Apan Qasem

Curriculum Vitae

apan@txstate.edu
userweb.cs.txstate.edu/~aq10
Apan-Qasem
Apan-Qasem
apanqasem
ApanQasem

Research Interests

Compilers and runtimes for heterogeneous computing; Performance modeling; Machine Learning for system software; Energy-efficient computing

Education

Jan 2008 **Rice University**, *Ph.D. in Computer Science*
Dissertation: *Automatic Tuning of Scientific Applications*
Advisors: Ken Kennedy and John Mellor-Crummey

May 2001 **Florida State University**, *M.S. in Computer Science*
Thesis: *Using a Swap Instruction to Coalesce Loads and Stores*
Advisor: David Whalley

May 1998 **Ohio Wesleyan University**, *B.A. in Computer Science and Economics*

Appointments

2018-present **Associate Chair**, *Computer Science Department*, Texas State University
2013-present **Associate Professor**, *Computer Science Department*, Texas State University
2016-2019 **Visiting Scholar**, *AMD Research, AMD, Austin*
2015 **Visiting Professor**, *School of Informatics, University of Edinburgh*
2012-2015 **Special Member, Graduate Faculty**, *University of Texas at San Antonio*
2007-2013 **Assistant Professor**, *Texas State University*

Honors and Awards

2013 NSF CAREER Award
2008 IBM Faculty Award
2012 Presidential Distinction Award, Excellence in Teaching, Texas State University
2013 College Achievement Award, Excellence in Research, Texas State University
2014 College Achievement Award, Excellence in Service, Texas State University
2015 College Achievement Award, Excellence in Teaching, Texas State University
2021 Best Paper, International Symposium on Cluster, Cloud and Internet Computing (CCGrid)
2019 Best Paper, Workshop on Education for High-Performance Computing (EduHPC)
2017 Best Paper, High Performance Computing and Communications (HPCC)
2010 Best Paper, International Conference on Network and Parallel Computing (NPC)
2015 Alpha Chi Favorite Professor, Alpha Chi National Honor Society (multiple)
2010 Intel Microgrant Award for Parallelism in the Classroom
2001 Phi Kappa Phi

Grants (▲ = active, ▲ = inter-disciplinary, ▲ = BPC; sorted by end date; most recent first)

▲ ▲ [NSF] *Expanding AI Curriculum and Infrastructure at Texas State University to Advance Inter-disciplinary Research and Grow a Diverse AI Workforce*, NSF Expand AI, $400K, co-PI, 2023-2024

▲ ▲ [DHHS] *NUEVA: Nutrition for Underserved Elderly Via Application*, Department of Health and Human Services (DHHS), $2.7M, co-PI, 2022-2027

▲ [THECB] *Accelerating Credentials Planning Grant*, Texas Higher Education Coordinating Board, $50K, co-PI, 2022

▲ ▲ [THECB] *Zeus Data Science Pathways*, Texas Higher Education Coordinating Board, $1.45M, co-PI, 2021-2023

▲ [AMD] *Investigating Performance Anomalies in Heterogeneous Memory Architectures*, AMD, $90K, PI, 2022- (Equipment Grant)

▲ [Intel] *Autotuning Hybrid Applications in Intel OneAPI*, Intel Donation, $22K, PI, 2022-2023

▲ [TXST] *High Performance Computing GPU Server for Machine Learning Research*, Materials Application Research Center (MARC), $49K, co-PI, 2022-

▲ ▲ [NSF] *Widening the CI Workforce On-ramp by Exposing Undergraduates to Heterogeneous Computing*, NSF OAC, $225K, PI, 2018-2023

▲ ▲ [CIC] *Data Collection for Inclusive Computing*, Center for Inclusive Computing, Northeastern University, $60K, co-PI, 2018-2023

[NSF] *Autotuning for Multicore and Manycore Architectures: An Enhanced Feedback-driven Approach*, NSF CAREER, $550K, PI, 2013-2018

▲ [NVIDIA] *Machine Learning Compilers for the GPU*, Equipment Grant, NVIDIA, > $30K (multiple grants since 2010), PI, 2017-

▲ [HAL] *Developing High School Students' Curiosity in Science and Technology in Capstone Design Day*, Halliburton Foundation, co-PI, $8.6K, 2014-2015

▲ [ONA] *TexasMusicViz: Music, Data And Storytelling*, Online News Association (ONA) Challenge Fund, co-PI, $35K, 2014-2015

[NSF] *II-NEW: Shared High Performance Data Center*, NSF CRI, $375K, co-PI, 2013-2016

[IBM] *High-performance QAP Solvers on the POWER8*, IBM, $22.5K, PI, 2017-2018

[NSF] *Preparing Computer Science Students for the Multicore Era: Teaching Parallel Computing in the Undergraduate Curriculum Early and Often*, NSF TUES, $149K, PI, 2012-2015

[SRC] *Power Consumption Based Multicore Task Scheduling and Load Balancing*, Semiconductor Research Consortium (SRC), $360K (collaboration with The University of Texas at Austin), co-PI, 2011-2014

