

## **Rajeev Thakur**

Argonne National Laboratory  
9700 S. Cass Avenue, Argonne, IL 60439  
Tel: (630) 252-7847, E-mail: thakur@anl.gov  
Web: www.mcs.anl.gov/~thakur

### **Education**

- Ph.D., Computer Engineering, Syracuse University, Syracuse, NY, 1995
- M.S., Computer Engineering, Syracuse University, Syracuse, NY, 1992
- B.E., Computer Engineering, University of Mumbai, India, 1990

### **Work Experience and Positions Held**

- Argonne National Laboratory, Argonne, IL
  - Deputy Director, Data Science and Learning Division, February 2018–present
  - Senior Computer Scientist, November 2011–present
  - Exascale Computing Project, Software Technology Lead for Programming Models and Runtimes, March 2016–present
  - Director of Software Technology, Exascale Computing Project, March 2016–Oct. 2017
  - Deputy Director, Mathematics and Computer Science Division, June 2012–Nov. 2016
  - Computer Scientist, April 2002–October 2011
  - Assistant Computer Scientist, May 1997–April 2002  
(on leave December 2000–December 2001)
  - Postdoctoral Researcher, May 1995–April 1997
- Computation Institute, University of Chicago, Chicago, IL
  - Senior Fellow, April 2010–present
  - Fellow, August 2006–March 2010
- Dept. of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL
  - Adjunct Professor, September 2009–August 2015
  - Adjunct Associate Professor, September 2006–August 2009
- Prismedia Networks (Acirro), Inc., San Jose, CA
  - Senior Scientist, December 2000–December 2001

### **Research Interests**

High-performance computing, scalable parallel I/O, parallel programming models, runtime systems, communication libraries

### **Honors and Awards**

- R&D 100 Award for “MPICH2: High Performance, Portable Implementation of the Message Passing Interface Standard for Parallel Computing,” 2005
- Selected to participate in Argonne’s Strategic Laboratory Leadership Program taught by the University of Chicago Booth School of Business, 2011
- Two outstanding papers at Euro PVM/MPI 2009
- Best Paper Award at the International Supercomputing Conference (ISC) 2009
- Finalist for Best Paper Award and Best Student Paper Award at SC08

- One outstanding paper at Euro PVM/MPI 2008
- Best Paper Award at the International Supercomputing Conference (ISC) 2008
- Two outstanding papers at Euro PVM/MPI 2007
- Two outstanding papers at Euro PVM/MPI 2006
- Best Poster Award, SC2003: High Performance Networking and Computing
- Best Paper Award, 4th Annual Linux Showcase and Conference, October 2000
- Best Student Paper Award at Supercomputing '92 in the category of performance measurement, November 1992

## Major Software Developed

1. **MPICH** I have been involved in the development of the MPICH implementation of MPI for many years. MPICH is used by thousands of users all over the world and forms the basis of vendor MPI implementations from Intel Cray, IBM, Microsoft, and other vendors. The majority of the largest supercomputers in the world run MPICH-based MPI implementations.
2. **ROMIO** I am the original developer of the ROMIO implementation of MPI-IO, which is used as the MPI-IO implementation in MPICH and other MPI implementations.

## International Standards Organizations

- I have been an active member of the MPI Forum since the mid-1990s. I participated in the definition of MPI-2 in 1995-97, particularly the interface for parallel I/O in MPI-2 (MPI-IO).
- More recently, I have participated in the MPI Forum for the definition of MPI 2.1, MPI 2.2, and MPI-3 standards. Together with Bill Gropp and Torsten Hoefler, I am co-editor of the One-Sided Communication chapter in MPI-3.

## Professional Activities

- Steering Committee, SC Conference Series, January 2015–present
  - Treasurer, ACM SIGHPC, July 2016–present
  - Member-at-Large, ACM SIGHPC, July 2013–June 2016
  - Steering Committee, Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS), 2016–present
  - Steering Committee, International Workshop on Data-Intensive Scalable Computing Systems (DISCS), 2012–2015
  - Editorial Advisory Board, Scientific Programming, November 2012–November 2014
  - Associate Editor, IEEE Transactions on Parallel and Distributed Systems, 2004–2007
- 
- Program Co-Chair, 24th European MPI Users Group Conference (EuroMPI/USA), Argonne, Illinois, September 2017
  - Program Committee, IEEE/ACM 17th International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2017), Madrid, Spain, May 2017
  - Awards Co-Chair, SC16: International Conference on High Performance Computing, Networking, Storage, and Analysis, Salt Lake City, Utah, November 2016
  - Program Committee, Second International Workshop on Extreme Scale Programming Models and Middleware (ESPM2 2016) at SC16, Salt Lake City, Utah, November 2016

- Tutorials Committee, International Supercomputing Conference (ISC'16), Frankfurt, Germany, June 2016
- Program Committee, International Conference on Supercomputing (ICS 2016), Istanbul, Turkey, June 2016
- Program Committee, 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2016), Chicago, Illinois, May 2016
- Program Committee, IEEE/ACM 16th International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2016), Cartagena, Columbia, May 2016
- Program Committee, First International Workshop on Extreme Scale Programming Models and Middleware (ESPM2 2015) at SC15, Austin, Texas, November 2015
- Program Committee, 24th International Conference on Parallel Architectures and Compilation Techniques (PACT 2015), San Francisco, California, October 2015
- Program Committee, 22nd European MPI Users Group Conference (EuroMPI), Bordeaux, France, September 2015
- Program Committee, 44th International Conference on Parallel Processing (ICPP 2015), Beijing, China, September 2015
- Program Committee, 21st International European Conference on Parallel and Distributed Computing (Euro-Par 2015), Vienna, Austria, August 2015
- Program Committee, 29th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2015), Hyderabad, India, May 2015
- Tutorials Committee, SC14: International Conference on High Performance Computing, Networking, Storage, and Analysis, New Orleans, Louisiana, November 2014
- Program Committee, 21st European MPI Users Group Conference (EuroMPI/Asia), Kyoto, Japan, September 2014
- Program Committee, 28th International Conference on Supercomputing (ICS 2014), Munich, Germany, June 2014
- Program Committee, 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2014), Phoenix, Arizona, May 2014
- Program Committee, Sixth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2013, Lyon, France, October 2013
- Program Committee, 20th European MPI Users Group Conference (EuroMPI), Madrid, Spain, September 2013
- Program Committee, 27th International Conference on Supercomputing (ICS 2013), Eugene, Oregon, June 2013
- Technical Program Chair, SC12: International Conference on High Performance Computing, Networking, Storage, and Analysis, Salt Lake City, Utah, November 2012
- Program Committee, 19th European MPI Users Group Conference (EuroMPI), Vienna, Austria, September 2012
- Program Committee, 41st International Conference on Parallel Processing (ICPP-11), Pittsburgh, Pennsylvania, September 2012
- Program Committee, Fifth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2012, Pittsburgh, Pennsylvania, September 2012
- Program Committee, 10th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA 2012), Madrid, Spain, July 2012
- Tutorials Committee, International Supercomputing Conference (ISC'12), Hamburg, Germany, June 2012

- Program Committee, IEEE International Symposium on Cluster Computing and the Grid (CC-Grid 2012), Ottawa, Canada, May 2012
- Technical Papers Co-Chair, SC11: International Conference on High Performance Computing, Networking, Storage, and Analysis, Seattle, Washington, November 2011
- Program Committee, 2011 IEEE International Symposium on Workload Characterization, Austin, Texas, November 2011
- Program Committee, IEEE International Conference on Cluster Computing, Austin, Texas, September 2011
- Program Committee, 18th European MPI Users Group Conference (EuroMPI), Santorini, Greece, September 2011
- Program Committee, 40th International Conference on Parallel Processing (ICPP-11), Taipei, Taiwan, September 2011
- Program Committee, Fourth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2011, Taipei, Taiwan, September 2011
- Program Committee, Workshop on Parallel and Distributed Systems: Testing, Analysis, and Debugging (PADTAD 2011), Toronto, Canada, July 2011
- Program Committee, 25th International Conference on Supercomputing (ICS 2011), Tucson, Arizona, June 2011
- Tutorials Committee, International Supercomputing Conference (ISC'11), Hamburg, Germany, June 2011
- Program Committee, IEEE International Symposium on Cluster Computing and the Grid (CC-Grid 2011), Newport Beach, California, May 2011
- Program Committee, 7th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI 2011), Denver, Colorado, May 2011
- Program Vice-Chair for communication networks, 17th IEEE International Conference on High Performance Computing (HiPC '10), Goa, India, December 2010
- System Software Area Co-Chair for Technical Papers, SC10: International Conference on High Performance Computing, Networking, Storage, and Analysis, New Orleans, Louisiana, November 2010
- Program Committee, 39th International Conference on Parallel Processing (ICPP-10), San Diego, California, September 2010
- Program Committee, Third International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2010, San Diego, California, September 2010
- Program Committee, 17th European MPI Users Group Conference (EuroMPI), Stuttgart, Germany, September 2010
- Program Committee, 5th IEEE International Conference on Networking, Architecture, and Storage (NAS 2010), Macau, Macau SAR, China, July 2010
- Program Committee, Workshop on Parallel and Distributed Systems: Testing, Analysis, and Debugging (PADTAD 2010), Trento, Italy, July 2010
- Program Committee, 24th International Conference on Supercomputing (ICS 2010), Tsukuba, Japan, June 2010
- Program Committee, 6th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI 2010), Incline Village, Nevada, May 2010
- Co-Chair, Student Research Symposium, 16th IEEE International Conference on High Performance Computing (HiPC '09), Cochin, India, December 2009

