

### RESEARCH INTEREST

High performance computing with a focus on the interaction among applications, numerical methods, data structures, algorithms, automatic performance tuning, and computer architectures. I am eager to pursue high performance sparse (multi-)linear algebra, solvers, and tensor decompositions for large-scale data analytics and domain applications on diverse computer architectures.

### EDUCATION

- 2013 – 2018 **Ph.D.**, *Georgia Institute of Technology*, Computational Science & Engineering, Advisor: Prof. Richard Vuduc.
- 2008 – 2013 **Doctor of Engineering**, *University of Chinese Academy of Sciences (UCAS)*, Computer Science, Advisors: Prof. Mingyu Chen and Guangming Tan.
- 2005 – 2008 **Bachelor of Sciences**, *Dalian University of Technology*, Computational Mathematics, In the Accelerated Student Program (2/180).

### PROFESSIONAL EXPERIENCE

- 2022- **Assistant Professor**, *Computer Science, North Carolina State University (NCSU)*, Raleigh, NC.
- 2021-2022 **Assistant Professor**, *Computer Science, College of William & Mary (W&M)*, Williamsburg, VA.
- 2018-2021 **Research Scientist**, *HPC Group, Pacific Northwest National Laboratory (PNNL)*, Richland, WA.
- 2014-2018 **Graduate Research Assistant**, *HPC Garage, Georgia Institute of Technology*, Atlanta, GA. Advisor: Prof. Richard Vuduc.
- 2016 **Summer Research Intern**, *IBM Thomas J. Watson Research Center*, Mentors: Dr. Jee Choi and Dr. Dong Chen.
- 2015 **Summer Research Intern**, *Intel Parallel Computing Research Lab*, Mentor: Dr. Mikhail Smelyanskiy.
- 2013-2014 **Graduate Research Assistant**, *HPC LAB, Georgia Institute of Technology*, Atlanta, GA. Advisor: Prof. David Bader.
- 2008-2013 **Graduate Research Assistant**, *CARCH, Institute of Computing Technology Chinese Academy of Sciences*, Atlanta, GA. Advisors: Prof. Guangming Tan and Prof. Mingyu Chen.

### HONORS AND AWARDS

- 2021 The 39th IEEE International Conference on Computer Design (ICCD'21) **Best Paper Award**
- 2019 Rising Stars for Women in Computational and Data Sciences. [\[Link\]](#)
- 2019 Principles and Practice of Parallel Programming. (PPoPP'19) **Best Paper Award Finalist**
- 2018 ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'18) **Best Student Paper Award.**[\[Link\]](#)[\[Link\]](#)
- 2018 SIAM ALA'18 Student Travel Grant.
- 2018 GaTech CoC Graduate Student Council Travel Grant.
- 2017 **IBM PhD Fellowship** for 2017-2018. [\[Link\]](#)
- 2017 Travel grant from ATIP Workshop, co-located with SC'17 [\[Link\]](#)
- 2017 Travel grant from IPAM for Big Data Meets Computation Workshop 2017 [\[Link\]](#)
- 2016 Selected students to attend IEEE WIE Women's Leadership Summit

- 2013 ZhuLiYueHua Award for the Excellent PhD Students of Chinese Academy of Sciences (Top 0.2%)
- 2011 Xia Peisu Scholarship of Institute of Computing Technology (Top 1%)
- 2011 Outstanding Research Assistant of the Computer Architecture Laboratory at UCAS
- 2010 Outstanding Student of the Computer Architecture Laboratory at UCAS