[DOE] *A Multi-Language Environment For Programmable Code Optimization and Empirical Tuning*, Department of Energy, $360K (collaboration with University of Texas at San Antonio and Lawrence Livermore National Laboratory), co-PI, 2009-2012

[IBM] *Software Support for Multicore Architectures*, Equipment Grant, IBM, $17K, PI, 2009

TXST *A Framework for Enhanced Feedback-driven Autotuning on Multicore and Manycore Architectures*, Sponsored Research at Texas State University, $45K, PI, 2009-2012

▲ [TXST] *Investigating the Challenges of Mapping OR Algorithms on Multi-core Systems*, Texas State University Research Enhancement Program, $10K, PI, 2009-2010

[DOE] *A Loop Transformation Tool for Improving Application Performance*, Department of Energy (DOE), subcontract from Rice University, $43K, PI, 2008-2010

[TXST] *Software Support for Better Utilization of the Shared-cache Architecture on Multi-core Systems*, Texas State University Research Enhancement Program, $8K, PI, 2008-2009

Publications (■ = conference, ■ = journal, ■ = book chapter)

■ [J14] Selena N. Ramirez, Cassandra M. Johnson, Apan Qasem, Alexandra E. Cagle, and Lesli Biediger-Friedman. Perception of nueva (nutrition for underserved elders via application) usability during alpha-testing. *Journal of Nutrition Education and Behavior*, 55(7, Supplement):56, 2023. Empowering Food Citizens: Together for Nutrition and Food Systems Transformation. Re-connect, Re-nourish, Re-inspire. . . .

■ [C69] Apan Qasem, Hartwig Anzt, Eduard Ayguade, Katharine Cahill, Ramon Canal, Jany Chan, Eric Fosler-Lussier, Fritz Gobel, Arpan Jain, Marcel Koch, Mateusz Kuzak, Josep Llosa, Raghu Machiraju, Xavier Martorell, Pratik Nayak, Shameema Oottikkal, Marcin Ostasz, Dhabaleswar K. Panda, Dirk Pleiter, Rajiv Ramnath, Maria-Ribera Sancho, Alessio Sclocco, Aamir Shafi, Hanno Spreeuw, Hari Subramoni, and Karen Tomko. Lightning Talks of EduHPC 2022. In *10th IEEE/ACM Workshop on Education for High Performance Computing, EduHPC@SC 2022, Dallas, TX, USA, November 14, 2022*, pages 43–50. IEEE, Nov 2022.

■ [C68] David Bunde, Kishwar Ahmed, Sridevi Ayloo, Tisha Brown-Gaines, Joel Fuentes, Vishwesh Jatala, Ruth Kurniawat, Isıl Oz, Apan Qasem, Philip J. Schielke, Mary C. Tedeschi, and Thomas Y. Yeh. Adopting Heterogeneous Computing Modules: Experiences from a ToUCH Summer Workshop. In *10th IEEE/ACM Workshop on Education for High Performance Computing, EduHPC@SC 2022, Dallas, TX, USA, November 14, 2022*. IEEE, Nov 2022.

■ [C66] Md Erfanul Haque Rafi, Kaylee Williams, and Apan Qasem. Raptor: Mitigating CPU-GPU False Sharing Under Unified Memory Systems. In *2022 IEEE 13th International Green and Sustainable Computing Conference (IGSC)*, pages 1–8, Oct 2022.

■ [C67] Md Erfanul Haque Rafi and Apan Qasem. Optimal Launch Bound Selection in CPU-GPU Hybrid Graph Applications with Deep Learning. In *2022 IEEE 13th International Green and Sustainable Computing Conference (IGSC)*, pages 1–7, Oct 2022.

■ [J13] Johana Di Girolamo, Jacob M. Hope, and Apan Qasem. Uncovering input-sensitive energy bottlenecks in oversubscribed GPU workloads. *Sustainable Computing: Informatics and Systems*, 35:100654, Sep 2022.

■ [C65] Apan Qasem. YODA: A pedagogical tool for teaching systems concepts. In Larry Merkle, Maureen Doyle, Judithe Sheard, Leen-Kiat Soh, and Brian Dorn, editors, *SIGCSE 2022: The 53rd ACM Technical Symposium on Computer Science Education, Providence, RI, USA, March 3-5, 2022, Volume 1*, pages 613–618. ACM, Mar 2022.

■ [C64] Apan Qasem and David P. Bunde. Heterogeneous computing for undergraduates: Introducing the touch module repository. In Larry Merkle, Maureen Doyle, Judithe Sheard, Leen-Kiat Soh, and Brian Dorn, editors, *SIGCSE 2022: The 53rd ACM Technical Symposium on Computer Science Education, Providence, RI, USA, March 3-5, 2022, Volume 2*, page 1201. ACM, Mar 2022.

■ [J12] Apan Qasem, David P. Bunde, and Philip Schielke. A module-based introduction to heterogeneous computing in core courses. *Journal of Parallel and Distributed Computing*, 158:56–66, Dec 2021.

