Benjamin J. Moseley, Associate Professor of Operations Research

CONTACT INFORMATION

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

http://www.andrew.cmu.edu/user/moseleyb/
moseleyb@andrew.cmu.edu

RESEARCH INTERESTS

Algorithms, operations research, machine learning, optimization, parallel and distributed algorithms, scheduling, network algorithms, and large data analysis

EDUCATION

University of Illinois: Urbana-Champaign: 2012 — Ph.D. in Computer Science
- Thesis: Online Scheduling Algorithms for Broadcasting and General Cost Functions
- Advisor: Chandra Chekuri

University of Illinois: Urbana-Champaign: 2007-2008 — Master of Science in Computer Science
- Thesis: Online Scheduling to Minimize the Maximum Delay Factor
- Advisor: Chandra Chekuri

University of Illinois: Urbana-Champaign: 2004-2006 — Bachelor of Science in Computer Science
Elgin Community College: 2002-2004

HONOURS AND AWARDS

NeurIPS Spotlight Paper 2023
Poets and Quants top 50 Undergraduate Professors 2021
NeurIPS Oral Presentation 2021(top 1% of submissions)
Google Faculty Research Award 2019
NSF CAREER Award 2018
Infor Faculty Award 2018
NeurIPS Spotlight Presentation 2018 (top 3.5% of submissions)
Carnegie Bosch Junior Faculty Chair 2018-2024
NIPS Oral Presentation 2017 (top 1.3% of submissions)
Simons-Berkeley Fellow 2016
Yahoo! Academic Career Enhancement (ACE) Award 2015
Google Faculty Research Award 2015
Best Paper Award IPDPS 2015
Best Paper Award SPAA 2013
Feng Chen Memorial Award 2011
KDD Oral Presentation 2011 (top 8% of submissions)
Best Student Paper Award SODA 2010
Outstanding Teaching Award Spring 2008
List of Teachers Ranked as Excellent Spring 2008

ACADEMIC FACULTY POSITIONS

Carnegie Mellon University, Pittsburgh, PA
Assoc. Prof. (with tenure) of Operations Research and Machine Learning — **2024– current**
Carnegie Bosch Assoc. Prof. (with tenure) of Operations Research and Machine Learning — **2022– 2024**
Carnegie Bosch Assoc. Prof. (without tenure) of Operations Research and Machine Learning — **2020– 2022**
Carnegie Bosch Assistant Professor of Operations Research and Machine Learning — **2018– 2020**

Washington University in St. Louis, St. Louis, MO
Assistant Professor of Computer Science and Engineering — **July 2014– December 2017**

Toyota Technological Institute at Chicago, Chicago, IL
Research Assistant Professor — **September 2012– July 2014**

UNIVERSITY LEADERSHIP POSITIONS

Co-Chair Tepper Strategic Planning **2023-2024**

Co-Chair Operations Research Faculty Recruiting **2021, 2022**

Ph.D Head, OR and Algorithms Combinatorics and Optimization (ACO) **2020- 2023**

CONSULTING AND VISITING POSITIONS

Relational AI, Berkeley, CA
Consulting Senior Scientist **July 2018– Current**

Simons-Berkeley , Berkeley, CA
Fellow for the program on Algorithms and Uncertainty **Fall 2016**

Yahoo! Research, New York, NY
Visiting Senior Research Scientist **May 2016 – June 2016**
Host: Maxim Sviridenko

Sandia National Laboratories, Albuquerque, NM
Faculty Summer Sabbatical **June 2013 – August 2013**
Host: Cindy Phillips

Yahoo! Labs, Santa Clara, CA
Research Scientist Intern **May 2010 – August 2010 and May 2011 – August 2011**
Host: Ravi Kumar

Great Lakes IT inc, Chicago, IL
Senior Programmer **December 2004 – August 2012**

CONFERENCE PAPERS

1. Benjamin Moseley, Aidin Niaparast, and R. Ravi
Putting Off the Catching Up: Online Joint Replenishment Problem with Holding and Backlog Costs
In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2025).
2. Sungjin Im, Benjamin Moseley, Hung Ngo and Kirk Pruhs
Polynomial Time Convergence of the Iterative Evaluation of Datalogo Programs
In Symposium on Principles of Database Systems (PODS 2025)
3. Benjamin Moseley, Heather Newman and Kirk Pruhs
Online k-Median with Consistent Clusters
In Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX 2024)
4. Samuel McCauley, Benjamin Moseley, Aidin Niaparast, and Shikha Singh
Incremental Topological Ordering and Cycle Detection with Predictions
In Proceedings of the 41st International Conference on Machine Learning (ICML 2024).
5. Sami Davies, Benjamin Moseley and Heather Newman
Simultaneously Approximating All lp-norms in Correlation Clustering
In the International Colloquium on Automata, Languages, and Programming (ICALP 2024).
6. Kunal Agrawal, Benjamin Moseley, Heather Newman, and Kirk Pruhs
Scheduling Out-Trees Online to Optimize Maximum Flow
In Proceedings of the 36th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2024).
7. Qingyun Chen, Sungjin Im, Chenyang Xu, Benjamin Moseley, and Ruilong Zhang
Sampling for Beyond-Worst-Case Online Ranking
In AAAI Conference on Artificial Intelligence (AAAI 2024).
8. Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, Sergei Vassilvitskii
Controlling Tail Risk in Online Ski-Rental
In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2024).

9. Benjamin Moseley, Heather Newman and Kirk Pruhs
The Public University Secretary Problem
In SIAM Symposium on Simplicity in Algorithms (SOSA 2024).

10. Sungjin Im, Benjamin Moseley, Hung Ngo and Kirk Pruhs
On the Convergence Rate of Linear Datalogo over Stable Semirings
In International Conference on Database Theory (ICDT 2024).

11. Samuel McCauley, Benjamin Moseley, Aidin Niaparast, and Shikha Singh
Online List Labeling with Predictions
In Advances in Neural Information Processing Systems (NeurIPS 2023).
Spotlight Paper

12. Sami Davies, Benjamin Moseley, and Heather Newman
Fast Combinatorial Algorithms for Min Max Correlation Clustering
In Proceedings of the 40th International Conference on Machine Learning (ICML 2023).

13. Sami Davies, Benjamin Moseley, Sergei Vassilvitskii, and Yuyan Wang
Predictive Flows for Faster Ford-Fulkerson
In Proceedings of the 40th International Conference on Machine Learning (ICML 2023).

14. Sungjin Im, Chenyang Xu, Benjamin Moseley, and Ruilong Zhang
Online State Exploration: Competitive Worst Case and Learning-Augmented Algorithms
In the European Conference on Machine Learning (ECML 2023)

15. Franziska Eberle, Anupam Gupta, Nicole Megow, Benjamin Moseley and Rudy Zhou
Configuration Balancing for Stochastic Requests
In Proceedings of the Integer Programming and Combinatorial Optimization (IPCO 2023).

16. Sungjin Im, Benjamin Moseley, Chenyang Xu and Ruilong Zhang
Online Dynamic Acknowledgement with Learned Predictions
In Proceedings of the IEEE International Conference on Computer Communications (INFOCOM 2023)

17. Qingyun Chen, Sungjin Im, Chenyang Xu, Benjamin Moseley, and Ruilong Zhang
Min-Max Submodular Ranking for Multiple Agents
In AAAI Conference on Artificial Intelligence (AAAI 2023).

18. Anupam Gupta, Benjamin Moseley, and Rudy Zhou
Minimizing Completion Times for Stochastic Jobs via Batched Free Times
In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2023).

19. Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii
Algorithms with Prediction Portfolios
In Advances in Neural Information Processing Systems (NeurIPS 2022).

