

Curriculum Vitae - Hang Liu

Assistant Professor
Stevens Institute of Technology
Hoboken, NJ 07030

202-216-8103
hliu77@stevens.edu
[Personal Homepage](#)

EDUCATION

- 2011 - 2017 **Ph.D. in High-Performance Computing**
Department of Electrical & Computer Engineering
The George Washington University
Adviser: H. Howie Huang
- 2007 - 2011 **B.E. in Software Engineering**
School of Software Engineering
Huazhong University of Science & Technology

EXPERIENCES

- 2022 - Present **Presidential Fellowship Assistant Professor**
Department of Electrical & Computer Engineering
Stevens Institute of Technology
- 2019 - 2022 **Assistant Professor**, Department of Electrical & Computer Engineering
Stevens Institute of Technology
- Summer 2019 & 2021 **Visiting Faculty**
Lawrence Berkeley National Laboratory
Host: Sherry X. Li & Aydin Buluc
- 2017 - 2019 **Assistant Professor**, Department of Electrical & Computer Engineering
University of Massachusetts Lowell
- 2011 - 2017 **Graduate Research Assistant**
The George Washington University
Adviser: H. Howie Huang
- Summer 2014 **Research Intern**
NEC Laboratories at America
Mentor: Cheng-Hong Li

HONORS & AWARDS

- 2022 IEEE CS TCHPC Early Career Researchers Award for Excellence in High Performance Computing (**One of the most prestigious awards for junior researchers in HPC**)
- 2022 - 2025 Presidential Fellow (2 awardees across the entire institute)
- 2022 3rd Prize at the 2022 Ansary Entrepreneurship Competition (Senior Design, Role: Advisor)
- 2022 Early Career Award for Research Excellence (2 awardees for the entire institute)
- 2021 NSF CAREER Award
- 2021 ECE Outstanding Research Award

- 2021&2019 Lawrence Berkeley National Laboratory SRP Fellowship
- 2020 One of the Best Papers in VLDB'20
- 2020 Excellent Teaching Evaluation Award
- 2019 NSF CRII Award
- 2019&2018 Champion of Graph Challenge Competition
- 2018 Best Dissertation Award, Electrical & Computer Engineering at GWU
- 2017 ICT Express Best Reviewer
- 2016 Phillip/Temofel Sprawcew Endowment Scholarship
- 2015 No. 1 Most Energy Efficient Graph Traversal at GreenGraph 500 (small graph category)

MEDIA COVERAGE

- July.2022 Interviewed by Dr. Marina Kraeva for [SC22 Mini-series: ECP Past Participants](#)
- May.2022 Quoted by a [CNET article](#) on AI and chip design
- Nov.2021 Quoted by a [Lifewire article](#) on fundamental value of data

RESEARCH - PUBLICATIONS

BOOK CHAPTERS

- 2018 Da Yan and **Hang Liu**. Parallel Graph Processing. In Encyclopedia of Big Data Technologies, *Springer*, 2018.

JOURNAL ARTICLES

- 2022 Yuede Ji, **Hang Liu**, Yang Hu and H. Howie Huang. iSpan: Parallel Identification of Strongly Connected Components with Spanning Trees. In *ACM Transaction on Parallel Computing (TOPC)*, 2022.
- 2021 Anil Gaihre, Xiaoye S. Li, and **Hang Liu**. GSOFA: Scalable Sparse LU Symbolic Factorization on GPUs. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2021.
- 2021 Santosh Pandey, Zhibin Wang, Sheng Zhong, Chen Tian, Lingda Li, Adolphy Hoise, Xiaoye S. Li, Caiwen Ding, Dong Li, Bolong Zheng and **Hang Liu**. TRUST: Triangle Counting on GPUs. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2021.
- 2020 Xu Xiang, **Hang Liu**, Tian Lan, Suresh Subramaniam, Howie Huang. Optimizing Job Reliability Through Contention-Free, Distributed Checkpoint Scheduling. In *IEEE Transactions on Network and Service Management (TNSM)*, 2020.
- 2017 Yunjie Zhao, Yiren Jian, Zhichao Liu, **Hang Liu**, Qin Liu, Chanyou Chen, Zhangyong Li, Lu Wang, H. Howie Huang, and Chen Zeng. [Network Analysis Reveals the Recognition Mechanism for Dimer Formation of Bulb-type Lectins](#). *Scientific Reports*, volume 7. Nature Publishing Group. 2017.
- 2016 Rajat Mittal, Jung Hee Seo, Vijay Vedula, Young J Choi, **Hang Liu**, H. Howie Huang, Saurabh Jain, Laurent Younes, Theodore Abraham, and Richard T George. [Computational Modeling of Cardiac Hemodynamics: Current Status and Future Outlook](#). In *Journal of Computational Physics (JCP)*. 305 (2016): 1065-1082.

