Stockholm - Session S1 - NoC/off-chip communication

B1. The CompSOC Design Flow for Virtual Execution Platforms

Designing a SoC for applications with mixed time-criticality is a complex and time-consuming task. Even when SoCs are built from components with known real-time properties, they still have to be combined and configured correctly to assert that these properties hold for the complete system, which is non trivial. Furthermore, applications need to be mapped to the available hardware resources and correctly integrated with the SoC's software stack, such that the real-time requirements of the applications are guaranteed to be satisfied. However, as systems grow in complexity, the design and verification effort increases, which makes it difficult to satisfy the tight time-to-market constraint. Design tools are essential to speed up the development process and increase profit. This paper presents the design flow for the CompSOC FPGA platform: a template for SoCs with mixed time-criticality applications. This work outlines how the development time of such a platform instance is reduced by means of its comprehensive tool flow, that aids a system designer in creating hardware, the associated software stack, and application mapping.

B2. A FLEXIBLE COMMUNICATION METHOD FOR MULTI-FPGA BASED DESIGNS

High throughput applications have been widely used for deployment onto FPGAs. As the requirements for performance increase so does the challenge of integrating more complex algorithms into these platforms. This paper presents a flexible, low resource usage method which can be configured to integrate large designs into multiple FPGA chips using Xilinx's high-speed serial interfaces that minimize the performance penalty due to chip-to-chip communication. The proposed solution also provides a way of sharing the same data link for multiple applications, which offers valuable support for distributed FPGA applications with a minimal overhead and configurable bandwidth division. The impact of using this method inside a new application should not require complex architectural changes. The design was successfully validated and then applied onto a set of open-source cores.

B3. An Improved Transmission scheme for Error-prone Inter-Chip Network-on-Chip Communication links implemented on FPGAs

Network-on-Chip (NoC) is an alternative to traditional busses for faster interconnect mechanism. The aim is to have infinite scalability, and this implies the possibility to extend the on-chip NoC communication protocol off-chip. To gain wholesome advantage of Network-on-Chip (NoC), off-chip extensions should also have similar communication throughput compared to the on-chip network. Faster data-rate is the single most demanded requirement of modern applications. There is a continuous drive to fulfill this escalating demand as much as possible. Two of the most prominent limiting factors in achieving this purpose are 'reduced accuracy' and 'protocol handling', especially in case of systems which do not have synchronous communication. Efficient optimizations are needed in multiple areas to upgrade the speed of data transfer. This paper presents an improved off-chip network solution to a slower and error-prone board-bridge part of a Network-on-Chip (NoC). The new solution increases the accuracy and speed of the plesiochronous off-chip extension to the NoC. The Network-on-Chip has 16 processor-nodes implemented on four interconnected plesiochronous Altera Stratix-II FPGA boards in 4x4 configuration in such a way that each board hosts a Quad-core NoC.

Stockholm - Session S2 - Power Optimization

B4. Energy proportional computing in Commercial FPGAs with Adaptive Voltage Scaling

Voltage and frequency adaptation can be used to create energy proportional systems in which energy usage adapts to the amount of work to be done in the available time. Closed-loop voltage and frequency scaling can also take into account process and temperature variations in addition to system load and this

removes a significant proportion of the margins used by device manufacturers. This paper explores the capabilities of commercial FPGAs to use closed-loop adaptive voltage scaling to improve their energy and performance profiles beyond nominal. An adaptive power architecture based on a modified design flow is created with in-situ detectors and dynamic reconfiguration of clock management resources. The results of deploying AVS in FPGAs shows power and energy savings exceeding 85% compared with nominal voltage operation at the same frequency or 100% better performance at nominal energy. The in-situ detector approach compares favorably with critical path replication based on delay lines since it avoids the need of cumbersome and error-prone delay line calibration.

