

Stephen Herbein

Computer Scientist, Graduate Researcher
University of Delaware

☎ 302.689.3894
✉ stephen@herbein.net
<https://herbein.net>

Education

- 2014-present **Ph.D., Computer Science**, *University of Delaware*, GPA 4.0.
Advisor: Dr. Michela Taufer
- 2016 **M.S., Computer Science**, *University of Delaware*, GPA 4.0.
- 2014 **Honors B.S. with Distinction, Computer Science**, *University of Delaware*.
GPA 3.932; Magna Cum Laude

Research Interests

- Next-generation resource management and job scheduling
- I/O-aware scheduling
- Benchmarking and auto-tuning parallel I/O

Research Experience

- 2014-present **University of Delaware**, *Graduate Research Assistant*.
Creating next-generation job schedulers for HPC clusters utilizing the Flux resource manager
- 2014-present **Lawrence Livermore National Laboratory**, *Research Intern*.
(Summers) **2014:** Developed a discrete event simulator for the Flux resource manager
2015: Developed an automatic job aggregator and hierarchical scheduler, resulting in a 4x speed up over the existing scheduler
2016: Added dynamicity to my hierarchical scheduler to eliminate resource fragmentation
2017: Integrated my hierarchical scheduler with the Uncertainty Quantification Pipeline (UQP), resulting in a 37% improvement in workload runtime
- 2013-2014 **University of Delaware**, *Undergraduate Research Assistant*.
Applied auto-tuning techniques to improve I/O performance of scientific applications on clusters with a Lustre filesystem
- Summer 2013 **Oak Ridge National Laboratory**, *Research Intern (DOE SULI)*.
Integrated the ADIOS I/O library into the QMCPack scientific application to improve performance on ORNL's supercomputer, Titan
- 2012-2013 **University of Delaware**, *Undergraduate Research Assistant*.
Improved the efficiency of GPU accelerated codes in scientific workflows on large-scale clusters
- Summer 2012 **University of Houston**, *Research Intern (NSF REU)*.
Optimized VolpexMPI library for use on large-scale clusters

Awards & Honors

- Summer 2016 Best Poster - Annual LLNL Poster Symposium
- Spring 2016 HPDC Student Travel Award for \$1000 - HPDC
- Spring 2016 Invited to attend and present at Salishan Conference - LLNL
- Fall 2014 ICPADS Student Travel Award for \$1000 - ICPADS/TCPP
- Spring 2014 Outstanding Senior Student Award - CIS Department
- Fall 2013 Undergraduate Student Research Competition Award - ACM
- Spring 2013 Outstanding Junior Student Award - CIS Department
- Fall 2012 General Honors Award - University of Delaware

Peer-Reviewed Publications

Conference Papers

- [1] Ryan McKenna, **Stephen Herbein**, Adam Moody, Todd Gamblin, and Michela Taufer. Machine Learning Predictions of Runtime and IO Traffic on High-end Clusters. In *Proceedings of the IEEE Cluster Conference*, page 1–4, September 2016. (Acceptance Rate: 57/162, 35%).
- [2] **Stephen Herbein**, Dong H. Ahn, Don Lipari, Thomas R.W. Scogland, Marc Stearman, Mark Grondona, Jim Garlick, Becky Springmeyer, and Michela Taufer. Scalable I/O-Aware Job Scheduling for Burst Buffer Enabled HPC Clusters. In *Proceedings of the 25th International Symposium on High-Performance Parallel and Distributed Computing, HPDC*, June 2016. (Acceptance Rate: 20/129, 15.5%).
- [3] **Stephen Herbein**, Ayush Dusia, Aaron Landwehr, Sean McDaniel, Jose Monsalve, Yang Yang, Seetharami R. Seelam, and Michela Taufer. Resource Management for Running HPC Applications in Container Clouds. In *Proceedings of 31st International Supercomputing Conference, ISC*, Leipzig, Germany, June 2016. (Acceptance Rate: 25/60, 40%).
- [4] Michael Matheny, **Stephen Herbein**, Norbert Podhorszki, Scott Klasky, and Michela Taufer. Using Surrogate-based Modeling to Predict Optimal I/O Parameters of Applications at the Extreme Scale. In *Proceedings of the 20th IEEE International Conference on Parallel and Distributed Systems, ICPADS*, pages 568–575, Dec 2014. (Acceptance Rate: 96/322, 29.8%).
- [5] **Stephen Herbein**, Michael Matheny, Matthew Wezowicz, Jaron Krogel, Jeremy Logan, Jeongnim H. Kim, Scott Klasky, and Michela Taufer. Performance Impact of I/O on QMCPack Simulations at the Petascale and Beyond. In *Proceedings of the 16th IEEE International Conference on Computational Science and Engineering, CSE*, pages 92–99, Dec 2013.
- [6] Samuel Schlachter, **Stephen Herbein**, Michela Taufer, Shuching Ou, Sandeep Patel, and Jeremy S. Logan. Efficient SDS Simulations on Multi-GPU Nodes of XSEDE High-End Clusters. In *Proceedings of the 9th IEEE International Conference on e-Science, eScience*, pages 116–123. IEEE Computer Society, Oct 2013. (Acceptance Rate: 41/98, 41.8%).

