Mao Lin

mlin59@ucmerced.edu https://www.linmao.org

Education

Ph.D. Student of Electrical Engineering and Computer Science — 08/2021 – present
University of California, Merced, CA, USA — Advisor: *Prof. Pengfei Su*

Master of Software Engineering — 09/2018 – 06/2021
Shandong University, Jinan, China — Advisor: *Prof. Lei Ju*

Bachelor of Computer Science and Technology — 09/2014 – 06/2018
Shandong University, Jinan, China

Research Interests

- Programming Languages
- Static and Dynamic Program Analysis
- High-performance/Parallel Computing (CUDA)
- Deep Learning

Experience

Ph.D. Intern of AI Accelerator Programming and Testing — 06/2022 – 08.2022
Pacific Northwest National Laboratory (PNNL), WA, USA — Mentor: *Dr. Ang Li*

- Develop a tool to pinpoint the floating-point data overflow of GPU-accelerated applications.

Publications

[TCAD'22] Du, Zelin, Qian Ling Zhang, **Mao Lin**, Shiqing Li, Xin Li, and Lei Ju. "A comprehensive memory management framework for CPU-FPGA heterogenous SoCs." IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (2022).

Project

Memory subsystem optimization for CPU-FPGA HMPSoCs — 08/2019 – 09/2020

- Focuses on the contentions occurring at both LLC and main memory levels on commercial off-the-shelf (COTS) CPU-FPGA heterogeneous multiprocessor system-on-chip (HMPSoC);
- Proposes a simulated annealing-based comprehensive memory management framework;
- Explores various cache optimization techniques, e.g., cache lockdown, cache partition, cache bypassing, to mitigate the memory resource contentions at both LLC and main memory levels;
- Improves overall system performance by 34% compared with state-of-the-art FPGA-only data allocation strategies on COTS Xilinx ZYNQ7020 HMPSoC.

Cache management for mobile games in CPU-GPU HMPSoC — 01/2019 – 12/2019

- Focuses on the energy inefficiencies incurred by memory subsystem on mobile devices;
- Adopts an eviction-buffer- and insertion-based cache management to reduce bandwidth saturation;
- Reduces the overall bandwidth usage by 11% compared with the LRU cache management strategy;
- Sponsored by a Chinese mobile phone company, and serves as a prototype of memory subsystem for their next-gen phone chip design.

UAV object detection system 08/2018 – 09/2019

- Adopts YOLOv2 neural network to achieve high performance and accuracy in tiny object detection, including cutting layers and fine-tuning layers;
- Employs multi-threading, pipelining, and half precision (fp16) for high inference speed;
- Deployed on the Robot Operating System to facilitate inter-module communication within the system;
- Achieves 90% accuracy and 90% recall, and reaches 30Hz for 1080 video input;
- Integrated into real aircraft by cooperating with a commercial UAV manufacturer.

AWARDS

- University admission scholarship *Dec. 2018*
- Graduate scholarship *Dec. 2019*
- Graduate scholarship *Dec. 2020*

TEACHING

- 2021 Fall, **Advanced Programming**, *CSE 024* *Teaching Assistant*
- 2021 Fall, **Intro to Object Orient Program**, *CSE 165* *Teaching Assistant*
- 2022 Spring, **Data Structure**, *CSE 030* *Teaching Assistant*

SKILLS

- Languages: C/C++; Python; Java; CUDA; Shell
- Platform: Linux; CPU-FPGA HMPSoCs; CPU-GPU HMPSoCs
- Framework: Darknet; TensorFlow
- Toolchains: Nsight Compute; Linux perf; GDB; Xilinx Vivado Suite