

The CASPER Collaboration for High-Performance Open Source Digital Radio Astronomy Instrumentation

Dan Werthimer¹ and 257 CASPER Collaborators at 68 institutions

¹University of California, Berkeley, United States – danw@ssl.berkeley.edu

Abstract

The Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) has revolutionized the cost and time scale for development of high performance radio astronomy signal processing instrumentation. We present open source designs (hardware, gateware, libraries, and software tools) for a variety of flexible high bandwidth instruments, including correlators, beamformers, spectrometers, VLBI, pulsar timing and transient search machines. The collaboration relies on a small number of modular, connectible, upgradable, hardware components and platform independent software libraries that can be reused and scaled as hardware capabilities expand.

1 Introduction

Traditional radio astronomy instrumentation is highly specialized, with custom instruments being designed and built for individual applications using specialized boards, backplanes, interconnections and monitor and control software. Each instrument takes three to ten years to design, construct, and debug, and by the time it is deployed, it has usually been made obsolete by the Moore's Law growth in the electronics industry.

The constant redesign and upgrade of custom radio astronomy instrumentation is fueled by the ever increasing need for high-performance real-time DSP computing power for applications such as beam forming, imaging, RFI mitigation, and wide-band, high-resolution spectroscopy. The next generation of radio telescopes (e.g., Allen Telescope Array (ATA), Combined Array for Research in Millimeter-wave Astronomy (CARMA), Epoch of Reionization (EoR) arrays, Atacama Large Millimeter Array (ALMA), Square Kilometer Array (SKA)), are being designed and built using large numbers of small antennas. With computation scaling as (bandwidth) \times ($N_{\text{telescopes}}$)², DSP systems for these projects, especially those requiring several Gigahertz of continuous RF bandwidth over hundreds of physical antennas, can require peta-operations per second.

Our collaboration seeks to shorten the instrument development cycle by designing modular, upgradeable hardware and a generalized, scalable architecture for combining this hardware into a signal processing instrument. Using GPU's Field-Programmable Gate Arrays (FPGAs), FPGA-based chip-independent signal processing libraries, and packetized data routed through commercially available switches, our proposed architecture looks much like a Beowulf cluster - with reconfigurable, modular computing hardware in place of CPUs. This architecture uses a small number of easily replaceable and upgradeable hardware modules connected with as many identical modules as necessary to meet the computational requirements of an application, known colloquially as "computing by the yard." In addition to having the advantage of built-in tolerance for hardware failure, as data can be routed to other identically reconfigured modules, this architecture can use multicast switches to allow many commensal experiments to run simultaneously on the same FPGA/GPU/CPU compute cluster (see Figure 1).

2 Hardware Modularity

Hardware modularity requires that a small number of components with consistent interfaces be connectable with an arbitrary number of identical components to meet the computing needs of an application. Additionally, upgrading/revising a component should not change the way in which components are combined in the system. A modular system architecture can provide orders of magnitude reduction in overall cost and design time and will closely track the early adoption of state-of-the-art IC fabrication by FPGA vendors.

CASPER has developed a variety of inexpensive open source ADC boards that plug into the iBob, Roach I or Roach II boards. Current CASPER ADC boards range from a 64 input 12 bit 64 Msps board, to a single 3 Gsps 8 bit ADC board (two of the 3 Gsps boards can be interleaved for 6 Gsps). The collaboration is currently developing a 20Gsps ADC board.

