LI TANG

505.667.2386 LTANG@LANL.GOV
MAIL STOP B287, LOS ALAMOS, NM 87545

EDUCATION

Ph.D. Computer Science and Engineering
UNIVERSITY OF NOTRE DAME, 2017
Thesis: "Performance and Energy Aware Workload Partitioning on Heterogeneous Platforms", Advisor: Dr. Xiaobo Sharon Hu

M.E. Computer Science
HUAZHONG UNIV. OF SCIENCE AND TECHNOLOGY, 2009

B.E. Computer Science
HUAZHONG UNIV. OF SCIENCE AND TECHNOLOGY, 2007

RESEARCH INTERESTS AND EXPERTISE

High-performance computing system co-design, from the perspective of efficient data movement across layers, and with a focus on combining performance/power modeling and analysis. Technical expertise in the areas of analytical performance/energy modeling, in-situ infrastructure, heterogeneous computing, and programming models.

PROFESSIONAL EXPERIENCE

LOS ALAMOS NATIONAL LABOATORY
NOV 2018 – Present
Los Alamos, NM

Scientist 2, Supervisor: Patrick McCormick

- Software infrastructure development for performant Julia based in-situ data analysis at large-scale on simulated weather data from E3SM.
- Performance Modeling of Machine Learning Accelerators.
- Performance analysis and modeling of asynchronous programming models and runtime by automating the task dependency abstraction and benchmark generation.
- Analytical performance modeling of HPC applications for GPU suitability identification by using Byfl, an LLVM based program analysis tool.

BROOKHAVEN NATIONAL LABOATORY
MAY 2017 – NOV 2018
Upton, NY

Research Associate (Post-Doc), Supervisor: Kerstin Kleese van Dam

- Runtime optimizations for in-situ/online data analysis and data reduction with large-scale dataset and enhancement for management, quantitative comparison, and selection of output from data reduction pipelines.
- Online performance profiling of complex scientific workflows and performance modeling for workflow management optimization.

SANDIA NATIONAL LABORATORIES
MAR 2012 – DEC 2012
Albuquerque, NM

Research Intern, Supervisor: Richard Barrett

- Developed new algorithms of accelerating a proxy of unstructured implicit finite element method applications on GPUs.
- Implemented the developed algorithms in CUDA, OpenCL and OpenACC and Intel Intrinsics. Evaluated the performance and energy of CPUs, GPUs and FPGA by using the ported implementations.

RESEARCH EXPERIENCE

UNIV. OF NOTRE DAME
JAN 2011 – MAY 2017
Notre Dame, IN

Graduate Research Assistant, Advisor: Xiaobo Sharon Hu

- Developed performance/energy models of heterogeneous platforms and built the PeaPaw and PeaPaw+ frameworks and guidelines to guide workload partitioning between CPU and GPU.
- Evaluated and modeled the performance and energy efficiency of integrating different accelerators (GPU, FPGA and Intel Phi) with CPU, and assessed the different choices for accelerator integration (discrete and integrated) and workload offloading (data- and code-partitioning).
- Developed several algorithms and mapping schemes of accelerating finite element method on GPUs.
- Developed an energy model for non-volatile multi-core CPUs.

NEW JERSEY INSTITUTE OF TECHNOLOGY
SEP 2009 – DEC 2010

Graduate Research Assistant, Advisor: Jie Hu

- Characterized the L1 data cache vulnerability of Chip-Multiprocessors for both snooping-based and directory-based cache coherence protocols by implementing an extended architectural vulnerability factor model in the M5 simulator. Enhanced the MESI protocol for higher reliability against soft errors.

HUAZHONG UNIV. OF SCIENCE AND TECHNOLOGY
JUN 2007 – MAY 2009
Wuhan, CHINA

Graduate Research Assistant, Advisor: Jingning Liu

- Implemented a MicroBlaze system with reconfigurable cryptographic coprocessors and low-latency task scheduling algorithm on Virtex 5 FPGA.

PUBLICATIONS

- **L. Tang** and Scott Pakin, "Cross-Level Characterization of Program Execution," *30th International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2022),* 2022.
- [Best Paper Award] S. Dutta, N.E. Klein, **L. Tang**, J.D. Wolfe, L.V. Roekel, J.J. Benedict, A. Biswas, E. Lawrence and N. Urban, "In Situ Climate Modeling for Analyzing Extreme Weather Events," *In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV 2021),* 2021.
- M. Grosskopf, E. Lawrence, A. Biswas, **L. Tang**, K. Rumsey, L.V. Roekel and N. Urban, "Estimating a Spatial Inference Model to a Climate Simulation In-Situ with Sparse Gaussian Processes," *In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV 2021),* 2021.
- [Best Paper Award] C. Kelly, S. Ha, K. Huck, H. van Dam, L. Pouchard, G. Matyasfalvi, **L. Tang**, N. D'Imperio, W. Xu, S. Yoo and K. Kleese van Dam, "Chimbuko: A Workflow-Level Scalable Performance Trace Analysis Tool," *In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV 2020),* 2020.
- J.M. Wozniaka, M. Doriera, R. Rossa, T. Shu, T. Kurc, **L. Tang**, N. Podhorszki and M. Wolf, "MPI Jobs within MPI Jobs: A Practical Way of Enabling Task-level Fault-Tolerance in HPC Workflows," *Future Generation Computer Systems (FGCS)*, 2019.

● L. Pouchard, S. Baldwin, T. Elsethaggen, S. Jha, B. Raju, E. Stephan, **L. Tang** and K. Kleese van Dam, "Computational reproducibility for scientific workflows at pre-exascale," *International Journal of High Performance Computing Applications (IJHPCA)*, 2019.

