

Rajeev Thakur

Computer Scientist
Mathematics and Computer Science Division
Argonne National Laboratory
9700 S. Cass Avenue, Argonne, IL 60439
thakur@mcs.anl.gov

Education

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

Appointments

- Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL
 - Computer Scientist, April 2002–present
 - 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
 - Fellow, August 2006–present
- Dept. of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL
 - Adjunct Associate Professor, September 2006–present
- Prismedia Networks, Inc., San Jose, CA
 - Senior Scientist, December 2000–December 2001

Awards

- Best Student Paper Award at Supercomputing '92 in the category of performance measurement, November 1992
- Best Paper Award, 4th Annual Linux Showcase and Conference, October 2000
- Best Poster Award, SC2003: High Performance Networking and Computing
- 2005 R&D 100 Award for “MPICH2: High Performance, Portable Implementation of the Message Passing Interface Standard for Parallel Computing”
- Two outstanding papers at Euro PVM/MPI 2006
- Two outstanding papers at Euro PVM/MPI 2007
- One outstanding paper at Euro PVM/MPI 2008
- Best Paper Award at the International Supercomputing Conference (ISC'08)

Professional Activities

- Associate Editor, IEEE Transactions on Parallel and Distributed Systems, 2004-2007
- Tutorials Vice-Chair, Supercomputing '96, Pittsburgh, Pennsylvania, November 1996
- Publicity Chair, Fifth Workshop on I/O in Parallel and Distributed Systems (IOPADS '97), San Jose, California, November 1997
- Program Committee, 4th International Conference on High Performance Computing (HiPC '97), Bangalore, India, December 1997
- Tutorials Vice-Chair, SC97: High Performance Networking and Computing, San Jose, California, November 1997
- Co-Guest Editor, Special Issue of the *International Journal of High Performance Computing Applications* on "I/O in Parallel Applications," 1998
- Vice-Chair, Parallel I/O Workshop at Euro-Par '98, Southampton, UK, September 1998
- Poster Exhibits Chair, SC98: High Performance Networking and Computing, Orlando, Florida, November 1998
- Program Committee, Heterogeneous Computing Workshop at IPPS '99, San Juan, Puerto Rico, April 1999
- General Co-Chair, Sixth Workshop on I/O in Parallel and Distributed Systems (IOPADS '99), Atlanta, Georgia, May 1999
- Two proposal-review panels at the National Science Foundation, 1999
- Program Committee, 6th International Conference on High Performance Computing (HiPC '99), Calcutta, India, December 1999
- Global Chair, Parallel I/O and Storage Technology track at the European Conference on Parallel Computing (Euro-Par 2000), Munich, Germany, August 2000
- Chair for Birds-of-a-Feather sessions, SC2000: High Performance Networking and Computing, Dallas, Texas, November 2000
- Program Committee, IEEE International Symposium on Cluster Computing and the Grid (CCGrid'2001), Brisbane, Australia, May 2001
- Editorial Board of Special Issue of *Calculateurs Parallèles* on Parallel I/O for Cluster Computing, 2001
- Vice Program Chair, Cluster 2002: IEEE International Conference on Cluster Computing, Chicago, Illinois, September 2002
- Program Committee, Workshop on Parallel I/O in Cluster Computing and Computational Grids at CCGrid 2003, Tokyo, Japan, May 2003
- Program Committee, IEEE International Conference on Cluster Computing, Hong Kong, December 2003
- Program Committee, International Conference on Parallel Processing, Montreal, Canada, August 2004
- Program Committee, Ninth Asia-Pacific Computer Systems Architecture Conference, Beijing, China, September 2004
- Program Committee, IEEE International Conference on Cluster Computing, San Diego, California, September 2004
- Program Committee, First Workshop on System Management Tools for Large-Scale Parallel Systems at IPDPS 2005, Denver, Colorado, April 2005
- Program Committee, Tenth Asia-Pacific Computer Systems Architecture Conference, Singapore, October 2005

