Curriculum Vitae

Alok N. Choudhary
John G. Searle Professor
Electrical Engineering and Computer Science
McCormick School of Engineering and Applied Sciences
Director, Center for Ultra-Scale Computing and Information Security
Professor, Technology Industry Management and Marketing, Kellogg School of Management
Northwestern University
Technological Institute , 2145 Sheridan Road
Evanston, IL 60208-3118
Phone: (847)-467-4129, Fax: (847)-467-4144
Email: choudhar@eecs.northwestern.edu
http://www.eecs.northwestern.edu/~choudhar

Education Record

- **Ph.D.**, University of Illinois, Urbana-Champaign, Illinois, Computer Engineering (ECE Department), August, 1989.
- **M.S.**, University of Massachusetts, Amherst, Massachusetts, Computer Engineering (ECE Department), February, 1986.
- **B.E. (w/Honors)**, Birla Institute of Technology and Science, Pilani, India, Electrical and Electronics Engineering, June, 1982.

Work Experience and Positions Held

- Professor, EECS Department, McCormick School of Engineering, Northwestern University, 2000-Present.
- Director of Center for Ultra-Scale Computing and Information Security, 2000-Present.
- Professor, Kellogg School of Management, Technology Industry Management and Marketing, Kellogg School of Management Northwestern University, 2000-Present.
- Chair, Electrical Engineering and Computer Science Department, 2007 – 2011.
- Chair, Computer Engineering and Systems Division, 2004-2007.
- Co-founder and VP research and technology, Accelchip Inc., 2000-2002.
- Guest appointments with Lawrence Livermore Lab and Argonne National Labs.
- Associate Professor, ECE Department, Northwestern University, 9/1/96 - 8/31/00.
- Associate Professor, Kellogg School of Management, Northwestern University, 3/1/99-9/1/00
- Associate Professor, ECE and CIS Department, Syracuse University, 7/1/93 - 8/31/96.
- Assistant Professor, ECE and CIS Department, Syracuse University, 8/89 - 6/93.
- Visiting Faculty, University of Vienna, 6/96 -7/96.
- Visiting Faculty, Intel Corp, Beaverton, OR, Summer 95&96.
- Visiting Scientist, IBM, Thomas J. Watson Research Center, Yorktown Heights, NY, Summer ‘91.
- Visiting Scientist, IBM, Thomas J. Watson Research Center, Yorktown Heights, NY, Summer ‘88.
- Visiting Scientist, IBM, Thomas J. Watson Research Center, Yorktown Heights, NY, Summer ‘87.
- Research Assistant, Coordinated Science Laboratory, University of Illinois, Urbana, 1/86 to 8/89.
- Teaching Assistant, Department of Electrical and Computer Engineering, University of Illinois, Urbana , 8/87 to 1/88.
- Research Assistant, Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, 8/84 to 1/86.
- Teaching Assistant, Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, 1/84 to 1/86.
- Systems Analyst and Designer, Tata Consultancy Services, New Delhi, India (Computer Consultants), 6/82 to 1/84.
- Engineer (Trainee), Department of Electronics, Defense Cell, Government of India, New Delhi, India, 12/81 to 6/82.
- Teaching Assistant, Birla Institute of Technology and Science, Pilani, India, 7/81 to 12/81.

Honors and Awards

1. National science Foundation, Young Investigator Award, 1993 (1993-1999).
2. Fellow of the IEEE (Institute of Electrical and Electronics Engineers) (2005).
3. Fellow of the ACM (Association of Computing Machinery) (2009).
4. Fellow of the AAAS (American Academy of Advancement of Science) (2009).
5. Member, National Academy of Science Committee on the Impact of High-End Computing on Science. The committee formed on behalf of US government Office of Science and Technology to understand the impact and recommend priorities for the federal government for next 10 years.
6. Excellence in research, teaching and service award, McCormick School of Engineering, Northwestern University, 2006.

7. IBM Faculty Development Award, 1994, 2005
8. Intel Research Council Award, 1993-1997, 2003-2005.
9. NSF Research Initiation Award, 1991 (1991-1993).
10. IEEE Engineering Foundation award (1990).
11. Received highest performance award while working with Tata Consultancy Services (For the year 1982-1983).
12. National Scholarship (Government of India, 1976-1982).

Current Research Interests

- High-Performance Computing: Scalable Architectures, System Software, Runtime Systems; Advanced architectures, High-Performance Input-Output, Data Mining.
- Business Intelligence, data mining, large-scale text mining and social network analysis, and applications to science, scientific discoveries, business, computational biology, medicine and health care, CRM and marketing, high-performance data management and information security.

Graduate Student Supervision

(MS graduates and current students are given in Appendix D).

***Ph.D. Students (Graduated)*:**

1. Ravi Ponnusamny, Ph.D., 1994 (business owner)
2. Rajeev Thakur, Ph.D., 1995 (Research Scientist, Argonne National Labs)
3. Rajesh Bordawekar, Ph.D., 1996 (IBM T. J. Watson Research Center)
4. Chaitanya Tumuluri, Ph.D., 1996 (Silicon Graphics)
5. Senthil Krishnamoorthy, Ph.D., 1996 (Intel Corporation)
6. Divyesh Jadav, Ph.D., 1997 (IBM Almaden Research Center)
7. Ms. Meena Kandaswamy, Ph.D., 1998 (Intel)
8. Sang-hue Oh, Ph.D., 1998 (Macromedia, Seoul, Korea)
9. Ms. Chutimet Srinilta, Ph.D., Sept. 1998 (Faculty, Thailand)
10. Rakesh Krishnaiyer, Ph.D., Sept. 1998 (Intel)
11. Ms. Jaechun No, Ph.D., 1999 (Faculty, Korea)
12. Wei-keng Liao, Ph.D., 1999 (Research Professor, Northwestern University)
13. Mahmut Kandemir, Ph.D., 1999 (Professor, Penn State)
14. Sanjay Goil. Ph.D., 1999 (Intel)
15. Malay Haldar , PhD 2002, (Company in India)
16. Sachin More, PhD, 2000 (EMC)
17. Anshuman Nayak , PhD 2002 (Company in India)
18. Xioui Shen, PhD, 2001 (Motorola)
19. Steve Chiu, PhD, 2004 (faculty, ISU)
20. Ms. Ying Liu, PhD, 2005 (Faculty, Chinese Academy of Science)
21. Jay Pisharath, PhD, 2005 (Intel)
22. Joseph Zambreno, PhD, 2006 (Faculty, Iowa State)
23. Jianwei Li, PhD, 2006 (Wallstreet, Bloomberg)
24. Avery Ching, PhD, 2007 (Yahoo)
25. Kenin Coloma, PhD 2007 (Advertising.com)
26. Berkin Ozisikyilmaz, PhD 2010 (Postdoc)
27. Ms. Yan Gao, PhD 2010
28. Ms. Arifa Nisar, PhD 2010 (NSF CIF Fellow Postdoc at UC Santa Cruz)
29. Ramanathan Narayanan PhD 2011 (Goldman Sachs)
30. Abhishek Das, PhD 2011 (Intel)
31. Prabhat Kumar, PhD 2012 (NVIDIA)
32. Sanchit Misra, PhD 2012 (Intel)

Research Grants and Contracts (Current)

Current (PI on All grants unless specified otherwise): Approximate total ($7.5M)

1. DOE, " Scalable and Power Efficient Data Analytics for Hybrid Exascale Systems," $705K. (2010-2013).
2. DOE, " DAMSEL: A Data Model Storage Library for Exascale Science," $890K. (2010-2013).
3. NSF Expedition, " Understanding Climate Change: A Data Driven Approach," 900K. (2010-2015).
4. NSF, “SDCI HPC: Improvement: Parallel I/O Software Infrastructure for Petascale Systems,” 1.2 Million (2007-2011).
5. NSF, “Scalable I/O Middleware and File System Optimizations for High-Performance Computing,” 510K, (2007-2011).
6. DOE, “Active Storage for Analytics Capabilities and I/O runtime system for Petascale Systems,” 630K, (2008-2011)
7. Department of Energy (DOE), Enabling Technologies for Scientific Data Management, $1,000,000. (2006-2011).
8. NSF- HECURA, " An Application Driven I/O Optimization Approach for PetaScale Systems and Scientific Discoveries," $320K (2010-2013).
9. NSF, " DC: Medium: Collaborative Research: ELLF: Extensible Language and Library Frameworks for Scalable and Efficient Data-Intensive Applications," 410K (2009-2012).
10. NSF, “Advanced Compiler Optimizations and Programming Language Enhancements for Petascale I/O and Storage,” 278K, (2008-2011).
11. DOE, "An Optimizing Compiler for Petascale Leadership Class Systems," $525K. (2009-2012)
12. NSF, " CT-M: Hardware Containers for Software Components - Detection and Recovery at the Hardware/Software Interface," $135K (2009-2012).
13. NSF, Collaborative Research: CRI – Scalable Benchmarks, Software and Data for Data Mining, Analytics and Scientific Discoveries, $284K (2006-2009)
14. NSF, Collaborative Research: Ultra-scalable system software and tools for data-intensive Computing, $200K (2004-2009)
15. NSF, ITR: Hardware/Software Codesign for Software protection, $425,000 (2003-2009)
16. NSF/DARPA ST-HEC program, Ultra-scalable system software and tools for data-intensive computing, $445,000 (2004-2009).
17. NSF, “High-performance techniques, designs and implementation of software infrastructure for change detection and mining,” $500K (2005-2009)
18. NSF, “Collaborative Research: Scalable I/O Middleware and File System Optimizations for High-Performance Computing” 520K, (2006-2009).

Publications (2000 – present)

Published over 350 papers in journals and conferences.

(Journal papers published prior to 2000 are listed in Appendix A).

Refereed Journal Papers (2000 – present)

1. Jason Scott Mathias, Ankit Agrawal, Joe Feinglass, Andrew J Cooper, David William Baker, Alok Choudhary, “Development of a 5 year life expectancy index in older adults using predictive mining of electronic health record data”, Journal of the American Medical Informatics Association, (04/2013).
2. Daniel Honbo, Amit Pande, Alok Choudhary, Submission "FPGA Architecture for Pairwise Statistical Significance Estimation", International Journal of High Performance Systems Architecture, (03/2013).

3. Bharath Pattabiraman, Stefan Umbreit, Wei-keng Liao, Alok Choudhary, Vassiliki Kalogera, Gokhan Memik, Frederic A. Rasio, “A Parallel Monte Carlo Code for Simulating Collisional N-body Systems”, The Astrophysical Journal Supplement, Volume 204, Issue 2, article id. 15, 16 pp. (02/2013).
4. Zhengzhang Chen, William Hendrix, Hang Guan, Issac K. Tetteh, Alok Choudhary, Fredrick Semazzi, Nagiza F. Samatova, “Discovery of extreme events-related communities in contrasting groups of physical system networks”, Data Mining and Knowledge Discovery, (09/2012).
5. Alok Choudhary, William Hendrix, Kathy Lee, Diana Palsetia, Wei-keng Liao, “Social Media Evolution of The Egyptian Revolution”, Communications of the ACM (05/2012).
6. Rob Latham, Chris Daley, Wei-keng Liao, Kui Gao, Rob Ross, Anshu Dubey and Alok Choudhary, “A case study for scientific I/O: improving the FLASH astrophysics code”, Comput. Sci. Disc. 5 (2012) 015001, (03/2012).
7. Chetan Kumar, Alok Choudhary, “A top-down approach to classify enzyme functional classes and sub-classes using random forest”, EURASIP Journal on Bioinformatics and Systems Biology 2012, (02/2012).
8. Arifa Nisar, Wei-keng Liao and Alok Choudhary, **“**Delegation-based I/O Mechanism for High Performance Computing Systems”, IEEE Transactions on Parallel and Distributed Systems, Vol. 23, No. 2, February 2012.
9. Yuhong Zhang, Sanchit Misra, Daniel Honbo, Ankit Agrawal, Wei-keng Liao, Alok N. Choudhary, “Accelerating Pairwise Statistical Significance Estimation for Local Alignment by Harvesting GPU’s Power”, BMC Bioinformatics (ICCABS ’11).
10. William Hendrix, Andrea M Rocha, Kanchana Padmanabhan, Alok Choudhary, Kathleen Scott, James R Mihelcic and Nargiza F Samatova, “DENSE: Efficient and Prior Knowledge-driven Discovery of Phenotype-associated Protein Functional Modules”, BMC Systems Biology, 2011, 5:172 .
11. Alok N. Choudhary, Daniel Honbo, Prabhat Kumar, Berkin Ozisikyilmaz, Sanchit Misra and Gokhan Memik. “Accelerating Data Mining Workloads: current approaches and future challenges in system architecture design”, WIREs Data Mining and Knowledge Discovery, January/February 2011.
12. Jack Dongarra, Alok Choudhary, etc. “The International Exascale Software Project roadmap”, International Journal of High Performance Computing Applications, January 2011.
13. Alok N. Choudhary, Daniel Honbo, Prabhat Kumar, Berkin Ozisikyilmaz, Sanchit Misra and Gokhan Memik. Accelerating data mining workloads: current approaches and future challenges in system architecture design. WIREs Data Mining and Knowledge Discovery, Vol 1, Jan/Feb 2011.
14. Sanchit Misra, Ankit Agrawal, Wei-keng Liao, Alok Choudhary. “Anatomy of a Hash-based Long Read Sequence Mapping Algorithm for Next Generation DNA Sequencing”, *Bioinformatics 2010*; doi: 10.1093/bioinformatics/btq648.
15. Ying Liu, Jianwei Li, Wei-keng Liao, Alok Choudhary, and Yong Shi. “High Utility Itemset Mining. In the International Journal of Information Technology and Decision Making”, Volume 9, Issue 6, pp. 905--934, 2010.
16. Florin Isaila, Francisco Javier Garcia Blas, Jesús Carretero, Wei-keng Liao, Alok Choudhary**.** “A Scalable Message Passing Interface Implementation of an Ad-Hoc Parallel I/O System”. International Journal of High Performance Computing Applications Volume 24 Issue 2, May 2010.
17. Ioan Raicu, Ian Foster, Mike Wilde, Zhao Zhang, Kamil Iskra, Peter Beckman, Yong Zhao, Alex Szalay, Alok Choudhary, Philip Little, Christopher Moretti, Amitabh Chaudhary & Douglas Thain, “Middleware support for many-task computing”. *Cluster Computing*, Vol 3, April 2010.
18. Narayanan, R. and Misra, S. and Lin, S. and Choudhary.A , “Mining Protein Interactions from Text Using Convolution Kernels”, New Frontiers in Applied Data Mining, Lecture Notes in Computer Science, 2010, Volume 5669/2010, 118-129, DOI: 10.1007/978-3-642-14640-4_9.
19. Jacqueline Chen, Alok Choudhary, Bronis de Supinski, Matthew DeVries, Evatt Hawkes, Scott Klasky, Wei-Keng Liao, Kwan-Liu Ma, John Mellor-Crummy, Norbert Podhorski, Ramanan Sankaran, Sameer Shende, Chun Sang Yoo. “Terascale Direct Numerical Simulations of Turbulent Combustion Using S3D”. Journal of Computational Science & Discovery, Volume 2, Number 015001, 2009.

20. Alok Choudhary, Wei-keng Liao, Kui Gao, Arifa Nisar, Robert Ross, Rajeev Thakur, and Robert Latham. “Scalable I/O and Analytics”. In the Journal of Physics: Conference Series, Volume 180, No. 012048 (10pp), August 2009.
21. Abhishek Das, David Nguyen, Joseph Zambreno, Gokhan Memik, Alok Choudhary. “An FPGA-based Network Intrusion Detection Architecture”, *IEEE Transactions on Information Forensics and Security* (**TIFS**), Volume 3, Issue 1, Mar. 2008.
22. W.K. Liao, K. Coloma, A. Choudhary, and L. Ward “Cooperative Client-side File Caching for MPI Applications”, *International Journal of High Performance Computing Applications*, Vol. 21, No. 2, 144-154, 2007.
23. S. W. Son, G. Chen, O. Ozturk, M. Kandemir, A. Choudhary. “Compiler-directed Energy Optimization for Parallel Disk Based Systems”, *IEEE Transactions on Parallel and Distributed Systems (TPDS)* 18(9):1241-1257, September 2007.
24. A. Das, S. Ozdemir, G. Memik, J. Zambreno, and A. Choudhary. “Microarchitectures for Managing Chip Revenues under Process Variations.” in *IEEE Computer Architecture Letters,* Volume 6, June 2007.
25. Gokhan Memik, Mahmut T. Kandemir, Wei-Keng Liao, Alok Choudhary. Multi-Collective I/O: A technique for exploiting inter-file access patterns *ACM Transactions on Storage* (**ToS**), Volume 2, Issue 3, pp. 349-369, August 2006.
26. Steve Chiu, Wei-keng Liao, and Alok Choudhary. Distributed Smart Disks for I/O-Intensive Workloads on Switched Interconnects. In the *Journal of Future Generation Computer Systems*, 22(5): 643-656, Apr., 2006.
27. J. Zambreno, D. Honbo, A. Choudhary, R. Simha, and B. Narahari. High-performance software protection using reconfigurable architectures. Proceedings of the IEEE, vol. 94, no. 2, pages 1-13, February 2006.
28. W. Liao, A. Choudhary, L Ward, N. Pandit. “Cooperative Write-behind Data Buffering for MPI I/O,” Lecture Notes in Computer Science, 2005, Volume 3666/2005, 102-109, DOI: 10.1007/11557265_17
29. Joseph Zambreno, Alok Choudhary, Rahul Simha, Bhagirath Narahari, and Nasir Memon. SAFE-OPS: An Approach to Embedded Software Security. *ACM Transactions on Embedded Computing Systems (TECS)*, 4(1):189–210, February 2005.
30. Wei-keng Liao, Alok Choudhary, Donald Weiner, and Pramod Varshney. Performance Evaluation of a Parallel Pipeline Computational Model for Space-Time Adaptive Processing. *Journal of Supercomputing*, 31(2):137–160, February 2005.
31. M. Kandemir, I. Kadayif, A. Choudhary, J. Ramanujam, and I. Kolcu., “Compiler-directed scratch pad memory optimization for embedded multiprocessors,” IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 12, No. 3, March 2004.
32. Chiu, S. C., W.-K. Liao, A. Choudhary, M. Kandemir. Processor-embedded Distributed Smart Disks for I/O-intensive Workloads: Architectures, Performance Models, and Evaluation. Journal of Parallel and Distributed Computing (JPDC) 64(3):427-446, March 2004.
33. X. Shen, W. Liao, A. Choudhary, G. Memik, and M. Kandemir. "A High Performance Application Data Environment for Large-Scale Scientific Computations", in *IEEE Transaction on Parallel and Distributed System* , Volume 14, Number 12, pp. 1262-1274, December 2003.
34. X. Shen, A. Choudhary, C. Matarazzo and P. Sinha, “A Distributed Multi-Storage Architecture and I/O Performance Prediction for Scientific Computing,” Cluster Computing 6, pp 189-200, 2003.
35. Nan Jiang, Jayaprakash Pisharath and Alok Choudhary. “Characterizing and Improving Energy-Delay Tradeoffs in Heterogeneous Communication Systems,” In Proc. of the IEEE International Symposium on Signals, Circuits and Systems (SCS), July 2003, IEEE Press.
36. Jayaprakash Pisharath, Nan Jiang, and Alok Choudhary. “Evaluation of Application-Aware Heterogeneous Embedded Systems for Performance and Energy Consumption,” In Proc. of the 9th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), May 2003, IEEE Press.
37. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “Reducing False Sharing and Improving Spatial Locality in a Unified Compilation Framework,” IEEE Transactions on Parallel and Distributed Systems (TPDS), April 2003.
38. Jaechun No, Rajeev Thakur, Alok Choudhary, “High -performance scientific data management system,” Journal of Parallel and Distributed Computing, volumn 63, issue 4, pages 434-447(Apr 2003).

