

Paul D. Marshall

Department of Computer Science
University of Colorado at Boulder
430 UCB
Boulder, CO 80309

Phone: 720.310.5212
Fax: 303.492.2844
paul.marshall@colorado.edu
<http://csc.cs.colorado.edu/~marshalp/>
U.S. Citizen

Education

Ph.D. Computer Science, University of Colorado at Boulder (expected May 2013)
M.S. Computer Science, University of Colorado at Boulder, 2010
B.A. Computer Science and Mathematics, Magna Cum Laude, Augustana College, Sioux Falls, 2006

Teaching Experience

Teaching Assistant, *CSCI 4576/5576: High Performance Scientific Computing*
Department of Computer Science, University of Colorado at Boulder, Fall 2008
Teaching Assistant, *CSCI 1300: Computer Science 1: Programming*
Department of Computer Science, University of Colorado at Boulder, Fall 2007

Research and Professional Positions

Research Assistant *Spring 2008–Present*
Computational Science Center *Department of Computer Science*
Henry M. Tufo *University of Colorado at Boulder*
Investigating algorithms and models to create efficient and responsive large-scale elastic computing environments that utilize infrastructure clouds. Also support scientific researchers by focusing on techniques and tools to simplify their workflows and increase efficiency.

Developer Intern II *Summer 2012*
Cloud Business Unit *Rackspace Hosting, Inc.*
Developer for Rackspace Cloud Databases, a database-as-a-service offering built on OpenStack Nova that provides a scalable relational database service.

Student Assistant *Spring 2010–Summer 2012*
Student Visitor *Spring 2008–Fall 2009*
Computer Science Section *National Center for Atmospheric Research*
Support and troubleshooting of NCAR's Asteroseismic Modeling Portal and Frost, an 8192-core Blue Gene/L supercomputer on the TeraGrid, a distributed and open scientific grid-computing infrastructure offering over 2.5 petaflops of compute capability.

Intern *Summers 2009, 2010, September 2010–May 2011*
Computation Institute *Argonne National Laboratory and*
Kate Keahey *University of Chicago*
Extended the open source Nimbus cloud computing toolkit to deploy backfill virtual machines on idle cloud resources to increase resource utilization without sacrificing the ability of the cloud to service on-demand requests. Also developed a prototype resource manager to dynamically extend existing physical clusters with cloud resources based on demand.

Cisco Choice Intern *Summer 2008*
Linux OS / Platform Team *Cisco Systems, Inc.*
Developed a benchmark tool to simulate the disk I/O of Cisco Unified Communications Manager (CUCM) on Linux, emphasized disk I/O performance in a virtual environment.

Programmer / Analyst
Information Technology Services

2005-2007
Augustana College

Software developer and administrator for a .NET-based Web portal. Along with programming responsibilities, assisted in basic system administration for AIX, Linux, and Windows servers.

NASA Space Grant Student Fellow
NASA EPSCoR Student Fellow
Daniel Swets

2003-2005
Summer 2005
Augustana College

Designed and developed a software suite with two other students to calculate seasonality metrics on NDVI satellite images. Also, designed and developed an application to calculate the Mann-Kendall statistical test on NDVI satellite images.

Journal Publications

J. Harris, J. Caporaso, J. Walker, J. Spear, N. Gold, C. Robertson, P. Hugenholtz, J. Goodrich, D. McDonald, D. Knights, P. Marshall, H. Tufo, R. Knight, and N. Pace, *Phylogenetic Stratigraphy in the Guerrero Negro Hypersaline Microbial Mat*, in the ISME Journal, 2012.

Magazine Articles

P. Marshall, H. Tufo, and K. Keahey, *High-Performance Computing and the Cloud: A Match Made in Heaven or Hell?*, in ACM XRDS, Spring 2013.

Conference and Workshop Publications

D. Duplyakin, P. Marshall, K. Keahey, H. Tufo, and A. Alzabarah, *Rebalancing in a Multi-Cloud Environment*, proceedings of the 4th Workshop on Scientific Cloud Computing and the 22nd ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), New York City, NY, June 2013 (in press).

P. Marshall, H. Tufo, K. Keahey, D. LaBissoniere, and M. Woitaszek, *Architecting a Large-Scale Elastic Environment: Recontextualization and Adaptive Cloud Services for Scientific Computing*, proceedings of the 7th International Conference on Software Paradigm Trends (ICSOFIT), Rome, Italy, July 2012.

P. Marshall, H. Tufo, and K. Keahey, *Provisioning Policies for Elastic Computing Environments*, in proceedings of the 9th High-Performance Grid and Cloud Computing Workshop and the 26th IEEE International Parallel and Distributed Processing Symposium (IPDPS), Shanghai, China, May 2012.

P. Marshall, K. Keahey, and T. Freeman, *Improving Utilization of Infrastructure Clouds*, proceedings of the 11th IEEE/ACM International Conference on Cluster, Cloud, and Grid Computing (CCGrid), Newport Beach, California, May 2011.

P. Marshall, K. Keahey, and T. Freeman, *Elastic Site: Using Clouds to Elastically Extend Site Resources*, proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud, and Grid Computing (CCGrid), Melbourne, Australia, May 2010.

P. Marshall, M. Oberg, N. Rini, T. Voran, and M. Woitaszek, *Virtual Clusters for Hands-On Linux Cluster Construction Education*, proceedings of the 11th LCI International Conference on High-Performance Clustered Computing, Pittsburgh, Pennsylvania, March 2010, Best Student Paper Award.

P. Marshall, M. Woitaszek, H. M. Tufo, R. Knight, D. McDonald, J. Goodrich, and J. Widmann, *Ensemble Dispatching on a Blue Gene/L for a Bioinformatics Knowledge Environment*, proceedings of the 2nd ACM Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS), Portland, Oregon, November 2009.

