

Richard (Yang) Peng
Associate Professor
Computer Science Department, Carnegie Mellon University
(Updated November 2024)

Contact

GHC 7115, 5000 Forbes Ave, Pittsburgh PA, 15213, USA
Email: yangp@cs.cmu.edu
Phone: (412) 268-9043 (office)

Contents

I	Earned Degrees	2
II	Employment History	2
III	Honors and Awards	2
IV	Research, Scholarship, and Creative Activities	3
A	Referred Publications and Submitted Articles	3
A.1	Published and Accepted Journal Articles	3
A.2	Conference Presentation with Proceedings (Refereed)	4
B	Other Refereed Material	10
D	Presentations	10
D.1	Invited Speaker at Workshops	10
D.2	Invited Talks at Seminars	11
E	Grants and Contracts	12
E.1	As Principle Investigator	12
F	Other Scholarly and Creative Accomplishments	13
V	Education	13
A	Courses Taught	13
B	Individual Student Guidance	14
B.1	Ph.D. Students	14
B.2	M.S. Students	14
B.3	Postdocs	14
B.4	Service on thesis or dissertation committees	15
C	Education Innovations and Other Contributions	15
VI	Service	16
A	Professional Contributions	16
A.1	Conference Program Committees	16
A.2	Conference Reviewer	16
A.3	Journal Reviewer	16

I Earned Degrees

- *Postdoctoral Instructor in Applied Mathematics*,
Department of Mathematics, Massachusetts Institute of Technology, 2013 - 2015,
Supervisor: Professor Jonathan A. Kelner.
- *Ph.D. in Computer Science*,
Carnegie Mellon University, 2009 - 2013,
Advisor: Professor Gary L. Miller,
Thesis: Algorithm Design using Spectral Graph Theory.
- *B. Math*,
University of Waterloo, 2006 - 2009,
Double major in Computer Science and Combinatorics & Optimization.

II Employment History

- *Associate Professor*, Carnegie Mellon University, August 2023 - present.
- *Visiting Associate Professor*, Carnegie Mellon University, August 2022 - July 2023.
- *Associate Professor*, University of Waterloo, January 2022 - July 2023.
- *Associate Professor*, Georgia Institute of Technology, August 2021 - December 2021.
- *Assistant Professor*, Georgia Institute of Technology, August 2015 - August 2021.
- *Adjunct Professor*,
 - Georgia Institute of Technology, January 2022 - present,
 - University of Waterloo, January 2021 - December 2021, August 2023 - present.
- *Visiting Researcher*, Microsoft Research Redmond, January - August 2019.
- *Visiting Professor*, Institute of Theoretical Computer Science, Shanghai University of Finance and Economics. Summers 2016 - 2018.
- *Research Intern*, Microsoft Research New England, Summer 2011.
- *Software Engineering Intern*, Google Seattle, April - July 2009.
- *Undergraduate Research Assistant*, University of Waterloo, May - August 2008.

III Honors and Awards

- Best Paper Award, IEEE 63rd Annual Symposium on Foundations of Computer Science, 2022, for paper titled “Maximum Flow and Minimum-Cost Flow in Almost-Linear Time”.
- Best Paper Award, ACM-SIAM Symposium on Discrete Algorithms, 2021, for paper titled “Solving Sparse Linear Systems Faster than Matrix Multiplication”.
- *CMU SCS Dissertation Award*, 2013.
- *Microsoft Research PhD Fellowship*, 2011.

