FPGA Technology Meets the Multicomputer Environment – Improving Processing Performance and Bandwidth in Advanced Naval Radar Missions

Sarah Leeper, Mercury Computer Systems, Inc., sleeper@mc.com
Mike Iaquinto, Lockheed Martin Maritime Systems & Sensors, michael.a.iaquinto@lmco.com
Richard Brooks, Mercury Computer Systems, Inc., rbrooks@mc.com
Joseph Finnivan, Mercury Computer Systems, Inc., jfinnivan@mc.com

Summary Abstract

The mission-critical demands for tomorrow's naval threats require ever increasing application enrichment and system performance. Increasing radar surveillance and target discrimination requirements, in parallel with integrated weapon system response and control functionality, requires a fresh architectural approach to system design. The architecture must attain the increased front-end I/O bandwidth and elevated throughput while maximizing signal processing performance for successful mission functionality. Today's general-purpose processors (GPPs), in a multicomputer environment, provide excellent flexible processing, but do not readily extend to the ultra-high I/O bandwidth required for growing front-end processing. This paper looks at the integrated field programmable gate array (FPGA) Compute Node (FCN) within a multicomputer environment as a solution to achieve these mission-critical goals.

Naval Mission of the Future

A part of the current mission of our naval cruiser and destroyers is the surveillance and response to airborne threats to the ship or fleet in theater. The heart of thise capability is the Aegis Weapon System. This incorporates advanced radar systems integrated with Mercury Computer Systems' signal processing systems.

Lockheed Martin has extended the traditional surveillance and response mission to include sea-based Ballistic Missile Defense (BMD) capability. This BMD capability requires real-time detection, tracking and discrimination performance against threat-representative targets and countermeasures. Architecting a combined, multi-mission-capable solution has been the goal of Lockheed Martin and Mercury Computer Systems' efforts in support of the U.S. Navy and Missile Defense Agency.

Lockheed Martin Maritime Systems & Sensors (MS2) is currently involved in the development of a digital signal processor (DSP) to support several advanced radar systems programs.

Requirements

The heart of radar operation, control and analysis systems is the digital signal processing system. However, with the complexities noted earlier, increasingly important is the capability to manipulate and pre-process data streams at extremely high-throughput rates. This led Lockheed Martin MS2 and Mercury Computer Systems to launch the development of an FPGA-based module integrated into a high-end multicomputer environment to handle the sensor I/O and front-end processing as well as the digital signal processing for these systems.

FPGA Compute Node (FCN) Module

The integrated FPGA module is required to input high-bandwidth raw data from a remote sensor, perform front-end algorithms in real time, handle a variety of processing scenarios that are controlled by an orthogonal real-time control interface, and distribute post-processed data to a GPP array.

It is known that some algorithms can run 10 times, and up to 20 times, faster on an FPGA than on a RISC processor. This incredible performance boost can attain far greater data manipulation speeds for the repetitive, less flexible decision-challenged pre-processing demands of the data streaming front end.

In support of these goals, the FCN module must:

- Utilize commercial off-the-shelf (COTS) based interfaces for sensor I/O, control, and data communication with GPPs
- Allow for high-bandwidth/low-latency communication between FPGAs within the system
- Incorporate fault detection and fault isolation into the basic design
- Provide a significant amount of gates within each FPGA that can be used to implement the front-end algorithms

- Provide for the use of commercially available FPGA development tools
- Provide a significant number of FPGAs to external RAM interfaces and adequate RAM for use by applications
- Provide a DMA-like capability for sending and receiving data from/to the GGP array
- Allow for fast download of FPGA bit streams to support low-latency reset
- Leverage intellectual property (IP) reuse standards originally developed by the ASIC community

Bringing it Together – An FCN Capability in a Multicomputer Environment

The team felt that what was needed was a system with the bandwidth and power of the PowerStream 7000 multi-computer system, and having integrated a new FCN FPGA-based capability with the high-bandwidth fabric access. This would provide the required multi-TeraOPS streaming I/O bandwidth, data manipulation, and pre-processing to be integrated somewhat seamlessly with TeraFLOP digital signal processing capability. This is what was needed as a foundation for a high-end, scalable radar multi-mission BMD solution.