## PUBLICATIONS

- SC22 Zhen Du, **Jiajia Li**, Yinshan Wang, Xueqi Li, Guangming Tan, Ninghui Sun. AlphaSparse: Generating High Performance SpMV Codes Directly from Sparse Matrices. The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC). 2022.
- CLUSTER22 Zheng Miao, **Jiajia Li**, Jon Calhoun, Rong Ge. BALA-CPD: BALanced and Asynchronous Distributed Tensor Decomposition. 24th IEEE Cluster. 2022.
- HPCA22 Cheng Tan, Nicolas Bohm Agostini, Tong Geng, Chenhao Xie, **Jiajia Li**, Ang Li, Kevin Barker, Antonino Tumeo. DRIPS: Dynamic Rebalancing of Pipelined Streaming Applications CGRAs. The 28th IEEE International Symposium on High-Performance Computer Architecture (HPCA). 2022.
- PPoPP22 Zhen Xie, Jie Liu, Yuchen Ma, **Jiajia Li**, Dong Li. LB-HM: Load Balance-Aware Data Placement on Heterogeneous Memory for Task-Parallel HPC Applications. Principles and Practice of Parallel Programming (PPoPP). 2022. (Poster, Accepted)
- LLVM-HPC21 Ruiqin Tian, Luanzheng Guo, **Jiajia Li**, Bin Ren, Gokcen Kestor. A High Performance Sparse Tensor Algebra Compiler in MLIR. The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), the Seventh Annual Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC). 2021.
- ICCD21 Cheng Tan, Tong Geng, Chenhao Xie, Nicolas Bohm Agostini, **Jiajia Li**, Ang Li, Kevin Barker, Antonino Tumeo. DynPaC: Coarse-Grained, Dynamic, and Partially Reconfigurable Array for Streaming Applications. The 39th IEEE International Conference on Computer Design (ICCD). 2021. (AR=24.4%) (**Best Paper Award**)
- HPEC21 Tong Geng, Chunshu Wu, Cheng Tan, Chenhao Xie, Anqi Guo, Pouya Haghi, Sarah Yuan He, **Jiajia Li**, Martin Herbordt, Ang Li. A Survey: Handling Irregularities in Neural Network Acceleration with FPGAs. IEEE High Performance Extreme Computing Conference (HPEC). 2021.
- ACDA21 Shruti Shivakumar, **Jiajia Li**, Ramakrishnan Kannan, Srinivas Aluru. Efficient Parallel Sparse Symmetric Tucker Decomposition for High-Order Tensors. SIAM Conference on Applied and Computational Discrete Algorithms (ACDA). 2021.
- ICS21 Jiawen Liu, Dong Li, Roberto Gioiosa, **Jiajia Li**. Athena: High-Performance Sparse Tensor Contraction Sequence on Heterogeneous Memory. International Conference on Supercomputing (ICS). 2021. (AR=25%)
- PPoPP21 Jiawen Liu, Jie Ren, Roberto Gioiosa, Dong Li, **Jiajia Li**. Sparta: High-Performance, Element-Wise Sparse Tensor Contraction on Heterogeneous Memory. Principles and Practice of Parallel Programming (PPoPP). 2021. (AR=21%)
- NVMW21 Jiawen Liu, Jie Ren, Roberto Gioiosa, Dong Li, **Jiajia Li**. Sparta: High-Performance, Element-Wise Sparse Tensor Contraction on Heterogeneous Memory. 12th Annual Non-Volatile Memories Workshop (NVMW). 2021.
- IISWC20 **Jiajia Li**, Mahesh Lakshminarasimhan, Xiaolong Wu, Ang Li, Catherine Olschanowsky, Kevin Barker. A Sparse Tensor Benchmark Suite for CPUs and GPUs. IEEE International Symposium on Workload Characterization (IISWC). 2020.
- MLSH20 Xiaolong Wu, Yang Yi, Dave (Jing) Tian, **Jiajia Li**. Generic, Sparse Tensor Core for Neural Networks. 1st International Workshop on Machine Learning for Software Hardware Co-Design (MLSH), in conjunction with the 29th International Conference on Parallel Architectures and Compilation Techniques (PACT). 2020.
- ToPC Eric Hein, Srinivas Eswar, Abdurrahman Yasar, **Jiajia Li**, Jeffrey S. Young, Tom Conte, Umit V. Catalyurek, Rich Vuduc, Jason Riedy, Bora Ucar. Programming Strategies for Irregular Algorithms on the Emu Chick. ACM Transactions on Parallel Computing. 2020.