- [C63] Blake Ford, Ziliang Zong, Apan Qasem, and Jelena Tesic. Migrating Software from x86 to ARM Architecture: An Instruction Prediction Approach. In *2021 IEEE International Conference on Networking, Architecture and Storage, NAS 2021, Riverside, CA, USA, October 24-26, 2021*, pages 1–8. IEEE, Oct 2021.

- [C62] Jacob M. Hope, Mikel Gjergji, Johana Di Girolamo, Marco A. Alvarez, and Apan Qasem. Characterizing input-sensitivity in tightly-coupled collaborative graph algorithms. In Laurent Lefèvre, Stacy Patterson, Young Choon Lee, Haiying Shen, Shashikant Ilager, Mohammad Goudarzi, Adel Nadjaran Toosi, and Rajkumar Buyya, editors, *21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, CCGrid 2021, Melbourne, Australia, May 10-13, 2021*, pages 287–296. IEEE, Aug 2021. [**Best Paper Award**].

- [J10] David P. Bunde, Apan Qasem, and Philip Schielke. Short modules for introducing heterogeneous computing: Conference tutorial. *J. Comput. Sci. Coll.*, 36(8):95–96, Apr 2021.

- [J11] David P. Bunde, Apan Qasem, and Philip Schielke. Short modules for introducing heterogeneous computing: Conference tutorial. *J. Comput. Sci. Coll.*, 36(6):95–96, Apr 2021.

- [C61] David P. Bunde, Apan Qasem, and Philip Schielke. Teaching about heterogeneous computing. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, page 1355, New York, NY, USA, Mar 2021. Proceedings of the 52nd ACM Technical Symposium on Computer Science Education, Association for Computing Machinery.

- [C60] Tanzima Sultana, Blake Allen, and Apan Qasem. Intelligent Data Placement on Discrete GPU Nodes with Unified Memory. In *Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques*, PACT '20, page 139–151, New York, NY, USA, 2020. Association for Computing Machinery.

- [C59] Apan Qasem. A Gentle Introduction to Heterogeneous Computing for CS1 Students. In *2019 IEEE/ACM Workshop on Education for High-Performance Computing (EduHPC) co-located with SC19*, Nov 2019. [**Best Paper Award**].

- [J9] Saeed Taheri, Apan Qasem, and Martin Burtscher. A tool for automatically suggesting source-code optimizations for complex GPU kernels. *CoRR*, abs/1910.07776, 2019.

- [C58] Jacob M. Hope, Trisha Nag, and Apan Qasem. Energy-efficient GPU graph processing with on-demand page migration. In *Tenth International Green and Sustainable Computing Conference, IGSC 2019, Alexandria, VA, USA, October 21-24, 2019*, pages 1–8, 2019.

- [C57] Md Syadus Sefat, Semih Aslan, Jeffrey W. Kellington, and Apan Qasem. Accelerating hotspots in deep neural networks on a capi-based FPGA. In *21st IEEE International Conference on High Performance Computing and Communications (HPCC19)*, pages 248–256, 2019.

- [C56] Apan Qasem, Clara Novoa, Chandra Kolla, and Samantha Coyle. High-Accuracy Scalable Solutions to the Dynamic Facility Layout Problem [extended abstract]. In *31st International Conference on High Performance Computing Networking, Storage and Analysis - Companion Volume (SC18)*, Nov 2018.

- [C55] Md Syadus Sefat, Semih Aslan, and Apan Qasem. Hardware Acceleration of CNNs with Coherent FPGAs [extended abstract]. In *31st International Conference on High Performance Computing Networking, Storage and Analysis - Companion Volume (SC18)*, Nov 2018.

- [B1] Apan Qasem. Modules for Teaching Parallel Performance Concepts. In Sushil K. Prasad, Anshul Gupta, Arnold Rosenberg, Alan Sussman, and Charles Weems, editors, *Topics in Parallel and Distributed Computing Enhancing the Undergraduate Curriculum: Performance, Concurrency, and Programming on Modern Platforms*, chapter 2, pages 59–77. Springer, 2018.

- [C54] Apan Qasem, Ashwin M. Aji, and Michael L. Chu. Investigating data layout transformations in Chapel. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops, IPDPS Workshops 2018, Vancouver, BC, Canada, May 21-25, 2018*, pages 915–924, 2018.

- [C51] Tiffany A. Connors and Apan Qasem. Automatically Selecting Profitable Thread Block Sizes for Accelerated Kernels. In *19th IEEE International Conference on High Performance Computing and Communications (HPCC17)*, pages 442–449, Dec 2017. [Acceptance Rate: 36%].

- [C50] Biplab Saha, Tiffany A. Connors, Saami Rahman, and Apan Qasem. A Machine Learning Approach to Automatic Creation of Architecture-sensitive Performance Heuristics. In *19th IEEE International Conference on High Performance Computing and Communications (HPCC17)*, pages 18–25, Dec 2017. [**Best Paper (1/176)**].

- [C49] Apan Qasem, Aswhin Aji, and Gregory Rodgers. Characterizing data organization effects on heterogeneous memory architectures. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, Mar 2017.