REFEREED CONFERENCE PROCEEDINGS

In total, 26 top publications (i.e., listed on csrankings.org): SC 7, HPDC 2, ICS 1, SIGMETRICS 1, SIGMOD 2, VLDB 1, ISCA 1, USENIX FAST 2, USENIX ATC 1, Eurosys 2, DAC 3, EMNLP 1, ACL 1, IJCAI 1.

- 2023 Chengying Huan, Shuaiwen Leon Song, Santosh Pandey, **Hang Liu**, Yongchao Liu, Baptiste Lepers, Charles He, Kang Chen, Jinlei Jiang and Yongwei Wu. “TEA: A General-Purpose Temporal Graph Random Walk Engine.” In *Proceedings of the European Conference on Computer Systems (Eurosys)*. ACM, 2023.
- 2022 Santosh Pandey, Lingda Li, Thomas Flynn, Adolfo Hoisie and **Hang Liu**. “Scaling Deep Learning-based Microarchitecture Simulation on GPUs.” In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE, 2022.
- 2022 Chengying Huan, Shuaiwen Leon Song, Yongchao Liu, Heng Zhang, **Hang Liu**, Charles He, Kang Chen, Jinlei Jiang, and Yongwei Wu. “T-GCN: A Sampling Based Streaming Graph Neural Network System With Hybrid Architecture.” In *30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2022.
- 2022 Heng Zhang, Lingda Li, **Hang Liu**, Donglin Zhuang, Rui Liu, Chengying Huan, Shuang Song et al. “Bring Orders into Uncertainty: Enabling Efficient Uncertain Graph Processing via Novel Path Sampling on Multi-Accelerator Systems.” In *Proceedings of the 36th ACM International Conference on Supercomputing (ICS)*, 2022.
- 2022 Lingda Li, Santosh Pandey, Thomas Flynn, **Hang Liu**, Noel Wheeler, and Adolfo Hoisie. “SimNet: Accurate and High-Performance Computer Architecture Simulation using Deep Learning.” In *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS)*, 2022.
- 2022 Shaoyi Huang, Dongkuan Xu, Ian En-Hsu Yen, Yijue Wang, Sung-En Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Sanguthevar Rajasekaran, **Hang Liu**, Caiwen Ding. “Sparse Progressive Distillation: Resolving Overfitting under Pretrain-and-Finetune Paradigm.” In *Proceeding of 60th Annual Meeting of the Association for Computational Linguistics (ACL), main conference*, 2022.
- 2022 Hongwu Peng, Shaoyi Huang, Shiyang Chen, Bingbing Li, Tong Geng, Ang Li, Weiwen Jiang, Wujie Wen, Jinbo Bi, **Hang Liu** and Caiwen Ding. “A Length Adaptive Algorithm-Hardware Co-design of Transformer on FPGA Through Sparse Attention and Dynamic Pipelining.” In *Proceedings of ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2022.
- 2022 Chengying Huan, **Hang Liu**, Mengxing Liu, Yongchao Liu, Changhua He, Kang Chen, Jinlei Jiang, Yongwei Wu and Shuaiwen Leon Song. “TeGraph: A Novel General-Purpose Temporal Graph Computing Engine.” In *Proceedings of 2022 IEEE 38th International Conference on Data Engineering (ICDE)*. ACM, 2022.
- 2021 Shiyang Chen, Shaoyi Huang, Santosh Pandey, Bingbing Li, Guang Gao, Long Zheng, Caiwen Ding and **Hang Liu**. E.T.: Rethinking Transformer Models on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. ACM, 2021.
- 2021 Anil Gaihre, Da Zheng, Scott Weitze, Lingda Li, Caiwen Ding, Shuaiwen Song and **Hang Liu**. Dr. Top-k: Delegate Centric Top-k Computation on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. ACM, 2021.

