Brian W. Barrett

CONTACT INFORMATION

Scalable System Software Group
Sandia National Laboratories
P.O. Box 5800 MS 1319
Albuquerque, NM 87185-1319

Phone: 505-284-2333
bwbarre@sandia.gov

RESEARCH INTERESTS

The design and implementation of high performance computing communication systems, including both advanced network interface designs and software communication paradigms. Research into advanced systems capable of supporting both traditional message passing applications and emerging graph-based informatics applications.

EDUCATION

Indiana University — Bloomington, IN
Ph.D. Computer Science — *March 2009*
Thesis: *One-Sided Communication for High Performance Computing Applications*
Advisor: Andrew Lumsdaine
Committee: Randall Bramley, Beth Plale, and Amr Sabry

Indiana University — Bloomington, IN
M.S. Computer Science — *August 2003*
Advisor: Andrew Lumsdaine

University of Notre Dame — Notre Dame, IN
B.S. Computer Science, cum Laude — *May 2001*

EXPERIENCE

Senior Member of Technical Staff — **Sandia National Laboratories**
October 2007 - present — *Albuquerque, New Mexico*
Research into advanced network design, particularly network interface adapters, for message passing. Research and development of large scale graph algorithms as part of the MTGL and PBGL graph libraries. Design of advanced computer architectures capable of supporting large scale graph informatics applications.

Technical Staff Member — **Los Alamos National Laboratory**
October 2006 - October 2007 — *Los Alamos, New Mexico*
Research and development work on the Open MPI implementation of the MPI standard. Focus on enhancements for the Road Runner hybrid architecture system, including high-performance heterogeneous communication.

Student Intern — **Los Alamos National Laboratory**
Summer 2006 — *Los Alamos, New Mexico*
Research and development work on the Open MPI implementation of the MPI standard. Implemented the MPI-2 one-sided specification within Open MPI. Co-developed a high performance point-to-point engine for interconnects that support MPI matching in the network stack. Implemented support for the Portals communication library within the new matching point-to-point engine.

Research Assistant — **Open Systems Laboratory**
Fall 2001 – Spring 2003, — *Indiana University, Bloomington*
Fall 2004 – present
Research work in high performance computing, particularly implementations of the Message Passing Interface (both LAM/MPI and Open MPI). Worked with the Parallel Boost Graph Library development team on extensions to the MPI one-sided interface to improve performance and scalability of the library's graph algorithms. Designed

and implemented support for the Red Storm / Cray XT platform, including support for the Catamount light-weight operating system and Portals communication library.

Programmer Analyst **Information Sciences Institute**
2003-2004 *University of Southern California*
Member of the Joint Experimentation on Scalable Parallel Processors team, extending large scale military simulation software to more efficiently utilize modern HPC clusters. Co-developed a new software routing infrastructure for the project, increasing scalability and failure resistance.

Student Intern **Sandia National Laboratories**
Summer 2002 *Albuquerque, New Mexico*
Worked with the Scalable Computing Systems organization on the parallel run-time environment for the Cplant clustering system. Developed a run-time performance metrics system for the Cplant MPI implementation.

Student Intern **Sandia National Laboratories**
Summer 2001 *Albuquerque, New Mexico*
Provided MPI support for the Alegra code development team. Investigated fault tolerance options for large scale MPI applications within the context of LAM/MPI.

Publications

Brian W. Barrett, Jonathan W. Berry, Richard C. Murphy, and Kyle B. Wheeler. Implementing a Portable Multi-threaded Graph Library: the MTGL on Qthreads. In *Proceedings of the 23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2009), Workshop on Multithreaded Architectures and Applications*, 2009.

Brian W. Barrett, Galen M. Shipman, and Andrew Lumsdaine. Analysis of Implementation Options for MPI-2 One-sided. In *Proceedings, 14th European PVM/MPI Users' Group Meeting*, Paris, France, September 2007.

Richard L. Graham, Ron Brightwell, Brian W. Barrett, George Bosilca, and Jelena Pješivac-Grbović. An Evaluation of Open MPI's Matching Transport Layer on the Cray XT. In *Proceedings, 14th European PVM/MPI Users' Group Meeting*, Paris, France, September 2007.

Galen M. Shipman, Ron Brightwell, Brian W. Barrett, Jeffrey M. Squyres, and Gil Bloch. Investigations on InfiniBand: Efficient Network Buffer Utilization at Scale. In *Proceedings, 14th European PVM/MPI Users' Group Meeting*, Paris, France, September 2007.

Richard L. Graham, Brian W. Barrett, Galen M. Shipman, Timothy S. Woodall and George Bosilca. Open MPI: A High Performance, Flexible Implementation of MPI Point-to-Point Communications. In *Parallel Processing Letters*, Vol. 17, No. 1, March 2007.

Ralph Castain, Tim Woodall, David Daniel, Jeff Squyres, and Brian W. Barrett. The Open Run-Time Environment (OpenRTE): A Transparent Multi-Cluster Environment for High-Performance Computing. In *Future Generation Computer Systems*. Accepted for publication.

