Edgar Solomonik

POSITION

Associate Professor, Siebel School of Computing and Data Science, University of Illinois, Urbana-Champaign (UIUC)

CONTACT INFORMATION

solomon2@illinois.edu
4229 Thomas M. Siebel Center, 201 North Goodwin Avenue, Urbana, IL 61801

WEB PAGE

`https://solomonik.cs.illinois.edu/`

PROFESSIONAL APPOINTMENTS

Associate Professor, Department of Computer Science, UIUC	August 2022 – present
Assistant Professor, Department of Computer Science, UIUC	August 2016 – August 2022
Postdoctoral Fellow, Department of Computer Science, ETH Zurich	August 2014 – August 2016

EDUCATION

Ph.D., Computer Science, University of California, Berkeley — August 2010 – August 2014
Advisor: James W. Demmel, Thesis: *Provably efficient algorithms for numerical tensor algebra*
B.S., Computer Science, UIUC — August 2008 – May 2010

AWARDS AND FELLOWSHIPS

Presidential Early Career Award for Scientists and Engineers (PECASE), 2025
Jack Dongarra Early Career Award, 2024
UIUC CS Department C.W. Gear Outstanding Junior Faculty Award, 2022
UIUC Dean's Award for Excellence in Research, 2021
UIUC List of Instructors Ranked as Excellent (Fall 2020, Spring 2022)
NSF CAREER Award, 2020
SIAM Activity Group on Supercomputing Early Career Prize, 2020
IEEE TCHPC Early Career Researchers Excellence Award in High Performance Computing, 2018
Alston S. Householder Prize XX, 2017
ETH Zurich Postdoctoral Fellowship, 2014-2016
Berkeley EECS Department David J. Sakrison Memorial Prize, 2014
NERSC Award for Innovative Use of High Performance Computing, 2013
ACM/IEEE-CS George Michael Memorial High Performance Computing Fellowship, 2013
Distinguished Paper Award, Euro-Par, 2011
Department of Energy Computational Science Graduate Fellowship (DOE CSGF), 2010-2014
Finalist for CRA Outstanding Undergraduate Research Award, 2010

TEACHING

CS 450 / CSE 401 / ECE 491 / MATH 450: Numerical Analysis, UIUC, Spring 2018, Fall 2018, Spring 2021, Spring 2023, Spring 2024, Fall 2026

CS 554 / CSE 512: Parallel Numerical Algorithms, UIUC, Fall 2017, Fall 2019, Fall 2021, Fall 2023, Spring 2026

CS 598-EVS: Tensor Computations, UIUC, Fall 2020, Spring 2022, Fall 2024

CS 598-EVS: Provably Efficient Algorithms for Numerical and Combinatorial Problems, UIUC, Spring 2020

CS 357 / MATH 357: Numerical Methods, UIUC, Spring 2017

CS 598-EVS: Communication Cost Analysis of Algorithms, UIUC, Fall 2016

ADVISING

Current postdoctoral mentees: Maryam Dehghan, Tong Ding

Current PhD advisees: Jack Weinstein, Qizhao Huang, Zhanrui Zhang, Nick Koskelo

Past Postdoctoral mentees: Raghavendra Kanakagiri

Past PhD advisees: Samah Karim (PhD 2022), Linjian Ma (PhD 2023), Edward Hutter (PhD 2024), Navjot Singh (PhD 2025), Yuchen Pang (PhD 2026)

Past BS and MS thesis advisees: Wentao Yang (MS 2022), Navjot Singh (MS 2020), Yiqing Zhou (BS 2020), Tianyi Hao (BS 2020), Youcef Hadjarab (BS 2020), Zecheng Zhang (BS 2019), Pavle Simonovic (BS 2018), Edward Hutter (DOE CSGF fellow, BS 2017), Tobias Wicky (BS 2016, MS 2017, ETH Zurich, co-advised with Torsten Hoefler)

PhD thesis committee member: Ping-Hsuan Tsai (2023), Shelby Lockhart (2023), James Stevens (2021), Hadi Asgharimoghaddam (2021), Si Zhang (2020), Amanda Bienz (2018)

SERVICE

Editorships: Associate Editor for ACM Transactions on Mathematical Software (TOMS); Associate Editor-in-Chief in Algorithms for Journal of Parallel and Distributed Computing (JPDC) 2023-2025

Organization: Workshop on Sparse Tensor Computations 2023 (lead organizer); Networks for Science Workshop 2025 (organizing committee); DRAGSTERS 2023, ARRAY 2017 (organizing committee); SIAM CSE 2023, SIAM PP 2020, 2022, PASC 2018 (minisymposium organizer)