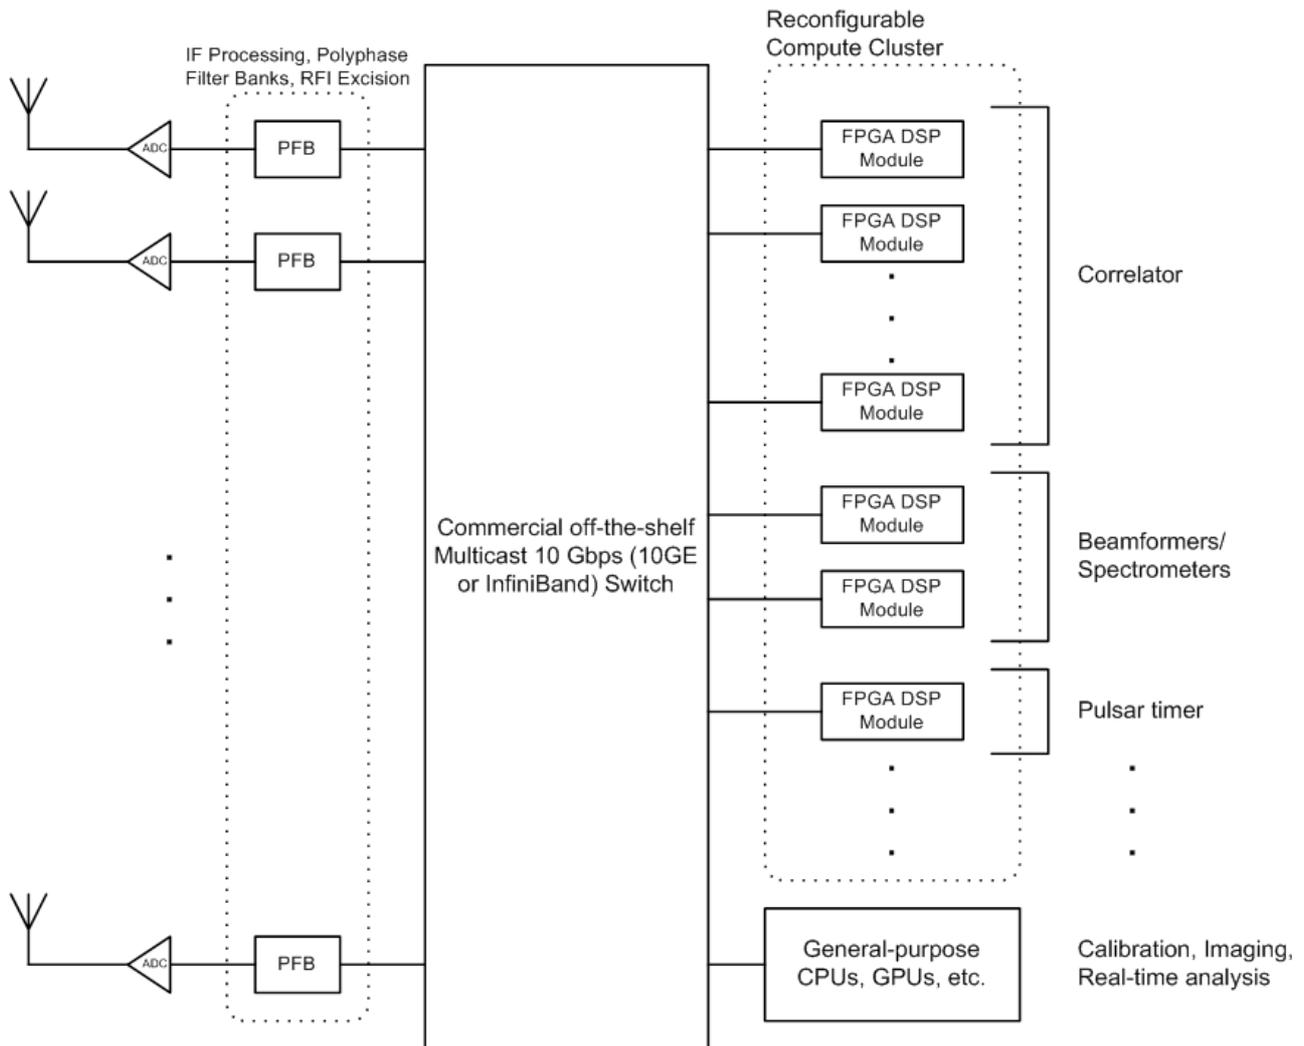


Figure 1: In this general architecture, RF, IF, or baseband signals are digitized, spectrally decomposed, packetized, and transmitted via 10 Gb ethernet protocol. Data is routed through a commercial multicast switch to an array of general purpose computational engines which can be dynamically partitioned for commensal applications.

3 Gateware and Hardware Reusability

The advantages of flexible, upgradeable hardware for radio astronomy signal processing cannot be realized without equally flexible signal processing libraries. Our group has designed a set of open-source libraries for the Simulink/Xilinx System Generator FPGA programming language [1]. These libraries abstract chip-specific components to provide high-level interfaces targeting a wide variety of devices. Signal processing blocks in these libraries are parametrized to scale up and down to arbitrary sizes, and to have selectable bit widths, latencies, and scaling. Though the design principles of parametrization and scalability have added complexity to the initial design of these libraries, it dramatically enhances their applicability and potential for longevity as hardware evolves. It also decreases testing time by allowing developers to debug scale models of systems that derive from the same parametrization code and are behaviorally similar to larger systems. Our current library has served as the starting point for a large number of instruments, and its flexibility has been demonstrated by the rapid development times achieved by these projects.