39. M. Kandemir, J. Ramanujam, and A. Choudhary, “Reducing False Sharing and Improving Spatial Locality”, *IEEE Transactions on Parallel and Distributed Systems, Vol. 14, No. 4, April 2003.*
40. M. Kandemir, A. Choudhary, and J. Ramanujam,” An I/O-conscious tiling strategy for disk-resident data sets,” Journal of Supercomputing, Vol 21, Number 3, March 2002, pp. 257-284.
41. Jaechun No, Sung-soon Park, Jesus Carretero Perez and Alok Choudhary, “Design and Implementation of a Parallel I/O Runtime System for Irregular Applications,” Journal of Parallel and Distributed Computing, volumn 62, issue 2, pages 173-326 (Feb 2002).
42. M. Kandaswamy, M. Kandemir, A. Choudhary, and D. Bernholdt, “An Experimental Evaluation of I/O Optimizations on Different Applications,” *IEEE Transactions on Parallel and Distributed Systems*, Vol. 13, No. 7, July 2002, pp. 728-744.
43. J. No, S. Park, J. Carretero Perez, and A. Choudhary, “Design and Implementation of a Parallel I/O Runtime System for Irregular Applications,” Journal of Parallel and Distributed Computing, Vol. 62, No. 2, February 2002, pp. 193-220.
44. M. Kandemir, J. Ramanujam, A. Choudhary, and P. Banerjee, “A Layout-Conscious Iteration Space Transformation Technique”, IEEE Transactions on Computers, Vol. 50, No. 12, December 2001.
45. G. Memik, M. Kandemir, and A. Choudhary, “Design and Evaluation of Smart Disk Cluster for DSS Commercial Workloads,” *Journal of Parallel and Distributed Computing* (Special Issue on Cluster and Network-Based Computing), Nov 2001.
46. M. Kandemir, P. Banerjee, A. Choudhary, J. Ramanujam, and E. Ayguade, “Static and Dynamic Locality Optimizations Using Integer Linear Programming”, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 12, No. 9, pp. 922-941, September 2001.
47. N. Shenoy, A. Choudhary, and P. Banerjee, “An Algorithm for Synthesis of Large Time-constrained Heterogeneous Adaptive Systems”, *ACM Transactions on the Design Automation of Electronic Systems*, Vol. 6, No. 2, April 2001.
48. S. Goil and A. Choudhary, “PARSIMONY: An Infrastructure for Parallel Multidimensional Analysis and Data Mining,” *Journal of Parallel and Distributed Computing* (Special Issue on High Performance Data Mining), Vol. 61, No. 3, March 2001.
49. M. Kandemir, A. Choudhary, P. Banerjee, J. Ramanujam, and N. Shenoy, “Minimizing Data and Synchronization Costs in One-Way Communication,” *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Vol. 11, No. 12, December 2000, pp. 1232-1251.
50. 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: the Journal of Networks, Software Tools and Applications*, Vol. 3, No. 1, pp. 45-60, 2000.
51. M. Kandemir, J. Ramanujam, and A. Choudhary, “Compiler Algorithms for Optimizing Locality and Parallelism on Shared and Distributed-Memory Machines”, *Journal of Parallel and Distributed Computing*, 60, pp. 924-965, 2000.
52. M. Kandemir, A. Choudhary, J. Ramanujam, and M. Kandaswamy, “A Unified Framework for Optimizing Locality, Parallelism, and Communication in Out-of-Core Computations”, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 11, No. 7, pp. 648-667, July 2000.
53. M. Kandemir, J. Ramanujam, and A. Choudhary, “Compiler Algorithms for Optimizing Locality and Parallelism on Shared and Distributed Memory Machines”, *Journal of Parallel and Distributed Computing*, 60, pp. 924-965, 2000.
54. A. Choudhary , W. Liao, D. Weiner, P. Varshney, R. Linderman, M. Linderman, and R. Brown, “Design, Implementation and Evaluation of Parallel Pipelined STAP on Parallel Computers”, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 36, Issue 2, April 2000, pp. 528-548.

Book: A. N. Choudhary and J. H. Patel, *Parallel Architectures and Parallel Algorithms for Integrated Vision Systems*, Kluwer Academic Publisher, Boston, MA, 1990.

Selected Patents

1. Hastings; Andrew B. (Eagan, MN), Rang; Anton B. (Houlton, WI), Choudhary; Alok N. (Chicago, IL), " Method and system for file access using a shared memory," U. S. Patent No. 7,757,049.
2. Hastings; Andrew B. (Eagan, MN), Choudhary; Alok N. (Chicago, IL), Coverston; Harriet G. (New Brighton, MN). " Method and system for collective file access using an mmap (memory-mapped file," U.S. Patent 7,647,471

3. Banerjee; Prithviraj (Glenview, IL), Choudhary; Alok (Chicago, IL), Haldar; Malay (Evanston, IL), Nayak; Anshuman (Evanston, IL), "Method and apparatus for automatically generating hardware from algorithms described in MATLAB," U. S. Patent **7,000,213.**
4. **Nayak; Anshuman** (Schaumburg, IL)**, Haldar; Malay** (Schaumburg, IL)**, *Choudhary; Alok*** (Chicago, IL)**, Saxena; Vikram** (Schaumburg, IL)**, Banerjee; Prithviraj** (Glenview, IL), "System for architecture and resource specification and methods to compile the specification onto hardware," U. S. Patent **7,376,939.**

Selected Book Chapters

1. Ankit Agrawal, Alok Choudhary, and Xiaoqiu Huang, "Sequence-Specific Sequence Comparison Using Pairwise Statistical Significance", In Software Tools and Algorithms for Biological Systems, Springer (in book series, Advances in Experimental Medicine and Biology, AEMB), 2010.
2. Robert Ross, Alok Choudhary, Garth Gibson, and Wei-keng Liao. Book Chapter 2: Parallel Data Storage and Access. In Scientific Data Management: Challenges, Technology, and Deployment, Chapman & Hall/CRC Computational Science Series, CRC Press, December 2009.
3. Avery Ching, Kenin Coloma, Arifa Nisar, Wei-keng Liao, and Alok Choudhary. Distributed File Systems. Wiley Encyclopedia of Computer Science and Engineering. Editor Benjamin W. Wah, John Wiley & Sons, Inc., January 2009.
4. Sanchit Misra, Ramanathan Narayanan, Daniel Honbo and Alok Choudhary. (Book Chapter) High Performance Distributed Data Mining. *Next Generation of Data Mining*, CRC Press, 2008.
5. Avery Ching, Kenin Coloma, Jianwei Li, and Alok Choudhary. High-performance techniques for parallel I/O. In Handbook of Parallel Computing: Models, Algorithms, and Applications. CRC Press, 2007.
6. Jianwei Li and Alok Choudhary and Nan Jiang and Wei-keng Liao, Mining Frequent Patterns by Differential Refinement of Clustered Bitmaps, Proceedings of the {SIAM} Int'l Conf. on Data Mining, April 2006.
7. Jianwei Li, Ying Liu, Wei-keng Liao, Alok Choudhary, Parallel Data Mining Algorithms for Association Rules and Clustering, Handbook of Parallel Computing: Models, Algorithms and Applications. Sanguthevar Rajasekaran and John Reif, ed., CRC Press, 2006.
8. Wei-keng Liao, Alok Choudhary, Kenin Coloma, Lee Ward, Eric Russell, and Neil Pundit. MPI Atomicity and Concurrent Overlapping I/O. Book chapter 10 in High Performance Computing: Paradigm and Infrastructure, pp. 203 -- 218, John Wiley & Sons Inc. November 2005.
9. I. Kadayif, M. Kandemir, A. Choudhary, M. Karakoy, N. Vijaykrishnan, and M. J. Irwin, “Compiler-directed communication energy optimizations for microsensor networks,” In Frontiers in Distributed Sensor Networks, edited by R. Brooks and R. Iyengar, CRC Press, 2003.
10. G. Memik, M. Kandemir, A. Choudhary, “Hardware/software techniques for improving cache performance in embedded systems”, Embedded Software for SoC, Edited by Jerraya, Yoo, When, Verkest, Kluwer, June 2003.
11. A. Choudhary and M. Kandemir, “Compilation for distributed-memory architectures”, The Compiler Design Handbook: Optimizations and Machine Code Generation, CRC Press, 2002.
12. H. Nagesh, S. Goil, and A. Choudhary, “Parallel Algorithms for Clustering High-dimensional Large-Scale Datasets,” in *Data Mining for Scientific and Engineering Applications*, Editors: R. Grossman, C. Kamath, P. Kegelmeyer, V. Kumar, and R. Namburu, Kluwer Academic Publishers, 2001.
13. Jones, D. Bagchi, S. Pal, A. Choudhary, and P. Banerjee, “PACT HDL: A C Compiler Targeting ASICs and FPGAs with Power and Performance Optimizations,” in *Power Aware Computing*, Editors: Rami Melhem and Bob Graybill, Kluwer Academic Publishers, 2001.
14. M. Kandemir and A. Choudhary, “I/O Programming Paradigms: Past and Future,” in *High Performance Mass Storage and Parallel I/O: Theory and Practice*, Editors: H. Jin, T. Cortes, and R. Buyya, IEEE Press, 2000.
15. M. Kandemir and A. Choudhary, “I/O Programming Paradigms: Past and Future”, *High Performance Mass Storage and Parallel I/O: Theory and Practice*, Editors: H. Jin, T. Cortes, and R. Buyya, IEEE Press, 2000.
16. M. Kandemir, J. Ramanujam, A. Choudhary, and P. Banerjee, “An Iteration Space Transformation Algorithm Based on Explicit Data Layout Representation for Optimizing Locality”, in *Languages and Compilers for Parallel Computing*, Editors: S. Chatterjee et al., Lecture Notes in Computer Science, Springer-Verlag, 1999.
17. M. Kandemir, J. Ramanujam, and A. Choudhary, “Improving Locality in Out-of-Core Computations Using Data Layout Transformations”, in *Languages, Compilers, and Run-Time Systems for*

Scalable Computers, Editors: O'Hallaron et al., Lecture Notes in Computer Science, Volume 1511, pp. 359-356, Springer-Verlag, 1998.

18. R. Thakur, A. Choudhary and R. Bordawekar, "Runtime Support for Out-of-Core Parallel Programs", *Input/Output in Parallel and Distributed Computer Systems*, Editors: R. Jain, J. Browne and J. Wirth, Kluwer Academic Publishers, 1996.
19. R. Bordawekar and A. Choudhary, "Issues in Compiling I/O Intensive Problems", *Input/Output in Parallel and Distributed Computer Systems*, Editors: R. Jain, J. Werth, and J. Browne, Kluwer Academic Publishers, 1996.
20. A. N. Choudhary and S. Ranka, "A Perspective on Parallel Processing in Computer Vision and Image Understanding", in *Parallel Processing for Artificial Intelligence*, pp.3-20, Editors: N. Kanal, V. Kumar, H. Kitano and C. Suttner, Elsevier/North Holland, 1994.

21. Z. Bozkus, A. Choudhary, G. Fox, T. Haupt and S. Ranka, "Compiling HPF for Distributed Memory MIMD Computers, *The Interaction of Compilation Technology and Computer Architecture"*, Editors: David Lilja and Peter Bird, Kluwer Academic Publishers, 1993.
22. A. N. Choudhary, J. H. Patel and N. Ahuja, "Architecture and Performance Evaluation of NETRA", in *Parallel Architectures and Algorithms for Image Understanding*, pp. 251-278, Editor: Prasanna Kumar, Academic Press, 1991.

Refereed Conference Papers (Year 2000 – present)

(Papers published prior to 2000 are listed in Appendix B).

1. Zhengzhang Chen, Yusheng Xie, Yu Cheng, Kunpeng Zhang, Ankit Agrawal, Wei-keng Liao, Nagiza Samatova, Alok Choudhary, "Forecast Oriented Classification of Spatio-Temporal Extreme Events" to appear in the proceedings of the 23rd International Joint Conference on Artificial Intelligence to be held in Beijing, China, on August 3-9, 2013.
2. Yusheng Xie, Zhengzhang Chen, Kunpeng Zhang, Yu Cheng, Ankit Agrawal, Wei-keng Liao, Alok Choudhary, "Detecting and Tracking Disease Outbreaks in Real-time through Social Media" to appear in the proceedings of the 23rd International Joint Conference on Artificial Intelligence to be held in Beijing, China, on August 3-9, 2013.
3. Bharath Pattabiraman, Stefan Umbreit, Wei-keng Liao, Frederic A. Rasio, Vassiliki Kalogera, Gokhan Memik, Alok Choudhary, "GPU-accelerated Monte Carlo Simulations of Dense Stellar System", in the proceedings of Innovative Parallel Computing, San Jose, CA, May 13-14 2012.
4. Zhengzhang Chen, John Jenkins, Jinfeng Rao, Alok Choudhary, Fredrick Semazzi, Anatoli V. Melechko, Vipin Kumar, Nagiza F. Samatova, "Automatic Detection and Correction of Multi-class Classification Errors Using System Whole-part Relationships," *To appear in the 13th SIAM International Conference on Data Mining* (*SDM 2013*), Austin, Texas, USA, May 2-4, 2013.
5. Yusheng Xie, Zhengzhang Chen, Kunpeng Zhang, Md. Mostofa Ali Patwary, Yu Cheng, Haotian Liu, Ankit Agrawal, Alok Choudhary, "Graphical Modeling of Macro Behavioral Targeting in Social Networks," *To appear in the 13th SIAM International Conference on Data Mining* (*SDM 2013*), Austin, Texas, USA, May 2-4, 2013.
6. William Hendrix, Md. Mostofa Ali Patwary, Ankit Agrawal, Wei-keng Liao, Alok Choudhary, "Parallel Hierarchical Clustering on Shared Memory Platforms", has been accepted for presentation and publication in the proceedings of 19th Annual International Conference on High Performance Computing, December 18-21, 2012 in Pune, India.
7. Seong Jo Kim, Seung Woo Son, Wei-keng Liao, Mahmut Kandemir, Rajeev Thakur, Alok Choudhary, " IOPin: Runtime Profiling and Optimization of Parallel I/O in HPC Systems", Parallel Data Storage Workshop (PDSW '12) of (Supercomputing, SC'12), November 10-16, 2012 in Salt Lake City, Utah.
8. Md. Mostofa Ali Patwary, Diana Palsetia, Ankit Agrawal, Wei-keng Liao, Fredrik Manne, and Alok Choudhary. "A New Scalable Parallel DBSCAN Algorithm Using the Disjoint-Set Data Structure", International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing, SC'12), November 10-16, 2012 in Salt Lake City, Utah.
9. Yu Cheng, Yusheng Xie, Kunpeng Zhang, Ankit Agrawal, Dan Honbo, Alok Choudhary. "On Active Learning in Hierarchical Classification". In *ACM CIKM*, October 2012.

10. Yu Cheng, Yusheng Xie, Kunpeng Zhang, Ankit Agrawal, Dan Honbo, Alok Choudhary, "CluChunk: Clustering Large Scale User-generated Content Incorporating Chunklet Information", (ACM SIGKDD 2012).
11. Yu Cheng, Yusheng Xie, Kunpeng Zhang, Ankit Agrawal, Dan Honbo, Alok Choudhary, "How Online Content is Received by Users in Social Media: A Case Study on Facebook.com Posts", (ACM SIGKDD 2012).
12. Diana Palsetia, Md. Mostofa Ali Patwary, Kunpeng Zhang, Kathy Lee, Christopher Moran, Yves Xie, Daniel Honbo, Ankit Agrawal, Wei-keng Liao, Alok Choudhary, "User-Interest based Community Extraction in Social Networks", (ACM SIGKDD 2012).
13. Yusheng Xie, Yu Cheng, Daniel Honbo, Kunpeng Zhang, Ankit Agrawal, Alok Choudhary, "Crowdsourcing Recommendations From Social Sentiment", (ACM SIGKDD 2012).
14. Yusheng Xie, Daniel Honbo, Kunpeng Zhang, Yu Cheng, Ankit Agrawal, Alok Choudhary, "VOXSUP: A Social Engagement Framework", (ACM SIGKDD 2012).
15. Kunpeng Zhang, Yu Cheng, Yusheng Xie, Ankit Agrawal, Alok Choudhary, "Sentiment Identification by Incorporating Syntactics, Semantics and Context Information", (SIGIR 2012).
16. A. Das, M. Schuchardt, N. Hardavellas, G. Memik, and A. Choudhary. “Dynamic Directories: Reducing On-Chip Interconnect Power in Multicores.” *(DATE 2012)*, Dresden, Germany.
17. Hongyu Gao, Yan Chen, Kathy Lee, Diana Palsetia and Alok Choudhary, “Towards Online Spam Filtering in Social Networks”, (NDSS 2012).
18. William Hendrix, Isaac Tetteh, Ankit Agrawal, Fredrick Semazzi, Wei-keng Liao, and Alok Choudhary,“Community Dynamics and Analysis of Decadal Trends in Climate Data”, (ICDM Climate 2011).
19. Ankit Agrawal and Alok Choudhary, “Identifying HotSpots in Lung Cancer Data Using Association Rule Mining”, (ICDM BioDM 2011).
20. Yu Cheng, Kunpeng Zhang, Yusheng Xie, Ankit Agrawal, Wei-keng Liao, and Alok Choudhary “Learning to Group Web Text Incorporating Prior Information”, (OEDM 2011).
21. Lalith Polepeddi, Ankit Agrawal, and Alok Choudhary, “Poll: A Citation-Text-Based System for Identifying High-Impact Contributions of an Article”, (ICDM DaMNet 2011).
22. Kunpeng Zhang, Yu Cheng, Yusheng Xie, Ankit Agrawal, Diana Palsetia, Kathy Lee, and Alok Choudhary, “SES: Sentiment Elicitation System for Social Media Data”, (ICDM-SENTIRE 2011).
23. Kathy Lee, Diana Palsetia, Mostofa Patwary, Ankit Agrawal, Alok Choudhary, and Ramanathan Narayanan, “Twitter Trending Topic Classification”, (OEDM 2011).
24. Kui Gao, Chen Jin, Alok Choudhary, Wei-keng Liao “Supporting Computational Data Model Representation with High-performance I/O in Parallel netCDF” (HiPC 2011).
25. Kunpeng Zhang, Yu Cheng, Wei-keng Liao, Alok Choudhary, “Mining Millions of Reviews: A Technique to Rank Products Based on Importance of Reviews ” (ICEC 2011), (Received Honourable Mention Award).
26. Chen Jin, Saba Sehrish, Wei-keng Liao, Alok Choudhary, Karen Schuchardt, “Improving the Average Response Time in Collective I/O” (EuroMPI 2011).
27. A. Agrawal, S. Misra, R. Narayanan, L Polepeddi, A. Choudhary, “A Lung Cancer Outcome Calculator Using Ensemble Data Mining on SEER Data” (BIOKDD11), 2011.
28. Isha Arkatkar, Alok Choudhary, John Jenkins and Nagiza Samatova. “Lessons Learned from Exploring the Backtracking Paradigm on the {GPU}”, 17th International European Conference on Parallel and Distributed Computing, August 2011.
29. Huseyin Sencan, Zhengzhang Chen, William Hendrix, Tatdow Pansombut, Frederick Semazzi, Alok Choudhary, Vipin Kumar, Nagiza Samatova and Anatoli Melechko,” Classification of Emerging Extreme Event Tracks in Multi-Variate Spatio-Temporal Physical Systems Using Dynamic Network Structures: Application to Hurricane Track Prediction”, (OJCAI 2011).
30. William Hendrix, Tatiana Karpinets, Byung-Hoon Park, Eric Schendel, Alok Choudhary, Nagiza F. Samatova. “Sensitive and Specific Identification of Protein Complexes in “Perturbed” Protein Interaction Networks from Noisy Pull-Down Data”. (IPDPS 2011)
31. Prabhat Kumar, Berkin Ozisikyilmaz, Wei-keng Liao, Gokhan Memik, and Alok Choudhary. “High Performance Data Mining Using R on Heterogeneous Platforms”. (MTAAP 2011).
32. Yuhong Zhang, Sanchit Misra, Daniel Honbo, Ankit Agrawal, Wei-keng Liao, Alok N. Choudhary, “Efficient pairwise statistical significance estimation for local sequence alignment using GPU”, (ICCABS 2011).

