

Kai (Tommy) Zhao

tzhao451@gmail.com 1(312)375-8295 Sunnyvale, CA, USA

Website: <http://pages.cs.wisc.edu/~kzhao32/>

E Education	University of Wisconsin-Madison, Madison, WI <i>Sept 2014 – May 2017</i> M. S. in Computer Science, <i>Dec 2015</i> 3.72/4.00 Overall GPA
	University of Illinois at Chicago, Chicago, IL <i>Sept 2010 – May 2014</i> B. S. in Electrical and Computer Engineering with Honors and with minors in CS, Math, and Biology 3.84/4.00 Overall GPA, 4.00 ECE major GPA, 4.00 CS minor GPA, 4.00 math minor GPA
W Work Experience	Software Engineer, Google, Sunnyvale, CA <i>July 2018 – present</i> <ul style="list-style-type: none"> • Hardware and software co-design for Google search to improve perf/cost • Evaluate performance of deep learning machine learning models for search • Evaluate fleet-wide resource utilization for all search clusters
	CPU Diagnostics Engineer, Advanced Micro Devices (AMD), Austin, TX <i>Aug 2017 – July 2018</i> <ul style="list-style-type: none"> • Developed software tool to generate test cases to validate processors post-silicon • Used instruction generators to stress micro-architectural features • Found, reported, and root caused silicon hardware bugs
	Teaching Assistant, University of Wisconsin-Madison <i>CS/ECE 354 Machine Organization and Programming</i> <i>Spring 2017</i> <i>CS 302 Introduction to Programming</i> <i>Summer 2016</i> <i>CS 368 Learning a Programming Language: MATLAB Programming</i> <i>Spring 2016</i> <i>CS/ECE 252 Introduction to Computer Engineering</i> <i>Fall 2015</i> <ul style="list-style-type: none"> • Created and graded homework, quizzes, and exams • Hosted office hours, answered questions on discussion board, and led review sessions • Wrote scripts to assist in grading programming assignments
	Lecturer, CS/ECE 252 Introduction to Computer Engineering; UW-Madison <i>Fall 2016</i> <ul style="list-style-type: none"> • Prepared and presented lectures for a class of 132 students • Created and graded exams and homework • Hosted office hours and answer questions on discussion board
	Embedded Engineer, Extreme Engineering Solution (X-ES), Middleton, WI <i>Jan 2016 – Sept 2016</i> <ul style="list-style-type: none"> • Implemented, tested, debugged, and released VxWorks for custom hardware • Resolved VxWorks issues involving spin locks, page faults, and virtual memory maps
	ECE Laboratory Research Assistant, University of Illinois at Chicago <i>June 2012 – June 2014</i> <ul style="list-style-type: none"> • Researched achieving fair resource distribution under connectivity constraints • Discussed and developed algorithms, implementations, visualizations, and simulations • Developed the simulation tool set to verify hypotheses and test algorithm efficiencies
	Programming Languages: C (proficient), C++, C#, Cilk, CSS, CUDA, F#, HTML, Java (proficient), JavaScript, LC-3 assembly, Lisp, MATLAB (proficient), MIPS assembly (proficient), MPI, OpenGL, OpenMP, Python, Ruby (proficient), Scheme, Scratch, SQL, Verilog, WebGL, and x86 assembly
P Projects	Exploring the parallel programming design space of Proximate, a multi-tile programmable hardware accelerator <i>Programming Multicore Processors</i> <ul style="list-style-type: none"> • Implemented parallel workloads on different parts of a Proximate compute tile • Investigated scalability and speedup to improve hardware design choices <i>Fall 2016</i>
	Tronsistor-32, a Nintendo Entertainment System (NES) Clone <i>Digital Engineering Laboratory</i> <ul style="list-style-type: none"> • Designed an Instruction Set Architecture (ISA) for computation and graphics • Used Verilog for FPGA programming to create a NES Clone and support the ISA • Used assembly language in the ISA to write Tron, DDR, and Pac-Man <i>Fall 2015</i>
	Bonsai, a distributed data collection and storage system for data processing <i>Advanced Operating Systems</i> <ul style="list-style-type: none"> • Scalable data collection in a ring structure by implementing a circular buffer for stored data with a token as synchronization logic • Real-time data processing in a tree structure by multicasting map requests to backends and reducing results towards the frontend <i>Fall 2014</i>