310s, and between ML310s and the PC's data acquisition cards; and self-synchronous for transfers between outer and inner nodes [16]. The ML310's receiver node interface uses the 60 MHz clock sourced by the S25 to synchronize data transfers. All other inner and outer node modules use the ML310 board's 125 MHz reference clock, which is divided by 2 for the beamforming datapath, and multiplied by 20 in the RocketIO transceivers to obtain a bit-rate clock for the serializer and deserializer. The transmitted signal encodes the data such that the clock can be recovered by the receiver. Brief resynchronizations referred to as clock corrections occur periodically, necessitating buffers on the transmitters and receivers.

Raw data collection mode allows the output from eight of the S25 boards to pass directly to the PCs. Each outer node receives data from two antennas at a combined rate of 30 MB/s, reformats the data, and passes it to an inner node using the Raw Data path shown in Figure 8. The in-

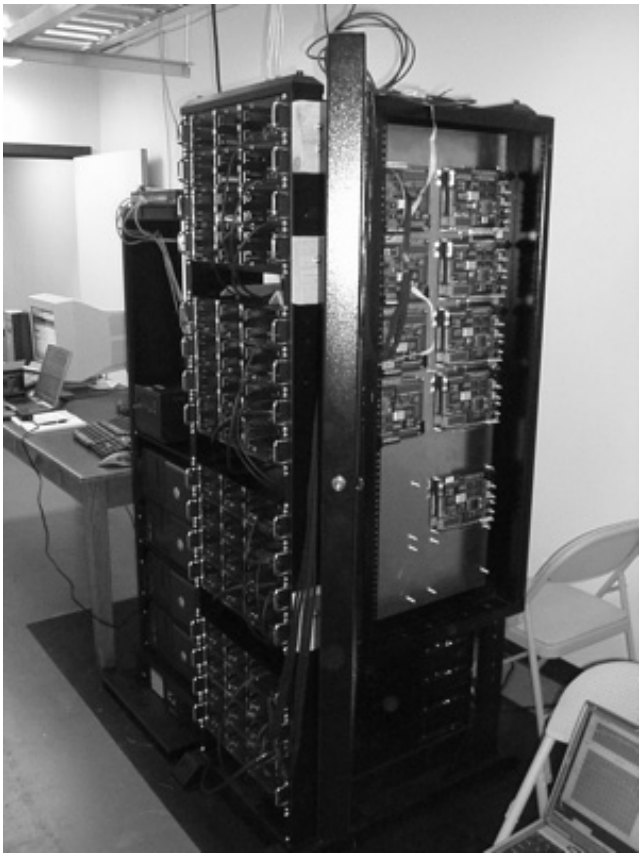

Figure 5. Inside ETA's equipment hut. The 16-node FPGA cluster occupies the middle rack. To the left is the PC cluster and tape drives, and to the right are the receiver nodes mounted inside a cabinet.

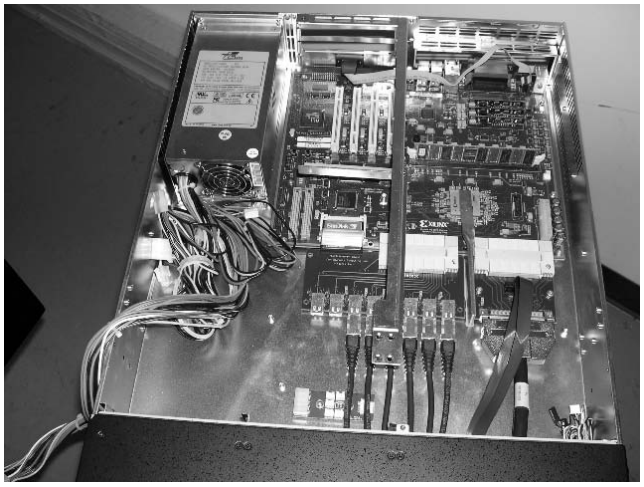

Figure 6. ML310 board and adapters. The ML310's form factor allows mounting inside PC ATX motherboard cases. This inner node has six InfiniBand cables connected to the adapter board on the bottom left of the ML310. The adapter board to the right provides a MICTOR connector to a 3 m "Blue Ribbon" brand coaxial cable coming from an S25 receiver node, and an Amphenol 80-pin right angle connector to a 7 ft black cable going to an EDT PCI CDa data acquisition card in a PC node. Between the two adapter boards is a copper bar conducting heat from the FPGA to the bottom of the case.

ner nodes collect data from two outer nodes, combine and format the input streams for storage on the PC cluster. In this mode each of the four PCs receives and store data at 60 MB/s, giving the system an aggregate recording rate of 240 MB/s. Data collections usually occur for roughly one hour, generating an 800 GB dataset. After offloading the data to 400 GB LTO-3 tapes [11], another acquisition can begin.

