SASINDU KANGARA MUDIYANSELAGE CURRICULUM VITAE

RESEARCH INTERESTS

- Parallel and Distributed Computing
- Hardware-Software Co-design
- Computer Architecture
- Tensor Algebra and Tensor Decomposition
- Performance Modeling

EDUCATION

Doctor of Philosophy in **Computer Engineering** *Summer, 2025*
University of Southern California, Los Angeles, CA **CGPA: 3.92/4.00**
Thesis: Hardware-aware Algorithm Design of Tensor Decomposition for CPU, GPU, and FPGA
Advisor: Viktor Prasanna Focus: High Performance Computing

Master of Science in **Electrical Engineering** *Spring, 2024*
University of Southern California, Los Angeles, CA **CGPA: 3.96/4.00**
Focus (Courses): Parallel Computing and Computer Architecture

B.Sc. Eng (Hons.) in **Electronic & Telecommunication Engineering** *Fall, 2018*
Univeristy of Moratuwa, Sri Lanka **CGPA: 3.75/4.20**

WORK EXPERIENCE

Senior Engineer, Architecture
Sep, 2025 - present
Tenstorrent Inc, Santa Clara, CA

Research Assistant, Ming Hsieh Department of Electrical and Computer Engineering
Aug, 2019 - Aug, 2025
University of Southern California, Los Angeles, CA

Teaching Assistant, Ming Hsieh Department of Electrical and Computer Engineering
Jan, 2020 - May 2023
University of Southern California, Los Angeles, CA

Teaching Assistant, Thomas Lord Department of Computer Science
Aug, 2019 - May 2020
University of Southern California, Los Angeles, CA

Research Intern, DEVCOM Army Research Office
May, 2022 - August, 2022
DEVCOM US Army Research Lab, Marina del Rey, CA

Engineer, Wave Computing
Feb, 2018 - July, 2019
Wave Computing, Santa Clara, CA

Visiting Intern, Research Institute for Nanodevice and Bio Systems
July, 2017 - Aug, 2017
Research Institute for Nanodevice and Bio Systems, Hiroshima University, Japan

EXPERIENCE

Performance Modeling and Workload Analysis on RISC-V CPUs *2025 Sep - Now*

Tenstorrent USA Inc.

- Benchmark development for emerging Agentic LLM and server platform workloads.
- Conduct competitive analysis of products, utilizing industry-standard benchmarks and internally deveped benchmarks to identify performance gaps
- Collaborate with architects and RTL designers to enhance micro-architectural features, focusing on performance and power efficiency
- Employ performance models and EDA frameworks to characterize workloads and predict performance under various real-world scenoarios

Accelerating Tensor Decomposition on CPU, GPU and FPGA *2021 June - Jan 2026*

USC, Supervised by Prof. Viktor Prasanna

- Develop parallel algorithms to reduce the total execution time of sparse Matricized Tensor Times Khatri-Rao Product
- Introduce novel tensor format to support the proposed parallel algorithms
- Develop custom designs on CPU, GPU, and FPGA

Hardware-Algorithm Co-design on Scalable Photonic SRAM-based In-Memory Computing Tensor Core (SPRINT) *2023 Dec - 2025 August*

USC, Supervised by Prof. Viktor Prasanna

- Hardware-Algorithm co-design to map the computations of Matricized Tensor Times Khatri-Rao Product to SPRINT hardware
- Develop performance models of SPRINT hardware to estimate the peak and sustained performance of the target hardware
- Work with the SPRINT hardware development team and advise them on bottlenecks of SPRINT hardware design based on the performance models

Automatic Target Recognition on Synthetic Aperture Radar Images using Graph Neural Networks (GNN for SAR-ATR) *2022 March - 2023 March*

USC, Supervised by Prof. Viktor Prasanna

- Develop Explainable AI models for GNNs to understand the important features in SAR images
- Introduce novel Graph Neural Network models for Automatic Target Recognition
- Human-in-the-loop to improve the accuracy of ML models

Shared Memory Controller for Heterogeneous Hardware Platforms *2020 Jan - 2021 July*

USC, Supervised by Prof. Viktor Prasanna

- Develop a Memory Controller that can handle incoming data traffic with irregular access requests from multiple types of hardware (e.g., CPU, GPU, and FPGA) using data routing/mapping techniques

Reconfigurable Co-Processor Architecture with Limited Numerical Precision to Accelerate Deep Neural Networks *2016 Dec - 2018 Jan*

University of Moratuwa, Supervised by Dr. Ajith Pasqual

- Design a scalable limited precision model independent reconfigurable co-processor

Software Defined Networking (SDN) Switch for Core-Networks on FPGA *2017 Jan - 2018 Feb*

University of Moratuwa, Supervised by Dr. Ajith Pasqual

- Implement a high-performance, scalable software-defined networking (SDN) switch fabric for the core-network environment. The implemented FPGA-based switch that is compatible with OpenFlow; the protocol for the Southbound Interface of SDN