Program committees: SC27 (Algorithms Track Co-Chair); NeurIPS 2024 (area track chair), IPDPS 2024 (program committee track chair); SPAA (2018, 2023, 2024, 2026), SC (2018, 2019, 2020, 2022, 2023, 2024, 2026), ICML 2024, 2026 (reviewer), HPC Asia (2024), ICLR 2024 (reviewer), NeurIPS 2023 (reviewer), ALENEX 2023, IPDPS (2018, 2019, 2021, 2022, 2023, 2025), IAˆ3 (2018, 2019, 2022), ICPP (2022, 2023), ISC (2019, 2021, 2023), EuroPar 2020, SIAM PP 2020, EuroMPI 2019, HIPC 2018, GABB 2017, PACT 2017

Fellowship application review: DOE CSGF (screening 2018-2026, selection committee 2021-2024)

PUBLICATIONS

61.[arXiv] Edgar Solomonik and Markus Bläser. *Simultaneous matrix multiplication by tensor contraction aggregation.* arXiv:2610.26563 [math.NA], October 2026.

60.[SISC] Ping-Hsuan Tsai, Paul Fischer, and Edgar Solomonik *Accelerating Galerkin reduced-order models for turbulent flows with tensor decomposition.* SIAM Journal on Scientific Computing, accepted for publication.

59.[NLAA] Navjot Singh, Edgar Solomonik, Xiaoye Sherry Li, and Yang Liu *Efficient Tensor Completion Algorithms for Highly Oscillatory Operators.* Numerical Linear Algebra with Applications 33(4), e70098, 2026.

58.[arXiv] Yuchen Pang and Edgar Solomonik *Simulating Clifford circuits with Gaussian elimination.* arXiv preprint arXiv:2511.06127, 2025.

57.[arXiv] Edgar Solomonik *Sparse LDL, LU, and inverse butterfly factorization via tree decomposition.* arXiv preprint arXiv:2504.20305, 2025.

56.[FoCM] Caleb Ju, Yifan Zhang, and Edgar Solomonik *Communication lower bounds for nested bilinear algorithms via rank expansion of Kronecker products.* Foundations of Computational Mathematics 25(1), 55–101, 2025.

55.[PhysRev] Louis Schatzki, Linjian Ma, Edgar Solomonik, and Eric Chitambar *Tensor rank and other multipartite entanglement measures of graph states.* Physical Review A 110(3), 032409, 2024.

54.[Quantum] Linjian Ma, Matthew Fishman, Edwin Miles Stoudenmire, and Edgar Solomonik. *Approximate contraction of arbitrary tensor networks with a flexible and efficient density matrix algorithm.* Quantum 8, 1580, 2024.

53.[SPAA] Raghavendra Kanakagiri and Edgar Solomonik *Minimum cost loop nests for contraction of a sparse tensor with a tensor network.* ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pp. 169–181, June 2024.

52.[Euro-Par] Andreas Irmler, Raghavendra Kanakagiri, Sebastian T. Ohlmann, Edgar Solomonik, and Andreas Grüneis *Optimizing distributed tensor contractions using node-aware processor grids.* Lecture Notes in Computer Science, Euro-Par, August 2023.

51.[arXiv] Caleb Ju, Serif Yesil, Mengyuan Sun, Chandra Chekuri, and Edgar Solomonik *Efficient parallel implementation of the multiplicative weight update method for graph-based linear programs.* arXiv:2307.03307 [cs.DC], July 2023.

50.[SISC] Navjot Singh and Edgar Solomonik. *Alternating Mahalanobis distance minimization for accurate and well-conditioned CP decomposition.* SIAM Journal on Scientific Computing 45(6), A2781–A2812, 2023.

49.[SC] Edward Hutter and Edgar Solomonik. *Application performance modeling via tensor completion.* ACM/IEEE Supercomputing Conference, November 2023.

48.[SPAA] Wentao Yang, Vipul Harsh, and Edgar Solomonik. *Optimal round and sample-size complexity for partitioning in parallel sorting.* ACM Symposium on Parallelism in Algorithms and Architectures, June 2023.

47.[ASPLOS] Toluwanimi O. Odemuyiwa, Hadi Asghari-Moghaddam, Michael Pellauer, Kartik Hegde, Po-An Tsai, Neal C. Crago, Aamer Jaleel, John D. Owens, Edgar Solomonik, Joel S. Emer, and Christopher W. Fletcher *Accelerating sparse data orchestration via dynamic reflexive tiling.* ACM International Conference on Architectural Support for Programming Languages and Operating Systems, March 2023.