- SC20-Poster Zheng Miao, Jon C. Calhoun, Rong Ge, **Jiajia Li**. Sparsity-Aware Distributed Tensor Decomposition. ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC). 2020. (Poster)
- WHPC20 Ruiqin Tian, **Jiajia Li**, Bin Ren, Gokcen Kestor. High-Performance Sparse Tensor Algebra Compiler. Women in High Performance Computing Workshop (WHPC), in conjunction with ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC). 2020. (Poster)
- HPEC20 Jesun Sahariar Firoz, Ang Li, **Jiajia Li**, Kevin Barker. On the Feasibility of Using Reduced-Precision Tensor Core Operations for Graph Analytics. IEEE High Performance Extreme Computing Conference (HPEC). 2020.
- PPoPP20 **Jiajia Li**, Mahesh Lakshminarasimhan, Xiaolong Wu, Ang Li, Catherine Olschanowsky, Kevin Barker. A Parallel Sparse Tensor Benchmark Suite on CPUs and GPUs. Principles and Practice of Parallel Programming (PPoPP). 2020. (Poster)
- TPDS Ang Li, Shuaiwen Leon Song, Jieyang Chen, **Jiajia Li**, Xu Liu, Nathan Tallent, Kevin Barker. Evaluating Modern GPU Interconnect: PCIe, NVLink, NV-SLI, NVSwitch and GPUDirect. IEEE Transactions on Parallel and Distributed Systems. 2020.
- SC19 Israt Nisa, **Jiajia Li**, Aravind Sukumaran-Rajam, Prashant Rawat, Sriram Krishnamoorthy, P. (Saday) Sadayappan. An Efficient Mixed-Mode Representation of Sparse Tensors. ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC). 2019.
- ICS19 **Jiajia Li**, Bora Ucar, Umit Catalyurek, Jimeng Sun, Kevin Barker, Richard Vuduc. Efficient and Effective Sparse Tensor Reordering. International Conference on Supercomputing (ICS). 2019.
- THPC **Jiajia Li**, Yuchen Ma, Xiaolong Wu, Ang Li, Kevin Barker. PASTA: A Parallel Sparse Tensor Algorithm Benchmark Suite. CCF Transactions on High Performance Computing. 2019.
- ParCo Jeffrey S. Young, Eric Hein, Srinivas Eswar, Patrick Lavin, **Jiajia Li**, Jason Riedy, Richard Vuduc, Thomas M. Conte. A Microbenchmark Characterization of the Emu Chick. Journal of Parallel Computing. 2019.
- PPoPP19 Ke Meng, **Jiajia Li**, Guangming Tan. A Pattern Based Algorithmic Autotuner for Graph Processing on GPUs. Principles and Practice of Parallel Programming (PPoPP). 2019. (**Best Paper Award Finalist**)
- IPDPS19 Israt Nisa, **Jiajia Li**, Aravind Sukumaran Rajam, Richard Vuduc, P. (Saday) Sadayappan. Load-balanced sparse MTTKRP on GPUs. IEEE International Parallel & Distributed Processing Symposium (IPDPS). 2019.
- TOPC Junhong Liu, Guangming Tan, Yulong Luo, **Jiajia Li**, Zeyao Mo, Ninghui Sun. An Autotuning Protocol to Rapidly Build Autotuners. ACM Transactions on Parallel Computing. 2019.
- THESIS **Jiajia Li**. Scalable Tensor Decompositions in High Performance Computing Environments. PhD Dissertation. Georgia Institute of Technology, Atlanta, GA, USA. July 2018.
- SC18 **Jiajia Li**, Jimeng Sun, Richard Vuduc. HiCOO: Hierarchical Storage of Sparse Tensors. ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC). 2018. (**Best Student Paper Award**)
- JPDC Yuchen Ma, **Jiajia Li**, Xiaolong Wu, Chenggang Yan, Jimeng Sun, Richard Vuduc. Optimizing Sparse Tensor Times Matrix on GPUs. Journal of Parallel and Distributed Computing (Special Issue on Systems for Learning, Inferencing, and Discovering). 2018.
- IPDPSW18 Eric Hein, Tom Conte, Jeffrey Young, Srinivas Eswar, **Jiajia Li**, Patrick Lavin, Richard Vuduc, Jason Riedy. An Initial Characterization of the Emu Chick. 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). 2018.
- PPoPP18 Yue Zhao, **Jiajia Li**, Chunhua Liao, Xipeng Shen. Bridging the Gap between Deep Learning and Sparse Matrix Format Selection. 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP). 2018.
- TOMS Guangming Tan, Junhong Liu, **Jiajia Li**. Design and Implementation of Adaptive SpMV Library for Multicore and Manycore Architecture. ACM Transactions on Mathematical Software. 2018.
- IPDPS17 **Jiajia Li**, Jee Choi, Ioakeim Perros, Jimeng Sun, Richard Vuduc. Model-Driven Sparse CP Decomposition for Higher-Order Tensors. 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS). 2017.

