

Shaden Smith

Curriculum Vitæ

August 2019

Address: Parallel Computing Laboratory
Intel SC12
3600 Juliette Ln
Santa Clara, CA 95054

Email: shaden@cs.umn.edu

WWW: <http://shaden.io/>

GitHub: <https://github.com/ShadenSmith>

Research Interests

My research is primarily concerned with the development of scalable, high performance algorithms for applications in data mining and machine learning. *Irregular* applications are of particular interest to me, such as those that operate on sparse graphs, matrices, and tensors. My work has focused on large-scale sparse tensor factorization and is culminated in SPLATT, an open source software toolkit for tensor factorization and related kernels. SPLATT has been scaled to over 16,000 cores and is actively used by academic, industry, and government researchers.

Education

- **UNIVERSITY OF MINNESOTA–TWIN CITIES** **Minneapolis, MN**
Ph.D., Computer Science. 2012 – 2019
Thesis title: “Algorithms for Large-Scale Sparse Tensor Factorization”
Thesis advisor: Dr. George Karypis
- **UNIVERSITY OF KENTUCKY** **Lexington, KY**
B.S., Computer Science. 2009 – 2012
Minor in Mathematics

Awards & Honors

- Intel Datacenter Engineering & Architecture Division Recognition Award, 2019.
- **ACM/IEEE-CS George Michael HPC Fellowship**, 2017. “For his work on efficient and parallel large-scale sparse tensor factorization for machine learning applications.” [Press release](#).
- **Distinguished Paper Award**, Euro-Par 2017. *Accelerating the Tucker Decomposition with Compressed Sparse Tensors*.
- HPEC GraphChallenge Finalist, 2017. *Truss Decompositions on Shared-Memory Parallel Systems*.
- HPEC GraphChallenge Finalist, 2017. *Exploring Optimizations on Shared-Memory Platforms for Parallel Triangle Counting Algorithms*.
- **Best Student Paper Finalist**, Supercomputing 2016. *An Exploration of Optimization Algorithms for High Performance Tensor Completion*.
- **Doctoral Dissertation Fellowship**, University of Minnesota, 2016-2017. Award covers 100% stipend, tuition, and fees for the academic year.
- Outstanding Graduating Senior Award, University of Kentucky, 2012. One graduating computer science student is selected each year.
- Student Symposium Winner, Lexmark Student Symposium, 2011. Awarded top prize in a symposium of student internship projects.
- Travel Awards: CLSAC '17, SIAM CSE'17, IPDPS'16, SIAM PP'16, IPDPS'15.

Professional Experience

- **INTEL LABS** **Santa Clara, CA**
Research Scientist – Parallel Computing Laboratory *January 2018 – Present*
 - Scaling up hardware, software, and algorithms for sparse problems in machine learning.
 - Technical lead of workload analysis for major hardware/software co-design effort.
 - Principal investigator for team researching streaming tensor algorithms and applications.

- **INTEL LABS**
Graduate Research Intern – Parallel Computing Laboratory
 – Optimized sparse tensor and graph algorithms on many-core processors.
 – Co-designed hardware/software to optimize sparse matrix and graph computations.

Santa Clara, CA
Summer 2017
- **LAWRENCE LIVERMORE NATIONAL LABORATORY**
Graduate Intern – High Energy Density Physics Department
 – Developed OpenACC implementation of LULESH, a hydrocode proxy application.
 – Provided regular feedback on compiler implementation and language design.

Livermore, CA
Summer 2013
- **LEXMARK INTERNATIONAL**
Graduate Intern – Material Technology Department
 – Developed on a massively parallel particle flow modeling engine in CUDA.

Lexington, KY
Summer 2012
- **LEXMARK INTERNATIONAL**
Undergraduate Intern – Material Technology Department
 – Designed and implemented a high-performance particle flow modeling engine in C++.
 – Project awarded 1st place prize in student symposium.

Lexington, KY
Summer 2011
- **UNIVERSITY OF KENTUCKY**
Research Assistant
 – Researched answer set programming and satisfiability solvers.