46.[SciPost] Yang Gao, Phillip Helms, Garnet Kin-Lic Chan, and Edgar Solomonik. *Automatic transformation of irreducible representations for efficient contraction of tensors with cyclic group symmetry.* SciPost Physics Codebases 10, February 2023.

45.[NeurIPS] Linjian Ma and Edgar Solomonik. *Cost-efficient Gaussian tensor network embeddings for tensor-structured inputs.* Conference on Neural Information Processing Systems, 2022.

44.[NeurIPS] Chaoqi Yang, Cheng Qian, Navjot Singh, Cao Xiao, M. Brandon Westover, Edgar Solomonik, and Jimeng Sun. *ATD: Augmenting CP tensor decomposition by self supervision.* Conference on Neural Information Processing Systems, 2022.

43.[SIMAX] Samah Karim and Edgar Solomonik. *Efficient preconditioners for interior point methods via a new Schur-complement-based strategy.* SIAM Journal on Matrix Analysis and Applications 43(4), 1680–1711, 2022.

42.[JPDC] Navjot Singh, Zecheng Zhang, Xiaoxiao Wu, Naijing Zhang, Siyuan Zhang, and Edgar Solomonik. *Distributed-memory tensor completion for generalized loss functions in Python using new sparse tensor kernels.* Journal of Parallel and Distributed Computing 169, 269–285, 2022.

41.[NLAA] Linjian Ma and Edgar Solomonik. *Accelerating alternating least squares for tensor decomposition by pairwise perturbation.* Numerical Linear Algebra with Applications 29(4), e2431, 2022.

40.[SIAM PP] Tim Baer, Raghavendra Kanakagiri, and Edgar Solomonik. *Parallel minimum spanning forest computation using sparse matrix kernels.* Proceedings of the SIAM Conference on Parallel Processing for Scientific Computing, February 2022.

39.[NeurIPS] Linjian Ma and Edgar Solomonik. *Fast and accurate randomized algorithms for low-rank tensor decompositions.* Conference on Neural Information Processing Systems, 2021.

38.[KDD] Chaoqi Yang, Navjot Singh, Cao Xiao, Cheng Qian, Edgar Solomonik, Jimeng Sun. *MTC: Multiresolution tensor completion from partial and coarse observations.* ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2021.

37.[SISC] Edgar Solomonik, James Demmel, and Torsten Hoefler. *Communication lower bounds of bilinear algorithms for symmetric tensor contractions.* SIAM Journal on Scientific Computing 43(5), A3328–A3356, 2021.

36.[SISC] Navjot Singh, Linjian Ma, Hongru Yang, and Edgar Solomonik. *Comparison of accuracy and scalability of Gauss-Newton and alternating least squares for CP decomposition.* SIAM Journal on Scientific Computing 43(4), C290–C311, 2021.

35.[IPDPS] Edward Hutter and Edgar Solomonik. *Accelerating distributed-memory autotuning via statistical analysis of execution paths.* IEEE International Parallel and Distributed Processing Symposium, May 2021.

34.[IPDPS] Linjian Ma and Edgar Solomonik. *Efficient parallel CP decomposition with pairwise perturbation and multi-sweep dimension tree.* IEEE International Parallel and Distributed Processing Symposium, May 2021.

33.[CMAM] Edgar Solomonik and James Demmel. *Fast bilinear algorithms for symmetric tensor contractions.* Computational Methods in Applied Mathematics 21(1), 211–231, 2021.

32.[SC] Yuchen Pang, Tianyi Hao, Annika Dugad, Yiqing Zhou, and Edgar Solomonik. *Efficient 2D tensor network simulation of quantum systems.* ACM/IEEE Supercomputing Conference, November 2020.

31.[SC] Ryan Levy, Edgar Solomonik, and Bryan K. Clark. *Distributed-memory DMRG via sparse and dense parallel tensor contractions.* ACM/IEEE Supercomputing Conference, November 2020.

30.[PACT] Linjian Ma, Jiayu Ye, and Edgar Solomonik. *AutoHOOT: Automatic High-Order Optimization for Tensors.* International Conference on Parallel Architectures and Compilation Techniques, October 2020.

29.[SIREV] Caleb Ju and Edgar Solomonik. *Derivation and analysis of fast bilinear algorithms for convolution.* SIAM Review 62(4), 743–777, 2020.