33. K. Zhang, R. Narayanan, A. Choudhary, “Voice of the Customers: Text Mining Online Customer Reviews for Product Feature-based Ranking”. In Workshop on Online Social Networks (WOSN '10), 2010.
34. G. Trajcevski, A. Choudhary, O. Wolfson, L. Ye and G. Li. "Uncertain Range Queries for Necklaces". 11th International Conference on Mobile Data Management, May 2010. **(Received Best Paper Award)**
35. G. Trajcevski, A. N. Choudhary and P. Scheuermann. “Sensing, Triggers and Mobile (Meta)Data”, M-PDMST Workshop (with the IEEE-MDM Conference), May 2010.
36. A.Choudhary, R.Narayanan, K.Zhang, “High-End Analytics and Data mining for Sustainable Competitive Advantage. In Proceedings of International Conference on Information Systems, Technology and Management” (ICISTM 2010), Bangkok.
37. S. W. Son, S. Lang, P. Carns, R. Ross, R. Thakur, B. Ozisikyilmaz, P. Kumar, W. Liao, A. Choudhary. Enabling Active Storage on Parallel I/O Software Stacks. In Proc. of the Symposium on Massive Storage Systems and Technologies. May 2010.
38. Ankit Agrawal, Alok Choudhary, and Xiaoqiu Huang, "Derived Distribution Points Heuristic for Fast Pairwise Statistical Significance Estimation", In Proc. ACM International Conference on Bioinformatics and Computational Biology (ACM-BCB) 2010, pp. 312-321.
39. Ankit Agrawal, Sanchit Misra, Daniel Honbo and Alok Choudhary, "MPIPairwiseStatSig: Parallel Pairwise Statistical Significance Estimation of Local Sequence Alignment", HPDC ECMLS 2010, pp. 470-476.
40. Sanchit Misra, Ramanathan Narayanan, Wei-keng Liao, Alok Choudhary and Simon Lin. “pFANGS: Parallel High Speed Sequence Mapping for Next Generation 454-Roche Sequencing Reads”**.** *Ninth IEEE International Workshop on High Performance Computational Biology (IPDPS 2010)*, April, 2010, Atlanta, GA
41. Sanchit Misra, Ramanathan Narayanan, Simon Lin and Alok Choudhary. FANGS: High Speed Sequence Mapping for Next Generation Sequencing Reads. *ACM Symposium on Applied Computing (ACM SAC)*, pp. 1539-1546, March 22-26, 2010, Sierre, Switzerland
42. Daniel Honbo, Ankit Agrawal and Alok Choudhary, "Efficient Pairwise Statistical Significance Estimation using FPGAs", BIOCOMP 2010, pp. 571-577.
43. Ankit Agrawal, Alok Choudhary and Xiaoqiu Huang, "Non-Conservative Pairwise Statistical Significance of Local Sequence Alignment Using Position-Specific Substitution Matrices", International Conference on Bioinformatics and Computational Biology (BIOCOMP) 2010, accepted.
44. Sanchit Misra, Ramanathan Narayanan, Wei-keng Liao, Alok Choudhary and Simon Lin. “pFANGS: Parallel High Speed Sequence Mapping for Next Generation 454-Roche Sequencing Reads”. 24th IEEE International Parallel & Distributed Processing Symposium, Workshops and Phd Forum (IPDPSW), pp 1-8, April, 2010, Atlanta, GA.
45. S. Ozdemir, Y. Pan, A. Das, G. Memik, G. Loh, A. Choudhary, " Quantifying and Coping with Parametric Variations in 3D-Stacked Microarchitectures," Design, Automation and Test in Europe (DATE), Dresden, Germany, March 2010.
46. A. Das, G. Memik, J. Zambreno, A. Choudhary, " Detecting/Preventing Information Leakage on the Memory Bus due to Malicious Hardware," of Design, Automation and Test in Europe (DATE), Dresden, Germany, March 2010.
47. Yan Gao, Alok Choudhary and Gang Hua, “A nonnegative sparsity induced similarity measure with application to cluster analysis of spam images”, IEEE International Conference on Acoustics, Speech and Signal Processing, 2010.
48. Yan Gao, Ming Yang and Alok Choudhary, “Semi supervised image spam hunter: a regularized discriminant EM approach”, International Conference on Advanced Data Mining And Applications, 2009.
49. Yan Gao and Alok Choudhary. “Active learning image spam hunter”. International Symposium on visual Computing, 2009.
50. N. Nakka, A. Choudhary, W. K. Liao, L. Ward, R. Klundt, M. I. Weston. "Detailed Analysis of I/O traces for large scale applications", *International Conference on High Performance Computing, HiPC 2009*, Cochin, India, Dec 2009.
51. G. Trajcevski and A. Choudhary. "Data-Aware Control => Efficient Traffic Management", Next Generation Data Mining Summit, October 2009.

52. Kui Gao, Wei-keng Liao, Arifa Nisar, Alok Choudhary, Robert Ross, and Robert Latham. “Using Subfiling to Improve Programming Flexibility and Performance of Parallel Shared-file I/O”. *International Conference on Parallel Processing*, Vienna, Austria, September 2009.
53. Kui Gao, Wei-keng Liao, Alok Choudhary, Robert Ross, and Robert Latham. “Combining I/O Operations for Multiple Array Variables in Parallel NetCDF”. In the Proceedings of the Workshop on Interfaces and Architectures for Scientific Data Storage, held in conjunction with the IEEE Cluster Conference, New Orleans, Louisiana, September 2009.
54. Ramanathan Narayanan, Bing Liu and Alok Choudhary. "Sentiment Analysis of Conditional Sentences." *Conference on Empirical Methods in Natural Language Processing (EMNLP-09)*. August 6-7, 2009. Singapore.
55. N. Nakka, A. Choudhary. "Failure data-driven selective node-level duplication to improve MTTF in High Performance Computing Systems," International Symposium on High Performance Computing, HPCS 2009, Kingston, Ontario, Canada, August 2009.
56. Alok Choudhary, Wei-keng Liao, Kui Gao, Arifa Nisar, Robert Ross, Rajeev Thakur, and Robert Latham. “Scalable I/O and Analytics”. In the Journal of Physics: Conference Series, Volume 180, No. 012048 (10pp), August 2009. (Proceedings of SciDAC conference, 14-18 June 2009, San Diego, California, USA).
57. Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, A. Choudhary. “Firefly: Illuminating Future Network-on-Chip with Nanophotonics “, *International Symposium on Computer Architecture* (**ISCA**), Austin, TX, June 2009.
58. P. Kumar, Y. Pan, J. Kim, G. Memik, A. Choudhary . Exploring Concentration and Channel Slicing in On-Chip Network Router, *International Symposium on Networks-on-Chip* (**NOCS**), San Diego, CA, May 2009.
59. Ramanathan Narayanan, Sanchit Misra, Simon Lin and Alok Choudhary. Mining Protein Interactions from Text using Convolution Kernels. In AIBDM 09, held in conjunction with Pacific Asia Conference on Knowledge Discovery and Data mining (PAKDD 2009), Bangkok, April 27-30, 2009.
60. Arifa Nisar, Waseem Ahmad, Wei-keng Liao, and Alok Choudhary. High Performance Parallel/Distributed Biclustering Using Barycenter Heuristic. In the Proceedings of SIAM International Conference on Data Mining, Sparks, Nevada, April 2009.
61. Y. Zhang, B. Ozisikyilmaz, G. Memik, J. Kim, A. Choudhary . Analyzing the Impact of On-Chip Network Traffic on Program Phases for CMPs. *International Symposium on Performance Analysis of Systems and Software* (**ISPASS**), Boston, MA, April 2009.
62. Florin Isaila, Francisco Javier Garcia Blas, Jesus Carretero, Wei-keng Liao, and Alok Choudhary. AHPIOS: An MPI-based Ad-hoc Parallel I/O System. In the Proceedings of 14th Intl Conference on Parallel and Distributed Systems, Melbourne, Victoria, Australia, December 2008.
63. Arifa Nisar, Wei-keng Liao, and Alok Choudhary. Scaling Parallel I/O Performance through I/O Delegate and Caching System In the Proceedings of *International Conference for High Performance Computing, Networking, Storage and Analysis* (supercomputing), Austin, Texas, November 2008.
64. Abhishek Das, Berkin Ozisikylmaz, Serkan Ozdemir, Gokhan Memik, Joseph Zambreno and Alok Choudhary, “Evaluating the Effects of Cache Redundancy on Profit,” in Proceedings of the 41st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-41), November 2008.
65. Wei-keng Liao, and Alok Choudhary. Dynamically Adapting File Domain Partitioning Methods for Collective I/O Based on Underlying Parallel File System Locking Protocols. In the Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, Texas, November 2008.
66. Wei-keng Liao, and Alok Choudhary. Dynamically Adapting File Domain Partitioning Methods for Collective I/O Based on Underlying Parallel File System Locking Protocols In the Proceedings of *International Conference for High Performance Computing, Networking, Storage and Analysis* (supercomputing), Austin, Texas, November 2008.
67. D. Bild, S. Misra, T. Chantem, P. Kumar, R. P. Dick, X. S. Hu, L. Shang, and A. Choudhary, "Temperature-Aware Test Scheduling for Multiprocessor Systems-On-Chip", In *Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design* (ICCAD '08). IEEE Press, Piscataway, NJ, USA, 59-66.

68. B. Ozisikyilmaz, G. Memik, A. Choudhary .Machine Learning Models to Predict Performance of Computer System Design Alternatives. *37th International Conference on Parallel Processing* (**ICPP**), Portland, OR, Sep. 2008.
69. Yan Gao, Ming Yang, Xiaonan Zhao, Bryan Pardo, Ying Wu, Thrasyvoulos N. Pappas, Alok Choudhary. Image spam hunter. In IEEE International Conference on Acoustics, Speech and Signal Processing, 2008.
70. Ramanathan Narayanan, Sanchit Misra, Simon Lin and Alok Choudhary. Mining Protein Interactions from Text using Convolution Kernels. In Intelligent Systems for Molecular Biology (ISMB)[Poster], July 2008.
71. Berkin Ozisikyilmaz, Gokhan Memik, and Alok Choudhary. Efficient System Design Space Exploration using Machine Learning Techniques. In *Proceedings Design Automation Conf.*, June 2008.
72. Alex Shye, Berkin Ozisikyilmaz, Arindam Mallik, Gokhan Memik, Peter Dinda, Robert Dick, and Alok Choudhary. Learning and Leveraging the Relationship between Architecture-Level Measurements and Individual User Satisfaction. In *Proceedings International Symposium Computer Architecture (ISCA)*, June 2008.
73. Ramanathan Narayanan, Sanchit Misra, Simon Lin and Alok Choudhary. Mining Protein Interactions from Text using Convolution Kernels. *Intelligent Systems for Molecular Biology* (ISMB), July 2008.
74. Sanchit Misra, Ramanathan Narayanan, Simon Lin and Alok Choudhary. HiSSS: A High Speed Sequence Search Method for Next Generation Sequences. *Intelligent Systems for Molecular Biology* (ISMB), July 2008.
75. Abhishek Das, Sanchit Misra, Sumeet Joshi, Joseph Zambreno, Gokhan Memik, and Alok Choudhary. An Efficient FPGA Implementation of Principal Component Analysis based Network Intrusion Detection System. *Proceedings Design, Automation & Test in Europe Conf.*, March 2008.
76. S. Pati, R. Narayanan, G. Memik, A. Choudhary, and J. Zambreno. "Design and Implementation of an FPGA Architecture for High-Speed Network Feature Extraction", *International Conference on Field-Programmable Technology (FPT)*, December 2007.
77. Avery Ching, Robert Ross, Wei-keng Liao, Lee Ward, and Alok Choudhary. Noncontiguous locking techniques for parallel file systems. *Supercomputing*, November 2007.
78. Wei-keng Liao, Avery Ching, Kenin Coloma, Arifa Nisar, Alok Choudhary, Jackie Chen, Ramanan Sankaran, and Scott Klanksy. Using MPI file caching to improve parallel write performance for large-scale scientific applications. *Supercomputing*, November 2007.
79. A. Das, S. Ozdemir, G. Memik and A. Choudhary. “Evaluating Voltage Islands in CMPs under Process Variations”. *International Conference on Computer Design (ICCD)*, Lake Tahoe, NV, October 2007.
80. A. Das, S. Ozdemir, G. Memik, J. Zambreno, and A. Choudhary. “Mitigating the Effects of Process Variations: Architectural Approaches for Improving Batch Performance”, *Workshop on Architectural Support for Gigascale Integration*, in conjunction with the 34th *International Symposium on Computer Architecture (ISCA)*, June 2007.
81. W.K. Liao, A. Choudhary, K. Coloma, and L. Ward “Implementation and Evaluation of Client-side File Caching for MPI-IO”, International Parallel and Distributed Processing Symposium, March 2007.
82. R. Narayanan, B. Ozisikyilmaz, G. Memik, A. Choudhary, J. Zambreno, Quantization Error and Accuracy-Performance Tradeoffs for Embedded Data Mining Workloads, *High Performance Data Mining Workshop* (**HPDM**), Beijing, China, May 2007.

83. R. Narayanan, D. Honbo, G. Memik, A. Choudhary, J. Zambreno , An FPGA Implementation of Decision Tree Classification. *Design, Automation, and Test in Europe* (**DATE**)*,* Nice, France, April 2007.
84. Wei-keng Liao, Avery Ching, Kenin Coloma, and Alok Choudhary. Improving MPI independent write performance using a two-stage write-behind buffering method. *NSF Next Generation Software Workshop, held in conjunction with the International Parallel and Distributed Processing Symposium*. March 2007.

85. Wei-keng Liao, Avery Ching, Kenin Coloma, Alok Choudhary, and Lee Ward. Implementation and evaluation of client-side file caching for MPI-IO. In *Proceedings of the International Parallel and Distributed Processing Symposium*. March, 2007.
86. Ying Liu, Alok Choudhary, Jianhong Zhou, Ashfaq Khokhar, A Scalable Distributed Stream Mining System for Highway Traffic Data, The 10th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD), September 2006, Berlin, Germany.
87. B. Ozisikyilmaz, R. Narayanan, J. Zambreno, G. Memik, A. Choudhary, "An Architectural Characterization Study of Data Mining and Bioinformatics Workloads", IEEE International Symposium on Workload Characterization (IISWC), San Jose, CA, October 2006.
88. R. Narayanan, B. Ozisikyilmaz, J. Zambreno, G. Memik, A. Choudhary, "MineBench: A Benchmark Suite for Data Mining Workloads", IEEE International Symposium on Workload Characterization (IISWC), San Jose, CA, October 2006.
89. K. Coloma, A. Ching, A. Choudhary, R. Ross, R. Thakur, and L. Ward, "A New Flexible MPI Collective I/O Implementation," IEEE International Conference on Cluster Computing, September 2006.
90. R. Simha, B. Narahari, J. Zambreno, and A. Choudhary. Secure execution with components from untrusted foundries. Advanced Networking and Communications Hardware Workshop, 2006.
91. Avery Ching, Wu-chun Feng, Heshan Lin, Xiaosong Ma, and Alok Choudhary. Exploring I/O strategies for parallel sequence database search tools with S3aSim. In Proceedings of the International Symposium on High Performance Distributed Computing, June 2006.
92. Peter Aarestad, Avery Ching, George Thiruvathukal, and Alok Choudhary. Scalable approaches for supporting MPI-IO atomicity. IEEE/ACM International Symposium on Cluster Computing and the Grid, May 2006.
93. J. Pisharath, J. Zambreno, B. Ozisikyilmaz, and A. Choudhary. Accelerating data mining workloads: current approaches and future challenges in system architecture design. 9th International Workshop on High Performance and Distributed Mining (HPDM), April 2006.
94. Jianwei Li, Alok Choudhary, Nan Jiang, and Wei-keng Liao. "Mining Frequent Patterns by Differential Refinement of Clustered Bitmaps". In Proceedings of Sixth SIAM International Conference on Data Mining, pages 294-305, April 2006.
95. Avery Ching, Alok Choudhary, Wei Keng Liao, Lee Ward, and Neil Pundit. Evaluating I/O characteristics and methods for storing structured scientific data. International Parallel and Distributed Processing Symposium, April 2006.
96. D. Nguyen, A. Das, G. Memik, A. Choudhary , A Reconfigurable Architecture for Network Intrusion Detection Using Principal Component Analysis, *Symposium on Field-Programmable Custom Computing Machines* (FCCM), Napa, California, April 2006 (poster presentation).
97. J. Zambreno, B. Ozisikyilmaz, J. Pisharath, G. Memik, and A. Choudhary. Performance characterization of data mining applications using MineBench. 9th Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW-9), February 2006.
98. D. Nguyen, G. Memik, and A. Choudhary , "A Reconfigurable Architecture for Network Intrusion Detection using Principal Component Analysis", 14th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), Monterey, CA, Feb. 2006.
99. Ying Liu, Wei-keng Liao, and Alok Choudhary. 2005. "A fast high utility itemsets mining algorithm". In *Proceedings of the 1st international workshop on Utility-based data mining* (UBDM '05). ACM, New York, NY, USA, 90-99. DOI=10.1145/1089827.1089839 http://doi.acm.org/10.1145/1089827.1089839.
100. Jay Pisharath and Alok Choudhary, "Design of a Hardware Accelerator for Density-Based Clustering Applications", International Conference on Application-specific Systems, Architectures and Processors (ASAP 2005).
101. S. W. Son, G. Chen, M. Kandemir, and A. Choudhary. Dynamic Compilation for Reducing Energy Consumption of I/O-intensive Applications, 18th International Workshop on Languages and Compilers for Parallel Computing (LCPC'05), Hawthorne, New York, October 2005.
102. Wei-keng Liao, Kenin Coloma, Alok Choudhary, and Lee Ward. Cooperative Write-Behind Data Buffering for MPI I/O. 12th European Parallel Virtual Machine and Message Passing Interface Conference (EURO PVM/MPI), Sorrento (Naples), Italy, September 2005.
103. D. Nguyen, G. Memik, S. Ogrenci Memik, and A. Choudhary, "Real-Time Feature Extraction for High Speed Networks", International Conference on Field Programmable Logic and Applications (FPL), Tampere, Finland, Aug. 2005.