IV Research, Scholarship, and Creative Activities

A Referred Publications and Submitted Articles

A.1 Published and Accepted Journal Articles

- Solving Sparse Linear Systems Faster than Matrix Multiplication. with Santosh Vempala. In Communications of the ACM, Vol. 67, no. 7, pp. 79-86, 2024.
- Distance queries over dynamic interval graphs. with Jingbang Chen, He Meng, Ian J. Munro, Kaiyu Wu, and Daniel J. Zhang. In Computational Geometry, Volume 122, p.102103, 2024.
- Almost-linear-time algorithms for maximum flow and minimum-cost flow. With Li Chen, Rasmus Kyng, Yang P. Liu, Maximilian Probst Gutenberg, and Sushant Sachdeva. In Communications of the ACM 66, no. 12, pp. 85-92, 2023.
- A Combinatorial Cut-Toggling Algorithm for Solving Laplacian Linear Systems. with Monika Henzinger, Billy Jin, and David P. Williamson. In Algorithmica 85 ,no. 12, pp. 3680-3716, 2023. Preliminary version in ITCS 2023. Available at: <https://arxiv.org/abs/2010.16316>.
- Towards Lightweight and Automated Representation Learning System for Networks. with Yuyang Xie, Jiezhong Qiu, Laxman Dhulipala, Wenjian Yu, Jie Tang, and Chi Wang. In IEEE Transactions on Knowledge and Data Engineering 35, no. 9, pp.9613-9627, 2023. Preliminary version in SIGMOD 2021. Available at: <https://keg.cs.tsinghua.edu.cn/jietang/publications/SIGMOD21-Qiu-et-al-LightNE.pdf>
- Graph Sparsification, Spectral Sketches, and Faster Resistance Computation via Short Cycle Decompositions. with Timothy Chu, Yu Gao, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang. In SIAM Journal on Computing (SICOMP) 52, no. 6, pp. FOCS18-85-FOCS18-157, 2020. Preliminary version in FOCS 2018. Invited to the special issue. Available at: <https://arxiv.org/abs/1805.12051>.
- Determinant-Preserving Sparsification of SDDM Matrices. with David Durfee, John Peebles, and Anup B. Rao. In SIAM Journal on Computing (SICOMP) 49, no.4, pp. FOCS17-350-FOCS17-408, 2020. Preliminary version in FOCS 2017. Invited to the special issue. Available at: <https://arxiv.org/abs/1705.00985>.
- Partitioning Well-Clustered Graphs: Spectral Clustering Works! with He Sun and Luca Zanetti. In SIAM Journal on Computing (SICOMP) 46, no. 2, pp 710-743, 2017. Preliminary version in COLT 2015. Available at <http://arxiv.org/abs/1411.2021>.
- Faster Spectral Sparsification and Numerical Algorithms for SDD Matrices. with Alex Levin and Ioannis Koutis. In ACM Transactions on Algorithms (TALG) 12, no.2, 17, 2016. Preliminary version appeared in STACS 2012. Available at: <http://arxiv.org/abs/1209.5821>.

- Approaching Optimality for Solving SDD Linear Systems
with Ioannis Koutis and Gary L. Miller.
In SIAM Journal on Computing (SICOMP) 43, no. 1, pp 337-354, 2014.
Preliminary version appeared in FOCS 2010. Invited to the special issue.
Available at <http://arxiv.org/abs/1003.2958>.
- Near linear Linear-Work Parallel SDD Solvers, Low-Diameter Decomposition, and Low-Stretch Subgraphs.
with Guy E. Blelloch, Anupam Gupta, Ioannis Koutis, Gary L. Miller, and Kanat Tangwongsan.
In Theory of Computing Systems 55, no. 3, pp. 521-554, 2014.
Preliminary version appeared in SPAA 2011.
Available at <http://arxiv.org/abs/1111.1750>.
- A Fast Solver for a Class of Linear Systems.
with Ioannis Koutis and Gary L. Miller.
In Communications of the ACM 55, no. 10, pp. 99-107, 2012.
- Efficient Triangle Counting in Large Graphs via Degree-based Vertex Partitioning.
with Mihail Kolountzakis, Gary Miller, and Charalampos Tsourakakis.
In Internet Mathematics 8, no. 1-2, pp. 161-185, 2012.
Preliminary version appeared in WAW 2010.
Available at <http://arxiv.org/abs/1011.0468>.
- Approximate Dynamic Programming for Fast Denoising of aCGH Data.
with Gary L. Miller, Russell Schwartz, and Charalampos E. Tsourakakis.
In ACM Journal of Experimental Algorithmics 16, no. 1.8, 2011.
Preliminary version appeared in SODA 2011.
Available at <http://arxiv.org/abs/1003.4942>.

A.2 Conference Presentation with Proceedings (Refereed)

- Scalable Algorithm for Finding Balanced Subgraphs with Tolerance in Signed Networks.
with Jingbang Chen, Qiuyang Mang, Hangrui Zhou, Yu Gao, and Chenhao Ma.
In KDD 2024.
Available at: <https://arxiv.org/abs/2402.05006>.
- Incremental Approximate Maximum Flow on Undirected Graphs in Subpolynomial Update Time.
with Jan van den Brand, Li Chen, Rasmus Kyng, Yang P. Liu, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford.
In SODA 2024.
Available at: <https://arxiv.org/abs/2311.03174>.
- A Deterministic Almost-Linear Time Algorithm for Minimum-Cost Flow. with Jan van den Brand, Li Chen, Rasmus Kyng, Yang P. Liu, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford.
In FOCS 2023.
Available at: <https://arxiv.org/abs/2309.16629>.
- Exponential Convergence of Sinkhorn Under Regularization Scheduling.
with Jingbang Chen, Li Chen, Yang P. Liu, and Arvind Ramaswami.
In ACDA 2023.
Available at: <https://arxiv.org/abs/2207.00736>

