

Bilge ACUN

✉ acun2@illinois.edu

🌐 <http://bilgeacun.com>

🔍 [Google Scholar Profile](#)

I work on building efficient, high performance and sustainable ML systems and datacenters at large. My research spans across systems, computer architecture and machine learning.

Affiliation

Jan 2019 | Research Scientist, META AI, USA
- Present | FAIR SysML team, supervisor: [Carole-Jean Wu](#)

Education

2012-2017 | Ph.D., UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, USA
Department of Computer Science
Dissertation: “Mitigating Variability in HPC Systems and Applications for Performance and Power Efficiency”. Advisor: [Prof. Laxmikant V. Kalé](#)

2008-2012 | B.S., BILKENT UNIVERSITY, TURKEY
Department of Computer Science
Salutatorian, second highest GPA among graduating class

2011 | Erasmus Exchange Program, UPPSALA UNIVERSITY, SWEDEN
Department of Computer Science

Honors and Awards

2021 [MLSys Conference Outstanding Paper Award](#)
2018 [ACM SigHPC Dissertation Award Honorable Mention](#)
2018 Selected as a [Heidelberg Laureate](#)
Heidelberg Laureate Forum brings together Turing Laureates with early career researchers.
2017 Selected as a [Rising Star in EECS](#), Stanford University
2017 [Kenichi Miura Award](#), UIUC
2017 [Illinois Innovation Prize Finalist](#), Technology Entrepreneur Center, UIUC
2012-13 [Saburo Muroga Endowed Fellowship](#), UIUC
2008-12 Bachelor of Science Full Tutition-Waiver with Stipend Merit Scholarship, Bilkent University

Grants and Scholarships

2017 [Google Grace Hopper Celebration Travel Grant](#), Google
2016 [AWARE \(Accelerating Women And underRepresented Entrepreneurs\)](#) Grant Winner, UIUC Research Park
2013 [Grace Hopper Celebration Travel Grant](#), Department of Computer Science, UIUC
2011 Erasmus Exchange Program Scholarship, Uppsala University, Sweden
2008-12 Bachelor of Science Full Tutition-Waiver with Stipend Merit Scholarship, Bilkent University

Professional Experience

2017 - 2019 | Research Staff Member at IBM T.J. WATSON RESEARCH CENTER, NY, USA
[Data Centric Systems Group](#)
- Distributed application performance optimization on world’s most powerful supercomputer: [Summit](#)
- Performance modeling and prediction for next generation cloud and machine learning platforms

AUGUST 2012 - AUGUST 2017 | Research Assistant at UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, IL, USA
[Parallel Programming Laboratory](#)
- Contributed to an open source C++ based parallel programming framework, [Charm++](#)
- Made Charm++ energy and power efficient, thermal and variation aware
- Optimized the performance of the award-winning petascale biomolecular simulation application [NAMD](#)
- Worked on network optimizations, load balancing, malleability for large scale supercomputers

JUNE 2016 - DECEMBER 2016	Research Intern at IBM T.J. WATSON RESEARCH CENTER, NY, USA Data Centric Systems Group - Implemented a neural network-based learning model for predicting core temperatures - Using the model, implemented runtime based proactive cooling methods, task scheduling techniques
JUNE 2014 - AUGUST 2014	Research Intern at LAWRENCE LIVERMORE NATIONAL LABORATORY, CA, USA Computation Division - Developed a parallel network simulator for performance prediction of production HPC applications - Simulated different network topologies to obtain better communication performance at scale
AUGUST 2013 - DECEMBER 2013	Teaching Assistant at UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, IL, USA <i>CS 598 - Parallel Programming with Migratable Objects</i> - Helped teach a graduate level special topics course on Charm++ , an adaptive and parallel runtime system which can be used to program multicore desktops, clusters, as well as petascale supercomputers - Prepared and graded class assignments, exams and projects

Undergraduate Work Experience

SUMMER 2011	Software Engineering Intern, MICROSOFT, Turkey
SUMMER 2011	Software Engineering Intern, TUBITAK UEKAE, Turkey (National Research Council of Turkey)
SUMMER 2010	Software Engineering Intern, MIDDLE EAST TECHNICAL UNIVERSITY, MODSIMMER, Turkey

Publications

★ 1500+ citations ★ 1 book chapter ★ 4 journal papers ★ 20 conference papers ★ 5 workshop papers

Book Chapter:

- [**CRC Press'17**] **B. Acun**, R. Buch, , L.V. Kalé, J. C. Phillips. “NAMD: Scalable Molecular Dynamics Based on the Charm++ Parallel Runtime System” chapter in “Exascale Scientific Applications: Scalability and Performance Portability”. 2017.