HEVC Video Encoder - SCC Extension on FPGA *2016 July - 2017 Feb*

ParaQum Technologies, Supervised by Dr. Ajith Pasqual

- Introduce resource-efficient real-time hardware design for the major SCC tools including Intra Block Copy, and Palette Coding

PUBLICATIONS

Thesis Work:

- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2025). HOPE: Hardware-Agnostic Parallel Algorithm to Accelerate MTTKRP for Sparse Tensor Decomposition on CPU and GPU. In IEEE Transactions on Parallel and Distributed Systems. [Submitted]
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2025). Accelerate MTTKRP for Tensor Decomposition on Sparse Tensors using GPUs. In ACM Transactions on Parallel Computing. [Submitted]
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2025). Accelerating MTTKRP for Sparse Tensor Decomposition on Multi-GPU Platform. In Proceedings of the 54st International Conference on Parallel Processing (ICPP '25).
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2024). Accelerating Sparse MTTKRP for Small Tensor Decomposition on GPU. In Proceedings of the 2024 Asilomar Conference on Signals, Systems, and Computers. **(Best Paper Candidate: 10/500)**
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2024). Sparse MTTKRP Acceleration for Tensor Decomposition on GPU. In Proceedings of the 21st ACM International Conference on Computing Frontiers.
- **Wijeratne, S.**, Wang, T. Y., Kannan, R., Prasanna, V. (2023). Accelerating Sparse MTTKRP for Tensor Decomposition on FPGA. In Proceedings of the 2023 ACM/SIGDA International Symposium on Field Programmable Gate Arrays.
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2023). Dynasor: A Dynamic Memory Layout for Accelerating Sparse MTTKRP for Tensor Decomposition on Multi-core CPU. In 35th IEEE International Symposium on Computer Architecture and High Performance Computing. **(Best Paper Candidate: 6/23)**
- **Wijeratne, S.**, Wang, T. Y., Kannan, R., Viktor, P. (2022). Towards Programmable Memory Controller for Tensor Decomposition. In 11th International Conference on Data Science, Technology and Applications (DATA'22).

Other:

- **Wijeratne, S.**, Sundar, S., Kaiser, Md., Jaiswal, A., Mathew, C., Jacob, A. P., Prasanna, V. (2024). Predictive Performance of Photonic SRAM-based In-Memory Computing for Tensor Decomposition. In 2024 IEEE High Performance Extreme Computing Conference (HPEC). **(Best Student Paper: 1/100)**
- **Wijeratne, S.**, Jaiswal, A., Jacob, A. P., Zhang, B., Prasanna, V. (2022). Performance modeling sparse MTTKRP using optical static random access memory on FPGA. In 2022 IEEE High Performance Extreme Computing Conference (HPEC).
- **Wijeratne, S.**, Zhang, B., Kannan, R., Prasanna, V., Busart, C. (2023). PAHD: Perception-Action based Human Decision Making using Explainable Graph Neural Networks on SAR images. In Automatic Target Recognition XXXIII.
- **Wijeratne, S.**, Pattnaik, S., Chen, Z., Kannan, R., Prasanna, V. (2021). Programmable fpga-based memory controller. In 2021 IEEE Symposium on High-Performance Interconnects (HOTI).
- **Wijeratne, S.**, Kannan, R., Prasanna, V. (2021). Reconfigurable Low-latency Memory System for Sparse Matricized Tensor Times Khatri-Rao Product on FPGA. In 2021 IEEE High Performance Extreme Computing Conference (HPEC).
- **Wijeratne, S.**, Ekanayake, A., Jayaweera, S., Ravishan, D., Pasqual, A. (2019). Scalable High Performance SDN Switch Architecture on FPGA for Core Networks. In Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays.
- **Wijeratne, S.**, Jayaweera, S., Dananjaya, M., Pasqual, A. (2018). Reconfigurable co-processor architecture with limited numerical precision to accelerate deep convolutional neural networks. In 2018 IEEE 29Th international conference on application-specific systems, architectures and processors (ASAP).
- Sunder, S., Kaiser A., **Wijeratne, S.**, Mathew, C., Prasanna, V., Jaiswal, A., Jacob, A. (2024). Scalable in-memory compute optical processor. In Smart Photonic and Optoelectronic Integrated Circuits 2025.
- Chen, P., Manjunath, P., **Wijeratne, S.**, Zhang, B., Prasanna, V. (2023). Exploiting On-chip Heterogeneity of Versal Architecture for GNN Inference Acceleration. In 33rd International Conference on Field-Programmable Logic and Applications.