- PACT17 Yue Zhao, **Jiajia Li**, Chunhua Liao, Xipeng Shen. Bridging the Gap between Deep Learning and Sparse Matrix Format Selection. The 26th International Conference on Parallel Architectures and Compilation Techniques (PACT). 2017. (Poster)
- PPoPP17 Xiuxia Zhang, Guangming Tan, Shuangbai Xue, **Jiajia Li**, Keren Zhou, Mingyu Chen. Understanding the GPU Microarchitecture to Achieve Bare-Metal Performance Tuning (PPoPP). 2017. (**Best Artifact Award**)
- SC16-IA3 **Jiajia Li**, Yuchen Ma, Chenggang Yan, Richard Vuduc. Optimizing Sparse Tensor Times Matrix on multi-core and many-core architectures. The sixth Workshop on Irregular Applications: Architectures and Algorithms (IA<sup>3</sup>), co-located with SC. 2016.
- SC15 **Jiajia Li**, Casey Battaglini, Ioakeim Perros, Jimeng Sun, Richard Vuduc. An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply. The International Conference for High Performance Computing, Networking, Storage and Analysis (SC) 2015
- SC15-EduHPC B. Neelima, **Jiajia Li**. Introducing high performance computing concepts into engineering undergraduate curriculum: a success story. Proceedings of the Workshop on Education for High-Performance Computing (EduHPC), co-located with SC. 2015
- THESIS **Jiajia Li**. Research on Sparse Matrix Vector Multiplication Auto-tuning Method. PhD Thesis. The University of Chinese Academy of Sciences, Beijing, China. July, 2013
- PLDI13 **Jiajia Li**, Guangming Tan, Mingyu Chen, Ninghui Sun. SMAT: An Input Adaptive Auto-Tuner for Sparse Matrix-Vector Multiplication. Programming Language Design and Implementation (PLDI) 2013
- ICS12 **Jiajia Li**, Xingjian Li, Guangming Tan, Mingyu Chen, Ninghui Sun. An Optimized Large-Scale Hybrid DGEMM Design for CPUs and ATI GPUs. International Conference on Supercomputing (ICS) 2012
- JCRD **Jiajia Li**, Xiuxia Zhang, Guangming Tan, Mingyu Chen. Study of Choosing the Best Storage Format of Sparse Matrix Vector Multiplication, Journal of Computer Research and Development. (IN CHINESE) 2012
- ICPADS10 **Jiajia Li**, Guangming Tan, Mingyu Chen. Automatically Tuned Dynamic Programming with an Algorithm-by-Blocks. The 16th International Conference on Parallel and Distributed Systems (ICPADS) 2010

## GRANTS

*As Principal Investigator*

- HiParTI **Jiajia Li** (PI), Ang Li, Ajay Panyala (Key Staffs). Application-Algorithm-Architecture Co-Design for Large-Scale, Sparse Tensor/Matrix Methods. PNNL Data-Model Convergence LDRD. Estimated Award: \$630K. Aug. 2019 - Oct. 2022.

*As Co-PI*

- SHARWK Stephen J. Young (PI), Sinan G. Aksoy, **Jiajia Li**. SHARWK: Scalable Hypergraph Analysis Via Random Walk Kernels. DOE EXPRESS: Randomized Algorithms for Extreme-Scale Science. Estimated Award: \$800K. 2021 - 2023.

*As Key Staff*

- CENATE Kevin Barker (PI), **Jiajia Li** (Key Staff of Tensor Task), et al. The Center for Advanced Technology Evaluation. DOE ASCR No. 66150. Oct. 2019 - Oct. 2022.
- NWChemEx Thom Dunning Jr, **Jiajia Li** (Key Staff of TAMM task, led by Sriram Krishnamoorthy), et al. Tackling Chemical, Materials, and Biomolecular Challenges in Exascale. DOE ECP Project. Oct. 2016 - Oct. 2023.
- ExaSGD Henry Huang (PI), **Jiajia Li** (Key Staff in GPU thrust, led by Slaven Peles), et al. Optimizing Stochastic Grid Dynamics at Exascale. DOE ECP Project. Oct. 2019 - Oct. 2020.
- CFA Ang Li (PI), Xiu Yang, Leon Song, **Jiajia Li** (Key Staff), Jesun Firoz. Computation-Flow-Architecture: Designing Non-Von-Neumann Architecture for Future Data-Centric Computing. PNNL Data-Model Convergence LDRD. Estimated Award: \$620K. Aug. 2019 - Oct. 2022.
- S-BLAS Ang Li (PI), **Jiajia Li** (Key Staff), Jesun Firoz. Designing Highly Scalable Sparse-BLAS Kernels for Complex Modern HPC Architectures. PNNL High Performance Data Analytics Program. Awarded: \$170K. Mar. 2019 - Mar. 2020.

---

## SOFTWARE

- HiParTI **A Hierarchical Parallel Tensor Infrastructure**, [\[Link\]](#).  
PASTA **A Parallel Sparse Tensor Algorithm Benchmark Suite**, [\[Link\]](#).  
ParTI! **A Parallel Tensor Infrastructure for Data Analysis**, [\[Link\]](#).  
AdaTM **Adaptive Tensor Memoization Algorithm for CP Decomposition**, [\[Link\]](#).  
InTensLi **Input-Adaptive and In-Place Dense Tensor-Times-Matrix Multiply**, [\[Link\]](#).  
SMAT **SpMV Auto-Tuner**, [\[Link\]](#).  
HDGEMM **A Hybrid DGEMM library on a Heterogeneous CPU-AMD GPU Architecture**, [\[Link\]](#).