Lexington, KY
May 2010 – July 2012
- **UNIVERSITY OF KENTUCKY**
Research Assistant – Center for Visualization and Virtual Environments February 2010 – September 2010
 – Developed a tool for digitally unwrapping ancient scrolls (EDUCE project).
 – Analyzed and manipulated high-resolution CT scan data.

Lexington, KY

Teaching Experience

- **Co-Instructor, INTRODUCTION TO PARALLEL COMPUTING**
Department of Computer Science and Engineering, University of Minnesota

Spring 2017
- **Teaching Assistant, INTRODUCTION TO PARALLEL COMPUTING**
Department of Computer Science and Engineering, University of Minnesota

Spring 2014, Spring 2015
- **Teaching Assistant, INTRODUCTION TO INTERNET PROGRAMMING**
Department of Computer Science and Engineering, University of Minnesota

Fall 2012, Spring 2013, Fall 2013
- **Teaching Assistant, INTRODUCTION TO ALGORITHM DESIGN & ANALYSIS**
Department of Computer Science, University of Kentucky

Fall 2011, Spring 2012

Publications

Book Chapters

- [B 1] Evangelia Christakopoulou, **Shaden Smith**, Mohit Sharma, Alex Richards, David C. Anastasiu, and George Karypis. Scalability and distribution of collaborative recommenders. In *Collaborative Recommendations: Algorithms, Practical Challenges and Applications*. World Scientific Publishing, Singapore, 2019 (forthcoming).
- [B 2] David C. Anastasiu, Jeremy Iverson, **Shaden Smith**, and George Karypis. Big data frequent pattern mining. In *Frequent Pattern Mining*. 2014.

Journals

- [J 1] **Shaden Smith**, Jongsoo Park, and George Karypis. HPC formulations of optimization algorithms for tensor completion. *Parallel Computing*, (74), 2018.
- [J 2] David C. Anastasiu, Evangelia Christakopoulou, **Shaden Smith**, Mohit Sharma, and George Karypis. Big data and recommender systems. *Novática: Journal of the Spanish Computer Scientist Association*, (240), 2016.

Refereed Conference & Workshop Proceedings

- [C 1] **Shaden Smith**, Kejun Huang, Nicholas D. Sidiropoulos, and George Karypis. Streaming tensor factorization for infinite data sources. *Proceedings of the 2018 SIAM International Conference on Data Mining (SDM)*, 2018.
- [C 2] Jee W. Choi, Xing Liu, **Shaden Smith**, and Tyler Simon. Blocking optimization techniques for sparse tensor computation. *32nd IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2018.
- [C 3] **Shaden Smith**, Xing Liu, Nesreen K. Ahmed, Ancy Sarah Tom, Fabrizio Petrini, and George Karypis. Truss decompositions on shared-memory parallel systems. *IEEE High Performance Extreme Computing Conference (HPEC)*, 2017. **GraphChallenge Finalist**.
- [C 4] Ancy Sarah Tom, Narayanan Sundaram, Nesreen K. Ahmed, **Shaden Smith**, Stijn Eyerman, Midhunchandra Kodyath, Ibrahim Hur, Fabrizio Petrini, and George Karypis. Exploring optimizations on shared-memory platforms for parallel triangle counting algorithms. *IEEE High Performance Extreme Computing Conference (HPEC)*, 2017. **GraphChallenge Finalist**.
- [C 5] **Shaden Smith** and George Karypis. Accelerating the Tucker decomposition with compressed sparse tensors. *European Conference on Parallel Processing (Euro-Par)*, 2017, **Distinguished Paper Award**.
- [C 6] **Shaden Smith**, Alec Beri, and George Karypis. Constrained tensor factorization with accelerated AO-ADMM. *46th International Conference on Parallel Processing (ICPP)*, 2017.
- [C 7] Michael Anderson, **Shaden Smith**, Narayanan Sundaram, Mihai Capotă, Zheguang Zhao, Subramanya Dullloor, Nadathur Satish, and Theodore L. Willke. Bridging the gap between HPC and Big Data frameworks. *Proceedings of the VLDB Endowment (PVLDB)*, 2017.
- [C 8] **Shaden Smith**, Jongsoo Park, and George Karypis. Sparse tensor factorization on many-core processors with high-bandwidth memory. *31st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2017.
- [C 9] **Shaden Smith**, Jongsoo Park, and George Karypis. An exploration of optimization algorithms for high performance tensor completion. *Proceedings of the 2016 ACM/IEEE Conference on Supercomputing (SC)*, 2016. **Finalist, Best Student Paper**.
- [C 10] **Shaden Smith** and George Karypis. A medium-grained algorithm for distributed sparse tensor factorization. *30th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2016.
- [C 11] **Shaden Smith** and George Karypis. Tensor-matrix products with a compressed sparse tensor. *Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms (IA3)*, 2015.
- [C 12] **Shaden Smith**, Niranjay Ravindran, Nicholas D. Sidiropoulos, and George Karypis. SPLATT: Efficient and parallel sparse tensor-matrix multiplication. *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2015.
- [C 13] Niranjay Ravindran, Nicholas D. Sidiropoulos, **Shaden Smith**, and George Karypis. Memory-efficient parallel computation of tensor and matrix products for big tensor decomposition. *Proceedings of the Asilomar Conference on Signals, Systems, and Computers*, 2014.
- [C 14] Yuliya Lierler, **Shaden Smith**, Miroslaw Truszczynski, and Alex Westlund. Weighted-sequence problem: ASP vs CASP and declarative vs problem-oriented solving. Springer, 2012.