- Tutorials Co-Chair, SC09: International Conference on High Performance Computing, Networking, Storage, and Analysis, Portland, Oregon, November 2009
- Technical Papers Committee, SC09: International Conference on High Performance Computing, Networking, Storage, and Analysis, Portland, Oregon, November 2009
- Program Committee, Third Conference on Partitioned Global Address Space Programming Models (PGAS 2009), Ashburn, Virginia, October 2009
- Program Committee, IEEE International Conference on Cluster Computing, New Orleans, Louisiana, September 2009
- Program Committee, 16th European PVM/MPI Users Group Conference (Euro PVM/MPI), Helsinki, Finland, September 2009
- Program Committee, Second International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2009, Vienna, Austria, September 2009
- Program Committee, 38th International Conference on Parallel Processing (ICPP-09), Vienna, Austria, September 2009
- Program Committee, Workshop on Parallel and Distributed Systems: Testing, Analysis, and Debugging (PADTAD 2009), Chicago, Illinois, July 2009
- Co-Chair, Student Research Symposium, 15th IEEE International Conference on High Performance Computing (HiPC '08), Bangalore, India, December 2008
- Program Committee, 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS'08), Melbourne, Australia, December 2008
- Technical Papers Committee, SC08: International Conference on High Performance Computing, Networking, Storage, and Analysis, Austin, Texas, November 2008
- Program Committee, First International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2) at ICPP 2008, Portland, Oregon, September 2008
- Program Committee, 37th International Conference on Parallel Processing (ICPP-08), Portland, Oregon, September 2008
- Program Committee, Thirteenth Asia-Pacific Computer Systems Architecture Conference, Taiwan, August 2008
- Poster/Presentation Chair, 14th IEEE International Conference on High Performance Computing (HiPC '07), Goa, India, December 2007
- Program Committee, IEEE International Conference on Cluster Computing, Austin, Texas, September 2007
- Program Committee, 2007 IEEE International Symposium on Workload Characterization, Boston, Massachusetts, September 2007
- Program Committee, Twelfth Asia-Pacific Computer Systems Architecture Conference, Seoul, Korea, August 2007
- Program Committee, Third Workshop on System Management Techniques, Services, and Processes at IPDPS 2007, Long Beach, California, March 2007
- Poster/Presentation Chair, 13th IEEE International Conference on High Performance Computing (HiPC '06), Bangalore, India, December 2006
- Program Committee, Twelfth International Conference on Parallel and Distributed Systems, Minneapolis, Minnesota, July 2006
- Program Committee, Eleventh Asia-Pacific Computer Systems Architecture Conference, Shanghai, China, June 2006

- Program Committee, Second Workshop on System Management Tools for Large-Scale Parallel Systems at IPDPS 2006, Rhodes Island, Greece, April 2006
- Poster/Presentation Chair, 12th IEEE International Conference on High Performance Computing (HiPC '05), Goa, India, December 2005
- Program Committee, Tenth Asia-Pacific Computer Systems Architecture Conference, Singapore, October 2005
- Program Committee, First Workshop on System Management Tools for Large-Scale Parallel Systems at IPDPS 2005, Denver, Colorado, April 2005
- Program Committee, IEEE International Conference on Cluster Computing, San Diego, California, September 2004
- Program Committee, Ninth Asia-Pacific Computer Systems Architecture Conference, Beijing, China, September 2004
- Program Committee, International Conference on Parallel Processing, Montreal, Canada, August 2004
- Program Committee, IEEE International Conference on Cluster Computing, Hong Kong, December 2003
- Program Committee, Workshop on Parallel I/O in Cluster Computing and Computational Grids at CCGrid 2003, Tokyo, Japan, May 2003
- Vice Program Chair, Cluster 2002: IEEE International Conference on Cluster Computing, Chicago, Illinois, September 2002
- Editorial Board of Special Issue of *Calculateurs Parallèles* on Parallel I/O for Cluster Computing, 2001
- Program Committee, IEEE International Symposium on Cluster Computing and the Grid (CC-Grid'2001), Brisbane, Australia, May 2001
- Chair for Birds-of-a-Feather sessions, SC2000: High Performance Networking and Computing, Dallas, Texas, November 2000
- Global Chair, Parallel I/O and Storage Technology track at the European Conference on Parallel Computing (Euro-Par 2000), Munich, Germany, August 2000
- Program Committee, 6th International Conference on High Performance Computing (HiPC '99), Calcutta, India, December 1999
- General Co-Chair, Sixth Workshop on I/O in Parallel and Distributed Systems (IOPADS '99), Atlanta, Georgia, May 1999
- Program Committee, Heterogeneous Computing Workshop at IPPS '99, San Juan, Puerto Rico, April 1999
- Poster Exhibits Chair, SC98: High Performance Networking and Computing, Orlando, Florida, November 1998
- Vice-Chair, Parallel I/O Workshop at Euro-Par '98, Southampton, UK, September 1998
- Co-Guest Editor, Special Issue of the *International Journal of High Performance Computing Applications* on "I/O in Parallel Applications," 1998
- Tutorials Vice-Chair, SC97: High Performance Networking and Computing, San Jose, California, November 1997
- Program Committee, 4th International Conference on High Performance Computing (HiPC '97), Bangalore, India, December 1997
- Publicity Chair, Fifth Workshop on I/O in Parallel and Distributed Systems (IOPADS '97), San Jose, California, November 1997
- Tutorials Vice-Chair, Supercomputing '96, Pittsburgh, Pennsylvania, November 1996

- Reviewer for *IEEE Transactions on Parallel and Distributed Systems*, *Journal of Parallel and Distributed Computing*, *Parallel Computing*, *Concurrency and Computation: Practice and Experience*, *Parallel Processing Letters*, *The Journal of Supercomputing*, *The Computer Journal*, *International Journal of High Performance Computing Applications*, and several conferences
- Served on Ph.D. thesis committees of Rakesh Krishnaiyer (1998), Chutimet Srinilta (1998), and Jaechun No (1999) at Syracuse University; Murali Vilayannur at Penn State University (2005); Sachin More (2000), Avery Ching (2007), Kenin Coloma (2007), and Arifa Nisar (2010) at Northwestern University; Yawei Li (2008) at Illinois Institute of Technology; Saba Sehrish (2010) at University of Central Florida

## Patent

1. “Method and Apparatus for Real-Time Parallel Delivery of Segments of a Large Payload File.” Inventors: Desmond Cho-Hung Chan, Siew Yong Sim, Rajeev Thakur, Trygve Isaacson, and William D. Gropp. US Patent 7076553, awarded July 11, 2006.

## Publications and Presentations

### Books

1. William Gropp, Torsten Hoefler, Rajeev Thakur, and Ewing Lusk, *Using Advanced MPI: Modern Features of the Message-Passing Interface*, MIT Press, November 2014.
2. William Gropp, Ewing Lusk, and Rajeev Thakur, *Using MPI-2: Advanced Features of the Message-Passing Interface*, MIT Press, November 1999. (Also translated into Japanese.)

## Refereed Journal Articles and Book Chapters

1. Antonios Kougkas, Hassan Eslami, Xian-He Sun, Rajeev Thakur, and William Gropp, “Re-thinking Key-Value Store for Parallel I/O Optimization,” *International Journal of High Performance Computing Applications*, 31(4):335–356, July 2017.
2. James Dinan, Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, Rajeev Thakur, “An Implementation and Evaluation of the MPI 3.0 One-Sided Communication Interface,” *Concurrency and Computation: Practice and Experience*, 28(17):4385–4404, December 2016.
3. Ashwin M. Aji, Lokendra S. Panwar, Feng Ji, Karthik Murthy, Milind Chabbi, Pavan Balaji, Keith R. Bisset, James Dinan, Wu-chun Feng, John Mellor-Crummey, Xiaosong Ma, and Rajeev Thakur, “MPI-ACC: Accelerator-Aware MPI for Scientific Applications,” *IEEE Transactions on Parallel and Distributed Systems*, 27(5):1401–1414, May 2016.
4. William Gropp and Rajeev Thakur, “Message Passing Interface,” *Programming Models for Parallel Computing*, MIT Press, Ch. 1, pp. 1–21, 2015.
5. Torsten Hoefler, James Dinan, Rajeev Thakur, Brian Barrett, Pavan Balaji, and William Gropp, “Remote Memory Access Programming in MPI-3,” *ACM Transactions on Parallel Computing*, Vol. 7, No. 2, pp. 9:1–9:26, July 2015.
6. Yong Chen, Yin Lu, Prathamesh Amritkar, Rajeev Thakur, and Yu Zhuang, “Performance Model Directed Data Sieving for High Performance I/O,” *Journal of Supercomputing*, 71(6):2066–2090, June 2015.

7. Seong Jo Kim, Yuanrui Zhang, Seung Woo Son, Mahmut Kandemir, Wei-keng Liao, Rajeev Thakur, Alok Choudhary, "IOPro: A Parallel I/O Profiling and Visualization Framework for High-Performance Storage Systems," *Journal of Supercomputing*, 71(3):840–870, March 2015.
8. Yin Lu, Yong Chen, Yu Zhuang, Jialin Liu, Rajeev Thakur, "Collective Input/Output under Memory Constraints," *International Journal of High Performance Computing Applications*, 29(1):21–36, February 2015.
9. James Dinan, Ryan E. Grant, Pavan Balaji, David Goodell, Douglas Miller, Marc Snir, and Rajeev Thakur, "Enabling Communication Concurrency Through Flexible MPI Endpoints," *International Journal of High Performance Computing Applications*, 28(4):390–405, November 2014.
10. Wei-keng Liao and Rajeev Thakur, "MPI-IO," *High Performance Parallel I/O*, CRC Press, Ch. 13, pp. 157–169, October 2014.
11. John Jenkins, James Dinan, Pavan Balaji, Tom Peterka, Nagiza F. Samatova, and Rajeev Thakur, "Processing MPI Derived Datatypes on Noncontiguous GPU-resident Data," *IEEE Transactions on Parallel and Distributed Systems*, 25(10):2627–2637, October 2014.
12. Anshu Dubey, Steve R. Brandt, Richard Brower, Merle Giles, Paul Hovland, Donald Q. Lamb, Frank Loffler, Boyana Norris, Brian W. O'Shea, Claudio Rebbi, Marc Snir, Rajeev Thakur, and Petros Tzeferacos, "Software Abstractions and Methodologies for HPC Simulation Codes on Future Architectures," *Journal of Open Research Software*, 2(1):e14, pp. 1–5, 2014.
13. Torsten Hoefler, James Dinan, Darius Buntinas, Pavan Balaji, Brian Barrett, Ron Brightwell, William Gropp, Vivek Kale, and Rajeev Thakur, "MPI+MPI: A New, Hybrid Approach to Parallel Programming with MPI Plus Shared Memory Computing," *Computing*, 95(12):1121–1136, December 2013.
14. Ganesh Gopalakrishnan, Robert M. Kirby, Stephen Siegel, Rajeev Thakur, William Gropp, Ewing Lusk, Bronis R. de Supinski, Martin Schulz, and Greg Bronevetsky, "Formal Analysis of MPI-Based Parallel Programs," *Communications of the ACM*, 54(12):82–91, December 2011.
15. Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, Torsten Hoefler, Sameer Kumar, Ewing Lusk, Rajeev Thakur, and Jesper Larsson Träff, "MPI on Millions of Cores," *Parallel Processing Letters*, 21(1):45–60, March 2011.
16. Torsten Hoefler, Rolf Rabenseifner, Hubert Ritzdorf, Bronis R. de Supinski, Rajeev Thakur, and Jesper Larsson Träff, "The Scalable Process Topology Interface of MPI 2.2," *Concurrency and Computation: Practice and Experience*, 23(4):293–310, March 2011.
17. P. Balaji, W. Feng, H. Lin, J. Archuleta, S. Matsuoka, A. Warren, J. Setubal, E. Lusk, R. Thakur, I. Foster, K. Shinpaugh, S. Coghlan, and D. Reed, "Global-scale Distributed I/O with ParaMEDIC," *Concurrency and Computation: Practice and Experience*, 22(16):2266–2281, November 2010.
18. Zhiling Lan, Jiexing Gu, Ziming Zheng, Rajeev Thakur, and Susan Coghlan, "A Study of Dynamic Meta-Learning for Failure Prediction in Large-Scale Systems," *Journal of Parallel and Distributed Computing*, 70(6):630–643, June 2010.
19. Jesper Larsson Träff, William D. Gropp, and Rajeev Thakur, "Self-Consistent MPI Performance Guidelines," *IEEE Transactions on Parallel and Distributed Systems*, 21(5):698–709, May 2010.
20. Jesper Larsson Träff, Andreas Ripke, Christian Siebert, Pavan Balaji, and Rajeev Thakur, and William Gropp, "A Pipelined Algorithm for Large, Irregular Allgather Problems," *International Journal of High Performance Computing Applications*, 24(1):58–68, Spring 2010.



21. Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, and Rajeev Thakur, "Fine-Grained Multithreading Support for Hybrid Threaded MPI Programming," *International Journal of High Performance Computing Applications*, 24(1):49–57, Spring 2010.
22. Pavan Balaji, Anthony Chan, William Gropp, Rajeev Thakur, and Ewing Lusk, "The Importance of Non-Data-Communication Overheads in MPI," *International Journal of High Performance Computing Applications*, 24(1):5–15, Spring 2010.
23. Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby, Rajeev Thakur, and William Gropp, "Formal Methods Applied to High Performance Computing Software Design: A Case Study of MPI One-Sided Communication Based Locking," *Software: Practice and Experience*, 40(1):23–43, January 2010.
24. Rajeev Thakur and William Gropp, "Test Suite for Evaluating Performance of Multithreaded MPI Communication," *Parallel Computing*, 35(12):608–617, December 2009.
25. Alok Choudhary, Wei-keng Liao, Kui Gao, Arifa Nisar, Robert Ross, Rajeev Thakur, and Robert Latham, "Scalable I/O and Analytics," *Journal of Physics: Conference Series (SciDAC 2009)*, Vol 180, 2009.
26. Pavan Balaji, Anthony Chan, Rajeev Thakur, William Gropp, and Ewing Lusk, "Toward Message Passing for a Million Processes: Characterizing MPI on a Massive Scale Blue Gene/P," *Computer Science – Research and Development*, 24(1-2):11–19, September 2009. (**Best Paper Award** at the International Supercomputing Conference (ISC) 2009)
27. P. Lai, P. Balaji, R. Thakur, and D. K. Panda, "ProOnE: A General Purpose Protocol Onload Engine for Multi- and Many-Core Architectures," *Computer Science – Research and Development*, 23(3-4):133–142, June 2009.
28. William Gropp and Rajeev Thakur, "Thread Safety in an MPI Implementation: Requirements and Analysis," *Parallel Computing*, 33(9):595–604, September 2007.
29. Robert Latham, Robert Ross, and Rajeev Thakur, "Implementing MPI-IO Atomic Mode and Shared File Pointers Using MPI One-Sided Communication," *International Journal of High Performance Computing Applications*, 21(2):132–143, Summer 2007.
30. Murali Vilayannur, Anand Sivasubramaniam, Mahmut Kandemir, Rajeev Thakur, and Robert Ross, "Discretionary Caching for I/O on Clusters," *Cluster Computing*, 9(1):29–44, January 2006.
31. Robert Ross, Rajeev Thakur, and Alok Choudhary, "Achievements and Challenges for I/O in Computational Science," *Journal of Physics: Conference Series (SciDAC 2005)*, (16):501–509, 2005.
32. Rajeev Thakur, William Gropp, and Brian Toonen, "Optimizing the Synchronization Operations in MPI One-Sided Communication," *International Journal of High Performance Computing Applications*, 19(2):119–128, Summer 2005.
33. Rajeev Thakur, Rolf Rabenseifner, and William Gropp, "Optimization of Collective Communication Operations in MPICH," *International Journal of High Performance Computing Applications*, 19(1):49–66, Spring 2005.
34. Rajeev Thakur, William Gropp, and Ewing Lusk, "ADIO: A Framework for High-Performance, Portable Parallel I/O," *Scalable Input/Output: Achieving System Balance*, MIT Press, Ch. 4, pp. 111–134, 2004.
35. Alok Choudhary, Mahmut Kandemir, Sachin More, Jaechun No, and Rajeev Thakur, "Collective I/O and Large-Scale Data Management," *Scalable Input/Output: Achieving System Balance*, MIT Press, Ch. 2, pp. 35–75, 2004.
36. Jaechun No, Rajeev Thakur, and Alok Choudhary, "High-Performance Scientific Data Management System," *Journal of Parallel and Distributed Computing*, 64(4):434–447, April 2003.

37. Rajeev Thakur and William Gropp, "Parallel I/O," *Sourcebook of Parallel Computing*, Morgan Kaufmann Publishers, Ch. 11, pp. 331–355, 2002.
38. Rajeev Thakur, William Gropp, and Ewing Lusk, "Optimizing Noncontiguous Accesses in MPI-IO," *Parallel Computing*, 28(1):83–105, January 2002.
39. Phillip Dickens and Rajeev Thakur, "Evaluation of Collective I/O Implementations on Parallel Architectures," *Journal of Parallel and Distributed Computing*, 61(8):1052–1076, August 1, 2001.
40. Dan Bonachea, Phillip Dickens, and Rajeev Thakur, "High-Performance File I/O in Java: Existing Approaches and Bulk I/O Extensions," *Concurrency and Computation: Practice and Experience*, 13(8-9):713–736, 2001.
41. A. Choudhary, M. Kandemir, J. No, G. Memik, X. Shen, W. Liao, H. Nagesh, S. More, V. Taylor, R. Thakur, and R. Stevens, "Data Management for Large-Scale Scientific Computations in High Performance Distributed Systems," *Cluster Computing*, 3(1):45–60, 2000.
42. Rajeev Thakur, Ewing Lusk, and William Gropp, "I/O in Parallel Applications: The Weakest Link," *International Journal of Supercomputer Applications and High Performance Computing*, 12(4):389–395, Winter 1998. (Guest Editors' Introduction)
43. Rajeev Thakur and Alok Choudhary, "An Extended Two-Phase Method for Accessing Sections of Out-of-Core Arrays," *Scientific Programming*, 5(4):301–317, Winter 1996.
44. Rajeev Thakur, Alok Choudhary, and J. Ramanujam, "Efficient Algorithms for Array Redistribution," *IEEE Transactions on Parallel and Distributed Systems*, 7(6):587–594, June 1996.
45. Rajeev Thakur, Alok Choudhary, Rajesh Bordawekar, Sachin More, and Sivaramakrishna Kuditipudi, "Passion: Optimized I/O for Parallel Applications," *IEEE Computer*, 29(6):70–78, June 1996.
46. Rajeev Thakur and Alok Choudhary, "Runtime Support for Out-of-Core Parallel Programs," *Input/Output in Parallel and Distributed Computer Systems*, Kluwer Academic Publishers, Ch. 6, pp. 147–165, 1996.
47. Rajeev Thakur, Ravi Ponnusamy, Alok Choudhary, and Geoffrey Fox, "Complete Exchange on the CM-5 and Touchstone Delta," *The Journal of Supercomputing*, 8(4):305–328, 1995.
48. Alok Choudhary and Rajeev Thakur, "Connected Component Labeling on Coarse Grain Parallel Computers: An Experimental Study," *Journal of Parallel and Distributed Computing*, 20(1):78–83, January 1994.
49. Ravi Ponnusamy, Rajeev Thakur, Alok Choudhary, Kishore Velamakanni, and Geoffrey Fox, "Experimental Performance Evaluation of the CM-5," *Journal of Parallel and Distributed Computing*, 19(3):192–202, November 1993.

## Refereed Proceedings

1. Yong Chen, Chao Chen, Yanlong Yin, Xian-He Sun, Rajeev Thakur, William Gropp, "Rethinking High Performance Computing System Architecture for Scientific Big Data Applications," in *Proceedings of the 14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA-16)*, August 2016. (**Best Paper Award**)
2. Hassan Eslami, Anthony Kougkas, Maria Kotsifakou, Theodoros Kasampalis, Kun Feng, Yin Lu, William Gropp, Xian-He Sun, Yong Chen, Rajeev Thakur, "Efficient Disk-to-Disk Sorting: A Case Study in the Decoupled Execution Paradigm," in *Proceedings of the 2015 International Workshop on Data-Intensive Scalable Computing Systems (DISCS)* (held in conjunction with SC15), November 2015.