Journals:

- [**IEEE Micro'24**] C.-J. Wu, **B. Acun**, R. Raghavendra, K. Hazelwood. “Beyond Efficiency: Scaling AI Sustainably”. 2024.
- [**IEEE Micro'21**] L. Wesolowski, **B. Acun**, V. Andrei, A. Aziz, G. Dankel, C. Gregg, X. Meng, C. Meurillon, D. Sheahan, L. Tian, J. Yang, P. Yu, K. Hazelwood. “Datacenter-Scale Analysis and Optimization of GPU Machine Learning Workloads”. 2021.
- [**IBM Journal of Research and Development'18**] **B. Acun**, D.J. Hardy, L.V. Kalé, K. Li, J.C. Phillips, J.E. Stone. “Scalable molecular dynamics with NAMD on the Summit system”. 2018.
- [**COMPUTER'16**] **B. Acun**, A. Langer, H. Menon, O. Sarood, E. Totonni, and L. V. Kalé. “Power, Reliability, Performance: One System to Rule Them All”. IEEE Computer, Energy Efficient Computing Special Issue. 2016.

Conferences:

- [**ICML'24**] S. Agarwal, **B. Acun**, B. Homer, M. Elhoushi, Y. Lee, S. Venkataraman, D. Papailiopoulos, C.-J. Wu. “CHAI: Clustered Head Attention for Efficient LLM Inference”. *International Conference on Machine Learning*. 2024.
- [**ACL'24**] M. Elhoushi, A. Shrivastava, D. Liskovich, B. Hosmer, B. Wasti, L. Lai, A. Mahmoud, **B. Acun**, S. Agarwal, A. Roman, A. Aly, B. Chen, C.-J. Wu. “Layer Skip: Enabling Early Exit Inference and Self-Speculative Decoding”. *Annual Meeting of the Association for Computational Linguistics*. 2024.
- [**ISPASS'24**] A. Golden, S. Hsia, F. Sun, **B. Acun**, B. Hosmer, Y. Lee, Z. DeVito, J. Johnson, G.-Y. Wei, D. Brooks, C.-J. Wu. “Generative AI Beyond LLMs: System Implications of Multi-Modal Generation”. *IEEE International Symposium on Performance Analysis of Systems and Software..* 2024.
- [**ISCA'24**] S. Hsia, A. Golden, **B. Acun**, N. Ardalani, Z. DeVito, G.Y. Wei, D. Brooks, C.-J. Wu. “MAD-Max Beyond Single-Node: Enabling Large Machine Learning Model Acceleration on Distributed Systems”. *International Symposium on Computer Architecture*. 2024.
- [**NeurIPS'23**] M. Mazumder, C. Banbury, X. Yao, B. Karlaš, W. G. Rojas, S. Diamos, G. Diamos, L. He, D. Kiela, D. Jurado, D. Kanter, R. Mosquera, J. Ciro, L. Aroyo, **B. Acun**, S. Eyuboglu, A. Ghorbani, E. Goodman,

T. Kane, C. R. Kirkpatrick, T.-S. Kuo, J. Mueller, T. Thrush, J. Vanschoren, M. Warren, A. Williams, S. Yeung, N. Ardalani, P. Paritosh, C. Zhang, J. Zou, C.-J. Wu, C. Coleman, A. Ng, P. Mattson, V.J. Reddi.

“DataPerf: Benchmarks for Data-Centric AI Development”. *Conference on Neural Information Processing Systems*. 2023.