- Poster/Presentation Chair, 12th IEEE International Conference on High Performance Computing (HiPC '05), Goa, India, December 2005
- Program Committee, Second Workshop on System Management Tools for Large-Scale Parallel Systems at IPDPS 2006, Rhodes Island, Greece, April 2006
- Program Committee, Eleventh Asia-Pacific Computer Systems Architecture Conference, Shanghai, China, June 2006
- Program Committee, Twelfth International Conference on Parallel and Distributed Systems, Minneapolis, Minnesota, July 2006
- Poster/Presentation Chair, 13th IEEE International Conference on High Performance Computing (HiPC '06), Bangalore, India, December 2006
- Program Committee, Third Workshop on System Management Techniques, Services, and Processes at IPDPS 2007, Long Beach, California, March 2007
- Program Committee, Twelfth Asia-Pacific Computer Systems Architecture Conference, Seoul, Korea, August 2007
- Program Committee, 2007 IEEE International Symposium on Workload Characterization, Boston, Massachusetts, September 2007
- Program Committee, IEEE International Conference on Cluster Computing, Austin, Texas, September 2007
- Poster/Presentation Chair, 14th IEEE International Conference on High Performance Computing (HiPC '07), Goa, India, December 2007
- Program Committee, Thirteenth Asia-Pacific Computer Systems Architecture Conference, Taiwan, August 2008
- Program Committee, 37th International Conference on Parallel Processing (ICPP-08), Portland, Oregon, September 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, SC08: International Conference on High Performance Computing, Networking, Storage, and Analysis, Austin, Texas, November 2008
- Program Committee, 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS'08), Melbourne, Australia, December 2008
- Co-Chair, Student Research Symposium, 15th IEEE International Conference on High Performance Computing (HiPC '08), Bangalore, India, December 2008
- Program Committee, Workshop on Parallel and Distributed Systems: Testing, Analysis, and Debugging (PADTAD 2009), Chicago, Illinois, July 2009
- 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), and Kenin Coloma (2007) at Northwestern University
- Referee 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

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

Books

1. William Gropp, Ewing Lusk, and Rajeev Thakur, *Using MPI-2: Advanced Features of the Message-Passing Interface*, MIT Press, November 1999.

Journal Articles and Book Chapters

1. 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.
2. 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.
3. 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.
4. 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.
5. 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.
6. 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.
7. 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.
8. 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)
9. 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.
10. 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.
11. 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.
12. Rajeev Thakur, William Gropp, and Ewing Lusk, "Optimizing Noncontiguous Accesses in MPI-IO," *Parallel Computing*, (28)1:83–105, January 2002.
13. Rajeev Thakur and William Gropp, "Parallel I/O," *Sourcebook of Parallel Computing*, Morgan Kaufmann Publishers, Ch. 11, pp. 331–355, 2002.
14. 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.
15. 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.

16. 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.
17. 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.
18. 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.
19. 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.
20. 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.
21. 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.
22. William Gropp and Rajeev Thakur, "Thread Safety in an MPI Implementation: Requirements and Analysis," *Parallel Computing*, (33)9:595–604, September 2007.

Refereed Proceedings

1. 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.
2. 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**)
3. Rajeev Thakur, Alok Choudhary, and Geoffrey Fox, "Complete Exchange on a Wormhole Routed Mesh," in *Proceedings of MASCOTS '94*, January 1994, pp. 131–135.
4. 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.
5. 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.
6. 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.
7. 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.
8. 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.
9. 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.
10. 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.
 11. Phillip Dickens and Rajeev Thakur, “A Performance Study of Two-Phase I/O,” in *Proceedings of Euro-Par '98*, September 1998, pp. 959–965.
 12. 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.
 13. 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.
 14. 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.
 15. 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.
 16. 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.
 17. 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.
 18. 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)**
 19. 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.
 20. 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.
 21. 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.
 22. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 2840, Springer, September 2003, pp. 257–267.
 23. 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.

24. 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.
25. 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.
26. Weihang Jiang, Jiuxing Liu, Hyun-Wook Jin, Dhableswar 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.
27. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3241, Springer, September 2004, pp. 57–67.
28. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3241, Springer, September 2004, pp. 68–76.
29. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3241, Springer, September 2004, pp. 87–96.
30. 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.
31. 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.
32. 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.
33. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3666, Springer, September 2005, pp. 415–424.
34. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3666, Springer, September 2005, pp. 119–128.
35. 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)*, Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3666, Springer, September 2005, pp. 84–93.