- Maximum Flow and Minimum-Cost Flow in Almost-Linear Time.
with Li Chen, Rasmus Kyng, Yang P. Liu, Maximilian Probst Gutenberg and Sushant Sachdeva.
In FOCS 2022. Best Paper Award.
Available at: <https://arxiv.org/abs/2203.00671>
- Faster Maxflow via Improved Dynamic Spectral Vertex Sparsifiers.
with Jan van den Brand, Yu Gao, Arun Jambulapati, Yin Tat Lee, Yang P. Liu, and Aaron Sidford.
In STOC 2022.
Available at: <https://arxiv.org/abs/2112.00722>
- Sparsified Block Elimination for Directed Laplacians.
with Zhuoqing Song.
In STOC 2022.
Available at: <https://arxiv.org/abs/2111.10257>
- Nested Dissection Meets IPMs: Planar Min-Cost Flow in Nearly-Linear Time.
with Sally Dong, Yu Gao, Gramoz Goranci, Yin Tat Lee, Sushant Sachdeva, and Guanghao Ye.
In SODA 2022.
Available at: <https://arxiv.org/abs/2205.01562>
- Fully Dynamic Electrical Flows: Sparse Maxflow Faster Than Goldberg-Rao.
with Yu Gao, and Yang P. Liu.
In FOCS 2021.
Available at: <https://arxiv.org/abs/2101.07233>
- Minor Sparsifiers and the Distributed Laplacian Paradigm.
with Sebastian Forster, Gramoz Goranci, Yang P. Liu, Xiaorui Sun, and Mingquan Ye
In FOCS 2021.
Available at: <https://arxiv.org/abs/2012.15675>
- ℓ_2 -norm Flow Diffusion in Near-Linear Time.
with Li Chen, and Di Wang.
In FOCS 2021.
Available at: <https://arxiv.org/abs/2105.14629>.
- Solving Sparse Linear Systems Faster than Matrix Multiplication.
with Santosh Vempala.
In SODA 2021. Best Paper Award.
Available at <https://arxiv.org/abs/2105.14629>
- Vertex Sparsification for Edge Connectivity.
with Parinya Chalermsook, Syamantak Das, Bundit Laekhanukit, Yunbum Kook, Yang P. Liu, Mark Sellke, and Daniel Vaz.
In SODA 2021.
Available at: <https://arxiv.org/abs/2007.07862>.
- Concentration Bounds for Co-occurrence Matrices of Markov Chains.
with Jiezhong Qiu, Chi Wang, Ben Liao, and Jie Tang.
In NeurIPS 2020.
Available at: <https://arxiv.org/abs/2008.02464>.
- Bipartite Matching in Nearly-linear Time on Moderately Dense Graphs.
with Jan van den Brand, Yin Tat Lee, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford,

Zhao Song, and Di Wang.
In FOCS 2020. Invited to the special issue (declined).
Available at: <https://arxiv.org/abs/2009.01802>

- Fast Dynamic Cuts, Distances and Effective Resistances via Vertex Sparsifiers.
with Li Chen, Gramoz Goranci, Monika Henzinger, and Thatchaphol Saranurak.
In FOCS 2020.
Available at: <https://arxiv.org/abs/2005.02368>
- A Deterministic Algorithm for Balanced Cut with Applications to Dynamic Connectivity, Flows, and Beyond.
with Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, and Thatchaphol Saranurak.
In FOCS 2020.
Available at: <https://arxiv.org/abs/1910.08025>
- Faster Graph Embeddings via Coarsening.
with Matthew Fahrbach, Gramoz Goranci, Sushant Sachdeva, and Chi Wang.
In ICML 2020.
Available at: <https://arxiv.org/abs/2007.02817>
- Flowless: Extracting Densest Subgraphs Without Flow Computations.
with Digvijay Boob, Yu Gao, Saurabh Sawlani, Charalampos E. Tsourakakis, Di Wang, and Junxing Wang.
In WWW 2020.
Available at: <https://arxiv.org/abs/1910.07087>
- Parallel Batched-Dynamic Graph Algorithms in Constant Rounds.
with Laxman Dhulipala, David Durfee, Janardhan Kulkarni, Saurabh Sawlani, and Xiaorui Sun.
In SODA 2020.
Available at: <https://arxiv.org/abs/1908.01956>.
- Fast, Provably convergent IRLS Algorithm for p-norm Linear Regression.
with Deeksha Adil, and Sushant Sachdeva.
In NeurIPS 2019.
Available at: <https://arxiv.org/abs/1907.07167>.
- Offline Dynamic Higher Connectivity.
with Bryce Sandlund, and Danny Sleator.
In WADS 2019.
Available at: <https://arxiv.org/abs/1708.03812>
- Flows in Almost Linear Time via Adaptive Preconditioning.
with Rasmus J. Kyng, Sushant Sachdeva, and Di Wang.
In STOC 2019.
Available at: <https://arxiv.org/abs/1906.10340>
- Fully Dynamic Vertex Spectral Sparsifiers and Applications.
with David Durfee, Yu Gao, and Gramoz Goranci.
In STOC 2019.
Available at: <https://arxiv.org/abs/1906.10530>
- Current Flow Group Closeness Centrality for Complex Networks.
with Huan Li, Liren Shan, Yuhao Yi, and Zhongzhi Zhang.

