

TUAN QUANG TA

471-B Rhodes Hall, Ithaca, NY 14853
qtt2@cornell.edu – <http://tuanqta.net>

RESEARCH INTERESTS

Vector Architectures, Many-Core Architectures & Heterogeneous Architectures

EDUCATION

Cornell University

2017 - Present

Doctor of Philosophy in Electrical & Computer Engineering
Advisor: Prof. Christopher Batten

University of Mississippi

2012 - 2016

Bachelor of Science in Computer Science
Minor in Mathematics for Engineering
Overall GPA: 4.0 / 4.0 — Summa Cum Laude

PUBLICATIONS

Khalid Al-Hawaj, **Tuan Ta**, Nick Cebry, Shady Agwa, Olalekan Afuye, Eric Hall, Courtney Golden, Alyssa Apsel, Christopher Batten. “EVE: Ephemeral Vector Engines.” In submission to *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023.

Tuan Ta, Khalid Al-Hawaj, Nick Cebry, Yanghui Ou, Eric Hall, Courtney Golden, and Christopher Batten. “big.VLITTLE: On-Demand Data-Parallel Acceleration for Mobile Systems on Chip.” *International Symposium on Microarchitecture (MICRO)*, 2022.

Christopher Batten, Pjotr Prins, Efraim Flashner, Arun Isaac, Ekaiz Zarraga, Erik Garrison, **Tuan Ta**. “The Case for Using Guix to Solve the gem5 Packaging Problem.” *gem5 Users’ Workshop held in conjunction with ISCA-49*, June. 2022.

Christopher Batten, Pjotr Prins, Efraim Flashner, Arun Isaac, Jan Nieuwenhuizen, Ekaiz Zarraga, **Tuan Ta**, Austin Rovinski, Erik Garrison. “The Case for Using Guix to Enable Reproducible RISC-V Software & Hardware.” *6th Workshop on Computer Architecture Research with RISC-V (CARRV) held in conjunction with ISCA-49*, June 2022.

Lingjun Zhu, **Tuan Ta**, Rossana Liu, Rahul Mathur, Xiaoqing Xu, Shidhartha Das, Ankit Kaul, Alejandro Rico, Doug Joseph, Brian Cline, Sung Kyu Lim. “Power Delivery and Thermal-Aware Arm-Based Multi-Tier 3D Architecture” *IEEE/ACM International Symposium on Low Power Electronics and Design*, July 2021.

Moyang Wang, **Tuan Ta**, Lin Cheng, and Christopher Batten. “Efficiently Supporting Dynamic Task Parallelism on Heterogeneous Cache-Coherent Systems.” *47th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, June 2020.

Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adrià Armejach, Nils Asmussen, Srikant Bharadwaj, Gabe Black, Gedare Bloom, Bobby R. Bruce, Daniel Rodrigues Carvalho, Jerónimo Castrillón, Lizhong Chen, Nicolas Derumigny, Stephan Diestelhorst, Wendy Elsasser, Marjan Fariborz, Amin Farmahini Farahani, ..., **Tuan Ta**, ..., Éder F. Zulian, “The gem5 Simulator: Version 20.0+.” *arXiv preprint arXiv:2007.03152*, 2020.

Tuan Ta, Xianwei Zhang, Anthony Gutierrez, and Bradford M. Beckmann. “Autonomous Data-Race-Free GPU Testing.” *IEEE International Symposium on Workload Characterization (IISWC)*, Nov. 2019.

David Troendle, **Tuan Ta**, and Byunghyun Jang. “A Specialized Concurrent Queue for Scheduling Irregular Workloads on GPUs.” *48th International Conference on Parallel Processing (ICPP)*, Aug. 2019.

Christopher Torng, Shunning Jiang, Khalid Al-Hawaj, Ivan Bukreyev, Berkin Ilbeyi, **Tuan Ta**, Lin Cheng, Julian Pucar, Ian Galton, and Christopher Batten. “A New Era of Silicon Prototyping in Computer Architecture Research.” *RISC-V Day Workshop at the 51st International Symposium on Microarchitecture*, Oct. 2018.

Tuan Ta, Lin Cheng, and Christopher Batten. “Simulating Multi-Core RISC-V Systems in gem5.” *2nd Workshop on Computer Architecture Research with RISC-V (CARRV)* held in conjunction with ISCA-45, June 2018.

Elliott Samuel, Raghu Raj Prasanna Kumar, Natasha Flyer, **Tuan Ta**, and Richard Loft. “Implementation of a Scalable, Performance Portable Shallow Water Equation Solver Using Radial Basis Function-Generated Finite Difference Methods.” *International Journal of High Performance Computing Applications (IJHPCA)*, 2018

Tuan Ta, David Troendle, Xiaoqi Hu, and Byunghyun Jang. “Understanding the Impact of Fine-Grained Data Sharing and Thread Communication on Heterogeneous Workload Development.” *16th IEEE International Symposium on Parallel & Distributed Computing (ISPDC)*, July 2017.

Tuan Ta, David Troendle, and Byunghyun Jang. “Thread Communication and Synchronization on Massively Parallel GPUs.” *A chapter in Advances in GPU Research and Practice book* edited by Hamid Sarbazi-Azad, 2016.

