

Hung-Wei Tseng

Associate Professor
Electrical and Computer Engineering
University of California, Riverside

(951)827-9347
htseng@ucr.edu
<https://intra.engr.ucr.edu/~htseng/>

Appointments

Associate Professor University of California, Riverside Department of Electrical and Computer Engineering Cooperating Faculty, Department of Computer Science and Engineering	2022-present
Assistant Professor University of California, Riverside Department of Electrical and Computer Engineering Cooperating Faculty, Department of Computer Science and Engineering	2019-2022
Assistant Professor North Carolina State University Department of Computer Science	2016-2019
Postdoctoral Scholar University of California, San Diego Department of Computer Science and Engineering Advisor: Steven Swanson	2014-2016

Education

Ph.D., Computer Science & Engineering University of California, San Diego Dissertation: Data-triggered threads Advisor: Dean Tullsen and Steven Swanson	2014
M.S., Computer Science National Taiwan University Dissertation: An Energy-Efficient Virtual Memory System with Flash Memory as the Secondary Storage Advisor: Chia-Lin Yang	2005
B.S., Computer Science National Taiwan University	2003

Refereed Journal and Conference Publications

1. Kuan-Chieh Hsu and Hung-Wei Tseng. Simultaneous and Heterogenous Multithreading. *IEEE Micro, Special Issue on the Top Picks from Computer Architecture Conferences*, 2024.
2. Kuan-Chieh Hsu and Hung-Wei Tseng. Simultaneous and Heterogenous Multithreading. In *The 56th IEEE/ACM International Symposium on Microarchitecture, MICRO 2023*, 2023.
3. Dongho Ha, Won Woo Ro, and Hung-Wei Tseng. Accelerating ML-adjacent Computation Using Tensor Processors. In *The 2023 ACM/IEEE International Symposium on Low Power Electronics and Design, ISLPED 2023*, 2023.
4. Yu-Chia Liu, Kuan-Chieh Hsu, and Hung-Wei Tseng. Rethinking Programming Frameworks for In-Storage Processing. In *60th Design Automation Conference, DAC 2023*, 2023.
5. Yunan Zhang, Po-An Tsai, and Hung-Wei Tseng. SIMD²: A Generalized Matrix Instruction Set for Accelerating Tensor Computation beyond GEMM. In *49th International Symposium on Computer Architecture, ISCA 2022*, 2022.

6. Gunjae Koo, Yunho Oh, Hung-Wei Tseng, Won Woo Ro, and Murali Annavaram. "FLIXR: Embedding Index into Flash Translation Layer in SSDs. *IEEE Transactions on Computers*, 2022.
7. Yu-Ching Hu, Yuliang Li, and Hung-Wei Tseng. TCUDB: Accelerating Database with Tensor Processors. In *the 2022 ACM SIGMOD/PODS International Conference on Management of Data*, SIGMOD 2022, 2022.
8. Yu-Chia Liu and Hung-Wei Tseng. NDS: N-Dimensional Storage. In *54th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 2021 (Best Paper Nomination), 2021.
9. Kuan-Chieh Hsu and Hung-Wei Tseng. Accelerating Applications using Edge Tensor Processing Units. In *SC: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, SC 2021, 2021.
10. Alec Rohloff, Zackary Allen, Kung-Min Lin, Joshua Okrend, Chengyi Nie, Yu-Chia Liu, and Hung-Wei Tseng. OpenUVR: Reducing Latency and Preserving Quality in Untethered Real-Time Virtual-Reality Gaming Platforms. In *27th IEEE Real-Time and Embedded Technology and Applications Symposium*, RTAS 2021 (Outstanding Paper Award), 2021.
11. Jinyoung Choi, Sergey Blagodurov, and Hung-Wei Tseng. Dancing in the Dark: Profiling in the Age of Tiered Memory. In *35th IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2021, 2021.
12. Abenezer Wudenhe and Hung-Wei Tseng. TPUPoint: Automatically Characterizing Hardware Accelerated Data Center Machine Learning Program Behavior. In *2021 IEEE International Symposium on Performance Analysis of Systems and Software*, ISPASS 2021, 2021.
13. Yu-Ching Hu, Murtuza Taher Lokhandwala, Te I, and Hung-Wei Tseng. Varifocal Storage: Dynamic Multi-Resolution Data Storage. *IEEE Micro, Special Issue on the Top Picks from Computer Architecture Conferences*, 2020.
14. Yu-Ching Hu, Murtuza Taher Lokhandwala, Te I, and Hung-Wei Tseng. Dynamic Multi-Resolution Data Storage. In *52th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 2019 (Best Paper Honorable Mention), 2019.
15. Kiran Kumar Matam, Gunjae Koo, Haipeng Zha, Hung-Wei Tseng, and Murali Annavaram. GraphSSD: Graph Semantics Aware SSD. In *46th International Symposium on Computer Architecture*, ISCA 2019, 2019.
16. Te I, Murtuza Lokhandwala, Yu-Ching Hu, and Hung-Wei Tseng. Pensieve: a Machine Learning Assisted SSD Layer for Extending the Lifetime. In *IEEE International Conference on Computer Design*, ICCD 2018, 2018.
17. Hung-Wei Tseng, Qianchen Zhao, Yuxiao Zhou, Mark Gahagan, and Steven Swanson. Morpheus: Exploring the Potential of Near-Data Processing for Creating Application Objects in Heterogeneous Computing. *SIGOPS Operating Systems Review*, volume 51(2):71–83, August 2018.
18. Gunjae Koo, Kiran Kumar Matam, Te I, Hema Venkata Krishna Giri Narra, Jing Liand Hung-Wei Tseng, Steven Swanson, and Murali Annavaram. Summarizer: Trading Bandwidth with Computing Near Storage. In *50th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 2017, 2017.
19. Yanqin Jin, Hung-Wei Tseng, Yannis Papakonstantinou, and Steven Swanson. Improving SSD lifetime with byte-addressable metadata. In *Proceedings of the International Symposium on Memory Systems*, MEMSYS 2017, pages 374–384, 2017.
20. Yanqin Jin, Hung-Wei Tseng, Steven Swanson, and Yannis Papakonstantinou. KAML: A Flexible, High-Performance Key-Value SSD. In *23th International Symposium on High Performance Computer Architecture*, HPCA 2017, 2017.