In WWW 2019.

Available at: <https://arxiv.org/abs/1802.02556>

- Iterative Refinement for ℓ_p -norm Regression.
with Deeksha Adil, Rasmus J. Kyng, and Sushant Sachdeva.
In SODA 2019.
Available at: <https://arxiv.org/abs/1901.06764>
- Solving Directed Laplacians in Nearly Linear Time through Sparse LU Factorizations.
with Michael B. Cohen, Jonathan Kelner, Rasmus Kyng, John Peebles, Anup B. Rao, and Aaron Sidford.
In FOCS 2018.
Available at <https://arxiv.org/abs/1811.10722>
- Graph Sketching Against Adaptive Adversaries Applied to the Minimum Degree Algorithm.
with Matthew Fahrbach, Gary L. Miller, Saurabh Sawlani, Junxing Wang, and Shen Chen Xu. In FOCS 2018.
Available at: <https://arxiv.org/abs/1804.04239>.
- Incomplete Nested Dissection.
with Rasmus J. Kyng, Robert Schwieterman, and Peng Zhang.
In STOC 2018.
Available at: <https://arxiv.org/abs/1805.09442>
- Parameterizing the Hardness of Binary Search Tree Access Sequences by Inversion Counts.
with Meng He and Yinzhan Xu.
In ANALCO 2018.
- Density Independent Algorithms for Sparsifying k-Step Random Walks. with Gorav Jindal, Pavel Kolev, and Saurabh Sawlani.
In APPROX 2017.
Available at: <https://arxiv.org/abs/1702.06110>.
- Almost-Linear-Time Algorithms for Markov Chains and New Spectral Primitives for Directed Graphs.
with Michael B. Cohen, Jonathan A. Kelner, John Peebles, Anup B. Rao. Aaron Sidford, and Adrian Vladu.
In STOC 2017. Invited to the special issue (declined).
Available at: <https://arxiv.org/abs/1611.00755>.
- A Framework for Analyzing Resparsification Algorithms.
with Rasmus J. Kyng, Jakub Pachocki, and Sushant Sachdeva.
In SODA 2017.
Available at: <https://arxiv.org/abs/1611.06940>.
- SPALS: Fast Alternating Least Squares via Implicit Leverage Scores Sampling.
with Dehua Cheng, Ioakeim Perros, and Yan Liu.
In NIPS 2016.
- An Empirical Study of Cycle Toggling Based Laplacian Solvers.
with Kevin Deweese, John R. Gilbert, Gary L. Miller, Hao Ran Xu, and Shen Chen Xu.
In CSC 2016.
Available at: <https://arxiv.org/abs/1609.02957>.

- Faster Algorithms for Computing the Stationary Distribution, Simulating Random Walks, and More.
with Michael B. Cohen, Jon Kelner, John Peebles, Aaron Sidford, and Adrian Vladu.
In FOCS 2016.
Available at <http://arxiv.org/abs/1608.03270>.
- On Fully Dynamic Graph Sparsifiers
with Ittai Abraham, David Durfee, Ioannis Koutis, and Sebastian Krinninger.
In FOCS 2016.
Available at: <http://arxiv.org/abs/1604.02094>.
- Simple and Scalable Constrained Clustering: A Generalized Spectral Method.
with Mihai Cucuringu, Ioannis Koutis, Sanjay Chawla, and Gary Miller.
In AISTATS 2016.
Available at http://www.math.ucla.edu/~mihai/consClust_AISTATS.pdf.
- Sparsified Cholesky and Multigrid Solvers for Connection Laplacians.
with Rasmus Kyng, Yin Tat Lee, Sushant Sachdeva, and Daniel A. Spielman.
In STOC 2016.
Available at: <http://arxiv.org/abs/1512.01892>.
- Faster and Simpler Width-Independent Parallel Algorithms for Positive Semidefinite Programming.
with Kanat Tangwongsan and Peng Zhang.
Updated 2016, Preliminary version in SPAA 2012.
Available at <http://arxiv.org/abs/1201.5135>.
- Approximate Undirected Maximum Flows in $O(\text{polylog}(n))$ Time.
In SODA 2016.
Available at <http://arxiv.org/abs/1411.7631>.
- Scalable Large Near-Clique Detection in Large-Scale Networks via Sampling
with Charalampos E. Tsourakakis, Michael Mitzenmacher, Jakub W. Pachocki, and Shen Chen Xu.
In KDD 2015.
Available at <http://www.cc.gatech.edu/~rpeng/MitzenmacherPPTX15.pdf>.
- Efficient Sampling for Gaussian Graphical Models via Spectral Sparsification.
with Dehua Cheng, Yu Cheng, Yan Liu, and Shang-Hua Teng.
In COLT 2015. Available at <http://arxiv.org/abs/1410.5392>.
- ℓ_p Row Sampling by Lewis Weights.
with Michael B. Cohen.
In STOC 2015.
Available at <http://arxiv.org/abs/1412.0588>.
- Improved Parallel Algorithms for Spanners and Hopsets.
with Gary L. Miller and Shen Chen Xu.
In SPAA 2015.
Available at <http://arxiv.org/abs/1309.3545>.
- Uniform Sampling for Matrix Approximation.
with Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford.
In ITCS 2015.
Available at <http://arxiv.org/abs/1408.5099>.