20. Benjamin Moseley, Kirk Pruhs, Clifford Stein, Rudy Zhou
A Competitive Algorithm for Throughput Maximization on Identical Machines
In Proceedings of the Integer Programming and Combinatorial Optimization (IPCO 2022).

21. Chenyang Xu and Benjamin Moseley
Learning-Augmented Algorithms for Online Steiner Tree
In AAAI Conference on Artificial Intelligence (AAAI 2022).

22. Kunal Agrawal, Michael Bender, Jonathan Berry, Rathish Das , Daniel DeLayo, Benjamin Moseley, Cynthia Phillips and Kenny Zhang
Automatic HBM Management: Models and Algorithms
In Proceedings of the 34th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2022)

23. Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii
Faster Matchings via Learned Duals
In Advances in Neural Information Processing Systems (NeurIPS 2021).
Oral Presentation (top 1% of submissions)

24. Silvio Lattanzi, Benjamin Moseley, Sergei Vassilvitskii, Yuyan Wang, and Rudy Zhou
Robust Online Correlation Clustering
In Advances in Neural Information Processing Systems (NeurIPS 2021).

25. Benjamin Berg, Mor Harchol-Balter, Benjamin Moseley, Weina Wang, Justin Whitehouse:
The Case for Phase-Aware Scheduling of Parallelizable Jobs
In Proceedings of the International Symposium on Computer Performance, Modeling, Measurements and Evaluation (Performance 2021).

26. Kefu Lu, Thomas Lavastida, Benjamin Moseley and Yuyan Wang
Scaling Average-Linkage via Sparse Cluster Embeddings
In Asian Conference on Machine Learning. (ACML 2021)

27. Mahmoud Abo Khamis, Ryan Curtin, Sungjin Im, Benjamin Moseley, Hung Ngo, Kirk Pruhs and Alireza Samadian
An Approximation Algorithm for the Matrix Tree Multiplication Problem
In Proceedings of the 246th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021).

28. Marilena Leichter, Benjamin Moseley and Kirk Pruhs
An Efficient Reduction of a Gammoid to a Partition Matroid
In Proceedings of the 29th Annual European Symposium on Algorithms (ESA 2021).

29. Thomas Lavastida, Benjamin Moseley, R. Ravi and Chenyang Xu
Learnable and Instance-Robust Predictions for Online Matching, Flows and Load Balancing
In Proceedings of the 29th Annual European Symposium on Algorithms (ESA 2021).

30. Anupam Gupta, Benjamin Moseley and Rudy Zhou
Structural Iterative Rounding for Generalized k-Median Problems
In the 49th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

31. Benjamin Moseley, Kirk Pruhs, Alireza Samadian and Yuyan Wang
Relational Algorithms for k-means Clustering
In the 49th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

32. Thomas Lavastida, Benjamin Moseley, R. Ravi and Chenyang Xu
Using Predicted Weights for Ad Delivery
In the 1st SIAM Conference on Applied and Computational Discrete Algorithms (ACDA 2021).

33. Jeremy Buhler, Thomas Lavastida, Kefu Lu, and Benjamin Moseley
A Scalable Approximation Algorithm for Weighted Longest Common Subsequence
In Proceedings of the International European Conference on Parallel and Distributed Computing (Euro-Par 2021).

34. Benjamin Moseley, Sergei Vassilvitskii, and Yuyan Wang
Hierarchical Clustering in General Metric Spaces using Approximate Nearest Neighbors
In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2021).

35. Mahmoud Abo-Khamis, Sungjin Im, Benjamin Moseley, Kirk Pruhs, Alireza Samadian
Instance Optimal Join Size Estimation
In Proceedings of the Latin and American Algorithms, Graphs and Optimization Symposium (LAGOS 2021).

36. Mahmoud Abo-Khamis, Sungjin Im, Benjamin Moseley, Kirk Pruhs, Alireza Samadian
Approximate Aggregate Queries Under Additive Inequalities
SIAM-ACM Symposium on Algorithmic Principles of Computer Systems (APoCS 2021).

37. Mahmoud Abo-Khamis, Sungjin Im, Benjamin Moseley, Kirk Pruhs, Alireza Samadian
A Relational Gradient Descent Algorithm For Support Vector Machine Training
SIAM-ACM Symposium on Algorithmic Principles of Computer Systems (APoCS 2021).

38. Sara Ahmadian, Alessandro Epasto, Marina Knittel, Ravi Kumar, Mohammad Mahdian, Benjamin Moseley, Philip Pham, Sergei Vassilvitskii and Yuyan Wang

Fair Hierarchical Clustering
In Advances in Neural Information Processing Systems (NeurIPS 2020).

39. Rathish Das, Kunal Agrawal, Michael Bender, Jonathan Berry, Benjamin Moseley and Cynthia Phillips
How to Manage High-Bandwidth Memory Automatically
In Proceedings of the 32nd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2020)

40. Benjamin Berg, Mor Harchol-Balter, Benjamin Moseley, Weina Wang, Justin Whitehouse:
Optimal Resource Allocation for Elastic and Inelastic Jobs
In Proceedings of the 32nd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2020)

41. Sungjin Im, Mashid Qaem, Benjamin Moseley, Xiaorui Sun, and Rudy Zhou
Fast Noise Removal for k-means Clustering
In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

42. Alireza Samadian, Kirk Pruhs, Benjamin Moseley, Sungjin Im, and Ryan Curtain
Unconditional Coresets for Regularized Loss Minimization
In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

43. Ryan Curtain, Benjamin Moseley, Hung Ngo, XuanLong Nguyen, Dan Olteanu, Maximillian Schleich
Rk-means: Fast Clustering for Relational Data
In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

44. Sungjin Im, Benjamin Moseley, Kamesh Munagala and Kirk Pruhs
Dynamic Weighted Fairness with Minimal Disruptions
In Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS and SIGMETRICS 2020)

45. Michael Dinitz and Benjamin Moseley
Scheduling for Weighted Flow and Completion Times in Reconfigurable Networks
In Proceedings of the IEEE International Conference on Computer Communications (INFOCOM 2020).

46. Benjamin Moseley and Yuyan Wang
An Objective for Hierarchical Clustering in Euclidean Space and its Connection to Bisecting K-means
In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2020).

47. Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii
Online Scheduling via Learned Weights
In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2020).

48. Shikha Singh, Sergey Madaminov, Michael Bender, Michael Ferdman, Ryan Johnson, Benjamin Moseley, Hung Ngo, Dung Nguyen, Soeren Olesen, Kurt Stirewalt, and Geoffrey Washburn.
A Scheduling Approach to Incremental Maintenance of Datalog Programs.
In Proceedings of the 34th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2020).

49. Giorgio Lucarelli, Benjamin Moseley, Nguyen Thang, Abhinav Srivastav and Denis Trystram
Online Non-preemptive Scheduling to Minimize Maximum Weighted Flow-time on Related Machines
In Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2019).

50. Ayan Chakrabarti and Benjamin Moseley
Backprop with Approximate Activations for Memory-efficient Network Training
In Advances in Neural Information Processing Systems, 2019 (NeurIPS 2019).

51. Shali Jiang, Benjamin Moseley, and Roman Garnett Cost Effective Active Search
In Advances in Neural Information Processing Systems, 2019 (NeurIPS 2019).

52. Benjamin Moseley and Maxim Sviridenko Submodular Optimization with Contention Resolution Extensions
In Proceedings of the 23rd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2019)

53. Benjamin Moseley
Scheduling to Approximate Minimization Objectives on Identical Machines
In the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).