21. Jing Li, Hung-Wei Tseng, Chunbin Lin, Steven Swanson, and Yannis Papakonstantinou. HippogriffDB: Balancing I/O and GPU Bandwidth in Big Data Analytics. *Proceedings of the VLDB Endowment*, volume 9(14), 2016.
22. Hung-Wei Tseng, Qianchen Zhao, Yuxiao Zhou, Mark Gahagan, and Steven Swanson. Morpheus: Creating Application Objects Efficiently for Heterogeneous Computing. In *43rd International Symposium on Computer Architecture*, ISCA 2016, 2016.
23. Yang Liu, Hung-Wei Tseng, Mark Gahagan, Jing Li, Yanqin Jin, Logan Gunthorp, Stephen Bates, and Steven Swanson. Hippogriff: Efficiently Moving Data in Heterogeneous Computing Systems. In *34th IEEE International Conference on Computer Design*, ICCD 2016, 2016.
24. Yang Liu, Hung-Wei Tseng, and Steven Swanson. Scale Up MapReduce with I/O-Oriented Scheduling for the GPU. In *34th IEEE International Conference on Computer Design*, ICCD 2016, 2016.
25. Hung-Wei Tseng and Dean M. Tullsen. CDTT: Compiler-generated data-triggered threads. In *20th International Symposium on High Performance Computer Architecture*, HPCA 2014, pages 650–661, 2014.
26. Leo Porter, Saturnino Garcia, Hung-Wei Tseng, and Daniel Zingaro. Evaluating student understanding of core concepts in computer architecture. In *18th ACM Conference on Innovation and Technology in Computer Science Education*, ITiCSE 2013, pages 279–284, 2013.
27. Hung-Wei Tseng, Laura M. Grupp, and Steven Swanson. Underpowering NAND flash: Profits and perils. In *48th Design Automation Conference*, DAC 2013, pages 1–6, 2013.
28. Hung-Wei Tseng and Dean M. Tullsen. Software data-triggered threads. In *ACM SIGPLAN 2012 Conference on Object-Oriented Programming, Systems, Languages and Applications*, OOPSLA 2012, pages 703–716, 2012.
29. Hung-Wei Tseng and Dean M. Tullsen. Eliminating redundant computation and exposing parallelism through data-triggered threads. *IEEE Micro, Special Issue on the Top Picks from Computer Architecture Conferences*, volume 32:38–47, 2012.
30. Hung-Wei Tseng, Laura M. Grupp, and Steven Swanson. Understanding the impact of power loss on flash memory. In *48th Design Automation Conference*, DAC 2011, pages 35–40, 2011.
31. Hung-Wei Tseng and Dean M. Tullsen. Data-triggered threads: Eliminating redundant computation. In *17th International Symposium on High Performance Computer Architecture*, HPCA 2011, pages 181–192, 2011.
32. Han-Lin Li, Chia-Lin Yang, and Hung-Wei Tseng. Energy-aware flash memory management in virtual memory system. *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)*, volume 16(8):952–964, 2008.
33. Hung-Wei Tseng, Han-Lin Li, and Chia-Lin Yang. An energy-efficient virtual memory system with flash memory as the secondary storage. In *2006 International Symposium on Low Power Electronics and Design*, ISLPED 2006, pages 418–423, 2006.
34. Shih-Hsien Yang, Hung-Wei Tseng, Eric Hsiao-Kuang Wu, and Gen-Huey Chen. Utilization based duty cycle tuning mac protocol for wireless sensor networks. In *IEEE Global Telecommunications Conference, 2005, GLOBECOM 2005*, pages 3258–3262, 2005.
35. Chia-Lin Yang, Hung-Wei Tseng, Chia-Chiang Ho, and Ja-Ling Wu. Software-controlled cache architecture for energy efficiency. *IEEE Transactions on Circuits and Systems for Video Technology*, volume 15(5):634–644, 2005.
36. Chia-Lin Yang, Alvin R. Lebeck, Hung-Wei Tseng, and Chien-Hao Lee. Tolerating memory latency through push prefetching for pointer-intensive applications. *Transactions on Architecture and Code Optimization (TACO)*, volume 1(4):445–475, 2004.