- [C52] Apan Qasem and Samuel Teich. Mitigating register pressure in GPU kernels for improved energy efficiency. In *Eighth International Green and Sustainable Computing Conference, IGSC 2017, Orlando, FL, USA, October 23-25, 2017*, pages 1–7. IEEE Computer Society, 2017.

- [C53] Apan Qasem and Samuel Teich. Evaluating the impact of data layout and placement on the energy efficiency of heterogeneous applications. In *Eighth International Green and Sustainable Computing Conference, IGSC 2017, Orlando, FL, USA, October 23-25, 2017*, pages 1–8. IEEE Computer Society, 2017.

- [C47] Claudia Alavardo, Dan Tamir, and Apan Qasem. Realizing energy-efficient thread affinity configurations with supervised learning. In *Sixth International Green and Sustainable Conference (IGSC15)*, Dec 2015.

- [C48] Tiffany Connors and Apan Qasem. Power-performance analysis of metaheuristic search algorithms on the gpu. In *International Workshop on Green Programming, Computing and Data Processing (GPCDP15)*, Dec 2015.

- [C46] Mario Gutierrez, Saami Rahman, and Apan Qasem. Neural network methods for fast and portable prediction of cpu power consumption. In *Sixth International Green and Sustainable Conference (IGSC15)*, Dec 2015.

- [C45] Biplab Kumar Saha, Saami Rahman, and Apan Qasem. Mltune: A tool-chain for automating the workflow of machine-learning based performance tuning (extended poster abstract). In *28th International Conference on High Performance Computing, Networking, Storage and Analysis (SC15)*, Nov 2015.

- [C44] Saami Rahman and Apan Qasem. Investigating prefetch potential on the xeon phi with autotuning (extended poster abstract). In *28th International Conference on High Performance Computing, Networking, Storage and Analysis (SC15)*, Nov 2015.
- [C42] Mario Gutierrez, Dan Tamir, and Apan Qasem. Evaluating neural network methods for pmc-based cpu power prediction. In *International Multi-Conference on Computing in the Global Information Technology (ICCGI)*, Oct 2015.
- [C43] Claudia Alavardo, Dan Tamir, and Apan Qasem. Energy-efficient thread migration via dynamic characterization of resource utilization. In *International Multi-Conference on Computing in the Global Information Technology (ICCGI)*, Oct 2015.
- [C41] Saami Rahman, Martin Burtscher, Ziliang Zong, and Apan Qasem. Maximizing hardware prefetch effectiveness with machine learning. In *17th IEEE International Conference on High Performance Computing and Communications (HPCC15)*, Aug 2015.
- [C40] Abhilash Chaparala, Clara Novoa, and Apan Qasem. Autotuning gpu-accelerated qap solvers for power and performance. In *17th IEEE International Conference on High Performance Computing and Communications (HPCC15)*, Aug 2015.
- [C39] Saeed Taheri, Apan Qasem, and Martin Burtscher. A tool for automatically suggesting source-code optimizations for complex gpu kernels. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, Jul 2015.
- [C38] Clara Novoa, Apan Qasem, and Abhilash Chaparala. A SIMD tabu search implementation for solving the quadratic assignment problem with gpu acceleration. In *Fourth Annual XSEDE Conference (XSEDE15)*, Jul 2015.
- [C37] Martin Burtscher, Wuxu Peng, Apan Qasem, Hongchi Shi, Dan Tamir, and Heather Thiry. A module-based approach to adopting the 2013 acm curricular recommendations on parallel computing. In *Proceedings of the 36th SIGCSE Technical Symposium on Computer Science Education (SIGCSE15)*, Mar 2015.
- [C36] Apan Qasem. Exposing undergraduates to parallel performance concepts with a three-module sequence. In *Workshop on Education for High-Performance Computing (EDUHPC14, educator workshop at SC14)*, Nov 2014.
- [C35] Abhilash Chaparala, Clara Novoa, and Apan Qasem. A simd solution for the quadratic assignment problem with gpu acceleration. In *Third Annual XSEDE Conference (XSEDE14)*, Jul 2014.
- [C32] Shwetha Shankar, Greg Lakomski, Claudia Alvarado, Richard Hay, Christopher Hyatt, Dan Tamir, and Apan Qasem. Power aware work stealing in homogeneous multicore systems. In *The Sixth International Conference on Future Computational Technologies and Applications*, 2014.
- [C34] Christopher Hyatt, Greg LaKomski, Claudia Alvarado, Richard Hay, Apan Qasem, and Dan Tamir. Power aware task matching and migration in heterogeneous processing environments. In *International Conference on Computational Science and Computational Intelligence (CSCI14)*, 2014.
- [C33] Claudia Alvarado, Dan Tamir, and Apan Qasem. Dynamic feedback-driven thread migration for energy-efficient execution of multithreaded workloads. In *Proceedings of the 16th Annual TECHCON Conference*, 2014.

