

Yi-Hua Chung

Ph.D. Student in ECE, University of Wisconsin-Madison

✉ yihua.chung@wisc.edu ☎ +1 (608)-692-3491 🌐 Personal Blog 📄 Google Scholar 🌐 Yi-Hua Chung

EDUCATION

Ph.D. Engineering in Electrical and Computer Engineering, 09/2023 – Present
University of Wisconsin-Madison, Advisor: Tsung-Wei Huang

- *Research Interests:*
 - GPU-parallel compiler optimization via equality saturation
 - GPU-accelerated EDA applications (logic simulation, gate sizing, logic synthesis)
- *Teaching:* Advanced Computer Architecture II
- *GPA:* 4.00/4.00 (Coursework completed)

M.S. Computer Science, National Taiwan University 02/2021 – 08/2022

- *Research Topic:* GPU-accelerated algorithms and system performance optimization for quantum computing simulations.
- *Master Thesis:* Enlarging Quantum Circuit Simulation and Analysis with Non-Volatile Memories
- *Teaching:* Computer Architecture
- *GPA:* 4.25/4.30 (Rank: 1/47)

B.S. Biomechatronics Engineering, National Taiwan University 09/2016 – 01/2021

- *Bachelor Thesis:* Development of a Small Intelligent Weather Station for Agricultural Applications

WORK EXPERIENCE

Student Researcher 08/2026 – present

Recursive Intelligence, CA; Manager: Anna Goldie, Mentor: Yi-Chen Lu

- Developing GPU-accelerated optimization techniques for AI-driven chip design workflows.

Graduate Research Assistant 08/2023 – present

Department of ECE, UW-Madison; Advisor: Prof. Tsung-Wei Huang

- Developed XEG, a GPU-parallel e-graph extractor, achieving a **243× speedup** over a state-of-the-art GPU-based e-graph extractor while maintaining comparable solution quality (ASPLOS'27).
- Developed SimPart, a GPU-parallel graph partitioner for logic simulation, achieving **23× faster** partitioning and **1.58× faster** GPU logic simulation runtime over the state-of-the-art method (Euro-Par'25).

Technical Intern; R&D Team, EDA Group 06/2024 – 12/2024

Synopsys Inc, CA; Manager: Malleswara Gupta, Mentor: Nahmsuk Oh

- Developed GPU-parallel kernel algorithms, achieving **38.13× speedup** over a 16-core CPU industrial sizer (Euro-Par'25).
- Integrated GPU kernels into a Synopsys Fusion Compiler to enable heterogeneous **CPU-GPU co-processing** for gate sizing.

Full-time Research Assistant 08/2022 – 03/2023

Institute of Information Science, Academia Sinica; Prof. Bo-Yin Yang

- Developed GPU-accelerated big-integer multiplication with NVIDIA-level performance for post-quantum cryptosystems.
- Accelerated NTRU and NTRU Prime lattice-based cryptosystems on Cortex-A72 with vectorized polynomial multipliers, achieving up to **6.7× faster** multiplications and **7.67× faster** key generation compared to state-of-the-art.

Research Assistant 07/2021 – 08/2022

Department of CSIE, NTU; Prof. Shih-Hao Hung

- Accelerated simulated quantum annealing on a GPU, achieving up to **86.6× speedup** over the state-of-the-art.
- Developed an NVM-based quantum circuit simulator and a circuit scheduler, achieving **1.2× speedup** over the QuEST simulator.
- Led a study group and assisted labmates on large-scale simulated quantum annealing on multi-GPUs.

Teaching Assistant, Advanced Computer Architecture II 01/2026 – 06/2026

Department of ECE, UW-Madison; Instructor: Prof. Joshua San Miguel

- Built a cross-platform Docker environment for gem5 simulation and benchmarking on Linux, macOS, and Windows.

Teaching Assistant, Computer Architecture 09/2021 – 01/2022

Department of CSIE, NTU; Instructor: Shih-Hao Hung

- Developed Verilog-based labs covering ALU, FPU, and pipelined RISC-V CPU designs.