37. Hung-Wei Tseng, Shih-Hsien Yang, Po-Yu Chuang, Eric Hsiao-Kuang Wu, and Gen-Huey Chen. An energy consumption analytic model for a wireless sensor MAC protocol. In *IEEE 60th Vehicular Technology Conference, VTC2004-Fall*, pages 4533–4537, 2004.
38. Chia-Lin Yang, Hung-Wei Tseng, and Chia-Chiang Ho. Smart cache: An energy-efficient D-cache for a software MPEG-2 video decoder. In *2003 Joint Conference of the Fourth International Conference on Information, Communications and Signal Processing, 2003 and Fourth Pacific Rim Conference on Multimedia, ICICS-PCM 2003*, pages 1660–1664, 2003.

Workshops

1. Hung-Wei Tseng and Dean M. Tullsen. Data-triggered multithreading for near data processing. In *1st Workshop on Near-Data Processing, WoNDP 2013*, 2013.

Tech Reports

1. Hung-Wei Tseng, Yang Liu, Mark Gahagan, Jing Li, Yanqin Jin, and Steven Swanson. Gullfoss: Accelerating and simplifying data movement among heterogeneous computing and storage resources. Technical Report CS2015-1015, Department of Computer Science and Engineering, University of California, San Diego technical report, 2015.

Patents

1. Murali Annamaram, Gunjae Koo, Kiran Kumar Matam, and Hung-Wei Tseng. Dynamic near-data processing control mechanism based on computer resource availability on solid-state disk platforms, October 1 2020. US Patent App. 16/650,758.

Grant and Funding

CNS Core: Small: Simultaneous and Heterogenous Multithreading

National Science Foundation

\$593,000.00. October, 2023 – June, 2026

Extending AMX Beyond AI

Intel Corp.

\$120,000.00. October, 22 –

CNS Core: Small: Re-engineering Applications for Tensor Processing Units

National Science Foundation

\$495,000.00. July, 2020 – June, 2024

CSR: Small: IOQL: an I/O Interface for Near-Data Processing

National Science Foundation

\$499,515.00. August, 2018 – August, 2022

A Drop-in Upgradable Near-Data Processing Architecture for Machine Learning Applications in Heterogeneous Computers

Facebook Research

\$50,000.00. March, 2017

CRII: CSR: Rethinking the FTL in SSDs – a file translation layer instead of a flash translation layer

National Science Foundation

\$174,998.00. March, 2017 – March, 2020

Awards and Competitions

ACM/IEEE MICRO Best Paper Nomination

2021

Outstanding Paper Award, the 27th IEEE Real-Time and Embedded Technology and Applications Symposium

2021

IEEE Micro’s top picks from computer architecture conferences, “Varifocal Storage: Dynamic Multi-Resolution Data Storage”

2020

ACM/IEEE MICRO Best Paper Honorable Mention

2019

Facebook Research Award

2018

IEEE Micro's top picks from computer architecture conferences, "Eliminating Redundant Computation and Exposing Parallelism through Data-Triggered Threads"	2012
HPCA best student paper nominee	2011
National Taiwan University Student Service Education Award	2003
ACM International Collegiate Programming Contest (ICPC) Asia Regional, Taipei Site. 4th Place	2000
Presidential Award of National Taiwan University	1999, 2000, 2003

Teaching Experience

Assistant/Associate Professor, Department of Electrical and Computer Engineering/Department of Computer Science and Engineering, University of California, Riverside

- CS 203: Advanced Computer Architecture 2023 Fall, 2023 Spring, 2022 Fall, 2021 Fall, 2020 Fall, 2019 Fall
- EE/CS 277: Data-Centric Computer Architecture 2024 Winter, 2022 Spring
- CS 202: Advanced Operating Systems 2022 Winter, 2021 Winter, 2020 Winter
- EE 260: Quantum Computer Architecture 2021 Spring
- EE/CS 120A: Logic Design 2020 Spring
- EE 260: Trends in Computer System Design 2020 Winter

