Aaron Vose

avose@aaronvose.net
(865) 851-6988
Falmouth, MA

High-Performance Computing (HPC) Engineer

Talented software engineer possessing comprehensive experience developing and optimizing software for a diverse range of computer systems. Experience spans a wide range, from low-level firmware and Linux device drivers for embedded systems to massively parallel high-performance computing (HPC) systems and supercomputers such as Titan, the largest computer in the world at launch. – Also possesses in-depth experience using machine learning to solve real-world science and industry problems at scale. Project lead for multiple machine learning research efforts analyzing graph datasets with deep neural networks and genetic algorithms which resulted in successful publications and formation of new product offerings.

Highlighted Skills

20+ years C experience. Python, Bash, C++, Fortran, Linux, GCC, Clang, GDB, OpenGL, GNU Make, LaTeX. POSIX, Pthreads, OpenMP, MPI, x86, AVX, ARM, SVE, SLURM, LSF, Xbyak JIT, PAPI, IPv4, Lustre, Git. Deep Neural Networks (TensorFlow, Keras, Sonnet): Convolutional, Dense, Graph. Genetic Algorithms Expertise.

Work Experience

Roche Diagnostics
High-Performance Computing (HPC) Engineer,
February 2024 – Present.
– Work on HPC data analysis pipeline supporting next-generation nanopore gene sequencing devices.

NanoSemi Inc., a MaxLinear Company
Principal Software / Machine Learning Engineer,
February 2020 – February 2024.
– Extend Intel's C++ OneDNN library to support novel deep neural network layer types on CPU platforms.
– Create optimized C/C++ implementations of computational kernels for novel neural network algorithms.
– Work with just-in-time compilation (JIT) for optimized x86-64 kernels in assembly using SSE and AVX.
– Develop low-level firmware and Linux kernel drivers in C for high-end 5G wireless infrastructure products.

Maxwell Labs
Senior HPC Engineer,
July 2022 – March 2023.
– Install, configure, and manage hybrid CPU & GPU computer cluster running MPI and SLURM.
– Create parallel data processing pipelines to convert protein and drug data into training sets for DNNs.
– Develop novel neural networks trained on molecular and protein data for drug discovery workflows.

Cray Inc., a Hewlett Packard Enterprise Company
Software Engineer and Machine Learning Engineer,
July 2013 – January 2020.
– Optimize codes for clusters and supercomputers like Titan, the fastest computer in the world at launch.
– Create OpenMP run-time library to evaluate performance of novel hardware synchronization primitives.
– Create novel genetic algorithms for neural network hyperparameter optimization as core of new products.
– Create novel deep neural networks written in Python and TensorFlow for analysis of graph datasets.
– Research for the Exascale Computing Project to help design the next generation of supercomputers.
– Optimize scientific benchmarks like XSBench, RSBench, and Lulesh for vectorization with ARM SVE.

Joint and National Institutes for Computational Science
Research Associate and Intern,
May 2010 – January 2012.
– Optimize I/O patterns produced by C++ bioinformatics tools such as NCBI's BLAST and PSI-BLAST; improvements enable the scaling of jobs to many thousands of nodes on the Kraken supercomputer.
– Optimize molecular docking software Dock6 to enable large-scale runs for drug discovery workflows.
– Write framework in C/C++ for job scheduling and I/O management with shared memory segments.

Academic History

University of Tennessee
Research Assistant, Graduate Teaching Assistant, and STARS Alliance Mentor,
August 2002 – July 2013.
– Complete numerous research projects, most of which resulted in peer-reviewed journal publications.
– Research projects in the area of ecology and evolutionary biology with a focus on adaptive radiation.

M.S., Computer Science, Summer 2013.
B.S., Computer Science, Spring 2010.
University of Tennessee, Knoxville.
Cum Laude

CCDA - Cisco Certified Design Associate, 2001.
CCNA - Cisco Certified Network Associate, 2000.

Selected Publications

Presentations

(2011) "**Petascale Genomic Sequence Search.**"
Rekepalli, B.; Vose, A. *Proceedings of The 11th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*, Newport Beach, California.

Books

(2017) "**Programming for Hybrid Multi/Manycore MPP Systems.**"
John Levesque and Aaron Vose, *Chapman and Hall / CRC.* Computational Science Series, ISBN 9781439873717, Taylor & Francis.

Papers

(2021) "**The Convergence of HPC, AI and Big Data in Rapid-Response to the COVID-19 Pandemic.**"
Sukumar, S.R., Balma, J.A., Rickett, C.D., Maschhoff, K.J., Landman, J., Yates, C.R., Chittiboyina, A.G., Peterson, Y.K., Vose, A., Byler, K. and Baudry, J. *Smoky Mountains Computational Sciences and Engineering Conference* (pp. 157-172). Springer, Cham.

(2020) "**Evaluating Gather and Scatter Performance on CPUs and GPUs.**"
Patrick Lavin, Jeffrey Young, Richard Vuduc, Jason Riedy, Aaron Vose, and Daniel Ernst *The International Symposium on Memory Systems (MEMSYS 2020), Washington, DC, USA*. ACM, New York, NY, USA.

(2019) "**PharML.Bind: Pharmacologic Machine Learning for Protein-Ligand Interactions.**"
Aaron D. Vose, Jacob Balma, Damon Farnsworth, Kaylie Anderson, and Yuri K. Peterson *arXiv.org* 1911.06105

(2019) "**Recombination of Artificial Neural Networks.**"
Aaron Vose, Jacob Balma, Alex Heye, Alessandro Rigazzi, Charles Siegel, Diana Moise, Benjamin Robbins, and Rangan Sukumar *arXiv.org* 1901.03900

(2018) "**Interactive Distributed Deep Learning with Jupyter Notebooks.**"
Steve Farrell, Aaron Vose, Oliver Evans, Matthew Henderson, Shreyas Cholia, Fernando Perez, Wahid Bhimji, Shane Canon, Rollin Thomas, and Prabhat *International Conference on High Performance Computing* (pp. 678-687) Springer, Cham.

(2016) "**An MPI/OpenACC implementation of a high-order electromagnetics solver with GPUDirect communication.**"
Otten, M., Gong, J., Mametjanov, A., Vose, A., Levesque, J., Fischer, P. and Min, M. *International Journal of High Performance Computing Applications.*

(2016) "**Massively parallel and linear-scaling algorithm for second-order Moller-Plesset perturbation theory applied to the study of supramolecular wires.**"
Kjaergaard, T., Baudin, P., Bykov, D., Eriksen, J.J., Ettenhuber, P., Kristensen, K., Larkin, J., Liakh, D., Pawlowski, F., Vose, A., and Wang, Y.M. *Computer Physics Communications.*

(2015) "**A case study of CUDA FORTRAN and OpenACC for an atmospheric climate kernel.**"
M. Norman, J. Larkin, A. Vose, and K. Evans *Journal of Computational Science.* Vol 9: 1-6.

(2014) "**Tri-Hybrid Computational Fluid Dynamics on DoE's Cray XK7, Titan.**"
A. Vose, B. Mitchell, J. Levesque 2014 *Cray User Group* (CUG).