28.[arXiv] Yifan Zhang and Edgar Solomonik. *On stability of tensor networks and canonical forms.* arXiv:2001.01191 [math.NA], January 2020.

27.[IPDPS] Maciej Besta, Raghavendra Kanakagiri, Harun Mustafa, Mikhail Karasikov, Gunnar Rätsch, Torsten Hoefler, and Edgar Solomonik. *Communication-efficient Jaccard similarity for high-performance distributed genome comparisons.* IEEE International Parallel and Distributed Processing Symposium, May 2020.

26.[MICRO] Kartik Hegde, Hadi Asghari-Moghaddam, Michael Pellauer, Neal Crago, Aamer Jaleel, Edgar Solomonik, Joel Emer, and Christopher W. Fletcher. *ExTensor: An accelerator for sparse tensor algebra.* IEEE/ACM International Symposium on Microarchitecture, October 2019.

25.[SPAA] Vipul Harsh, Laxmikant Kale, and Edgar Solomonik. *Histogram sort with sampling.* ACM Symposium on Parallelism in Algorithms and Architectures, June 2019.

24.[IPDPS] Edward Hutter and Edgar Solomonik. *Communication-avoiding CholeskyQR2 for rectangular matrices.* IEEE International Parallel and Distributed Processing Symposium, May 2019.

23.[SC] Edgar Solomonik, Maciej Besta, Flavio Vella, and Torsten Hoefler. *Scaling betweenness centrality using communication-efficient sparse matrix multiplication.* ACM/IEEE Supercomputing Conference, November 2017.

22.[arXiv] Edwin Pednault, John A. Gunnels, Giacomo Nannicini, Lior Horesh, Thomas Magerlein, Edgar Solomonik, Erik W. Draeger, Eric T. Holland, Robert Wisnieff. *Pareto-efficient quantum circuit simulation using tensor contraction deferral.* arXiv:1710.05867 [quant-ph], October 2017.

21.[SPAA] Edgar Solomonik, Grey Ballard, James Demmel, and Torsten Hoefler. *A communication-avoiding parallel algorithm for the symmetric eigenvalue problem.* ACM Symposium on Parallelism in Algorithms and Architectures, 2017.

20.[HPDC] Maciej Besta, Michał Podstawski, Linus Groner, Edgar Solomonik, and Torsten Hoefler. *To push or to pull: on reducing communication and synchronization in graph computations.* 26th ACM Symposium on High Performance Parallel and Distributed Computing, Washington DC, 2017.

19.[IPDPS] Tobias Wicky, Edgar Solomonik, and Torsten Hoefler. *Communication-avoiding parallel algorithms for solving triangular systems of linear equations.* IEEE International Parallel and Distributed Processing Symposium, 2017.

18.[IPDPS] Maciej Besta, Florian Marending, Edgar Solomonik and Torsten Hoefler. *SlimSell: A vectorizable graph representation for breadth-first search.* IEEE International Parallel and Distributed Processing Symposium, 2017.

17.[TOPC] Edgar Solomonik, Erin Carson, Nicholas Knight, and James Demmel. *Trade-offs between synchronization, communication, and computation in parallel linear algebra computations.* ACM Transactions on Parallel Computing 3(1), Article 3, 2016.

16.[arXiv] Edgar Solomonik and Torsten Hoefler. *Sparse tensor algebra as a parallel programming model.* arXiv:1512.00066 [cs.MS], 2015.

15.[JPDC] Edgar Solomonik, Devin Matthews, Jeff R. Hammond, John F. Stanton, and James Demmel. *A massively parallel tensor contraction framework for coupled-cluster computations.* Journal of Parallel and Distributed Computing 74(12), 3176–3190, 2014.

14.[SPAA] Edgar Solomonik, Erin Carson, Nicholas Knight, and James Demmel. *Tradeoffs between synchronization, communication, and computation in parallel linear algebra computations.* ACM Symposium on Parallelism in Algorithms and Architectures, 2014.

13.[IPDPS] Grey Ballard, James Demmel, Laura Grigori, Mathias Jacquelin, Hong Diep Nguyen, and Edgar Solomonik. *Reconstructing Householder vectors from Tall-Skinny QR.* IEEE International Parallel and Distributed Processing Symposium, 2014.

12.[IPDPS] Edgar Solomonik, Devin Matthews, Jeff Hammond, and James Demmel. *Cyclops Tensor Framework: reducing communication and eliminating load imbalance in massively parallel contractions.* IEEE International Parallel and Distributed Processing Symposium, 2013.