Invited Talks & Posters

- [U 1] **Shaden Smith**. Invited talk. Dagstuhl Seminar, Tensor Computations: Applications and Optimization.

- [U 2] **Shaden Smith**. Invited talk. AI and Tensor Factorization for Physical Chemical and Biological Systems, September 2019.
- [U 3] **Shaden Smith**. Invited talk. KDD Workshop on Tensor Methods for Emerging Data Science Challenges (TMEDSC'19), August 2019.
- [U 4] **Shaden Smith**, Jongsoo Park, and George Karypis. HPC formulations of optimization algorithms for tensor completion. SIAM Conference on Computational Science and Engineering (CSE), Minisymposium: High-Performance Tensor Computation and its Applications, February 2019.
- [U 5] **Shaden Smith**. Scaling up sparse tensor factorization: Challenges and lessons learned. Sandia Technical Seminar, June 2018.
- [U 6] **Shaden Smith**. Large-scale sparse tensor factorization: Scale up using these three simple rules! ThRee-way methods In Chemistry And Psychology (TRICAP), invited talk, June 2018.
- [U 7] **Shaden Smith**. Scalability challenges in large-scale sparse tensor factorization. University of California, Santa Barbara (UCSB), ECE 594BB: Selected Topics in High-Dimensional Tensor Data Analysis, invited lecture, May 2018.
- [U 8] **Shaden Smith** and George Karypis. Accelerating the Tucker decomposition with compressed sparse tensors. SIAM Conference on Parallel Processing for Scientific Computing (PP), Minisymposium: Tensor Decomposition for High Performance Data Analytics, March 2018.
- [U 9] Ancy Sarah Tom, **Shaden Smith**, and George Karypis. Triangle counting and truss decomposition on modern parallel architectures (*talk given by Ancy Sarah Tom*). SIAM Conference on Parallel Processing for Scientific Computing (PP), Minisymposium: Architecture-Aware Graph Analytics, March 2018.
- [U 10] **Shaden Smith**. Algorithms for large-scale sparse tensor factorization. Chesapeake Large-Scale Analytics Conference (CLSAC), invited poster, October 2017.
- [U 11] **Shaden Smith**, Jongsoo Park, and George Karypis. An exploration of optimization algorithms for high performance tensor completion. SIAM Conference on Computational Science and Engineering (CSE), Minisymposium: Tensor Decompositions: Applications and Efficient Algorithms, February 2017.
- [U 12] **Shaden Smith**. Algorithms for large-scale sparse tensor factorization. University of Minnesota Doctoral Dissertation Fellowship (DDF) showcase, poster, April 2017.
- [U 13] **Shaden Smith** and George Karypis. High performance sparse tensor factorization. Intel Labs, invited talk, April 2016.
- [U 14] **Shaden Smith** and George Karypis. SPLATT: Enabling large-scale sparse tensor analysis. Workshop on Algorithms for Modern Massive Data Sets (MMDS), poster, April 2016.
- [U 15] **Shaden Smith**, Jongsoo Park, and George Karypis. An exploration of optimization algorithms for high performance tensor completion. The 9th International Workshop on Parallel Matrix Algorithms and Applications (PMAA), Minisymposium: Sparse Matrix and Tensor Computations, July 2016.
- [U 16] **Shaden Smith** and George Karypis. Efficient factorization with compressed sparse tensors. SIAM Conference on Parallel Processing for Scientific Computing (PP), Minisymposium: Parallel Algorithms for Tensor Computations, March 2016.
- [U 17] **Shaden Smith** and Peter Robinson. LULESH and openacc: To exascale and beyond!!! PGI OpenACC Workshop, September 2013.

