

NAWAB ALI

395 Dreese Laboratories
2015 Neil Avenue
Columbus, OH 43210-1277
<http://www.cse.ohio-state.edu/~alin>

Office: +1 - 614 - 592 - 1987
Home: +1 - 614 - 592 - 1987
Fax: +1 - 614 - 292 - 2911
alin@cse.ohio-state.edu

Education

PhD, Computer Science and Engineering

The Ohio State University, 09/2004 – present
Dissertation: *Rethinking I/O in High-Performance Computing Environments*
Advisor: Dr. P. Sadayappan
Expected Graduation Date: December, 2009

MS, Computer Science

University of Cincinnati, 09/2002 – 07/2004
Thesis : *Improving Paging Performance of Memory Intensive Applications*
Advisor: Dr. Yiming Hu

MSc. (Tech.), Information Systems

Birla Institute of Technology and Science, Pilani, India, 09/1997 – 07/2001
Graduated with Honors, First Division

Employment

Graduate Teaching Associate, Computer Science and Engineering

Faculty: Dr. Gojko Babic
Period: 06/2009 – present
Course: CSE 230

Resident Associate, Argonne National Laboratory

Supervisor: Dr. Robert Ross
Period: 09/2008 – 06/2009
Research Work: *Common HEC I/O Forwarding Scalability Layer*

Research Aide, Argonne National Laboratory

Supervisor: Dr. Robert Ross
Period: 06/2008 – 09/2008
Research Work: *Common HEC I/O Forwarding Scalability Layer*

Graduate Research Associate, Ohio Supercomputer Center

Faculty: Dr. P. Sadayappan
Period: 03/2007 – 03/2009
Research Work: *Applicability of Object-Based Storage Devices in Parallel File Systems*

Graduate Teaching Associate, Computer Science and Engineering

Faculty: Michelle Mallon
Period: 09/2006 – 03/2007
Course: CSE 200

Graduate Research Associate, Computer Science and Engineering

Faculty: Dr. Mario Lauria

Period: 09/2004 – 09/2005

Research Work: *SRB Enabled MPI-IO Library for Access to Remote Storage*

Member Technical Staff, Sun Microsystems, India

Period: 07/2001 – 08/2002

Project: *Sun iPlanet Portal Server*

Awards and Recognitions

- University Graduate Scholarship, University of Cincinnati, 2002 – 2004.

Refereed Conference Publications

1. **N. Ali**, P. Carns, K. Iskra, D. Kimpe, S. Lang, R. Latham, R. Ross, L. Ward, P. Sadayappan. “*Scalable I/O Forwarding Framework for High-Performance Computing Systems*”, IEEE International Conference on Cluster Computing, New Orleans, LA, August 2009.
2. **N. Ali**, A. Devulapalli, D. Dalessandro, P. Wyckoff, P. Sadayappan. “*Revisiting the Metadata Architecture of Parallel File Systems*”, Proceedings of the 3rd International Petascale Data Storage Workshop, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Austin, TX, November 2008.
3. **N. Ali**, A. Devulapalli, D. Dalessandro, P. Wyckoff, P. Sadayappan. “*An OSD-based Approach to Managing Directory Operations in Parallel File Systems*”, IEEE International Conference on Cluster Computing, Tsukuba, Japan, September 2008.
4. A. Devulapalli, D. Dalessandro, P. Wyckoff, **N. Ali**, P. Sadayappan. “*Integrating Parallel File Systems with Object-Based Storage Devices*”, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Reno, NV, November 2007.
5. A. Devulapalli, D. Dalessandro, P. Wyckoff, **N. Ali**. “*Attribute Storage Design for Object-based Storage Devices*”, 24th IEEE Conference on Mass Storage Systems and Technologies, San Diego, CA, September 2007.
6. **N. Ali**, M. Lauria. “*Improving the Performance of Remote I/O using Asynchronous Primitives*”, 15th IEEE International Symposium on High Performance Distributed Computing, Paris, France, June 2006.
7. **N. Ali**, M. Lauria. “*SEMPAR: High-Performance Remote Parallel I/O over SRB*”, 5th IEEE/ACM International Symposium on Cluster Computing and the Grid, Cardiff, UK, May 2005.

