Gokulan Ravi

ravig@purdue.edu | **https://gokulan97.github.io**

OBJECTIVE

Research engineer interested in all aspects of modern CPU design, implementation and verification

EDUCATION

Ph.D. in Electrical & Computer Engineering Jan 2021 - Present
Purdue University, West Lafayette, USA GPA: 3.87/4.00

B.Tech + M.Tech (Dual Degree) in Computer Science & Engineering Jul 2015 - May 2020
Indian Institute of Technology Madras, Chennai, India GPA: 8.23/10.00

SKILLS

Relevant Courses: Computer Architecture, Parallel Computer Architecture, GPU Programming, Concurrent Programming, Operating Systems, Compilers

Languages: C/C++, Java, Python (PyTorch), Bluespec System Verilog

RESEARCH PROJECTS

Formal verification of modern processors, *Purdue University* May 2021 - present
- Applying domain knowledge from computer architecture to formal verification techniques

Hardware accelerator for Graph Neural Networks, *Purdue University* Jan 2021 - May 2021
- Aimed at exploiting ultra sparsity of graphs to accelerate GNNs
- Reported that GPUs SpMM is a good fit and no further opportunity remained to be explored

Systolic-array based co-processor for Deep Neural Network inference Jun 2019 - May 2020
Masters thesis, IIT Madras
- Led the implementation of a parameterized systolic-array based co-processor in Bluespec
- Built an end-to-end working solution with a prototype compiler based on TVM

MOESI Cache Coherence for in-order processor, *IIT Madras* May - Jun 2018, Jan - May 2019
Summer Research Project (Snoop-based protocol), Senior Design project (Directory-based protocol)
- Modified memory subsystem of a 5-stage RISC-V core to enable multicore extensions
- Designed the finite state machine and implemented the entire cache controller in Bluespec
- Implemented and customized Tilelink-Cacheable (TL-C) bus protocol to support P2P transactions

GPU kernels for weight sharing in CNNs, *Course project, IIT Madras* Oct 2018 - Feb 2019
- Proposed an optimization to improve convolution by exploiting redundancy of quantized weights
- Measured overheads in execution and reported 1.3x slowdown

PROFESSIONAL SERVICE

Member of Artifact Evaluation Committee, MICRO 54 Aug 2021
Reproduced and verified results of one paper in MICRO 54's proceedings

PROFESSIONAL EXPERIENCE

Graduate Teaching Assistant Aug 2019 - Present
- Purdue University - Operating Systems, Object Oriented Programming, Compilers
- IIT Madras - Computer System Design, Systems for Deep Learning

Software Engineering Internship, Quiklo, Bengaluru May 2017 - Jul 2017
- Built image processing pipeline to extract bank transactions from images with 99% accuracy

EXTRACURRICULAR ACTIVITIES

Head, Web Operations, Shaastra 2018, IIT Madras Mar 2017 - Jan 2018
- Led a team of 20 to build the software stack of the technical symposium
- Unified different user-facing applications, streamlined internal operations, increased registrations by 80%

Instructor, Webops Club, Center for Innovation, IIT Madras May 2017 - Dec 2019
Organized various workshops in web, mobile app and Python development for undergraduates