- 2021 Yijue Wang, Chenghong Wang, Zigeng Wang, Shanglin Zhou, **Hang Liu**, Jinbo Bi, Caiwen Ding, Sanguthevar Rajasekaran. Against Membership Inference Attack: Pruning is All You Need. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*. 2021.
- 2021 Geng Yuan, Payman Behnam, Zhengang Li, Ali Shafiee, Sheng Lin, Xiaolong Ma, **Hang Liu**, Xuehai Qian, Mahdi Bojnordi, Yanzhi Wang, and Caiwen Ding. FORMS: Fine-grained Polarized ReRAM-based In-situ Computation for Mixed-signal DNN Accelerator. In *Proceedings of the 46th International Symposium on Computer Architecture (ISCA)*, 2021
- 2021 Hongwu Peng, Shiyang Chen, Zhepeng Wang, Junhuan Yang, Scott A. Weitze, Tong Geng, Ang Li, Jinbo Bi, Minghu Song, Weiwen Jiang, **Hang Liu** and Caiwen Ding. Optimizing FPGA-based Accelerator Design for Large-Scale Molecular Similarity Search (Special Session Paper). In *IEEE/ACM International Conference On Computer Aided Design (ICCAD)* 2021 Nov 1 (pp. 1-7). IEEE.
- 2021 Zhen Xie, Wenqian Dong, Jiawen Liu, **Hang Liu** and Dong Li. Tahoe: Tree Structure-Aware High Performance Inference Engine for Decision Tree Ensemble on GPU. In *Proceedings of the European Conference on Computer Systems (Eurosys)*. ACM, 2021.
- 2020 Santosh Pandey, Lingda Li, Adolffy Hoisie, Xiaoye S. Li and **Hang Liu**. C-SAW: A Framework for Graph Sampling and Random Walk on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE, 2020.
- 2020 Bolong Zheng, Xi Zhao, Lianggui Weng, Nguyen Quoc Viet Hung, **Hang Liu** and Christian S. Jensen. PM-LSH: A Fast and Accurate LSH Framework for High-Dimensional Approximate NN Search. In *Proceedings of the VLDB Endowment (VLDB)*. 2020. **One of the best papers in VLDB '20.**
- 2020 Bingbing Li, Zhenglun Kong, Tianyun Zhang, Ji Li, Zhengang Li, **Hang Liu**, Caiwen Ding. Efficient Transformer-based Large Scale Language Representations using Hardware-friendly Block Structured Pruning. In *Proceedings of ACL Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- 2020 Md Hafizul Islam Chowdhuryy, **Hang Liu**, Fan Yao, BranchSpec: Information Leakage Attacks Exploiting Speculative Branch Instruction Executions. In *Proceedings of the 38th IEEE International Conference on Computer Design (ICCD)*, 2020.
- 2020 Linnan Wang, Wei Wu, Junyu Zhang, **Hang Liu**, George Bosilca, Maurice Herlihy, and Rodrigo Fonseca. FFT-based Gradient Sparsification for the Distributed Training of Deep Neural Networks. In *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, pp. 113-124. 2020.
- 2020 Runbin Shi, Yuhao Ding, Xuechao Wei, He Li, **Hang Liu**, Hayden So, and Caiwen Ding. FTDL: A Tailored FPGA-Overlay for Deep Learning with High Scalability. In *ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2020.
- 2020 Bingbing Li, Santosh Pandey, Haowen Fang, Yanjun Lyv, Ji Li, Jieyang Chen, Mimi Xie, Lipeng Wan, **Hang Liu**, and Caiwen Ding. FTRANS: Energy-Efficient Acceleration of Transformers using FPGA. In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 175-180. 2020.
- 2020 Shilong Wang, Da Li, Hengyong Yu and **Hang Liu**. ELDA: Efficient LDA on GPUs (short paper) In *In Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP)*. 2020.

