International Journal Of Applied Research In Science And Engineering
Volume 1, Issue 6, October-2016 *ISSN: 2456-124X*

MULTI-CORE PROCESSING WITH INTERCONNECTION NETWORKS IN MODERN ERA

Dr. Ajay Mathur[1], Ashish Sharma[2]
[1]HoD (C.S.E.) Govt. Polytechnic College Jodhpur, [2]HoD (C.S.E)Jodhpur National University

Abstract: Customizable processors that perform intensive data processing are designed to provide programmability in the performance-intensive data plane of the system-on-chip (SoC) design. Not only do they combine the capabilities of a DSP and a CPU, but they can be customized to maximize efficiency for your target application. Certainly, more competition usually leads to new innovation and markets. Customizing your processor design makes the processor unique, which makes it much harder for competitors to copy your ideas. You get a version of a processor that no one else can buy. No one else can get the matching software tool chain unless you provide it to them, so no one can use the processor's custom additions in your ASIC unless you allow it. When used with your software development tools, your optimized processor will get better performance, operate at lower required clock rates, and consume less energy than the industry-standard, fixed-ISA microprocessors.

I. INTRODUCTION

While processors are often used for the control functions in system-on-chip (SoC) designs, designers turn to RTL blocks for many data-intensive functions that control processors can't handle. However, RTL blocks take a long time to design and even longer to verify, and they are not programmable to handle multiple standards or designs. The most common embedded microprocessor architectures—such as the ARM, MIPS, and PowerPC processors—were developed in the 1980s for stand-alone microprocessor chips. These general-purpose processor architectures, or CPUs, are good at executing a wide range of algorithms with a focus on control code, but SoC designers often need more performance in critical datapath portions of their designs than these microprocessorarchitectures can deliver. To bridge this performance gap, the two most-used approaches are to run the general-purpose processor at a higher clock rate (thus extracting more performance from the same processor architecture), or to hand-design acceleration hardware that offloads some of the processing burden from the processor. Running a general-purpose processor core at a high clock rate incurs a power and area penalty, and designing acceleration hardware takes additional development time, not just for the design but for verification of the new acceleration hardware. In fact, verification can consume as much as 70% to 80% of the total design time. Customizable processors can achieve high performance and lower energy consumption, save time, and provide design flexibility versus hand-coded RTL hardware or a general-purpose processor.

II. CONCEPTOF CUSTOMIZABLE PROCESSOR

When based on a single processor architecture customizable processor can use a common development flow with tools that scale from tiny microcontrollers and digital signal controllers to high-performance, real-time controllers and DSPs. Processors can implement wide, parallel, and complex data path operations that closely match those used in custom RTL hardware. The equivalent data paths are implemented by augmenting the base processor's integer pipeline with additional execution units, registers, and other functions developed by the chip architect for a target application. You get this customization with an automated flow that requires no extra processor verification and keeps the development tools and simulation models updated with every change. The processor is ultimately delivered as synthesizable RTL code, ready for integration into an FPGA prototype or SoC design, so they fit easily into existing design flows. The result is a new processor, not a processor with bolted-on coprocessors. The new instructions and registers are available to the firmware programmer via the same compiler and assembler that target the processor's base instructions and register set. The instruction extensions greatly accelerate the processor's performance on the targeted algorithm and the controlling firmware can be written in a high-level language for easy development and maintenance. Nvidia has been an ARM licensee, combining ARM's CPU technology with its own graphics technology to create its Tegra line of mobile processors. Nvidia has demonstrated that it had ported its Kepler GPU architecture so that it would work in systems with ARM CPUs. Recently it has licensed GPU core as well as rights to its visual computing intellectual property, so customers can create their own GPUs. New licensing plans seem to be primarily aimed at other mobile processor vendors & at growing embedded market by focusing on how some technologies can operate with little as half a watt of power.IBM has also announced that it will offer its Power technology, typically used by the company in its chips for high-endd servers and mainframes for development. Along with Google, Mellanox, Nvidia and Tyan, IBM are forming the OpenPower Consortium aimed at extending the Power architecture and server, networking storage, and graphics technology around it to create solutions aimed at very large data centers. Intel has also talked about creating semi-custom versions of server chips for its giant customers. The most sophisticated customers, for which specialized chipsare designed, tend to expend in least for testing software on new platforms when compared to the cost of actually running the data center.

III. DESIGNING OF CUSTOMIZABLE PROCESSORS

The level of processor and chip customization varies with the workload, data center design & even cooling solutions. An example would be a cooling system in data center that can allow customers to run processors at a higher frequency. Customers usually give information about the application they are running, the accelerators they need, the performance and power consumption levels they are looking to hit. Intel who has started building system on chip (SoC) designs combines CPU with other accelerators, I/O, graphics and other processing units. Advanced Micro Devices has created custom chips based on its CPU & graphics architectures largely for non-server ventures like Playstation & Xbox. Latestprocessors uses specialized language like Tensilica Instruction Extensions (TIE) for customization& define the desired processor instructions and I/O. This language closely resembles a simplified version of Verilog.The hardware description language that ASIC designers already know. TIE enables you to extend the base RISC architecture for specific tasks, and permits the high-level specification of new datapath functions as new processor instructions, registers, register files, I/O ports, and FIFO queue interfaces. A TIE description is both simpler and much more concise than RTL because it omits all sequential logic descriptions including FSM descriptions and initialization sequences. These complex items are more easily developed in firmware. It is also unnecessary to describe the logic structure of these new functions and registers. The Xtensa Processor Generator (XPG) infers structure from the functional TIE description and creates new processor hardware that is guaranteed correct-by-construction. The XPG creates an RTL description of the processor and generates tailored versions of all necessary software development tools including the compiler, assembler, debugger, and instruction set simulator, as well as a Cor System C simulation model of the processor and EDA synthesis scripts. No manual work is required to match software development tools and processor. This new processor hardware is automatically blended into the Xtensaprocessor's base architecture, creating a seamless fusion of base architecture and task-specific ISA extensions