- [C30] Martin Burtscher, Wuxu Peng, Apan Qasem, Hongchi Shi, and Dan Tamir. Preparing computer science students for an increasingly parallel world: Teaching parallel computing early and often (extended poster abstract). In *26th International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, Nov 2013.
- [C31] Martin Burtscher, Wuxu Peng, Apan Qasem, Hongchi Shi, Dan Tamir, and Heather Thiry. Integrating parallel computing into the undergraduate curriculum at texas state university: Experiences from the first year. In *Workshop on Parallel, Distributed, and High-Performance Computing in Undergraduate Curricula (EDUPDHPC13, educator workshop at SC13)*, Nov 2013.
- [C29] Saami Rahman, Richard Hay, and Apan Qasem. Enhancing learning-based autotuning with composite and diagnostic feature vectors (extended poster abstract). In *26th International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, Nov 2013.
- [C28] Hammad Rashid, Richard Hay, Clara Novoa, and Apan Qasem. Algorithmic choice in optimization problems: A performance study (extended poster abstract). In *26th International Conference on High Performance Computing, Networking, Storage and Analysis (SC13)*, Nov 2013.
- [C25] Shwetha Shankar, Dan Tamir, and Apan Qasem. Towards an operating system based framework for energy-efficient scheduling of parallel workloads. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, Jul 2013.
- [C26] Christopher R Hyatt, Greg R. LaKomski, Dan Tamir, and Apan Qasem. Power aware task matching and migration in heterogeneous processing environments. In *Proceedings of the 15th Annual TECHCON Conference*, 2013.
- [C27] Jim Holt, George Bazzera, Apan Qasem, Jason Miller, and Henry Hoffmann. A pattern language for adaptive parallel software. In *Proceedings of the 20th International Conference on Pattern Languages of Programs*, 2013.
- [J8] Apan Qasem and Joshua Magee. Improving TLB performance on current chip multiprocessor architectures through demand-driven superpaging. *Software Practice and Experience (SPE)*, 43(6):750–729, 2013.
- [C24] Apan Qasem, Michael Jason Cade, and Dan Tamir. Improved energy efficiency for multi-threaded kernels through model-based autotuning. In *Proceedings of the 2012 IEEE Green Technology Conferenc (GTC12)*, pages 1–12, Apr 2012.
- [C23] Swapneela Unkule, Christopher Shaltz, and Apan Qasem. Automatic restructuring of GPU kernels for exploiting inter-thread data locality. In *Proc. Int'l. Conf. on Compiler Construction (CC12)*, pages 21–40, Mar 2012.
- [C22] Apan Qasem. Efficient execution of time-step computations with pipelined parallelism and inter-thread data locality optimizaitions. In *Proceedings of the 2012 PPOPP International Workshop on Programming Models and Applications for Multicore and Manycores (PMAM12)*, pages 27–35, Feb 2012.
- [C21] Apan Qasem and Dan Tamir. Memory performance diagnosis through feedback synthesis. In *Proceeding of the Workshop on Feedback-Directed Compiler Optimization for Multi-Core Architectures (COMA12 a HIPEAC workshop)*, pages 5–10, Jan 2012.

- [J7] Santosh Sarangkar and Apan Qasem. Mats: A model-driven adaptive tuning system for parallel workloads. *Journal of Parallel and Cloud Computing (JPCC)*, 1(2):50–64, 2012.
- [J6] Apan Qasem. High-level language extensions for fast execution of pipeline-parallelized code on current chip multi-processor systems. *International Journal of Programming Languages and Applications (IJPLA)*, 2(3):1–12, 2012.
- [J3] Hammad Rashid, Clara Novoa, Mark McKenney, and Apan Qasem. Efficient parallel solutions to the integral knapsack problem on current chip-multiprocessor systems. *International Journal of Parallel, Emergent and Distributed Systems (IJPEDS)*, 27(1):19–44, 2012.
- [J4] Apan Qasem. Autotuning strategies for reducing synchronization costs in multithreaded kernels. *Journal of Systems and Software*, 2(4):152–165, 2012.
- [J5] Apan Qasem. Architectural considerations for compiler-guided unroll-and-jam of cuda kernels. *American Journal of Computer Architecture*, 1(2):12–20, 2012.
- [C19] Swapneela Unkule and Apan Qasem. Register pressure aware code transformations on GPU. In *24th International Conference on High Performance Computing Networking, Storage and Analysis - Companion Volume (SC11)*, pages 19–20, Nov 2011.
- [C20] Faizur Rahman, Qing Yi, and Apan Qasem. Understanding stencil code performance on multicore architectures. In *Conf. Computing Frontiers (CF11)*, pages 30–45, 2011.
- [C18] Clara Novoa, Apan Qasem, Hammad Rashid, and Mark McKenney. Dynamic programming solutions for the integral knapsack problem on multicore architectures, (extended abstract). In *11th INFORMS Computing Society Conference, (ICS11)*, 2011.
- [C17] Santosh Sarangkar and Apan Qasem. Intelligent feedback for fast and effective autotuning, (extended poster abstract). In *23rd International Conference on High Performance Computing, Networking, Storage and Analysis - Companion Volume (SC10)*, Nov 2010.
- [C13] Hammad Rashid, Clara Novoa, and Apan Qasem. An evaluation of parallel knapsack algorithms on multicore architectures. In *Proceedings of the 2010 International Conference on Scientific Computing (CSC10)*, pages 230–235, Jul 2010.
- [C12] Santosh Sarangkar and Apan Qasem. Restructuring parallel loops to curb false sharing on multicore architectures. In *24th IEEE International Symposium on Parallel Distributed Processing Workshops and Phd Forum (IPDPSW10)*, pages 1–7, 2010.
- [C15] Apan Qasem. Locality-conscious superpaging for improved tlb behavior of stencil computations. In *Proceedings of the 2010 International Conference on High Performance Computing Systems (HPCS10)*, 2010.
- [C14] Qing Yi, Santosh Sarangkar, and Apan Qasem. Improving autotuning effciency and portability through feedback diagnostics. In *Proceedings of the Fifth International Workshop on Automatic Performance Tuning (iWAPT10)*, 2010.
- [C11] Apan Qasem, Jichi Guo, Faizur Rahman, and Qing Yi. Exposing tunable parameters in multi-threaded numerical code. In *Network and Parallel Computing, IFIP International Conference, (NPC10)*, pages 46–60, 2010.
- [C16] Qing Yi, Jichi Guo, and Apan Qasem. Evaluating the role of optimization-specific search heuristics in effective autotuning (short paper). In *23rd International Workshop Languages and Compilers for Parallel Computing (LCPC10)*, 2010.