54. Sungjin Im, Benjamin Moseley, Kirk Pruhs and Manish Purohit
Matroid Coflow Scheduling
In the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).

55. Silvio Lattanzi, Thomas Lavastida, Kefu Lu, and Benjamin Moseley
A Framework for Parallelizing Hierarchical Clustering Methods
In the European Conference on Machine Learning (ECML 2019) .

56. Mahmoud Abo Khamis, Ryan Curtin, Benjamin Moseley, Hung Ngo, Long Nguyen, Dan Olteanu and Maximilian Schleich
On Functional Aggregate Queries with Additive Inequalities
In Proceedings of the 38th ACM Symposium on Principles of Database Systems (PODS 2019).

57. Kunal Agrawal, I-Ting Angelina Lee, Jing Li, Kefu Lu and Benjamin Moseley
Practically Efficient Scheduler for Minimizing Average Flow Time of Parallel Jobs
In Proceedings of the 33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2019).

58. Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley and Roman Garnett
Efficient Nonmyopic Batch Active Search
In Advances in Neural Information Processing Systems, 2018 (NeurIPS 2018).
Spotlight Presentation (top 168 out of 4856 submissions)

59. Giorgio Lucarelli, Benjamin Moseley, Nguyen Kim Thang, Abhinav Srivastav and Denis Trystram
Online Non-Preemptive Scheduling to Minimize Weighted Flow-time on Unrelated Machines
In Proceedings of the 26th Annual European Symposium on Algorithms (ESA 2018).

60. Giorgio Lucarelli, Benjamin Moseley, Nguyen Kim Thang, Abhinav Srivastav and Denis Trystram
Online Non-preemptive Scheduling on Unrelated Machines with Rejections
In Proceedings of the 30th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2018).

61. Kunal Agrawal, Jing Li, Kefu Lu and Benjamin Moseley
Scheduling Parallelizable Jobs Online to Maximize Throughput
In Proceedings of the 15th Latin American Theoretical Informatics Symposium (LATIN 2018).

62. Benjamin Moseley and Joshua Wang
Approximation Bounds for Hierarchical Clustering: Average-Linkage, Bisecting K-means, and Local Search
In Advances in Neural Information Processing Systems, 2017 (NIPS 2017).
Oral Presentation (top 40 out of 3240 submissions)

63. Sungjin Im, Benjamin Moseley, Kirk Pruhs and Clifford Stein
An $O(\log \log m)$-competitive Algorithm for Online Machine Minimization
In Proceedings of the Real-Time Systems Symposium, 2017 (RTSS 2017).

64. Sungjin Im, Benjamin Moseley, Kirk Pruhs and Clifford Stein
Minimizing Maximum Flow Time on Related Machines via Dynamic Posted Pricing
In Proceedings of the 25th Annual European Symposium on Algorithms (ESA 2017).

65. Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, Roman Garnett
Efficient Nonmyopic Active Search
In Proceedings of the 34th International Conference on Machine Learning (ICML 2017).

66. Kunal Agrawal, Jing Li, Kefu Lu and Benjamin Moseley
Scheduling Parallelizable Jobs Online to Maximize Throughput
In Proceedings of the 29th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2017). Brief Announcement.

67. Sungjin Im, Benjamin Moseley, and Xiaorui Sun
Efficient Massively Parallel Methods for Dynamic Programming
In Proceedings of the Symposium on Theory of Computing (STOC 2017)

68. Shalmoli Gupta, Ravi Kumar, Kefu Lu, Benjamin Moseley, and Sergei Vassilvitskii
Local Search Methods for k-Means with Outliers
In Proceedings of the International Conference on Very Large Data Bases (VLDB 2017)

69. Gustavo Malkomes, Kefu Lu, Blakeley Hoffman, Roman Garnett, Benjamin Moseley and Richard Mann
Cooperative Set Function Optimization Without Communication or Coordination
In Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)

70. Varun Gupta, Benjamin Moseley, Marc Uetz and Qiaomin Xie
Stochastic Online Scheduling on Unrelated Machines
In Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO 2017).

71. Sungjin Im, Benjamin Moseley and Shi Li
Breaking $1 - 1/e$ Barrier for Non-preemptive Throughput Maximization
In Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO 2017).

72. Sungjin Im, and Benjamin Moseley
Fair Scheduling via Iterative Quasi-Uniform Sampling
In Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017).

73. Sungjin Im, Janardhan Kulkarni, Benjamin Moseley and Kamesh Munagala
A Competitive Flow Time Algorithm for Heterogeneous Clusters under Polytope Constraints
In Proceedings of the 19th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2016).

74. Sungjin Im and Benjamin Moseley
General Profit Scheduling and the Power of Migration on Heterogeneous Machines
In Proceedings of the 28th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2016).

75. Kunal Agrawal, Jing Li, Kefu Lu and Benjamin Moseley
Scheduling Parallelizable Jobs Online to Minimize Maximum Flow Time
In Proceedings of the 28th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2016).

76. Shaurya Ahuja, Kefu Lu, and Benjamin Moseley
Partitioned Feasibility Tests for Sporadic Tasks on Heterogeneous Machines
In Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2016).

77. Kunal Agrawal, Jing Li, Kefu Lu, and Benjamin Moseley
Scheduling Parallel DAG Jobs Online to Minimize Average Flow Time
In Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016).

78. Gustavo Malkomes, Matt Kusner, Wenlin Chen, Kilian Weinberger and Benjamin Moseley
Fast Distributed k-Center Clustering with Outliers on Massive Data
In Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS 2015)

79. Roozbeh Ebrahimi, Samuel McCauley and Benjamin Moseley
Scheduling Parallel Jobs Online with Convex and Concave Parallelizability
In Proceedings of the 13th Workshop on Approximation and Online Algorithms (WAOA 2015)

80. Michael A. Bender, Jonathan Berry, Simon D. Hammond, Branden Moore, Benjamin Moseley and Cynthia A. Phillips
k-Means Clustering on Two-Level Memory Systems
In The International Symposium on Memory Systems (MEMSYS 2015).

81. Sungjin Im and Benjamin Moseley
Weighted Reordering Buffer Improved via Variants of Knapsack Covering Inequalities
In The 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015).

82. Noa Avigdor-Elgrabli, Sungjin Im, Benjamin Moseley and Yuval Rabani
On the Randomized Competitive Ratio of Reordering Buffer Management with Non-Uniform Costs
In The 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015).

83. Sungjin Im, Janardhan Kulkarni and Benjamin Moseley
Temporal Fairness of Round Robin: Competitive Analysis for Lk-norms of Flow Time
In Proceedings of the 27th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2015).

84. Sungjin Im and Benjamin Moseley
Scheduling in Bandwidth Constrained Tree Networks
In Proceedings of the 27th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2015).

85. Sungjin Im and Benjamin Moseley
Fast and Better Distributed MapReduce Algorithms for k-Center Clustering
In Proceedings of the 27th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2015). Brief Announcement.

86. Michael A. Bender, Jonathan W Berry, Simon Hammond, Karl Hemmert, Samuel McCauley, Branden Moore, Benjamin Moseley, Cynthia A Phillips, David Resnick, and Arun Rodrigues
Two-Level Main Memory Co-Design: Multi-Threaded Algorithmic Primitives, Analysis, and Simulation
In Proceedings of the 29th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2015).
Awarded Best Paper

87. Sungjin Im, Benjamin Moseley and Kirk Pruhs
Stochastic Scheduling of Heavy-tailed Jobs
In Proceedings of the 32nd Symposium on Theoretical Aspects of Computer Science (STACS 2015).

88. Sungjin Im, Shi Li, Benjamin Moseley, and Eric Torng
A Dynamic Programming Framework for Non-Preemptive Scheduling Problems on Multiple Machines
In Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015).