- 2020 Runbin Shi, Yuhao Ding, Xuechao Wei, **Hang Liu**, So Hayden, and Caiwen Ding. FTDL: An FPGA-Tailored Architecture for Deep Learning Applications (short paper). In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*. ACM, 2020.
- 2020 Yuede Ji, **Hang Liu**, Howie Huang. SwarmGraph: Analyzing Large-Scale In-Memory Graphs on GPUs. In the IEEE International Conference on High Performance Computing and Communications (**HPCC**), 2020.
- 2019 Santosh Pandey, Xiaoye S. Li, Aydin Buluc, Jiejun Xu and **Hang Liu**. H-INDEX: Hash-Indexing for Parallel Triangle Counting on GPUs. In *GraphChallenge*. 2019. **Awarded Champion**.
- 2019 Daniel Giger and **Hang Liu**. An Efficient Parallel Algorithm for Dominator Detection (ACM Undergraduate Poster Competition). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. 2019.
- 2019 **Hang Liu** and H. Howie Huang. SIMD-X: Programming and Processing of Graph Algorithms on GPUs. In *Proceedings of the 2019 USENIX Conference on Usenix Annual Technical Conference (USENIX ATC)*. USENIX Association. 2019.
- 2019 Anil Gaihre, Zhenlin Wu and **Hang Liu**. XBFS: eXploring Dynamic Optimizations for Breadth-First Search on GPUs. In *proceedings of the 28th international symposium on High-performance parallel and distributed computing (HPDC)*. ACM. 2019.
- 2019 Bibek Bhattarai, **Hang Liu** and H. Howie Huang. CECI: Compact Embedding Cluster Index for Scalable Subgraph Matching. In *Proceedings of ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, 2019.
- 2019 Eric Finnerty, Zach Sherer, Yan Luo and **Hang Liu**. Dr. BFS: Data Centric Breadth-First Search on FPGAs. In *56th ACM/ESDA/IEEE Design Automation Conference (DAC)*. IEEE. 2019.
- 2019 Hao Jin, Chen Xu, Yan Luo, Peilong Li, **Hang Liu** and Chunyang Hu. A Blockchain based Approach for Secure and Privacy-Preserving Medical Data Sharing. In *IFIP Networking Conference (IFIP Networking)*. IEEE. 2019. (WIP)
- 2019 Jialing Zhang, **Hang Liu** and Seung Woo Son. Efficient Encoding and Reconstruction of HPC Datasets for Checkpoint/Restart. In *35th Symposium on Mass Storage Systems and Technologies (MSST)* (pp. 1-12). IEEE. 2019.
- 2019 Zach Sherer, Eric Finnerty, Yan Luo and **Hang Liu**. Software and Hardware Co-Optimized BFS on FPGAs. In *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*. ACM, 2019.
- 2018 Anil Gaihre, Yan Luo and **Hang Liu**. [Do Bitcoin Users Really Care About Anonymity: An Analysis of the Bitcoin Transaction Graph](#). In *Proceedings of IEEE International Conference on Big Data (BigData)*. IEEE, 2018.
- 2018 **Hang Liu**, Yang Hu and H. Howie Huang. [High-Performance Triangle Counting on GPUs](#). In *GraphChallenge*. 2018. **Awarded Champion**.
- 2018 Yang Hu, **Hang Liu** and H. Howie Huang. [TriCore: Parallel Triangle Counting on GPUs](#). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE, 2018.
- 2018 Yuede Ji, **Hang Liu** and H. Howie Huang. [iSpan: Parallel Identification of Strongly Connected Components with Spanning Trees](#). In *Proceedings of the International Conference*

