Parallel and Distributed Computing, 1999-2000

Externally Funded Research Projects

CHIMAERA: Architecture, Compilers, and Configuration Management for Reconfigurable Computing for Mass-Market Computing

P. Banerjee*, S. Hauck, and M. Sarrafzadeh

Sponsor: Defense Advanced Research Projects Agency

This research project seeks to develop a complete adaptive solution for general-purpose computing systems. It includes research in adaptive computer architectures, configuration-management techniques, high-level compilation, mapping algorithms, template-based physical design, software / algorithms optimized for adaptive systems, and defense applications of adaptive computing. A key specific aim is to produce a general-purpose computing paradigm simple enough to become the standard mass-market computing model, yet with enough power to enable several orders-of-magnitude speedup for defense applications. This will yield a future defense computing solution with the price and availability advantages of commercial off-the-shelf hardware.

MATCH: A MATLAB Compilation Environment for Adaptive Computing Systems

P. Banerjee* and A. N. Choudhary*

Sponsor: Defense Advanced Research Projects Agency

Adaptive computing systems constitute a new class of computing and communication technology composed of configurable hardware capable of system-level adaptation. Such systems are often built out of combinations of microprocessor-based embedded systems, specialized digital signal processors (DSP's), and field-programmable gate arrays (FPGA's). The objective of this project is to make it easier for users to develop efficient codes for adaptive computing systems. As part of this project, we are developing a compiler that allows input of a user's applications written in the high-level language, MATLAB, and generates efficient low-level code that runs on commercial off-the-shelf FPGA's, embedded processors, and DSP's. Our specific aims include:

---

*Denotes Parallel and Distributed Computing faculty member(s), listed in alphabetical order.

- development of a hardware testbed consisting of FPGA's, embedded processors, and DSP's;
- development of a basic compiler for mapping a given MATLAB application onto this heterogeneous target;
- investigation of automated parallelization and mapping techniques;
- design and support of compiler directives;
- development of library functions and applications of interest to DOD;
- development of faster algorithms for compilation.

PACT: Power-Aware Architectural and Compilation Techniques
P. Banerjee*, M. Sarrafzadeh, A. N. Choudhary*, A. Moshovos*, and H. Yuen
Sponsor: Defense Advanced Research Projects Agency

The objective of the PACT Project is to develop power-aware architectural techniques and associated compiler and CAD tool support. Specific goals are:

- Develop novel architectural and compiler concepts at various levels that can reduce the total energy consumption in specific applications by factors of 10-100X over conventional, non-power-aware architectures.
- Develop compiler techniques to automate the process of generating efficient code that is within a factor of two of the best manual approach with respect to optimizing power under performance and resource constraints.
- Demonstrate the usefulness of the compiler and architectural concepts on some real applications.

PANTHER: A High-Performance Distributed-Computing Infrastructure
P. Banerjee*, A. N. Choudhary*, M. Sarrafzadeh, P. Scheuermann*, and V. Taylor*
Sponsor: National Science Foundation

Specific aims of this project include:

- Explore using high-speed networking and computing to investigate file-systems and data-management issues for high-performance distributed computing.
- Investigate the parallel-programming support of networks of high-speed workstations and personal computers as an alternative to stand-alone parallel computers.
- Investigate high-performance CAD of electronic systems in a heterogeneous environment.
- Develop a Web-based CAD computing center that takes advantage of high-speed networking.
- Explore new instructional techniques that take advantage of high bandwidth and high speed.

This project provides for the acquisition of 50 high-performance Hewlett-Packard C-180 UNIX workstations, 20 medium-performance B-132 UNIX workstations; 3 J-280 UNIX fileservers, an 8-processor Silicon Graphics Origin 2000 distributed shared-memory multiprocessor, and four CISCO systems LS1010 ATM switches. All of the machines are connected via an OC-3 ATM network operating at 155 Mbps.