- Solving SDD Linear Systems in Nearly $m \log^{1/2} n$ Time.
with Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Anup B. Rao and Shen Chen Xu.
In STOC 2014.
This paper is a merger of the following two manuscripts on arXiv:
 - Preconditioning in Expectation.
with Michael B. Cohen, Rasmus Kyng, Jakub W. Pachocki, and Anup Rao.
Available at <http://arxiv.org/abs/1401.6236>
 - Stretching Stretch.
with Michael B. Cohen, Gary L. Miller, Jakub W. Pachocki, and Shen Chen Xu. Available at <http://arxiv.org/abs/1401.2454>.
- An Efficient Parallel Solver for SDD Linear Systems.
with Daniel A. Spielman.
In STOC 2014.
Available at <http://arxiv.org/abs/1311.3286>.
- Solving 1-Laplacians of Convex Simplicial Complexes in Nearly Linear Time: Collapsing and Expanding a Topological Ball.
with Michael B. Cohen, Brittany Terese Fasy, Gary L. Miller, Amir Nayyeri, and Noel Walkington.
In SODA 2014.
Available at <http://www.cc.gatech.edu/~rpeng/CohenFMNPW14.pdf>.
- Fully Dynamic $(1 + \epsilon)$ -Approximate Matchings.
with Manoj Gupta.
In FOCS 2013.
Available at <http://arxiv.org/abs/1304.0378>.
- Iterative Row Sampling.
with Mu Li and Gary L. Miller.
In FOCS 2013.
Available at <http://arxiv.org/abs/1211.2713>.
- Parallel Graph Decompositions Using Random Shifts.
with Gary L. Miller and Shen Chen Xu.
In SPAA 2013.
Available at <http://arxiv.org/abs/1307.3692>.
- Runtime Guarantees for Regression Problems.
with Hui Han Chin, Aleksander Madry, and Gary L. Miller.
In ITCS 2013.
Available at <http://arxiv.org/abs/1110.1358>.
- Approximate Maximum Flow on Separable Undirected Graphs.
with Gary L. Miller.
In SODA 2013.
Available at <http://arxiv.org/abs/1210.5227>.
- Faster Approximate Multicommodity Flow Using Quadratically Coupled Flows.
with Jonathan A. Kelner and Gary L. Miller.
In STOC 2012.
Available at <http://arxiv.org/abs/1202.3367>.

- A Nearly mlogn Time Solver for SDD Linear Systems.
with Ioannis Koutis and Gary L. Miller.
In FOCS 2011.
Available at <http://arxiv.org/abs/1102.4842>.
- Linear-Work Greedy Parallel Approximate Set Cover and Variants. with Guy E. Blelloch and Kanat Tangwongsan.
In SPAA 2011.
Available at <http://www.cc.gatech.edu/~rpeng/BlellochPT11.pdf>.

B Other Refereed Material

- Numerical Difficulties of Combinatorial Preconditioning.
with Kevin Dewese, Serban Stan, and Haoran Xu.
Preliminary version presented at CSC20.
- A Study of Performance of Optimal Transport.
with Yihe Dong, Yu Gao, Ilya Razenshteyn, and Saurabh Sawlani.
Preliminary version presented at CSC20.
Available at: <https://arxiv.org/abs/2005.01182>.