- [C10] Joshua Magee and Apan Qasem. A case for compiler-driven superpage allocation. In *Proceedings of the 47th Annual Southeast Regional Conference, (ACMSE09)*, 2009.
- [C9] Michael Jason Cade and Apan Qasem. Balancing locality and parallelism on shared-cache mulit-core systems. In *11th IEEE International Conference on High Performance Computing and Communications (HPCC09)*, pages 188–195, 2009.
- [C7] Apan Qasem. Evaluating an early-stop criterion and a statistical pruning strategy of the optimization search space. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, pages 506–510, Jul 2008.
- [J2] Apan Qasem and Ken Kennedy. Model-guided empirical tuning of loop fusion. *International Journal of High Performance Systems Architecture (IJHPSA)*, 1(3):183–198, 2008.
- [C8] Qing Yi and Apan Qasem. Exploring the optimization space of dense linear algebra kernels. In *21st International Workshop Languages and Compilers for Parallel Computing (LCPC08)*, pages 343–355, 2008.
- [C6] Apan Qasem and Ken Kennedy. Pruning the optimization search space using architectureaware cost models. In *Proceedings of the First Workshop on Statistical and Machine Learning Approaches Applied to Architecture and Compilation (SMART07)*, Feb 2007.
- [C5] Apan Qasem and Ken Kennedy. Profitable loop fusion and tiling using model-driven empirical search. In *Proceedings of the 20th Annual International Conference on Supercomputing (ICS)*, pages 249–258, Jul 2006.
- [J1] Apan Qasem, Ken Kennedy, and John M. Mellor-Crummey. Automatic tuning of whole applications using direct search and a performance-based transformation system. *The Journal of Supercomputing*, 36(2):183–196, 2006.
- [C4] Apan Qasem and Ken Kennedy. A cache-conscious profitability model for empirical tuning of loop fusion. In *18th International Workshop on Languages and Compilers for Parallel Computing, (LCPC)*, pages 106–120, 2005.
- [C3] Apan Qasem, Ken Kennedy, and John Mellor-Crummey. Automatic tuning of whole applications using direct search and a performance-based transformation system. In *Proceedings of the Los Alamos Computer Science Institute 5th Annual Symposium (LACSI04)*, 2004.
- [C2] Robert Fowler, John Mellor-Crummey, Guohua Jin, and Apan Qasem. A source-to-source loop transformation tool (extended poster abstract). In *Proceedings of the Los Alamos Computer Science Institute 3rd Annual Symposium (LACSI02)*, 2002.
- [C1] Apan Qasem, David B. Whalley, Xin Yuan, and Robert van Engelen. Using a swap instruction to coalesce loads and stores. In *7th International Euro-Par Conference Parallel Processing, (EuroPar01)*, pages 235–240, 2001.

Invited Talks

Jun 2023 *Building Machine Learning Compilers for Next-Generation HPC Systems*, Research Colloquium, University of Liberal Arts, Dhaka, Bangladesh

Jun 2023 *High-Performance Computing for Analytics and Data Science*, Data Analytics Showcase, Texas State University, San Marcos, TX

Oct 2019 *Compilers and Runtimes for Heterogeneous Memory Systems*, Research Colloquium, University of Texas at Arlington, Arlington, TX

Oct 2019 *Adaptive Data Placement in Heterogeneous Memory Architectures*, Sustainable Pathways Workshop, Lawrence Berkeley National Laboratory, Berkeley, CA

Oct 2015 *Enhancing Learning-based Autotuning With Feedback Diagnostics*, Research seminar, University of Edinburgh, Edinburgh, United Kingdom

Jun 2015 *Enhancing Learning-based Autotuning With Feedback Diagnostics*, Research Seminar, AMD Research, Austin, TX