89. Sungjin Im, Benjamin Moseley, Kirk Pruhs and Eric Torng
Competitively Scheduling Tasks with Intermediate Parallelizability
In Proceedings of the 26th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2014).

90. Antonios Antoniadis, Neal Barcelo, Daniel Cole, Kyle Fox, Benjamin Moseley, Michael Nugent and Kirk Pruhs
Packet Forwarding Algorithms in a Line Network
In Proceedings of the 11th Latin American Theoretical Informatics Symposium (LATIN 2014).

91. Antonios Antoniadis, Sungjin Im, Ravishankar Krishnaswamy, Vishwanath Nagarajan, Benjamin Moseley, Kirk Pruhs and Cliff Stein
Hallucination Helps: Energy Efficient Virtual Circuit Routing
In Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014).

92. Sungjin Im and Benjamin Moseley
New Approximations for Reordering Buffer Management
In Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014).

93. Kyle Fox, Sungjin Im, Janardhan Kulkarni and Benjamin Moseley
Online Non-clairvoyant Scheduling to Simultaneously Minimize All Convex Functions
In Proceedings of the 16th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2013).

94. Ravi Kumar, Benjamin Moseley, Sergei Vassilvitskii and Andrea Vattani
Fast Greedy Algorithms in MapReduce and Streaming
In Proceedings of the 25th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2013).
Awarded Best Paper

95. Sungjin Im and Benjamin Moseley
Online Batch Scheduling for Flow Objectives
In Proceedings of the 25th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2013). Brief Announcement.

96. Arpita Ghosh, Satyen Kale, Kevin Lang and Benjamin Moseley
Bargaining for Revenue Shares on Tree Trading Networks
In Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013).

97. Benjamin Moseley, Kirk Pruhs and Cliff Stein
The Complexity of Scheduling for p-norms of Flow and Stretch
In Proceedings of the 16th Conference on Integer Programming and Combinatorial Optimization (IPCO 2013).

98. Kyle Fox, Sungjin Im and Benjamin Moseley
Energy Efficient Scheduling of Parallelizable Jobs
In Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013).

99. Neal Barcelo, Sungjin Im, Benjamin Moseley and Kirk Pruhs
Shortest-Elapsed-Time-First on a Multiprocessor
In Proceedings of the Mediterranean Conference on Algorithms (MedAlg 2012).

100. Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar and Sergei Vassilvitskii
Scalable K-Means++
In Proceedings of the International Conference on Very Large Data Bases (VLDB 2012).

101. Kevin Lang, Benjamin Moseley and Sergei Vassilvitskii
Handling Forecast Errors while Bidding for Display Advertising
In Proceedings of the International Conference on World Wide Web (WWW 2012).

102. Sungjin Im, Benjamin Moseley and Kirk Pruhs
Online Scheduling with General Cost Functions
In Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)

103. Anupam Gupta, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley and Kirk Pruhs
Scheduling Heterogeneous Processors Isn't As Easy As You Think
In Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)

104. Alina Ene, Sungjin Im and Benjamin Moseley
Fast Clustering using MapReduce
In Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2011)

105. Silvio Lattanzi, Benjamin Moseley, Siddharth Suri and Sergei Vassilvitskii
Filtering: A Method for Solving Graph Problems in MapReduce
In Proceedings of the 23rd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2011)

106. Benjamin Moseley, Anirban Dasgupta, Ravi Kumar and Tamas Sarlos
On Scheduling in Map-Reduce and Flow-Shops
In Proceedings of the 23rd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2011)

107. Kyle Fox and Benjamin Moseley
Online Scheduling on Identical Machines using SRPT
In Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)

108. Jeff Edmonds, Sungjin Im and Benjamin Moseley
Online Scalable Scheduling for the ℓ_k-norms of Flow Time Without Conservation of Work
In Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)

109. Sungjin Im and Benjamin Moseley
An Online Scalable Algorithm for Minimizing ℓ_k-norms of Weighted Flow Time on Unrelated Machines
In Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)

110. Chandra Chekuri, Avigdor Gal, Sungjin Im, Samir Khuller, Jian Li, Richard McCutchen, Benjamin Moseley and Louiqa Raschid
New Models and Algorithms for Throughput Maximization in Broadcast Scheduling
In Proceedings of the 8th Workshop on Approximation and Online Algorithms (WAOA 2010)

111. Anupam Gupta, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley and Kirk Pruhs
Scheduling Jobs with Varying Parallelizability to Reduce Variance
In Proceedings of the 22nd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA 2010)

112. Sungjin Im and Benjamin Moseley
An Online Scalable Algorithm for Average Flow Time in Broadcast Scheduling
In Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)
Awarded Best Student Paper

113. Chandra Chekuri, Sungjin Im and Benjamin Moseley
Longest Wait First for Broadcast Scheduling
In Proceedings of the 7th Workshop on Approximation and Online Algorithms (WAOA 2009)

114. Chandra Chekuri, Sungjin Im and Benjamin Moseley
Minimizing Maximum Response Time and Delay Factor in Broadcast Scheduling
In Proceedings of the 17th Annual European Symposium on Algorithms (ESA 2009)

115. Chandra Chekuri and Benjamin Moseley
Online Scheduling to Minimize the Maximum Delay Factor.
In Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2009)

REFEREED JOURNAL PAPERS

116. Lin An, Andrew Li, Benjamin Moseley and R. Ravi
The Nonstationary 'ndor with (and without) Predictions
Manufacturing & Service Operations Management (MSOM). Accepted Feb 2025.

117. Sungjin Im, Benjamin Moseley, Hung Ngo and Kirk Pruhs
Polynomial Time Convergence of the Iterative Evaluation of Datalogo Programs
In Proceedings of the ACM on Management of Data (PACMMOD). Accepted Sept. 2024.

118. Franziska Eberle, Anupam Gupta, Nicole Megow, Benjamin Moseley, and Rudy Zhou.
Configuration balancing for stochastic requests.
Mathematical Programming. Accepted August 2024.

119. Anupam Gupta, Benjamin Moseley, and Rudy Zhou.
Structural iterative rounding for generalized k-median problems.
Mathematical Programming. Accepted July 2024.

120. Benjamin Moseley, Kirk Pruhs, Clifford Stein, and Rudy Zhou
A Competitive Algorithm for Throughput Maximization on Identical Machines
Mathematical Programming, Jan. 2024

121. Benjamin Moseley and Joshua Wang
Approximation Bounds for Hierarchical Clustering: Average-Linkage, Bisecting K-means, and Local Search
Journal of Machine Learning Research (JMLR). 24: 1:1-1:36 (2023)

122. Benjamin Berg, Justin Whitehouse, Benjamin Moseley, Weina Wang, and Mor Harchol-Balter
The Case for Phase-Aware Scheduling of Parallelizable Jobs
SIGMETRICS Performance Evaluation Review. 49(3): 65-66 (2022)

123. Benjamin Moseley, Ruilong Zhang and Shanjiawen Zhao
Online Scheduling of Parallelizable jobs in the Directed Acyclic Graphs and Speed-up Curves Models
Theoretical Computer Science. 938: 24-38 (2022).

124. Marilena Leichter, Benjamin Moseley and Kirk Pruhs
On the Impossibility of Decomposing Binary Matroids
Operations Research Letters. 50(5): 623-625 (2022)

125. Benjamin Moseley and Shai Vardi
The Efficiency-Fairness Balance of Round Robin Scheduling
Operations Research Letters. 50(1): 20-27 (2022)

126. Sungjin Im, Benjamin Moseley, and Rudy Zhou
The Matroid Cup Game
Operations Research Letters. 49(3): 405-411 (2021)

127. Giorgio Lucarelli, Benjamin Moseley, Nguyen Kim Thang, Abhinav Srivastav and Denis Trystram
Online Non-preemptive Scheduling on Unrelated Machines with Rejections
TACM Transactions on Parallel Computing.8(2): 9:1-9:22 (2021)
Special Issue for best papers at SPAA 2018.