- Zhang, B., **Wijeratne, S.**, Kannan, R., Prasanna, V., Busart, C. (2023). Graph neural network based SAR automatic target recognition with human-in-the-loop. In Algorithms for Synthetic Aperture Radar Imagery XXX.
- Zhang, B., **Wijeratne, S.**, Kannan, R., Prasanna, V., Busart, C. (2023). Graph Neural Network for Accurate and Lowcomplexity SAR ATR. In 15th International Conference on Advanced Geographic Information Systems, Applications, and Services.
- Zhang, B., **Wijeratne, S.**, Kannan, R., Prasanna, V., Busart, C. (2023). How can Human-in-the-loop Improve the Performance of SAR ATR? A Reinforcement Learning Based Approach. In International Radar Conference 2023. **(Best Paper Candidate: 5/200)**
- Zhang, B., Jaiswal, A., Mathew, C., Lakkireddy, R., Jacob, A., **Wijeratne, S.**, Prasanna, V. (2022). A high throughput parallel hash table on fpga using xor-based memory. In 2020 IEEE High performance extreme computing conference (HPEC).
- Ye, T., Kuppannagari, S. R., De Rose, C. A., **Wijeratne, S.**, Kannan, R., Prasanna, V. (2022). Estimating the Impact of Communication Schemes for Distributed Graph Processing. In 2022 21st International Symposium on Parallel and Distributed Computing (ISPDC).
- Zhang, R., **Wijeratne, S.**, Yang, Y., Kuppannagari, S. R., Prasanna, V. (2020). A high throughput parallel hash table on fpga using xor-based memory. In 2020 IEEE High performance extreme computing conference (HPEC).
- Senanayake, R., Liyanage, N., **Wijeratne, S.**, Atapattu, S., Ekanayake, A., Pasqual, A. (2017). High performance hardware architectures for Intra Block Copy and Palette Coding for HEVC screen content coding extension. In 2018 IEEE international conference on application-specific systems, architectures and processors (ASAP).

TEACHING EXPERIENCE

- Teaching Assistant for Introduction to Digital Circuits (EE 354L), Ming Hsieh Department of Electrical and Computer Engineering, USC (Instructor: Gandhi Puvvada)
  6 semesters (Fall 2020 - Spring 2023)
- Teaching Assistant for Accelerated Computing using FPGAs (EE 599), Ming Hsieh Department of Electrical and Computer Engineering, USC (Instructor: Viktor Prasanna)
  1 semester (Spring 2020)
- Teaching Assistant for Introduction to Computer Systems (CSCI 356), Thomas Lord Department of Computer Science, USC (Instructor: Marco Paolieri)
  2 semesters (Fall 2019 - Spring 2020)

MENTORING EXPERIENCE

Undergraduate and Graduate Student Mentor, FPGA/Parallel Computing Lab
Recruited, mentored, and guided multiple undergraduate and graduate students while working on research at FPGA/Parallel Computing Lab, Ming Hsieh Department of Electrical and Computer Engineering, USC: Ruizhi Zhang, Sanket Pattnaik, Zhiyu Chen, Paul Chen, Pavan Manjunath, Priyadarshan Murugan, Ruiqi Cao, and Chao Gao. Led to 3 conference publications.

Mentor - Viterbi Mentorship Program
Actively participated in USC Viterbi Mentorship Program for 2 semesters.

Intern and New Graduate Mentor at Wave Computing (Sri Lankan Branch)
Recruited, mentored, and guided multiple interns and new college graduates over the 1 year I was at Wave Computing.

AWARDS AND ACHIEVEMENTS

- Best Student Paper: 2024 IEEE High Performance Extreme Computing Conference (HPEC)
- ACM Travel Grant Recipient of ISFPGA 2023 Conference

- USC Ming Hsieh Department of Electrical Engineering, Outstanding Poster Award- Taming Sparse Tensor Decomposition: Unified Acceleration Framework, ECE 14th Annual Research Festival, Fall 2024
- Best Paper Nominee: (1) Asilomar Conference on Signals, Systems, and Computers 2024, (2) SBAC-PAD 2023, (3) International Radar Conference 2023
- Ming Hsieh Department of Electrical and Computer Engineering PhD Scholar Finalist 2023-2024
- Mahapola merit scholarship to pursue undergraduate studies by Department of Examinations, Sri Lanka
- All Sri Lanka Rank 81 (out of ∼80,000) in Undergraduate, University Entrance Examination
- All Sri Lanka Rank 16 (out of ∼400,000) in General Ordinary Level Examination

ACADEMIC AND VOLUNTEERING SERVICE

- Local arrangement co-chair of the 31st IEEE International Symposium On Field-Programmable Custom Computing Machines (FCCM 2023)
- Web chair of the 29th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2022)
- Publicity co-chair of the 29th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2022)
- Publicity chair of the 28th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2021)
- Artifact evaluator of the 31st IEEE International Symposium On Field-Programmable Custom Computing Machines (FCCM 2023)
- Journal Reviewer for IEEE Transactions on Parallel and Distributed Systems (TPDS), Microprocessors and microsystems, Pattern Recognition, and ACM Transactions on Architecture and Code Optimization
- Program Committee member of 34th IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM 2026) and Computing Frontiers 2026 (CF '26)
- Committee member of SPEC OSG Cloud (contributor to SPEC Server Platform 2026 Benchmark)