Journal Papers

- [7] **Stephen Herbein**, Sean McDaniel, Norbert Podhorszki, Jeremy Logan, Scott Klasky, and Michela Taufer. Performance characterization of irregular I/O at the extreme scale. *Parallel Computing*, 51:17 – 36, Jan 2016. Special Issue on Parallel Programming Models and Systems Software for High-End Computing.
- [8] Sam Schlachter, **Stephen Herbein**, Shuching Ou, Jeremy S. Logan, Sandeep Patel, and Michela Taufer. Pursuing Coordinated Trajectory Progression and Efficient Resource Utilization of GPU-Enabled Molecular Dynamics Simulations. *IEEE Design Test*, 31(1):40–50, Feb 2014.

Workshop Papers

- [9] Robert Searles*, **Stephen Herbein***, and Sunita Chandrasekaran. A Portable, High-Level Graph Analytics Framework Targeting Distributed, Heterogeneous Systems. In *Third Workshop on Accelerator Programming Using Directives (WACCPD)*, Super Computing, Salt Lake City, UT, November 2016. *Equal contribution.
- [10] **Stephen Herbein**, Scott Klasky, and Michela Taufer. Benchmarking the Performance of Scientific Applications with Irregular I/O at the Extreme Scale. In *2014 43rd International Conference on Parallel Processing Workshops, ICCPW*, pages 292–301, Sept 2014.

Posters

- Best Poster Candidate** [11] **Stephen Herbein**, Tapasya Patki, Dong H. Ahn, Don Lipari, Tamara Dahlgren, David Domyancic, and Michela Taufer. Fully Hierarchical Scheduling: Paving the Way to Exascale Workloads. In *Proceedings of the 29th ACM/IEEE International Conference for High Performance Computing and Communications conference*, SC, Denver, CO, November 2017. (Acceptance Rate: 98/169 58.0%.
- [12] **Stephen Herbein**. Exploring the Trade-off Space of Hierarchical Scheduling for Very Large HPC Centers. In *Proceedings of the 27th ACM/IEEE International Conference for High Performance Computing and Communications conference*, SC, Austin, TX, November 2015.
- [13] Sean McDaniel, **Stephen Herbein**, and Michela Taufer. A Two-Tiered Approach to I/O Quality of Service in Docker Containers. In *Proceedings of the IEEE International Conference on Cluster Computing*, CLUSTER, pages 490–491, Sept 2015.
- [14] **Stephen Herbein**. Enabling Fine-grained Gathering of Scientific Data in QMCPack Simulations on Titan. In *Proceedings of the 25th ACM/IEEE International Conference for High Performance Computing and Communications conference*, SC, Denver, CO, November 2013.
- [15] Matthew Wezowicz, Michael Matheny, **Stephen Herbein**, Jeremy Logan, Jeognim Kim, Jaron Krogel, Scott Klasky, and Michela Taufer. Predictions of Large-scale QMCPack I/Os on Titan using Skel. In *Proceedings of the 25th ACM/IEEE International Conference for High Performance Computing and Communications conference*, SC, Denver, CO, November 2013.

Undergraduate Senior Thesis

- [16] **Stephen Herbein**. *Benchmarking and Auto-tuning I/O Intensive Applications at the Extreme Scale*. Thesis, University of Delaware, Newark, DE, May 2014.

Invited Talks

- December, 2015 **IBM Thomas J. Watson Research Center**,
Invited Talk: Towards Resource QoS In Container Clouds.
6th Student Workshop on Cloud and Data Services

Professional Services & Activities

Committee Member

- 2017 Lead Student Volunteer - Information Booth at SuperComputing (SC)

Conference (Sub-)Reviewer

- 2017 SC, IPDPS
2016 PARCO, Cluster, COM-HPC, SBAC-PAD
2015 WoC

Other

- 2012-2016 Student Volunteer at SuperComputing (SC)

Memberships

- Association for Computing Machinery (ACM)
- Institute of Electrical and Electronics Engineers (IEEE)

Skills

- General Python, C, C++, MPI, Spark, OpenMP, CUDA, OpenCL, Java, git, svn, CMake.
Python-specific numpy, scipy, scikit-learn, matplotlib, pyspark, multiprocessing, requests, beautifulsoup
Web Development Server-side: PHP, MySQL; Client-side: HTML, CSS, Javascript.

Open-Source Software Projects

- flux-core Next-generation resource manager
- flux-sched Next-generation batch job scheduler