3. Swann Perarnau, Rajeev Thakur, Kamil Iskra, Ken Raffanetti, Franck Cappello, Rinku Gupta, Pete Beckman, Marc Snir, Henry Hoffmann, Martin Schulz, and Barry Rountree, "Distributed Monitoring and Management of Exascale Systems in the Argo Project," in *Proceedings of the 15th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*, June 2015.
4. Yanlong Yin, Antonios Kougkas, Kun Feng, Hassan Eslami, Yin Lu, Xian-He Sun, Rajeev Thakur, and William Gropp, "Rethinking Key-Value Store for Parallel I/O Optimization," in *Proceedings of the 3rd International Workshop on Data-Intensive Scalable Computing Systems (DISCS)* (held in conjunction with SC14), November 2014.
5. Chao Chen, Yong Chen, Kun Feng, Yanlong Yin, Hassan Eslami, Rajeev Thakur, Xian-He Sun, and William D. Gropp, "Decoupled I/O for Data-Intensive High Performance Computing," in *Proceedings of the 7th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, September 2014.
6. Xin Zhao, Pavan Balaji, William Gropp, and Rajeev Thakur, "Optimization Strategies for MPI-Interoperable Active Messages," in *Proceedings of the 13th IEEE International Conference on Scalable Computing and Communication (ScalCom 2013)*, December 2013. (**Best Paper Award**)
7. Xin Zhao, Pavan Balaji, William Gropp, and Rajeev Thakur, "MPI-Interoperable Generalized Active Messages," in *Proceedings of the 19th IEEE International Conference on Parallel and Distributed Systems (ICPADS'13)*, December 2013.
8. James Dinan, Pavan Balaji, David Goodell, Douglas Miller, Marc Snir, and Rajeev Thakur, "Enabling MPI Interoperability Through Flexible Communication Endpoints," in *Proceedings of the 20th European MPI Users' Group Meeting (EuroMPI 2013)*, September 2013.
9. Antonio J. Peña, Ralf G. Correa Carvalho, James Dinan, Pavan Balaji, Rajeev Thakur, and William Gropp, "Analysis of Topology-Dependent MPI Performance on Gemini Networks," in *Proceedings of the 19th European MPI Users' Group Meeting (EuroMPI 2013)*, September 2013.
10. Yin Lu, Yong Chen, Rajeev Thakur, and Yu Zhuang, "Memory-Conscious Collective I/O for Extreme Scale HPC Systems," in *Proceedings of the International Workshop on Runtime and Operating Systems for Supercomputers (ROSS 2013) at ICS 2013*, June 2013.
11. Palden Lama, Yan Li, Ashwin M. Aji, Pavan Balaji, James Dinan, Shucui Xiao, Yunquan Zhang, Wu-chun Feng, Rajeev Thakur, and Xiaobo Zhou, "pVOCL: Power-Aware Dynamic Placement and Migration in Virtualized GPU Environments," in *Proceedings of the 33rd International Conference on Distributed Computing Systems (ICDCS)*, July 2013.
12. Ashwin M. Aji, Lokendra S. Panwar, Feng Ji, Milind Chabbi, Karthik Murthy, Pavan Balaji, Keith R. Bisset, James Dinan, Wu-chun Feng, John Mellor-Crummey, Xiaosong Ma, and Rajeev Thakur, "On the Efficacy of GPU-Integrated MPI for Scientific Applications," in *Proceedings of the 22nd ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC'13)*, June 2013.
13. Ashwin Aji, Pavan Balaji, James Dinan, Wuchun Feng, and Rajeev Thakur, "Synchronization and Ordering Semantics in Hybrid MPI+GPU Programming," in *Proceedings of the 3rd International Workshop on Accelerators and Hybrid Exascale Systems (AsHES) at IPDPS 2013*, May 2013.
14. Yanlong Yin, Jibing Li, Jun He, Xian-He Sun, and Rajeev Thakur, "Pattern-Direct and Layout-Aware Replication Scheme for Parallel I/O Systems," in *Proceedings of the 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2013)*, May 2013.

15. Xin Zhao, Darius Buntinas, Judicael Zounmevo, James Dinan, David Goodell, Pavan Balaji, Rajeev Thakur, Ahmad Afsahi, and William Gropp, "Towards Asynchronous, MPI-Interoperable Active Messages," in *Proceedings of the 13th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2013)*, May 2013.
16. Seong Jo Kim, Seung Woo Son, Wei-keng Liao, Mahmut Kandemir, Rajeev Thakur, and Alok Choudhary, "IOPin: Runtime Profiling of Parallel I/O in HPC Systems," in *Proceedings of the 7th Parallel Data Storage Workshop (PDSW)* (held in conjunction with SC12), November 2012.
17. Yong Chen, Chao Chen, Xian-He Sun, William D. Gropp, and Rajeev Thakur, "A Decoupled Execution Paradigm for Data-Intensive High-End Computing," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2012)*, September 2012.
18. Jun He, Xian-He Sun, and Rajeev Thakur, "KNOWAC: I/O Prefetch via Accumulated Knowledge," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2012)*, September 2012.
19. John Jenkins, James Dinan, Pavan Balaji, Nagiza F. Samatova, and Rajeev Thakur, "Enabling Fast, Noncontiguous GPU Data Movement in Hybrid MPI+GPU Environments," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2012)*, September 2012.
20. James Dinan, David Goodell, William Gropp, Rajeev Thakur, and Pavan Balaji, "Efficient Multithreaded Context ID Allocation in MPI," in *Proceedings of the 19th European MPI Users' Group Meeting (EuroMPI 2012)*, September 2012.
21. Torsten Hoefler, James Dinan, Darius Buntinas, Pavan Balaji, Brian Barrett, Ron Brightwell, William Gropp, Vivek Kale, and Rajeev Thakur, "Leveraging MPI's One-Sided Communication Interface for Shared-Memory Programming," in *Proceedings of the 19th European MPI Users' Group Meeting (EuroMPI 2012)*, September 2012.
22. Hui Jin, Jiayu Ji, Xian-He Sun, Yong Chen, and Rajeev Thakur, "CHAIIO: Enabling HPC Applications on Data-Intensive File Systems," in *Proceedings of the 2012 International Conference on Parallel Processing*, September 2012.
23. Ashwin Aji, James Dinan, Darius Buntinas, Pavan Balaji, Wu-chun Feng, Keith Bisset, and Rajeev Thakur, "MPI-ACC: An Integrated and Extensible Approach to Data Movement in Accelerator-Based Systems," in *Proceedings of the 14th IEEE International Conference on High Performance Computing and Communications (HPCC-2012)*, June 2012.
24. Feng Ji, Ashwin M. Aji, James Dinan, Darius Buntinas, Pavan Balaji, Rajeev Thakur, Wu-Chun Feng, and Xiaosong Ma, "DMA-Assisted, Intranode Communication in GPU Accelerated Systems," in *Proceedings of the 14th IEEE International Conference on High Performance Computing and Communications (HPCC-2012)*, June 2012.
25. Yin Lu, Yong Chen, Prathamesh Amritkar, Rajeev Thakur, and Yu Zhuang, "A New Data Sieving Approach for High Performance I/O," in *Proceedings of the 7th International Conference on Future Information Technology (FutureTech'12)*, June 2012. (**Best Paper Award**)
26. Huaiming Song, Hui Jin, Jun He, Xian-He Sun, and Rajeev Thakur, "A Server-Level Adaptive Data Layout Strategy for Parallel File Systems," in *Proceedings of the 2012 International Workshop on High Performance Data Intensive Computing (HPDIC 2012)* (held in conjunction with IPDPS 2012), May 2012.
27. Shucai Xiao, Pavan Balaji, Qian Zhu, Rajeev Thakur, Susan Coghlan, Heshan Lin, Gaojin Wen, Jue Hong and Wu-chun Feng, "VOCL: An Optimized Environment for Transparent Virtualization of Graphics Processing Units," in *Proceedings of 2012 Innovative Parallel Computing:*

*Foundations & Applications of GPU, Manycore, and Heterogeneous Systems (InPar 2012)*, May 2012.

28. Shucaï Xiao, Pavan Balaji, James Dinan, Qian Zhu, Rajeev Thakur, Susan Coghlan, Heshan Lin, Gaojin Wen, Jue Hong, and Wu-Chun Feng, "Transparent Accelerator Migration in a Virtualized GPU Environment," in *Proceedings of the 12th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2012)*, May 2012.
29. Yanlong Yin, Surendra Byna, Huaiming Song, Xian-He Sun, and Rajeev Thakur, "Boosting Application-Specific Parallel I/O Optimization Using IOSIG," in *Proceedings of the 12th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2012)*, May 2012.
30. Huaiming Song, Yanlong Yin, Xian-He Sun, Rajeev Thakur, and Samuel Lang, "Server-Side I/O Coordination for Parallel File Systems," in *Proceedings of SC11: International Conference on High Performance Computing, Networking, Storage, and Analysis*, November 2011.
31. Jun He, Huaiming Song, Xian-He Sun, Yanlong Yin, and Rajeev Thakur, "Pattern-Aware File Reorganization in MPI-IO," in *Proceedings of the 6th Parallel Data Storage Workshop (PDSW)* (held in conjunction with SC11), November 2011.
32. David Goodell, William Gropp, Xin Zhao, and Rajeev Thakur, "Scalable Memory Use in MPI: A Case Study with MPICH2," in *Proceedings of the 18th European MPI Users' Group Meeting (EuroMPI 2011)*, September 2011.
33. William Gropp, Torsten Hoefler, Rajeev Thakur, and Jesper Larsson Träff, "Performance Expectations and Guidelines for MPI Derived Datatypes," in *Proceedings of the 18th European MPI Users' Group Meeting (EuroMPI 2011)*, September 2011.
34. Seung Woo Son, Samuel Lang, Robert Latham, Robert Ross, and Rajeev Thakur, "Reliable MPI-IO through Layout-Aware Replication," in *Proceedings of the 7th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI 2011)*, May 2011.
35. Yong Chen, Xian-He Sun, Rajeev Thakur, Philip C. Roth, and William D. Gropp, "LACIO: A New Collective I/O Strategy for Parallel I/O Systems," in *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2011)*, May 2011.
36. Huaiming Song, Yanlong Yin, Xian-He Sun, Rajeev Thakur and Samuel Lang, "A Segment-Level Adaptive Data Layout Scheme for Improved Load Balance in Parallel File Systems," in *Proceedings of the 11th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2011)*, May 2011.
37. David Goodell, Pavan Balaji, Darius Buntinas, Gábor Dózsa, William Gropp, Sameer Kumar, Bronis R. de Supinski, and Rajeev Thakur, "Minimizing MPI Resource Contention in Multi-threaded Multicore Environments," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2010)*, September 2010.
38. Yong Chen, Xian-He Sun, Rajeev Thakur, Huaiming Song, and Hui Jin, "Improving Parallel I/O Performance with Data Layout Awareness," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2010)*, September 2010.
39. Jayesh Krishna, Pavan Balaji, Ewing Lusk, Rajeev Thakur, and Fabian Tillier, "Implementing MPI on Windows: Comparison with Common Approaches on Unix," in *Proceedings of the 17th European MPI Users' Group Meeting (EuroMPI 2010)*, September 2010.
40. Torsten Hoefler, William Gropp, Rajeev Thakur, and Jesper Larsson Träff, "Toward Performance Models of MPI Implementations for Understanding Application Scaling Issues," in *Proceedings of the 17th European MPI Users' Group Meeting (EuroMPI 2010)*, September 2010.