The basic beamforming mode uses all 12 outer nodes, combining multiple streams of data to form patrol beams for improved sensitivity. Although the sky could be tessellated with up to 12 independent beams, only 8 beams are implemented due to hardware resource limits, and poor sensitivity (because of antenna pattern roll-off) near the horizon. Each outer node again receives data from two antennas at a rate of 30 MB/s, or 360 MB/s aggregate. The outer nodes multiply each antenna input stream with the corresponding beamforming coefficients and sum the results to produce eight single-pol beams. Four of the beams are sent to one inner node and the other beams are sent to an adjacent inner node as illustrated in Figure 8. At this stage each Aurora connection is transmitting 120 MB/s for an aggregate bandwidth of 2.88 GB/s to the inner nodes. Each connection has an available bandwidth of 240 MB/s (5.76 GB/s aggregate), increasable to 360 MB/s (8.64 GB/s aggregate) by selecting a faster RocketIO reference clock although the bit error rate (BER) will increase. Each inner node combines a four-beam input from six outer nodes. The six input streams are synchronized, with corresponding samples summed, shifted and rounded (to avoid a DC bias) before transmission to the PCs. This reduction allows each PC to record four single-pol beams at 60 MB/s, or an aggregate 240 MB/s. For each polarization, one PC contains data for beams 1 to 4 and the other PC contains data for beams 5 to 8.

Since propagation speed through the ionized interstellar medium varies with the frequency of the wave, pulse dispersion or smearing results. In order to dedisperse the signal, the observing band is split into narrow channels and the detected signal from each channel is delayed by a different amount before summation to obtain the total power sig-

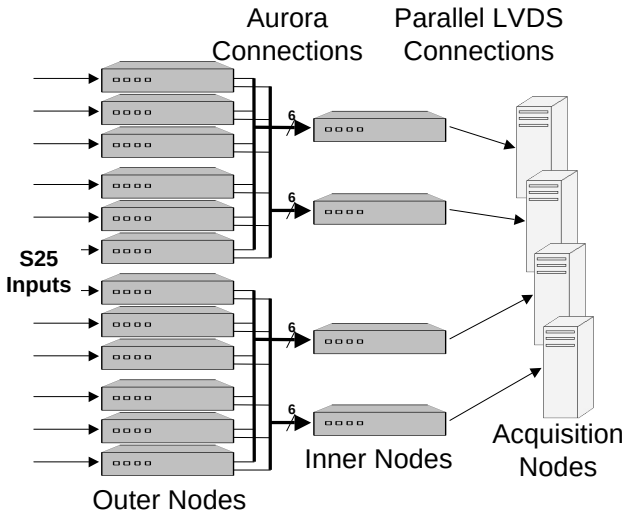

Figure 7. Aurora and LVDS physical links.

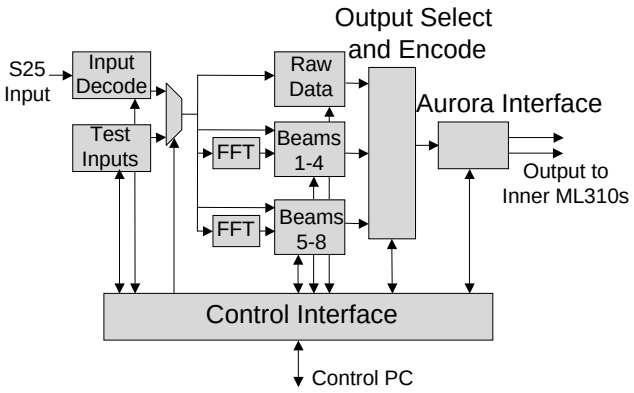

Figure 8. ML310 outer node implementation.

nal. Even on a fast workstation, dedispersion requires over twenty-four hours for each hour of data collected due to the large number of FFTs. Because it is difficult to completely automate pulse detection, the main focus is on speeding up dataset post-processing. The FPGA cluster can optionally perform a real-time, 1024-point FFT of the 18 MHz-wide passband into 1024 17.6 KHz-wide spectral channels with 66 μs integration. Beamforming is applied on a per-channel basis, leaving the software to apply a range of dispersion measures.