B. House, P. Marshall, M. Oberg, H. M. Tufo, and M. Woitaszek, *Grid Service Hosting on Virtual Clusters*, proceedings of the 9th IEEE/ACM International Conference on Grid Computing (Grid 2008), Tsukuba, Japan, September 2008.

P. Marshall and D. Swets, *Developing Parallel Algorithms for Seasonal Metrics Extraction*, proceedings of the South Dakota Academy of Science, Sioux Falls, South Dakota, April 2005.

Presentations

Elastic Environments for Scientific Computing in the Cloud, Bioinformatics Supergroup, BioFrontiers, University of Colorado at Boulder, March 2013.

Scientific Computing in the Cloud, presented in CSCI 4576/5576: High-Performance Scientific Computing, Department of Computer Science, University of Colorado at Boulder, December 2012.

Dynamically Extend Your Cluster with Infrastructure Clouds, presented at Infrastructure Clouds and Elastic Services tutorial at the 2012 ACM/IEEE Conference on High Performance Computing, Networking, and Storage (SC12), Salt Lake City, Utah, November 2012.

Architecting a Large-Scale Elastic Environment: Recontextualization and Adaptive Cloud Services for Scientific Computing, 7th International Conference on Software Paradigm Trends (ICSOFT), Rome, Italy, July 2012.

Provisioning Policies for Elastic Computing Environments, 9th High-Performance Grid and Cloud Computing Workshop and the 26th IEEE International Parallel and Distributed Processing Symposium (IPDPS), Shanghai, China, May 2012.

How to Stop Worrying and Get Involved with Open Source, Department of Computer Science, Carroll College, Helena, Montana, March 2012.

Improving Utilization of Infrastructure Clouds, 11th IEEE/ACM International Conference on Cluster, Cloud, and Grid Computing (CCGrid), Newport Beach, California, May 2011.

How to Dynamically Extend Your Site with Hundreds of Amazon EC2 Instances, Argonne National Laboratory, Chicago, Illinois, July 2010.

Elastic Site: Using Clouds to Elastically Extend Site Resources, 10th IEEE/ACM International Conference on Cluster, Cloud, and Grid Computing (CCGrid), Melbourne, Australia, May 2010.

Piled Higher and Deeper: Life in Graduate School, Department of Computer Science, Carroll College, Helena, Montana, March 2010.

Virtual Clusters for Hands-On Linux Cluster Construction Education, 11th LCI International Conference on High-Performance Clustered Computing, Pittsburgh, Pennsylvania, March 2010.

Ensemble Dispatching on a Blue Gene/L for a Bioinformatics Knowledge Environment, 2nd ACM Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS), Portland, Oregon, November 2009.

Statistical Software for Time Series Analysis, South Dakota Academy of Sciences Annual Meeting, Chamberlain, South Dakota, April 2006.

Developing Parallel Algorithms for Seasonality Analysis, South Dakota Academy of Sciences Annual Meeting, Sioux Falls, South Dakota, April 2005.

Tutorials

Using and Building Infrastructure Clouds for Science, presented at the 2011 ACM/IEEE Conference on High Performance Computing, Networking, and Storage (SC11).

FutureGrid: An Introduction to Nimbus, presented at the 11th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2011).

Linux Cluster Construction, presented at the 2009 and 2010 ACM/IEEE Conference on High Performance Computing, Networking, Storage (SC09 and SC10).

Quick Start Linux Cluster Construction, presented at the 10th and 11th LCI International Conference on High-Performance Clustered Computing (LCI 2009 and LCI 2010).

Posters

Nimbus Elastic Scaling in the Clouds, the 2nd International Conference on Cloud Computing Technology and Science, Indianapolis, Indiana, December 2010.

Developing Parallel Algorithms for Seasonality Analysis, Council on Undergraduate Research Posters on the Hill, Washington, DC, April 2005.

Seasonality Metrics Extraction from Smoothed NDVI Images, Augustana College Symposium, Sioux Falls, South Dakota, April 2005.

Seasonality Metrics Extraction from Smoothed NDVI Images, South Dakota State Legislature Undergraduate Student Research Activities Poster Session, Pierre, South Dakota, February 2005.

Software

Contributor to RedDwarf, an open source database-as-a-service solution built on OpenStack Nova.
<http://github.com/stackforge/reddwarf>

Contributor to Nimbus, an open source Infrastructure-as-a-Service (IaaS) cloud computing toolkit.
<http://www.nimbusproject.org/>

Awards and Honors

Best Student Paper Award, The 11th LCI International Conference on High-Performance Clustered Computing for *Virtual Clusters for Hands-On Linux Cluster Construction Education*

CS Department Travel Award for presentations, Spring 2010, Fall 2011, Spring 2012

Clive Fraser Baillie Memorial Travel Award, CU-Boulder, Spring 2010, Fall 2011, Fall 2012

Graduate Student Research and Community Development Award, CU-Boulder, Spring 2009

Dean's Graduate Assistantship, CU-Boulder, Fall 2007

University Fellowship, CU-Boulder, Fall 2007

College of Engineering Fellowship, CU-Boulder, Fall 2007

Dean's List, Augustana College

Federated Insurance Scholarship

Professional Service and Activities

Contributor to the Science Clouds blog, 2010–Present
<http://scienceclouds.org/blog/>

Reviewer for the European Conference on Parallel Processing (Euro-Par), 2010

Augustana Computer Organization, Co-Founder, Augustana College, 2004–2006

Student Newspaper (*The Mirror*), Web Editor, Augustana College, 2003–2006

Information Technology Services Committee, Student Representative, Augustana College, 2005–2006

Strategic Planning Task Force Committee, Student Representative, Augustana College, 2005