128. Sungjin Im, Benjamin Moseley, and Kirk Pruhs
The Matroid Intersection Cover Problem
Operations Research Letters. 49(1): 17-22 (2021)

129. Varun Gupta, Benjamin Moseley, Marc Uetz and Qiaomin Xie
Corrigendum: Greed Works - Online Algorithms For Unrelated Machine Stochastic Scheduling
Mathematics of Operations Research (MOR). 46(3): 1230-1234 (2021).

130. Mahmoud Abo Khamis, Ryan Curtin, Benjamin Moseley, Hung Ngo, Long Nguyen, Dan Olteanu and Maximilian Schleich
On Functional Aggregate Queries with Additive Inequalities
ACM Transactions on Database Systems (TODS 2020). 45(4): 17:1-17:41 (2020) ***Special Issue for best papers at PODS 2019.***

131. Sungjin Im, Benjamin Moseley and Shi Li
Breaking $1 - 1/e$ Barrier for Non-preemptive Throughput Maximization
SIAM Journal of Discrete Math (SIDMA). 34(3): 1649-1669 (2020).

132. Sungjin Im, and Benjamin Moseley
Fair Scheduling via Iterative Quasi-Uniform Sampling
SIAM Journal on Computing (SICOMP). 658-680 (2020).

133. Bryce Bagley, Blake Bordelon, Benjamin Moseley, and Ralf Wessel
Pre-Synaptic Pool Modification (PSPM): A Supervised Learning Procedure for Recurrent Spiking Neural Networks
PLoS One. 2020; 15(2): e0229083.

134. Sungjin Im, Benjamin Moseley, Kamesh Munagala, and Kirk Pruhs
Dynamic Weighted Fairness with Minimal Disruptions
Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS) . 4(1): 19:1-19:18 (2020).

135. Antonios Antoniadis, Sungjin Im, Ravishankar Krishnaswamy, Vishwanath Nagarajan, Benjamin Moseley, Kirk Pruhs and Cliff Stein
Hallucination Helps: Energy Efficient Virtual Circuit Routing
SIAM Journal on Computing (SICOMP). 49(1): 37-66 (2020).

136. Kyle Fox, Sungjin Im, Janardhan Kulkarni and Benjamin Moseley
Online Non-clairvoyant Scheduling to Simultaneously Minimize All Convex Functions

Algorithmica 81(9): 3746-3764 (2019).

137. Varun Gupta, Benjamin Moseley, Marc Uetz and Qiaomin Xie
Greed Works - Online Algorithms For Unrelated Machine Stochastic Scheduling
Mathematics of Operations Research (MOR). 45:2, 497-516 (2018). There is a subtle error in the result. See the corrigendum published at Mathematics of Operations Research where we show slightly looser bounds.

138. Kyle Fox, Sungjin Im and Benjamin Moseley. Energy Efficient Scheduling of Parallelizable Jobs. *Theoretical Computer Science. 726: 30-40. 2018.*

139. Roozbeh Ebrahimi, Samuel McCauley and Benjamin Moseley. Scheduling Parallel Jobs Online with Convex and Concave Parallelizability. *Theory of Computing Systems. 62(2): 304-318 (2018).*
Special Issue for best papers at WAOA 2015.

140. Shalmoli Gupta, Ravi Kumar, Kefu Lu, Benjamin Moseley, and Sergei Vassilvitskii
Local Search Methods for k-Means with Outliers
In Proceedings of the International Conference on Very Large Data Bases 10(7): 757-768 (2017).

141. Michael A. Bender, Jonathan W Berry, Simon Hammond, Karl Hemmert, Samuel McCauley, Branden Moore, Benjamin Moseley, Cynthia A Phillips, David Resnick, and Arun Rodrigues
Two-Level Main Memory Co-Design: Multi-Threaded Algorithmic Primitives, Analysis, and Simulation
Journal of Parallel and Distributed Computing. 102: 213-228 (2017).
Special Issue for best papers at IPDPS 2015

142. Sungjin Im, Benjamin Moseley, Kirk Pruhs and Eric Torng
Competitively Scheduling Tasks with Intermediate Parallelizability
In ACM Transactions on Parallel Computing
Special Issue for best papers at SPAA 2014*. 3(1): 4 (2016).*

143. Ravi Kumar, Benjamin Moseley, Sergei Vassilvitskii and Andrea Vattani
Fast Greedy Algorithms in MapReduce and Streaming
In ACM Transactions on Parallel Computing. 2(3): 14 (2015).
Special Issue for best papers at SPAA 2013

144. Benjamin Moseley
Scheduling to Minimize Energy and Flow Time in Broadcast Scheduling
In Journal of Scheduling. 18(1): 107-118 (2015).

145. Sungjin Im, Benjamin Moseley and Kirk Pruhs
Online Scheduling with General Cost Functions
In SIAM Journal on Computing. 43(1): 126-143 (2014).

146. Daniel Cole, Sungjin Im, Benjamin Moseley and Kirk Pruhs
Speed Scaling for Total Stretch Plus Energy
In Operations Research Letters. 40(3): 180-184 (2012).

147. Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar and Sergei Vassilvitskii
Scalable K-Means++
In Proceedings of the International Conference on Very Large Data Bases. 5(7): 622-633 (2012).

148. Chandra Chekuri, Sungjin Im and Benjamin Moseley
On Scheduling to Minimize Maximum Response Time and Maximum Delay Factor
In Theory of Computing. 8(1): 165-195 (2012).
Special Issue in honor of Rajeev Motwani

149. Sungjin Im and Benjamin Moseley
An Online Scalable Algorithm for Average Flow Time in Broadcast Scheduling
In ACM Transactions on Algorithms. 8(4): 39 (2012)

BOOK AND TUTORIALS

150. Sungjin Im, Ravi Kumar, Silvio Lattanzi, Benjamin Moseley and Sergei Vassilvitskii
Massively Parallel Computation: Algorithms and Applications
In Foundations and Trends in Optimization: Vol. 5: No. 4, pp 340-417. 2023.

151. Sungjin Im, Benjamin Moseley and Kirk Pruhs
A Tutorial on Amortized Local Competitiveness in Online Scheduling
In SIGACT News. 42(2): 83-97 (2011).

WORKING JOURNAL SUBMISSIONS

152. Lin An, Andrew Li, Benjamin Moseley and Gabriel Visotsky
Online Resource Allocation with Predictions under Unknown Arrival Model
In Submission to Management Science. Major Revision.

153. Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, Sergei Vassilvitskii
Controlling Tail Risk in Online Ski-Rental
In Submission to Operations Research.

154. Sami Davies, Benjamin Moseley and Heather Newman
Fast Combinatorial Algorithms for Simultaneously Approximating All ℓ_p-norms in Correlation Clustering
In Submission to Mathematics of Operations Research.