104. Wei-keng Liao, Kenin Coloma, Alok Choudhary, Lee Ward, Eric Russel, and Sonja Tideman. Collective Caching: Application-Aware Client-Side File Caching. *14th International Symposium on High Performance Distributed Computing (HPDC)*, July 2005.
105. Jayaprakash Pisharath, Wei-keng Liao, and Alok Choudhary. Design and Evaluation of Database Layouts for MEMS-Based Storage Systems. International Database Engineering and Applications Symposium (IDEAS), July 2005.
106. Kenin Coloma, Alok Choudhary, Wei-keng Liao, Lee Ward, and Sonja Tideman. DAChe: Direct Access Cache System for Parallel I/O. 19th International Conference on Supercomputing (ICS), June 2005. **Received best paper award.**
107. S. W. Son, G. Chen, M. Kandemir, and A. Choudhary, Exposing Disk Layout to Compiler for Reducing Energy Consumption of Parallel Disk Based Systems, Symposium on Principles and Practice of Parallel Programming (PPoPP'05), Chicago, IL, June 2005.
108. Kripishankar Mohan, Bhagirath Narahari, Rahul Simha, Paul Ott, Alok Choudhary, and Joseph Zambreno. Performance Study of a Compiler/Hardware Approach to Embedded Systems Security. IEEE International Conference on Intelligence and Security Informatics (ISI), pp. 543–548, May 2005.
109. Olga Gelbart, Paul Ott, Bhagirath Narahari, Rahul Simha, Alok Choudhary, and Joseph Zambreno. CODESSEAL: A Compiler/FPGA Approach to Secure Applications. IEEE International Conference on Intelligence and Security Informatics (ISI), pp. 530–535, May 2005.
110. Y. Liu, W. Liao and A. Choudhary, “A Two-Phase Algorithm for Fast Discovery of High Utility Itemsets”, PAKDD, May 2005.
111. Joseph Zambreno, Dan Honbo, and Alok Choudhary. Exploiting Multi-Grained Parallelism in Reconfigurable SBC Architectures. *IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*, April 2005.
112. S. W. Son, M. Kandemir and A. Choudhary. Software-directed Disk Power Management for Scientific Applications, 19th International Parallel and Distributed Processing Symposium (IPDPS'05), Denver, Colorado, April 2005.
113. S. C. Chiu, A. Choudhary, and M. Kandemir, Fault Recovery Designs for Processor-embedded Distributed Storage Architectures with Intensive DB Workloads, 22nd IEEE/13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST'05), Monterey, CA, April 2005.
114. Steve Chiu, Wei-keng Liao, and Alok Choudhary. Processor-Embedded Distributed MEMS-Based Storage Systems for High-Performance I/O. *18th International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 91–100, April 2004.
115. Jayaprakash Pisharath, Alok Choudhary, and Mahmut Kandemir. Energy Management Schemes for Memory-Resident Database Systems. *13th ACM International Conference on Information and Knowledge Management (CIKM)*, pp. 218–227, November 2004.
116. Y. Liu, J. Pisharath, W.-K. Liao, G. Memik, A. Choudhary. Performance Evaluation and Characterization of Scalable Data Mining Algorithms, In Proc. of *Parallel and Distributed Computing and Systems* (**PDCS**), San Francisco, CA, Sept. 2004
117. K. Coloma, A. Choudhary, W. Liao, L. Ward, E. Russell, and N. Pundit, "Scalable High-level Caching for Parallel I/O", International Parallel and Distributed Processing Symposium, April, 2004.
118. Joseph Zambreno, Alok Choudhary, Rahul Simha, and Bhagi Narahari. Flexible Software Protection using HW/SW Codesign Techniques. *Design, Automation, and Test in Europe (DATE)*, pp. 636–641, February 2004.
119. A. Ching, A. Choudhary, W. Liao, R. Ross, and W. Gropp, "Efficient Structured Access in Parallel File Systems", 2003 IEEE International Conference on Cluster Computing, December, 2003.
120. Jianwei Li, Wei-keng Liao, Alok Choudhary, Robert Ross, Rajeev Thakur, William Gropp, and Rob Latham. "Parallel netCDF: A Scientific High-Performance I/O Interface", 15th SuperComputing conference, November 2003.
121. W. Liao, A. Choudhary, K. Coloma, G. Thiruvathukal, L. Ward, E. Russell, and N. Pundit, "Scalable Implementations of MPI Atomicity for Concurrent Overlapping I/O", International Conference on Parallel Processing, October, 2003.
122. G. Chen, M. Kandemir, I. Kolcu, and A. Choudhary, “Exploiting on-chip data transfers for improving performance of chip-scale multiprocessors”, International Conference on Parallel and Distributed Computing (Euro-Par'03), Klagenfurt, Austria, August 2003.

123. I. Kadayif, M. Kandemir, A. Choudhary, and M. Karakoy, “An energy-oriented evaluation of communication optimizations for microsensor networks”, International Conference on Parallel and Distributed Computing (Euro-Par'03), Klagenfurt, Austria, August 2003.
124. A. Ching, A. Choudhary, K. Coloma, W. Liao, R. Ross, and W. Gropp,"Noncontiguous Access through MPI-IO", 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid, May, 2003.
125. Ying Liu, Wei-keng Liao, and Alok Choudhary, “Design and Evaluation of a Parallel HOP Clustering Algorithm for Cosmological Simulation”, IPDPS, April 2003.
126. G. Memik, M Kandemir, A. Choudhary, and I Kadayif, “An integrated approach for improving cache behavior”, 6th Design Automation and Test in Europe Conference (DATE'03), Munich, Germany, March, 2003.
127. S. Chiu, W. Liao, A. Choudhary, “Design and Evaluation of Distributed Smart Disk Architecture for I/O-Intensive Workloads”, LNCS, pp 230-241, Springer-Verlag 2003.
128. K. Basu, A. Choudhary, J. Pisharath, and M. Kandemir, “Power protocol: reducing power dissipation on off-chip data buses”, 35th Annual International Symposium on Microarchitecture (MICRO-35), Istanbul, Turkey, November 2002.
129. J. Zambreno, M. Kandemir, and A. Choudhary, “Enhancing Compiler Techniques for Memory Energy Optimizations”, Second International Workshop on Embedded Software (EMSOFT’02), October 2002.
130. M. Kandemir, I. Kadayif, A. Choudhary, and J. Zambreno, “Optimizing Inter-Nest Data Locality”, International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES’02), Grenoble, France, October 2002.
131. J. Pisharath and A. Choudhary, “An Integrated Approach to Reducing Power Dissipation in Memory Hierarchies”, International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES’02), Grenoble, France, October 8-11, 2002.
132. A. Jones, D. Bagchi, S. Pal, X. Tang, A. Choudhary, and P. Banerjee, “PACT HDL: A C Compiler with Power and Performance Optimizations,” International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES’02), Grenoble, France, October 2002.
133. A. Ching, A. Choudhary, W. Liao, R. Ross, and W. Gropp, "Noncontiguous I/O through PVFS", IEEE International Conference on Cluster Computing, September, 2002.
134. J. Li, W.K. Liao, and A. Choudhary, “I/O Analysis and Optimization for an AMR Cosmology Application”, IEEE International Conference on Cluster Computing, Chicago, IL, September 2002.
135. I. Kadayif, M. Kandemir, and A. Choudhary, “A Hybrid Strategy Based on Data Distribution and migration for optimizing memory locality”, 15th Workshop on Languages and Compilers for Parallel Computing (LCPC'02), College Park, Maryland, July 25-27, 2002.
136. M. Kandemir and A. Choudhary, “Compiler-directed scratch pad memory hierarchy design and management”, Design Automation Conference (DAC'02), New Orleans, LA, June 2002.
137. M. Kandemir, J. Ramanujam, and A. Choudhary, “Exploiting shared scratch-pad memory space in embedded multiprocessor systems”, Design Automation Conference (DAC'02), New Orleans, LA, June 2002.
138. X. Shen and A. Choudhary, “MSI/O: A Distributed Multi-Storage I/O System for High Performance Data Intensive Computing”, IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid), Berlin, Germany, May 21-24, 2002.
139. M. Kandemir and A. Choudhary, “Compiler-directed I/O optimization”, Parallel and Distributed Processing Symposium (IPDPS'02), Fort Lauderdale, Florida, April 15-19, 2002.
140. A. Nayak, M. Haldar, A. Choudhary, and P. Banerjee, “Accurate Area and Delay Estimators for FPGAs”, Proc. Design Automation and Test in Europe (DATE-2002), Paris, France, March 2002.
141. G. Memik, M. Kandemir, and A. Choudhary, “Exploiting Inter-File Access Patterns Using Multi-Collective I/O”, USENIX Conference on File and Storage Technologies (FAST’02), Monterey, CA, January 28-29, 2002.
142. N. E. Crosbie, M. Kandemir, I. Kolcu, J. Ramanujam, and A. Choudhary, “Strategies for Improving Data Locality in Embedded Applications”, Proc. 7th Asia and South Pacific Design Automation Conference & 15th International Conference on VLSI Design (VLSI Design/ASPDAC’02), Bangalore, India, January 7-11, 2002.
143. M. Haldar, A. Nayak, A. Choudhary, and P. Banerjee, “A System for Synthesizing Optimized FPGA Hardware from MATLAB”, International Conference on Computer Aided Design (ICCAD 2001), San Jose, CA, November 4-8, 2001.

144. X. Shen and A. Choudhary, "DPFS: A Distributed Parallel File System", *Proc. 2001 International Conference on Parallel Processing*, Valencia, Spain, September 3-7, 2001.
145. P. Banerjee, M. Haldar, A. Nayak, and A. Choudhary, "Overview of the MATCH Compiler for Compiling MATLAB Programs into Hardware", *Proc. NASA Earth Science Technology Conference*, Washington, DC, August 2001.
146. X. Shen, W. Liao, and A. Choudhary, "An Integrated Graphical User Interface for High Performance Distributed Computing", *Proc. International Database Engineering and Applications Symposium (IDEAS)*, Grenoble, France, July 2001.
147. Nayak, M. Haldar, A. Choudhary, and P. Banerjee, "Precision and Error Analysis of MATLAB Applications During Automated Hardware Synthesis for FPGAs", *Proc. Design Automation and Test in Europe (DATE 2001),* Berlin, Germany*,* March 2001.
148. M. Haldar, A. Nayak, A. Choudhary, and P. Banerjee, "FPGA Hardware Synthesis from MATLAB Utilizing Optimized IP Cores", *Ninth ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, San Jose, CA, February 2001.
149. X. Shen, W. Liao, and A. Choudhary, "Remote I/O Optimization and Evaluation for Tertiary Storage Systems Through Storage Resource Broker", *IASTED Applied Informatics*, Innsbruck, Austria, February 2001.
150. M. Haldar, A. Nayak, A. Choudhary, and P. Banerjee, "Automated Synthesis of Pipelined Designs on FPGAs for Signal and Image Processing Applications Described in MATLAB", *Asia Pacific Design Automation Conference (ASP-DAC)*, Tokyo, Japan, February 2001.
151. A. Moshovos, G. Memik, and A. Choudhary, "JETTY: Reducing Snoop-Induced Power Consumption in Small-Scale, Bus-Based SMP Systems", *Seventh International Symposium on High Performance Computer Architecture*, Monterrey, Mexico, January 19-24, 2001.
152. M. Haldar, A. Nayak, N. Shenoy, A. Choudhary, and P. Banerjee, "FPGA Hardware Synthesis from MATLAB", *VLSI Design Conference*, Bangalore, India, January 2001.
153. N. Shenoy, P. Banerjee, A. Choudhary, and M. Kandemir, "Efficient Synthesis of Array Intensive Computations on FPGA Based Accelerators", *VLSI Design Conference*, Bangalore, India, January 2001.
154. W. Liao, X. Shen, and A. Choudhary, "Meta-Data Management System for High Performance Large-Scale Scientific Data Access", *7th International Conference on High Performance Computing*, December 2000.
155. M. Haldar, A. Nayak, A. Choudhary, and P. Banerjee, "Scheduling Algorithms for Automated Synthesis of Pipelined Designs on FPGAs for Applications Described in MATLAB", *International Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES 2000)*, San Jose, CA, November 2000.
156. S. More and A. Choudhary, "Scheduling Queries for Tape-Resident Data", *European Conference on Parallel Computing*, 2000.
157. M. Haldar, A. Nayak, A. Kanhere, P. Joisha, N. Shenoy, A. Choudhary, and P. Banerjee, "MATCH Virtual Machine: An Adaptive Runtime System to Execute MATLAB in Parallel", *International Conference on Parallel Processing (ICPP-2000)*, Toronto, CANADA, August 21-24, 2000.
158. G. Memik, M. Kandemir, and A. Choudhary, "Design and Evaluation of Smart Disk Architectures for Commerical Workloads", *International Conference on Parallel Processing (ICPP 2000)*, Toronto, CANADA, Aug. 21-24, 2000.
159. H. Nagesh, S. Goil, and A. Choudhary, "PMAFIA: A Scalable Parallel Subspace Clustering Algorithm for Massive Datasets", *International Conference on Parallel Processing (ICPP 2000)*, Toronto, CANADA, Aug. 21-24, 2000.
160. X. Shen and A. Choudhary, "A Distributed Multi-Storage Resource Architecture and I/O Performance Prediction for Scientific Computing", *High Performance Distributed Computing Conference (HPDC 2000)*, Pittsburgh, PA, Aug. 1-4, 2000.
161. A. Nayak, M. Haldar, A. Kanhere, P. Joisha, N. Shenoy, A. Choudhary, and P. Banerjee, "A Library Based Compiler to Execute MATLAB Programs on a Heterogeneous Platform", *ISCA 13th International Conference on Parallel and Distributed Computing Systems (PDCS-2000)*, Las Vegas, NV, Aug. 8-10, 2000.
162. G. Memik, M. Kandemir, and A. Choudhary, "Design and Evaluation of a Compiler-Directed I/O Technique", *European Conference on Parallel Computing (Euro-Par 2000)*, Munich, GERMANY, August 2000.

163. P. Joisha, A. Kanhere, P. Banerjee, N. Shenoy, and A. Choudhary, “Handling Context-Sensitive Syntactic Issues in the Design of a Front-end for a MATLAB Compiler”, *Proc. ACM Array Programming Languages Conference (APL-Berlin-2000)*, Berlin, GERMANY, July 24-27, 2000.
164. X. Shen, G. Thiruvathukal, W. Liao, A. Choudhary, and A. Singh, “A Java Graphical User Interface for Large-Scale Scientific Computations in Distributed Systems”, *Fourth International Conference on High Performance Computing in Asia-Pacific Region (HPC-ASIA2000)*, Beijing, CHINA, May 14-17, 2000.
165. X. Shen, W. Liao, A. Choudhary, G. Memik, M. Kandemir, S. More, G. Thiruvathukal, and A. Singh, “A Novel Application Development Environment for Large-Scale Scientific Computations”, 2000 International Conference on Supercomputing, Santa Fe, NM, May 8-11, 2000.
166. W. Liao, D. Weiner, A. Choudhary, and P. Varshney, “Design and Evaluation of I/O Strategies for Parallel Pipelined STAP Applications”, 14th International Parallel and Distributed Processing Symposium (IPDPS 2000), Cancun, MEXICO, May 1-5, 2000.
167. P. Banerjee, N. Shenoy, A. Choudhary, S. Hauck, M. Haldar, P. Joisha, A. Jones, A. Kanhere, A. Nayak, S. Periyacheri, M. Walkden, and D. Zaretsky, “A MATLAB Compiler for Distributed Heterogeneous Reconfigurable Computing Systems”, International Symposium on FPGA Custom Computing Machines (FCCM-2000), Napa Valley, CA, April 2000.
168. N. Shenoy, A. Choudhary, and P. Banerjee, “A System-Level Synthesis Algorithm with Guaranteed Solution Quality”, Design Automation and Test in Europe (DATE 2000), Paris, FRANCE, March 27-30, 2000.
169. M. Haldar, A. Nayak, A. Choudhary, and P. Banerjee, “Parallel Algorithms for FPGA Placement”, Great Lakes Symposium on VLSI (GVLSI 2000), Chicago, IL, March 2000.
170. G. Memik, M. Kandemir, and A. Choudhary, “APRIL: A Run-Time Library for Tape Resident Data”, *8th NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies and 17th IEEE Symposium on Mass Storage Systems*, Baltimore, MD, March 2000.
171. S. More and A. Choudhary, “Tertiary Storage Organization for Large Multidimensional Datasets”, *8th NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies and 17th IEEE Symposium on Mass Storage Systems*, 2000.

Talks: Invited & Keynote (Talks prior to 2003 are listed in appendix C.)