● Q. Lou, I. Palit, **L. Tang**, A. Horvath, M.T. Niemier and X.S. Hu, "Application-level Studies of Cellular Neural Network-based Hardware Accelerators," *arXiv preprint arXiv:1903.06649*, 2019.

● L. Pouchard, K. Huck, G. Matyasfalvi, D. Tao, **L. Tang**, H. van Dam and S. Yoo, "Prescriptive provenance for streaming analysis of workflows at scale," *2018 New York Scientific Data Summit (NYSDS 2018)*, 2018.

● L. Pouchard, S. Baldwin, T. Elsethagen, C. Gamboa, S. Jha, B. Raju, E. Stephan, **L. Tang** and K. Kleese van Dam, "Use Cases of Computational Reproducibility for Scientific Workflows at Exascale," *arXiv:1805.00967*, 2018.

● R. Perricone, **L. Tang**, M.T. Niemier and X.S. Hu, "Exploiting Non-Volatility for Information Processing," *Proc. of the ACM Great Lakes Symposium on VLSI (GLSVLSI 2017)*, 2017.

● **L. Tang**, X.S. Hu, R.F. Barrett and J. Cook, "PeaPaw: Performance and Energy Aware Partitioning of Workload on Heterogeneous Platforms," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, Volume 22, Issue 3, March 2017.

● **L. Tang**, X.S. Hu and R.F. Barrett, "PerDome: A Performance Model for Heterogeneous Computing Systems," *23rd High Performance Computing Symposium (HPC 2015)*, 2015.

● **L. Tang**, X.S. Hu and R.F. Barrett, "Energy Efficiency in Heterogeneous Computing Platforms: A miniFE Case Study," *Sandia National Laboratories Technical Report*, SAND-2014-20215, 2014.

● **L. Tang**, X.S. Hu, D.Z. Chen, M.T. Niemier, R.F. Barrett, S.D. Hammond and M.Y. Hsieh, "GPU Acceleration of Data Assembly in Finite Element Methods and Its Energy Implications," *The 24th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP 2013)*, 2013.

● R.F. Barrett, S. Borkar, S.S. Dosanjh, S.D. Hammond, M.A. Heroux, X.S. Hu, J.Luitjens, S.Parker, J. Shalf and **L. Tang**, "On the Role of Co-design in High Performance Computing," *Transition of HPC Towards Exascale Computing, E.H. D¡'Hollander et al. (Eds.)*, 2013.

● **L. Tang**, S. Wang, J. Hu and X.S. Hu, "Characterizing the L1 Data Cache's Vulnerability to Transient Errors in Chip-Multiprocessors," *Proc. of the IEEE Computer Society Annual Symposium on VLSI (ISVLSI 2011)*, 2011.

PRESENTATIONS AND POSTERS

● **L. Tang**, S. Dutta and N.E. Klein, "In-Situ Data Analysis with Julia for E3SM at Large Scale," *JuliaCon 2021,* Virtual Conference, 2021.

● **L. Tang**, and S. Yoo, "Application of Machine Learning to Memory Subsystem Characterization," *Workshop on Modeling & Simulation of Systems and Applications (ModSim 2018),* Seattle, Washington, 2018.

● **L. Tang**, X.S. Hu and R. Barrett, "PerDome: A Performance Model for Heterogeneous Computing Systems," *Great Chicago Area Systems Research Workshop (GCASR 2016),* Chicago, Illinois, 2016.

● **L. Tang**, "Performance and Energy Aware Workload Partitioning on Heterogeneous Platforms," *Supercomputing 2016 (SC 2016),* Doctoral showcase, Salt Lake City, Utah, 2016.

- **L. Tang**, X.S. Hu and R.F. Barrett, "Performance and Energy Comparisons between FPGA and GPU Implementations of Data Assembly," *SIAM Parallel Processing for Scientific Computing*, Portland, Oregon, 2014.
- X.S. Hu and **L.Tang**, "Exploring the Energy and Performance Landscape of FEM Execution," *SIAM Parallel Processing for Scientific Computing*, Savannah, Georgia, 2012.
- **L. Tang**, X.S. Hu, D.Z. Chen and M. Niemier, "Is It Worth It to Use GPU If Energy Is the Concern?" *2011 Salishan Conference on High-Speed Computing*, Salishan, Oregon, 2011.

GRANTS

- **Dates of the project:** 05/2022-09/2022**, Role:** Principal Investigator**, Funding:** $40K**, Project title:** Ocean Modeling in the Language of Tensors**, Funding organization:** DOE, Los Alamos National Laboratory
- **Dates of the project:** 03/2021-09/2021**, Role:** Principal Investigator**, Funding:** $60K**, Project title:** Performance Modeling of Scientific Codes on Machine Learning Accelerators**, Funding organization:** DOE, Los Alamos National Laboratory

AWARDS

- **Alamos Awards Program (LAAP) Award, 2022**
- **Best Paper Award, ISAV 2021**
- **Best Paper Award, ISAV 2020**
- **Supercomputing Doctoral Showcase, 2016**
- **R&D 100, 2013**

PROFESSIONAL SERVICES

- **Editorial position:** IEEE Transactions on Parallel and Distributed Systems (TPDS) Review Board
- **Program committee and reviewer:** FGCS, ICCCN, IPDRM, CLUSTER, ISVLSI, JOPLE, JSA, NAS, NPC, PMAM, RTSS, SC, SNTA, ZUSC
- Mentoring: NVIDIA GPU Hackathon 2017, 2018