

Dorian Cecil Arnold, Ph.D.

☎ +1 505 688 1897 • ✉ dorian.arnold@gmail.com • 🌐 computerscience.emory.edu/~darnold
🔗 2jEFz5EAAAAJ • in [drdorianarnold](#) • 🆔 0000-0002-9425-1884

Technical Interests: Computer System Software R&D, including AI- and data-driven approaches.

Other Interests: University Leadership; Computing and STEM Education; Inclusive Workforce Development.

Scientist, Architect, Manager, Administrator, Leader!

- 30 Years Distributed, Runtime, and Operating Systems Software R&D, including HPC Systems, Fault-tolerance, Debugging and Performance Analysis, and Middleware.
- 18+ Years Small and large team Management with many direct reports; Developing and executing strategic vision; Developing and managing multi-million dollar budgets, including six and seven figure contracts.
- 8 Years Directing a complex multi-department graduate program; Overseeing 200% and 100% growth in students and faculty, respectively; Student Recruitment, Admission and Progress; Curricular Revisions; and Professional Development.

Technical Proficiencies and Highlights

- ★★★★ C/C++, Linux/Unix, Performance Tools, Telemetry, Job Schedulers, Resource Managers
- ★★★★☆ Java, Python, MPI, Debuggers, Software Testing, Parallel File Systems, Computer System Simulation
- ★★☆☆ OpenMP, Matlab, Fortran, Scientific and Numerical Libraries, Scrum, Agile Development
- ☆☆☆☆ Large Language Models, Kubernetes/Docker, Workflow Orchestration
- 70+ refereed research papers and one patent
- 2400+ citations. h-index: 26; i40-index: 13; i10-index: 41
- 2 R&D Top 100 awards
- \$13M of collaborative sponsored research, including \$3.4M of PI funding
- Partnerships strong industry, DOE lab, and university collaborations, including \$1.6M corporate funding
- Software production software packages deployed on leadership platforms across the world
- Renown Distinguished ACM Speaker (2017-2020); IEEE Senior Member; many Invited Keynote/Plenary Talks

Professional Appointments

- 2017 – **Associate Professor**, *Computer Science*, Emory University
- 2018 – **Director of Graduate Studies**, *Computer Science and Informatics*, Emory University
- 2009 – 2017 **Assistant/Associate Professor**, *Computer Science*, University of New Mexico
- 2013 **Summer Faculty**, *Department of Scalable System Software*, Sandia National Labs
- 2011 – 2013 **Affiliate Research Scientist**, *Ultrascale Systems Research Center*, New Mexico Consortium
- 2009 – 2010 **Visiting Scientist**, *Ctr. Applied Scientific Computing*, Lawrence Livermore National Lab
- 2006 **Technical Scholar**, *Ctr. Applied Scientific Computing*, Lawrence Livermore National Lab
- 1999 – 2001 **Research Associate**, *Innovative Computing Laboratory*, University of Tennessee

Select Projects and Project Supervision

- Software Synthesis Statistical and AI-based methods for synthesizing *Application Analogs*, proxies, benchmarks, traces, or other workload representatives of real applications. (Ongoing SNL collaboration.)

Telemetry	Practical monitoring optimizations, via data-driven principles for configuring spatial (which sensors) and temporal (sampling frequency) monitoring and data reductions. (Ongoing TACC collaboration.)
Resilience	Studied various aspects of fault-tolerance, primarily using software prototypes and simulation to evaluate checkpoint/restart protocols under various configurations, workloads, and platform environments. This project resulted in many high-impact research publications and the training of multiple PhD, MS, and BS students. (Past collaboration with Sandia National Labs and University of Tennessee.)
Software Testing	Managed team of software engineers and students in the design and development of production software regression and specification compliance tests for Cray Programming Environment suite, including C, C++ and Fortran tests for Cray Compilers, MPI, OpenMP, SHMEM, and CUDA. (2014 – 2020, supported by Hewlett Packard Enterprise, formerly Cray Inc.)
Overlay Networks	Founding researcher/developer of the multicast/reduction network, a tree-based overlay network for flexible, reliable data multicast and aggregation. MRNet has improved by orders of magnitude the scalability of many tools and applications and is used at laboratories and universities worldwide.
Scalable Debugging	Founding researcher/developer of the stack trace analysis tool (STAT), which uses MRNet to organize application stack traces into spatiotemporal call graph prefix trees for scalable debugging. STAT won an R&D Magazine Top 100 award in 2011 and remains in production use on large HPC systems today.
Grid Computing	Technical lead and developer of the network-enabled solvers (NetSolve), a client-server middleware for remote access to complex hardware and software resources. NetSolve has been used in many real applications across many domains and won an R&D Magazine Top 100 award, in 2000.
Rollback Recovery	Developed the Checkpoint Library for Unix Based Systems (CLUBS), for user-transparent checkpoint/restart, integrating optimizations like efficient file formats and exploiting copy-on-write semantics.
PhD Theses	W. Schonbein (2020), "Intelligent Networks for High Performance Computing" S. Gutiérrez (2018), "Adaptive Parallelism for Coupled, Multithreaded Message-Passing Programs" D. Ibtsham (2017), "Improving Large Scale Application Performance via Data Movement Reduction" T. Groves (2017), "Characterizing and Improving Power and Performance in HPC Networks"
MS Theses	Y. Wang (2016), "Generating Traces of Application Behavior Using Generative Adversarial Networks" B. Klemme (2016), "Software Testing of Parallel Programming Frameworks" L. Frey (2014), "Introducing Modularity into Complex Software Test Suite Frameworks" J. Goehner (2011), "Efficiently Bootstrapping Extreme Scale Software Systems"
BS Theses	M. Lafrance (2026), "On Checkpoint/Restart Viability for Exascale High Performance Computing" C. Chance (2022), "Zoom Audio Transcription Accuracy for African American Vernacular English" H. Feng (2022), "PADL: Persistent Application Data Library" Y. Jiang (2022), "Synthetic Generation of Datasets With Complex Attributes"

Select Impactful Papers

- [SC11] Ferreira et al and Arnold. Evaluating ... Process Replication Reliability for Exascale Systems.
- [SC08] Lee, Ahn, Arnold et al. Lessons Learned at 208k: Towards Debugging Millions of Cores.
- [IPDPS07] Arnold et al. Stack Trace Analysis for Large Scale Applications.
- [SC03] Roth, Arnold, and Miller. MRNet: A Software-based Multicast/reduction Network for Scalable Tools.
- [ICPP12] Ibtsham, Arnold et al. On ... Compression for Reducing the Overheads of Checkpoint/restart ...

Professional Preparation

- 2008 **Ph.D., Computer Sciences**, *University of Wisconsin*, Advisor: Dr. Barton Miller
- 1998 **M.S., Computer Science**, *University of Tennessee*, Advisor: Dr. James Plank
- 1996 **B.S. summa cum laude, Mathematics, Computer Science**, *Regis University*
- 1994 **A.S., Mathematics, Physics, Chemistry**, *St. John's Junior College (Belize)*