May 2015 *Autotuning for Exascale*, Research Seminar, Qualcomm, Austin, TX

Jun 2015 *Integrating Parallel Computing into the Undergraduate Curriculum : Efforts at Texas State University*, HPC Workshop with Applications in Big Data, Prairie View A&M University, Prairie View, TX

May 2015 *Autotuning for Exascale*, Qualcomm Inc., Technical Talk, Austin, TX

Nov 2014 *Singularity in Movies*, STEM Learning Community Event, Texas State University, San Marcos, TX

Mar 2014 *Adding PDC within a Six-course Subset of the CS major*, SIGCSE Special Session, SIGCSE 2014, Atlanta, GA

Apr 2013 *Autotuning for Exascale*, Industrial Advisory Board Meeting, Texas State University, San Marcos, TX

Oct 2009 *Automatic Tuning of Scientific Code for Multicore Architectures*, Discrete Math Seminar, Department of Mathematics, Texas State University, San Marcos, TX

Jun 2009 *Compiler's Role in Efficient Multicore Software Development*, IBM Research Meeting, IBM, Austin, TX

Apr 2009 *Computer Science Involvement in HPC Research at Texas State*, Texas State HPC Meeting, Texas State University, San Marcos, TX

Nov 2008 *Data Locality and Parallelism on CMPs with Shared Caches*, Industrial Advisory Board Meeting, Texas State University, San Marcos, TX

Feb 2008 *Model-based Autotuning*, Computer Science Seminar, Department of Computer Science, Texas State University, San Marcos, TX

Dec 2007 *Model-guided Automatic Performance Tuning*, Workshop in Memory of Ken Kennedy, Computer and Information Technology Institute, Houston, TX

Apr 2007 *Automatic Tuning of Whole Applications*, Computer Science Seminar, Department of Computer Science, Rice University, Houston, TX

Students

2022 **Kaylee Williams**
BS, Texas State University
Towards a Green Future: Energy Efficient Conversational AI on the Edge
First Employment: MS at UT Austin

2021 **Samantha Coyle**
BS, Texas State University
High-Accuracy Scalable Solutions to the Dynamic Facility Layout Problem
First Employment: Intel

2018 **Md. Syadus Sefat**
MS, Texas State University
Design And Performance Analysis Of Hardware Accelerators for Deep Neural Networks on Heterogeneous Platforms
First Employment: PhD student at Virginia Tech

2017 **Tiffany Connors**
BS, Texas State University
Automatically Selecting Profitable Thread Block Sizes Using Machine Learning
First Employment: Lawrence Livermore

2016 **Biplab Kumar Saha**
MS, Texas State University
Towards a Framework for Automating the Workflow for Building ML-based Performance Tuning Models
First Employment: SageGlass

2016 **Shuvabrata Saha**
MS, Texas State University
A Multi-Objective Autotuning Framework for the Java Virtual Machine
First Employment:

2015 **Saami Rahman**
MS, Texas State University
An Exploration into the Effectiveness of Prefetching on Program Performance with the Help of an Autotuning Model
First Employment: TSNS

2014 **Claudia Alvarado**
MS, Texas State University
Dynamic Feedback-Driven Thread Migration For Energy-efficient Execution Of Multithreaded Workloads
First Employment: Intel

2014 **Richard Hay**
MS, Texas State University
Machine Learning Based DVFS For Energy Efficient Execution Of Multithreaded Workloads
First Employment: Jwaala Inc.

2011 **Swapneela Unkule**
MS, Texas State University
Exploiting Shared-Memory Reuse Through Source-Level Transformations of CUDA Kernels
First Employment: atsec information security

2011 **Santosh Sarangkar**
MS, Texas State University
PTUNE: A Tool For Online Autotuning With Optimizing Compiler
First Employment: Intel

2011 **Suman Vara**
MS, Texas State University
Evaluating Shared-cache Performance With Microbenchmarks And Reuse Distance Analysis
First Employment: John Deere

2010 **Hammad Rashid**
MS, Texas State University
Parallel Knapsack Algorithms On Multicore Architectures
First employment: VMWare

2008 **Joshua Magee**
MS, Texas State University
Automated Compiler Driven Superpage Allocation and its Applications
First Employment: Sony Playstation

2008 **Michael J Cade**
MS, Texas State University
Balancing Data Locality and Parallelism for Improved Application Performance on Multicore Platforms
First employment: IBM

Teaching

2017- *CS7331: High Performance Computing*, Texas State
2010- *CS3339: Computer Architecture*, Texas State
2020- *CS4379V: Fundamentals of Algorithms*, Texas State
2010-19 *CS1428: Foundations of Computer Science - Honors*, Texas State
2008-17 *CS5331: Crafting Compilers*, Texas State, graduate course
2008-17 *CS4318: Program Translators*, Texas State
2014 *MC4382: Coding and Data Skills for Communicators*, School of Journalism, Texas State
2008,10-12 *CS4350: Unix Systems Programming*, Texas State
2007-09 *CS1428: Foundations of Computer Science*, Texas State, Enrollment : 85
2003 *CS512: Advanced Compiler Construction*, Rice University, graduate course, TA
2003 *Theory of Automata and Formal Languages*, Rice University, TA
2002 *Computer Organization*, Rice University, TA
2002 *Applied Algorithms and Data Structures*, Rice University, TA
2001 *Intermediate Programming*, Rice University, TA
2001 *Theory of Automata*, North South University
2001 *Compiler Construction*, North South University
2000 *Micro Applications for Business*, Florida State University, Instructor, Enrollment: 150
1999 *Computer Literacy*, Florida State University, TA