The outer nodes buffer 1024 samples which are streamed into the FFT block. FFT computations requires 1024 clock cycles, after which the output is streamed to the beamforming stage as indicated in Figure 9. Each beam must multiply the 1024 frequency domain values with corresponding complex coefficients. For the two input streams and the eight beams formed, each outer node contains a table with 16384 coefficients loaded through the UART. After the multiplication, corresponding beams are summed and multiplexed for transmission to the inner nodes as in the basic beamforming mode. A useful feature of the FFT mode is that frequency bins can be selectively included or removed, enabling the system to record only the frequencies of interest and extend viewing time.

As shown in Figure 10, the main functions of the inner node are to synchronize the streams from the outer nodes, combine them, and format the output for storage on a PC. The streams are synchronized with a bit field indicating the start of a vector. Corresponding entries from a set of vectors are summed and shifted to extract the appropriate bits. A 32-bit value encodes the 4-bit vector start flag and either one or two complex 14-bit samples, depending on the mode. Two 32-bit streams are multiplexed and output to a PC. This allows a sustained 60 MB/s transfer to disk of either four antenna streams or beams, and a system maximum of sixteen antenna streams or single-pol beams. Sixteen (i.e. eight dual-pol) beams allows a majority of the viewable sky to be observed with maximum sensitivity.

The outer nodes utilize approximately 20,000 (70%) of the logic elements, 60 (44%) of the BRAMs, and 88 (65%) of the multipliers in the 2VP30 FPGA. By contrast, less than 28 (5%) of the IOBs are required. Both inner and outer nodes use six of the eight available gigabit transceivers, necessitating a lightweight protocol such as Aurora. Excluding the FFT, each outer node performs 2 GMAC/s while beamforming is active, for a system aggregate of 24 GMAC/s.

2.3 PC Cluster

Four cost-effective Dell SC430 servers running Red Hat Enterprise Linux are used to record the data streams. Each SC430 is equipped with a 2.8 GHz Pentium D (dual core) processor, 1GB of RAM, 300 GB and 147 GB SCSI drives, a 250 GB SATA drive, and an EDT PCI CDa data acquisition card for connecting to the FPGA cluster inner nodes. The three hard drives are software-configured as a RAID Level 0 partition with an ext2 (non-journaling) file system. After extensive testing it was determined that the system could continuously record at 60 MB/s for one hour without error. The data collected is partitioned into 200 files, each containing 1 GB.

Following acquisition, the 200 GB datasets are written to tape for transport from PARI. Hot swapping of recording media is not useful since the ETA system is unattended, and tape jukeboxes are prohibitively expensive. Two of the SC430s have 400 GB (uncompressed) Quantum LTO-3 external tape drives, and eventually all SC430s will have tape drives. The LTO-3 drives automatically apply LZ-based compression to records that decrease in size, but compression is ineffective on ETA datasets. It may be possible to