- for High Performance Computing, Networking, Storage and Analysis (SC). IEEE, 2018.
- 2018 Nai Xia, Tian Chen, Yan Luo, **Hang Liu** and Xiaoliang Wang. [UKSM: Swift Memory Deduplication via Hierarchical and Adaptive Memory Region Distilling](#). In *16th USENIX Conference on File and Storage Technologies (USENIX FAST)*. 2018.
- 2017 Aekeyung Moon, Jaeyoung Kim, Jialing Zhang, **Hang Liu** and SeungWoo Son. [Understanding the Impact of Lossy Compressions on IoT Smart Farm Analytics](#). In *IEEE BigData Workshop on Big Data Analytics for Internet of Things*, 2017.
- 2017 **Hang Liu** and H. Howie Huang. [Graphene: Fine-Grained IO Management for Graph Computing](#). In *15th USENIX Conference on File and Storage Technologies (USENIX FAST)*. 2017.
- 2016 **Hang Liu**, H. Howie Huang, and Yang Hu. [iBFS: Concurrent Breadth-First Search on GPUs](#). In *Proceedings of ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, 2016.
- 2015 **Hang Liu** and H. Howie Huang. [Enterprise: Breadth-First Graph Traversal on GPUs](#). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. ACM, 2015.
- 2014 Yu Xiang, **Hang Liu**, Tian Lan, H. Howie Huang, and Suresh Subramaniam. [Optimizing Job Reliability via Contention-free, Distributed Scheduling of VM Checkpointing](#). In *Proceedings of the 2014 ACM SIGCOMM workshop on Distributed cloud computing (DCC)*. ACM, 2014.
- 2014 H. Howie Huang and **Hang Liu**. [Big Data Machine Learning and Graph Analytics: Current State and Future Challenges](#). In *International Conference on Big Data (BigData)*. IEEE, 2014.
- 2013 **Hang Liu**, Jung-Hee Seo, Rajat Mittal, and H. Howie Huang. [GPU-Accelerated Scalable Solver for Banded Linear Systems](#). In *International Conference on Cluster Computing (CLUSTER)*. IEEE, 2013.
- 2012 **Hang Liu**, Jung-Hee Seo, Rajat Mittal, and H. Howie Huang. [Matrix Decomposition Based Conjugate Gradient Solver for Poisson Equation \(short paper\)](#). In *Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE, 2012.

RESEARCH - GRANTS

- NSF *Collaborative Research: SHF: Medium: Co-optimizing Spectral Algorithms and Systems for High-Performance Graph Learning*
- Role: Co-PI (PI: Zhuo Feng);
 - Total: \$800,000; Personal share: \$400,000 (50%);
 - Sponsor: National Science Foundation;
 - 2022.5 - 2026.4.
- NSF *CAREER: A Framework for Graph Sampling and Random Walk on GPUs*
- Role: Sole PI;
 - Total: \$584,001;
 - Sponsor: National Science Foundation (#2046102);
 - 2021.1 - 2025.12.
- NSF *CRII: SHF: Expediting Subgraph Matching on GPUs*

- Role: Sole PI;
- Total: \$190,000;
- Sponsor: National Science Foundation (#2000722);
- Period: 2019.09 - 2022.01.