- Invited talk “Big Data and Exascale Systems — Challenges and Opportunities”, 1st International Workshop on Strategic Development of High Performance Computers, Tsukuba, Japan, 03/2013.
- Invited talk "Action-Based Connections: Big Data Analytics in Social Media", at a Social Media Week-Chicago event hosted by Northwestern University's Medill School of Journalism on 09/2012. Link to News Article on Northwestern's Kellogg web site .
- Invited Keynote Address Department of Energy Computational Science Graduate Fellowship conference (DOE CSGF 2012), Arlington, Virginia, 07/2012.
- Keynote speaker, “Discovering Knowledge from Massive Social Networks and Science Data - Next Frontier for HPC” International Symposium on Cluster Cloud and Grid Computing (CCGrid 2012), Ottawa, Canada, 05/2012.
- Keynote speaker, “Discovering Knowledge from Massive Social Networks and Science Data - Next Frontier for HPC” International Conference on High Performance Computing (HiPC 2011), Bangalore, India 12/2011.
- “Discovering Knowledge from Massive Data and Social Networks”, Distinguished Lecture Series, Boston University, 09/2011.
- "Developing Scalable and Power-Efficient Data Mining Kernels", Understanding Climate Change Workshop, University of Minnesota, 08/2011.
- “Discovering Knowledge from Massive Social and Science Network Unstructured Data”, Data Intensive Computing Workshop, The City University of New York, 07/2011.
- “Exascale Supercomputing: Challenges and Approaches to Data Intensive Computing and Analytics,” University of Tsukba, Japan, Feb 8, 2011.
- “High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery”, AIST, Tsukuba, and Tokyo, Feb 9, 2011.
- “Exascale Supercomputing: Challenges and Approaches to Data Intensive Computing and Analytics,” Tokyo Institute of Technology, Tokyo, Japan, Feb 17, 2011.

- “Exascale Supercomputing: Challenges and Approaches to Data Intensive Computing and Analytics,” University of Tokyo, Tokyo, Japan, Feb 16, 2011.
- “Exascale Supercomputing: Challenges and Approaches to Data Intensive Computing and Analytics,” Tokyo Supercomputing Center, Tokyo, Japan, Feb 15, 2011.
- “Exascale Supercomputing: Challenges and Approaches to Data Intensive Computing and Analytics,” RIKEN, AICS, Kobe Supercomputing Center, Japan, Feb 10, 2011.
- “High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery”, Univ. of Florida, Dec 2010.
- Keynote Talk, “High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery" The 2010 International Conference on Multimedia Information Networking and Security (MINES 2010),” Nanjing, China, Nov 2010.
- Keynote Talk, “High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery" The 2010 Workshop Multimedia Technologies,” Anyang, China, Nov 2010.
- Invited Talk, “"High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery", Dean’s Grand Challenge Lectures in Medicine and Engineering (May 2010)
- Keynote Talk, “Computational Thinking and Privacy in the 21st Century”, at Computer Associate’s Innovation Tour, Dec 2007
- Invited Talk, “Large-Scale Scientific Knowledge Discovery: Problems and Potential Approaches,” Next Generation Data Mining 2007 (organized by NSF), October 2007.
- Invited Talk, “Scalable I/O Middleware,” Oak-Ridge National Labs Fall Creek Conference, Nashville, TN, September, 2007.
- Invited Talk, “Challenges and Opportunities in an Information Dominated Age for Large-Scale Scientific Knowledge Discovery,” Georgia Tech, April 2007.
- Invited/Keynote talk,"Are Data Mining Applications Different from Traditional Applications? An Architecture Perspective" The 17th European Conference on Machine Learning and the 10th European Conference on Principles and Practice of Knowledge Discovery in Databases, Berlin, Germany, Sept 2006
- Keynote talk, “A Case for using High-level Knowledge for Scalable I/O Optimizations,” IEEE International Conference on Cluster Computing, Barcelona, Spain, Sept 2006.
- Invited Talk, “Architecture, High-Performance Computing and Information Assurance Research at Northwestern University,” Intel Corp. Bangalore, India, July 2006.
- Panel talk, “Future System Software Issues for Supercomputing,” IBM Bluegene Workshop, Tokyo, Japan, April 2006.
- Panel talk, “Complexity in Future High-Performance Storage Systems,” SOS workshop, Maui, March, 2006.
- Keynote talk, “Scientific Data Management :Science in an Information-Dominated Age” International Conference on Computational Science, Atlanta, May 2005.
- Keynote talk, “Ultra-Scale I/O, Challenges and Prospects.” At SOS Conference, Davos, Switzerland, March 2005.
- Keynote talk, “A Case for On-Line Analysis in Scientific Simulations”, European Simulation Conference, Oct 2003, , Naples, Italy.
- Keynote talk, “Scientific Data Management “, Richard Tapia Celebration of Diversity in Computing Conference, October 2003.
- IIT, Chicago, Distinguished Lecture, “Scalable Scientific Data Management”, Sept. 2003
- Rice University, “Scalable Scientific Data Management”, May 2003.
- Lawrence Livermore national laboratories, “Scientific Data Management: A Case for using High-level Information for I/O Optimizations”, May 2003.
- KAIST, Korea, “Compilation and Architecture Techniques for Power Aware Systems and Software Protection”, May 2003.
- University of Madrid, Spain, “Power-Aware Computing”, March 2003.
- University of Madrid, Spain, “Scalable I/O”, March 2003.
- Indian Institute of Sciences, Bangalore, India, “Parallel and Scalable Data Mining”, Jan 2003.

Professional Activities

Journal Editorial

- *Associate Editor of IEEE Transactions on Parallel and Distributed (1998-2004)*
- *Subject Area Editor of Journal of Parallel and Distributed Computing (1993-2007)*
- *Editor of International Journal on High-Performance Computing and Networking*
- *International Journal of Web Services*
- *Guest Editor, IEEE Computer (top computer science magazine)*
- *Guest Editor, IEEE Parallel and Distributed Technology*

Conferences and Workshops (Recent activities).

Prior activities are given in Appendix E.

- Program Committee, International Conference on Data Mining, Brussels, Belgium (ICDM 2012).
- International Advisory Committee, International Symposium on Parallel Applications, Madrid, Spain (ISPA 2012).
- International Advisory Committee, International Conference on Data Mining, Computational and Business Intelligence, Bhubaneshwar, Orissa, India (DMCBI 2012).
- Program Committee, Architectures and Systems for Big Data Workshop, Galveston Island, TX (ASBD 2011).
- Tutorials Committee, Supercomputing 2011, Seattle WA (SC11, HPC).
- General Chair, International Conference on High Performance Computing & Communications 2011, Banff, Canada (HPCC11).
- Program Committee, Data Cloud 2011, Anchorage, AK (IPDPS 2011)
- Chair, Reaching Exascale in this Decade, Goa, India (HiPC Workshop 2010).
- Program Committee, GPUs and Scientific Applications Workshop, Vienna, Austria (PACT 2010)
- Co-Chair, Petascale Data Analytics on Clouds 2010 (PDAC-10, ACM/IEEE SC10).
- Vice Chair, International Conference on Data Mining (ICDM 2010).
- Program Co-Chair, International Conference of Parallel Processing (ICPP 2009).
- PC Vice Chair, International Conference on Data Mining, (ICDM 2009).
- Panelist, Several (more than 20) Panels in National Science Foundation
- Served on Program Committees of
    - IPDPS 2011, Several SC 2011 Workshops
    - GPUScA 2010, EMCL2010, SDM 2010, HPDC 2010 workshop.
    - KDD2009, SDM 2009, Supercomputing 2009, SC tutorials 2009, PAKDD 2009
    - KDD2008, ICDM 2008, Supercomputing 2009.
    - Micro 2007, SDM 2007
    - IPDPS, 2005, 2007
    - ICWS, 2004, 2005, 2006
    - PDM 2006
    - Hot AC, 2006

Recent Media Coverage:

- Interviewed by Tasty Trade, a real financial network (interview).
- Appeared in the McCormick Fall Magazine article on the Data Age.
- VivaKi Inks Deal with The Echo System & Voxsup for Exclusive Access to Social Data & Buying Platform. Link to Announcement at Yahoo!
- Speaker at a Social Media Week-Chicago event hosted by Northwestern University's Medill School of Journalism on Sept. 28th, 2012. Presented "Action-Based Connections: Big Data Analytics in Social Media", Link to News Article on Northwestern's Kellogg web site .

- Appeared in Northwestern Research Newsletter, September 2012.
- Invited to give Keynote address at the Department of Energy Computational Science Graduate Fellowship (DOE CSGF), Arlington, VA, 7/2012.
- Keynote speaker at the International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2012), Ottawa, Canada, 5/2012.
- International Exascale Software Project, Kobe, Japan, 04/2012, (Agenda).
- Quoted in the Advanced Scientific Computing Research (ASCR) Discovery Article, "HPC must ramp up efficiency to deal with extreme-data flood", 03/2012, (Article).
- Received an invitation from the White House and will attend the Big Data Event on March 29th, 2012 in Washington DC. The purpose of the Big Data Event is to address the challenges and opportunities relating to "Big Data". Specifically, the Big Data Event provides the opportunity to learn about several new federal efforts relating to big data and to provide a discussion forum where leaders from government, academia and industry can work together on this important challenge. SDAV (Scalable Data Management, Analysis and Visualization) was formally announced March 29th as part of the Obama Administration's "Big Data Research and Development Initiative".
- Recognition of DOE (Department of Energy) supported research through the NSF (National Science Foundation) Expeditions program on climate science problems. The results of this research will be featured in an advertisement by North Carolina State University in the Chronicle of Higher Education and will be published as an article on the DOE (Department of Energy) website. Dr. Choudhary and Dr. Nagiza Samatova have been invited to present this work to the ASCAC (Advanced Scientific Computing Advisory Committee) of the DOE, 03/2012 (presentation).
- International Conference on High Performance Computing (HiPC 2011) Keynote speaker, Bangalore, India 12/2011
  "Discovering Knowledge from Massive Social Networks and Science Data - Next Frontier for HPC", (program) (abstract).
- Distinguished Lecture Series, Boston University, ”Discovering Knowledge from Massive Data and Social Networks”, 09/2011.
- Understanding Climate Change Workshop, University of Minnesota, "Developing Scalable and Power-Efficient Data Mining Kernels", 08/2011
- Data Intensive Computing Workshop, The City University of New York, “Discovering Knowledge from Massive Social and Science Network Unstructured Data”, 07/2011.
- Pulse of the Tweeters related media:
  - Pulse of The Tweeters website (Making Twitter more relevant)
  - DiscoveryNews - Can Twitterblitzing Start an Effective Twitter Trend? 08/2011
  - “Influencing the Influencers”, Beverly Macy, Pulse of The Tweeters in Huffington Post , (05/2011)
  - “Tens of thousands defy violent reprisals to protest in Cairo’s Tahrir Square”, Dr. Choudhary on NPR/WBEZ Worldview , (02/2011)
  - “Celeb tweets make impact on top stories in 2010”, Patricia Reaney, Pulse of The Tweeters in Reuters , (01/2011).
  - “Most Influential Tweeters of 2010”, ABC Chicago, Dr. Choudhary's interview on ABC 7 , (01/2011).
  - Alok Choudhary spoke with a journalist in Bogotá, Colombia about his research where he found that celebrities on Twitter with millions of followers are mostly ignored. Columbian National Radio interview (11/2010)
  - “Website Ranks Most Influential Tweeters”, New York Times, NU Article , (09/2010)
  - “Most influential Twitter users? Pulse of the Tweeters will tell you.” Article from thenextweb.com Social Media (09/2010)

- Northwestern University Electrical Engineering and Computer Science Department
  Meet the faculty series**:** Alok N. Choudhary
  High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery
- Deans Grand Challenge Lectures in Medicine and Engineering: Alok Choudhary
  High-Performance Data Mining: An Essential Paradigm for Knowledge Discovery

Selected Past Research grants:

- DARPA/DOE, “High-Productivity Languages and Runtime Systems,” 400K (2005-2007)
- Department of Energy (DOE), Enabling Technologies for Scientific Data Management, $1,117,000. (2001-2007).
- NSF NGS program, Dynamic Runtime and Compilation Support for I/O-Intensive Applications, $336,000 (2004-2007)
- DOE (Sandia National Laboratories) , I/O Optimizations for Large-Scale systems, $175K / per year (2003- 2006).
- AFOSR, “Evolvable Approaches to Software Verification and Validation,” 225K, 2005-2008.
- Chicago Biomedical Consortium, “Bioinformatics and Proteomics,” 75K, 2006-2007.
- NSF (Co-PI), “SoD-TEAM: Robust System Design Under Weak Component Assumptions,” 200K (2006-2008).
- IBM Faculty Partnership award, $40K per year.
- Intel Corporation, Characterizing Scalable Data Mining Kernels/Primitives on SMPs, $49,658 per year and Equipment for $35,000. (2003-2006)
- NASA, GSRP: High-Performance and Fault-tolerant system software for Large-scale Applications, $48,000 (2003-2006)
- DOE, Parallel I/O Platform for High-Performance Data Mining system, $30,000 (2003-2006)
- DOE, Efficient Implementation for Overlapping File Access in MPI-IO, $50,000 (2006-2007)
- DOE (LLNL), Caching and Parallel I/O optimizations in MPI-I/O, $100,000 (2005-2006)
- Lawrence Livermore National Laboratories, On-Line Data Mining and Analysis, $100,000
- DOE, High Performance Data Management, Access and Storage Techniques for Tera-scale Scientific Applications, $876,000
- NSF Young Investigator Award, “Compiler and Runtime Optimizations for Distributed Memory Machines,” 1993-1999, $62,500 per year.
- NSF, Scalable I/O Management and Access Optimizations for Scientific Applications for High-Performance Computing, $99,827
- NSF (Subcontract from University of Wyoming), “Interoperable Data Files for High-Performance Computing,” $152,867, 7/1/97 - 6/30/00.
- Department of Energy (ASCI Level 2), “Large High-Performance Data Mangement, Access and Storage Techniques for Tera-scale Scientific Applications,” $1,065,658, 10/98 - 10/01.
- NSF (CISE Research Infrastructure Program) (CO-PI with P. Banerjee et. al.), “A Distributed High-Performance Computing Infrastructure,” $906,512, 9/1/97 - 8/31/02.
- DARPA (Co-PI), “A MATLAB Compilation Environment for Adaptive Computing,” $1,855,662, 1998-2001.
- DARPA (Co-PI), “PACT: Power-Aware Architectural and Compilation Techniques,” $1,958,000, 2000-2003.
- DARPA (Subcontract from Caltech), “Scalable I/O Initiative” $155,000 per year 1995-1999.
- NSF, “System Software Support for Input-Output on Parallel Computing,” $190,000, 11/1/96 - 4/30/99.
- Intel Corporation, “Modeling and Evaluation of I/O Architectures in Servers,” $45,000, 1995-96.
- Intel Corporation, “Modeling and Evaluation of I/O Architectures in Servers,” $45,000, 1996-97.
- Intel Corporation, “Modeling and Evaluation of I/O Architectures in Servers,” $45,000, 1997-98.
- Air Force Systems Command (Rome Labs) (Subcontract form Syracuse University), “Design, Development, Benchmarking and Evaluation of Parallel Applications,” $118,870 per year, 12/24/96 - 1/31/99.
- Sandia National Laboratories, “Runtime Libraries for Large-Scale Parallel I/O,” $74,986, 1/15/97 - 1/14/98.
- Department of Energy (with Argonne National Laboratory), “Hybrid Technology Multi-threaded Computer Architecture for Petaflops Computing,” $48,800, 10/1/97 - 9/30/98.
- Intel Corporation, “I/O Architectures for Multi-Media Servers,” $40,000, 1995-96.
- Intel Corporation, “Compiler and Runtime Support for Parallel I/O,” matching funds for NYI, $38,000, 1995-96.
- Rome Labs, “Software Specification and Verification of Real-Time Systems of Parallel Computers, $99,600, 1994-95.

- Intel Corporation, “Compiler and Runtime Support for Parallel I/O,” matching funds for NYI, $38,000, 1994-95.
- IBM Faculty Development Award, $30,000, 1994-95.
- NASA-CESDIS, “High Performance Input-Output Support for Four Dimensional Data Assimilation, $50,000 per year for three years, 1993-96.
- NSF Research Initiation, “Design, Analysis, Modeling and Evaluation of Hierarchical Caches for Multiprocessors,” $69,936, 1991-94.
- ARPA (Defense Advance Research Project Agency), “Fortran 77D and Fortran 90D: Scalable and Protable Software Modules for SIMD and MIMD Distributed Memory Parallel Computers” (with G. Fox), $1,729,000, 1991-94.
- ONR, “Software Environments and Scalable High Level Data Structure Extensions of Fortran for Irregular Loosely Synchronous Problems on SIMD and MIMD Parallel Computers” (with G. Fox), $50,000 per year, 1992-95.
- NSF/NASA, “Software Environments and Scalable High Level Data Structure Extensions of Fortran for Irregular Loosely Synchronous Problems on SIMD and MIMD Parallel Computers” (with G. Fox), $150,00 per year, 1992-95.
- IBM, “Development of Software and Applications for Parallel Processors” (with G. Fox), $143,823, 1992.
- IEEE Engineering Foundation, “Parallel Processing for Computer Vision,” $20,000, 1990-91.
- Unrestricted Award from Texas Instruments, “Parallel Processing,” $15,000, 1991.
- Rome Labs, “Design and Analysis of Optical Interface Message Processor (OPTIMP)” (with S. Hariri and Q. Wang Song), $100,000, 1991-92.

Selected Industry Advising, Consulting, Collaboration and Technology Transfer

(past few years)

- Co-founder and VP of research and technology (1999-2002) of Accelchip Inc. (Accelchip was eventually acquired by Xilinx).
    - Tools and software for synthesizing hardware on FPGAs from high-level languages such as MATLAB in particular for applications in Digital Signal Processing (DSP), Communication and embedded systems.
- Co-founder of Nimkathana Inc. (HPC, data mining and analytics consulting)
- Board of Directors of C3Research
- Consultant and advisor, Sun Microsystems
- Consultant, Microsoft
- Consultant, Lucent
- Consultant to SPSS
- Consultant to Sony
- Consultant, Futuretrading
- Technical advisory board, Xtreme data Inc
- Advisor/consultant Netteza Inc.
    - Scalable Data warehouses
    - Architectures and algorithms for scalable data analysis, OLAP and decision support (applications in Retail, communication, marketing and other fields)
- Technical Advisory Board and consultant to Virtual Gold Inc.
    - Data mining applications (e.g., sports, entertainment, CRM, communication)
    - Multidimensional databases and analysis
- Consultant ZS Associates
    - Datawarehouse design for pharmaceutical applications
- IBM
    - Database systems design
    - Algorithm development
    - Performance analysis

    - High-performance systems
- Intel
    - Multimedia Server systems design and analysis
    - High-performance I/O and Storage Architectures
    - Data Mining
- Technical Advisory Board, Prismmedia and Acciro Inc
    - Distributed file systems for multimedia
- CTO summit attendee and collaborator with Teradata (a division of NCR)
- High-performance file systems and I/O software from research projects adopted by various public domain software systems and companies
- Technology transfer to Portland group Inc. (now part of ST Microelectronics) of Compiler Research project
    - Department of defense research project on compiling for high-performance parallel machines, when I was at Syracuse University
- Tata Consultancy Services (1982-1984)
    - Systems Analyst and Designer (database design and applications)
- Granted several patents and filed several patents.