Assistant Professor, Department of Computer Science, North Carolina State University

- CSC 456: Computer Architecture and Multiprocessors 2019 Spring, 2018 Spring
- CSC 501: Operating Systems Principles 2018 Fall, 2017 Fall, 2016 Fall
- CSC 236: Computer Organization and Assembly Language for Computer Scientist 2017 Spring

Visiting Assistant/Associate Professor/Lecturer, Department of Computer Science and Engineering, University of California, San Diego

- CSE 142: Computer Architecture: Software Perspective 2023 Summer, 2022 Summer
- CSE 142L: Software Projects of Computer Architecture 2023 Summer, 2022 Summer
- CSE 141: Introduction to Computer Architecture 2020 Summer, 2019 Summer I/II, 2016 Summer, 2016 Spring, 2014 Summer, 2012 Summer
- CSE 141L: Project of Computer Architecture 2020 Summer, 2019 Summer I/II, 2016 Summer, 2014 Summer, 2012 Summer

Teaching Assistant, Department of Computer Science and Engineering, University of California, San Diego

- CSE 141: Introduction to Computer Architecture 2009 Spring, 2009 Fall, 2010 Summer
- CSE 240A: Principle of Computer Architecture 2010 Winter
- CSE 8A: Introduction to Computer Science: Java 2009 Fall
- CSE 141L: Project of Computer Architecture 2009 Summer

Other Working Experience

Visiting Professor , Google, Sunnyvale	Jan. 2023–March. 2023
Researcher , Intel Labs, Santa Clara	July 2022–Sep. 2022
Research Assistant , University of California, San Diego	Sep. 2007–Sep. 2014
Intern Software Engineer , Intel Labs, Santa Clara	June 2013–Sep. 2013
Software Engineer , Academia Sinica	Jan. 2007–Aug. 2007
Intern Software Engineer , Streaming 21 Inc.	July 2000–Aug 2000

Conference and Workshop Organization

General Co-Chair , The 14th Non-Volatile Memories Workshop (NVMW)	2023
Artifact Evaluation Chair , The 29th IEEE Symposium on High Performance Computer Architecture (HPCA)	2023
Track Program Chair , IEEE International Conference on Computer Design (ICCD)	2022
Track Program Chair , The 2022 ACM/SIGAPP Symposium on Applied Computing (SAC)	2022
Steering Committee Member , Non-Volatile Memories Workshop	2022, 2021, 2020, 2019
Finance Chair , The 26th International Symposium on High-Performance Computer Architecture (HPCA)	2020
Local Arrangements Co-chair , 2018 IEEE International Symposium on Workload Characterization (IISWC)	2018
Registration Chair , The 45th International Symposium on Computer Architecture (ISCA)	2018
Web Chair , ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)	2018
Program Co-chair , Non-Volatile Memories Workshop	2018, 2017
Web Chair , The 23rd IEEE Symposium on High Performance Computer Architecture (HPCA)	2017
Web Chair , The 46th IEEE/ACM International Symposium on Microarchitecture (MICRO)	2016
Submission Chair , The 19th IEEE Symposium on High Performance Computer Architecture (HPCA)	2013

Program Committee Member

Program Committee Member , IEEE/ACM International Symposium on Computer Architecture (ISCA)	2023, 2022
Program Committee Member , IEEE International Symposium on High-Performance Computer Architecture (HPCA)	2023, 2021, 2020
Program Committee Member , IEEE/ACM International Symposium on Microarchitecture (MICRO)	2022
Program Committee Member , IEEE International Conference on Computer Design (ICCD)	2021–2017
Program Committee Member , IEEE International Symposium on Workload Characterization (IISWC)	2021, 2018
Program Committee Member , IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)	2017
Program Committee Member , Non-Volatile Memory Systems and Applications Symposium (NVMSA)	2017
Program Committee Member , Workshop on Parallel Programming Models - Special Edition on Task Parallelism	2014–2016
Program Committee Member , Non-Volatile Memories Workshop (NVMW)	2016

Reviewers and Editors

Associate Editor , ACM Transactions on Architecture and Code Optimization (TACO)	2022–
Associate Editor , IEEE Computer Architecture Letters (CAL)	2022–
Reviewer , IEEE Transactions on Very Large Scale Integration Systems	2012–
Reviewer , IEEE Computer Architecture Letters	2016–
Reviewer , IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems	2018–
Reviewer , IEEE Transactions on Computers	2018–
Reviewer , IEEE Transactions on Embedded Computing Systems	2021–

External Reviewer, IEEE/ACM International Symposium on Microarchitecture

2021, 2020, 2019

External Reviewer, International Symposium on Computer Architecture

2020, 2017