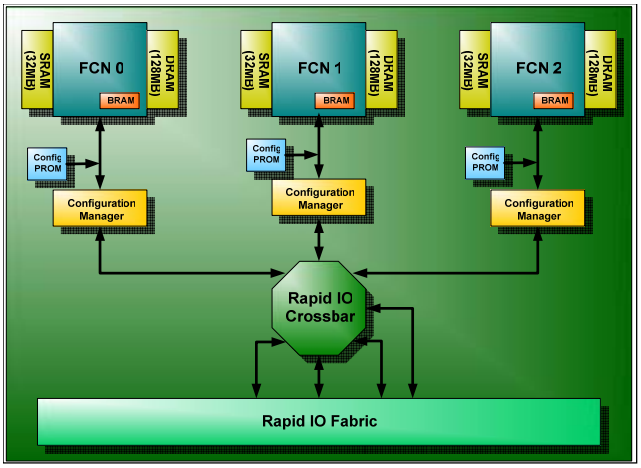

Figure 1. Basic block diagram of the field programmable gate array (FPGA) Compute Node (FCN) module within the PowerStream 7000.

Streaming I/O

Critical to the success of the Aegis radar programs is the ability to manage high-bandwidth, continuous input streams across multiple channels directly into the heart of the FPGA CNs. Dual 10-Gigabit/sec Ethernet (10 GigE) ports and IP protocol was selected for this approach.

The 10 GigE design required extensive CAD and signal modeling and simulation to ensure signal and data transfer integrity. At these speeds, even very low bit error rates (BER) can generate staggering data and message errors with significant throughput and data integrity problems. With a known vendor part having 10-12 BER specifications, it meant the team needed to drive extreme measures and exhaustive analysis and testing to minimize BER. As a naval environment requires remote signal input distances of 100 meters, special "aqua fiber" (2000 MHz-Km bandwidth multi-mode fiber) was selected and tested at 300m distances. Significant optimization of the pre-emphasis settings on the FPGA multi-gigabit transceiver (MGT) mesh links, and tuning of the register defaults in the key PHY communications chip (Physical 10 Gbit Xilinx XAUI-to-XFI parallel-to-serial signal conversion interface) made a huge impact on the BER. The final test demonstrated an extremely low BER of 10-16 at a 95% confidence level.

Table 1. Shown is the corresponding BER for a level of continuous error free testing. Residual BER test time, at 10 Gbaud, 95% confidence level.

Residual Bit Error Rate	Test Time
1E-10	3.00 sec
1E-11	30.0 sec
1E-12	300 sec
1E-13	49.9 min
1E-14	8.32 hrs
1E-15	3.47 days
1E-16	34.7 days
1E-17	346 days
1E-18	9.50 yrs

One of the key processing needs for scalable radar processing is the movement in massively parallel fashion of data from the FPGA processing nodes to the multi-computer processing nodes. Figure 2 depicts the performance of simultaneous DMA data transfers from FPGA CNs to PowerPC CNs nodes going over the RapidIO fabric. Different transfer buffer sizes are depicted, with the larger buffer sizes giving the highest transfer rate.

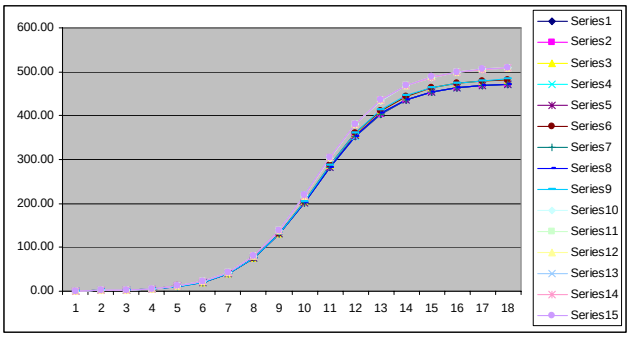

Figure 2. Performance of simultaneous DMA data transfers.

Finally, some of the key design features targeted for this integrated FCN multicomputer system approach include:

- Flexible FPGA bit-stream loading
- FCN onboard 240 GigaOPS per module
- QDR SRAM and 128 MB RLDRAM local to each FPGA to maximize FCN effectiveness and dataset-sizing flexibility
- Chip-to-chip 5x MGT mesh link at 3.125 GHz for larger local dataset manipulation across FPGAs
- Board-to-board 4x MGT copper mesh link at 3.125 GHz for successive operation within the FPGA environment
- FCN module 60 Gbytes/sec bisection bandwidth for massive, balanced I/O
- Low-Voltage Differential Signal (LVDS)

These issues will be addressed more fully at HPEC.