- DOE *GPU Accelerated Symbolic Factorization for SuperLU*
- Role: Sole PI;
 - Total: \$320,000;
 - Sponsor: Department of Energy (Lawrence Berkeley National Laboratory);
 - Period: 2019.09 - 2023.08.
- Air Force *Real-Time Image Stitching on FPGAs*
- Role: Hang Liu (PI), Kevin Lu and Min Song;
 - Total: \$224,999;
 - My share: \$179,999;
 - Sponsor: Air Force (Circle Optics);
 - Period: 2020.06 - 2022.08.
- DOE *SIMNET: Deep Learning Accelerated Microarchitectural Simulator*
- Role: Sole PI;
 - Total: \$174,581;
 - Sponsor: Department of Energy (Brookhaven National Laboratory);
 - Period: 2020.06 - 2022.08.
- LittleLights *Knowledge Graph Assisted Scalable Adaptive Learning for LittleLights.AI*
- Role: Sole PI;
 - Total: \$50,725;
 - Sponsor: LittleLights.AI (Industry);
 - Period: 2018.08 - 2019.08.
- Amazon *Graph Mining at Extreme Scale*
- Role: Sole PI;
 - Total: \$68,000 (cloud credit);
 - Sponsor: Amazon AWS;
 - Period: 2018.07 - 2019.07.
- Intel *Real-Time Deep Learning on FPGAs*
- Role: Sole PI;
 - Total: Stratix 10 FPGA (worth \$10,000);
 - Sponsor: Amazon AWS;
 - Period: 2019.07.
- Nvidia *Expediting Asynchronous Graph Analytics on GPUs*
- Role: Sole PI;
 - Total: Quadro P6000 GPU (worth \$4,500);
 - Sponsor: Nvidia;
 - Period: 2018.07.
- Nvidia *Expediting Graph Mining on GPUs*
- Role: Sole PI;
 - Total: Titan Xp GPU (worth \$1,200);
 - Sponsor: Nvidia;
 - Period: 2017.10.
- Xilinx *Expediting Transformer Models on FPGAs*
- Role: Sole PI;
 - Total: Two Xilinx Alveo U280 FPGAs (worth \$17,146.92);
 - Sponsor: Xilinx;
 - Period: 2021.10.

RESEARCH - Invited Talks

- Aug.2022 Presenting the Overview of High-Performance Data Analytics Lab to NJ Secretary of Education for Higher Education (Dr. Brian Bridges)
Stevens Institute of Technology
- Dec.2021 Hardware and Software Co-designed Data Analytics
Jilin University, China
- Dec.2021 E.T.: Rethinking Transformer Models on GPUs
Stevens Institute of Technology
- Nov.2021 High-Performance Frameworks for Graph Sampling and Random Walk on GPUs
Stevens Institute of Technology
- Jan.2021 My NSF CAREER Proposal Writing Experience
Stevens Institute of Technology
- Dec.2020 How to Write a Technical Paper for ECE Student Professional Development Workshop
Stevens Institute of Technology
- Oct.2020 High-Performance Graph Sampling and Random Walk on GPUs
Brookhaven National Laboratory
- July.2019 SIMD-X: Programming and Processing of Graph Algorithms on GPUs
USENIX ATC, Renton, WA
- July.2019 Hardware Accelerated Data Science
Lawrence Berkeley National Laboratory, Berkeley, CA
- June.2019 Dr. BFS: Data Centric Breadth-First Search on FPGAs
DAC, Las Vegas, NV
- June.2019 Hardware Accelerated Data Analytics
Samsung Research Forum, San Jose, CA
- June.2019 Hardware Accelerated Graph Computing, Mining and Learning
HRL Laboratories, Malibu, CA
- April.2019 Graph Computing: System, Application and Future Directions
Massachusetts Institute of Technology, Cambridge, MA
- Mar.2019 Hardware Accelerated Data Analytics
Stevens Institute of Technology, Hoboken, NJ
- July.2018 High-Performance Graph Computing on GPUs
Nvidia Research, Westford, MA
- Feb.2018 Novel Techniques for Graph Algorithm Acceleration
Brown University, Providence, RI
- Feb.2017 Novel Techniques for Graph Algorithm Acceleration
University of Massachusetts Lowell, Lowell, MA
- Feb.2017 Novel Techniques for Graph Algorithm Acceleration
University of North Carolina Charlotte, Charlotte, NC
- Jan.2017 Novel Techniques for Graph Algorithm Acceleration
Clemson University, Clemson, SC
- Feb.2017

Graphene: Fine-Grained IO Management for Graph Computing
USENIX FAST, San Jose, CA