155. Benjamin Berg, Mor Harchol-Balter, Benjamin Moseley, and Weina Wang,
Asymptotically Optimal Scheduling of Multiple Parallelizable Job Classes
In Submission to Stochastic Systems

FUNDING

- Office of Naval Research: Optimization for Submarine Scheduling. Total: $643,000. Moseley, PI. 01/01/2025- 12/30/2027
- Office of Naval Research: Optimization for Submarine Scheduling. Total: $400,000. Moseley, PI. 08/01/2022- 7/30/2025
- NSF: Collaborative Research: AF: Small: Foundations of Algorithms Augmented with Predictions. Total: $250,000. Moseley, PI. 01/01/2022- 12/31/2024
- NSF: OE: Optimal Scheduling for Parallelizable Jobs. Total: $550,000. Mor Harchol-Balter, PI. Moseley, co-PI. 01/01/2020-12/31/2022
- Google Faculty Research Award: "Algorithmic Techniques for Fast Scalable Hierarchical Clustering". Moseley, PI. 2019. Total: $81,631 (Gift).
- NSF: CAREER: Pushing the Theoretical Limits of Scalable Distributed Algorithms. Moseley, PI. Total $500,000. 7/01/2019-6/30/2024.
- Infor Faculty Award. Moseley, PI. 2018. Total: $100,000. (Gift).
- NSF: SPX: Collaborative Research: Harnessing the Power of High-Bandwidth Memory via Provably Efficient Parallel Algorithms. Moseley, PI. Total: $249,999. 09/15/17-8/31/2021.
- NSF: AITF: Applied Algorithmic Foundation for Scheduling Multiprogrammed Parallelizable Workloads. Moseley, PI. Co-PIs Kunal Agrawal (WashU) and Angelina Lee (WashU). Total: $650,000 . 10/1/17-9/30/21.
- NSF: AF Small: Collaborative Research: Algorithmic and Computational Frontiers of MapReduce for Big Data Analysis. Moseley, PI. 2016. Total: $252,767. 7/1/16-6/30/19.
- NSF: REU Site: Big Data Analytics. Sanmay Das, PI. Moseley, Senior Personnel. 2016. Total: $359,111. 4/2/16-3/31/19.
- Yahoo! Academic Career Enhancement (ACE) Award. Moseley, PI. 2015. Total $10,000 (Gift)
- Google Faculty Research Award: "Fast Distributed Algorithms for Clustering Data". Moseley, PI. 2015. Total: $31,845 (Gift).

PHD STUDENTS (THESIS SUPERVISOR)

- Jody Zhou 2023.
- Helia Niaparasat. 2023.
- Aidin Niaparasat. 2021. With R. Ravi.
- Lin An. Entered 2021. With Andrew Li.
- Heather Newman (Math ACO). Entered 2020.

- Rudy Zhou. Entered Fall 2018. Thompson Dissertation Award Award Winner. First Position: Postdoc CMU.
- Yuyan Wang. Entered Fall 2017. Graduated 2022. First Position: Google
- Thomas Lavastida. Entered Fall 2017. Graduated 2022. First position: Jindal School of Business, University of Texas Dallas .
- Kefu Lu. Entered Fall 2014. Graduation August 2019. First position: Assistant Professor at Washington and Lee University.

LONG TERM PHD STUDENT VISITORS

- Ruilong Zhang (City University of Hong Kong). Fall 2021-Spring 2022
- Chenyang Xu (Zhejiang University). Fall 2019 - Spring 2021
- Marilena Leichter (University of Munich). Spring 2020-Spring 2021 [virtual during COVID]

UNDERGRADUATE ADVISEES

- Helen Zhao. CMU, Computer Engineering. 2021.
- Matt Abbott. 2017 Summer REU.
- Bryce Bagley. 2017 Spring independent study and Summer. → PhD Stanford in Biophysics.
- Geoff Converse. 2016 Summer REU. → PhD University of Iowa in Math.
- Alyssa Shofner. 2016 Summer REU.
- Blakeley Hoffman. 2015 Summer REU and independent study Fall 2015. → Student at MIT CS.
- Roman Blum. 2015 research project. → PhD Brown University in CS.
- Diqiu Zhou. 2015 research project.

PHD THESIS COMMITTEES

- Alexander Lindermayr (Universitat Bremen, Mathematics and Computer Science)
- Hao-Ting Wei (Columbia, IEOR)
- Mik Zlatin (CMU, OR)
- David Stalfa (Northeastern University, CS)
- Mirmahdi Rahgoshay (University of Alberta, CS)
- Benjamin Berg (CMU, CS).
- Melda Korkut (CMU, OR).
- Marilena Leichter (University of Munich, Operations Research)
- Ziye Tang (CMU, OR)
- Alireza Samadian Zakaria (University of Pittsburgh, CS)
- Ryo Kimura (CMU, OR). First position: Wayfair.
- Shali Jiang (Washington University in St. Louis, CS). First position: Facebook.
- Matt Kusner (Washington University in St. Louis, CS). First position: Postdoc at the Alan Turing Institute. Currently an Assistant Professor at Oxford

INVITED PRESENTATIONS

2024 Cornell IEOR
2024 Cornell Tech
2024 TTIC and the University of Chicago
2024 Columbia Business School and IEOR
2024 Penn State
2024 International Symposium on Mathematical Programming (ISMP)
2024 Johns Hopkins University
2024 Worksop on New Challenges in Scheduling
2023 Technical University of Munich
2023 ICALP Workshop on Online Algorithms
2022 Washington University in St. Louis
2022 Simons Workshop on Quantifying Uncertainty: Stochastic, Adversarial and Beyond
2022 University of Pittsburgh
2022 STOC Workshop on Algorithms with Predictions
2021 Mindshare by CMU
2021 Yahoo! Research Seminar
2021 The Scheduling Seminar (schedulingseminar.com)
2021 Facebook Operations Research Seminar
2021 SQUALL Seminar CMU

2021 Machine Learning for Algorithms Workshop
2020 Combinatorics, Game Theory, and Optimisation seminar London School of Economics (LSE)
2020 AGCO (Algorithms, Games, Combinatorics and Optimization)
joint seminar of Pontificia Universidad Catolica de Chile and Universidad de Chile
2020 Dagstuhl Seminar on Networking and Scheduling
2020 Operations Research Center MIT
2019 TTIC Workshop on Learning-Based Algorithms
2018 Johns Hopkins University
2018 Carnegie Mellon University, SQUALL Seminar
2018 Open Problem Presentation at the TTIC Summer Workshop on Data Center Scheduling
2018 TTIC Summer Workshop on Data Center Scheduling from Theory to Practice
2018 International Symposium on Mathematical Programming (ISMP)
2018 Carnegie Mellon University, Machine Learning Seminar
2018 Carnegie Mellon University, Operations Research Seminar
2018 Carnegie Mellon University, Combinatorial Optimization Seminar
2017 George Washington University
2017 University of Chicago
2017 Carnegie Mellon University
2017 Columbia University
2017 Saint Louis University
2016 Simons Institute at UC-Berkeley
2016 Stanford University
2016 LogicBlox
2016 Technical University of Munich
2015 Workshop at STOC 2015. Algorithmic Frontiers of Massively Parallel Computation
2014 University of California Merced
2014 University of Pittsburgh
2014 Dartmouth College
2014 Virginia Tech
2014 Columbia University
2013 Northwestern University
2013 Washington University
2013 University of Illinois at Chicago
2013 Dagstuhl seminar on scheduling
2013 University of Chicago
2012 University of Minnesota
2012 Yahoo
2012 Google
2012 Oregon State University
2012 University of Buffalo
2012 Toyota Technological Institute at Chicago
2012 Sandia National Laboratories
2012 Shonan meeting on large-scale distributed computation

Conference and Workshop Presentations

- "Incremental Topological Ordering and Cycle Detection with Predictions". MAPSP, Kolding Denmark. 2024.
- "Faster Matchings via Learned Duals". MAPSP, Oropa Italy. 2022.
- "Submodular Optimization with Contention Resolution Extensions". APPROX, Boston MA. 2019.
- "Online Stochastic Scheduling on Unrelated Machines". INFORMS. Seattle, WA.
- "Online Scheduling with Learned Weights". Workshop on Frontiers of Scheduling Theory. Sanya, China.
- "Scheduling to Approximate Minimization Objectives on Identical Machines". ICALP, 2019. Patras, Greece.
- " Matroid Coflow Scheduling". ICALP, 2019. Patras, Greece.
- "Scheduling to Approximate Minimization Objectives on Identical Machines." Renesse Netherlands.

