BEE2: a multi-purpose computing platform for radio telescope signal processing applications

Robert Brodersen, Chen Chang, John Wawrzynek, Dan Werthimer
{rb,chenzh,johnw}@eecs.berkeley.edu, danw@ssl.berkeley.edu
EECS, University of California, Berkeley
Space Science Laboratory, SETI Institue

Abstract

We propose to design and build the BEE2 multi-purpose computing platform, which will offer a general FPGA-based hardware architecture and software design methodology that target a range of real-time radio telescope signal processing applications, such as wide band spectrometer and large antenna array correlator, while providing drastic reduction in overall cost and design time.

The Problem

Existing radio astronomy instrumentation is highly specialized, with custom complex dedicated instruments built for individual applications, such as spectroscopy, pulsar searching/timing, and SETI. Each instrument not only takes years (3~5 in practice) to design, construct, and debug, but also becomes obsolete rapidly as the electronics capabilities grow, especially in the realm of digital signal processing, where the performance doubles every 18 months (Moore's Law).

Traditionally to increase the collecting area, thus sensitivity of the radio telescope antenna, gigantic single dish antennas are constructed, such as the Arecibo telescope. However, due to the dominating cost of steel for constructing the telescope, Arecibo has remained to be the single largest collecting area telescope in the world for over 40 years. In light of the digital computing revolution in the past several decades, currently constructing and future large collecting area radio telescope designs (such as ATA and SKA) has been moving in favor of a large array (hundreds to thousands) of small diameter (6~12m) antennas, where the steel cost can be balanced with the electronics cost. The antennas can be physically spread cross large geographic area, providing extremely long and various length baselines, hence better angular resolution. Beamforming on antenna array also enable multiple currently observations to different regions of the sky. With the recent advancements in RF technology, the new telescope arrays are capable of observing up to 11 GHz of continuous bandwidth from near DC, rather than the traditional approach of narrow (up to a couple 100 MHz) bands around a few fixed RF center frequencies spread across from DC to 50 GHz.

Despite the numerous advantages of using the antenna array radio telescope, this approach poses significant signal processing challenges. To form proper images, all the signals from the antennas need to be correlated with each other, thus require $O\left(N^2\right)$ computing power. To achieve one square kilometer collecting area, over 4000 antennas are required (each 12m diameter), if the whole 11 GHz need to be correlated, the computational requirement is on the order of 10^{17} operations per second, which is over 10,000 times the capability of current fastest supercomputer, the Earth Simulator. Clearly, novel computing approach is required to achieve

the computation throughput at reasonable cost. Furthermore, radio telescopes are typically designed to operate well over 30 years. It is not necessary to invest all the electronics up front to meet the bandwidth requirements, but rather upgrade the electronics every several years and gradually increase the total bandwidth observable. The approach can achieve the best price-performance ratio overall, by leveraging the exponentially decreasing cost of semiconductor technology. However, each redesign of the system require significant engineering time and cost. Currently, with all specialized instrumentation, the NRE cost of the hardware and software design significantly overweight the actual cost of the electronic hardware. Therefore, a modular, flexible, and efficient computing architecture is necessary to best leverage the cost benefit provided by Moore's Law, while achieving the I/O bandwidth and real-time processing requirements.

The BEE2 Solution

The BEE2 computing platform is based on FPGA technology, which provides high computational density and reconfigurability. Although not as power or area efficient as ASIC for a given silicon area, FPGA designs does not require expensive CAD tools or the long design cycles of ASIC, nor does it need the million dollar custom masks set. On the other hand, state-of-the-art FPGA chips can cost a couple of thousand dollars, which are about 5 times more expensive than Intel microprocessors. However, with several hundreds dedicated 18-bit multipliers and design clock frequency of over 150MHz, a Xilinx Virtex 2 Pro 70 FPGA chip can performance up to 48 billion MAC (16 bit) operations per second, which is over 50 times more than the latest Intel Pentium 4 processors, hence making FPGA 10 times more cost effective than microprocessors. Furthermore, due to the microprocessor architecture, such as hardware controlled cache, bus-based I/O subsystem, and limited memory channels (up to 2 per Pentium 4 processor), it is extremely difficult to meet the real-time high throughput processing requirement from the radio astronomy applications. On the contrary, in the BEE2 system, each FPGA has 4 independent memory channels (aggregating to 12.8 GBps throughput), and 80 Gbps direct point-to-point I/O bandwidth to/from external source/sink.