---

## INVITED TALKS AND PRESENTATIONS

- 2021 **How to Do Successful Career Change**, (*Invited*). *ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'21) Women in HPC workshop (WHPC)*, Online, Nov, 2021.  
**High-Performance Sparse Tensor Operations in HiParTI Library**, (*Invited*). *CCF HPC China 2021*, Online, Oct, 2021.  
**Sparsity-Aware Distributed Tensor Decomposition**, (*Invited*). *2021 SIAM Conference on Applied Linear Algebra (SIAM-LA21)*, New Orleans, Louisiana, May, 2021.  
**High-Performance Sparse Tensor Operations in HiParTI Library**, (*Invited*). *SPCL\_Bcast* organized by *Scalable Parallel Computing Laboratory (SPCL)*, *ETH*, Online, Feb 25, 2021.  
**Tensor Methods and Emerging Applications to the Physical and Data Sciences**, (*Invited*). *2021 IPAM Long Program (TM2021)*, Los Angeles, CA, Mar-Jun, 2021.  
**Sparta: High-Performance, Element-Wise Sparse Tensor Contraction on Heterogeneous Memory**, (*Invited*). *2021 SIAM Conference on Computational Science and Engineering (SIAM-CSE21)*, Fort Worth, TX, March, 2021.
- 2020 **PASTA: A Parallel Sparse Tensor Algorithm Benchmark Suite**, (*Invited*). *SIAM Conference on Parallel Processing for Scientific Computing (SIAM-PP20)*, Seattle, WA, February, 2020.
- 2019 **Sparse Tensor Library based on HiCOO Format**, (*Invited*). *AI and Tensor Factorization for Physical, Chemical, and Biological Systems Workshop*, Santa Fe, NM, September, 2019.  
**Efficient and Effective Sparse Tensor Reordering**, (*Paper*). *International Conference on Supercomputing (ICS19)*, Phoenix, AZ, June, 2019.  
**Sparse Tensor Algebra and its Relations to Matrix and Graph Problems**, (*Invited*). *SIAM Conference on Computational Science and Engineering (SIAM-CSE19)*, Spokane, WA, February, 2019.  
**A Sparse Tensor Format and a Benchmark Suite**, (*Invited*). *Workshop on Compiler Techniques for Sparse Tensor Algebra*, Cambridge, MA, January, 2019.
- 2018 **HiCOO: Hierarchical Storage of Sparse Tensors**, (*Paper*). *ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC18)*, Dallas, TX, November, 2018.  
**Scalable Tensor Decompositions in High Performance Computing Environments**, *Ph.D. Thesis Defense*, Atlanta, GA, July, 2018.  
**ParTI!: A Parallel Tensor Infrastructure on HPC Platforms**, (*Invited*). *The 8th edition of the multidisciplinary conference on ThRee-way methods In Chemistry And Psychology (TRICAP18)*, Angel Fire, NM, June, 2018.  
**Parallel Sparse Tensor Decompositions using HiCOO Format**, (*Invited*). *SIAM Conference on Applied Linear Algebra (SIAM-ALA18)*, Hong Kong, May, 2018.  
**HiCOO: Hierarchical Storage of Sparse Tensors**, (*Invited*). *SIAM Conference on Parallel Processing for Scientific Computing (SIAM-PP18)*, Tokyo, Japan, March, 2018.
- 2017 **HiCOO: A Hierarchical Sparse Tensor Format for Tensor Decompositions**, (*Poster*). *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC17)*, Denver, CO, November, 2017..