Our library provides parameterized implementations of a variety of signal processing algorithms useful to radio astronomy instrumentation, including: (1) digital down-conversion with a runtime-selectable mixing frequency and

digital filtering to allow flexibility and control in selecting passband shapes and adjusting fine delays (2) Polyphase Filter Banks (PFBs) for efficient implementation of a bank of evenly spaced, decimating FIR filters. Our core is parametrized to use selectable windowing functions, allowing adjustment of the out-of-band rejection and passband ripple/rolloff (3) a radix-2 biphase pipelined FFT capable of analyzing two independent, complex data streams using a fraction of the FPGA resources of commercial designs, including extensions for analyzing bandwidths higher than the native clock rate of an FPGA and for performing two real FFTs with each half of a biphase FFT, and (4) a cross-multiplier for use in our FX correlator architecture that computes all antenna cross-multiples within a frequency channel, and can handle multiple frequencies that are multiplexed into the core as dictated by processor bandwidth. Our library also includes a wide variety of useful data manipulation components, including data permuters, integrators, 10 GbE packetizers, summers, buffers, and serializers.



Figure 2: General purpose ROACH FPGA/CPU board, shown with a pair of dual channel 1 Gbps ADC boards

4 Spectrometers

Spectrometer applications have been a central target of the CASPER effort since its inception. Indeed the first application of CASPER hardware/gateway was a SETI spectrometer for use by NASA's Jet Propulsion Laboratory on Deep Space Network (DSN) station DSS-13. The application required analysis of a selectable 200MHz band at sub-2 Hz resolution for identification of anomalous narrow band emission. This 128-Million channel spectrometer built by CASPER was implemented on a first-generation ADC and IBOB module connected to one BEE2 module.

The technology and expertise garnered from the original SETI spectrometer design, especially the demonstration of DSP libraries and board interconnects, rapidly catalyzed the development of dozens of other spectrometers used all over the world for a diverse variety of applications. Just a few examples include a 13 Beam ANTF Parkes Telescope Pulsar

Spectrometer, FASR Kurtosis Spectrometer at NJIT, GAVERT (DSN), NRAO GUPPI Pulsar Machine, CNRS Orleans Nan cay Pulsar Machine, and the Allen Telescope Array fly's eye transient and pulsar instruments.

5 Correlators

A major problem in the design of large antenna array systems has been the routing of high bandwidth data. Each cross-correlation element must receive data from every antenna, and many processors are needed to handle the aggregate data rate; in the most straightforward arrangement, this leads to an unmanageable number of interconnections. CASPER has developed a packetized switch architecture capable of performing this antenna/frequency data transposition using commercial switches and cables [2]. We have developed a general purpose correlator architecture that uses standard 10-Gbit Ethernet switches to pass data between flexible hardware modules containing FPGA chips. These chips are programmed using open-source signal processing libraries we have developed to be flexible, scalable, and chip-independent. This work reduces the time and cost of implementing correlators with a wide range of capabilities, and facilitates upgrading to new generations of processing technology. CASPER has been responsible for several correlator deployments, including a 16-antenna, 200-MHz bandwidth, full Stokes parameter application deployed on the Precision Array for Probing the Epoch of Reionization, and a 32-antenna single polarization instrument deployed on Italy's Medicina Array.

One of the first correlator systems developed by CASPER was the "Pocket Correlator"--a single board system that includes F and X engines on a single board for correlating and accumulating 4 input signals. Each input is sampled at up to 800 Msps, and a digital down-converter extracts part of the digitized band. This system can act as a 2-antenna, full Stokes correlator, or as a 4-antenna single polarization correlator. The Pocket Correlator is valuable as a simple, stand-alone instrument, and for board verification in larger packetized systems.

The success of these early versions of the CASPER packetized correlator has drawn many interested parties to become involved in CASPER correlator development, including India's GMRT, South Africa's MeerKAT, CARMA, the ATA, SKADS, ISI and FASR. Plans for these arrays and for current collaborating projects (PAPER and Medicina) are driving future correlator development.

7 Workshops, Website and Online Training

The central mode of collaboration and sharing documentation about CASPER technology is through electronic communication, including the CASPER website, mailing lists and video teleconferences. CASPER also holds workshops; this year we are holding mini-workshops in Taipei (ASIAA), and Beijing (Beijing University); the main annual CASPER workshop will be held in Pune, India in October. For more information, see casper.berkeley.edu

References

1. A. Parsons, D. Backer, D. Werthimer, M. Wright, "A New Approach to Radio Astronomy Signal Processing: Packet Switched, FPGA-based, Upgradeable, Modular Hardware and Reusable, Platform-Independent Signal Processing Libraries", *Proceedings of the XXXth General Assembly of the International Union of Radio Science, Boulder, Colorado*, January 2006.
2. A. Parsons, D. Backer, A. Siemion, H. Chen, D. Werthimer, P. Droz, T. Filiba, J. Manley, P. MacMahon, A. Parsa, D. MacMahon and M. Wright, "A scalable correlator architecture based on modular FPGA hardware, reusable gateway, and data packetization", *Publications of the Astronomical Society of the Pacific* **120** (2008), pp. 1207–1221.
3. <http://casper.berkeley.edu>