Aaditya G. Landge

Graduate Student pursuing M.S. in Computing at The University of Utah;
Academic Collaborator at Lawrence Livermore National Laboratory (LLNL), Livermore, CA

✉ *aaditya@sci.utah.edu*
www.sci.utah.edu/∼aaditya
https://www.linkedin.com/in/adityalandge

Summary

- Currently conducting research and development at LLNL, Livermore and based in San Francisco Bay Area
- Applied research and developer experience in large-scale data analysis, high performance computing, topological data analysis, parallel programming, GPU computing and information visualization – developed and deployed parallel topological data analysis techniques on supercomputers for analyzing large real-world data sets
- Working knowledge of machine learning techniques
- Experience profiling and debugging large parallel codes – developed tool for visualization and analysis of performance metrics
- Strong written and oral communication skills with several research publications and conference presentations

Education

2011 - 2015(Exp) **M.S. in Computing (Data Analysis and Management)**,
School of Computing, The University of Utah, Salt Lake City, *GPA: 3.84/4.0*.
Advisor: Prof. Valerio Pascucci

Thesis: Large Scale In-Situ Topological Analysis Using Segmented Merge Trees: Performance, Scalability and Power Efficiency

Relevant Coursework: Machine Learning, Data Mining, Advanced Algorithms, Computational Topology, Information Visualization, Scientific Visualization, Advanced Scientific Computing, Linear Models

2005 -'09 **Bachelor of Engineering (Computer Science)**,
Vishwakarma Institute of Technology, University of Pune, Pune, India, *GPA: 7.9/10.0*.

Experience

2011 - current **Graduate Research Assistant**, *Scientific Computing and Imaging (SCI) Institute, The University of Utah*,
Academic Collaborator, *Lawrence Livermore National Laboratory*, Livermore, CA.

- Large-scale parallel data analysis algorithms for scientific simulations
  - Developed large scale parallel topological data analysis techniques for in-situ (on-the-fly) topological feature extraction of scientific simulations by adding less than 1% overhead and a 10x increase in analysis frequency
  - Developed a C/C++/MPI-based library for parallel topological analysis of scientific simulations that is scalable to 16,000 nodes (256,000 CPU cores) and deployed on the world's fastest supercomputers (previous known implementation scaled to only 4,000 cores)
  - Developed theoretical proofs and formal descriptions of above techniques requiring conceptual understanding of topology
- Analysis and visualization of performance data obtained from profiling massively parallel simulation codes
  - Developed a module for a visualization tool called Boxfish to visualize network traffic on a supercomputer interconnect using Python, PyOpenGL that helped design better communication strategies in parallel applications

Summer '12,'13, '14,'15 **Research Intern (Data Analysis Group)**, *Lawrence Livermore National Laboratory*, Livermore, CA.

- Topological analysis of scientific simulations
  - Developed scalable algorithm for parallel merge tree computation for topological analysis using C++/MPI
  - Gained experience in working with large parallel high performance computing systems
- Analysis of complex performance data obtained after profiling large scale parallel simulations
  - Performed canonical correlation analysis (CCA) of performance data and simulation output to identify performance metrics that are impacted by several scientific data fields in the simulation

2009 -'11 **System Software Engineer (CUDA Profiler Team)**, *NVIDIA Corporation*, Pune, India.

- Responsibilities involved developing functionality in the CUDA driver (like exposing hardware performance counters and GPU machine code instrumentation of CUDA kernels) to support key profiler features
- Gained deep understanding and knowledge of the GPU architecture, and CUDA performance optimizations
- Conducted several CUDA workshops and training seminars

2008 -'09 **System Software Intern (Video Team)**, *NVIDIA Corporation*, Pune, India.

- Developed a Windows application using C/C++/CUDA that performed intelligent scaling (resizing) of digital images to reduce the space required for storage
- Developed a hand gesture recognition system — Ported parts of the OpenCV AdaBoost supervised learning routines to CUDA to accelerate the classifier training process – reduced training time from several days to within a day

Programming Skills

C, C++, Large scale distributed computing using MPI, GPU Computing - CUDA, Python

Academic Projects

Fall '13 **Machine Learning: Scalable Domain Adaptation via Intelligent Sampling**.