Compiler and Runtime Optimization Techniques for Parallel Programming
A. N. Choudhary*
Sponsor: National Science Foundation (Young Investigator Award)

This project develops compiler and run-time optimization techniques for scalable parallel programming. In particular, the objectives are to develop compiler techniques for parallel I/O and

locality optimizations, and run-time techniques for data-redistribution and memory-hierarchy optimizations. Further, this research involves the development of fundamental models and compilation techniques for out-of-core computations.

High-Performance Data-Management, Access, and Storage Techniques for Tera-Scale Scientific Applications

A. Choudhary*, P. Banerjee*, and V. Taylor*

Sponsor: Department of Energy ASCI Level-2 Grant

Emerging large-scale scientific experiments and simulations require the storage, management, efficient access, and analysis of hundreds of gigabytes to hundreds of terabytes of data. Current data-management and analysis techniques do not satisfy these needs in term of performance, scalability, ease of use, and interfaces. This project is developing a scalable, high-performance, data-management system (SHPDM) to provide support for data management, query capability, and high-performance accesses to large datasets stored in hierarchical storage systems such as HPSS. SHPDM will provide the flexibility of conventional databases for indexing, searches, management of objects, and creating and keeping histories and trails of data accesses. It will also provide high-performance access methods and optimizations (pooled striping, prefetching, caching, collective I/O) for accessing large-scale data-objects found in scientific computations.

Interoperable Data Files for High-Performance Computing

A. N. Choudhary*

Sponsor: National Science Foundation / University of Wyoming

This project is developing data-management techniques that enhance both interoperability and high-performance parallel access to scientific data. Specific objectives include the development of software support for: (1) “portable” data files that provide user-defined abstractions; (2) automated filtering and conversion techniques capable of extracting the “meaning” of a data file and presenting it in a form compatible with a host application; (3) improved access techniques that permit parallel access to the files while preserving the basic abstractions; (4) run-time techniques that incorporate collective I/O, data reuse, and prefetching strategies; and (5) legacy files.

Modeling and Evaluation of I/O Architecture in Servers

A. N. Choudhary*

Sponsor: Intel Corporation

This project is developing an object-oriented simulation environment to study various architectural configurations for servers. We have developed detailed server-workload models from typical application domains such as OLTP, video-on-demand, decision support etc. These models are used to study the performance, reliability, and availability of server features. While the primary focus is to study different approaches for the server’s I/O subsystem, the modeling is general enough to study other server components including the processing nodes and the communications network.

Automated Reasoning Tools for Verification

L. J. Henschen*

Sponsor: Motorola Center for Communications

A verification system for telephony software is being developed based on the ACL2 system. As ACL2 is both a formal language for representing assertions about software and an executable language, our system can both prove results about the software being modeled and execute that software. Our system will have a library of techniques for telephony software verification. A user interface will allow arbitrary mixing of ACL’s descriptions of functions with actual C implementations of those functions to form an executable module. The verified ACL2 representation will serve as an oracle for generating correct output against which to compare C-module results.

Challenges in CISE: Metacomputing Environments for Optimization (MetaNEOS)
J. Nocedal*
Sponsor: National Science Foundation

The metacomputing paradigm has emerged as an economical way to harness the power of large, distributed collections of computers, making use of compute cycles that would otherwise be wasted. The goal of this project is to use metacomputing platforms to enable the solution of very large optimization problems that arise in science, engineering, and economics. By combining distributed hardware, network hardware / software infrastructure, and optimization algorithms and modeling tools, we aim to produce environments powerful enough to tackle optimization problems of unprecedented size and complexity.

Large Scale Nonlinear Programming
J. Nocedal*
Sponsor: National Science Foundation

Our goals in this project are threefold. First, we are developing new techniques for solving large nonlinear programming problems. The new algorithms are being designed to be robust even for highly ill-conditioned problems, and are being tailored to very large problems. Second, we are developing a new method for solving problems in which derivatives are not available. Finally, we are creating a NEOS server for mixed-integer nonlinear programming (MINLP).