D Presentations

D.1 Invited Speaker at Workshops

- November 2022, “Maximum Flow and Minimum-Cost Flows in Almost-Linear Time”, Alumni Talk at the ICPC World Finals Dhaka.
- August 2022, “Maximum Flow and Minimum-Cost Flow in Almost-Linear Time”, Chinese Computing Federation Advanced Disciplines Lectures 128 (virtual).
- June 2022, “Data Structures in Optimization”, Dynamic Algorithms: Recent Advances and Applications at STOC 2022.
- May 2022, “A Survey of Fast Flow Algorithms”, Highlights of Algorithms 2022.
- August 2021, “Fully Dynamic Effective Resistance”, Workshop on Algorithms for Large Data (online).
- July 2020, “High Performance Linear System Solvers with Focus on Graph Laplacians”, Minisymposium “Linear Algebraic Tools for Graph Computation” at SIAM MDS2020 (virtual).
- September 2019, “Graph Algorithms and Batched Processing”, DIMACS workshop on Randomized Numerical Linear Algebra, Statistics, and Optimization.
- July 2019, “Almost-linear time algorithms for Markov chains and new spectral primitives for directed graphs”, Approximation Algorithms for Combinatorial Scientific Computing at ICIAM 2019.
- May 2018, “Batch-Dynamic Graph Algorithms”, Fine Grained Approximation Algorithms and Complexity (FG-APX 2019).
- Oct 2018, “Merging the Continuous and Discrete”, Workshop on Laplacian Paradigm 2.0 at FOCS 2018 .
- Sep 2018, “Scalable Algorithmic Primitives for Data Science”, Simons Workshop on Randomized Numerical Linear Algebra and Applications.

- June 2017, “Determinant Preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees”, Theory @ Nanjing 2017.
- Mar 2017, “High Performance Solvers for Linear Systems in Graph Laplacians”, NSF Algorithms in the Field PI meeting.
- Feb 2017, “High Performance Solvers for Linear Systems in Graph Laplacians”, SIAM Conference on Computer Science and Engineering (CSE) 2017.
- Sep 2016, “Parallel Graph Algorithms”, 5th Workshop on Advances in Distributed Graph Algorithms.
- Aug 2016, “Algorithm Frameworks Based on Adaptive Sampling”, Banff International Research Station Workshop on Algebraic and Spectral Graph Theory.
- July 2016, “ ℓ_p Row Sampling by Lewis Weights”, NII Shonan Meeting on Recent Advances in Randomized Numerical Linear Algebra.
- July 2016, “Algorithm Frameworks Based on Adaptive Sampling”, IAS/Park City Mathematics Institute Summer Session 2016.
- Jan 2016, “Algorithm Frameworks Based on Structure Preserving Sampling”, UC San Diego Workshop on Big Graphs.
- July 2015, “Approximate Undirected Maximum Flows in $O(m\text{polylog}(n))$ Time”, 22nd International Symposium on Mathematical Programming 2015.

D.2 Invited Talks at Seminars

- December 2022, “Maximum Flow and Minimum-Cost Flow in Almost-Linear Time”, Harvard Theory of Computing Seminar.
- August 2022, “Maximum Flow and Minimum-Cost Flow in Almost-Linear Time”, Shanghai Center for Mathematical Sciences Combinatorics Seminar (virtual).
- December 2021, “Faster Network Flow Algorithms”, Chinese Optimization Society seminar (virtual).
- October 2021, “Solving Sparse Linear Systems Faster than n^ω ”, Georgia State University.
- July 2021, “Solving Sparse Linear Systems Faster than Matrix Multiplication”, AriC (Arithmetic and Computing) seminar, Ecole Normale Supérieure de Lyon.
- April 2021, “Solving Sparse Linear Systems Faster than Matrix Multiplication”, MIT Theory Seminar.
- January 2021, “Solving Sparse Linear Systems Faster than Matrix Multiplication”, UT Austin CS Theory Seminar.
- October 2020, “The Numerics of Solving Sparse Linear Systems Faster than Matrix Multiplication”, ETH Zurich Algorithms and Complexity Seminar.
- July 2020, “A Study of Performance of Optimal Transport”, ITCS seminar at SHUFE (virtual).
- October 2019, “Graph Algorithms and Batched Processing”, University of Toronto Theory Seminar.
- April 2019, “Fully Dynamic Spectral Vertex Sparsifiers and Applications”, TCS+.