6. [ASPLOS’23] [IEEE Micro Top Picks Honorable Mention] **B. Acun**, B. Lee, F. Kazhmiaka, K. Maeng, M. Chakkaravarthy, U. Gupta, D. Brooks, C.-J. Wu. “Carbon Explorer: A Holistic Approach for Designing Carbon Aware Datacenters”. *ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 2023.
7. [ASPLOS’23] [IEEE Micro Top Picks Honorable Mention] S. Hsia, U. Gupta, **B. Acun**, N. Ardalani, P. Zhong, G.-Y. Wei, D. Brooks, C.-J. Wu. “MP-Rec: Hardware-Software Co-design to Enable Multi-path Recommendation”. *ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 2023.
8. [MLSys’22] C.-J. Wu, R. Raghavendra, U. Gupta, **B. Acun**, N. Ardalani, K. Maeng, G. Chang, F. Aga, J. Huang, C. Bai, M. Gschwind, A. Gupta, M. Ott, A. Melnikov, S. Candido, D. Brooks, G. Chauhan, B. Lee, H.-H. Lee, B. Akyildiz, M. Balandat, J. Spisak, R. Jain, M. Rabbat, K. Hazelwood. “Sustainable AI: Environmental Implications, Challenges and Opportunities”. *Conference on Machine Learning and Systems*. 2022.
9. [ASPLOS’22] G. Sethi, **B. Acun**, N. Agarwal, C. Kozyrakis, C. Trippel, C.-J. Wu. “RecShard: Statistical Feature-Based Memory Optimization for Industry-Scale Neural Recommendation”. *ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 2022.
10. [HPCA’22] W. Xiong, L. Ke, D. Jankov, M. Kounavis, X. Wang, E. Northup, J. A. Yang, **B. Acun**, C.-J. Wu, P. T. P. Tang, G. E. Suh, X. Zhang, H.-G. S. Lee. “SecNDP: Secure Near-Data Processing with Untrusted Memory”. *IEEE International Symposium on High-Performance Computer Architecture*. 2022.
11. [MLSys’21][Outstanding Paper Award] C. Yin, **B. Acun**, Carole-Jean Wu, Xing Liu. “TT-Rec: Tensor Train Compression for Deep Learning Recommendation Models”. *Conference on Machine Learning and Systems*. 2021.
12. [HPCA’21] **B. Acun**, M. Murphy, X. Wang, J. Nie, C.-J. Wu, K. Hazelwood. “Understanding Training Efficiency of Deep Learning Recommendation Models at Scale”. *IEEE International Symposium on High-Performance Computer Architecture*. 2021.
13. [IGSC’19] **B. Acun**, K. Chandrasekar, and L. V. Kale. “Fine-Grained Energy Efficiency Using Per-Core DVFS with an Adaptive Runtime System”. *International Green and Sustainable Computing Conference*. 2019.
14. [HPCA’19] **B. Acun**, A. Buyuktosunoglu, E. K. Lee, Y. Park. “Power-Aware Heterogeneous Node Assembly”. *IEEE International Symposium on High-Performance Computer Architecture*. 2019.
15. [HiPC’17] **B. Acun**, E. K. Lee, Y. Park, L. V. Kalé. “Support for Power Efficient Proactive Cooling Mechanisms”. *International Conference on High Performance Computing*. 2017.
16. [ICS’16] **B. Acun**, P. Miller, L. V. Kalé. “Variation Among Processors Under Turbo Boost in HPC Systems”. *International Conference on Supercomputing*. 2016.
17. [HiPC’14] A. Gupta, **B. Acun**, O. Sarood, L. V. Kalé. “Towards Realizing the Potential of Malleable Jobs”. *International Conference on High Performance Computing*. 2014.
18. [SC’14] **B. Acun**, A. Gupta, N. Jain, A. Langer, H. Menon, E. Mikida, X. Ni, M. Robson, Y. Sun, E. Toton, L. Wesolowski, L. V. Kalé. “Parallel Programming with Migratable Objects: Charm++ in Practice”. *Supercomputing*. 2014.
19. [CLUSTER’13] H. Menon, **B. Acun**, SG De Gonzalo, O. Sarood, L. V. Kalé. “Thermal-Aware Automated Load Balancing for HPC Applications.” *IEEE International Conference on Cluster Computing*. 2013.
20. [ISCIS’13] **B. Acun**, A. Başpınar, E. Oğuz, M.İ. Saraç, F. Can. “Topic Tracking Using Chronological Term Ranking”. *International Symposium on Computer and Information Sciences*. 2013.

Workshops:

1. [CCAI, NeurIPS’23] **B. Acun**, B. Morgan, H. Richardson, N. Steinsultz, C.-J. Wu. “Unlocking the Potential of Renewable Energy Through Curtailment Prediction”. *Tackling Climate Change with Machine Learning*. 2023.

