Solving extreme FPGA high-density computing with AdvancedTCA

By Greg Tiedemann

A new initiative recently authorized by the Federal Communications Commission (FCC) teams satellites with earth-bound communications systems to significantly improve coverage and capacity for voice and data communications. The key technology to enable this coexistence is antenna beamforming, which adaptively forms hundreds of data channels from a satellite, while canceling interference from ground communication systems, via a combination of an array of antennas on the satellite and very sophisticated ground-based computational systems. With this new system, very high-bandwidth data or voice services can be directed to an area on the ground and received by mobile terminals, including small handheld devices. This provides a host of new communications services to out-of-the-way places, or after disasters like Hurricane Katrina.

Here Greg outlines how Mercury Computer Systems leveraged AdvancedTCA to enable a high-speed communications infrastructure – one that delivers the massive computing requirements that make this ground-based beam former a reality.

Leveraging the broad AdvancedTCA ecosystem

Our customer challenged Mercury with solving this difficult beamforming application while using a standards-based system compute platform. To meet the aggressive deployment costs and scheduled objectives, Mercury proposed leveraging the AdvancedTCA ecosystem, to take advantage of I/O and processing infrastructure to handle teraOPS of processing on terabits of data. While a proprietary system could have handled the requirements, it would have denied the customer the economic advantages of adopting a standard and using it to reduce the project risk and shorten time to market. Mercury selected AdvancedTCA as the standards-based platform for three important reasons:

- I/O – AdvancedTCA has unprecedented I/O capacity in the front and rear panels as well as the backplane
- Compute – the form factor is well-balanced teraOPS in a single subrack, though cooling remains a challenging design problem
- The need for a telecom-ready ecosystem to support high availability requirements and remote management

AdvancedTCA's modularity allowed the customer to scale all this I/O and compute capacity to various opportunities. Mercury is using every aspect of modularity in AdvancedTCA from Rear Transition Modules (RTMs) that route the antenna data streams into and out of the system to AdvancedMCs serving as host processors and hard drives.

In addition, the AdvancedTCA Intelligent Platform Management Interface (IPMI) infrastructure can be leveraged to accomplish management, upgrades, monitoring system health, and reporting alarms. These tasks can all be done speaking the same language, although there are different "dialects" within IPMI. Working with our customer and suppliers, Mercury is managing the tremendous amount of information that one can collect from the AdvancedTCA infrastructure and providing the ability to control and act on that information (for example, alarms), gaining a distinct advantage over any other commercially available technology.

> "Mercury worked closely with its chassis partner to achieve more than 200 W per slot."

A high-density FPGA computing solution

Figure 1 shows the two major components of the beam former that Mercury developed in close collaboration with our customer. Maximizing satellite receive power with beam shaping, which enables more antenna gain and less interference, as well as leveraging existing low-power wireless devices, requires:

- 300 Gbps of continuous I/O capacity in each direction. This implies 600 Gbps of intrasystem, bidirectional capacity. This translates to 25 Gbps in each direction per FPGA board with 12 boards in each shelf.
- 15 teraOPS of continuous computing per shelf. The beamforming is accomplished using 28 Xilinx Virtex4 SX-55 FPGAs at 400 MHz.

The Analog Conversion Unit (ACU) and the Beamformer Computational Unit (BCU) make up the central part of the beamforming system.

The ACU comprises 12 Analog Conversion Engines (ACEs) and two host processor modules. Each host processor module has a Gigabit Ethernet base switch and a Pentium M processor and a hard drive that are plugged into AdvancedMC sites. The system is hosted from one slot, with the other slot serving as a backup. The BCU consists of 12 Beamformer Conversion Engines (BCEs), which serve as compute blades for the BCU, and two host processors. Figure 2 shows a BCE block diagram and how traffic flows in through the RTM until being routed out to the FX60s.