July.2016 iBFS: Concurrent Breadth-First Search on GPUs
SIGMOD, San Francisco, CA

Nov.2015 Enterprise: Breadth-First Graph Traversal on GPUs
SC, Austin, TX

TEACHING

2021 Spring CPE 517-A: Digital and Computer Systems Architecture
Stevens Institute of Technology (evaluation: 3.83/4)

2020 Fall CPE 360-A: Computational Data Structure and Algorithms
Stevens Institute of Technology (evaluation: 3.94/4)
Excellent teaching evaluation award

2020 Spring CPE 517-A: Digital and Computer Systems Architecture
Stevens Institute of Technology (evaluation: 3.13/4)

2019 Fall CPE 517-A: Digital and Computer Systems Architecture
Stevens Institute of Technology (evaluation: 2.7/4)

2019 Spring EECE 7110: High-Performance Computing on GPUs
University of Massachusetts Lowell

2018 Fall EECE4810/EECE5730: Operating Systems
University of Massachusetts Lowell

2018 Spring EECE 7110: High-Performance Computing on GPUs
University of Massachusetts Lowell

STUDENT ADVISING & MENTORING

PHD STUDENTS

2018 Anil Gaihre (2018.01 - Present)
Dissertation Topic: High-Performance Data Analytics Systems

2019 Santosh Pandey (2019.01 - Present)
Dissertation Topic: Machine Learning Expedited Computer Architecture Simulation

Shiyang Chen (2019.08 - Present)
Dissertation Topic: Re-thinking Machine Learning Models on Emerging Accelerators

2020 Scott Weitze (2020.08 - Present)
Dissertation Topic: Accelerating Data Mining With Graphics Processing Units

2021 Zhanfu Yang (2021.01 - Present)
Dissertation Topic: High-Performance Graph Mining Systems

2022 Lang Zhu (2022.09 - Present)
Dissertation Topic: Reconfigurable High-Performance Computing Systems

MASTER STUDENTS

2020 Zehui Xie (2020.01 - 2020.12) *Outstanding Master's Research Project Award*

Yufeng Liu (2020.01 - 2020.12)

Neel Haria, *Intern at Jabil* (2020.05 - 2020.12)

Ghaith Arar, *Intern at Jabil* (2020.05 - 2020.12)

Yupeng Cao & Yunxiang Yang (2020.01 - 2020.12)
ECE Honors Summer Research Program 3rd Place

2021 Runbang Hu (2021.09-2022.05)

2022 Lang Zhu (2022.03 - 2022.05)

UNDERGRADUATE STUDENTS

2018 Daniel Giger (2018.07 - 2019.08)

Topic: An Efficient Parallel Algorithm for Dominator Detection
ACM Undergraduate Poster Competition

2020 Jared Kantor & Chris Waldt (2020.05 - 2020.12)

Topic: [5G Phased Array Calibration](#)
Senior Design Semi-Final & Intern at Jabil

Shivam Sheth (2020.07 - 2021.05)

Topic: Graph Computing Assisted Latency Critical Job Scheduling on Supercomputers

2021 Jie Dai, *Intern at Jabil* (2021.06 - present)

Abdullah Hyder (2021.09 - present)

Grant Simmons (2021.09 - present)

Pridhvi Myneni (2021.09 - present)

2022 Christian O'Connell (2022.03 - 2022.05)

Justin Young (2022.03 - 2022.05)

Matthew Jaworski (2022.05 - present)

Kamen Kresnitchki (2022.05 - present)

