Experimental Analysis of Multi-FPGA Architectures over RapidIO for Space-Based Radar Processing

Chris Conger, David Bueno, and Alan D. George {conger, bueno, george}@hcs.ufl.edu

High-performance Computing and Simulation (HCS) Research Laboratory Department of Electrical and Computer Engineering, University of Florida

Abstract

Increasingly powerful radiation-hardened FPGAs, ASICs, and conventional processors along with high-performance embedded interconnect technologies are helping to enable the on-board processing of real-time, high-resolution radar data on satellites and other space platforms. The streaming nature of most Space-Based Radar (SBR) algorithms calls for highly-efficient data and control networks, both locally within the processing nodes as well as across the system interconnect. Furthermore, the customized architecture of FPGAs and ASICs allows for unique features, enhancements, and communication options to support such applications. Using a real-time Ground Moving Target Indicator (GMTI) application as a case study, this research investigates low-level architectural design considerations on a custom-built RapidIO testbed. We present experimentally gathered results supported by high-resolution simulations of scaledup systems to provide insight into the relationship between the external communication network and the local memory architecture for these challenging HPEC applications.

Introduction

In previous work, we have shown using simulation [1-3] that on-board processing of high-resolution radar data requires parallel processing platforms providing high throughput to all processing nodes and efficient processing engines to keep up with strict real-time deadlines. In this work, we conduct an experimental study of cutting-edge system architectures for SBR based upon FPGA processing nodes and the RapidIO high-performance interconnect for embedded computing. Hardware processing provides the clock-cycle efficiency necessary to enable high-performance computing with lower-frequency, radiation-hardened components. Compared to bus-based designs, packet-switched interconnects such as RapidIO substantially increase the scalability, robustness, and network performance of future embedded systems, and the small footprint and fault tolerance inherent to RapidIO make it an ideal fit for use with FPGAs in space systems.

GMTI is used to track moving targets from air or space, and may have a strict real-time deadline for processing each consecutive radar return [1]. Radar satellites have a potentially wide field of view looking down from space, but maintaining resolution at that viewpoint results in very large data sets. Our design of the algorithm is composed of a sequence of several very common signal and image processing kernels that we study individually as well as together as a whole application. The data input and output requirements of the processing engines, as well as the data paths provided by the interconnect and the local memory hierarchies, are key design parameters that affect the ultimate system performance. By experimenting with various node architectures and communication schedules, we are able to identify critical design aspects as well as suggest enhancements or alternative approaches to improve performance for SBR applications.

Experimental research is conducted on a RapidIO testbed constructed in our laboratory using Xilinx FPGAs, development boards, and IP cores. In addition, we have constructed a collection of custom-built enhancements including a RapidIO switch board and additional processing cores. Each node of the 4-node testbed consists of a Virtex-II Pro 40 FPGA (containing two PowerPC405 processor cores), 128 MB of DDR SDRAM, and a single RapidIO endpoint (8-bit parallel, 250 MHz DDR). Xilinx currently makes radiation-hardened FPGA devices similar to those used in our testbed, and we will show the benefit of leveraging the potential efficiency of hardware processing to achieve higher performance in frequencyand power-limited environments such as space. The RapidIO switch is a 4-port Tundra Tsi500 and is used to connect each node of the testbed as well as assist in Our extensive RapidIO performance measurement. simulation environment [2] is also significantly enhanced to accurately capture the performance of our hardware testbed through the integration of high-fidelity FPGA and memory subsystem models with our existing RapidIO network model.

Description of Testbed Architecture

Our testbed architecture is designed to mimic a realistic satellite processing platform in order to provide a relevant foundation for our experiments. Each node of the testbed is a combination of multiple processing engines, noncached internal processing memory (SRAM) and external storage memory (SDRAM), a control network, and a network interface. Using RapidIO for the communication fabric provides a high-performance and scalable interconnect, favored by both Honeywell and SEAKR [4] for their next-generation space systems. Like typical production-level systems, the network fabric must contend with the internal processing engines for access to each node's SDRAM, and the internal SRAM for each processing engine is much smaller than the node's data set at any one time. The PowerPC405s in the Virtex-II Pro FPGAs are relatively low-performance and have been found to be ideally suited for controlling system operation (e.g. DMA transfers and co-processor operations), leaving data processing to the highly-efficient hardware processing

engines implemented in the reconfigurable fabric of each FPGA.