As Figure 1 indicates, connectivity between the ACU and BCU takes place via a Fiber-Optic Rear Transition Module (FOM). A front panel analog SMA connector brings the data into

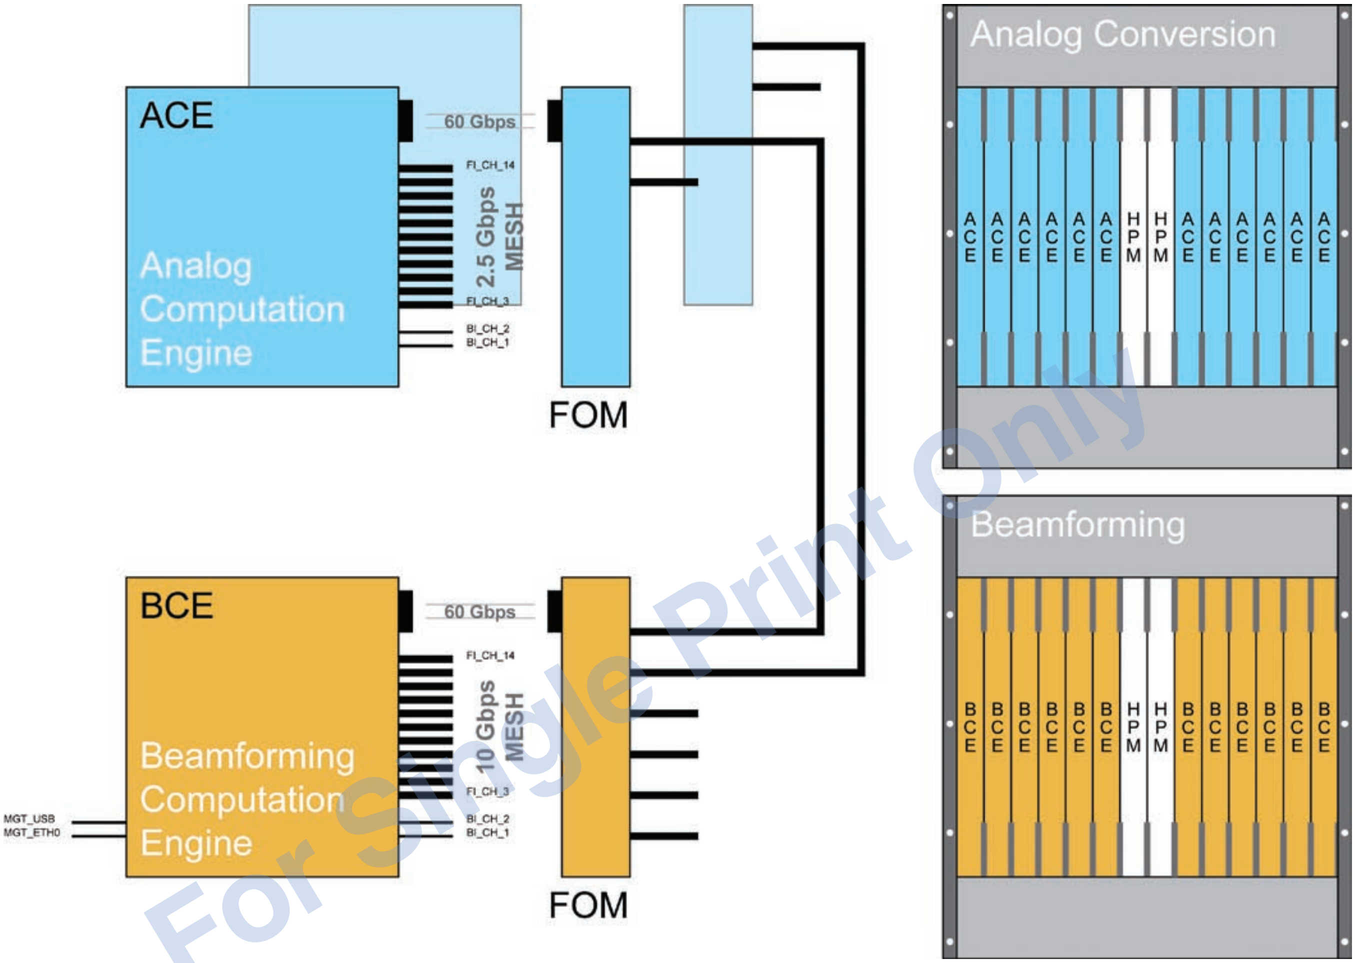

Figure 1

the ACE board from the satellite, where it is converted to digital. The data is sent out of the ACE rear transition module fiber optical (XFP) connectors and transferred to the RTM of the BCE.