Teaching, Curriculum development and Interdisciplinary programs:

I have been involved in developing multidisciplinary programs with many schools and within schools of engineering. These programs include Computational biology and bioinformatics (with Science and Medical School), Masters of Information technology (with Kellogg School of Management), Analytical Marketing and Customer Relationship management (with Kellogg School of Management), Business Intelligence (Kellogg School of Management), PhD cluster in “Computation in physical Sciences” (with School of Arts and Sciences), PhD cluster in “Predictive Science and Engineering Design (PSED)” (with School of Arts and Sciences), and currently working on programs with Medical School, Medill School of Journalism and Communication.

Courses:

- Graduate/Undergraduate Level
    - Operating Systems
    - Computer Organization
    - Computer Architecture
    - Advanced Computer Architecture
    - Database Systems
    - Information Security and Assurance
    - Parallel Computing
- Undergraduate Level
    - Object Oriented Design and C++
    - Digital Logic Design
    - Computer Architecture projects
- Graduate Level
    - Distributed Computing
    - Memory Systems for Multiprocessors and Scalable Architectures
    - Advanced Computer Architecture
    - Parallel Computing Systems
    - Distributed Databases and Information Systems
    - Information Assurance and Security

    - Advanced Data Mining
- Business/Technology in Business Courses
    - **Academic Director,** Executive program on Managing Customer Relationships for Profit, Kellogg School of Management
    - Customer Relationship Management and Analytical Marketing, Kellogg School of Management, NWU
    - Custom Executive Program for Microsoft and Sony.
    - Strategic Data and Knowledge Management for Business Intelligence at Kellogg School of Management, NWU
    - Business Intelligence Technologies and Customer Relationship Management at Kellogg School of Management, NWU
    - Distributed Information Systems at Masters of Information Technology Program at NWU
- Tutorials at Conferences
    - Memory Systems for Multiprocessors and Scalable Architectures, at ISCA 91, Supercomputing 91, 92.
    - Compiling for Distributed Memory Machines, at SHPCC 94, ISCA 94.
    - Issues in Designing Parallel I/O Systems, at ISCA 97,98, Supercomputing 97.
- Program Development
    - Revised Computer Engineering curriculum as chair and director of the computer engineering curriculum
    - Founding member of the bioinformatics program
    - Development of Interdisciplinary programs with Kellogg School of management, School of medicine
    - Developed an interdisciplinary center for information security and assurance

Appendix A. Journal Papers Published Prior to 2000

1. M. Kandemir, P. Banerjee, A. Choudhary, J. Ramanujam, N. Shenoy, “A Global Communication Optimization Technique Based on Data Flow Analysis and Linear Algebra”, ACM Transactions on Programming Languages and Systems (TOPLAS), Vol. 21, No. 6, November 1999.
2. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “A Matrix-Based Approach to Global Locality Optimization,” *Journal of Parallel and Distributed Computing, Special Issue on Compilation and Architectural Support for Parallel Applications*, Vol. 58, No. 2, Aug. 1999.
3. D. Jadav, A. Choudhary, and P. Berra, “Techniques for Increasing the Stream Capacity of a High-Performance Multimedia Server,” *IEEE Transactions on Knowledge and Data Engineering*, Vol. 11, No. 2, March/April 1999.
4. M. Kandemir, J. Ramanujam, and A. Choudhary, “Improving Cache Locality by a Combination of Loop and Data Transformations,” *IEEE Transactions on Computers*, Vol. 48, No. 2, Feb. 1999.
5. M. Kandemir, A. Choudhary, N. Shenoy, P. Banerjee, and J. Ramanujam, “A Linear Algebra Framework for Automatic Determination of Optimal Data Layouts,” *IEEE Transactions on Parallel and Distributed Systems*, Vol. 10, No. 2, Feb. 1999.
6. M. Kandaswamy, M. Kandemir, A. Choudhary and D. Bernholdt, “An Experimental Study to Analyze and Optimize Hartree-Fock Application's I/O With PASSION,” *International Journal of High Performance Computing Applications*, Vol. 12, No. 4, pp. 411-439, Winter 1998.
7. M. Kandemir, J. Ramanujam, R. Bordawekar, and A. Choudhary, “Compilation Techniques for Out-of-Core Parallel Computations,” *Parallel Computing*, 24(3-4): 597-628, June 1998.
8. M. Kandemir, A. Choudhary, J. Ramanujam and M. Kandaswamy, “Locality Optimization Algorithms for Compilation of Out-of-Core Codes, *Journal of Information Science and Engineering*, 14(1): 107-138, March 1998.
9. S. Adve, D. Burger, R. Eigenmann, A. Rawsthorne, M. Smith, C. Gebotys, M. Kandemir, D. Lilja, A. Choudhary, J. Fang, P. Yew, “Changing Interaction of Compiler and Architecture,” *IEEE Computer*, Vol. 30, No. 12, pp. 51-58, Dec. 1997.
10. I. Foster, D. Kohr, Jr., R. Krishnaiyer, and A. Choudhary, “A Library-Based Approach to Task Parallelism in a Data-Parallel Language,” *Journal of Parallel and Distributed Computing*, Vol. 45, No. 2, pp. 148-158, September 1997.
11. S. Goil and A. Choudhary, “High Performance OLAP and Data Mining on Parallel Computers,” *Journal of Data Mining and Knowledge Discovery (Special Issue on Scalable High-Performance Computing for KDD)*, Vol. 1, No. 4, pp. 391-417, 1997.
12. D. Jadav, C. Srinilta, A. Choudhary, and P. B. Berra, “An Evaluation of Design Trade-Offs in a High-Performance, Media-On-Demand Server,” *ACM Multimedia Systems Journal*, Vol. 5, Jan. 1997, pp. 53-68.
13. R. Thakur and A. Choudhary, “An Extended Two-Phase Method for Accessing Sections of Out-of-Core Arrays,” *Journal of Scientific Programming*, Vol. 5, No. 4, pp. 301-317, Winter 1996.
14. A. Choudhary, R. Thakur, R. Bordawekar, S. More, and S. Kutipidi, “PASSION: Optimized I/O for Parallel Applications,” *IEEE Computer*, (29) 6: 70-78, June 1996 (PASSION: PArallel and Scalable Software for Input-Output).
15. R. Thakur, A. Choudhary and J. Ramanujam, “Efficient Algorithms for Array Redistribution,” *IEEE Trans. on Parallel and Distributed Systems*, (7) 6:587-594, June 1996.
16. D. Jadav, C. Srinilta, A. Choudhary, and P. B. Berra, “Techniques for Scheduling I/O in a High Performance Multimedia-on-Demand Server,” *Journal of Parallel and Distributed Computing*, Vol. 30, No. 2, Nov. 1995.
17. R. Ponnusamy, J. Saltz, A. Choudhary, Y.-S. Hwang, and G. Fox, “Runtime Support and Compilation Methods for User-Specified Irregular Data Distributions,” *IEEE Transactions on Parallel and Distributed Systems*, Vol. 6, No. 8, August, 1995.
18. M. Harry, J. del Rosario, and A. Choudhary, “The Design of VIP-FS: A Virtual, Parallel File System for High Performance Parallel and Distributed Computing,” *ACM Operating Systems Review*, Vol. 29, No. 3, July 1995.
19. A. Choudhary, I. Foster, and R. Stevens, “Multimedia Applications and High-Performance Computing” (Guest Editors' Introduction), *IEEE Parallel and Distributed Technology*, Summer 1995.
20. D. Jadav and A. Choudhary, “Designing and Implementing High Performance Media-on-Demand Servers,” *IEEE Parallel and Distributed Technology*, Summer 1995.

21. D. Reed, C. Catlett, A. Choudhary, D. Kotz and M. Snir, “Parallel I/O: Getting Ready for Prime Time,” *IEEE Parallel and Distributed Technology*, Summer 1995.
22. R. Ponnusamy, J. Saltz, A. Choudhary, Yuan-Shin Hwang, R. Das, and G. Fox, “Supporting Irregular Distributions Using Data-Parallel Languages,” *IEEE Parallel and Distributed Technology*, Vol. 3, No. 1, Spring 1995.
23. R. Thakur, A. Choudhary, R. Ponnusamy and G. Fox, “Complete Exchange on the CM-5 and Touchstone Delta,” *The Journal of Supercomputing*, 1995.
24. S. Krishnamoorthy and A. Choudhary, “A Scalable Distributed Shared Memory Architecture,” *Journal of Parallel and Distributed Computing*, Vol. 22, No. 3, pp. 547-554, Sept. 1994.
25. Bozkus, Alok Choudhary, Geoffrey C. Fox, T. Haupt, S. Ranka, and M. Y. Wu, “Compiling Fortran 90D/HPF for Distributed Memory MIMD Computers,” *Journal of Parallel and Distributed Computing*, April 1994, pp. 15-26.
26. A. N. Choudhary, B. Narahari, D. M. Nicol and R. Simha, “Optimal Processor Assignment for Pipeline Computations,” *IEEE Transactions on Parallel and Distributed Systems*, April 1994.
27. J. M. del Rosario and A. Choudhary, “High Performance I/O for Massively Parallel Computers: Problems and Prospects,” *IEEE Computer*, March 1994, pp. 59-68.
28. R. Ponnusamy, N. Mansour, A. Choudhary and G. Fox, “Graph Contraction for Mapping Data on Parallel Computers: A Quality-Cost Tradeoff,” *Journal of Scientific Programming*, Vol 3, pp. 73-82, 1994.
29. A. Choudhary and S. Krishnamoorthy, “Evaluation of Multilevel Caches for Shared Memory Multiprocessors,” *Journal of Computer and Software Engineering*, Vol. 2, No. 1, 1994, pp. 87-110.
30. A. Choudhary and R. Thakur, “Connected Component Labeling on Course Grain Parallel Computers: An Experimental Study,” *Journal of Parallel and Distributed Computing*, Vol. 20, No. 1, Jan. 1994, pp. 78-83.
31. A. N. Choudhary, G. Fox, S. Hiranandani, K. Kennedy, C. Koelbel, S. Ranka and C. W. Tseng, “Unified Compilation of Fortran 77D and 90D,” *Letters on Programming Languages and Systems (LOPLAS)*, 2(1-4):95-114, March-December 1993.
32. A. Choudhary, J. H. Patel and N. Ahuja, “NETRA: A Hierarchical and Partitionable Architecture for Computer Vision,” *IEEE Transactions on Parallel and Distributed Computing*, Oct. 1993, pp. 1092-1104.
33. R. Ponnusamy, R. Thakur, A. Choudhary, K. Velamakanni, Z. Bozkus and G. Fox, “Experimental Performance Evaluation of the CM-5,” *Journal of Parallel and Distributed Computing*, September 1993, Vol. 19, No. 1, pp. 192-202.
34. A. Choudhary, B. Narahari, and R. Krishnamurti, “An Efficient Heuristic Scheme for Dynamic Remapping of Parallel Computations,” *Parallel Computing*, Vol. 19, pp. 621-632, 1993.
35. A. N. Choudhary, “Parallel I/O Systems,” *Journal of Parallel and Distributed Computing*, Jan. 1993., Vol 17, pp. 1-3.
36. W. Song, S. Hariri and A. Choudhary, “Design and Analysis of an Optical Communications Processor,” *Journal of Parallel and Distributed Computing*, Vol. 17, 1993, pp. 222-229.
37. A. Choudhary, G. Fox, S. Ranka, S. Hiranandani, K. Kennedy, C. Koelbel, and J. Saltz, “Software Support for Irregular and Loosely Synchronous Problems,” *International Journal of Computing Systems in Engineering*, Dec. 1992, Vol. 3, pp. 43-52.
38. A. N. Choudhary and S. Ranka, “Mesh and Pyramid Algorithms for Iconic Indexing,” *Journal of Pattern Recognition*, Sept. 1992, Vol. 25, No. 9, pp. 1061-1067.
39. S. Hariri, A. N. Choudhary and S. Sarikaya, “Architectural Support for Designing Fault-Tolerant Open Distributed Systems,” *IEEE Computer*, June 1992, Vol. 25, No. 6, pp. 50-62.
40. W. Song, S. Hariri and A. N. Choudhary, “An Optical Interface Message Processor for Fiber Communication Networks,” *Journal of Optics Communications*, Vol. 91, pp. 304-11, 1992.
41. A. N. Choudhary and R. Ponnusamy, “Run-Time Data Decomposition for Parallel Implementation of Image Processing and Computer Vision Tasks,” *Journal of Concurrency, Practice and Experience*, June 1992, Vol. 4, No. 4, pp. 313-334.
42. A. Choudhary and S. Ranka, “Parallel Processing for Computer Vision and Image Understanding,” *IEEE Computer*, Feb. 1992, pp. 7-11.
43. A. N. Choudhary and R. Ponnusamy, “Parallel Implementation and Evaluation of a Motion Estimation System Algorithm using Several Data Decomposition Strategies,” *Journal of Parallel and Distributed Computing*, Jan. 1992, Vol. 14, pp. 50-65.

44. A. Choudhary and S. Krishnamoorthy, “Shared Memory Multiprocessor Simulations to Study Dynamic Characteristics of Two Level Caches,” *International Journal of Computer Simulation*, Vol. 2, No. 1, Jan. 1992, pp. 1-19.
45. A. N. Choudhary and R. Ponnusamy, “Implementation and Evaluation of Hough Transform Algorithms on a Shared Memory Multiprocessor,” *Journal of Parallel and Distributed Computing*, June 1991, Vol. 12, pp. 178-188.
46. A. Choudhary, W. H. Kohler, J. Stankovic and D. F. Towsley, “A Modified Priority Based Probe Algorithm for Distributed Deadlock Detection and Resolution,” *IEEE Transactions on Software Engineering*, Vol. 15, No. 1, January 1989, pp. 10-17.

Appendix B. Refereed Conference Papers Prior to 2000

1. W. Liao, A. Choudhary, D. Wiener, and P. Varshney, “I/O Implementation and Evaluation of Parallel Pipelined STAP on High Performance Computers,” *Proc. 6th International Conference on High Performance Computing (HiPC'99)*, Calcutta, INDIA, Dec. 1999.
2. S. Periyacheri, A. Nayak, A. Jones, N. Shenoy, A. Choudhary, and P. Banerjee, “Library Functions in Reconfigurable Hardware for Matrix and Signal Processing Operations in MATLAB,” *Proc. IASTED Parallel and Distributed Computing and Systems (PDCS'99)*, Cambridge, Nov. 1999.
3. C. Srinilta and A. Choudhary, “Multi-Pool Caching in Continuous Media Server,” *Proc. Conference on Multimedia Modeling (MMM'99)*, Ottawa, CANADA, Oct. 1999, pp. 267-282.
4. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “On Reducing False Sharing While Improving Locality on Shared Memory Multiprocessors,” *Proc. 1999 International Conference on Parallel Architectures and Compilation Techniques (PACT'99)*, Newport Beach, CA, October 12-16, 1999.
5. M. Kandemir, A. Choudhary, and J. Ramanujam, “Compiler Optimizations for I/O-Intensive Computations,” *Proc. 1999 International Conference on Parallel Processing (ICPP'99)*, Aizu, JAPAN, Sept. 1999, pp. 164-171.
6. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “A Framework for Interprocedural Locality Optimization Using Both Loop and Data Layout Transformations,” *Proc. 1999 International Conference on Parallel Processing (ICPP'99)*, Aizu, JAPAN, Sept. 1999, pp. 95-102.
7. M. Kandemir, A. Choudhary, and J. Ramanujam, “I/O-Conscious Tiling for Disk-Resident Data Sets,” *Proc. 5th International Euro-Par Conference (Euro-Par'99), Parallel Processing*, Toulouse, FRANCE, August-September 1999.
8. S. Goil and A. Choudhary, “Efficient Parallel Classification Using Dimensional Aggregates,” *Proc. Workshop on Large-Scale Parallel Data Mining, KDD'99*, San Diego, CA, Aug. 1999.
9. 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,” *Proc. 8th IEEE International Symposium on High Performance Distributed Computing*, Redondo Beach, CA, Aug. 1999.
10. S. Goil and A. Choudhary, “A Parallel Scalable Infrastructure for OLAP and Data Mining,” *Proc. 1999 International Database Engineering and Applications Symposium (IDEAS'99)*, Montreal, CANADA, Aug. 1999.
11. S. Goil and A. Choudhary, “An Infrastructure for Scalable Parallel Multidimensional Analysis,” *Proc. 11th International Conference on Scientific and Statistical Database Management (SSDBM11)*, Cleveland, OH, July 1999.
12. M. Kandemir, P. Banerjee, A. Choudhary, J. Ramanujam, and E. Ayguade, “An ILP Approach for Optimizing Cache Locality,” *Proc. 1999 ACM International Conference on Supercomputing (ICS'99)*, Rhodes, GREECE, June 1999.
13. M. Kandemir, A. Choudhary, and J. Ramanujam, “Restructuring I/O-Intensive Computations for Locality,” Proc. 7th International Conference on High-Performance Computing and Networking (HPCN Europe 1999), Amsterdam, THE NETHERLANDS, April 1999.
14. J. No, J. Carretero, and A. Choudhary, “High Performance Parallel I/O Schemes for Irregular Applications on Clusters of Workstations,” *Proc. of 7th International Conference on High-Performance Computing and Networking (HPCN Europe 1999)*, Amsterdam, THE NETHERLANDS, April 1999.