Experimental Setup

A complete description of each experimental study will appear in the full presentation, but for this abstract a condensed version follows. The main kernels that make up our GMTI application include Pulse Compression (PC), Doppler Processing, Beamforming, Constant False Alarm Rate (CFAR) detection, and distributed corner turns. Independent variables for experimentation include the number of processing engines in each node, the allocation of processing tasks among the nodes, and the algorithm decomposition and communication schedule. In addition, the flexible architecture of the FPGA will be leveraged to examine variations and enhancements to common design features such as the interface to the SDRAM or network controller, or the processing engine topologies. Examples of such enhancements include adding multiple parallel request queues/ports on the external memory controller, or enabling direct communication of data from internal processing memories to the network controller.



Figure 1: Kernel processing time vs. varied clock frequency

Results

Figure 1 displays kernel execution time for four combinations of RapidIO clock frequencies and DDR SDRAM clock frequencies on a two-node distributed corner turn followed by CFAR processing on a 512 KB data cube. These clock frequencies were chosen based on available clock generation hardware on our development boards. The results show performance is currently bound by the speed of the DDR SDRAM interface, where improvements in the DDR SDRAM interface yield significant speedup for the overall application because both communication and computation require heavy access to the shared DDR SDRAM. We are in the process of integrating a commercial DDR SDRAM controller for our final configuration that will substantially improve the performance of the memory subsystem. The full presentation will include results with larger data cube sizes along with Pulse Compression, Doppler Processing, and Beamforming kernels. In addition, we will investigate

other corner turn variations and architectural design tradeoffs.

Conclusions

The use of RapidIO as well as customized FPGAs in a dual-paradigm computing design approach suggests an attractive infrastructure for next-generation, highperformance satellite processing systems. This new work employs a combination of experimental and simulative research techniques that leverage one another to be more effective. The experimental testbed provides a realistic foundation upon which to experiment and investigate practical design challenges, as well as experimentally verify simulation results. The simulative research allows us to accurately predict system performance for largescale, realistic system sizes that are impractical for experimentation. Results show that system performance is very sensitive to main memory throughput, and confirm our assertion in [2] regarding the importance of doublebuffering and intelligent communication scheduling for SBR applications. Our work shows that scalable, highperformance parallel processing platforms can be realized in the frequency-limited, radiation-hardened constraints of space. Future directions for the experimental research include expanding the testbed by enhancing memory and communication capabilities (e.g. Serial RapidIO), and considering additional SBR applications. Future simulative work will include different traffic classes such as high-throughput data traffic and latency-sensitive control traffic with strict delivery deadlines.

Acknowledgements

We wish to thank Honeywell Defense and Space Electronic Systems (DSES) in Clearwater, FL for support of this research. We also extend thanks to Xilinx for their generous donation of hardware and IP cores, Tundra for donating the RapidIO switches in our testbed, as well as MLDesign Tech. for the donation of simulation software.

References

- D. Bueno, C. Conger, A. Leko, I. Troxel, and A. George, "Virtual Prototyping and Performance Analysis of RapidIObased System Architectures for Space-Based Radar," *Proc. High-Performance Embedded Computing (HPEC) Workshop*, MIT Lincoln Lab, Lexington, MA, Sep. 28-30, 2004.
- [2] D. Bueno, C. Conger, A. Leko, I. Troxel, and A. George, "RapidIO-based Space System Architectures for Synthetic Aperture Radar and Ground Moving Target Indicator," *Proc. High-Performance Embedded Computing (HPEC) Workshop*, MIT Lincoln Lab, Lexington, MA, Sep. 20-22, 2005.
- [3] D. Bueno, A. Leko, C. Conger, I. Troxel, and A. George, "Simulative Analysis of the RapidIO Embedded Interconnect Architecture for Real-Time, Network-Intensive Applications," Proc. 29th IEEE Conf. on Local Computer Networks (LCN) via the IEEE Workshop on High-Speed Local Networks (HSLN), Tampa, FL, Nov. 16-18, 2004.
- [4] S. Vaillancourt, "Space Based Radar On-Board Processing Architecture," Proc. 26th IEEE Aerospace Conference, Big Sky, MT, Mar. 5-12, 2005.