Research Assistant Opening in the area of Parallel Computing Systems and Architectures in Fall 2017

My name is Yonghong Yan, and I will be joining CSE department from Fall 2017 as assistant professor. My areas of research include parallel and high performance computing, parallel programming models and compilers, computer architectures and systems. The research and development in my research group support computer, computational and data scientists by accelerating performance improvement, application development and energy efficiency improvement for their applications. We develop new parallel programming models based on OpenMP, compiler and runtime systems using Clang/LLVM. We improve and use software and hardware tools (PAPI, Pin, HPCToolkit, Lttng, Eclipse, etc) and modeling technology to deepen the understanding of parallel execution and resource utilizations of parallel applications. We explore new hardware architecture (GPU, FPGA, RISC-V-based architecture) for applications and we collaborate with scientists to apply our solutions to real world problems. We are also expanding our research horizons to the emerging computing paradigm including neuromorphic computing, quantum computing, cognitive technologies such as AI architecture and applications. More information can be found from https://passlab.github.io .

There are multiple openings for research assistants in my research group. Though the positions are for fall 2017, I am also happy to work with you during the summer remotely to get things started. We are looking for highly motivated graduate and undergraduate students who are interested in computer system and in developing innovative solutions of compiler and system software, as well as of computer architectures for parallel and high performance computing. Specific areas of research and development include:

1. Compiler analysis and transformation, loop transformation, optimization and code generation for GPUs, FPGAs and vector instructions, all based on Clang/LLVM;

2. Parallel runtime for heterogeneous systems with deep and complex memory hierarchy based on LLVM OpenMP runtime for NUMA, Xeon Phi and Hybrid accelerator systems;

3. Performance analysis and tools for parallel applications based on OpenMP OMPT interface, Lttng, Eclipse Tracecompass, Intel Pintool/Pinplay and PAPI;

4. Exploration and development of programming/compiler support, application and algorithm support for emerging architectures such as FPGA, neuromophic and quantum computers;

5. Exploration and development of new applications for NVIDIA Jetson TX1/TX2 embedded CPU/GPU credit-sized computer, including using it with drones, 3-D cameras and other sensors, and the use of its high-efficient computer vision and deep learning capability;

6. Computer architecture including a) the study and exploration of recent AI chips such as Google TPU, Intel Nervana, Graphcores, neuromophic architecture, etc, and b) RISC-V based extensions and development.

For performing productively in those areas, knowledge of C/C++ programming, data structures, compiler, computer architecture and operating system with Linux are required. Knowledge and experience of parallel programming (PThreads, OpenMP, Cilkplus, CUDA, MPI, etc) or computer architecture design such as using Verilog/VHDL, and RISC-V are not required, but preferred. Preference will be given to full-time Ph.D. students, thesis-option master students, and senior undergraduate students.

Within the group, you will work with a motivated team to solve cutting-edge performance problem in parallel computing. You will have opportunity of exposing yourself through peer-review publications, attendance and presentation in professional conferences worldwide, and internship in government labs and industry.

If you are interested, drop me an email at yanyh15@gmail.com with your resume/CV and transcript (no need official) that shows the core CS/ECE courses you took, as well as your area(s) of interest from the above 6 areas. We will schedule a time to talk if there is a fit of your background and interest.