PUBLICATIONS

- **Yi-Hua Chung**, Hao-Hsiang Hsiao, Cheng-Hsiang Chiu, Boyang Zhang, Che Chang, Joshua San Miguel, and Tsung-Wei Huang, "XEG: A GPU-Parallel Algorithm for Efficient and Effective E-Graph Extraction," *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Heraklion, Crete, Greece, 2027

- Chih-Chun Chang, **Yi-Hua Chung**, Wan-Luan Lee, Boyang Zhang, and Tsung-Wei Huang, "SSTA-X: GPU-Accelerated First-Order Block-Based Statistical Static Timing Analysis," *IEEE Transactions on Computer-aided Design of Integrated Circuits and Systems (TCAD)*, to appear, 2026
- Shui Jiang, Wen Cheng, **Yi-Hua Chung**, Tsung-Yi Ho, and Tsung-Wei Huang, "QDP: Worst-case Fidelity-aware Qubit Mapping and Routing using Dynamic Programming," *IEEE Transactions on Computer-aided Design of Integrated Circuits and Systems (TCAD)*, to appear, 2026
- Boyang Zhang, Chih-Chun Chang, **Yi-Hua Chung**, Che Chang, Cheng-Hsiang Chiu, Aditya Das Sarma, and Tsung-Wei Huang, "G-STAR: GPU-Accelerated Statistical Static Timing Analysis using Level-by-level Replication," *International European Conference on Parallel and Distributed Computing (Euro-Par)*, Pisa, Italy, 2026
- Che Chang, **Yi-Hua Chung**, Cheng-Hsiang Chiu, Wan-Luan Lee, Boyang Zhang, Ulf Schlichtmann, Ing-Chao Lin, Xiangyao Yu, and Tsung-Wei Huang, "G-PathGen: An Efficient GPU-Parallel k-Critical Path Generation Algorithm", *ACM International Conference on Supercomputing (ICS)*, Belfast, Northern Ireland, UK, 2026
- **Yi-Hua Chung**, Shui Jiang, Wan Luan Lee, Yanqing Zhang, Haoxing Ren, Tsung-Yi Ho, and Tsung-Wei Huang, "SimPart: A Simple Yet Effective Replication-aided Partitioning Algorithm for Logic Simulation on GPU," *International European Conference on Parallel and Distributed Computing (Euro-Par)*, Dresden, Germany, 2025
- **Yi-Hua Chung**, Nahmsuk Oh, Malleswara Gupta Balabhadra Naga Venkata, Aditya Shiledar, Sudipto Kundu, Vishal Khandelwal, and Tsung-Wei Huang, "Accelerating Gate Sizing using GPU," *International European Conference on Parallel and Distributed Computing (Euro-Par) PhD Symposium*, Dresden, Germany, 2025
- Wan-Luan Lee, Shui Jiang, Dian-Lun Lin, Che Chang, Boyang Zhang, **Yi-Hua Chung**, Ulf Schlichtmann, Tsung-Yi Ho, and Tsung-Wei Huang, "iG-kway: Incremental k-way Graph Partitioning on GPU," *ACM/IEEE Design Automation Conference (DAC)*, San Francisco, CA, 2025
- Cheng-Hsiang Chiu, Wan-Luan Lee, Boyang Zhang, **Yi-Hua Chung**, Che Chang, and Tsung-Wei Huang, "A Task-parallel Pipeline Programming Model with Token Dependency," *Workshop on Asynchronous Many-Task Systems and Applications (WAMTA)*, St. Louis, MO, 2025
- Shui Jiang, **Yi-Hua Chung**, Chih-Chun Chang, Tsung-Yi Ho, and Tsung-Wei Huang, "BQSim: GPU-accelerated Batch Quantum Circuit Simulation using Decision Diagram," *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Rotterdam, Netherlands, 2025
- Boyang Zhang, Che Chang, Cheng-Hsiang Chiu, Dian-Lun Lin, Yang Sui, Chih-Chun Chang, **Yi-Hua Chung**, Wan-Luan Lee, Zizheng Guo, Yibo Lin, and Tsung-Wei Huang, "iTAP: An Incremental Task Graph Partitioner for Task-parallel Static Timing Analysis," *IEEE/ACM Asia and South Pacific Design Automation Conference (ASP-DAC)*, Tokyo, Japan, 2025
- Che Chang, Boyang Zhang, Cheng-Hsiang Chiu, Dian-Lun Lin, **Yi-Hua Chung**, Wan-Luan Lee, Zizheng Guo, Yibo Lin, and Tsung-Wei Huang, "PathGen: An Efficient Parallel Critical Path Generation Algorithm," *IEEE/ACM Asia and South Pacific Design Automation Conference (ASP-DAC)*, Tokyo, Japan, 2025
- Chen, Han-Ting, **Yi-Hua Chung**, Vincent Hwang, and Bo-Yin Yang. "Algorithmic Views of Vectorized Polynomial Multipliers–NTRU," In *International Conference on Cryptology in India*, pp. 177-196. Cham: Springer Nature Switzerland, 2023
- Chen, Han-Ting, **Yi-Hua Chung**, Vincent Hwang, Chi-Ting Liu, and Bo-Yin Yang. "Algorithmic Views of Vectorized Polynomial Multipliers for NTRU and NTRU Prime (Long Paper)," *Cryptology ePrint Archive*, Report 2023/541, 2023.
- **Yi-Hua Chung**. "Enlarging Quantum Circuit Simulation and Analysis with Non-Volatile Memories," Master's thesis, National Taiwan University, 2022
- **Yi-Hua, Chung**, Cheng-Jhih Shih, and Shih-Hao Hung. "Accelerating Simulated Quantum Annealing with GPU and Tensor Cores," In *International Conference on High Performance Computing*, pp. 174-191. Cham: Springer International Publishing, 2022
- **Yi-Hua, Chung**, Huang Jun-Fu, Hu Yuan-Chen, and Huang Chen-Kang. "Development of a Small Intelligent Weather Station for Agricultural Applications," *Advances in Technology Innovation* 6, no. 2 (2021): 74