- Reduced the training set by using soft clustering based sampling of unlabeled data for Domain Transfer Multiple Kernel Learning (DTMKL) technique achieving a 3x speedup in training time with less than 1% reduction in accuracy
- Applied the above technique to obtain classifiers for detecting spam emails for different users using labeled data from a single user (Email dataset : http://www.ecmlpkdd2006.org/challenge.html)

Awards and Achievements

Aug '15 Outstanding Poster Award, LLNL Student Poster Symposium, 2015

Apr '15 One of the six fully-funded invited students at The Salishan Conference on High Speed Computing, 2015

Oct '14 Visiting Scholar at Extreme Computing Research Center, King Abdulah University of Science Technology, hosted by Prof. David Keyes

Aug '13 Fully-funded invited participant at Argonne Training Program on Extreme Scale Computing (ATPESC 2013), Chicago, IL

Nov '11 Best Poster – SCIx-2011, SCI Institute, Salt Lake City, UT

Jan '09 2nd Place in Software Project Competition for Hand Gesture Recognition System at IIT-Bombay, Mumbai, India

Publications

Peer-Reviewed Papers

[1] **A. G. Landge**, P.-T. Bremer, A. Gyulassy, and V. Pascucci. Notes on the distributed computation of merge trees on cw-complexes. In *Topology-Based Methods in Visualization and Analysis 2015*, TopoInVis'15, 2015. (accepted for workshop presentation).

[2] **A. G. Landge**, V. Pascucci, A. Gyulassy, J. C. Bennett, H. Kolla, J. Chen, and P.-T. Bremer. In-situ feature extraction of large scale combustion simulations using segmented merge trees. In *Proc. of the International Conf. for High Performance Computing, Networking, Storage and Analysis*, SC '14, 2014.

[3] M. Gamell, I. Rodero, M. Parashar, J. C. Bennett, H. Kolla, J. Chen, P.-T. Bremer, **A. G. Landge**, A. Gyulassy, P. McCormick, S. Pakin, V. Pascucci, and S. Klasky. Exploring power behaviors and trade-offs of in-situ data analytics. In *Proc. of the International Conf. on High Performance Computing, Networking, Storage and Analysis*, SC '13, 2013.

[4] **A. G. Landge**, J. A. Levine, K. E. Isaacs, A. Bhatele, T. Gamblin, M. Schulz, S. H. Langer, P.-T. Bremer, and V. Pascucci. Visualizing network traffic to understand the performance of massively parallel simulations. *IEEE Trans. on Visualization and Computer Graphics (TVCG)*, Dec 2012. (InfoVis'12).

[5] A. Bhatele, T. Gamblin, S. H. Langer, P.-T. Bremer, E. W. Draeger, B. Hamann, K. E. Isaacs, **A. G. Landge**, J. A. Levine, V. Pascucci, M. Schulz, and C. H. Still. Mapping applications with collectives over sub-communicators on torus networks. In *Proc. of the International Conf. on High Performance Computing, Networking, Storage and Analysis*, SC '12, 2012.

Posters

[1] **A. G. Landge**, P.-T. Bremer, and V. Pascucci. Scalability and power efficiency of in-situ analysis workflows. *The Salishan Conference on High Speed Computing*, Apr 2015. INVITED POSTER.

[2] K. E. Isaacs, **A. G. Landge**, T. Gamblin, P.-T. Bremer, V. Pascucci, and B. Hamann. Exploring performance data with boxfish. *SC12, Salt Lake City, UT*, Nov 2012.

[3] **A. G. Landge**, J. A. Levine, K. E. Isaacs, A. Bhatele, T. Gamblin, M. Schulz, S. H. Langer, P.-T. Bremer, and V. Pascucci. Interactive linked visualizations for performance analysis of heterogeneous computing clusters. *GPU Technology Conference, San Jose, CA*, May 2012.

[4] **A. G. Landge**, J. A. Levine, K. E. Isaacs, A. Bhatele, T. Gamblin, M. Schulz, S. H. Langer, P.-T. Bremer, and V. Pascucci. Performance visualization of large scale simulations. *SCIx, Scientific Computing and Imaging Institute, Salt Lake City, UT*, Nov 2011. BEST POSTER.