Tuan Ta, Kyoshin Choo, Eh Tan, Byunghyun Jang, and Eunseo Choi. “Accelerating DynEarthSol3D on Tightly Coupled CPU-GPU Heterogeneous Processors.” *Computers & Geosciences Journal*, June 2015.

ACADEMIC EXPERIENCE

Batten Research Group
Graduate Research Assistant

Aug, 2017 - Present
Cornell University

- *big.VLITTLE architecture* – taking a lead in developing a novel on-demand reconfigurable next-generation vector architecture for heterogeneous mobile systems on chip with minimal area overheads
- *Ephemeral vector engines* – contributed to developing a compute-in-memory SRAM-based solution to mitigate vector execution’s area overhead in general-purpose compute systems
- *CIFER chip tapeout in collaboration with Princeton University* – led a team of graduate and undergraduate students at Cornell University to develop and integrate many-core compute tiles into a heterogeneous system on chip taped out using an advanced technology node
- *Heterogeneous cache coherence* – contributed to developing heterogeneous cache coherence to enable cooperative execution of different compute cores in future heterogeneous many-core systems
- *Development of gem5* – supported simulating multi-thread RISC-V workloads, developed an in-order CPU model, and contributed to modeling an integrated vector execution unit and a decoupled vector engine in gem5

HEROES Research Group
Research Associate

Jun, 2016 - Mar, 2017
University of Mississippi

- Studied different coherence protocols for tightly coupled CPU/GPU systems
- Took leading efforts in preparing a funded NSF grant proposal that aims at advancing the memory subsystem of tightly coupled CPU/GPU heterogeneous processors
- Characterized the performance of multiple CPU-GPU cooperation paradigms in fine-grained data sharing CPU-GPU systems

HEROES Research Group
Undergraduate Research Assistant

Dec, 2013 - May, 2016
University of Mississippi

- Designed concurrent data structures including linked list, queue and d-ary heap for CPU-GPU heterogeneous systems
- Accelerated an unstructured mesh-based simulator, DynEarthSol3D, used to study the long-term deformation of Earth's lithosphere on GPU using OpenCL

INDUSTRY RESEARCH EXPERIENCE

Arm Research
Research Intern

Jun, 2021 - Aug, 2021
Remote from Ithaca, NY

- Explored the potential of using Arm's Scalable Matrix Extension (SME) in accelerating sparse matrix & vector computation in emerging graph and machine learning workloads
- Proposed new instructions to Arm SME to increase the performance of such sparse matrix & vector computation

Arm Research
Research Intern

Jun, 2020 - Aug, 2020
Remote from Ithaca, NY

- Modeled wafer-scale many-core systems in gem5
- Explored the design space of scaling future many-core systems using multiple chiplets and state-of-the-art 3D integration technology

AMD Research
Research Engineering Co-op

Mar, 2017 - Jul, 2017
Bellevue, WA

- Modeled a next-generation AMD GPU's memory subsystem in gem5 simulator
- Evaluated the performance of emerging data-parallel workloads in the next-generation AMD GPU
- Developed a random testing methodology for cache coherence and weak consistency model in modern GPU and heterogeneous CPU-GPU systems

National Center for Atmospheric Research
Undergraduate Research Intern

Mar, 2015 - Aug, 2015
Boulder, CO

- Parallelized and accelerated Shallow Water Equations (SWE) using Radial Basis Function Finite Difference (RBF-FD) method on multi-core CPUs and GPUs using different programming models
- Analyzed a trade-off between portability and performance of OpenCL programming model in SWE using RBF-FD in comparison with OpenMP and CUDA on multi-core CPUs and GPUs

TEACHING EXPERIENCE

Cornell University
Lead Graduate Teaching Assistant
ECE-2400 Computer Systems Programming

Fall, 2018

University of Mississippi
Undergraduate Lab & Teaching Assistant

2013 - 2016

HONORS & AWARDS

Jacobs Scholar Fellowship
A one-semester fellowship given to PhD students by School of Electrical and Computer Engineering

Cornell University, 2017

Summer Student of the Tour

AMD, 2017

An award recommended by managers and mentors at AMD to recognize significant contributions of outstanding student interns and co-ops to the company

Taylor Medal

University of Mississippi, 2016

The highest academic honor conferred by the University of Mississippi

Outstanding Computer Science Student Awards University of Mississippi, 2014, 2015 & 2016

Annual awards recommended by faculty members in Department of Computer and Information Science

Computer Science SAP Scholarship

University of Mississippi, 2015

An award recommended by faculty members in Department of Computer and Information Science

Second place in a Student Poster Competition at RMACC conference

2015

Won a trip to Supercomputing Conference (SC15) in Austin, TX in Nov 2015

International Undergraduate Student Scholarship

University of Mississippi, 2014

A two-year scholarship awarded by Office of International Programs to outstanding international students on account of their academic excellence and achievement

Academic Excellence Scholarship

University of Mississippi, 2012 – 2016

A four-year full-tuition scholarship awarded by Office of International Programs to a handful of outstanding undergraduate students at University of Mississippi

John G. Adler Engineering Scholarship

University of Mississippi, 2012 – 2016

A four-year scholarship awarded by School of Engineering to a handful of undergraduate students at University of Mississippi