TALKS & PRESENTATIONS

- "SimPart: A Simple Yet Effective Replication-aided Partitioning Algorithm for Logic Simulation on GPU," Oral Presentation, Euro-Par (Dresden, Aug 2025)
- "Accelerating Gate Sizing using GPU," Elevator Pitch & Poster Presentation, Euro-Par PhD Symposium (Dresden, Aug 2025)
- "Scalable Code Generation for RTL Simulation of Deep Learning Accelerators with MLIR," Oral Presentation (on behalf of first author), Euro-Par (Dresden, Aug 2025)
- "GPU-Accelerated Gate Sizing in Synopsys EDA Tool," Intern Final Presentation, Synopsys EDA Group (Sunnyvale, Aug 2024)
- "Introduction to GPU Computing for EDA," Invited Talk, Synopsys EDA Group (Sunnyvale, Jun 2024)
- "Accelerating Simulated Quantum Annealing with GPU and Tensor Cores," Oral Presentation, ISC (Hamburg (Remote), May 2022)
- "Development of a Small Intelligent Weather Station for Agricultural Applications," Oral Presentation, IMETI (Taichung, Oct 2020)

AWARDS

- **Department of ECE The Gerald Holdridge Teaching Excellence Award**, 2026
- **ACM/IEEE DAC Young Student Fellowship**, 2024 & 2025
- **NTUEE-1975 Innovation and Entrepreneurship Fund Award**
- **2022 Future Tech Awards**, National Science and Technology Council, R.O.C.
- **Best Paper Award**, 9th International Multi-Conference on Engineering and Technology Innovation 2020
- **Outstanding Performance Award**, NTU-IBM Q System 2020 Q-Camp, Hackathon 2020

SKILLS

C/C++, CUDA C++, OpenMP, ARM Intrinsic, ARM Assembly, Linux, Shell — *Expert* |

Python, C#, Qiskit, JavaScript, WebGL — *Experienced*