Each BCE has 10 FPGAs. In the past customers might have used ASICs, but the advances in FPGA technology have opened the aperture on the possibility of deploying FPGAs for high-end compute demands. An FPGA-based solution aids with time to market and offers the customer the ability to continue to tune applications after the system is deployed.

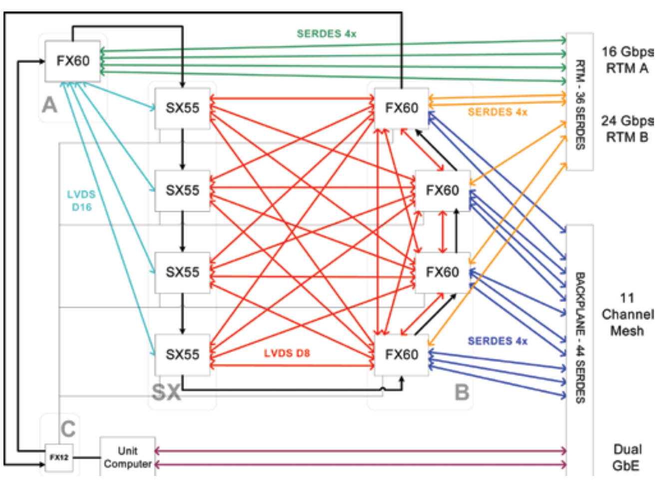

Figure 2

Opting to go with more than 100 FPGAs in one 14-slot chassis (Figure 3) meant pushing the AdvancedTCA 200 W per slot thermal limits, and we are satisfied with how AdvancedTCA is performing. Mercury worked closely with its chassis partner to achieve more than 200 W per slot. We could have put fewer FPGAs on the board to avoid cooling challenges, but that would have meant additional systems. Early in architecting the solution we had five separate chassis, but succeeded in compressing that to two chassis, thus needing to cool more per slot.

Figure 3

FPGA communications firmware is critical

Mercury developed an FPGA communications infrastructure firmware specifically for this application, making it possible to switch synchronously among the 100 or so FPGAs in real time, with very low latency so that voice quality is maintained.

The FPGA firmware delivered with the system enables the user to focus on the application by handling all the communication between the FPGAs in the chassis and the chassis-to-chassis communication. Within the IP there is an application socket, which presents all the I/O resources to the user after data capture and parallelization has occurred. This application socket is where the user enters their IP and utilizes the communications and control structure provided.

The firmware has built-in link checkers for all communication links within the system. This allows in-field system diagnostics to be run without any additional firmware. Also included within the firmware is the command and control infrastructure, which allows the processing subsystem to communicate to all the FPGA registers. The API enables software engineers to immediately start developing the application to reduce time to market.

The communications' infrastructure is customized for beamforming applications, enabling reconfigurable segmentation and re-assembly of I/O streams in time and space. There are many possible configurations for the IP, which will solve a variety of application domains.

We are also deploying the most advanced platform management software for remote operation, administration, and management of FPGA applications. Our goal is making applications highly available and highly reliable, and improving the serviceability and visibility of field-deployed applications.

Carrier Grade Linux deployed on processors throughout the system makes up the system's Linux Support Package.

More applications

Mercury has found in developing this system that this architecture could open up a number of other applications for AdvancedTCA. For example, beamforming applications are similar to some defense and commercial radar applications. Now a springboard to other applications that involve deploying massive numbers of FPGAs exists. We can do much more with the AdvancedTCA infrastructure than we were initially anticipating, and that bodes well for future growth and a long lifetime for AdvancedTCA.

Greg Tiedemann *is director of business development and systems engineering for Mercury Computer Systems' Communications Computing Segment. Prior to joining Mercury, he spent nearly 10 years with Ericsson. He has a BS in Mechanical Engineering from Tri-State University, Indiana.*

For more information, contact Greg at:

Mercury Computer Systems, Inc.
199 Riverneck Road • Chelmsford, MA 01824
978-967-1664 • Fax: 978-256-3599
gtiedemann@mc.com
www.mc.com