Technical Reports

1. **N. Ali**, P. Carns, K. Iskra, D. Kimpe, S. Lang, R. Latham, R. Ross, L. Ward, P. Sadayappan. “*Scalable I/O Forwarding Framework for High-Performance Computing Systems*”, The Ohio State University, Tech. Rep. OSU-CISRC-4/09-TR07, 2009.

2. **N. Ali**, A. Devulapalli, D. Dalessandro, P. Wyckoff, P. Sadayappan. “*Revisiting the Metadata Architecture of Parallel File Systems*”, The Ohio State University, Tech. Rep. OSU-CISRC-7/08-TR42, 2008.
3. A. Devulapalli, D. Dalessandro, **N. Ali**, P. Wyckoff. “*Attribute Storage Design for Object-based Storage Devices*”, The Ohio State University, Tech. Rep. OSU-CISRC-9/08-TR45, 2008.

Posters

1. **N. Ali**, A. Devulapalli, D. Dalessandro, P. Wyckoff. “*Redesigning Parallel File Systems Using Object-based Storage Devices*” National Science Foundation Collaborative Expedition Workshop #74, Arlington, VA, June 10, 2008.
2. **N. Ali**, P. Sadayappan. “*Rethinking I/O in High-Performance Computing Environments*” TCPP PhD Forum, IEEE International Parallel and Distributed Processing Symposium, Miami, FL, 2008.

Invited Talks

1. “*Rethinking I/O in High-Performance Computing Environments*”, Sandia National Laboratories, February 2009, Albuquerque, NM.
2. “*Revisiting the Metadata Architecture of Parallel File Systems*”, Proceedings of the 3rd International Petascale Data Storage Workshop, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Austin, TX, November 2008.
3. “*An OSD-based Approach to Managing Directory Operations in Parallel File Systems*”, LANS Seminar, Argonne National Laboratory, September 2008, Argonne, IL.
4. “*SEMPAR: High-Performance Remote Parallel I/O over SRB*”, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Tampa, FL, November 2006.

Research Projects

1. **Common HEC I/O Forwarding Scalability Layer:** (2008 – present) Future leadership-class machines will incorporate millions of processing elements, a significant increase in scale over existing systems. At the same time, storage device access rates are not increasing substantially and the access rate of the I/O system as a whole is not keeping pace with increases in raw compute power. Commercial parallel file systems, designed for competitiveness in the larger business market, are likewise being stretched to address leadership computing scales. A comprehensive software solution is needed to bridge the gap between processing trends and I/O systems so that leadership-class machines can most efficiently leverage the available storage resources. This solution must address the extreme degree of parallelism on the compute side, match the lightweight design of operating systems on these platforms, and aggregate I/O operations in order to better manage interaction with commercial parallel file systems. This project aims to design, build, and distribute a scalable, unified high-end computing I/O forwarding software layer.

2. **Applicability of Object-Based Storage Devices in Parallel File Systems:** (2007 – 2008) As storage systems evolve, the block-based design of today’s disks is becoming inadequate. As an alternative, Object-Based Storage Devices (OSDs) offer a view where the disk manages data layout and keeps track of various attributes about data objects. By moving functionality that is traditionally the responsibility of the host OS to the disk, it is possible to improve overall performance and simplify management of a storage system. The capabilities of OSDs will also permit performance improvements in parallel file systems, such as further decoupling metadata operations and thus reducing metadata server bottlenecks. This project examines the feasibility of OSDs for use in parallel file systems, in particular, discovering techniques to accommodate this high-performance usage model. We are currently investigating multiple aspects of the mismatch between the needs of a parallel file system, specifically PVFS2 and the capabilities of OSDs. Some of the open research topics in this area include mapping data to objects, metadata management, transport, caching, and reliability. This work is being done in collaboration with researchers at the Ohio Supercomputer Center.
3. **SRB Enabled MPI-IO Library for Access to Remote Storage:** (2004 – 2007) One of the challenges in high-performance computing is to provide users with reliable, remote data access in a distributed, heterogeneous environment. The increasing popularity of high-speed wide-area networks and centralized data repositories lead to the possibility of direct, high-speed access to remote data sets from within a parallel application. SEMPLAR is a remote, parallel I/O library that combines the standard programming interface of MPI-IO with the remote storage functionality of the SDSC Storage Resource Broker. SEMPLAR relies on parallel TCP streams to maximize the remote data throughput in a design that preserves the parallelism of the access all the way from the storage to the application.

Memberships

- Student member, Institute of Electrical and Electronics Engineers (IEEE).