11.[IPDPS] Edgar Solomonik, Aydin Buluç, and James Demmel. *Minimizing communication in all-pairs shortest paths.* IEEE International Parallel and Distributed Processing Symposium, 2013.

10.[IPDPS] Michael Driscoll, Evangelos Georganas, Penporn Koanantakool, Edgar Solomonik, and Katherine Yelick. *A communication-optimal n-body algorithm for direct interactions.* IEEE International Parallel and Distributed Processing Symposium, 2013.

9. [VECPAR] Edgar Solomonik and James Demmel. *Matrix multiplication on multidimensional torus networks.* Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013.

8. [SC] Evangelos Georganas, Jorge Gonzalez-Dominguez, Edgar Solomonik, Yili Zheng, Juan Tourino and Katherine Yelick. *Communication avoiding and overlapping for numerical linear algebra.* ACM/IEEE Supercomputing Conference, 2012.

7. [UCB] Edgar Solomonik, Jeff Hammond, and James Demmel. *A preliminary analysis of Cyclops Tensor Framework.* Technical Report UCB/EECS-2012-29, University of California, Berkeley, 2012.

6. [SC] Edgar Solomonik, Abhinav Bhatele, and James Demmel. *Improving communication performance in dense linear algebra via topology aware collectives.* ACM/IEEE Supercomputing Conference, 2011.

5. [Euro-Par] Edgar Solomonik and James Demmel. *Communication-optimal parallel 2.5D matrix multiplication and LU factorization algorithms.* Lecture Notes in Computer Science, Euro-Par, 2011.

4. [Enc. Par. Comp.] Laxmikant V. Kalé and Edgar Solomonik. *Sorting.* Encyclopedia of Parallel Computing, Springer, David Padua, Ed., 2011.

3. [IJHPCA] Abhinav Bhatele, Lukasz Wesolowski, Eric Bohm, Edgar Solomonik, and Laxmikant V. Kale. *Understanding application performance via micro-benchmarks on three large supercomputers: Intrepid, Ranger and Jaguar.* International Journal of High Performance Computing Applications 24(4), 411–427, 2010.

2. [IPDPS] Edgar Solomonik and Laxmikant V. Kale. *Highly scalable parallel sorting.* IEEE International Parallel and Distributed Processing Symposium, 2010.

1. [ParaPLOP] Vivek Kale and Edgar Solomonik. *Parallel sorting pattern.* Workshop on Parallel Programming Patterns, 2010.

FUNDING

1. [DARPA MOCHA] Compiler 2.0: Revolutionizing Compiler Construction with Neuro-Symbolic Machine Learning co-PI (PI Michael Gordon, co-PIs Charith Mendis, Saman Amarasinghe, Jonathan Ragan-Kelley, Adam Chlipala, Armando Solar-Lezama, Michael Carbin), 2025–2028.

2. [DOD AFOSR] Tensor Approaches for Simulating Kinetic Systems. co-PI, $1,210,000 (PI Jingmei Qiu, co-PIs Daniele Venturi, Jingwei Hu, Lexing Ying), 2024-2028.

3. [NSF DMS] Conference: Workshop on Sparse Tensor Computations. PI, participant funding $43,463 (Co-PI Grey Ballard), 2023-2024.

4. [DOE ASCR] Sparsitute: A Mathematical Institute for Sparse Computations in Science and Engineering. co-PI, $1,250,000 (PI Aydin Buluc, co-PIs Ramakrishnan Kannan, Ariful Azad, Grey Ballard, David Gleich), 2022-2027.

5. [NSF MPS] QLCI-CI: NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks. senior personnel, total amount $7,700,000 (PI Brian DeMarco), 2020-2025.

6. [NSF OAC] PPoSS: Planning: A cross-layer approach to accelerate large-scale graph computations on distributed platforms. co-PI, total amount $250,000 (PI Josep Torrellas, Co-PIs Sasa Misailovic, Chandra Chekuri), 2020-2021.

7. [NSF OAC] CAREER: Next-generation infrastructure for tensor computations. PI $500,000, 2020-2025.

8. [NSF OAC] Collaborative Research: Frameworks: Scalable modular software and methods for high-accuracy materials and condensed phase chemistry simulation. PI, total amount $1,854,820 (co-PIs Lucas Wagner, Garnet Chan, Timothy Berkelbach), 2019-2022.

9. [NSF CCF] RAISE/TAQS: Quantum simulation of materials and molecules using quantum computation. co-PI, total amount $1,000,000 (PI Fernando Brandao, co-PIs Garnet Chan, Austin Minnich), 2018-2021.