

The IMData Approach to Accelerating Data Intensive Workloads

Marcus Völz, Nils Asmussen, Hermann Härtig

Technische Universität Dresden, Dresden, Germany,
{voelp,nils,haertig}@os.inf.tu-dresden.de

Abstract. Having started operational work in August 2012, the ESF young researcher group IMData seeks to develop new integrated mechanisms for accelerating data intensive workloads in heterogeneous many-core systems. This extended abstract and the accompanying poster summarizes the observations that motivate the project and the approach we are going to take.

Accelerator architectures such as GPGPUs, IBM Cell, and signal processors achieve their performance and energy advantages by offering a vast number of very simple and specialized cores. The tremendous core count of these architectures stems in part from a sacrifice of even the most fundamental hardware features that interacting applications and operating system kernels require. However, for not frequently interacting workloads, this tradeoff pays off in gigantic speedups.

Comparing different Alpha generations, Kumar et al. [2] already found that little less than five simple cores occupy the same area as a large out-of-order core motivating a combination of few big and many little cores in single-ISA heterogeneous systems. However, the chip area requires to support an operating system surprised us. We configured two Tensilica Xtensa cores [1] for a 400MHz 65nm Low Power Process Technology using the Xtensa Xplorer Tools. One of it is used in the Tomahawk MPSoC accelerator for Software Defined Radio [3]. The other has all features enabled that modern general purpose operating systems require (interrupts, exceptions, ...). Adding OS support to a core (notice, we stick to the same ISA and application-level execution profile) increases the core's area by almost 100% from 0.095 mm² to 0.171 mm² (excluding caches). Therefore, just by removing hardware support for the operating system, thread parallelism could be doubled. Additionally, replacing the

32 KB tightly coupled SRAM with two 16 KB caches for instructions and data, multiplies chip area by almost another factor of two (from 0.299 mm^2 vs. 0.516 mm^2).

The research direction, which we derive from this observation is therefore how to enable complex, interactive and data intensive workloads on accelerator platforms without sacrificing the benefits obtained from hardware specialization. The immediate questions that arise are: “How can we control applications without having full hardware support for running an operating system kernel beneath them?”, “How can applications on different cores interact with each other?”, “How can they synchronize on shared data and communicate results?”, “How can critical parts of applications be isolated?”, and finally, “How can all this be achieved without reintroducing the area-expensive hardware mechanisms of general purpose architectures?”.

At the hardware level, the starting point will be the Tomahawk processor built by TU Dresden’s mobile communications group. On top of it, the L4 microkernels NOVA and Fiasco.OC from TUD’s operating systems group will provide efficient and light-weight communication and synchronization mechanisms to the data intensive applications, which we draw from in-memory databases and from other research fields of TUD’s database and data analytics group. In addition, our probabilistic model checking group will support us by looking into new techniques for investigating functional and non-functional properties of our algorithms and system components.

References

1. Inc., T.: Xtensa customizable processors. <http://www.tensilica.com/products/xtensa-customizable.htm> (2011)
2. Kumar, R., Farkas, K.I., Jouppi, N.P., Ranganathan, P., Tullsen, D.M.: Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction. In: Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture. pp. 81–. MICRO 36, IEEE Computer Society, Washington, DC, USA (2003), <http://dl.acm.org/citation.cfm?id=956417.956569>
3. Limberg, T., Winter, M., Bimberg, M., Klemm, R., Tavares, M.B., Ahlendorf, H., Matus, E., Fettweis, G., Eisenreich, H., Ellguth, G., Schlüssler, J.: A Heterogeneous MPSoC with Hardware Supported Dynamic Task Scheduling for Software Defined Radio. In: 46th Design Automation Conference (DAC’09). San Francisco, USA (267 - 317 2009)