- August 2018, “Faster Computations of Effective Resistances”, Machine Learning and Optimization Seminar at MSR Redmond.
- May 2018, “Fully Dynamic Effective Resistances”, Theory Lunch at MSR Redmond.
- June 2017, “Determinant Preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees”, ITCS seminar at SHUFE.
- Nov 2016, “Almost-Linear-Time Algorithms for Markov Chains and New Spectral Primitives for Directed Graphs”, Duke University Algorithms Seminar.
- Oct 2016, “Directed Spectral Sparsification and Laplacian Solvers in Almost Linear Time”, UT Austin CS Theory Seminar.
- June 2016, “Sparsified Matrix Algorithms for Graph Laplacians”, ITCS seminar at SHUFE.
- Mar 2016, “Sparsified Matrix Algorithms for Graph Laplacians”, UC Irvine Applied & Computational Mathematics Seminar .
- Oct 2015, “Algorithm Frameworks Based on Structure Preserving Sampling”, UC Berkeley AMPLab Seminar.

E Grants and Contracts

E.1 As Principle Investigator

- Title of Project: High Performance Graph Algorithms and Data Structures
Agency: NSERC
Total Amount: 300,000 CAD
Role: PI
Collaborators: none
Period of Contract: 4/1/2022 - 1/31/2024
Candidate’s Share: 60,300 CAD (due to shortened duration).
- Title of Project: CAREER: Scalable Algorithmic Primitives for Data Science
Agency: NSF
Total Amount: 456,546 USD
Role: PI
Collaborators: none
Period of Contract: 7/1/2019 - 8/31/2024
Candidate’s Share: 100% (456,546 USD)
- Title of Project: AF: Small: New Algorithmic Primitives for Directed Graphs: Sparsification and Preconditioning
Agency: NSF
Total Amount: 450,000 USD
Role: PI
Collaborators: none
Period of Contract: 7/1/2017 - 6/30/2020
Candidate’s Share: 100% (450,000 USD)
- Title of Project: AitF: Collaborative Research: High Performance Linear System Solvers with Focus on Graph Laplacians

Agency: NSF
Total Amount: 800,000 USD
Role: PI
Collaborators: John Gilbert (co-PI, UCSB), Gary L. Miller (co-PI, CMU)
Period of Contract: 9/1/2016 - 8/31/2020
Candidate's Share: 33% (266,666 USD)

F Other Scholarly and Creative Accomplishments

- Co-organizer of Banff International Research Station workshop ‘Perspectives on Matrix Computations: Theoretical Computer Science Meets Numerical Analysis’ in March 2023. Website at <https://www.birs.ca/events/2023/5-day-workshops/23w5108>.
- Co-organizer of workshop ‘Laplacian Paradigm 2.0’ at FOCS 2019 in October 2019. Website at <https://sachdevasushant.github.io/laplacian2.0/>.
- Co-organizer of Dagstuhl Seminar 18241 ‘High-Performance Graph Algorithms’ in June 2018. Website at <https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=18241>.
- Organized session ‘High Performance Spectral Algorithms’ at the 2017 SIAM Annual Meeting. Program at http://meetings.siam.org/session/dsp_programsess.cfm?SESSIONCODE=62686.

V Education

A Courses Taught

- Winter 2024, 2025: 15-151 Mathematical Foundations for Computer Science, CMU, 50 students.
- Fall 2023, 2024: 15-495/795 Topics in Algorithmic Problem solving, CMU, 5 students.
- Fall 2024: 15-359/659 Probability & Computing: Randomized Algorithms and Markov Chains, CMU, 15 students.
- Winter 2023, 15-210 Parallel and Sequential Data Structures and Algorithms, CMU, co-taught with Guy Blelloch, 210 students.
- Winter 2022, CS 341 Algorithms, University of Waterloo, co-taught with Arne Storjohann, 2 out of 4 sections of 320 students total.
- Summer 2021, CS 234 Data Types and Structures, University of Waterloo, 190 students.
- Spring 2021, CS 6550 Design and Analysis of Algorithms, Georgia Tech, 35 students.
- Fall 2020, CS 4540 Advanced Algorithms, Georgia Tech, 81 students.
- Spring 2020, CS 4510 Automata and Complexity, Georgia Tech, 290 students.
- Fall 2019, CS 7510, Graph Algorithms, Georgia Tech, 31 students.
- Fall 2018, CS 7540, Spectral Algorithms, Georgia Tech, co-taught with Di Wang, 25 students.
- Fall 2018, CS 4510 Automata and Complexity, Georgia Tech, 181 students.
- Fall 2017, CS 3510 Design and Analysis of Algorithms, Georgia Tech, 312 students.

- Spring 2017, CS 7540, Spectral Algorithms, Georgia Tech, 21 students.
- Fall 2016, CS 3510 Design and Analysis of Algorithms, Georgia Tech, 82 students.
- Fall 2015, CS 8803-SA Sampling Algorithms, Georgia Tech, 21 students.
- Fall 2014, 6.046 Algorithm Design, M.I.T., co-taught with Costis Daskalakis and Dana Moshkovitz, 210 students.

