

**Objective****Computer Science Conference, Journal and Workshop Reviewer Position****Education****North Carolina State University**

Raleigh, NC, U.S.

*PhD in Computer Science**Fall 2013 - Fall 2017*

- Research: Realized a distributed and high-performance data staging framework on supercomputers
- GPA: 3.917
- Advisor: Dr. Samatova

**Peking University**

Master in Software Engineering, Beijing China, September 2008 - May 2011

**ZZULI**

Bachelor in Computer Science, Zhengzhou China, September 2004 - May 2008

**Candidate Statement****Research Field**

My research focuses on scientific data management on supercomputers. The mission is to provide domain scientists with effective and efficient tools to analyze the increasingly large datasets generated by scientific simulations on modern supercomputers. Specifically, my dissertation is about constructing a data staging framework to enhance runtime Adaptive Mesh Refinement (AMR) data analysis performance on supercomputers. This work can run over thousands of compute nodes on supercomputers, handle up to 16,384 parallel access with good scalability, and orchestrate large volume of data across multiple memory hierarchies on supercomputers. It is capable of optimizing AMR data access performance by tracking data access patterns and performing prefetching, and distributing data in a network-topology-aware manner to reduce network access latency.

**Specialized Fields**

Given more than four years of experience in constructing high performance, data intensive, parallel and distributed systems, I have deep knowledge in computer memory, storage IO and networking. I also have years industry's intern and full time experience in the areas of enterprise software development and cloud computing.

**Published Papers****BigData 2016***First Author*

- Exploring Memory Hierarchy and Network Topology for Runtime AMR Data Sharing Across Scientific Applications

**ICPP 2016***Coauthor*

- In situ Storage Layout Optimization for AMR Spatio-temporal Read Accesses

**CCGrid 2016***First Author*

- AMRZone: A Runtime AMR Data Sharing Framework For Scientific Applications

**CCGrid 2016***Coauthor*

- Usage Pattern-Driven Dynamic Data Layout Reorganization

**HPC 2016***Coauthor*

- AMR-aware In Situ Indexing and Scalable Querying

**CLUSTER 2015***First Author*

- Exploring Memory Hierarchy to Improve Scientific Data Read Performance

## Wenzhao Zhang

### Work Experience

#### Renaissance Computing Institute

Chapel Hill, NC, U.S.

*Software Engineer; GENI Team; Intern; C++/Java, Linux*

*May 2017 - Current*

- Added new functions to Apache Mesos source code.
- Developed web services, which can schedule a framework over multiple Mesos clusters.
- Deployed Mesos over a cloud, with Docker applications.

#### Lawrence Berkeley National Lab

Berkeley, CA, U.S.

*Research Assistant; SDM Team; Intern; C, Linux*

*May 2016 - August 2016*

- Enabled an in-memory data staging framework to use SSDs as an overflow space, on supercomputers.
- Improved data access performance for the framework by up to 25%, via data prefetching and network-topology-aware data distribution.

#### Lawrence Berkeley National Lab

Berkeley, CA, U.S.

*Research Assistant; SDM Team; Intern; C, Linux*

*May 2015 - August 2015*

- Developed an in-memory data staging framework on supercomputers.  
It can run over thousands of nodes, and process up to 16,000 parallel requests with good scalability.

#### Samsung

Beijing, China

*Software Engineer; PaaS Cloud Team; Full Time; Java, Linux*

*August 2011 - June 2013*

- Debugged and developed new features for CloudFoundry.
- Developed a web-based management system for the cloud platform.
- Developed new services for the cloud platform.

#### Oracle

Beijing, China

*Software Engineer; Tuxedo CCE Team; Intern; Java, Linux*

*June 2009 - October 2010*

- Designed and updated a web-based workflow management system for the CCE team.
- Developed and executed test-cases for Tuxedo and related products on various OSs.