41. Gábor Dózsa, Sameer Kumar, Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, Joseph Ratterman, and Rajeev Thakur, “Enabling Concurrent Multithreaded MPI Communication on Multicore Petascale Systems,” in *Proceedings of the 17th European MPI Users’ Group Meeting (EuroMPI 2010)*, September 2010.
42. Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, Jayesh Krishna, Ewing Lusk, and Rajeev Thakur, “PMI: A Scalable Parallel Process Management Interface for Extreme-Scale Systems,” in *Proceedings of the 17th European MPI Users’ Group Meeting (EuroMPI 2010)*, September 2010.
43. James Dinan, Pavan Balaji, Ewing Lusk, P. Sadayappan, and Rajeev Thakur, “Hybrid Parallel Programming with MPI and Unified Parallel C,” in *Proceedings of the 2010 ACM International Conference on Computing Frontiers*, May 2010.
44. Seung Woo Son, Samuel Lang, Philip Carns, Robert Ross, Rajeev Thakur, Berkin Ozisikyilmaz, Prabat Kumar, Wei-Keng Liao, and Alok Choudhary, “Enabling Active Storage on Parallel I/O Software Stacks,” in *Proceedings of the 26th IEEE Symposium on Massive Storage Systems and Technologies*, May 2010.
45. Tom Peterka, David Goodell, Robert Ross, Han-Wei Shen, and Rajeev Thakur, “A Configurable Algorithm for Parallel Image Compositing Applications,” in *Proceedings of SC09: International Conference on High Performance Computing, Networking, Storage, and Analysis*, November 2009.
46. Vinod Tipparaju, William Gropp, Hubert Ritzdorf, Rajeev Thakur, and Jesper Larsson Träff, “Investigating High Performance RMA Interfaces for the MPI-3 Standard,” in *Proceedings of the 2009 International Conference on Parallel Processing*, September 2009.
47. Saba Sehrish, Jun Wang, and Rajeev Thakur, “Conflict Detection Algorithm to Minimize Locking For MPI-IO Atomicity,” in *Proceedings of the 16th European PVM/MPI Users’ Group Meeting (Euro PVM/MPI 2009)*, September 2009.
48. Sriram Ananthakrishnan, Michael DeLisi, Sarvani Vakkalanka, Anh Vo, Ganesh Gopalakrishnan, Robert Kirby, and Rajeev Thakur, “How Formal Dynamic Verification Tools Facilitate Novel Concurrency Visualizations,” in *Proceedings of the 16th European PVM/MPI Users’ Group Meeting (Euro PVM/MPI 2009)*, September 2009.
49. Anh Vo, Sarvani Vakkalanka, Jason Williams, Ganesh Gopalakrishnan, Robert Kirby, and Rajeev Thakur, “Sound and Efficient Dynamic Verification of MPI Programs with Probe Non-Determinism,” in *Proceedings of the 16th European PVM/MPI Users’ Group Meeting (Euro PVM/MPI 2009)*, September 2009.
50. Robert Ross, Robert Latham, William Gropp, Ewing Lusk, and Rajeev Thakur, “Processing MPI Datatypes outside MPI,” in *Proceedings of the 16th European PVM/MPI Users’ Group Meeting (Euro PVM/MPI 2009)*, September 2009. (**outstanding paper**)
51. Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, Sameer Kumar, Ewing Lusk, Rajeev Thakur, and Jesper Larsson Träff, “MPI on a Million Processors,” in *Proceedings of the 16th European PVM/MPI Users’ Group Meeting (Euro PVM/MPI 2009)*, September 2009. (**outstanding paper**)
52. G. Santhanaraman, P. Balaji, K. Gopalakrishnan, R. Thakur, W. Gropp, and D. K. Panda. “Natively Supporting True One-sided Communication in MPI on Multi-core Systems with InfiniBand,” in *Proceedings of the 9th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2009)*, May 2009.
53. Anh Vo, Sarvani Vakkalanka, Michael Delisi, Ganesh Gopalakrishnan, Robert M. Kirby, and Rajeev Thakur, “Formal Verification of Practical MPI Programs,” in *Proceedings of the 14th*

*ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2009)*, February 2009.

54. Anthony Chan, Pavan Balaji, William Gropp, and Rajeev Thakur, "Communication Analysis of Parallel 3D FFT for Flat Cartesian Meshes on Large Blue Gene Systems," in *Proceedings of the 15th International Conference on High Performance Computing (HiPC '08)*, December 2008.
55. Pavan Balaji, Sitha Bhagvat, Rajeev Thakur, and Dhabaleswar Panda, "Sockets Direct Protocol for Hybrid Network Stacks: A Case Study with iWARP over 10G Ethernet," in *Proceedings of the 15th International Conference on High Performance Computing (HiPC '08)*, December 2008.
56. Surendra Byna, Yong Chen, Xian-He Sun, Rajeev Thakur, and William Gropp, "Parallel I/O Prefetching Using MPI File Caching and I/O Signatures," in *Proceedings of SC08: International Conference on High Performance Computing, Networking, Storage, and Analysis*, November 2008.
57. Yong Chen, Surendra Byna, Xian-He Sun, Rajeev Thakur, and William Gropp, "Hiding I/O Latency with Pre-execution Prefetching for Parallel Applications," in *Proceedings of SC08: International Conference on High Performance Computing, Networking, Storage, and Analysis*, November 2008. **(finalist for both best student paper award and best paper award)**
58. P. Balaji, A. Chan, W. Gropp, R. Thakur, and E. Lusk, "Non-Data-Communication Overheads in MPI: Analysis on Blue Gene/P," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008. **(outstanding paper)**
59. Jesper Larsson Träff, Andreas Ripke, Christian Siebert, Pavan Balaji, and Rajeev Thakur, and William Gropp, "A Simple, Pipelined Algorithm for Large, Irregular All-gather Problems," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.
60. William Gropp, Dries Kimpe, Robert Ross, Rajeev Thakur and Jesper Larsson Träff, "Self-Consistent MPI-IO Performance Requirements and Expectations," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.
61. Pavan Balaji, Darius Buntinas, David Goodell, William Gropp, and Rajeev Thakur, "Toward Efficient Support for Multithreaded MPI Communication," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.
62. Sarvani Vakkalanka, Michael DeLisi, Ganesh Gopalakrishnan, Robert M. Kirby, Rajeev Thakur, and William Gropp, "Implementing Efficient Dynamic Formal Verification Methods for MPI Programs," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.
63. Subodh Sharma, Sarvani Vakkalanka, Ganesh Gopalakrishnan, Robert M. Kirby, Rajeev Thakur, and William Gropp, "A Formal Approach to Detect Functionally Irrelevant Barriers in MPI Programs," in *Proceedings of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.
64. P. Balaji, W. Feng, H. Lin, J. Archuleta, S. Matsuoka, A. Warren, J. Setubal, E. Lusk, R. Thakur, I. Foster, D. S. Katz, S. Jha, K. Shinpaugh, S. Coghlan, and D. Reed, "Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer," in *Proceedings of the International Supercomputing Conference (ISC'08)*, June 2008. **(Best Paper Award)**
65. P. Balaji, W. Feng, J. Archuleta, H. Lin, R. Kettimuttu, R. Thakur and X. Ma, "Semantics-based Distributed I/O for mpiBLAST," in *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2008)*, February 2008. (short paper)

66. P. Balaji, W. Feng, S. Bhagvat, D. K. Panda, R. Thakur, and W. Gropp, "Analyzing the Impact of Supporting Out-of-Order Communication on In-Order Performance with iWARP," in *Proceedings of SC07: International Conference on High Performance Computing, Networking, Storage, and Analysis*, November 2007.
67. Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby, Robert Palmer, Rajeev Thakur, and William Gropp, "Practical Model Checking Method for Verifying Correctness of MPI Programs," in *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007 pp. 344–353.
68. Robert Latham, William Gropp, Robert Ross, and Rajeev Thakur, "Extending the MPI-2 Generalized Request Interface," in *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007 pp. 223–232.
69. Jesper Larsson Träff, William Gropp, and Rajeev Thakur, "Self-Consistent MPI Performance Requirements," in *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007 pp. 36–45. (**outstanding paper**)
70. Rajeev Thakur and William Gropp, "Test Suite for Evaluating Performance of MPI Implementations That Support MPI\_THREAD\_MULTIPLE," in *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007, pp. 46–55. (**outstanding paper**)
71. William Gropp and Rajeev Thakur, "Revealing the Performance of MPI RMA Implementations," in *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007, pp. 272–280.
72. Prashasta Gujrati, Yawei Li, Zhiling Lan, Rajeev Thakur, and John White, "A Meta-Learning Failure Predictor for BlueGene/L Systems," in *Proceedings of the 2007 International Conference on Parallel Processing*, September 2007.
73. Pavan Balaji, Sitha Bhagvat, Dhableswar Panda, Rajeev Thakur, and William Gropp, "Advanced Flow-Control Mechanisms for the Sockets Direct Protocol over InfiniBand," in *Proceedings of the 2007 International Conference on Parallel Processing*, September 2007.
74. P. Balaji, D. Buntinas, S. Balay, B. Smith, R. Thakur, and W. Gropp, "Nonuniformly Communicating Noncontiguous Data: A Case Study with PETSc and MPI," in *Proceedings of the 21st IEEE International Parallel and Distributed Processing Symposium (IPDPS 2007)*, March 2007.
75. Subhash Saini, Dale Talcott, Rajeev Thakur, Panagiotis Adamidis, Rolf Rabenseifner, and Robert Ciotti, "Parallel I/O Performance Characterization of Columbia and NEC SX-8 Superclusters," in *Proceedings of the 21st IEEE International Parallel and Distributed Processing Symposium (IPDPS 2007)*, March 2007.
76. Kenin Coloma, Avery Ching, Alok Choudhary, Wei-keng Liao, Rob Ross, Rajeev Thakur, and Lee Ward, "A New Flexible MPI Collective I/O Implementation" in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2006)*, September 2006.
77. Robert Latham, Robert Ross, and Rajeev Thakur, "Can MPI Be Used for Persistent Parallel Services?" in *Proceedings of the 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006)*, September 2006, pp. 275–284.
78. William Gropp and Rajeev Thakur, "Issues in Developing a Thread-Safe MPI Implementation," in *Proceedings of the 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006)*, September 2006, pp. 12–21. (**outstanding paper**)
79. Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby, Rajeev Thakur, and William Gropp, "Formal Verification of Programs That Use MPI One-Sided Communication," in *Proceedings*