B Individual Student Guidance

B.1 Ph.D. Students

- Alicia Stepin, Fall 2024 - present at CMU.
- Junzhao Yang, Fall 2024 - present at CMU.
- Hoai-An Nguyen, co-advised with David Woodruff, Fall 2023 - present at CMU. <https://hoaiannguyen.com/>
- Jingbang Chen, Winter 2023 - present at UWaterloo. <https://chenjb1997.github.io/>
- Li Chen, Fall 2019 - Spring 2023 at Georgia Tech, <https://lic225.github.io/>.
- Yu Gao, Fall 2017 - Summer 2022 at Georgia Tech, <https://sites.google.com/view/ygao2606/home>. Huawei Summer 2022 - present.
- Saurabh Sawlani, Fall 2016 - Spring 2020 at Georgia Tech, <https://www.cc.gatech.edu/~ssawlani3/> Postdoc at Carnegie Mellon University, Fall 2020 - Summer 2021. SoundHound, Summer 2021 - present.
- David Durfee, Fall 2015 - Fall 2018 at Georgia Tech, <https://www.cc.gatech.edu/~ddurfee3/>. LinkedIn, Winter 2019 - present.
- Peng Zhang, Fall 2015 - Summer 2018 at Georgia Tech, <https://sites.google.com/site/pengzhang27182/>. Postdoc at Yale University, Fall 2018 - Summer 2021. Rutgers Univeristy, Fall 2021 - present.

B.2 M.S. Students

- Jingbang Chen, Fall 2020 - Fall 2022 at Georgia Tech.
- Qian Lyu, Fall 2017 at Georgia Tech.
- Robert Schwieterman, Fall 2017 - Spring 2018 at Georgia Tech.

B.3 Postdocs

- Li Chen, September 2023 - February 2024 at CMU,
- Di Wang, January 2018 - May 2019 at Georgia Tech.

B.4 Service on thesis or dissertation committees

- Praneeth Kacham
- Taisuke Yasuda
- Karthik Murali
- Yu Gao
- Matthew Farhbach
- Shijie Xie
- Junxing Wang
- David Durfee
- Tung Mai
- Saurabh Sawlani
- Chi Ho Yuen
- Yan Wang
- Peng Zhang

C Education Innovations and Other Contributions

With programming competitions, which are algorithmic problem solving based outreach programs.

- Scientific Committee of the 2024 ATHKA programming contest in Saudi Arabia.
- Asian-Pacific Informatics Olympiad:
 - Scientific Committee (chair), 2022.
 - Technical Committee, 2024.
- Invited speaker at the Chinese National Olympiad in Informatics Winter Camp, 2021, 2022.
- Lead trainer for the North America Programming Camp (NAPC), 2019 - 2022.
- Coach for USA Computing Olympiad (USACO), 2006 - 2017.
- Problem setter for International Olympiad in Informatics (IOI):
 - Host Scientific Committee: 2008, 2009, 2010;
 - external task submitter: 2011, 2013;
 - International Scientific Committee (elected): 2014 - 2018 (chair in 2018).
- Coach of teams competing in the International Collegiate Programming Contest (ICPC):
 - CMU 2010 - 2013 (co-coach), 2021 - 2024 (co-coach), 2024 - present (coach).
 - Georgia Tech 2019 - 2021 (coach).
 - M.I.T. 2013 - 2015 (coach).
 - University of Waterloo 2008 - 2009 (co-coach).

VI Service

A Professional Contributions

A.1 Conference Program Committees

FOCS 2025, ALENEX 2025, ICALP 2024, SODA 2024, FOCS 2023, ACDA 2023, SODA 2023, ICALP 2022, SOSA 2022, ESA 2021, ACDA 2021, ESA 2021, ITCS 2021, ISAAC 2020, FAW 2020, WADS 2019, FOCS 2019, FAW 2019, NCTCS 2018, ESA 2018, SODA 2018, RANDOM 2017, WADS 2017, SPAA 2017, APPROX 2016, FOCS 2015.

A.2 Conference Reviewer

CCC 2019, 2017, 2012; COLT 2018, 2016; ESA 2019, 2018, 2016, 2015, 2013; FOCS 2024,2023,2022, 2021, 2020, 2018, 2017, 2016, 2015, 2014, 2013; ICALP 2018, 2017, 2016; NeurIPS 2019, 2018, 2016; PODC 2019, 2017, 2016; SOCG 2018, 2015; SODA 2025,2022, 2021, 2020, 2019, 2017, 2016, 2015, 2013; SPAA 2014; STOC 2024,2023, 2022, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014; SWAT 2022.

A.3 Journal Reviewer

ALGO, JACM, RSA, SICOMP, SIDMA, SIMAX, SIOPT, SISC, TCS, TKDD, TOC.