36. 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.
37. 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.
38. 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.
39. 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.
40. 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. (**selected as outstanding paper**)
41. 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. (**selected as outstanding paper**)
42. 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.
43. 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.
44. 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.
45. 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.
46. Pavan Balaji, Sitha Bhagvat, Dhabaleswar 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.
47. 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.
48. 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.
49. 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. (**selected as outstanding paper**)

50. 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. **(selected as outstanding paper)**
51. 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.
52. Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby, Robert Palmer, Rajeev Thakur, and William Gropp, "Practical Model Checking Method for Verifying Correctness of MPI Programs," *Proceedings of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007 pp. 344–353.
53. 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," *Proceedings of SC07*, November 2007.
54. 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 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, February 2008. (short paper)
55. 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," *Proceedings of the International Supercomputing Conference (ISC'08)*, June 2008. **(Best Paper Award)**
56. 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.
57. 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.
58. 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.
59. 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.
60. 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.
61. 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. **(selected as outstanding paper)**
62. Pavan Balaji, Sitha Bhagvat, Rajeev Thakur, and Dhableswar Panda, "Sockets Direct Protocol for Hybrid Network Stacks: A Case Study with iWARP over 10G Ethernet," in *Proceedings of the International Conference on High Performance Computing (HiPC)*, December 2008.
63. 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 International Conference on High Performance Computing (HiPC)*, December 2008.

Proceedings

1. 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.
2. 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.
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. 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.
5. 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.

Technical Reports

1. 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.
2. 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.
3. 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.
4. 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.
5. 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.

Tutorials

1. “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)
2. “Tuning MPI Applications for Peak Performance,” at SC98: High Performance Networking and Computing, Orlando, Florida, November 1998. (with William Gropp and Ewing Lusk)
3. “Tuning MPI Applications for Peak Performance,” at SC99: High Performance Networking and Computing, Portland, Oregon, November 1999. (with William Gropp and Ewing Lusk)

4. "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)
5. "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)
6. "Parallel I/O in MPI-2," in tutorial on Porting to PACI Architectures, NPACI All-Hands Meeting, San Diego, California, March 2002.
7. "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)
8. "Using MPI for Peak Performance," SHARCNET Workshop, Univ. of Western Ontario, London, Ontario, Canada, October 2003. (with Brian Toonen)
9. "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)
10. "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)
11. "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)
12. "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)
13. "Introduction to Parallel I/O and MPI-IO," 11th Annual Computing Institute, San Diego Supercomputer Center, San Diego, California, July 2005.
14. "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)
15. "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)
16. "MPI on the Grid," at the 6th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2006), Singapore, May 2006. (with William Gropp)
17. "Introduction to MPI-IO and MPI-2," 12th Annual SDSC Summer Institute, San Diego Supercomputer Center, San Diego, California, July 2006.
18. "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)
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," at SC07: International Conference on High Performance Computing, Networking, Storage, and Analysis, Reno, Nevada, November 2007. (with William Gropp, Ewing Lusk, and Robert Ross)

21. "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)
22. "Advanced MPI," at the Workshop on Trends, Technologies and Collaborative Opportunities in High Performance and Grid Computing (WTTC2008), Bangkok, Thailand, June 2008.
23. "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)

Invited Talks

1. "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.
2. "MPI-IO: A Standard, Portable API for High-Performance Parallel I/O," Extreme Linux Workshop, Monterey, California, June 1999.
3. "MPICH on Clusters: Future Directions," Linux Supercluster Users Conference, Albuquerque, New Mexico, September 2000.
4. "MPI: Enabling Portable, High-Performance Parallel Computing," SHARCNET Workshop, University of Western Ontario, London, Ontario, Canada, October 2003.
5. "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.