2. [HotCarbon'22] **B. Acun**, B. Lee, F. Kazhmiaka, A. Sundarrajan, K. Maeng, M. Chakkaravarthy, D. Brooks, C.-J. Wu. "Carbon Dependencies in Datacenter Design and Management". *HotCarbon: Workshop on Sustainable Computer Systems Design and Implementation*. 2022.
3. [E2SC, SC'16] **B. Acun**, E. K. Lee, Y. Park, L. V. Kalé. "Neural Network-Based Task Scheduling with Preemptive Fan Control". *International Workshop on Energy Efficient Supercomputing*. 2016.
4. [VarSys, IPDPS'16] **B. Acun**, L. V. Kalé. "Mitigating Processor Variation with Dynamic Load Balancing". *IEEE International Workshop on Variability in Parallel and Distributed Systems*. 2016.
5. [PADABS, EUROPAR'15] **B. Acun**, N. Jain, A. Bhatele, L. V. Kalé. "TraceR: A Parallel Trace Replay Tool for Studying Interconnection Networks". *Workshop on Parallel and Distributed Agent-Based Simulations*. 2015.

Patents

1. E. K. Lee, **B. Acun**, Y. Park, P. W. Coteus. "Learning-based Thermal Estimation in Multicore Architecture". *US Patent No: 11,334,398*. 2022.
2. **B. Acun**. "Device Identification via Chip Manufacturing Related Fingerprints". *US Patent No: 11,205,018*. 2021.
3. C.C. Yang, G. Cong, **B. Acun**, A. Morari. "Locality aware data loading for machine learning". *US Patent No: 11,093,862*. 2021.
4. **B. Acun**, E. K. Lee, Y. Park. "Job Scheduling Based on Node and Application Characteristics". *US Patent No: 10,725,834*. 2020.
5. **B. Acun**, E. K. Lee, Y. Park. "Power Efficiency Aware Node Component Assembly". *US Patent No: 10,831,252*. 2020.
6. E. K. Lee, **B. Acun**, Y. Park, A. Morari, A. Buyuktosunoglu. "Variation-Aware Intra-Node Power Shifting Among Different Hardware Components" *US Patent No: 10,761,583*. 2020.
7. **B. Acun**, E. Candan. "Systems and Methods for Computer Input". *US Patent Application No: 15/252,163*. (Pending patent). 2016.

Invited Talks and Other Presentations

- Invited Talk | [E2MC](#), Energy Efficient Machine Learning and Cognitive Computing Workshop, ASPLOS. 2024.
- Invited Talk | Stanford University, MLSys Seminar Series, "[Sustainable Datacenters for AI](#)". 2022.
- Guest Lecture | Carnegie Mellon University, "[ML with Large Datasets](#)" Course. 2022.
- Invited Talk | "A Holistic Approach for Datacenter Sustainability". *SIGEnergy WeCan*. 2022.
- Podcast | "May Earth Day", [Meta Tech Podcast](#). 2022.
- Invited Talk | "Large-Scale Deep Learning Recommendation Models at Facebook". *NVIDIA GTC*. 2021.
- Panelist | "MLPerf: A Benchmark for Machine Learning BoF Session" *Supercomputing Conference (SC)*. 2020.
- Invited Talk | "Characteristics of Facebook's Deep Learning Recommendation Training Models". *MLBench at ISPASS*. 2020.
- Panelist | "Introduction to HPC Research - HPC Experiences for Undergraduates" at *Supercomputing (SC)*. 2018.
- Invited Talk | "Energy Efficient Smart Programming Languages and Supercomputers." [CS Seminar](#) at *Bilkent University, Turkey*. 2018.
- PhD Forum | "Mitigating Variability in HPC Systems and Applications for Performance and Power Efficiency." *Doctoral Showcase Presentation, PhD Forum, Supercomputing Conference (SC)*. 2017.
- Invited Talk | "Energy Efficient Smart Programming Languages and Data Centers". *Hacettepe University, Turkey*. 2017.
- Invited Talk | "Neural Network-Based Power Optimizations in Runtime." *15th Annual Charm++ Workshop and Applications*. 2017.
- Demo | "Innovation Showcase: Illinois Innovation Prize Finalist". *Entrepreneurship Forum, UIUC*. 2017.