- ParTI!: A Parallel Tensor Infrastructure**, *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC17) ATIP workshop on International Exascale and Next-Generation Computing Programs*, Denver, CO, November, 2017.
- Model-Driven Sparse CP Decomposition for Higher-Order Tensors**, (Invited). *SIAM Annual Meeting (SIAM-AN17)*, Pittsburgh, PA, July, 2017.
- Model-Driven Sparse CP Decomposition for Higher-Order Tensors**, (Paper). *The 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS17)*, Orlando, FL, May, 2017.
- Non-negative Sparse Tensor Decomposition on Distributed Memory Systems**, (Poster). *SIAM Conference on Computational Science and Engineering (SIAM-CSE17)*, Atlanta, GA, February, 2017.
- ParTI!: A Parallel Tensor Infrastructure for Data Analysis**, (Poster). *Institute for Pure & Applied Mathematics Big Data Meets Computation Workshop (IPAM-DMC17)*, Los Angeles, CA, January, 2017.
- 2016 **Optimizing Sparse Tensor Times Matrix on Multi-core and Many-core Architectures**, (Paper). *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC16), the sixth Workshop on Irregular Applications: Architectures and Algorithms (IA<sup>3</sup>)*, Salt Lake City, UT, November, 2016.
- Model-driven Sparse CP Decomposition for High-Order Tensors**, (Poster). *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC16), Women in HPC: Diversifying the HPC Community (WHPC) Workshop*, Salt Lake City, UT, November, 2016.
- ParTI!: A Parallel Tensor Infrastructure for Data Analysis**, (Paper). *The Conference on Neural Information Processing Systems, the Workshop on Learning with Tensors: Why Now and How? (NeurIPS16 Tensor-Learn)*, Barcelona, Spain, December, 2016.
- An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply**, (Invited). *SIAM Conference on Parallel Processing for Scientific Computing (SIAM-PP16)*, Paris, France, April, 2016.
- 2015 **An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply**, (Paper). *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, Austin, TX, November, 2015.
- Introducing high performance computing concepts into engineering undergraduate curriculum: a success story**, (Paper). *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC15), the Workshop on Education for High-Performance Computing (EduHPC)*, Austin, TX, November, 2015.
- 2013 **Research on Sparse Matrix Vector Multiplication Auto-tuning Method**, *Ph.D. Thesis Defense*, Beijing, China, July, 2013.
- SMAT: An Input Adaptive Auto-Tuner for Sparse Matrix-Vector Multiplication**, (Paper). *Programming Language Design and Implementation (PLDI13)*, Seattle, WA, June, 2013.
- 2012 **An Optimized Large-Scale Hybrid DGEMM Design for CPUs and ATI GPUs**, (Paper). *International Conference on Supercomputing (ICS12)*, Venice, Italy, June, 2012.
- 2010 **Automatically Tuned Dynamic Programming with an Algorithm-by-Blocks**, (Paper). *The 16th International Conference on Parallel and Distributed Systems (ICPADS10)*, Shanghai, China, December, 2010.

## ORGANIZING ACTIVITIES

- 2022 Program Co-Chair of the 26th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS'22)
- Program Chair of the Emerging Parallel and Distributed Runtime Systems and Middleware Workshop (IPDRM), held in conjunction with IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC'22) <https://ipdrm.github.io>
- 2021 Finance Chair of International Conference on Supercomputing (ICS'21) <https://ics21.github.io/>
- Session Chair of ACM Capital Region Celebration of Women in Computing (CAPWIC'21) <https://capwic.org/conferences/2021/>

- Artifact Evaluation Chair of the 2021 edition of the International Conference on Languages Compilers, Tools and Theory of Embedded Systems (LCTES'21) <https://pldi21.sigplan.org/home/LCTES-2021>
- 2020 Publicity Chair of International Conference on Parallel Architectures and Compilation Techniques (PACT'20) <https://pact20.cc.gatech.edu>
- Finance and Session Chair of the Principles and Practice of Parallel Programming 2020 (PPoPP'20) <https://ppopp20.sigplan.org>
- Publicity Chair of the 17th Annual IFIP International Conference on Network and Parallel Computing (NPC'20) <http://www.ncic.ac.cn/npc2020>
- Publicity Chair of the Emerging Parallel and Distributed Runtime Systems and Middleware Workshop (IPDRM), held in conjunction with IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC'20) <https://ipdrm.github.io>
- 2019 Proceeding Chair of the Emerging Parallel and Distributed Runtime Systems and Middleware Workshop (IPDRM), held in conjunction with IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC'19) <https://ipdrm.github.io>
- Co-Chair of the 25th International European Conference on Parallel and Distributed Computing (EuroPar'19) <https://2019.euro-par.org>
- Web and Social Media Chair of International Conference on Parallel Architectures and Compilation Techniques (PACT'19) <https://pactconf.org>
- Co-Chair of The First International Workshop on the Intersection of High Performance Computing and Machine Learning (HPCaML'19), held in conjunction with International Symposium on Code Generation and Optimization (CGO'19) <http://hpc.pnl.gov/hpcaml19/>
- Co-Organizer of SIAM Conference on Computational Science and Engineering (SIAM CSE'19) Minisymposium "High Performance Sparse Matrix, Tensor, and Graph Kernels" <http://hpc.pnl.gov/siamcse19/>