Seminars and Colloquia

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

Other Presentations

1. "Scheduling Regular and Irregular Communication Patterns on the CM-5," Supercomputing '92, Minneapolis, Minnesota, November 1992.
2. "Runtime Array Redistribution in HPF Programs," Scalable High Performance Computing Conference, Knoxville, Tennessee, May 1994.
3. "Compiler and Runtime Support for Out-of-Core HPF Programs," International Conference on Supercomputing, Manchester, UK, July 1994.
4. "PASSION Runtime Library for Parallel I/O," Scalable Parallel Libraries Conference, Mississippi State University, Starkville, Mississippi, October 1994.
5. "I/O Characterization of a Portable Astrophysics Application," Scalable I/O Workshop at Supercomputing '95, San Diego, California, December 1995.
6. "An Abstract-Device Interface for Parallel I/O," Annual Meeting of the Center for Research on Parallel Computation, Argonne, Illinois, May 1996.
7. "Portable Implementation of MPI-IO," Panel discussion on MPI-IO at the MPI Developers Conference and Users Group Meeting, Notre Dame, Indiana, July 1996.
8. "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.
9. "An Abstract-Device Interface for Implementing Portable Parallel-I/O Interfaces," 6th Symposium on the Frontiers of Massively Parallel Computation, Annapolis, Maryland, October 1996.
10. "Implementing High-Level Interfaces on the SIO Low-Level API," Scalable I/O Roundtable at Supercomputing '96, Pittsburgh, Pennsylvania, November 1996.
11. "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.
12. "MPI-IO and Its Implementation," 2nd Midwest Workshop on High-Performance Systems, University of Illinois at Chicago, Chicago, May 1997.
13. "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.
14. "ROMIO: A High-Performance, Portable MPI-IO Implementation," Scalable I/O Initiative Sponsors Review meeting, Rice University, Houston, Texas, September 1997.
15. "A Case for Using MPI's Derived Datatypes to Improve I/O Performance," SC98: High Performance Networking and Computing, Orlando, Florida, November 1998.
16. "Achieving High-Performance with MPI-IO," SIO-ASCI Meeting, Livermore, California, January 1999.
17. "Data Sieving and Collective I/O in ROMIO," 7th Symposium on the Frontiers of Massively Parallel Computation, Annapolis, Maryland, February 1999.
18. "High-Performance Parallel I/O in MPI," MPICH-LANL meeting, Los Alamos, New Mexico, March 1999.
19. "On Implementing MPI-IO Portably and with High Performance," Sixth Workshop on I/O in Parallel and Distributed Systems, Atlanta, Georgia, May 1999.
20. "How Can File Systems Be Designed to Better Support MPI-IO?" Scalable Global Parallel File System Meeting, Santa Fe, New Mexico, September 1999.
21. "MPI I/O Perspectives," Panel Session on Scalable I/O, Cluster Computing Workshop, New Orleans, Louisiana, March 2000.

22. "OS, Message Passing, and Runtime Tools," Panel Session on OS, Message Passing, and Runtime Tools, Cluster Computing Workshop, New Orleans, Louisiana, March 2000.
23. "Chiba City: An Open Source Computer Science Testbed," Cluster Computing Workshop, New Orleans, Louisiana, March 2000.
24. "Parallel I/O Efforts at Argonne: A Brief Introduction," Microsoft Corporation, Redmond, Washington, April 2002.
25. "I/O on Clusters," 7th Workshop on Distributed Supercomputing (SOS7), Durango, Colorado, March 2003.
26. "Improving the Performance of Collective Operations in MPICH," 10th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2003), Venice, Italy, September 2003.
27. "Improving the Performance of MPI Derived Datatypes by Optimizing Memory-Access Cost," IEEE International Conference on Cluster Computing (Cluster 2003), Hong Kong, December 2003.
28. "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.
29. "I/O Consistency Semantics and Performance Implications," 9th Workshop on Distributed Supercomputing (SOS9), Davos, Switzerland, March 2005.
30. "MPI-IO Needs and Requirements," DARPA HPCS I/O Workshop, Washington, DC, July 2005.
31. "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.
32. "Implementing Byte-Range Locks Using MPI One-Sided Communication," 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005), Sorrento, Italy, September 2005.
33. "Opportunities in Parallel I/O for Scientific Data Management," Scientific Data Management SciDAC All-Hands Meeting, Raleigh, North Carolina, October 2005.
34. "Automatic Memory Optimizations for Improving MPI Derived Datatype Performance," 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006), Bonn, Germany, September 2006.
35. "Open Issues in MPI Implementation," 12th Asia-Pacific Computer Systems Architecture Conference (ACSAC 2007), Seoul, Korea, August 2007.
36. "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.
37. "The Future of Deep Computing System Software," IBM Deep Computing Institute External Advisory Board Meeting, San Diego, CA, February 2008.