- of the 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006), September 2006, pp. 30–39. (**outstanding paper**)
80. Surendra Byna, Xian-He Sun, Rajeev Thakur, and William Gropp, "Automatic Memory Optimizations for Improving MPI Derived Datatype Performance," in *Proceedings of the 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006)*, September 2006, pp. 238–246.
  81. Jonghyun Lee, Robert Ross, Scott Atchley, Micah Beck, and Rajeev Thakur, "MPI-IO/L: Efficient Remote I/O for MPI-IO via Logistical Networking," in *Proceedings of the 20th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2006)*, April 2006.
  82. Ernie Chan, Robert van de Geijn, William Gropp, Rajeev Thakur, "Collective Communication on Architectures that Support Simultaneous Communication over Multiple Links," in *Proceedings of the ACM SIGPLAN 2006 Symposium on Principles and Practice of Parallel Programming (PPoPP 2006)*, March 2006.
  83. H. Yu, R. K. Sahoo, C. Howson, G. Almasi, J. G. Castanos, M. Gupta J. E. Moreira, J. J. Parker, T. E. Engelsiepen, R. Ross, R. Thakur, R. Latham, and W. D. Gropp, "High Performance File I/O for the BlueGene/L Supercomputer," in *Proceedings of the 12th International Symposium on High-Performance Computer Architecture (HPCA-12)*, February 2006.
  84. Robert Latham, Robert Ross, Rajeev Thakur, and Brian Toonen, "Implementing MPI-IO Shared File Pointers without File System Support," in *Proceedings of the 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005)*, September 2005, pp. 84–93.
  85. Rajeev Thakur, Robert Ross, and Robert Latham, "Implementing Byte-Range Locks Using MPI One-Sided Communication," in *Proceedings of the 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005)*, September 2005, pp. 119–128.
  86. William Gropp and Rajeev Thakur, "An Evaluation of Implementation Options for MPI One-Sided Communication," in *Proceedings of the 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005)*, September 2005, pp. 415–424.
  87. Robert Ross, Robert Latham, William Gropp, Rajeev Thakur, and Brian Toonen, "Implementing MPI-IO Atomic Mode Without File System Support," in *Proceedings of the 5th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2005)*, May 2005.
  88. Surendra Byna, Xian-He Sun, William Gropp, and Rajeev Thakur, "Predicting Memory-Access Cost Based on Data-Access Patterns," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2004)*, September 2004.
  89. Jonghyun Lee, Xiaosong Ma, Robert Ross, Rajeev Thakur, and Marianne Winslett, "RFS: Efficient and Flexible Remote File Access for MPI-IO," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2004)*, September 2004.
  90. Rob Latham, Rob Ross, and Rajeev Thakur, "The Impact of File Systems on MPI-IO Scalability," in *Proceedings of the 11th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2004)*, September 2004, pp. 87–96.
  91. Weihang Jiang, Jiuxing Liu, Hyun-Wook Jin, Dhableswar K. Panda, Darius Buntinas, Rajeev Thakur, and William Gropp, "Efficient Implementation of MPI-2 Passive One-Sided Communication on InfiniBand Clusters," in *Proceedings of the 11th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2004)*, September 2004, pp. 68–76.
  92. Rajeev Thakur, William Gropp, and Brian Toonen, "Minimizing Synchronization Overhead in the Implementation of MPI One-Sided Communication," in *Proceedings of the 11th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2004)*, September 2004, pp. 57–67.

93. Weihang Jiang, Jiuxing Liu, Hyun-Wook Jin, Dhabaleswar K. Panda, William Gropp, and Rajeev Thakur, "High Performance MPI-2 One-Sided Communication over Infiniband," in *Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2004)*, April 2004.
94. Murali Vilayannur, Robert B. Ross, Philip H. Carns, Rajeev Thakur, Anand Sivasubramaniam, and Mahmut Kandemir, "On the Performance of the POSIX I/O Interface to PVFS," in *Proceedings of the 12th Euromicro Conference on Parallel, Distributed, and Network-based Processing*, February 2004, pp. 332–339.
95. Surendra Byna, William Gropp, Xian-He Sun, and Rajeev Thakur, "Improving the Performance of MPI Derived Datatypes by Optimizing Memory-Access Cost," in *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2003)*, December 2003, pp. 412–419.
96. Jianwei Li, Wei-keng Liao, Alok Choudhary, Robert Ross, Rajeev Thakur, William Gropp, Rob Latham, Andrew Siegel, Brad Gallagher, and Michael Zingale, "Parallel netCDF: A High-Performance Scientific I/O Interface," in *Proceedings of SC2003: High Performance Networking and Computing*, November 2003.
97. Rajeev Thakur and William Gropp, "Improving the Performance of Collective Operations in MPICH," in *Proceedings of the 10th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2003)*, September 2003, pp. 257–267.
98. Murali Vilayannur, Anand Sivasubramaniam, Mahmut Kandemir, Rajeev Thakur, and Robert Ross, "Discretionary Caching for I/O on Clusters," in *Proceedings of the 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003)*, May 2003, pp. 96–103.
99. Jaechun No, Rajeev Thakur, Dinesh Kaushik, Lori Freitag, and Alok Choudhary, "A Scientific Data Management System for Irregular Applications," in *Proceedings of the Eighth International Workshop on Solving Irregular Problems in Parallel (Irregular 2001)*, April 2001.
100. Jaechun No, Rajeev Thakur, and Alok Choudhary, "Integrating Parallel File I/O and Database Support for High-Performance Scientific Data Management," in *Proceedings of SC2000: High Performance Networking and Computing*, November 2000.
101. Philip H. Carns, Walter B. Ligon III, Robert B. Ross, and Rajeev Thakur, "PVFS: A Parallel File System for Linux Clusters," in *Proceedings of the 4th Annual Linux Showcase and Conference, Atlanta*, October 2000, pp. 317–327. (**Best Paper Award**)
102. Phillip Dickens and Rajeev Thakur, "An Evaluation of Java's I/O Capabilities for High-Performance Computing," in *Proceedings of the ACM 2000 Java Grande Conference*, June 2000, pp. 26–35.
103. A. Choudhary, M. Kandemir, H. Nagesh, J. No, X. Shen, V. Taylor, S. More, and R. Thakur, "Data Management for Large-Scale Scientific Computations in High Performance Distributed Systems," in *Proceedings of the Eighth IEEE International Symposium on High Performance Distributed Computing*, August 1999, pp. 263–272.
104. Rajeev Thakur, William Gropp, and Ewing Lusk, "On Implementing MPI-IO Portably and with High Performance," in *Proceedings of the Sixth Workshop on I/O in Parallel and Distributed Systems*, May 1999, pp. 23–32.
105. Phillip Dickens and Rajeev Thakur, "Improving Collective I/O Performance Using Threads," in *Proceedings of the 13th International Parallel Processing Symposium and 10th Symposium on Parallel and Distributed Processing*, April 1999, pp. 38–45.
106. Rajeev Thakur, William Gropp, and Ewing Lusk, "Data Sieving and Collective I/O in ROMIO," in *Proceedings of the 7th Symposium on the Frontiers of Massively Parallel Computation*, February 1999, pp. 182–189.

107. Rajeev Thakur, William Gropp, and Ewing Lusk, "A Case for Using MPI's Derived Datatypes to Improve I/O Performance," in *Proceedings of SC98: High Performance Networking and Computing*, November 1998.
108. Phillip Dickens and Rajeev Thakur, "A Performance Study of Two-Phase I/O," in *Proceedings of Euro-Par '98*, September 1998, pp. 959–965.
109. Rajeev Thakur, William Gropp, and Ewing Lusk, "An Abstract-Device Interface for Implementing Portable Parallel-I/O Interfaces," in *Proceedings of the 6th Symposium on the Frontiers of Massively Parallel Computation*, October 1996, pp. 180–187.
110. Rajeev Thakur, William Gropp, and Ewing Lusk, "An Experimental Evaluation of the Parallel I/O Systems of the IBM SP and Intel Paragon Using a Production Application," in *Proceedings of the 3rd International Conference of the Austrian Center for Parallel Computation with special emphasis on Parallel Databases and Parallel I/O*, Lecture Notes in Computer Science 1127, Springer-Verlag, September 1996, pp. 24–35.
111. Kevin Roe, Rajeev Thakur, Thong Dang, and Edward Bogucz, "Implementation of a 3D Mixing Layer Code on Parallel Computers," in *Proceedings of AIAA 6th International Aerospace Planes and Hypersonics Technologies Conference*, April 1995.
112. Rajeev Thakur, Rajesh Bordawekar, and Alok Choudhary, "Compiler and Runtime Support for Out-of-Core HPF Programs," in *Proceedings of International Conference on Supercomputing*, July 1994, pp. 382–391.
113. Rajeev Thakur, Alok Choudhary, and Geoffrey Fox, "Runtime Array Redistribution in HPF Programs," in *Proceedings of Scalable High Performance Computing Conference*, May 1994, pp. 309–316.
114. Rajeev Thakur, Rajesh Bordawekar, and Alok Choudhary, "Compilation of Out-of-Core Data Parallel Programs for Distributed Memory Machines," in *Proceedings of the Workshop on I/O in Parallel Computer Systems at IPPS '94*, April 1994, pp. 54–72.
115. Rajeev Thakur and Alok Choudhary, "All-to-All Communication on Meshes with Wormhole Routing," in *Proceedings of 8th International Parallel Processing Symposium*, April 1994, pp. 561–565.
116. Rajeev Thakur, Alok Choudhary, and Geoffrey Fox, "Complete Exchange on a Wormhole Routed Mesh," in *Proceedings of MASCOTS '94*, January 1994, pp. 131–135.
117. Ravi Ponnusamy, Rajeev Thakur, Alok Choudhary, and Geoffrey Fox, "Scheduling Regular and Irregular Communication Patterns on the CM-5," in *Proceedings of Supercomputing '92*, November 1992, pp. 394–402. **(Best Student Paper Award in the category of performance measurement)**
118. Alok Choudhary and Rajeev Thakur, "Evaluation of Connected Component Labeling Algorithms on Shared and Distributed Memory Multiprocessors," in *Proceedings of 6th International Parallel Processing Symposium*, March 1992, pp. 362–365.

### Other Conference Proceedings

1. Rajeev Thakur and William Gropp, "Open Issues in MPI Implementation," in *Proceedings of the 12th Asia-Pacific Computer Systems Architecture Conference (ACSAC 2007)*, August 2007, pp. 327–338.
2. Rajeev Thakur, Alok Choudhary, Rajesh Bordawekar, Sachin More, and K. Sivaram, "PASSION Runtime Library for the Intel Paragon," in *Proceedings of the Intel Supercomputer User's Group Conference*, June 1995.

3. Rajeev Thakur, Rajesh Bordawekar, Alok Choudhary, Ravi Ponnusamy, and Tarvinder Singh, "PASSION Runtime Library for Parallel I/O," in *Proceedings of the Scalable Parallel Libraries Conference*, October 1994, pp. 119–128.
4. Zeki Bozkus, Alok Choudhary, Geoffrey Fox, Tomasz Haupt, Sanjay Ranka, Rajeev Thakur, and J.C. Wang, "Scalable Libraries for High Performance Fortran," in *Proceedings of Scalable Parallel Libraries Conference*, October 1993, pp. 67–75.
5. Ishfaq Ahmad, Rajesh Bordawekar, Zeki Bozkus, Alok Choudhary, Geoffrey Fox, Kanchana Parasuram, Ravi Ponnusamy, Rajeev Thakur, and Sanjay Ranka, "Fortran 90D Intrinsic Functions on Distributed Memory Machines: Implementation and Scalability," in *Proceedings of 26th Hawaii International Conference on System Sciences*, January 1993.