15. W. Liao, A. Choudhary, D. Weiner, and P. Varshney, “Multi-Threaded Design and Implementation of Parallel Pipelined STAP on Parallel Computers with SMP Nodes,” *Proc. 13th International Parallel Processing Symposium, 10th Symposium on Parallel and Distributed Processing (IPPS/SPDP 1999)*, San Juan, PUERTO RICO, April 1999.
16. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “A Graph Based Framework to Detect Optimal Memory Layouts for Improving Data Locality,” *Proc. 13th International Parallel Processing Symposium, 10th Symposium on Parallel and Distributed Processing (IPPS/SPDP 1999)*, San Juan, PUERTO RICO, April 1999.
17. S. Goil and A. Choudhary, “Design and Implementation of a Scalable Parallel System for Multidimensional Analysis,” *Proc. 13th International Parallel Processing Symposium & 10th Symposium on Parallel and Distributed Processing (IPPS/SPDP'99)*, San Juan, PUERTO RICO, April 1999.
18. A. Choudhary and M. Kandemir, “System-Level Meta-Data for High Performance Data Management,” *Proc. Third IEEE Meta-Data Conference*, Bethesda, MD, April 1999.
19. J. Carretero, J. No, and A. Choudhary, “Optimizing I/O for Irregular Applications on Distributed-Memory Machines,” *Proc. ACPC'99*, Salzburg, Austria, Feb. 1999.
20. J. No, J. Carretero, and A. Choudhary, “Optimizations to Provide High-Performance Parallel I/O for Irregular Applications,” *Proc. Applied Informatics'99*, Innsbruck, Austria, Feb. 1999.
21. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “Improving Locality Using Loop and Data Transformations in an Integrated Framework,” *MICRO-31*, Dallas, TX, Dec. 1998.
22. S. More and A. Choudhary, “Extended Collective i/O for Efficient Retrieval of Large Objects,” *Proc. 5th International Conference on High Performance Computing (HiPC'98)*, Chennai, India, Dec. 1998.
23. S. Goil and A. Choudhary, “High Performance Multidimensional Analysis and Data Mining,” *Proc. High Performance Networking and Computing Conference (SC'98)*, Orlando, FL, Nov. 1998.
24. S. Goil and A. Choudhary, “High Performance Multidimensional Analysis of Large Datasets,” *ACM First International Workshop on Data Warehousing and OLAP (DOLAP'98)* (in conjunction with *CIKM'98*), Washington, DC, Nov. 1998.
25. M. Kandemir, A. Choudhary, J. Ramanujam, and P. Banerjee, “A Matrix-Based Approach to the Global Locality Optimization Problem,” *Proc. International Conference on Parallel Architectures and Compilation Techniques (PACT'98)*, Paris, FRANCE, Oct. 1998.
26. J. Carretero, W. Zhu, X. Shen, and A. Choudhary, “MiPFS: A Multimedia Integrated Parallel File System,” *International Joint Conference on Information Systems*, Raleigh, NC, Oct. 1998.
27. J. Carretero, J. No, and A. Choudhary, “Parallel I/O for Irregular Applications on Distributed Memory Machines,” *IX Jornadas de Paralelismo*, San Sebastian, Spain, Sept. 1998.
28. R. Krishnaiyer, Ian Foster, and A. Choudhary, “Performance of a Remote I/O Library in High-Performance Distributed Computing Environments,” *Proc. International Conference on Parallel and Distributed Computing and Systems (PDCS'98)*, Chicago, IL, Sept. 1998.
29. M. Kandemir, A. Choudhary, J. Ramanujam, N. Shenoy, and P. Banerjee, “Enhancing Spatial Locality Using Data Layout Optimizations,” *Proc. Euro-Par'98 (Workshop on Automatic Parallelisation)*, Southampton, UK, September 1998.
30. J. Carretero, A. Choudhary, J. No, and P. Chen, “Experimental Evaluation of COMPASSION: a Parallel Runtime System for Irregular Applications,” *Irregular'98*, Lawrence Livermore National Lab, Berkley, CA, Aug. 1998.
31. J. Carretero, J. No, P. Chen, and A. Choudhary, “COMPASSION: A Parallel I/O Runtime System Including Chunking and Compression for Irregular Applications,” *Proc. IRREGULAR'98*, Berkeley, CA, Aug. 1998.
32. M. Kandemir, N. Shenoy, P. Banerjee, J. Ramanujam, A. Choudhary, “Minimizing Data and Synchronization Costs in One-Way Communication,” *Proc. 1998 International Conference on Parallel Processing (ICPP'98)*, Minneapolis, MN, Aug. 1998.
33. M. A. Kandaswamy, M. Kandemir, A. Choudhary, and D. Bernholdt, “Performance Implications of Architectural and Software Techniques on I/O-Intensive Applications,” *Proc. 1998 International Conference on Parallel Processing (ICPP)*, Minneapolis, MN, Aug. 1998.
34. S. Park, J. No, J. Carretero, and A. Choudhary, “Software Caching in a Parallel I/O Runtime System to Support Irregular Applications,” *HPCS'98*, Alberta, Canada, 1998.

35. S. Goil and A. Choudhary, “On Scalable Parallel Computation of the Multidimensional Data Cube,” *Proc. International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'98)*, Las Vegas, NV, July 1998.
36. M. Kandemir, A. Choudhary, N. Shenoy, P. Banerjee, and J. Ramanujam, “A Hyperplane Based Approach for Optimizing Spatial Locality in Loop Nests,” *Proc. 1998 ACM International Conference on Supercomputing (ICS)*, Melbourne, Australia, July 1998.
37. M. Kandemir, A. Choudhary, and J. Ramanujam, “Improving Locality in Out-of-Core Computations Using Data Layout Transformation,” *Proc. 4th Workshop on Languages, Compilers, and Run-Time Systems for Scalable Computers (LCR)*, Pittsburgh, PA, May 1998.
38. J. Carretero, J. No, S. Park, P. Chen, and A. Choudhary, “COMPASSION: A Parallel I/O Runtime System Including Chunking and Compression for Irregular Applications,” *Proc. HPCN'98*, Amsterdam, THE NETHERLANDS, April 1998.
39. A. Choudhary, W. Liao, D. Weiner, P. Varshney, R. Linderman, and M. Linderman, “Design, Implementation and Evaluation of Parallel Pipelined STAP on Parallel Computers,” *Proc. International Parallel Processing Symposium (IPPS/SPDP'98)*, Orlando, FL, Mar.-Apr. 1998.
40. M. Kandemir, P. Banerjee, A. Choudhary, J. Ramanujam, and N. Shenoy, “A Generalized Framework for Global Communication Optimization,” *Proc. International Parallel Processing Symposium (IPPS/SPDP'98)*, pp. 69-73, Orlando, FL, Mar.-Apr. 1998.
41. J. No, S. Park, J. Carretero, P. Chen, and A. Choudhary, “Design and Implementation of a Parallel I/O Runtime System for Irregular Applications,” *Proc. 12th International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP'98)*, Orlando, FL, March-April, 1998.
42. S. Goil and A. Choudhary, “High Performance Data Mining Using Data Cubes on Parallel Computers,” *Proc. 12th International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP'98)*, Orlando, March-April, 1998.
43. C. Srinilta and A. Choudhary, “Performance Enhancement Using Intraserver Caching in a Continuous Media Server,” *Proc. 8th International Workshop on Research Issues in Data Engineering: Continuous-Media Databases and Applications (RIDE'98)*, Orlando, FL, Feb. 1998.
44. S. Goil and A. Choudhary, “Parallel Data Cube Construction for High Performance On-Line Analytical Processing,” *Proc. 4th International Conference on High Performance Computing (HiPC'97)*, Bangalore, India, Dec. 1997.
45. M. Kandemir, M. Kandaswamy, and A. Choudhary, “Global I/O Optimizations for Out-of-Core Computations,” *Proc. High-Performance Computing Conference (HiPC'97)*, Bangalore, India, Dec. 1997.
46. J. No and A. Choudhary, “Techniques to Provide Run-Time Support for Solving Irregular Problems,” *Proc. ICPADS'97*, Seoul, Korea, Dec. 1997.
47. M. Kandemir, A. Choudhary, J. Ramanujam, and M. Kandaswamy, “A Unified Compiler Algorithm for Optimizing Locality, Parallelism and Communication in Out-of-Core Computations,” *Proc. 5th Workshop on I/O in Parallel and Distributed Systems (IOPADS'97)*, pp. 79-92, San Jose, CA, Nov. 1997.
48. M. Kandemir, J. Ramanujam, and A. Choudhary, “Compiler Algorithms for Optimizing Locality and Parallelism on Shared and Distributed Memory Machines,”*Proc. International Conference on Parallel Architectures and Compilation Techniques (PACT'97)*, pp. 236-247, San Francisco, CA, Nov. 1997.
49. M. Kandaswamy, M. Kandemir, A. Choudhary, and D. Bernholdt, “Optimization and Evaluation of Hartree-Fock Application's I/O With PASSION,” *Proc. SC'97 Conference (formerly Supercomputing 1997)(SC '97)*, Nov. 1997
50. J. No and A. Choudhary, “Runtime Library for Parallel I/O for Irregular Applications,” *Proc. PARCO'97*, Bonn, Germany, pp. 437-440, Sept. 1997.
51. M. Kandemir, J. Ramanujam, and A. Choudhary, “Improving the Performance of Out-of-Core Computations,” *Proc. International Conference on Parallel Processing*, pp. 128-136, Bloomingdale, IL, August 1997.
52. S. Chaudhry and A. Choudhary, “Time Dependent Priority Scheduling for Guaranteed QOS Systems,” *Proc. 6th International Conference on Computer Communications and Networks*, Las Vegas, NV, Sept. 1997.

53. M. Kandemir, J. Ramanujam, and A. Choudhary, “Optimizing Out-of-Core Computations Using Chain Vectors,” *Proc. Euro-Par'97, Workshop on Parallel Languages,, Lecture Notes in Computer Science*,Volume 1300, pp. 601-608, Passau, Germany, Aug. 1997.
54. M. Kandemir, J. Ramanujam, and A. Choudhary, “A Compiler Algorithm for Optimizing Locality in Loop Nests,” *Proc. 11th ACM Intl. Conference in Supercomputing*, pp. 269-278, Vienna, Austria, July 1997.
55. M. Kandemir, A. Ramanujam, and R. Bordawekar, “Optimizing Out-of-Core Computations in Uniprocessors,” *Proc. Workshop on Architecture - Compiler Interaction, 3rd HPCA*, San Antonio, TX, Feb. 1997, also published in *Newsletter of the Technical Committee on Computer Architecture (TCCA)*, pp. 25-27, June 1997.
56. C. Srinilta, D. Jadav, and A. Choudhary, “Design and Evaluation of a Data Storage and Retrieval Strategies in a Distributed Memory Continuous Media Server,” *Proc. International Parallel Processing Symposium*, Geneva, Switzerland, April 1997.
57. S. More, A. Choudhary, I. Foster, and M. Q. Xu, “MTIO A Multi-Threaded Parallel I/O System,” *Proc. International Parallel Processing Symposium*, Geneva, Switzerland, April 1997.
58. M. Kandemir, R. Bordawekar, and A. Choudhary, “Data Access Reorganizations in Compiling Out-of-Core Data Parallel Programs on Distributed Memory Machines,” *Proc. Intl. Parallel Processing Symposium*, Geneva, Switzerland, April 1997.
59. M. Kandemir, A. Choudhary, and R. Bordawekar, “I/O Optimizations for Compiling Out-of-Core Programs on Distributed-Memory Machines,,” *Proc. PP'97 (Eighth SIAM Conference on Parallel Processing for Scientific Computing)*, Minneapolis, MN, March 1997.
60. M. Kandemir, R. Bordawekar, A. Choudhary and J. Ramanujam, “A Unified Tiling Approach for Out-of-Core Computations,” *Proc. Sixth International Workshop on Compilers for Parallel Computers (CPC'96)*, pp. 323-334, Aachen, Germany, Dec. 1996.
61. I. Foster, D. R. Kohr, Jr., R. Krishnaiyer and A. Choudhary, “Communicating Data Parallel-Tasks: An MPI Library for HPF,” *Proc. International Conference on High Performance Computing*, Trivandrum, India, 1996.
62. I. Foster, D. R. Kohr, Jr., R. Krishnaiyer, and A. Choudhary, “Double Standards: Bringing Task Parallelism to HPF via the Message Passing Interface,” *Proc. of Supercomputing '96*.
63. R. Bordawekar, A. Choudhary, and J. Ramanujam, “A Framework for Integrated Communication and I/O Placement,” *Proc. of Euro-Par'96 (Parallel Processing)*, Lyon, France, Aug. 1996.
64. C. Tumuluri and A. Choudhary, “Scalable Software Latency Hiding Schemes: Evaluation of the Poststore and Prefetch Options,” *Proc. of Euro-Par'96 (Parallel Processing)* Lyon, France, Aug. 1996.
65. C. Tumuluri, C. K. Mohan, and A. Choudhary, “GST Networks: Learning Emergent Spatio-Temperal Correlations,” *Proc. International Conference on Neural Networks*, June 1996.
66. M. Arunachalam, A. Choudhary, and B. Rullman, “Implementation and Evaluation of Prefetching in the Intel Paragon Parallel File System,” *IPPS,* April 1996.
67. R. Bordawekar and A. Choudhary, “Communication Strategies for Out-of-Core Programs on Distributed Memory Machines,” *Proc. IEEE International Conference on High Performance Computing*, New Delhi, India, Dec. 1995.
68. D. Jadav, C. Srinilta and A. Choudhary, “Input/Output Scheduling Tradeoffs in a High Performance Media-on-Demand Server,” *Proc. 2nd IEEE International Conference on High Performance Computing*, New Delhi, India, Dec. 1995.
69. S. Chaudhry, M. Razziudin and A. Choudhary, “On Guaranteed Bandwidth Channels, *International Conference on Network Protocols*, Tokyo, Japan, Nov. 1995.
70. R. Bordawekar and A. Choudhary, “Communication Strategies for Out-of-Core Parallel Programs,” *International Conference on Supercomputing*, Barcelona, Spain, July, 1995.
71. R. Bordawekar, A. Choudhary, K. Kennedy, C. Koelbel and M. Paleczny, “A Model and Compilation Strategy for Out-of-Core Data Parallel Programs,” *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, July 1995.
72. D. Jadav, C. Srinilta, A. Choudhary and P. B. Berra, “Design and Evaluation of Data Access Strategies in a High Performance Multimedia-on Demand Server,” *International Conference on Multimedia Computing and Systems*, May 1995.
73. M. Arunachalam, A. Choudhary and B. Rullman, “A Prefetching Prototype for the Parallel File System on the Paragon,” (poster), *SIGMETRICS 1995*, Ottawa, Canada, May 1995.

74. M. Harry, J. M. del Rosario and A. Choudhary, “VIP-FS: A Virtual Parallel File System for High-Performance Parallel and Distributed Computing,” *International Parallel Processing Symposium*, April 1995.
75. B. Avalani, A. Choudhary, I. Foster and R. Krishnaiyer, “Integrating Task and Data Parallelism Using Parallel I/O Techniques,” *International Workshop on Parallel Processing*, Dec. 1994.
76. R. Thakur, R. Bordawekar and A. Choudhary, “Compiler and Runtime Support for Out-of-Core HPF Programs,” *Proc. of International Conference on Supercomputing*, Manchester, England, July 1994.
77. R. Thakur and A. Choudhary, “All-to-All Communication on Meshes with Wormhole Routing,” *Proc. of Int. Parallel Processing Symposium 1994*.
78. R. Thakur, A. Choudhary and G. Fox, “Runtime Array Redistribution in HPF programs,” *Proc. of Scalable High Performance Computing Conference*, May 1994.
79. I. Foster, B. Avalani, A. Choudhary and M. Xu, “A Compilation System that Integrates High Performance Fortran and Fortran M,” *Proc. of Scalable High Performance Computing Conference 94*.
80. R. Thakur and A. Choudhary, “Complete Exchange on a Wormhole Routed Mesh,” *Proc. of MASCOTS '94*, January 1994.
81. Z. Bozkus, A. Choudhary, G. C. Fox, T. Haupt, and S. Ranka, “Compiling Distribution Directives in a Fortran 90D Compiler,” *Fifth IEEE Symposium on Parallel and Distributed Processing*, Dec. 1993.
82. Z. Bozkus, A. Choudhary, G. C. Fox, and S. Ranka, “Performance Comparison of the CM-5 and Intel Touchstone Delta for Data-Parallel Operations,” *Fifth IEEE Symposium on Parallel and Distributed Processing*, Dec. 1993.
83. R. Bordawekar, A. Choudhary and J. M. del Rosario, “Design and Evaluation of Primitives for Parallel I/O,” *Supercomputing '93*, November 1993, Portland, OR. (**Received best paper in the category of Systems**)
84. A. Choudhary, C. Koelbel, and M. Zosel, “High Performance Fortran: Implementor and Users Workshop,” *Proc. of Supercomputing'93*, Portland, OR, Nov. 1993.
85. Z. Bozkus, A. Choudhary, G. C. Fox, T. Haupt, and S. Ranka, “Fortran 90D/HPF Compiler for Distributed Memory MIMD Computers: Design, Implementation, and Performance Results,” *Supercomputing '93*, Portland, OR, November 1993.
86. R. Bordawekar, J. del Rosario, and A. Choudhary, “Design and Evaluation of Primitives for Parallel I/O,” *Proc. Supercomputing'93*, Portland, OR, November 1993.
87. R. Ponnusamy, J. Saltz and A. Choudhary, “Runtime Compilation Techniques for Data Partitioning and Communication Schedule Reuse,” *Supercomputing '93*, Portland, OR, November 1993.
88. S. Costicoglou, M. Podgorny and A. Choudhary, “On Benchmarking the Oracle Parallel Server on nCUBE2,” *IEEE Workshop on Advances in Parallel and Distributed Systems*, Oct. 1993.
89. Z. Bozkus, A. Choudhary, G. C. Fox, T. Haupt, and S. Ranka, “A Compilation Approach for Fortran 90D/HPF Compilers on Distributed Memory MIMD Computers,” *Sixth Annual Languages and Compilers for Parallelism Workshop*, Portland, OR, August 1993.
90. R. Bordawekar, A. Choudhary and J. M. del Rosario, “An Experimental Performance Evaluation of Touchstone Delta Concurrent File System,” *International Conference on Supercomputing*, Tokyo, July 1993.
91. R. Ponnusamy, N. Mansour, and A. Choudhary, “Graph Contraction for Physical Optimization Methods: A Quality-Cost Tradeoff for Mapping Data on Parallel Computers,” *International Conference on Supercomputing*, Tokyo, July 1993.
92. E. Suh, B. K. Lee, R. Martino, B. Narahari and A. Choudhary, “Parallel Computation of Solvent Accessible Surface Area of Protein Molecules,” *Seventh Parallel Processing Symposium*, April 1993.
93. R. Ponnusamy, N. Mansour, A. Choudhary and G. Fox, “Mapping Realistic Data Sets on Parallel Computers,” *Seventh Parallel Processing Symposium*, April 1993.
94. A. N. Choudhary, G. Fox, S. Hiranandani, K. Kennedy, C. Koelbel, S. Ranka and C. W. Tseng, “Compiling Fortran 77D and 90D for MIMD Distributed-Memory Machines,” *Frontiers'92: The Fourth Symposium on the Frontiers of Massively Parallel Computation*, Oct. 1992.
95. R. Ponnusamy, R. Thakur, A. Choudhary, and G. Fox, “Scheduling Regular and Irregular Communications Patterns on the CM-5,” *Supercomputing '92*, **received the best student paper award in “Performance Measurements”**.