Nonlinear Programming: Algorithms and Software Environments
J. Nocedal*
Sponsor: U.S. Department of Energy

Projects here range from the design of new algorithms and the study of their convergence properties to the development of production-level software. One of the main goals is to devise new interior-point methods for large-scale nonlinear programming. Much attention will also be given to the study of new automatic preconditioners for the conjugate gradient method, and to their application in both optimization calculations and finite-element models.

Prototype System for Transparent Replication of HTTP Service
P. Scheuermann*
Sponsor: Usenix Association

This project deals with a number of performance-optimization issues applied to a client / server architecture based on smart clients which are dynamically downloaded as mobile code into user applications. We are implementing a mechanism called batch resource transmission that reduces the HTTP overhead by sending not only the HTML resource requested by the client, but also the resources embedded in the HTML resource that are not cached by the client. We are investigating the use of checksums as a mechanism for identifying embedded resources and for validating the content of the client's cache.

Putting Log Data to Work: Mass Storage Information Systems
P. Scheuermann*
Sponsor: National Aeronautics and Space Administration

This study is part of the Mass Storage System Performance Analysis project at NASA Goddard. The goal is to capture useful information about the activity of the mass data storage and delivery system from system logs. Currently, these logs are large, ill-structured, and inaccessible for querying. Therefore, the overall performance of the mass-storage system, as well as the statistics about user access patterns, cannot be easily determined. We are investigating the use of data-warehousing and data-mining techniques to obtain easy access to summaries of the log data, as well as to search for interesting patterns of user activity in the storage system.

Scalable and Adaptive File Management in Networks of Workstations

P. Scheuermann*

Sponsor: National Aeronautics and Space Administration

This project is developing a prototype system to provide a scalable architecture for replication on the Web. The new client / server architecture allows for user-transparent geographical replication of Web services. This system is legacy-friendly, i.e., it requires no changes to the existing Web infrastructure. The clients are downloaded as applets to commercially available browsers, while the Web servers are extended with servlets. Both server and client machines take active roles: server machines determine the resources to be replicated / migrated and the best locations for these resources; while client machines choose the best server to submit their requests.

SCALAR: Scalable Architecture for Replication on the Web

P. Scheuermann*

Sponsor: National Science Foundation

The wide use of the World Wide Web and the ease of access to the Internet are changing drastically our requirements in distributed systems. The exponentially growing scale of use of the Web is stressing the capacity of the Internet and leads to poor performance and low reliability of Web service. Replication is one of the major techniques that have been proposed to overcome these problems. This project focuses on the development of a prototype system that will provide a scalable architecture for replication on the Web. A new client/server architecture that will allow for user-transparent geographic replication is investigated. Unlike other proposed Web architectures, both clients and servers take active roles in this architecture. Server machines determine the documents to be replicated/migrated and best locations for the documents, while client machines choose the best server to submit their requests based on estimates of the servers' prior response times. Heuristic algorithms based on clustering techniques will be developed in order to determine sets of documents that be beneficially migrated/replicated as a unit.

Alliance Team B: Distributed Computing

V. Taylor*

Sponsor: National Science Foundation / National Center for Supercomputing Applications (NCSA)

This project is part of an overall NSF PACI award to NCSA, called the *Alliance*. We plan to complete two major tasks: (1) development of performance models for the cosmology application executed on distributed systems; and (2) integration of these models into the PART tool, which performs data partitioning for distributed systems. Our initial focus will be on the cosmology application because it builds upon the expertise of the researchers at Northwestern in the area of mesh-based applications.