## Technical Reports

1. Jack Dongarra, Pete Beckman, Terry Moore, Patrick Aerts, Giovanni Aloisio, Jean-Claude Andre, David Barkai, Jean-Yves Berthou, Taisuke Boku, Bertrand Braunschweig, Franck Cappello, Barbara Chapman, Xuebin Chi, Alok Choudhary, Sudip Dosanjh, Thom Dunning, Sandro Fiore, Al Geist, Bill Gropp, Robert Harrison, Mark Hereld, Michael Heroux, Adolfo Hoisie, Koh Hotta, Zhong Jin, Yutaka Ishikawa, Fred Johnson, Sanjay Kale, Richard Kenway, David Keyes, Bill Kramer, Jesus Labarta, Alain Lichnewsky, Thomas Lippert, Bob Lucas, Barney Maccabe, Satoshi Matsuoka, Paul Messina, Peter Michielse, Bernd Mohr, Matthias S. Mueller, Wolfgang E. Nagel, Hiroshi Nakashima, Michael E Papka, Dan Reed, Mitsuhsa Sato, Ed Seidel, John Shalf, David Skinner, Marc Snir, Thomas Sterling, Rick Stevens, Fred Streit, Bob Sugar, Shinji Sumimoto, William Tang, John Taylor, Rajeev Thakur, Anne Trefethen, Mateo Valero, Aad van der Steen, Jeffrey Vetter, Peg Williams, Robert Wisniewski, and Kathy Yelick, "The International Exascale Software Project Roadmap," *International Journal of High Performance Computing Applications*, 25(1):3–60, February 2011.
2. Tarek El-Ghazawi, Francois Cantonnet, Proshanta Saha, Rajeev Thakur, Rob Ross, and Dan Bonachea, "UPC-IO: A Parallel I/O API for UPC," Version 1.0, Technical Report, High Performance Computing Laboratory, George Washington University, July 2004.
3. Rajeev Thakur, Ewing Lusk, and William Gropp, "Users Guide for ROMIO: A High-Performance, Portable MPI-IO Implementation," Technical Memorandum ANL/MCS-TM-234, Mathematics and Computer Science Division, Argonne National Laboratory, revised September 2000.
4. Rajeev Thakur, Ewing Lusk, and William Gropp, "I/O Characterization of a Portable Astrophysics Application on the IBM SP and Intel Paragon," Preprint MCS-P534-0895, Mathematics and Computer Science Division, Argonne National Laboratory, August 1995.
5. Rajesh Bordawekar, Rajeev Thakur, and Alok Choudhary, "Data Access Reorganizations in Compiling Out-of-Core Data Parallel Programs on Distributed Memory Machines," NPAC Technical Report SCCS-622, Syracuse University, revised September 1994.
6. Alok Choudhary, Rajesh Bordawekar, Michael Harry, Rakesh Krishnaiyer, Ravi Ponnusamy, Tarvinder Singh, and Rajeev Thakur, "PASSION: Parallel and Scalable Software for I/O," NPAC Technical Report SCCS-636, Syracuse University, September 1994.

## Invited Talks

1. "Exascale Computing Project: Software Technology Perspective," Keynote at JST/CREST International Symposium on Post Petascale System Software, Tokyo, Japan, December 2017.

2. "MPI-IO: A Retrospective," 25th Anniversary of MPI Workshop, Argonne, Illinois, September 2017.
3. "Exascale Computing Project: Software Technology Perspective," Keynote at 15th Annual Workshop on Charm++ and its Applications, Champaign, Illinois, April 2017.
4. "Exascale Computing Project: Software Technology Perspective," Keynote at ExaMPI 2016 Workshop at SC16, Salt Lake City, Utah, November 2016.
5. "Future of MPI," DOE Programming Environments Workshop, Rockville, Maryland, March 2015.
6. "Future Node Architectures and their Implications for MPI," Workshop on Clusters, Clouds, and Data for Scientific Computing, Dareizé, France, September 2014.
7. "Evolving MPI to Address the Challenges of Exascale Systems," 20th European MPI Users Group Meeting (EuroMPI 2013), Madrid, Spain, September 2013.
8. "Update on MPI and OS/R Activities at Argonne," Ninth Workshop of the INRIA-Illinois Joint Laboratory on Petascale Computing, Lyon, France, June 2013.
9. "Challenges for Communication Libraries and Runtime Systems at Exascale," Workshop on Clusters, Clouds, and Data for Scientific Computing, Dareizé, France, September 2012.
10. "Recent Activities in Programming Models and Runtime Systems at ANL," Seventh Workshop of the INRIA-Illinois Joint Laboratory on Petascale Computing, Rennes, France, June 2012.
11. "Challenges in Scaling MPI to Exascale," First Workshop of the MCS-ISCAS Joint Laboratory for Parallel Processing and Computing Technologies, Beijing, China, May 2012.
12. "Challenges in Scaling MPI to Exascale," Sixth Workshop of the INRIA-Illinois Joint Laboratory on Petascale Computing, Champaign, Illinois, November 2011.
13. "Future Directions in MPI," Workshop on Clusters, Clouds, and Grids for Scientific Computing, Asheville, North Carolina, September 2010.
14. "MPI at Exascale," SciDAC 2010 conference, Chattanooga, Tennessee, July 2010.
15. "Programming Models for Large-Scale Parallel Machines," Workshop on Trends, Technologies and Collaborative Opportunities in High Performance Computing (WTTC09), Mae Fah Luang University, Chiang Rai, Thailand, May 2009.
16. "Programming Models for Large-Scale Parallel Machines," Workshop on Trends, Technologies and Collaborative Opportunities in High Performance and Grid Computing (WTTC2008), Phuket, Thailand, June 2008.
17. "MPI: Enabling Portable, High-Performance Parallel Computing," SHARCNET Workshop, University of Western Ontario, London, Ontario, Canada, October 2003.
18. "MPICH on Clusters: Future Directions," Linux Supercluster Users Conference, Albuquerque, New Mexico, September 2000.
19. "MPI-IO: A Standard, Portable API for High-Performance Parallel I/O," Extreme Linux Workshop, Monterey, California, June 1999.
20. "Portable Parallel I/O on Cluster Computers," Workshop on High Performance Cluster Computing at the 11th ACM International Conference on Supercomputing, Vienna, Austria, July 1997.

### **Seminars and Colloquia**

1. "Future of MPI and MPICH," Électricité de France, Paris, France, June 2015.
2. "Future Directions in MPI," Fermilab Computing Division Seminar, February 2011.
3. "Programming Extreme Scale Parallel Computers," Department of Electrical Engineering and Computer Science, University of Central Florida, July 2010.

4. "MPI on Millions of Cores," LANS seminar, Argonne National Laboratory, January 2010.
5. "Multithreaded MPI for Multicore Systems," IBM Research Multicore Seminar Series, May 2008.
6. "Programming Models for Next Generation Parallel Machines," IBM T. J. Watson Research Center, Yorktown Heights, NY, March 2008.
7. "MPICH2: Recent Developments and Future Directions," University of Chicago Computation Institute Brown Bag Seminar Series, June 2006.
8. "MPI: Enabling Portable, High-Performance Parallel Programming," Mahidol University, Bangkok, Thailand, May 2006.
9. "The MPI Standard: Design, Implementation, and Verification," University of Utah School of Computing, Salt Lake City, Utah, November 2005.
10. "MPI-IO: Design, Implementation, and Performance," Compaq Computer Corporation, Nashua, New Hampshire, March 2002.
11. "Scalable Systems Software for Large Linux Clusters," IBM T. J. Watson Research Center, Yorktown Heights, New York, September 2000.
12. "Portable Parallel I/O with MPI-IO," Sandia National Laboratories, Albuquerque, New Mexico, January 1998.
13. "Portable Parallel I/O with MPI-IO," Hewlett-Packard, Richardson, Texas, October 1997.
14. "ROMIO: A High-Performance, Portable MPI-IO Implementation," Aurora Project Workshop, European Centre for Parallel Computing at Vienna (VCPC), Austria, July 1997.
15. "A Framework for High-Performance, Portable Parallel I/O," Department of Computer Science, Indiana University-Purdue University at Indianapolis, March 1997.

### **Other Presentations**

1. "ECP Software Technology Update," DOE ASCR Computer Science PI Meeting, Bethesda, Maryland, March 2017.
2. "Evolving MPI to Address the Challenges of Exascale Systems," DOE X-Stack Portfolio PI Meeting, Berkeley, CA, March 2013.
3. "Future Directions in MPI," FLASH HEDP Exascale Codesign Center Meeting, Argonne, IL, February 2011.
4. "Implementing MPI on Windows: Comparison with Common Approaches on Unix," 17th European MPI Users' Group Meeting (EuroMPI 2010), Stuttgart, Germany, September 2010.
5. "MPI on a Million Cores," Scientific Data Management SciDAC All-Hands Meeting, Davis, California, October 2009.
6. "Conflict Detection Algorithm to Minimize Locking For MPI-IO Atomicity," 16th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2009), Helsinki, Finland, September 2009.
7. "MPI on a Million Processors," 16th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2009), Helsinki, Finland, September 2009.
8. "The Future of Deep Computing System Software," IBM Deep Computing Institute External Advisory Board Meeting, San Diego, CA, February 2008.
9. "Test Suite for Evaluating Performance of MPI Implementations That Support MPI\_THREAD\_MULTIPLE," 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007), Paris, France, September 2007.
10. "Open Issues in MPI Implementation," 12th Asia-Pacific Computer Systems Architecture Conference (ACSAC 2007), Seoul, Korea, August 2007.