The system architecture of BEE2 is analogous to NOW clusters, with work stations replaced by BEE2 modules, and Ethernet replaced by Infiniband. Each BEE2 module packs 5 Virtex 2 Pro 70 FPGA chips and 20 DDR2 DRAM DIMM modules on to a single 14" by 15" PCB. The number of processing components (FPGA and DRAM) is mainly limited by the board size, cost, and power budget of 250 watt, such that 10 of such blade boards can fit into a single 19" wide 8U rack mount chassis. Each BEE2 module provide up to 18 Infiniband 4X connectors (10 Gbps full duplex each). Any one of these can be connected to either the ADC frontend board from the radio antennas, or the commercial non-blocking packet switched Infiniband crossbar, which provide up to 144 Infiniband 4X connections in a 10U rack mount chassis, aggregating to 2.88 Tbps throughput with port-to-port latency under 1 microsecond. At the current price of less than $400 per port, the Infiniband switch has the best price-throughput ratio among all commercial network switch technologies.

BEE2 system enables a range of different radio astronomy applications to share a common pool of computing resource connected through the Infiniband switch. As shown in the diagram below, the overall signal processing are divided into two parts: the frontend individual antenna spectral channel formation and the backend radio astronomy applications. After the analog RF signal from each antenna is digitized, a packet of 4KB data is formed with a relative time stamp labeling the relative phase with respect to other antennas. The signal data packets are streamed

into a bank of frontend BEE2 processing modules, where the input IF band is divided into sub-channels implemented using polyphase filter banks. After which, the output spectral data is reordered by the channel and stored in a circular buffer, waiting for remote DMA access by the backend BEE2 computing modules through the Infiniband switch. Therefore the frontend spectral processing modules are completely decoupled from the backend computing.

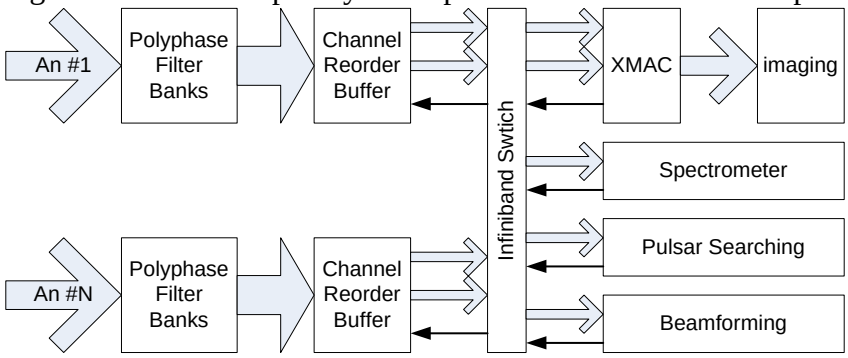

Multiple radio astronomy applications can be running simultaneously on the backend processing modules, such as spectrometer, correlator, pulsar timing/search, and beam forming. One or more applications can be present at one time, and the same BEE2 processing modules can be reconfigured in second to switch from one application to another. Using the Infiniband remote DMA access specification, all applications can be designed using the same data access interface. With identical processing modules on both front and backend computing, the system reliability can be greatly improved and system down time reduced by avoiding faulty processing modules.

In all, by standardizing and modularizing the signal processing module and interconnect architecture, the BEE2 system can be very cost effective computing solution to the radio astronomy applications. With only one digital processing board to be design, the BEE2 like system architecture can survive many generations of silicon technology upgrade, while minimally affecting the software and applications used. Using the Simulink based dataflow diagram design methodology, with proper interconnect and memory subsystem abstraction, the applications designed for the original BEE2 system only require simple recompilation to be upgraded to new BEE2 like hardware systems to take advantage of the new FPGA, memory, or interconnect technology.