Professional Service

Program Committee

International Conference on Parallel Processing (ICPP) IEEE International Conference on Cloud Computing (IEEE CLOUD) International Symposium on Code Generation and Optimization (CGO), IEEE International Conference on Cluster Computing (CLUSTER) International Conference on Compiler Construction (CC) IEEE/ACM Intl Symposium on Cluster, Cloud and Grid Computing (CCGrid), Workshop on Education for High Performance Computing (EduHPC) IEEE Intl Conf on High Performance Computing and Communications (HPCC), Intl Conf on Network and Parallel Computing (NPC), Intl Conf on High Performance Computing (HIPC), Intl Conf on Parallel and Distributed Computing, Applications and Technologies (PDCAT), Intl Workshop on GPUs and Scientific Applications (GPUScA), Workshop on Statistical and Machine learning approaches to ARchitectures and compilation (SMART), Intl Workshop on Automatic Performance Tuning, (iWAPT)

Reviewer

NSF CNS review panel, NSF/IEEE-TCPP PDC curriculum, Oak Ridge Associated Universities grant proposal, Book Chapter: Topics in Parallel and Distributed Computing, ACM Transactions on Architectural and Code Optimization (TACO), ACM Transactions on Embedded Computing (TECS), IEEE Transactions on Paralllel and Distributed Systems (TPDS), Journal of Supercomputing (JOS), Journal of Parallel and Distributed Computing (JPDC), ACM Conference on Programming Language Design and Implementation (PLDI), Parallel Architectures and Compilation Techniques (PACT), ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP), European Conference on Paralllel and Distributed Computing (EUROPAR), Cluster Computing (CC)

Workshops and Tutorials

2023 Program Chair, Workshop on Education for High-Performance Computing, EduHPC23

2022 Program Co-Chair, Workshop on Education for High-Performance Computing, EduHPC22

2022 Co-Organizer, Heterogeneous Computing for Undergraduates: Introducing the ToUCH Module Repository, SIGCSE22

2022 Co-Organizer, ToUCH Training Workshop, TXST and Knox College

2021 Co-Organizer, Teaching about Heterogeneous Computing, SIGCSE21

2021 Co-Organizer, ToUCH Training Workshop, TXST, Knox College and Concordia

2021 Co-Organizer, Short modules for introducing heterogeneous computing: conference tutorial, CCSC-Southwestern

2021 Co-Organizer, Short modules for introducing heterogeneous computing: conference tutorial, CCSC-Central Plains

2014 Co-Organizer, Regional Workshop in teaching parallel computing, TXST

University Service

2021-2022	Chair, Non-tenure Line Hiring Committee
2021-	Member, PhD Prior Learning Assessment
2021-	Member, Student Assistant Assignment Committee
2020-	Member, Space Committee
2019-	ACM Student Advisory Board, Texas State University
2018	Texas Higher Education Coordinating Board's Field of Study (FOS) Advisory Committee
2018-	Member, Presidential Awards Nomination Committee
2017-	Course Coordinator, CS3339: Computer Architecture
2015	Texas Higher Education Coordinating Board's Academic Course Guide Committee
2015	Speaker, New Scholar Orientation
2014-	Coordinator, Adjunct Faculty Pool
2014	Speaker, STEM Learning Community
2014-2017	Academic Governance Committee
2014-	Academic Affairs Computing Priority Committee
2014-	Faculty Advisor, Bangladesh Student Association
2014	Speaker, Admitted Student STEM Faculty Panel
2014	Honors College Scholarship Committee
2014	Faculty Search Committee
2014-	Personnel Committee
2013-	Member, Tenure Track Faculty Search Committee
2013	Graduate College Scholarship Committee
2013	Charter Member, Texas State Chapter of Phi Kappa Phi
2013-2015	H-LSAMP Faculty Advisor
2012-	Class Schedule Coordinator
2012-	Program Assessment and Review Committee
2012-	Departmental ABET Committee
2012-	Undergraduate Faculty Advisor (Summer)
2012-2014	Judge, REU Research Poster Session
2012	Departmental Liaison, Intel Internship Program
2012	Reviewer, Student Undergraduate Research Fund (SURF) Applications
2012	Staff Search Committee
2011	Judge, Women in Science and Engineering Conference (WISE) Poster Competition
2011	Judge, HEB University Developer Challenge
2010	Departmental Representative, Annual CAHSI meeting
2010	Judge and Session Chair, International Research Conference for Graduate Students
2008-2015	Course Coordinator, CS1428: Foundations of Computer Science
2008-2020	Course Coordinator, CS4318 Program Translators
2008	Student Commencement Speaker Committee
2007-	Undergraduate Curriculum Committee
2007-	Graduate Comprehensive Exam