LOCA SERIES. Reading Club. 2nd session

Objectives

The general objective of LOCA Technical sessions is to value the knowledge, talent and experience of the center's workers and researchers, and support the creation of a multidisciplinary work team based on the strengths of each and every one of its members.

To contribute to a sustainable and organic growth of the hardware design area at BSC, we propose organizing Internal LOCA Tech sessions. A monthly event whose objectives are: 1) promotion of internal BSC talent (Tech-Talk sessions), and 2) continuous learning/development mechanism (Reading club sessions).
? Tech-Talk session: Technical presentations on specific topics with the aim of putting the spotlight on the most senior and post-doc community's background and know-how, and also making visible the research and/or interests areas of work of each of those members.
? Reading club session: Technical sessions in the form of a presentation with the objectives of: 1) actively involving members with less experience in research activities, and 2) offering a "light-weight" training/retraining mechanism for all members that allows them to keep up to date with the state of the art.

Topic and Presenter: Fast Behavioural RTL Simulation of 10B Transistor SoC Designs with Metro-MPI, Guillem Lopez Paradis, High Performance Domain-Specific Architectures, CS

Abstract: Chips with tens of billions of transistors have become today's norm. These designs are straining our electronic design automation tools throughout the design process, requiring ever more computational resources. In many tools, parallelisation has improved both latency and throughput for the designer's benefit. However, tools largely remain restricted to a single machine and in the case of RTL simulation, we believe that this leaves much potential performance on the table.
We introduce Metro-MPI to improve RTL simulation for modern 10 billion transistor-scale chips. Metro-MPI exploits the natural boundaries present in chip designs to partition RTL simulations and leverage High Performance Computing (HPC) techniques to extract parallelism. For chip designs that scale in size by exploiting latency-insensitive interfaces like networkson-chip and AXI, Metro-MPI offers a new paradigm for RTL simulation scalability. Our implementation of Metro-MPI in OpenPiton+Ariane delivers 2.7 MIPS of RTL simulation throughput for the first time on a design with more than 10 billion transistors and 1,024 Linux-capable cores, opening new avenues for distributed RTL simulation of emerging system-on-chip designs. Compared to sequential and multithreaded RTL simulations of smaller designs, Metro-MPI achieves up to 135.98× and 9.29× speedups. Similarly, for a representative regression run, Metro-MPI reduces energy consumption by up to 2.53× and 2.91×.

Topic and Presenter: Approximate Arithmetic Circuits for Energy-Efficient DNN Object Detection, Jordi Fornt Mas, Synthesis and Physical design of ICs, CS

Abstract: Approximate computing has recently resurfaced in relevance as a way to improve the efficiency of deep neural network (DNN) accelerators. In this talk, we will present the results of our evaluation of

different approximate multiplier and adder circuits in the context of deep neural networks for object detection, using an approximate version of the YOLOv3 model for assessing the trade-offs of each multiplier and adder architecture. We also combine approximate multipliers with approximate adders and study their interaction as a multiply-add unit. Using the best performing approximate multiply-add unit in our study, we are able to save 80% of power consumption while only decreasing the mean Average Precision by 0.46% with respect to the exact YOLOv3 implementation.

Chairs: Xabier Abancens, European Exascale Accelerator Research Engineer, Computer Sciences

Speakers

Presenters: Guillem Lopez Paradis, High Performance Domain-Specific Architectures and Jordi Fornt Mas, Synthesis and Physical design of ICs, Computer Sciences
Chairs: Xabier Abancens, European Exascale Accelerator Research Engineer, Computer Sciences
Barcelona Supercomputing Center - Centro Nacional de Supercomputación

---