One of first design for Power architecture was announced by IBM as Power 8which the company plans to announce at Hot Chips conference later this month and will begin shipping in 2014. Power 8 includes a new advanced I/O bus, known as the Coherent Attached Processor Interface (CAPI), which IBM says will make it easier to combine Power cores with other system components for heterogeneous computing. The idea is to allow organizations to easily link multiple Power CPUs with Nvidia GPUs in a way that makes sense & eventually allow for specialized processors that could create an alternative to standard Intel servers. Soft-core processors are becoming increasingly common in modern technology. A soft-core processor is a programmable processor that can be synthesized to a circuit, typically integrated into a larger system existing on anapplication-specific integrated circuit (ASIC)or field-programmable gate array (FPGA). Popular commercially available soft-cores include ARM, Tensillica, Microblaze, and Nios.However when comparing processors, a focus on total energy consumption is the key. Too often, designers fixate on a static "milliWatts per megahertz" (mW/MHz) number while ignoring the total energy consumption of the workload. There is an implicit assumption in this type of thinking that all RISC processors deliver about the same performance per clock cycle. This assumption does not apply to processors that have been customized for a specific task or application. Soft core'sDoE[3][7] experimental design is based on four main principles:

- Randomization means that the experiments are performed in a random order.
- Replication means that each experiment is repeated.
- Blocking is the process of grouping different experiments into a group and running those experiments in thesame environment.
- Orthogonality is the process of creating the experiments so that one factor canbe analyzed independently ofthe other factors.

IV. BENEFITS OF CUSTOMIZABLE PROCESSORS

Adding a few custom instructions, increases the processor core's size, which in turn increases the average power dissipation per clock cycle (increases the mW/MHz rating). However, if the custom instructions dramatically cut the total clock cycles required to perform a given workload, then the total energy consumed (power-per-cycle multiplied by total cycle time) can be substantially reduced. A 20% increase in power dissipated per clock cycle, offset by a 3X speed up in task execution, actually reduces energy consumption by 60%. This reduction in required task-execution cycles allows the system either to spend much more time in a low-power sleep state or to reduce the processor's clock frequency and core operating voltage, leading to further reductions in both dynamic and leakage power. The well-established Design of Experiments (DoE) paradigm can be exploited to tune a microprocessor soft-core to an application, yielding 6x-17x speedup compared to a base core. Those speedups are 3x-6x better than obtained by a previous non-DoE- based core tuning approach. The key benefit of DoE is the multi-factor analysis, proven to yield near-maximum information from a given small number of experimental runs.

REFERENCES

[1] Givargis, T., F. Vahid. Platune: A Tuning Framework for System-on-a Chip Platforms. IEEE Transactions on Computer Aided Design, Vol. 21, No. 11, Nov. 2002, pp. 1317-1327

[2] Kumar. R., D. Tullsen, P. Ranganathan, N. Jouppi, K. Farkas. Single-ISA Heterogeneous Multi-core Architectures for Multithreaded Workload Performance. In31st International Symposium on Computer Architecture, ISCA-31, June 2004.

[3] McLean, R., V. Anderson, Applied Factorial and Fractional Designs. Marcel Dekker, Inc. New York, New York, 1984.

[4] R. Kumar, D. Tullsen, P. Ranganathan, N. Jouppi, and K. Farkas, "Singleisa heterogeneous multi-core

architectures for multithreaded workload performance," in ACM SIGARCH Computer Architecture News, vol. 32 ,no. 2. IEEE Computer Society, 2004, p. 64.

[5] Yiannacouras, P., J. G. Steffan, J. Rose. Application-Specific Customization of Soft Processor Microarchitecture.FPGA 2006.

[6] Sheldon, D., R. Kumar, R. Lysecky, F. Vahid, D. Tullsen. Application-Specific Customization of Paramaterized FPGA Soft-Core Processors.Intl. Conf. on Computer-Aided Design (ICCAD), 2006.

[7] Petersen, R., Design and Analysis of Experiments. Mercel Dekker Inc. New York, New York, 1985.

[8] D. Pham, S. Asano, M. Bolliger, M. Day, H. Hofstee, C. Johns,J. Kahle, A.Kameyama,J. performance issues on both single core and multi-core architecture," Technical report, University of Virginia, Department of Computer Science, Charlottesville, Tech. Rep.,2007.

[9] R. Ubal, J. Sahuquillo, S. Petit, P. L´opez, Z. Chen, and D. Kaeli, "The multi2sim simulation framework.

[10] K. Hwang, Advanced computer architecture. Tata McGraw-Hill Education, 2003.

[11] R. Ubal, J. Sahuquillo, S. Petit, and P. L´opez, "Multi2sim: A simulation framework to evaluate multicore-multithread processors," in IEEE 19th International Symposium on Computer Architecture and High Performance computing, page (s), 2007, pp. 62–68.