## PEER REVIEW ACTIVITIES

- 2022 PC Member of the International Conference on Supercomputing 2022 (ICS'22)
- PC Member of ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'22) in "Machine Learning with HPC" area
- PC Member of the 36th IEEE International Parallel & Distributed Processing Symposium (IPDPS'22)
- PC Member of ISC High Performance 2022 (ISC'22)
- PC Member of IEEE CLUSTER 2022
- PC Member of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'22)
- PC Member of the SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP'22)
- 2021 PC Member of ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'21) in Algorithm area
- PC member for IEEE Transactions on Parallel and Distributed System (TPDS), special session on "Parallel and Distributed Computing Techniques for AI, ML and DL" (2nd Edition)
- PC Member of the 2021 edition of the International Conference on Languages Compilers, Tools and Theory of Embedded Systems (LCTES'21) <https://pldi21.sigplan.org/home/LCTES-2021>
- PC Member of the International Conference on Parallel Processing 2021 (ICPP'21)
- PC Member of IEEE International Conference on Cluster Computing (Cluster'21) in Algorithm track
- Review Editor on Journal of Applied Sciences
- PC Member of the International Conference on Supercomputing 2021 (ICS'21) (Review 10 papers)
- PC Member of the 41st IEEE International Conference on Distributed Computing Systems (ICDCS'21) in the Cloud Computing and Data Centers and Distributed Fault Tolerance and Dependability tracks (Review 10 papers)

- Technical Papers Program Committee Member of ISC High Performance (ISC'21) in Machine Learning, AI, and Emerging Technologies track (Review 5 papers)
- 2020 PC Member of the 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS'20) in Algorithms track (Review 9 papers)
- PC Member of ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'20) in Programming Systems area (Review 7 papers)
- PC member for IEEE Transactions on Parallel and Distributed System (TPDS), special session on "Parallel and Distributed Computing Techniques for AI, ML and DL" (Review 14 papers)
- PC Member of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'20) (Review 10 papers)
- PC Member of the International Conference on Parallel Processing 2020 (ICPP'20) (Review 5 papers)
- Review Editor on the Editorial Board of Mathematics of Computation and Data Science, Frontiers in Applied Mathematics and Statistics (Review 4 papers)
- Review Editor on the Editorial Board of Algorithms (Review 1 paper)
- PC Member of the 29th International Conference on Parallel Architectures and Compilation Techniques (PACT'20) ACM Student Research Competition (SRC) (Review 2 posters)
- PC member of the 6th Workshop on Machine Learning in High Performance Computing Environments (MLHPC'20), held in conjunction with ACM/IEEE International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'20) (Review 1 paper)
- PC Member of the 17th Annual IFIP International Conference on Network and Parallel Computing (NPC'20) (Review 7 papers)
- Research Poster Committee Member of the ISC High Performance 2020 (ISC'20)
- PC Member of the MLBench workshop, held in conjunction with 2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'20) (Review 1 paper)
- 2019 Technical Program Committee Member of the Algorithms track of the 26th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'19) (Review 2 papers)
- Travel Grant PC Member of the ACM Symposium on High-Performance Parallel and Distributed Computing 2019 (HPDC'19) (Review 29 applications)
- PC Member of the Workshop on Tensor Methods for Emerging Data Science Challenges (TMEDSC), held in conjunction with the 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'19) (Review 2 papers)
- PC member of the 25th International European Conference on Parallel and Distributed Computing (EuroPar'19). (Review 4 papers)
- PC member of the First International Workshop on the Intersection of High Performance Computing and Machine Learning (HPCaML'19), held in conjunction with International Symposium on Code Generation and Optimization (CGO'19) <http://hpc.pnl.gov/hpcaml19/>. (Review 1 paper)
- 2017-2019 PC Member of the International Conference on Advanced Engineering Computing and Applications in Sciences (ADVCOMP'19) (Review 1 paper)
- 2013-2019 PC member of Parallel Algorithm Track of "National Annual Conference on High Performance Computing (HPC China)". (Review 41 papers)
- 2018 External PC member of The 32nd ACM International Conference on Supercomputing (ICS'18). (Review 6 papers)
- PC Member of the 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS'18). (Review 6 papers)
- 2017 PC member of Student Research Competition (SRC) of The 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'18). (Review 3 papers)