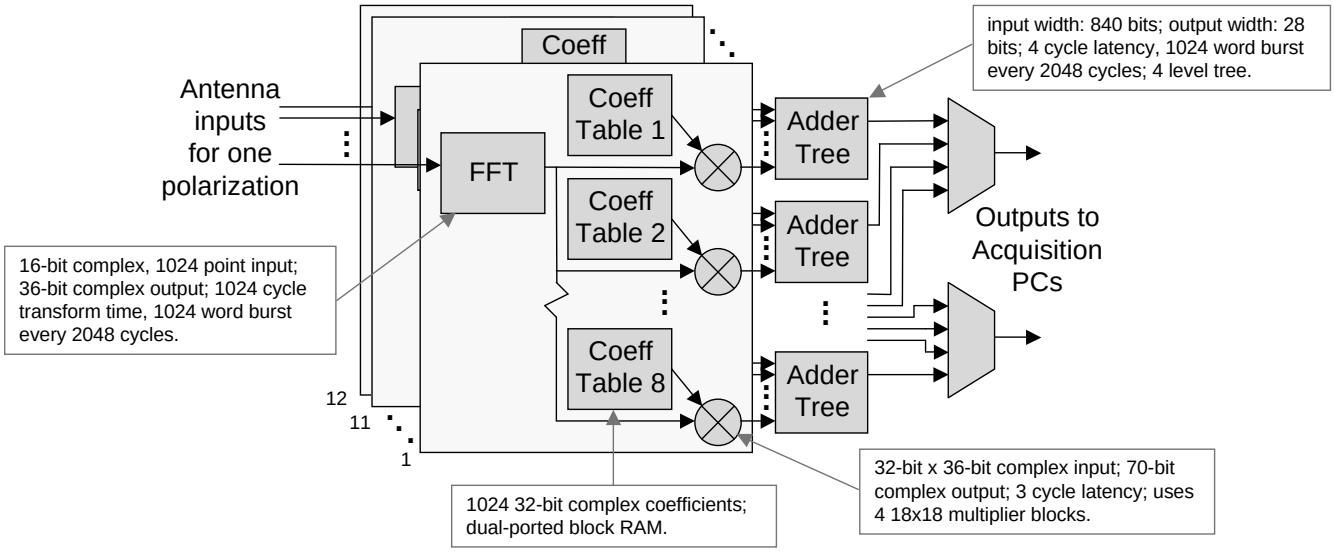

Figure 9. Beamforming across spectral channels.

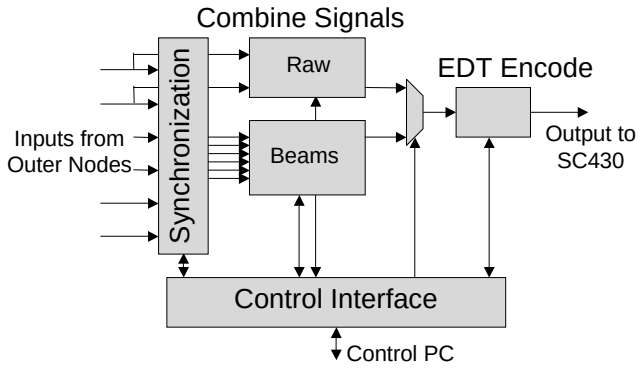

Figure 10. ML310 inner node implementation.

stream data directly to tape without first writing to disk since the LTO-3 drives supports sustained transfer rates of up to 68 MB/s, and the tape automatically speeds up or slows down to match the incoming data rate. LTO-4 tape drives are now available, which can record 800 GB (uncompressed) at up to 120 MB/s.

A Fedora Linux-based control PC is accessible through the Internet and allows external control and observation of the FPGA and PC clusters. A Gigabit Ethernet switch connects the control PC and the SC430s. The control PC accesses each ML310 through two Lava Link Octopus-550 serial PCI cards [15]. Each card has eight RS-232 ports giving a total of 16 serial connections. Scripts run on the control PC start and terminate acquisitions, download new configurations and coefficients to the ML310 nodes, and record system status information. Test modes are also available to check system integrity.

3 System Validation

To enable BER testing, synthetic data can be sourced either from the S25 or ML310 nodes. The data (typically counter or numerically-controlled oscillator output) are verified before and after every communication link in the system, and any errors are tallied on each ML310. The serial connection between each ML310 and the control PC allows for internal state to be read or written. C programs convert user-generated text files to a bit string defining the operating mode and data such as beamforming coefficients, or from a bit string to text files describing the internal state. Readback also returns the status of the Aurora links, indicates whether buffers have ever overflowed, and confirms the beamforming coefficients were written correctly. Error counts may be queried after a test to facilitate on-site or remote error rate testing and diagnosis, or after a data acquisition to check that the integrity of the data collected. In addition, synthetic data tests may be invoked at the end of a data collection script. Through the detailed information received from the ML310s, problems can be quickly isolated to specific connections, cables or boards. These scan paths allowed each half of the FPGA cluster to be connected and tested in less than a day.