MAPSP 2019.

- "Submodular Optimization with Contention Resolution Extensions". Informs 2018. Phoenix, AZ.
- "Fair Scheduling via Iterative Quasi-Uniform Sampling". SODA 2017. Barcelona, Spain.
- " General Profit Scheduling and the Power of Migration on Heterogeneous Machines". SPAA 2016. Monterey, CA.
- "Scheduling Parallel DAG Jobs to Minimize the Average Flow Time". New Challenges in Scheduling Theory. Aussois, France.
- "Scheduling in Bandwidth Constrained Tree Networks". Portland, OR. SPAA 2015.
- "Fast and Better Distributed MapReduce Algorithms for k-Center Clustering". Portland, OR. SPAA 2015.
- "Online Non-clairvoyant Scheduling to Simultaneously Minimize All Convex Functions". La Roche-en-Ardenne, Belgium. MAPSP 2015.
- "Two-Level Main Memory Co-Design: Multi-Threaded Algorithmic Primitives, Analysis, and Simulation". Hyderabad, India. IPDPS 2015. Best Paper plenary presentation.
- "Competitively Scheduling Tasks with Intermediate Parallelizability". Prague, Czech Republic. SPAA 2014.
- " New Approximations for Reordering Buffer Management". Portland, OR. SODA 2014.
- "Bargaining for Revenue Shares on Tree Trading Networks". Beijing, China. IJCAI 2013.
- "Fast Greedy Algorithms in MapReduce and Streaming". Montreal, Canada. SPAA 2013. Best paper presentation.
- "Online Batch Scheduling for Flow Objectives". Montreal, Canada. SPAA 2013.
- "The Complexity of Scheduling for p-norms of Flow and Stretch". Valparaiso, Chile. IPCO 2013.
- "Scheduling Heterogeneous Processors Isn't As Easy As You Think". Kyoto, Japan. SODA 2012.
- "Fast Clustering using MapReduce". Chicago, IL. Midwest Theory Day 2012 at Northwestern.
- "Filtering: A Method for Solving Graph Problems in MapReduce". San Jose, CA. SPAA 2011.
- "On Scheduling in Map-Reduce and Flow-Shops ". San Jose, CA. SPAA 2011.
- "An Online Scalable Algorithm for Minimizing ℓ_k-norms of Weighted Flow Time on Unrelated Machines". San Francisco, CA. SODA 2011.
- "An Online Scalable Algorithm for Average Flow Time in Broadcast Scheduling". Austin, TX. SODA 2010. Best Student Paper presentation.
- "Online Scheduling to Minimize the Maximum Delay Factor".
  - New York, NY. SODA 2009.
  - Chicago, IL. Midwest Theory Day 2008 at Northwestern.

INSTRUCTION

Carnegie Mellon University: Numbers in the right margin are student evaluations of instructor/course quality (maximum 5.0/5.0).

Mini 1 2024	46-880: Probability and Statistics (Section 1)	MSBA	48 Students	(4.56/4.56)
Mini 1 2024	45-750: Probability and Statistics (Section 3)	MBA	33 Students	(4.70/4.80)
Mini 1 2024	45-750: Probability and Statistics (Section 2)	MBA	32 Students	(4.75/4.56)
Mini 1 2024	45-750: Probability and Statistics (Section 1)	MBA	31 Students	(4.72/4.89)
Spring 2024	70-257: Optimization for Business (Section 1)	undergrad	53 Students	(4.43/4.25)
Spring 2024	70-257: Optimization for Business (Section 2)	undergrad	28 Students	(4.71/4.76)
Mini 2 2023	47-842: Network Optimization 2	grad	5 Students	(NA/NA)
Mini 1 2023	46-880: Probability and Statistics (Section 1)	MSBA	48 Students	(4.44/4.30)
Mini 1 2023	45-750: Probability and Statistics (Section 3)	MBA	33 Students	(4.73/4.59)
Mini 1 2023	45-750: Probability and Statistics (Section 2)	MBA	32 Students	(4.48/4.39)
Mini 1 2023	45-750: Probability and Statistics (Section 1)	MBA	31 Students	(4.45/4.41)
Spring 2023	70-257: Optimization for Business (Section 1)	undergrad	31 Students	(4.42/4.33)
Spring 2023	70-257: Optimization for Business (Section 2)	undergrad	59 Students	(4.41/4.18)
Mini 2 2022	46-880: Probability and Statistics (Section 1)	MSBA	28 Students	(4.5/4.5)
Mini 2 2022	46-880: Probability and Statistics (Section 2)	MSBA	24 Students	(4.32/4.23)
Mini 2 2022	47-842: Network Optimization 2	grad	5 Students	(NA/NA)
Mini 1 2022	45-750: Probability and Statistics (Section 3)	MBA	33 Students	(4.41/4.32)
Mini 1 2022	45-750: Probability and Statistics (Section 2)	MBA	35 Students	(4.62/4.52)
Mini 1 2022	45-750: Probability and Statistics (Section 1)	MBA	36 Students	(4.65/4.5)
Spring 2022	70-257: Optimization for Business (Section 1)	undergrad	71 Students	(4.72/4.46)
Spring 2022	70-257: Optimization for Business (Section 2)	undergrad	58 Students	(4.68/4.46)
Mini 2 2021	46-880: Intro Probability and Statistics	MSBA	54 Students	(4.48/4.48)
Mini 2 2021	47-842: Advanced Graph Theory	grad	2 Students	(NA/NA)
Mini 1 2021	45-750: Probability and Statistics (Section 3)	MBA	36 Students	(4.22/4.22)
Mini 1 2021	45-750: Probability and Statistics (Section 2)	MBA	35 Students	(4.10/4.13)
Mini 1 2021	45-750: Probability and Statistics (Section 1)	MBA	36 Students	(4.24/4.21)
Spring 2021	70-257: Optimization for Business (Section 1)	undergrad	61 Students	(4.95/4.88)
Spring 2021	70-257: Optimization for Business (Section 2)	undergrad	55 Students	(4.74/4.81)
Mini 4 2021	47-842: Algorithms for Massive Data Science	grad	8 Students	(NA/NA)
Fall 2020	70-467: Machine Learning for Business Analytics	undergrad	55 Students	(4.52/4.52)
Mini 2 2020	47-842: Advanced Graph Theory	grad	8 Students	(NA/NA)
Spring 2020	70-257: Optimization for Business (Section 1)	undergrad	57 Students	(4.55/4.52)
Spring 2020	70-257: Optimization for Business (Section 2)	undergrad	34 Students	(4.87/4.67)
Mini 2 2019	47-842: Advanced Grapfh Theory	grad	5 Students	(NA/NA)
Mini 4 2019	47-842: Algorithms for Massive Data Science	grad	5 Students	(NA/NA)
Spring 2019	70-257: Optimization for Business (Section 1)	undergrad	77 Students	(4.55/4.45)
Spring 2019	70-257: Optimization for Business (Section 2)	undergrad	53 Students	(4.67/4.5)
Spring 2018	70-257: Optimization for Business (Section 1)	undergrad	52 Students	(4.48/4.33)
Spring 2018	70-257: Optimization for Business (Section 2)	undergrad	54 Students	(4.42/4.5)

Washington University: Numbers in the right margin are student evaluations of instructor/course quality (maximum 7.0/7.0).