Project Plan

The goal of the project is to demonstrate BEE2 as a cost efficient FPGA-based reconfigurable computing solution that addresses a range of radio telescope signal processing applications. The plan is to achieve the overall objective in three steps:

1. Design an extensible hardware architecture to best match the processing requirements.
2. Provide a convenient and efficient hardware abstraction layer over the physical resources.
3. Demonstrate the performance and design process of selected applications on the BEE2 system.

Step one has been well under way with the BEE2 board in the PCB layout phase, and expected to be fabricated in early Q4 of 2004. Two prototype BEE2 boards will be constructed and tested by the end of 2004. Any minor board revision will be combined with the initial batch production of 8~10 boards, which can take place as early as Jan 2005. Future production of BEE2 modules will be dependent on funding availability and application demands.

All demonstration applications will be designed using the existing BEE design flow, which uses Matlab/Simulink and Xilinx System Generator library. System level routing currently is

handled by the BEE router, which target specifically to the BEE hardware platform. To support the BEE2 hardware platform, additional stream interconnect library need to be constructed to provide technology independent abstraction over the physical connection links. In addition, stream RAM interface library is needed to provide convenient usage model of the external DRAM DIMMs. Inter-chip level partition is left to user, but tool flow enhancements are planed to facilitate stamping of identical or near identical subsystems, as well as relative placement instead of absolute locations.

Two demonstration applications have been selected for physical implementation: SETI spectrometer and ATA 32 antenna correlator. From initial analysis, each BEE2 modules can support up to 2 billion channel spectrometer with input bandwidth of 1 GHz dual polarization 4 bit I/Q, thus providing a spectral resolution of less than 0.5 Hz. 16 Gbps input signal enters the center FPGA, and then split into four 250 MHz 8 bit I/Q channels, one to each of the four corner FPGA chips. Each corner FPGA performs a 64K tap, 8K channel (decimation) polyphase filter bank, followed by a matrix transpose to reorder the signal by spectral channel, then fed into 64K point pipelined biplex FFT, and finally the spectrum power is calculated, and data points that exceed a preset threshold times higher than the local average are reported to a backend server, through Ethernet on the center FPGA. Each FPGA performs 37 GMult/s and 50.5 GAdd/s, and requires 5.8 Mbits on-chip RAM, 4GB off-chip DRAM and 4GBps DRAM bandwidth. All of the requirements are within the design spec of BEE2 system.

The 32 antenna ATA correlator is based on the FX correlator architecture, where the input 100MHz dual polarization 4 bit I/Q signal is first divided into 1024 channels using 8K tap polyphase filter bank, and then reordered according to spectral channel and stored on the off-chip DRAM as circular buffer. Each 100MHz IF requires one BEE2 module (F module) for the polyphase filter bank and channel reorder buffer. The cross MAC unit require 205 GMAC/s for all 32 antenna 100MHz IF band, which can be implement with a single BEE2 module (X module). The cross connection between the F and X module will be implemented through a single 24-port Infiniband switch, which provide 12 port to each F/X module. Through the remote DMA interface, the X module pulls the necessary time-stamped data from the F module. To support all 4 IF bands (100MHz each), four 24-port Infiniband switches can be used, or a single 96-port Infiniband switch.

Both of the demonstration will be implemented on the physical BEE2 hardware system, and tested using additional BEE2 modules emulating antenna signals from the ADC. Although the actual deployment of BEE2 system at the radio telescope site is beyond the scope of the project, if resources permit, the SETI spectrometer design will be physically tested at the Arecibo telescope.

Broader impact of the research

In addition to the two demonstration applications above, the BEE2 system can also performance well on a wide range of radio astronomy application, from the more traditional pulsar timing and searching, to emerging applications such as antenna array beamforming and real-time adaptive RFI identification and rejection. Besides radio astronomy applications, the BEE2 system can target various communication, DSP, and image processing applications, such as cognitive radio, hyperspectral image processing, autonomous vehicle navigation; as well as even applications traditionally left in the realm of microprocessor based compute clusters, such as bioinformatics and scientific computation.