Christopher Gottbrath, Brian Barrett, Bill Gropp, Ewing Rusty Lusk, and Jeff Squyres. An Interface to Support the Identification of Dynamic MPI 2 Processes for Scalable Parallel Debugging. In *Proceedings, 13th European PVM/MPI Users' Group Meeting*, Bonn, Germany, September 2006.

Richard L. Graham, Brian W. Barrett, Galen M. Shipman, and Timothy S. Woodall. Open MPI: A High Performance, Flexible Implementation of MPI Point-To-Point Communications. In *Proceedings, Clusters and Computational Grids for cientific Computing*, Flat Rock, North Carolina, September 2006.

Richard L. Graham, Galen M. Shipman, Brian W. Barrett, Ralph H. Castain, and George Bosilca. Open MPI: A High Performance, Heterogeneous MPI. In *Proceedings, Fifth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks*, Barcelona, Spain, September 2006.

Brian W. Barrett, Ron Brightwell, Jeffrey M. Squyres, and Andrew Lumsdaine. Implementation of Open MPI on the XT3. *Cray Users Group* 2006, Lagano, Switzerland, May 2006.

Brian W. Barrett, Jeffrey M. Squyres, and Andrew Lumsdaine. Implementation of Open MPI on Red Storm. Technical report LA-UR-05-8307, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, October 2005.

B. Barrett, J. M. Squyres, A. Lumsdaine, R. L. Graham, and G. Bosilca. Analysis of the Component Architecture Overhead in Open MPI. In *Proceedings, 12th European PVM/MPI Users' Group Meeting*, Sorrento, Italy, September 2005.

R. H. Castain, T. S. Woodall, D. J. Daniel, J. M. Squyres, B. Barrett, and G. E. Fagg. The Open Run-Time Environment (OpenRTE): A Transparent Multi-Cluster Environment for High-Performance Computing. In *Proceedings, 12th European PVM/MPI Users' Group Meeting*, Sorrento, Italy, September 2005.

Brian Barrett and Thomas Gottschalk. Advanced Message Routing for Scalable Distributed Simulations. In *Proceedings, Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Orlando, FL 2004.

Edgar Gabriel, Graham E. Fagg, George Bosilca, Thara Angskun, Jack J. Dongarra, Jeffrey M. Squyres, Vishal Sahay, Prabhanjan Kambadur, Brian Barrett, Andrew Lumsdaine, Ralph H. Castain, David J. Daniel, Richard L. Graham, and Timothy S. Woodall. Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, Budapest, Hungary, September 2004.

T.S. Woodall, R.L. Graham, R.H. Castain, D.J. Daniel, M.W. Sukalski, G.E. Fagg, E. Gabriel, G. Bosilca, T. Angskun, J.J. Dongarra, J.M. Squyres, V. Sahay, P. Kambadur, B. Barrett, and A. Lumsdaine. Open MPI's TEG Point-to-Point Communications Methodology: Comparison to Existing Implementations. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, Budapest, Hungary, September 2004.

Brian W. Barrett. Return of the MPI Datatypes. ClusterWorld Magazine, MPI Mechanic Column, 2(6):34–36, June 2004.

Brian Barrett, Jeff Squyres, and Andrew Lumsdaine. Integration of the LAM/MPI environment and the PBS scheduling system. In *Proceedings, 17th Annual International Symposium on High Performance Computing Systems and Applications*, Quebec, Canada, May 2003.

John Mugler, Thomas Naughton, Stephen L. Scott, Brian Barrett, Andrew Lumsdaine, Jeffrey M. Squyres, Benoit des Ligneris, Francis Giraldeau, and Chokchai Leangsuksun.

OSCAR Clusters. In *Proceedings of the Ottawa Linux Symposium (OLS'03)*, Ottawa, Canada, July 23-26, 2003.

Sriram Sankaran, Jeffrey M. Squyres, Brian Barrett, Andrew Lumsdaine, Jason Duell, Paul Hargrove, and Eric Roman. The LAM/MPI Checkpoint/Restart Framework: System-Initiated Checkpointing. In *LACSI Symposium*, October 2003.

Thomas Naughton, Stephen L. Scott, Brian Barrett, Jeffrey M. Squyres, Andrew Lumsdaine, Yung-Chin Gang, and Victor Mashayekhi. Looking inside the OSCAR cluster toolkit. Technical report in PowerSolutions Magazine, chapter HPC Cluster Environment, Dell Computer Corporation, November 2002.

Software

LAM/MPI (`http://www.lam-mpi.org/`) Open source implementation of the MPI standard.

Open MPI (`http://www.open-mpi.org/`) High performance open source implementation of the MPI standard, developed in collaboration by the developers of LAM/MPI, LA-MPI, and FT-MPI.

Mesh-based routing infrastructure for the RTI-s implementation of the HLA discrete event simulation communication infrastructure, providing plug-in replacement to the existing tree-based routing infrastructure.

Honors and Awards

Department of Energy High Performance Computer Science fellowship, 2001–2003.

Service

Secretary, Computer Science Graduate Student Association, Indiana University, 2002–2003

President, Notre Dame Linux Users Group, University of Notre Dame, 2000–2001