Polytech network form for PhD Research Grants from the China Scholarship Council

This document describes the PhD subject and supervisor proposed by the French Polytech network of 14 university engineering schools. Please contact the PhD supervisor by email or Skype for further information regarding your application.

Supervisor information	
Family name	Lionel
First name	Torres
Email	lionel.torres@lirmm.fr
Web reference	www.lirmm.fr/~torres
Lab name	LIRMM : Laboratory of Informatic, Robotic and Microelectronic of Montpellier
Lab web site	http://www.lirmm.fr
Polytech name	Polytech Montpellier
University name	University of Montpellier
Country	France

PhD information	
Title	FPGA-accelerated computer architecture simulation to rethink data storage and movement in many-core systems
Main topics regards to CSC list (3 topics at maximum)	Structure of new computer systems Techniques of simulation and application

Required skills in science and engineering	Good knowledge of HDL and FPGA design, fluent programming skills and basic knowledge of computer architecture.

Subject description (two pages maximum)

Innovation in computer architectures is paramount to the mobile and datacenter computing revolutions. Since building an actual computer is hugely time-consuming and expensive, architects have long resorted to software simulators, such as gem5, to explore and evaluate novel computer architecture alternatives. However, the growing upheaval in complexity resulting from multicore architectures is rendering software simulation impractical: simulating in software a modern multicore computer is about a million times slower than real hardware (a 1-second execution becomes more than a 10-days simulation!). Importantly, this problem will only grow worse in the future since software simulators will invariably become slower as the number of cores keeps on increasing.

In this thesis, we will leverage FPGA acceleration to produce a computer architecture simulator tailored for the exploration of memory subsystem architectures for emerging memory technologies, such as the Spin-Transfer Torque Magnetic Random Access Memory (STT-MRAM). The FPGA acceleration will enable the simulation of many core systems executing representative applications in a reasonable time. However, not all the architecture object of the simulation will be implemented as FPGA hardware. Instead, we will run in software specific parts of the memory system functionality, such as memory controllers, on programmable soft-processors. The latter will enable selective flexibility on the architectural aspects of interest (i.e., the memory subsystem) avoiding the need to regenerate the FPGA bitstream -- a process that takes multiple hours in such complex designs.

Essential challenges of the thesis include:

1) Identification of architectural functionality to port onto software. To achieve maximum flexibility, one would want to port as much functionality as possible onto software. However, on the extreme, this will take us back to the slow software-based simulator we started from. Thus, the partitioning has to be done carefully to maintain the speedup of a pure FPGA simulator to the largest extent. Accordingly, suitable targets may include functionality of Last-Level-Cache (LLC) controllers (e.g., cache replacement policy) and main memory controllers.

2) Design of the interphase between the hardware and the soft-processors. To minimize the overhead, the communication between the hardware and the soft-processors has to be optimized. We will use streaming interphases, circular buffers, and smart controllers to enable as much as possible the parallel execution of the soft-processors and the hardware computer simulator.

3) Definition of a suitable programming abstraction to write code for the soft-processors. To facilitate the use of the simulator, we will define a high-level programming interphase for the end-user. This should at least support ANSI-C and include a library of commonly used functionality (e.g., communication methods between the hardware and the soft-processor) but can also be taken a leap forward by designing a simple Domain-Specific Language (DSL) embedded on C++.

4) Instruction-Set Architecture (ISA) extensions of the soft-processor to minimize the impact of software execution in the overall simulation time. Initially, we will start with basic MicroBlaze soft-processors, which is a soft microprocessor core designed for Xilinx. After profiling this basic implementation, we will identify the performance bottlenecks and propose extensions (i.e., multi-core and new hardware instructions) to minimize the gap with respect to a pure FPGA simulator.

5) Co-design of adequate SST-MRAM simulation models. Main memory and cache memory architectures will be modeled in the FPGA as a combination of hardware and software. We will include models based on the emerging STT technology as well as on current technology (e.g., SRAM and SDRAM), which will be used as reference.

The framework developed in this thesis will be based on RISC-V processors. More concretely, we will build on the Berkeley Rocket Chip, which is a parameterized SoC generator written in Chisel and the MIDAS simulation framework. Rocket Chip generates general-purpose in-order and out-of-order RISCV processor cores. Complementary, MIDAS automatically produces an FPGA-accelerated simulator from a Chisel-based RTL design. For the memory device modeling, we will rely on a strong partnership with a world-leading research and innovation hub in nanoelectronics and digital technologies with a strong background on STT-MRAM.

The resulting framework will enable to quickly simulate variations in the memory subsystem of many-core architectures while maintaining to a large extent the ease of use and programmability of current software simulators. Therefore, the use of the simulator will not require hardware design expertise, which is a very uncommon skill among computer architects.

The selected candidate will work together with an international group of postdocs, Ph.D. students, and faculty members from Polytech Montpellier and CNRS at the Montpellier Laboratory of Informatics, Robotics, and Microelectronics (LIRMM), France. LIRMM is a world-class university laboratory, known internationally for its expertise in a number of areas connected to computer architecture/computer sciences (Embedded, HPC, Pervasive), digital hardware (MPSoC, reconfigurable architectures, and hardware security) and emerging technologies (design of hybrid MRAM / CMOS devices). Altogether, LIRMM provides a state-of-the-art diverse, inclusive, and stimulating research environment.

Related publications:

[1] Evaluation of RISC-V RTL with FPGA-Accelerated Simulation. Donggyu Kim, Chris Celio, David T. Biancolin, Jonathan Bachrach, and Krste Asanović. In Proceedings of the Workshop on Computer Architecture Research with RISC-V (CARRV), 2017.

[2] Main memory organization trade-offs with DRAM and STT-MRAM options based on gem5-NVMain simulation frameworks. Manu Komalan, Oh Hyung Rock, Matthias Hartmann, Sushil Sakhare, Christian Tenlladoy, José Ignacio Gómezy, Gouri Sankar Kar, Arnaud Furnemont, Francky Catthoor, Sophiane Sennix, David Novo, Abdoulaye Gamatie, and Lionel Torres. In Proceedings of the Design, Automation and Test in Europe Conference and Exhibition (DATE), 2018.

[3] Firesim: FPGA-accelerated cycle-exact scale-out system simulation in the public cloud. Sagar Karandikar, Howard Mao, Donggyu Kim, David Biancolin, Alon Amid, Dayeol Lee, Nathan Pemberton, Emmanuel Amaro, Colin Schmidt, Aditya Chopra, Qijing Huang, Kyle D Kovacs, Borivoje Nikolic, Randy Katz, Jonathan Bachrach, and Krste Asanovic . In Proceedings of the 45th Annual International Symposium on Computer Architecture (ISCA), 2018.