B5. Energy and Performance Exploration of Accelerator Coherency Port Using Xilinx ZYNQ

Cooperation of CPU and hardware accelerator to accomplish computational intensive tasks, provides significant advantages in run-time speed and energy. Efficient management of data sharing among multiple computational kernels can rapidly turn into a complicated problem. The Accelerator coherency port (ACP) emerges as a possible solution by enabling hardware accelerators to issue coherent accesses to the memory space. In this paper, we quantify the advantages of using ACP over the traditional method of sharing data on the DRAM. We select the Xilinx ZYNQ as target and develop an infrastructure to stress the ACP and high-performance (HP) AXI interfaces of the ZYNQ device. Hardware accelerators on both of HP and ACP AXI interfaces reach full duplex data processing bandwidth of over 1.6GBytes/s running at 125MHz on a XC7Z020-1C device. The effect of background DRAM and cache traffic on the performance of accelerators is analyzed. For a sample image filtering task, the cooperative operation of CPU and ACP accelerator (CPU-ACP) gains a speed-up of 1.2X over CPU and HP acceleration (CPU-HP). In terms of energy efficiency, an improvement of 2.5nJ (>20%) is shown for each byte of processed data. This is the first work which represents detailed practical comparisons on the speed and energy efficiency of various processor-accelerator memory sharing techniques in a configurable heterogeneous platform.

B6. Tracking the Pipelining-Power Rule along the FPGA Technical Literature

This work reviews the contributions of power-oriented pipelining over the last two decades, and adds up-to-date results on 65 nm and 45 nm FPGAs. The data show that power consumption can be reduced by a factor between 0.1 and 0.8 using different levels of pipelining. More than 34 experiments, developed in 12 laboratories in 8 countries during 17 years are summarized. This paper also contributes to this research topic adding updated results for Altera 65-nm Cyclone III and Xilinx 45-nm Spartan-6 devices.

Copenhagen Session C1 - Optimization/Design Examples

B1. Teaching System-on-Chip design with FPGAs

This paper presents our experiences in using FPGA in teaching System-on-Chip design in Tampere University of Technology. We had a major reform on our courses and, most notably, chose a common HW platform which is used in 11 courses. It has proved good that most exercises are mandatory and bonus points are awarded for good work. In order to manage the schedules, larger projects have been partitioned by the teachers into smaller tasks and pairwork is allowed. Automated testbenches, reuse, startup examples were very useful. As a result, we observed increased motivation among students and better learning outcomes. The schedule slippages were reduced, although both teachers and students still underestimate the required time and effort. Moreover, we introduce 15 student projects where FPGA platform was also used. Some of the most innovative topics were suggested by students themselves, such as games. In the future, more effort is needed is finalizing the project works for easier reuse and setting up a common repository.

B2. Fast and Energy Efficient AdaBoost Classifier

The paper presents a new concept of creating an energy and computation effective AdaBoost classifier systems. The presented method is mainly novel in the way how it divides and accelerates an AdaBoost classifier into two parts – a pre-processing and a post-processing unit. Pre-processing unit is designed to process a major part of the computational operations of the AdaBoost algorithm but it is also helps in an energy savings.

B3. High performance 3-Dimensional Heterogeneous Tree-based FPGA Architectures (HT-FPGA)

We describe the design and exploration methodology to optimize 3-dimensional (3D) Heterogeneous Tree-based FPGAs (HT-FPGAs) by introducing a break-point at a particular tree level interconnect to optimize the speed, power consumption and area. The ability of the flow to decide a horizontal or vertical partitioning of the programmable tree network based on design specifications is a defining feature. The break-point of the vertically partitioned tree is designed to balance the placement of logic blocks and switch blocks into multiple tiers while the horizontally partitioned tree is designed to optimize the interconnect delay of the programmable tree network. We finally evaluate the performance, area and power of the proposed 3D HT-FPGA using the newly developed flow and show that vertical and horizontally partitioned 3D stacked HT-FPGA improves speed by 16% and 55% respectively compared to 2D planar design.

B4. Resource efficient implementation of a 10Gb/s radio receiver baseband in FPGA

Data-rate of wireless links are increasing fast, particularly when new carrier bands with very high bandwidth becomes available. Utilizing the full bandwidth for a single carrier facilitates very large baud-rates. When the baud-rates approaches or exceeds Gbaud, the implementation of the digital baseband is no longer a simple extension of existing methods. In the present paper we propose a resource efficient implementation of digital baseband for multi-Gbaud rates in a standard FPGA utilizing Xilinx Simulink-based System generator design and verification tool.