ETA's systolic architecture makes data retransmission difficult, and error-correcting codes would be a significant overhead given the number of communication links. However, digital system errors are unacceptable even though

RFI is a much larger concern. All data errors originating in the ETA system have occurred in the 2.5 Gbps Aurora channels. Synchronization errors (manifesting as detectable buffer overflows) have never been observed. Careful design and signal integrity analysis of the adapter boards, and transmitter pre-emphasis and differential voltage swing adjustments, have resulted in observed BER less than the InfiniBand maximum of 10^{-12}. In a 200 GB dataset, the number of data bit errors is generally 0, with a maximum of 2 observed. The main source of errors are the connections between a cable and an adapter board. Outer nodes transmit only two channels and make each channel available on three of the adapter board connectors to facilitate bypassing a connection generating errors. Inner nodes use all six input channels, and errors are usually remedied by reseating cable connections. Enabling the RocketIO transceiver's CRC feature and checking if any CRC errors have occurred during an observation should address inter-board data integrity concerns.

4 Current Status

ETA development began in August 2005, with a demonstration of direct sampling from the first four dipoles three months later. One commissioning test was confirmation of the diurnal variation shown in Figure 11. The first 200 GB raw data acquisition through S25 and ML310 nodes to a PC occurred in April 2006. A lightning strike in July 2006 damaged preamplifiers, but most have since been repaired. Over 30 hours (6 TB) of raw data has been collected by the end of February 2007. FFT beamforming mode was added in July 2007. Two independent dedispersion software workbenches have been developed, and we see a role for undergraduate students in applying these tools to the dataset archives. An analysis of how frequently pulses are detected and their dispersion measures should provide valuable insights to the distribution of the progenitor objects, some of which may be a component of dark matter [2].

Mitigation of RFI is an important consideration for any instrument operating at low frequencies. Although dedispersion tends to suppress narrowband RFI, and FFT bins may be selectively blanked to reject noise in certain frequency ranges, an additional approach to RFI rejection is anticoincidence. A portable second instrument (ETA2) is being assembled at a site more than 300 km from PARI. ETA2 still sees the same section of the sky as ETA but is not affected by ETA's local RFI. Although each station records time accurately via GPS, there is no beamforming across stations since data collection is not synchronized during the observation. During dataset post-processing, pulses detected simultaneously at both ETA and ETA2 would rule out local RFI sources.

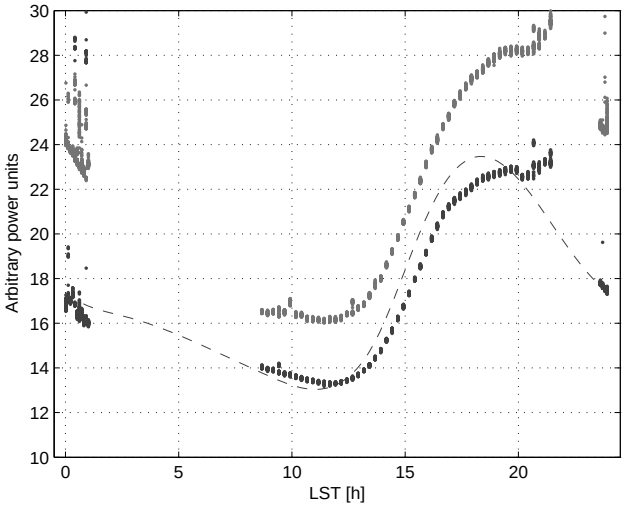

Figure 11. Diurnal variation in a 5 MHz bandwidth centered at 45 MHz. The dotted line is predicted power due to galactic noise, while the black and gray scatters are measured power from a core array dipole and the east outrigger respectively.

5 Conclusions

FPGAs with integrated multi-gigabit serial transceivers are ideal platforms for the signal processing requirements of radio telescope arrays. Synchronous streaming computations may be implemented over many boards without requiring system-wide clock distribution. As in ETA, the physical communication topology can be tailored to the dataflow requirements, simplifying the interfaces and development effort. In contrast, software processors and their communication protocols are difficult to use in this environment due to variable latencies.

Data recording is the bottleneck for the ETA back-end. One of the biggest design challenges was finding a low-cost PC interface and configuration to record data continuously at 60 MB/s for one hour. The first attempt used a Gigabit Ethernet link between each ML310 and PC, but PCI driver overheads overwhelm the ML310 FPGA's 300 MHz PowerPC processor. Positive developments for data streaming and recording performance are multi-core processors allowing better overlap of computation and I/O, and high capacity solid state disks that avoid seek latencies.