11. "Automatic Memory Optimizations for Improving MPI Derived Datatype Performance," 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006), Bonn, Germany, September 2006.
12. "Opportunities in Parallel I/O for Scientific Data Management," Scientific Data Management SciDAC All-Hands Meeting, Raleigh, North Carolina, October 2005.
13. "Implementing Byte-Range Locks Using MPI One-Sided Communication," 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005), Sorrento, Italy, September 2005.
14. "An Evaluation of Implementation Options for MPI One-Sided Communication," 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005), Sorrento, Italy, September 2005.
15. "MPI-IO Needs and Requirements," DARPA HPCS I/O Workshop, Washington, DC, July 2005.
16. "I/O Consistency Semantics and Performance Implications," 9th Workshop on Distributed Supercomputing (SOS9), Davos, Switzerland, March 2005.
17. "Minimizing Synchronization Overhead in the Implementation of MPI One-Sided Communication," 11th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2004), Budapest, Hungary, September 2004.
18. "Improving the Performance of MPI Derived Datatypes by Optimizing Memory-Access Cost," IEEE International Conference on Cluster Computing (Cluster 2003), Hong Kong, December 2003.
19. "Improving the Performance of Collective Operations in MPICH," 10th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2003), Venice, Italy, September 2003.
20. "I/O on Clusters," 7th Workshop on Distributed Supercomputing (SOS7), Durango, Colorado, March 2003.
21. "Parallel I/O Efforts at Argonne: A Brief Introduction," Microsoft Corporation, Redmond, Washington, April 2002.
22. "Chiba City: An Open Source Computer Science Testbed," Cluster Computing Workshop, New Orleans, Louisiana, March 2000.
23. "OS, Message Passing, and Runtime Tools," Panel Session on OS, Message Passing, and Runtime Tools, Cluster Computing Workshop, New Orleans, Louisiana, March 2000.
24. "MPI I/O Perspectives," Panel Session on Scalable I/O, Cluster Computing Workshop, New Orleans, Louisiana, March 2000.
25. "How Can File Systems Be Designed to Better Support MPI-IO?" Scalable Global Parallel File System Meeting, Santa Fe, New Mexico, September 1999.
26. "On Implementing MPI-IO Portably and with High Performance," Sixth Workshop on I/O in Parallel and Distributed Systems, Atlanta, Georgia, May 1999.
27. "High-Performance Parallel I/O in MPI," MPICH-LANL meeting, Los Alamos, New Mexico, March 1999.
28. "Data Sieving and Collective I/O in ROMIO," 7th Symposium on the Frontiers of Massively Parallel Computation, Annapolis, Maryland, February 1999.
29. "Achieving High-Performance with MPI-IO," SIO-ASCI Meeting, Livermore, California, January 1999.
30. "A Case for Using MPI's Derived Datatypes to Improve I/O Performance," SC98: High Performance Networking and Computing, Orlando, Florida, November 1998.
31. "ROMIO: A High-Performance, Portable MPI-IO Implementation," Scalable I/O Initiative Sponsors Review meeting, Rice University, Houston, Texas, September 1997.

32. "ROMIO: A High-Performance, Portable MPI-IO Implementation," Scalable I/O Initiative All-Hands Meeting, Center for Advanced Computing Research, California Institute of Technology, June 1997.
33. "MPI-IO and Its Implementation," 2nd Midwest Workshop on High-Performance Systems, University of Illinois at Chicago, May 1997.
34. "Portable Implementation of Parallel-I/O APIs," Minisymposium on Parallel I/O at the Eighth SIAM Conference on Parallel Processing for Scientific Computing, Minneapolis, Minnesota, March 1997.
35. "Implementing High-Level Interfaces on the SIO Low-Level API," Scalable I/O Roundtable at Supercomputing '96, Pittsburgh, Pennsylvania, November 1996.
36. "An Abstract-Device Interface for Implementing Portable Parallel-I/O Interfaces," 6th Symposium on the Frontiers of Massively Parallel Computation, Annapolis, Maryland, October 1996.
37. "An Experimental Evaluation of the Parallel I/O Systems of the IBM SP and Intel Paragon Using a Production Application," 3rd International Conference of the Austrian Center for Parallel Computation with special emphasis on Parallel Databases and Parallel I/O, Klagenfurt, Austria, September 1996.
38. "Portable Implementation of MPI-IO," Panel discussion on MPI-IO at the MPI Developers Conference and Users Group Meeting, Notre Dame, Indiana, July 1996.
39. "An Abstract-Device Interface for Parallel I/O," Annual Meeting of the Center for Research on Parallel Computation, Argonne, Illinois, May 1996.
40. "I/O Characterization of a Portable Astrophysics Application," Scalable I/O Workshop at Supercomputing '95, San Diego, California, December 1995.
41. "PASSION Runtime Library for Parallel I/O," Scalable Parallel Libraries Conference, Mississippi State University, Starkville, Mississippi, October 1994.
42. "Compiler and Runtime Support for Out-of-Core HPF Programs," International Conference on Supercomputing, Manchester, UK, July 1994.
43. "Runtime Array Redistribution in HPF Programs," Scalable High Performance Computing Conference, Knoxville, Tennessee, May 1994.
44. "Scheduling Regular and Irregular Communication Patterns on the CM-5," Supercomputing '92, Minneapolis, Minnesota, November 1992.

## Tutorials

1. "Advanced MPI Programming," at SC17: International Conference on High Performance Computing, Networking, Storage, and Analysis, Denver, Colorado, November 2017. (with Pavan Balaji, William Gropp, and Torsten Hoefler)
2. "Advanced MPI Programming," at SC16: International Conference on High Performance Computing, Networking, Storage, and Analysis, Salt Lake City, Utah, November 2016. (with Pavan Balaji, William Gropp, and Torsten Hoefler)
3. "Advanced MPI Programming," at SC15: International Conference on High Performance Computing, Networking, Storage, and Analysis, Austin, Texas, November 2015. (with Pavan Balaji, William Gropp, and Torsten Hoefler)
4. "Advanced MPI Programming," at SC14: International Conference on High Performance Computing, Networking, Storage, and Analysis, New Orleans, Louisiana, November 2014. (with Pavan Balaji, William Gropp, and Torsten Hoefler)



5. "Advanced MPI Programming," at SC13: International Conference on High Performance Computing, Networking, Storage, and Analysis, Denver, Colorado, November 2013. (with Pavan Balaji, James Dinan, and Torsten Hoefler)
6. "Advanced MPI," at SC12: International Conference on High Performance Computing, Networking, Storage, and Analysis, Salt Lake City, Utah, November 2012. (with William Gropp, Ewing Lusk, and Robert Ross)
7. "Advanced MPI," at the 19th European MPI Users' Group Meeting (EuroMPI 2012), Vienna, Austria, September 2012. (with William Gropp and Ewing Lusk)
8. "Advanced MPI," at SC11: International Conference on High Performance Computing, Networking, Storage, and Analysis, Seattle, Washington, November 2011. (with William Gropp, Ewing Lusk, and Robert Ross)
9. "MPI and Parallel I/O," CHPC 2011 Winter School at The University of the Witwatersrand, Johannesburg, South Africa, July 2011.
10. "Advanced MPI," at SC10: International Conference on High Performance Computing, Networking, Storage, and Analysis, New Orleans, Louisiana, November 2010. (with William Gropp, Ewing Lusk, and Robert Ross)
11. "MPI and Parallel I/O," PRACE Autumn School on High Performance Computing, Barcelona Supercomputing Center, Barcelona, Spain, October 2010.
12. "Programming in MPI for Performance and MPI at Exascale," CScADS Summer Workshop on Leadership-class Machines, Petascale Applications, and Performance Strategies, Snowbird, Utah, July 2010
13. "Advanced MPI," at SC09: International Conference on High Performance Computing, Networking, Storage, and Analysis, Portland, Oregon, November 2009. (with William Gropp, Ewing Lusk, and Robert Ross)
14. "Advanced MPI," at the Workshop on Trends, Technologies and Collaborative Opportunities in High Performance Computing (WTTC09), Bangkok, Thailand, May 2009.
15. "Advanced MPI," at SC08: International Conference on High Performance Computing, Networking, Storage, and Analysis, Austin, Texas, November 2008. (with William Gropp, Ewing Lusk, and Robert Ross)
16. "Advanced MPI," at the Workshop on Trends, Technologies and Collaborative Opportunities in High Performance and Grid Computing (WTTC2008), Bangkok, Thailand, June 2008.
17. "Parallel I/O in Practice," at SC07: International Conference on High Performance Computing, Networking, Storage, and Analysis, Reno, Nevada, November 2007. (with Robert Ross, Robert Latham, and Bill Loewe)
18. "Advanced MPI," at SC07: International Conference on High Performance Computing, Networking, Storage, and Analysis, Reno, Nevada, November 2007. (with William Gropp, Ewing Lusk, and Robert Ross)
19. "Parallel I/O in Practice," at SC06: International Conference on High Performance Computing, Networking, Storage, and Analysis, Tampa, Florida, November 2006. (with Robert Ross, Robert Latham, and Bill Loewe)
20. "Advanced MPI: I/O and One-Sided Communication," at SC06: International Conference on High Performance Computing, Networking, Storage, and Analysis, Tampa, Florida, November 2006. (with William Gropp, Ewing Lusk, and Robert Ross)
21. "Introduction to MPI-IO and MPI-2," 12th Annual SDSC Summer Institute, San Diego Supercomputer Center, San Diego, California, July 2006.
22. "MPI on the Grid," at the 6th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2006), Singapore, May 2006. (with William Gropp)

23. "Parallel I/O in Practice," at SC|05: International Conference on High Performance Computing, Networking, and Storage, Seattle, Washington, November 2005. (with Robert Ross, Robert Latham, and Bill Loewe)
24. "Advanced MPI: I/O and One-Sided Communication," at SC|05: International Conference on High Performance Computing, Networking, and Storage, Seattle, Washington, November 2005. (with William Gropp, Ewing Lusk, and Robert Ross)
25. "Introduction to Parallel I/O and MPI-IO," 11th Annual Computing Institute, San Diego Supercomputer Center, San Diego, California, July 2005.
26. "Advanced MPI: I/O and One-Sided Communication," at SC2004: High Performance Computing, Networking, and Storage Conference, Pittsburgh, Pennsylvania, November 2004. (with William Gropp, Ewing Lusk, and Robert Ross)
27. "Using MPI-2: Advanced Features of the Message-Passing Interface," at the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2004), Chicago, Illinois, April 2004. (with Robert Ross)
28. "Using MPI-2: Advanced Features of the Message-Passing Interface," at the Second International Workshop on Grid and Cooperative Computing (GCC 2003), Shanghai, China, December 2003. (with Robert Ross)
29. "Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface Standard," at SC2003: High Performance Networking and Computing, Phoenix, Arizona, November 2003. (with William Gropp, Ewing Lusk, and Robert Ross)
30. "Using MPI for Peak Performance," SHARCNET Workshop, Univ. of Western Ontario, London, Ontario, Canada, October 2003. (with Brian Toonen)
31. "Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface Standard," at SC2002: High Performance Networking and Computing, Baltimore, Maryland, November 2002. (with William Gropp, Ewing Lusk, and Robert Ross)
32. "Parallel I/O in MPI-2," in tutorial on Porting to PACI Architectures, NPACI All-Hands Meeting, San Diego, California, March 2002.
33. "Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface," at SC2001: High Performance Networking and Computing, Denver, Colorado, November 2001. (with William Gropp, Ewing Lusk, and Robert Ross)
34. "Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface," at SC2000: High Performance Networking and Computing, Dallas, Texas, November 2000. (with William Gropp and Ewing Lusk)
35. "Tuning MPI Applications for Peak Performance," at SC99: High Performance Networking and Computing, Portland, Oregon, November 1999. (with William Gropp and Ewing Lusk)
36. "Tuning MPI Applications for Peak Performance," at SC98: High Performance Networking and Computing, Orlando, Florida, November 1998. (with William Gropp and Ewing Lusk)
37. "Introduction to Performance Issues in Using MPI for Communication and I/O," at the 7th IEEE International Conference on High Performance Distributed Computing, Chicago, Illinois, July 1998. (with William Gropp)