2014-Now Reviewer of the Transactions on Parallel and Distributed Systems(TPDS) (11 papers), IEEE Transactions on Neural Networks and Learning Systems(TNNLS) (3 papers), ACM Transactions on Knowledge Discovery from Data (2 papers), Transactions on Knowledge and Data Engineering (TKDE) (1 paper), Transactions on Computers (1 paper), ACM Transactions on Mathematical Software (TOMS) (1 paper), Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD) (1 paper), IEEE Access (1 paper), Journal of Parallel and Distributed Computing (JPDC) (2 papers), ACM Transactions on Parallel Computing (TOPC) (3 papers), Parallel Computing Journal (ParCo) (4 papers), CCF Transactions on High Performance Computing (THPC) (8 papers), Frontiers of Computer Science (1 paper), Algorithmica Journal (2 papers), Journal of Low Power Electronics and Applications (1 paper), Songklanakarin Journal of Science and Technology (SJST) (3 papers), International Journal of High Performance Computing Applications (2 papers), the 47th International Conference on Parallel Processing (ICPP'18) (1 paper), the 48th International Conference on Parallel Processing (ICPP'19) (1 paper), the 21st IEEE International Conference on Parallel and Distributed Systems (ICPADS'15) (2 papers), the 2020 International Conference on Memory Systems (MEMSYS'20) (1 paper), the 29th ACM International Conference on Information and Knowledge Management (CIKM'20) (1 paper), ACM Transactions on Embedded Computing Systems (TECS) (1 paper), ACM Transactions on Architecture and Code Optimization (TACO) (1 paper)

## OTHER ACTIVITIES

2021-Now Colloquium Committee Member at W&M.  
2021-Now Awards and Prizes Committee Member at W&M.  
2019-2020 STEM Ambassador of the Office of STEM Education at PNNL.  
2014-2018 Organizer of Hot CSE seminar, a PhD academic seminar in GT CSE.  
2013-2017 Volunteer librarian of Repetitive Stress Injury (RSI) Lending Library of GT College of Computing.  
Apr 2016 Volunteer judge for Undergraduate Research Opportunities Program 11th Annual Undergraduate Research Spring Symposium.  
2013-2015 Volunteer reviewer of "President's Undergraduate Research Awards (PUMA)" and "National Center for Women & IT (NCWIT) Award"  
2008-2009 Vice Minister of Academic Study of Student Union at UCAS

## MENTORING EXPERIENCE

Ph.D. Tianjie (Horio) Hu, William & Mary, USA. 2021-Now  
Students Yuchen (Sam) Ma, William & Mary, USA. 2021-Now  
Zhen Du, Institute of Computing Technology Chinese Academy of Sciences, China. 2021-Now  
Shruti Shivakumar, Georgia Institute of Technology, USA. 2020-Now  
Zheng Miao, Clemson University, USA. 2019-Now  
Jiawen Liu, University of California Merced, USA. 2020-2021  
Israt Nisa, Ohio State University, USA. 2018-2020  
Ke Meng, Institute of Computing Technology Chinese Academy of Sciences, China. 2018-2019  
Srinivas Eswar, Georgia Institute of Technology, USA. 2017-2018  
Yue Zhao, North Carolina State University, USA. 2017-2018  
Junhong Liu, Institute of Computing Technology Chinese Academy of Sciences, China. 2016  
Xiuxia Zhang, Institute of Computing Technology Chinese Academy of Sciences, China. 2016  
Subramanya Dulloor, Georgia Institute of Technology, USA. 2015  
Masters Elijah Mas, William & Mary, USA. 2021-Now  
Junghyun Kim, Georgia Institute of Technology, USA. 2017-2018  
Undergraduates Zengmin Xiao, William & Mary, USA. 2022-Now  
Yuchen Ma, Hangzhou Dianzi University, China. 2016-2018  
Nicholas Liu, Georgia Institute of Technology, USA. 2017

## TEACHING EXPERIENCE

- Spring 2022 Instructor of “[Undergrad] Algorithms (CSCI 303-01)” at W&M
- Fall 2021 Instructor of “[Grad] Accelerating Deep Learning (CSCI 780-02)” at W&M
- Spring 2017 Teaching Assistant of “Intro to High-Performance Computing (OMSCS) (CSE 6220)” at Georgia Tech
- Fall 2014 Teaching Assistant of “High-Performance Computing: Tools and Applications (CSE 6230)” at Georgia Tech
- May 2012 Teaching Assistant of “Parallel Computer Architecture” class of Dragonstar Project at University of Chinese Academy of Sciences
- Jun 2012 Instructor of “Parallel Computing on GPUs using CUDA” Training at Sun Yat-sen University

## PROFESSIONAL MEMBERSHIPS

- Member of the Association of Computing Machinery (ACM)
- Member of the Association of Computing Machinery (ACM-Women)
- Member of the Society for Industrial and Applied Mathematics (SIAM) and SIAM Activity Group on Supercomputing
- Member of the Institute of Electrical and Electronics Engineers (IEEE)