Our goal of obtaining results within the first year necessitated the use of commercial-off-the-shelf (COTS) FPGA boards. This approach distinguishes ETA from many other radio telescope arrays using custom boards. The design of complex, high speed boards can squander FPGA development time and cost advantages. Although a single board

integrating the receiver and beamforming functions is desirable, the effort required to design a custom PCB would be excessive even if no respins were required. The low-cost evaluation boards meet the ETA system specifications remarkably well, and all interfacing is accomplished with standard cables and simple adapter boards containing only connectors and traces.

An additional benefit of COTS is straightforward upgrades to new FPGA families. ETA2 can use a newer generation of FPGA evaluation boards: the ML410 contains a Virtex-4 FX60 while retaining the PC ATX form factor and personality module interfaces [13], and the Stratix-II DSP Development Kit contains a 2S60 device and preserves the MICTOR interface [8]. Evaluation boards further extend the low cost and rapid development virtues of FPGAs, with the new boards costing roughly the same as their predecessors despite having significantly improved logic capacity and performance. The university price for these boards usually reflects a donation or discount of the FPGA and configuration components.

ETA invites comparisons with BEE2, a general-purpose, scalable, FPGA-based system for large-scale emulation and DSP [4]. Both suit stream-based computations, although a BEE2 board is significantly more powerful with 5 2VP70 FPGAs, 20 GB of DDR2 DRAM, and 18 10-Gbps serial interfaces allowing connections to other BEE2 boards, to an InfiniBand or 10 Gbps Ethernet packet-switched network, or to analog interfaces. On the other hand, ETA's use of evaluation boards is cost effective for the size of the system, and also allows quicker migration to new FPGA families. A larger scale version of ETA might justify more powerful compute platforms such as BEE2, or servers with FPGA modules plugged into some of the processor slots [6]. It may also be feasible to have FPGAs stream data directly to a parallel set of SATA disks. Regardless of whether FPGAs and/or multicore processors are used, built-in interfaces for high speed serial protocols are as essential as computational power.

6 Acknowledgments

This work was supported by the National Science Foundation under Grant No. AST-0504677, a SCHEV ETF grant, and by the Pisgah Astronomical Research Institute. Graduate student Vivek Venugopal performed data streaming tests on the Aurora links and data acquisition PCs, and Michael Kavic is investigating theoretical connections between PBH explosions and large extra dimensions. PARI personnel constructed the antenna masts and underground coaxial cable system. We are grateful for early access to the ML310 board and the assistance of Vince Eck, Rick LeBlanc, Punit Kalra and Saeid Mousavi in Xilinx's Systems Engineering Group.

References

[1] http://www.xilinx.com/aurora/aurora_technology.htm.
[2] D. Blais, C. Kiefer, and D. Polarski. Can primordial black holes be a significant part of dark matter? *Physics Letters B*, 535:11, 2002.
[3] http://precisionint.com.
[4] C. Chang, J. Wawrzynek, and R. W. Brodersen. BEE2: a high-end reconfigurable computing system. *IEEE Des. Test. Comput.*, 22(2):114–125, Mar. 2005.
[5] D.R. Lorimer, M. Bailes, M.A. McLaughlin, D.J. Narkevic, and F. Crawford. A bright millisecond radio burst of extragalactic origin. To appear in *Science*, 2007. Preprint available at http://www.arxiv.org/abs/0709.430.
[6] http://drccomputer.com/.
[7] http://altera.com/products/devkits/altera/kit-dsp_stratix.html.
[8] http://altera.com/products/devkits/altera/kit-dsp-2S60.html.
[9] http://www.edt.com/pcicda.html.
[10] ETA project website: http://www.ece.vt.edu/swe/eta.
[11] http://www.lto-technology.com.
[12] http://www.xilinx.com/products/boards/ml310/current/.
[13] http://www.xilinx.com/ml410-p/.
[14] R. Narayan, B. Paczynski, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *Astrophysical Journal*, 395:L83–L86, Aug. 1992.
[15] http://www.lavalink.com/index.php?id=232.
[16] http://www.xilinx.com/publications/books/serialio/serialio-book.pdf.
[17] S.W. Ellingson, J.H. Simonetti, and C.D. Patterson. Design and evaluation of an active antenna for a 29–47 MHz radio telescope array. *IEEE Trans. Antennas Propagat.*, 55(3):826–831, Mar. 2007.
[18] S.W. Hawking. Black hole explosions? *Nature*, 248:30–31, 1974.