Poster	“An Assistive Computer Interaction Technology for People with Carpal Tunnel Syndrome or Chronic Ergonomic Problems”. <i>Chittenden Symposium on Assistive Technologies in Health, UIUC</i> . 2017.
Tutorial	“Charm++ Hands-on Tutorial.” <i>Annual Charm++ Workshop and Applications</i> 2015 - 2017.
Poster	“Thermal-Aware Task Scheduling with Neural Network Based Modeling.” <i>Summer Intern Poster Symposium at IBM T.J. Watson Research Center</i> . 2016.
Invited Talk	“Parallel Programming with Migratable Objects: Charm++.” <i>IBM T.J. Watson Research Center</i> . 2016.
Invited Talk	“Processor Variation in Large Scale Supercomputers.” <i>14th Annual Charm++ Workshop and Applications</i> . 2016.
Invited Talk	“TraceR: A Parallel Scalable Network Simulator.” <i>13th Annual Charm++ Workshop and Applications</i> . 2015.
Invited Talk	“Malleable Jobs: Shrink and Expand with Charm++.” <i>13th Annual Charm++ Workshop and Applications</i> . 2015.
Tutorial	“Charm++ Hands-on Tutorial.” <i>Argonne Training Program on Extreme-Scale Computing, ATPESC</i> . 2015.
Invited Talk	“TraceR: A Parallel Trace Replay Tool for HPC Network Simulations.” <i>CODES Workshop, Argonne National Laboratory</i> . 2015.
Poster	“Scalable Trace-driven Parallel Network Simulation.” <i>LLNL Summer Intern Poster Symposium</i> . 2014.

Professional Activities

Organizational Activities

- Program Chair for the HPC for ML Track, *Supercomputing Conference*. 2024.
- Program Committee Member, *MLSys Conference*. 2022 & 2023 & 2024.
- Program Committee Member, *IPDPS Conference*. 2022 & 2023.
- Program Committee Member, *ICML DataPerf Workshop*. 2022.
- Program Committee member, *ICPP Conference*, Performance Track. 2021.
- Technical Program Committee Member, *ICCD Conference*, Systems Track. 2019.
- Program Committee Member, Women in High Performance Computing Workshop (WHPC at SC). 2018.
- Workshop Organizer, “Ethics in Computing” at *Heidelberg Forum*. 2018.
- Conference Co-organizer, *WeSTEM* (Women Empowered in STEM) Conference by Society of Women Engineers (SWE). 2014.

Peer-Reviews

- External Reviewer, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*. 2022.
- Reviewer, *Data Centric AI Workshop, NeurIPS*. 2021.
- Reviewer, *NeurIPS Datasets and Benchmarks Track*. 2021.
- Reviewer, *Systems for Machine Learning Research Proposals, Facebook*. 2019.
- External Reviewer, *IEEE Transactions on Cloud Computing (TCC)*. 2019.
- External Reviewer, *Latin America High Performance Computing Conference (CARLA)*. 2019.

Other Activities

- Member of DataPerf Working Group, MLCommons. <https://dataperf.org>. 2021-2.
- Technical Session Chair, “Efficient Training” at *MLSys*. 2022.
- Technical Session Chair, “Data Centers” & “Energy Efficiency and Measurements” sessions, *IGSC Conference*. 2019.
- Technical Session Chair, “Interfaces Session”, at *15th Annual Charm++ Workshop and Applications*. 2017.

- Technical Session Chair, “Applications Session”, at *14th Annual Charm++ Workshop and Applications*. 2016.
- Graduate Student Mentor for incoming PhDs, *Department of Computer Science at UIUC*. 2013-2015.

Professional Affiliations

ACM, IEEE	Member
SYSTEMS, ANITA BORG INSTITUTE	Member
SOCIETY OF WOMEN ENGINEERS (SWE)	Member
WOMEN IN AI INFRASTRUCTURE AT FACEBOOK	Co-lead (2019-2020)
WOMEN IN COMPUTER SCIENCE (WCS) AT UIUC	Member (2012-17)
TURKISH STUDENT ASSOCIATION (TSA) AT UIUC	Board Member (2014-17)
WOMEN IN HIGH PERFORMANCE COMPUTING (WHPC)	Member
GRADUATE SOCIETY OF WOMEN ENGINEERS (GRADSWE) AT UIUC	Digital Media Coordinator (2013-15)
COMPUTER RESEARCH ASSOCIATION (CRA), GRADUATE WOMEN COHORT	Member and Alumni

Programming Skills

C/C++, Python, Unix Shell, Java, R, Assembly, Arduino, LaTeX, HTML, SQL

Parallel & Distributed Programming: Charm++, MPI, OpenMP, Pthreads, Spark

Machine Learning: PyTorch, Caffe2