DEFENSE/PROPOSAL COMMITTEE

Aug.2021 Proposal Committee for Ali Aghdaei

Aug.2021 Proposal Committee for Ying Zhang

Dec.2021 Defense Committee for Fangzhou Wang

Dec.2021 Proposal Committee for Yuandong Cyrus Liu

April.2022 Proposal Committee for Yifan Wang

June.2022 Proposal Committee for Xuting Tang

PROFESSIONAL ACTIVITIES

JOURNAL EDITORSHIP

2021 Associate Editor: Journal of BigData: Theory and Practice;

CONFERENCE ORGANIZER

- 2022 Program Co-Chair: The workshop on Graph Techniques for Adversarial Activity Analytics 2022 (*GTA³ 2022*)
- 2021 Program Co-Chair: The workshop on Graph Techniques for Adversarial Activity Analytics 2021 (*GTA³ 2021*)
- 2021 Session chair for High performance Graph Algorithms at SC '21
- 2021 Session chair for Cloud and Distributed Computing Exhibition Forum
- 2019 Session chair for Session 4: Scalable Graph Processing at HPDC '19

TECHNICAL PROGRAM COMMITTEE

- 2023 The IEEE International Parallel & Distributed Processing Symposium (*IPDPS* Program Committee Chair's Team)
- 2022 IEEE International Conference on Big Data (*BigData*)
- 2022 The International Conference for High Performance Computing, Networking, Storage, and Analysis (*SC*)
- 2022 The ACM Symposium on Principles and Practice of Parallel Programming (*PPoPP*)
- 2021 Best Paper Selection Committee at SC '21
- 2022 The IEEE International Parallel & Distributed Processing Symposium (*IPDPS*)
- 2021 The International Conference for High Performance Computing, Networking, Storage, and Analysis (*SC*)
- 2021 The IEEE International Parallel & Distributed Processing Symposium (*IPDPS*)
- 2021 The ACM International Symposium on High-Performance Parallel and Distributed Computing (*HPDC*)
- 2020 The International Conference for High Performance Computing, Networking, Storage, and Analysis (*SC*)
- 2020 The IEEE International Parallel & Distributed Processing Symposium (*IPDPS*)
- 2020 ACM International Symposium on High-Performance Parallel and Distributed Computing (*HPDC*)
- 2020 The IEEE International Conference on Distributed Computing Systems (*ICDCS*)
- 2020 SIAM Workshop on Combinatorial Scientific Computing (*CSC*)
- 2019 IEEE International Conference on Big Data (*BigData*)
- 2019 The ACM International Symposium on High-Performance Parallel and Distributed Computing (*HPDC*)
- 2018 The ACM International Symposium on High-Performance Parallel and Distributed Computing (*HPDC*)
- 2018 The IEEE International Parallel & Distributed Processing Symposium (*IPDPS*)

JOURNAL REVIEWER

- 2021 IEEE TC, IEEE TPDS, IEEE TCAD, IEEE ToCC
- 2020 IEEE TC, IEEE TPDS, IEEE TCAD, IEEE ToCC, IEEE TSC

2019 IEEE TC, IEEE TPDS, IEEE TOPC, IEEE TM, Elsevier Neurocomputing, IEEE TPPNA
2018 IEEE TC, IEEE TPDS, IEEE TOPC
2017 IEEE TC, IEEE TPDS

PANELIST

June.2022 National Science Foundation (NSF)
Mar.2022 National Science Foundation (NSF)
Mar.2022 National Science Foundation (NSF)
Feb.2021 National Science Foundation (NSF)
Jan.2021 National Science Foundation (NSF)
June.2020 National Science Foundation (NSF)
April.2020 National Science Foundation (NSF)
April.2019 National Science Foundation (NSF)
April.2018 National Science Foundation (NSF)

INTERNAL SERVICES

2023 Undergraduate recruitment and orientation committee
2023 Graduate recruitment and orientation committee
2022 Research computing committee, *The HPC infrastructure construction at Stevens Institute of Technology*
2022 Chair of Department Award Committee, *Department of Electrical & Computer Engineering, Stevens Institute of Technology*
2021 Strategic Planning Committee, *Department of Electrical & Computer Engineering, Stevens Institute of Technology*
2020 - 2021 Graduate Student Recruitment Committee, *Department of Electrical & Computer Engineering, Stevens Institute of Technology*
2019 - 2021 Master Student Advisor, *Department of Computer Science, Stevens Institute of Technology*
2020 - 2021 Faculty candidate interview, *CS/ECE, Stevens Institute of Technology*