Fall 2017	CSE 240: Logic and Discrete Mathematics	undergrad	174 Students	(6.04/5.95)
Fall 2017	CSE 544: Special Topics	grad	3 Students	(NA/NA)
Spring 2017	CSE 581: Approximation Algorithms	grad	29 Students	(6.62/6.69)
Fall 2015	CSE 240: Logic and Discrete Mathematics	undergrad	174 Students	(5.67/5.18)
Spring 2015	CSE 240: Logic and Discrete Mathematics	undergrad	67 Students	(6.52/6.33)
Fall 2014	CSE 581: Approximation Algorithms	grad	15 Students	(6.22/6.00)

EDUCATIONAL ACTIVITIES

- Developed a Minor in Business Analytics and Optimization with Andrew Li (CMU).
- Developed the first offering of the course Machine Learning for Business with Andrew Li (CMU).

- Developed several courses on business analytics, machine learning and optimization for executive education at CMU.

Professional Activities

- Editorial Boards
  - Associate Editor for Operations Research Letters 2017-
  - Associate Editor for IEEE Transactions on Knowledge and Data Engineering (TKDE) 2018-2023
- Guest Editor
  - Special Issue of Journal of Scheduling devoted to selected papers from the 2019 Workshop on Models and Algorithms for Planing and Scheduling Problems (MAPSP 2019)
  - Special Issue of Journal of Scheduling devoted to selected papers from the 2018 Seminar on New Challenges in Scheduling Theory
  - Special Issue of Journal of Scheduling devoted to selected papers from the 2016 Seminar on New Challenges in Scheduling Theory
  - Special Issue of Journal of Scheduling devoted to selected papers from the 2014 Seminar on New Challenges in Scheduling Theory
- Program committee chair for:
  - Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP 2024)
- Area Chair for:
  - Neural Information Processing Systems (NeurIPS) 2020, 2021,2022, 2023, 2024
  - International Conference on Machine Learning (ICML) 2021,2022, 2023, 2024
  - International Conference on Learning Representations (ICLR) 2020, 2021,2022, 2023, 2024
- Program committee member for:
  - ACM Principles of Database Systems (PODS 2026)
  - European Symposium on Algorithms (ESA 2025)
  - Integer Programming and Combinatorial Optimization (IPCO 2025)
  - ACM Symposium on Parallelism in Algorithms and Architectures 2024 (SPAA 2024)
  - Symposium on Theoretical Aspects of Computer Science (STACS 2023)
  - Workshop on Approximation and Online Algorithms (WAOA 2023)
  - SIAM Conference on Applied and Computational Discrete Algorithms (ACDA 2023)
  - ACM Symposium on Parallelism in Algorithms and Architectures 2022 (SPAA 2022)
  - SIAM Conference on Applied and Computational Discrete Algorithms (ACDA 2022)
  - ACM-SIAM Symposium on Discrete Algorithms (SODA 2022)
  - NeurIPS Workshop on Databases and AI 2021
  - Workshop on Approximation and Online Algorithms (WAOA 2021)
  - ACM Symposium on Parallelism in Algorithms and Architectures 2021 (SPAA 2021)
  - Scandinavian Workshop on Algorithm Theory (SWAT 2020)
  - European Conference on Parallel and Distributed Computing (Euro-Par 2020, Global Chair of Topic Theory and Algorithms for Parallel and Distributed Processing)
  - Workshop on Approximation and Online Algorithms (WAOA 2019)
  - International Symposium on Algorithms and Computation 2019 (ISAAC 2019)
  - International Conference on Machine Learning (ICML 2019)
  - Workshop on Models and Algorithms for Planning and Scheduling Problems 2019 (MAPSP 2019)
  - IEEE International Parallel & Distributed Processing Symposium 2019 (IPDPS 2019)
  - International Conference on Web Search and Data Mining (WSDM 2018)
  - Conference on Artificial Intelligence (AAAI 2018)
  - ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)
  - Knowledge Discovery, Data Mining, and Data Science Research (KDD 2017, Senior Program committee)

- European Symposium on Algorithms Track B (ESA 2017)
- International Conference on Parallel Processing (ICPP 2017)
- International Conference on Artificial Intelligence and Statistics (AISTATS 2017)
- IEEE International Parallel & Distributed Processing Symposium 2017 (IPDPS 2017)
- International Conference on Web Search and Data Mining (WSDM 2016)
- International Conference on High Performance Computing, Data and Analytics (HiPC 2016)
- Workshop on Approximation and Online Algorithms (WAOA 2016)
- IEEE International Parallel & Distributed Processing Symposium 2016 (IPDPS 2016)
- ACM Symposium on Parallelism in Algorithms and Architectures 2016 (SPAA 2016)
- Workshop on Models and Algorithms for Planning and Scheduling Problems 2015 (MAPSP 2015)
- International Computing and Combinatorics Conference 2015 (COCOON 2015)
- International Computing and Combinatorics Conference 2014 (COCOON 2014)
- International Symposium on Algorithms and Computation 2013 (ISAAC 2013)
- Workshop on Approximation and Online Algorithms (WAOA 2013)

- Organizer:
  - Dagstuhl on Fairness in Scheduling 2025. Warden, Germany.
  - Dagstuhl on Learned Predictions for Data Structures and Running Time 2025. Warden, Germany.
  - Dagstuhl on Beyond Worst-Case Analysis in Scheduling 2023. Warden, Germany.
  - Workshop on Frontiers of Scheduling Theory 2020. Aussois, France.
  - Workshop on Frontiers of Scheduling Theory 2019. Sanya, China.
  - Workshop on New Challenges in Scheduling Theory 2018. Aussois, France.
  - Workshop on New Challenges in Scheduling Theory 2016. Aussois, France.
  - Workshop on New Challenges in Scheduling Theory 2014. Aussois, France.
- Reviewer or referee for:
  - Symposium on Foundations of Computer Science (FOCS)
  - Symposium on the Theory of Computing (STOC)
  - Symposium on Discrete Algorithms (SODA)
  - Neural Information Processing Systems (NIPS)
  - International Colloquium on Automata, Languages and Programming (ICALP)
  - European Symposium on Algorithms (ESA)
  - International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)
  - International European Conference on Parallel and Distributed Computing (Euro-Par)
  - Workshop on Approximation and Online Algorithms (WAOA)
  - Symposium on Theoretical Aspects of Computer Science (STACS)
  - Latin American Theoretical Informatics Symposium (LATIN)
  - Innovations in Theoretical Computer Science (ITCS)
  - Integer Programming and Combinatorial Optimization (IPCO)
  - ACM Symposium on Parallel Algorithms and Architectures (SPAA)
  - ACM Transactions on Algorithms (TALG)
  - Operations Research
  - Management Science
  - Algorithmica

- Journal of Scheduling (JoS)
- IEEE Transactions on Knowledge and Data Engineering (TKDE)
- Concurrency and Computation: Practice and Experience
- Networks
- IEEE Transactions on Signal Processing
- Theory of Computing Systems
- Parallel Computing
- Mathematics of Operations Research
- Journal of the ACM
- SIAM Journal of Computing
- ACM Transactions on Algorithms
- Math Programming
- International Journal of Information and Communication Technology (IJICT)
- IEEE Transactions on Signal Processing
- Journal of Combinatorial Optimization
- Journal of Computer Science and Technology (JCST)
- Journal of Computer and System Sciences (JCSS)
- Transactions on Cloud Computing
- Many more