96. R. Ponnusamy, A. Choudhary and G. Fox, “Communication Overhead on CM5: An Experimental Performance Evaluation,” *Frontiers '92: The Fourth Symposium on the Frontiers of Massively Parallel Computation*, Oct. 1992.
97. A. Choudhary and S. Krishnamoorthy, “ An Evaluation of Set-Associativity in Multi-Level Caches for Shared Memory Multiprocessors,” *Proc. of Parallel Architectures and Languages Europe*, Paris, June 1992.
98. R. Ponnusamy, J. Saltz, C. Koelbel, and A. N. Choudhary, “A Run-time Data Mapping Scheme for Irregular Problems,” *Scalable High-performance Computing Conference*, April 1992.
99. A. Choudhary and R. Thakur, “Evaluation of Connected Component Labeling Algorithms on Shared and Distributed Memory Multiprocessors,” *Sixth Parallel Processing Symposium*, March 1992.
100. A. Choudhary and S. Ranka, “Iconic Indexing on Meshes and Pyramids,” *ACM International Conference on Supercomputing*, Cologne, Germany, June 17-21, 1991.
101. A. Choudhary and R. Ponnusamy, “Performance Evaluation for Hough Transform on a Shared Memory Multiprocessor System,” *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Maui, Hawaii, June 3-6, 1991.
102. A. N. Choudhary and R. Ponnusamy, “Guided Scheduling Schemes for Image Understanding Tasks for Shared and Distributed Memory Multiprocessors,” *5th International Parallel Processing Symposium*, pp. 260-264, April, 1991.
103. A. Choudhary, S. Hariri, W. Song, Partha Banerjee and S. Ranka, “Optical Switching and Routing Architectures for Fiber-Optic Communication Networks,” *15th Annual Conference on Local Computer Networks*, Minneapolis, MN, Oct. 1990.
104. A. Choudhary and J. H. Patel, “Performance Evaluation of Clusters of NETRA: An Architecture for Computer Vision Systems,” *1990 International Conference on Parallel Processing*, St. Charles, IL, Aug. 1990.
105. A. Choudhary, M. K. Leung, T. S. Huang and J. H. Patel, “Parallel Implementation and Evaluation of Motion Estimation System Algorithms on a Distributed Memory Multiprocessor Using Knowledge Based Mappings,” *10th International Conference on Pattern Recognition*, Atlantic City, NJ, June 17-21, 1990.
106. A. Choudhary, S. Das, J. H. Patel and N. Ahuja, “A Reconfigurable and Hierarchical Parallel Processing Architecture: Performance Results for Stereo Vision,” *10th International Conference on Pattern Recognition*, Atlantic City, NJ, June 17-21, 1990.
107. A. Choudhary, “Performance of Vision Algorithm on Multiple Clusters of NETRA,” *Fourth Annual Parallel Processing Symposium*, Fullerton, CA, April 1990.
108. A. N. Choudhary, “Cost of Distributed Deadlock Detection : A Performance Study,” *Sixth International Conference on Data Engineering* (nominated for best paper), Los Angeles, CA, February 1990.
109. A. N. Choudhary and J. H. Patel, “Load Balancing and Task Decomposition Techniques for Parallel Implementation of Integrated Vision Systems Algorithms,” *Supercomputing 89*, Reno, Nevada, November 14-19, 1989.
110. A. Choudhary, M. K. Leung, T. S. Huang and J. H. Patel, “Point Matching in a Time Sequence of Stereo Image Pairs and Its Parallel Implementation on a Multiprocessor,” *IEEE Workshop on Visual Motion*, Irvine, CA, March 1989, pp. 321-328.
111. A. Choudhary, S. Das, J. H. Patel and N. Ahuja, “Surface Reconstruction from Stereo Images: An Implementation on a Hypercube Multiprocessor,” *The Fourth Conference on Hypercube Concurrent Computers and Applications*, Monterrey, CA, March 1989.
112. A. Choudhary and J. H. Patel, “A Parallel Processing Architecture for an Integrated Vision System,” *17th Annual International Conference on Parallel Processing*, Aug. 1988, pp. 383-388.
113. A. N. Choudhary, W. H. Kohler, J. Stankovic and D. F. Towsley, “A Priority based Probe Algorithm for Distributed Deadlock Detection and Resolution,” *7th International Conference on Distributed Computing Systems*, W. Berlin, W. Germany, pp. 162-169, September, 1987.

Appendix C. Invited or Keynote talks prior to 2003.

- Kellogg School of Management, “Lessons Learned in CRM
  Panel on Nov 14, 2002.
- Sandia National Laboratories, “Scientific Data Management”, Aug 2002.

- Invited Speaker, University of Paris, Dauphine, “Scalable Multidimensional Data Mining,” September 2001.
- Invited Speaker, University of Madrid, “Scalable Data Mining,” March 2002.
- Invited Speaker, DOE Salishan Conference, “High-Performance Data Mining and Analysis for Scientific Discoveries,” 2001.
- Invited speaker,.”Power-Aware compilation and Synthesis for Heterogeneous Embedded Systems,” NASA Jet Propulsion Laboratory, Pasadena, CA, June 2001.
- Keynote Speaker, Conference on High Speed Computing, “Scalable Scientific Data Management,” Gleneden Beach, OR, April 2001.
- “Scalable Data management for Large-Scale Scientific Computing,” CALTECH, Feb, 1999.
- “Recent Results in High-Performance I/O,” Sandia National Laboratory, Albuquerque, NM, Aug. 1998.
- “High-Performance Data Mining,” Technical Strategies to Beat Your Competition by the Year 2000, New York, NY, Oct. 1997.
- “On Languages and Libraries,” Panel at ICPP '97, Aug. 1997.
- “PASSION: Optimized Parallel I/O for High-Performance Computers,” Lawrence Livermore National Laboratories, Livermore, CA, Aug. 1997.
- “PASSION: Optimized Parallel I/O for High-Performance Computers,” University of Alberta, Edmonton, Canada, Aug. 1997.
- “The Design of the PASSION System,” Workshop on High-Performance Fortran, Vienna, July 1997.
- “Runtime Systems for I/O in Irregular Applications on HPCs,” Sandia National Labs, June 1997.
- “Where Should Collective I/O Be Performed? Runtime Systems or File Systems,” SIAM, Parallel Processing Conference, March 1997.
- “High-Performance On Line Analytical Processing,” The George Washington University, March 1997.
- “PASSION Runtime System,” Sandia National Laboratories, Albuquerque, NM, May 1996.
- “PASSION Runtime System,” University of Minnesota, Jan. 1996.
- “Models, Language and Compiler Support for Input-Output for I/O Intensive Applications,” SPDP Workshop on Models and Strategies for Input-Output, Oct. 1995.
- “Increasing stream-capacity of a Multimedia-on-Demand Server Using Dynamic Buffer Allocation Strategies,” Intel SSD, August, 1995.
- “The PASSION System,” IBM, Kingston, SP-2 Division, May 1995.
- “Runtime Support for Parallel I/O,” UCLA, April 1995.
- “Runtime Support for Parallel I/O,” UCSD, April 1995.
- “The PASSION System,” Carnegie-Mellon University, March 1995.
- “Design of the PASSION System,” Intel SSD, Feb. 1995.
- “Scheduling Input-Output in High-Performance Servers for Multimedia-on-Demand,” CASE Center Fall Workshop, Minnowbrook, NY, Sept. 1994.
- “Scheduling Input-Output in High-Performance Servers for Multimedia-on-Demand,” Intel SSD, Beaverton, OR, Aug 1994.
- “Scheduling Input-Output in High-Performance Servers for Multimedia-on-Demand,” IBM, Poughkeepsie, NY, July 1994.
- “Software Support for High-Performance Computers,” IBM T. J. Watson Research Center, Aug. 1994.
- “Compiler and Runtime Support for Parallel I/O,” University of Vienna, July 1994.
- “Integrating Task and Data Parallelism,” University of Vienna, July 1994.
- “Compiling Out-of-Core Problems on Distributed Memory Machines,” IBM, T. J. Watson Research Center, Feb. 1994.
- “Compiler and Runtime Support for Parallel I/O,” Intel SSD, Feb 1994.
- “Compiling F90D/HPF on Distributed Memory Machines,” CDAC, India, March 1994.
- “Integrating Task and Data Parallelism,” Caltech, Jan. 1994.
- “Runtime Support for Parallel I/O,” Caltech, Jan. 1994.
- “A Fortran 90D/HPF Compiler for Distributed Memory Machines,” Intel SSD, Beaverton, OR, July 16, 1993.

- "A Fortran 90D/HPF Compiler for Distributed Memory Machines," IBM Almaden Research Center, San Jose, CA, June 4, 1993.
- "Parallel Architectures and Software : State-of-the-art and Future Directions," Dual-Use Technology Conference, SUNY, Utica, May 25, 1993.
- "Compiling For Distributed Memory Machines," National Institute of Health, Bethesda, MD, May 19, 1993.
- "Runtime and Compiler Support for Implementing Programs on Distributed Memory Parallel Computers," National Institute of Health, Bethesda, MD, May 19, 1993.
- "A Fortran 90D/HPF Compiler for Distributed Memory Machines," IBM Yorktown Heights, April 2, 1993.
- "A Fortran 90D/HPF Compiler for Distributed Memory Machines," CALTECH, Feb. 9, 1993.
- "Runtime Support for Parallel I/O," CRPC, Rice University, Jan. 26, 1993.
- "Scalable Parallel I/O," CALTECH, Jan. 19, 1993.

Appendix D. Students and Postdocs

M.S. Students (Graduated) – ***Note: most of the graduated PhD students also obtained M.S. degree with thesis)*:**

1. Sriram Raghavendran, M.S., 1998
2. Chutinet Srinilta, M.S. (Thesis), 1995
3. Sivaram Kudatipidi, M.S. (Thesis), 1996
4. Sachin More, M.S. (Thesis), 1996
5. Balaji Thiagrajan, M.S., 1995
6. Bhaven Avalani, M.S., 1995
7. Rajesh Bordawekar, M.S. (Thesis),1994
8. Kohinoor Basu (Intel), 2001
9. Debrarata Bagchi, 2001
10. Satrajit Pal, 2001
11. Gokhan Memik, 2000
12. Jay Pisharath, 2002
13. Joe Zambreno, 2002
14. Chetan Kumar (2008)

***Current Ph.D. Students*:**

1. Ms. Kathy Lee
2. Dan Honbo (NSF Fellowship)
3. Ms. Chen Jin
4. Ms. Diana Palsetia
5. Kunpeng Zhang
6. Yu Cheng
7. Bharath Pattabiraman
8. Yves Xie
9. Esteban Rangel
10. Reda Al-Bahrani

Current Postdocs/Research Faculty

- Ankit Agrawal
- William Hendrix
- Wei-keng Liao
- Mostofa Patwary
- Ms. Saba Sehrish

- Seung Woo Son
- Zhengzhang Chen

Appendix E. Past Professional activities

- Program VIce-Chair, International Conference on High-Performance Computing, Dec. 1999.
- General Chair, Fifth International Workshop on I/O in Parallel and Distributed Systems (IOPADS), Nov. 1997, San Jose, CA.
- Associate Editor of *IEEE Transactions on Parallel and Distributed Systems*
- Subject Area Editor of *Journal of Parallel and Distributed Computing* .
- Program Chair, Fourth International Workshop on I/O in Parallel and Distributed Systems (IOPADS), with FCRC, May 1996, Philadelphia.
- Program Co-Chair, International Conference on Parallel Processing, 1993.
- Tutorials Chair, International Parallel Processing Symposium, 1994.
- Program Committees
    - HiperIO, 2006
    - SSDBM 2004
    - ISPAN 2005
    - ISPA 2004
    - IEEE Conference on Web Services, 2003
    - SDM04
    - ICPP 2004
    - ESMc Conference, Naples on 27 - 29 October 2003.
    - DMKD03
    - HCW 03
    - HPDM03
    - SDM03
    - IPSN'04
    - HPCA-8, 2002
    - ACM Data Mining, 2002
    - IPDPS 2002, High Performance Computational Biology(December 2001),
    - INTERACT-6 (2002),
    - SDM'2001 (Scientific Data Mining),
    - Sixth International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'02),
    - 5th International High Performance Data Mining Workshop (April 2002),
    - WDAS-2002
    - IEEE Data Mining.
    - Program Committee, 4th International Workshop, CANPC 2000, Toulouse, FRANCE, January 2000.
    - *High-Performance Computer Architectures, 1996.*
    - *International Parallel Processing Symposium, 1996, 1997, 1998.*
    - *International Conference on Multimedia Computing Systems, 1995.*
    - *International Conference on Distributed Computing Systems, 1995.*
    - *International Conference on High-Performance Computing, 1995.*
    - *Symposium on Frontiers of MPPs, 1995.*
    - *International Symposium on Computer Architecture, 94.*
    - *International Conference on Computer Vision and Pattern Recognition,* 1992.
    - *First International Conference on High-Performance Distributed Computing, 1992.*
    - *Minnowbrook Workshop on Software Engineering for Parallel Computers,* July, 1991.
- Member of the Association of Computing Machinery, Member of the IEEE Computer Society.
- Chair, Panel on Software Tools for Parallel Computers, *Minnowbrook Workshop on Software Engineering for Parallel Computers,* July, 1991.
- Session Chair, *Fifth International Symposium on Parallel Processing,* April-May, 1991.
- Session Chair, *Sixth International Parallel Processing Symposium,* March, 1991.

- Member of the High-Performance Fortran Forum, An Industry, Academia and Government Committee to define standard for High-Performance Fortran
- Served as a Referee for: International Symposium on Computer Architecture, International Conference on Parallel Processing, IEEE Transactions on Computers, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, International Conference on Distributed Computing Systems, IEEE Computer Magazine, IEEE Transactions on ASSP, Journal of Computer Vision, Graphics and Image Processing, Conference on Fault Tolerant Computing Systems, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Knowledge and Data Engineering, Journal of Parallel and Distributed Computing, International Parallel Processing Symposium, IBM Journal of Research and Development.

Appendix F: Selected Projects and Impact

The following is a list of some recent major projects, results and their impact.

Pulse of the Tweeters: As a part of our data mining projects on massive datasets, and our algorithms for social networking and sentiment analysis at a massive scale, this project entails the development of robust software, which in real-time performs sentiment analysis of massive number of tweets from Twitter, identifies influencers and ranks them in real-time as well as identifies similar conversations and topics and clusters them. This robust and scalable software, which operates on constantly increasing data of sizes in Terabytes is available to public via www.pulseofthetweeters.com which is used by tens of thousands of people worldwide. This project has been written up in New York Times and more than 100 regular and online publications in more than 25 languages worldwide.

Parallel I/O: As a part of this research, we have developed several fundamental techniques as well as system software for scaling the I/O performance of supercomputing systems. We developed techniques such as two-phase I/O, data sieving, user-level scalable caching at runtime using global access pattern information, persistent file domains and others. This work not only influenced a lot of research by many groups world-wide in the field of scalable I/O but these techniques also form the foundation of optimizations in many widely used software for HPC such MPI-IO (distributed as part of MPICH from Argonne National Laboratories and our optimizations have been incorporated). This software or its derivation by companies is used by thousands of users worldwide. Parallel Furthermore, many techniques have been incorporated into parallel file systems and runtime software provided by companies such as IBM, Intel, Cray, Sun and others.

Parallel NetCDF (PnetCDF): Parallel-NetCDF is a library providing high-performance I/O while still maintaining file-format compatibility with Unidata's NetCDF. This is a production quality software, primarily developed at Northwestern University in collaboration with Argonne National Labs, which incorporates highly scalable optimizations for massively parallel supercomputing and corresponding storage systems. PNetCDf software package has enabled scientists and engineers worldwide to scale their applications for storage and I/O which is very critical to obtain high performance from modern supercomputers. Application domains that use PNetCDF include Climate Modeling, Weather Forecasting, Fusion, Astrophysics, Chemistry and many more applications. PnetCDF is freely made available in the public domain from http://cucis.ece.northwestern.edu/projects/PNETCDF/ .

NUMinebench: Scalable Data Analytics and Mining: This research entails the design and development of scalable data mining algorithms, software and its applications for business, security, bioinformatics, medicine and scientific knowledge discoveries. The goal of this research is the scalability, applications and use of large-scale data mining algorithms. In the recent past, we have developed scalable algorithms in high-dimensional clustering (which is used by many scientists and is also adapted into commercial products such as Oracle), hierarchical and density-based scalable clustering (used by scientists in cosmology and astrophysics), Utility associations rules mining (used to compute value-based association rules in business), and scalable classification. **Furthermore, recently, we developed a benchmark of data mining and bioinformatics algorithms, called "NU-Minebench", which is currently a collection of fifteen sequential and parallel algorithms and software** (http://cucis.ece.northwestern.edu/projects/DMS/MineBench.html). This robust software benchmark has

been downloaded more than by several thousand users worldwide and is used by industry, researchers, and academics to design and develop new architectures, new systems, and scalable algorithms. Almost all major computer and systems companies have started to use this benchmark.

MATCH - MATLAB Compiler for Reconfigurable Computing.
Digital signal processing and image processing applications are typically written in the MATLAB programming language, and are typically executed on general purpose DSP processors. However, recently DSP algorithms are being mapped onto Reconfigurable Field Programmable Gata Arrays (FPGAs) for performance and reconfigurability reasons. However, to map DSP algorithms onto FPGAs, users are required to manually translate MATLAB programs onto languages such as VHDL or Verilog. As part of the MATCH project, we developed the MATCH compiler that takes MATLAB programs and automatically parallelizes it and maps it a heterogeneous environment of off-the-shelf embedded processors, digital signal processors, and FPGAs. This project was supported by DARPA.

Our group commercialized this project by forming a Company called Accelchip. Accelchip provided products including A MATLAB compiler, a Simulink compiler and libraries to automatically synthesize reconfigurable hardware from Algorithms written in these languages onto FPGAs. Accelchip was acquired by Xilinx Inc. a few years ago.

OTHER Projects:
Along with co-inventors, Alok Choudhary designed a cache snoop filtering protocol called "Jetty", which or a derivative thereof, is used in almost all modern multicore processors (e.g., designed and built by Intel, AMD, IBM etc.). The protocol was published in a paper and was made available in the public domain.

In the mid 90s, Along with Prof. Geoffrey Fox and other collaborators, Alok Choudhary led a team that designed and prototyped an High-Performance Fortran Compiler for Supercomputers. That technology was transferred to Portland Group Inc., and has been the basis for most of the compiler products provided by PGI Inc. for supercomputers. PGI Inc was subsequently acquired by STMicroelectronics.