- [U 18] **Shaden Smith** and Jerry Fish. 2010: A GPU odyssey. Lexmark Celebrate Success Seminar, August 2012.
- [U 19] **Shaden Smith** and Jerry Fish. Particle flow modeling or: How I learned to stop worrying and love DEM. Lexmark Student Symposium, August 2011.

Software

- [S 1] **Shaden Smith** and George Karypis. SPLATT: the Surprisingly Parallel spArse Tensor Toolkit. <http://cs.umn.edu/~splatt/>, 2015.
- [S 2] **Shaden Smith**, Jee W. Choi, Jiajia Li, Richard Vuduc, Jongsoo Park, Xing Liu, and George Karypis. FROSTT: The Formidable Repository of Open Sparse Tensors and Tools. <http://frostdt.io/>, 2017.
- [S 3] Jeremy Cohen, Rasmus Bro, Tamara Kolda, and **Shaden Smith**. TensorWorld: a community hub for tensor-related software development and discussion. <http://tensorworld.org/>, 2018.

Professional Service

- *Panelist:*
 - Student Success Stories, at *Supercomputing 2016*
 - Experiencing HPC for Undergraduates, at *Supercomputing 2016*
- *Program Committee Member:*
 - IEEE Cluster (CLUSTER'19)
 - KDD Workshop on Tensor Methods for Emerging Data Science Challenges (TMEDSC'19)
 - International Workshop on the Intersection of High Performance Computing and Machine Learning (HPCaML'19)
 - International Conference on Parallel Processing (ICPP'19)
 - IEEE Fifth International Workshop on High Performance Big Graph Data Management, Analysis, and Mining (BigGraphs'18)
 - IEEE Cluster (CLUSTER'18)
 - International Conference on Big Data (Industry & Government) (BigData'18)
 - The Eighth Workshop on Irregular Applications: Architectures and Algorithms (IA3'18)
 - The Seventh Workshop on Irregular Applications: Architectures and Algorithms (IA3'17)
- *Journal reviewer:*
 - Journal of Parallel and Distributed Computing (JPDC)
 - Algorithmica
 - Transactions on Parallel and Distributed Systems (TPDC)
 - International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)
- *Conference reviewer:*
 - Very Large Data Bases (VLDB)
 - International Conference on Data Mining (ICDM)
 - SIAM International Conference on Data Mining (SDM)
 - International Conference on Parallel Processing (ICPP)
 - International Parallel and Distributed Processing Symposium (IPDPS)
 - International Conference on Big Data (BigData)
 - International Conference on Data Science and Advanced Analytics (DSAA)

- *Workshop reviewer:*
 - Knowledge Representation and Automated Reasoning (RCRA)
- *Benchmark Contributor:* Answer Set Programming Competition (ASPCOMP)
Co-authored the following benchmarks:
 - ASPCOMP 2013 – *Weighted-Sequence Problem*
 - ASPCOMP 2011 & 2013 – *Hanoi Tower*
- *Student President:* Association for Computing Machinery (ACM)– University of Kentucky Student Chapter (Fall 2011 – Spring 2012)

Membership in Professional Organizations

- Association for Computing Machinery (ACM)
 - ACM Committee on Women (ACM-W)
 - ACM Special Interest Group on High Performance Computing (SIGHPC)
- Institute of Electrical and Electronics Engineers (IEEE)
- Society for Industrial and Applied Mathematics (SIAM)