# Charles Block

coblock2@illinois.edu

## Education

2022-Present **Ph.D. Candidate, Computer Science**, *University of Illinois*, Urbana-Champaign
NSF Graduate Research Fellow, advised by Dr. Josep Torrellas. Research topics include computer architecture and distributed systems.

2022-2024 **M.S. Computer Science**, *University of Illinois*, Urbana-Champaign

2018-2022 **B.S. Electrical & Computer Engineering with Honors**, *University of Texas*, Austin

## Industry Experience

Summer 2025 **NVIDIA Corp.**, *Deep Learning Training Performance Intern*
- Characterized performance bottlenecks in large-scale distributed LLM training.
- Identified and resolved inefficiencies in communication/GEMM overlap in TransformerEngine.
- Investigated relationships between power and performance on the latest generation of GPUs.

Summer 2024 **Apple Inc**, *CPU Architecture/Performance Intern*
- Investigated novel memory ordering techniques and their impact on next-generation CPUs.
- Developed enhancements to trace-based performance simulators and tooling.
- Identified and addressed performance simulation bugs to accurately model RTL.

Summer 2022 **Intel Corporation**, *Design Validation Intern*
- Supported CPU validation in the Advanced Architecture Development Group.
- Debugged core RTL & software infrastructure to address performance and power issues.
- Improved run-time performance of long- and often-running validation jobs by 20x.

Summer 2021 **Microsoft Corporation**, *Software Development Intern (Azure)*
- Initiated development of a distributed service to supplement PostgreSQL databases.
- Developed a new distributed, in-memory storage backend for PostgreSQL.
- Developed software using the C and Rust programming languages.

Summer of 2019 & 2020 **Amazon Robotics**, *Firmware Development Intern*
- Supported firmware development for the Amazon Scout robotics project.
- Integrated various sensors, MCUs, and SoCs using both embedded Linux and bare-metal C.
- Developed/evaluated a real-time object detection and tracking system on an embedded SoC.

Summer 2018 **TyRex Group, LTD**, *Engineering Intern*

## Academic & Leadership Experience

2023-Present **i-acoma Group at UIUC**, *Graduate Researcher*
- Working with Professor Josep Torrellas on HPC architectures and distributed systems.
- Supercomputer-scale distributed algorithms for sparse matrix kernels.
- Multi-agent reinforcement learning for multi-core architectures.

2021-2022 **Lu Research Group at UT Austin**, *Research Assistant*
- Worked with Professor Nanshu Lu to develop wearable low-power biomedical sensors.
- Developed measurement circuitry for experiments with novel pressure sensor technology.

2018-2022 **Longhorn Racing - Solar Vehicle Team**
- Oversaw manufacturing of composite materials, electronic systems, and structural components.
- Led development of embedded hardware and software for vehicle control and power.
- Developed lessons & mentored other students for designing circuits and software.

Spring of 2020 & 2021 **Intro to Embedded Systems at UT Austin (EE319K)**, *Teaching Assistant*
- Assisted in teaching and lab supervision for a freshman embedded systems course.

## Awards

2022 **NSF Graduate Research Fellowship**, *National Science Foundation*

2022 **Wing Kai Cheng Fellowship**, *University of Illinois Urbana-Champaign*

2020 & 2021 **Dr. Ariane L. Beck and Mr. Eric Sebesta Endowed Scholarship**, *UT Austin ECE*

## Publications

### Conference & Journal Publications

[1] **C. Block**, G. Gerogiannis, and J. Torrellas, "Micro-mama: Multi-agent reinforcement learning for multicore prefetching," in *58th IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '25. ACM, Oct. 2025.

[2] G. Gerogiannis, D. Merkouriadis, **C. Block**, A. Zulfiqar, F. Tofalos, M. Shahbaz, and J. Torrellas, "Netsparse: In-network acceleration of distributed sparse kernels," in *58th IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '25. ACM, Oct. 2025.

[3] C. Sudusinghe, G. Gerogiannis, D. Lenadora, **C. Block**, J. Torrellas, and C. Mendis, "COGNATE: Acceleration of sparse tensor programs on emerging hardware using transfer learning," in *Forty-second International Conference on Machine Learning (ICML '25)*, 2025.

[4] L. Scalco de Vasconcelos, Y. Yan, P. Maharjan, S. Kumar, M. Zhang, B. Yao, H. Li, S. Duan, E. Li, E. Williams, S. Tiku, P. Vidal, R. S. Solorzano-Vargas, W. Hong, Y. Du, Z. Liu, F. Iwane, **C. Block**, A. T. Repetski, P. Tan, P. Wang, M. G. Martín, J. d. R. Millán, X. He, and N. Lu, "On-scalp printing of personalized electroencephalography e-tattoos," *Cell Biomaterials*, 2024.

[5] I. Ranawaka, M. T. Hussain, **C. Block**, G. Gerogiannis, J. Torrellas, and A. Azad, "Distributed-memory parallel algorithms for sparse matrix and sparse tall-and-skinny matrix multiplication," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '24)*. IEEE Press, 2024.

[6] V. Suresh, B. Mishra, Y. Jing, Z. Zhu, N. Jin, **C. Block**, P. Mantovani, D. Giri, J. Zuckerman, L. P. Carloni, and S. V. Adve, "Mozart: Taming taxes and composing accelerators with shared-memory," in *Proceedings of the 2024 International Conference on Parallel Architectures and Compilation Techniques (PACT '24)*, 2024, p. 183–200.

[7] K.-H. Ha, Z. Li, S. Kim, H. Huh, Z. Wang, H. Shi, **C. Block**, S. Bhattacharya, and N. Lu, "Stretchable hybrid response pressure sensors," *Matter*, vol. 7, no. 5, 2024.

[8] **C. Block**, G. Gerogiannis, C. Mendis, A. Azad, and J. Torrellas, "Two-face: Combining collective and one-sided communication for efficient distributed spmm," in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24), Volume 2*, 2024.

## Workshop Publications

[9] **C. Block**, P. Palacios Almendros, A. Farrell, G. Gerogiannis, and J. Torrellas, "Performance-driven composite prefetching with bandits," in *Fourth Data Prefetching Champsionship at HPCA 2026*, ser. DPC4, 2 2026.

[10] C. Sudusinghe, G. Gerogiannis, D. Lenadora, **C. Block**, J. Torrellas, and C. Mendis, "Automated data selection for efficient cost model training to optimize sparse matrix kernels on emerging hardware accelerators," in *The Exploration in AI Today Workshop at ICML 2025*, 2025.