Integration of Computer Architecture and Parallel Programming Tools Into Computer Science and Engineering

V. Taylor*, A. N. Choudhary*, J. Fortes and R. Eigenmann (Purdue University),
L. Vidal and J. J. Chen (Chicago State University)

Sponsor: National Science Foundation / Purdue University

Three institutions (Purdue, Northwestern, and Chicago State University) are involved in this project. The goal is to integrate computer architecture and parallel programming tools into computer science and engineering curricula. We do this using a network computing infrastructure called the Purdue University Network Computing Hubs (PUNCH). Existing courses are being enriched with both content and experiments not previously possible. Northwestern is involved with integrating several tools into two computer architecture courses. All three universities are developing educational materials and are installing tools for shared use through PUNCH.

Prophesy: A Hierarchical Tool for Modeling and Analyzing Parallel Scientific Applications
V. Taylor*
Sponsor: National Aeronautics and Space Administration

Currently, there exists a large gap between the theoretical peak performance and the actual achieved performance of advanced parallel computers when running large scientific simulations. This gap stems in part from an incomplete understanding of how the parallel-system features impact the performance of the applications. Detailed performance models aid in understanding the relationship between the parallel system and the application. This project studies the use of several performance models, including nonlinear models, to aid in understanding and predicting application performance.

Prophesy: A Performance Network for Developing a Computer Algebra
V. Taylor* and R. Stevens (University of Chicago and Argonne National Laboratory)
Sponsor: National Science Foundation

Currently, a large gap exists between the peak performance of high-performance systems and the achieved performance when running large scientific simulations. The large performance gap stems partly from a poor understanding of realistic, large applications and the lack of understanding of how the computer system features impact the performance of the applications. This project addresses this problem by developing a framework, called "Prophesy" that facilitates the development of very-fine to coarse-grain analytical performance models. Further, Prophesy will enable the development of an algebra that identifies how elementary performance models of system components and application algorithms should be composed to reflect the performance of the application executing on a parallel system.

Young Investigator Award
V. Taylor*
Sponsor: National Science Foundation

This grant supports the PI's research in three key areas related to the performance of parallel scientific applications: (1) development of the mesh-partitioning tool, "PART," for distributed systems; (2) investigation of methods to improve the performance of parallel shortest-path applications; and (3) development of a tool for setting up performance models.

Design of Assembly Operations for Recycling Disassembly
C. H. Wu*, D. T. Lee, and S. M. Mok, D. Hong, and T. Babin (Motorola)
Sponsor: Motorola Center for Communications

This project is developing automatic assembly operations for recycling disassembly at Motorola. We are developing a Virtual Assembly and Disassembly (VIRAD) system that can be integrated into CAD/CAM software to permit designers to evaluate products for assembly and disassembly efficiencies. The framework of the system divides a product's life cycle into three phases: CAD design, the VIRAD environment, and manufacturing and de-manufacturing processes. The processes of automatic assembly and disassembly have become key technologies in recycling for protecting our environment. By having automatic disassembly processes to de-manufacture products at a lower cost, recycling efforts can be accelerated and made profitable.

Book

J. Nocedal* and S. Wright, *Numerical Optimization*. New York: Springer-Verlag, 1999.

Book Section / Chapter

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 Computers*, S. Chatterjee et al., eds., *Lecture Notes in Computer Science*, Springer-Verlag, 1999.

Journals Edited

P. Banerjee*, Assoc. Editor, *IEEE Trans. Computers.*

P. Banerjee*, Assoc. Editor, *IEEE Trans. Parallel and Distributed Systems.*

P. Banerjee*, Assoc. Editor, *J. Parallel and Distributed Computing.*

A. N. Choudhary*, Assoc. Editor, *IEEE Trans. Parallel and Distributed Systems.*

A. N. Choudhary*, Subject Area Editor, *J. Parallel and Distributed Computing.*

L. J. Henschen*, Assoc. Editor, *J. Automated Reasoning.*

L. J. Henschen*, Assoc. Editor, *Software Engineering/Knowledge Engineering Journal.*

J. Nocedal*, Co-Editor, *Mathematical Programming Journal A.*

J. Nocedal*, Assoc. Editor, *SIAM J. Optimization.*

P. Scheuermann*, Assoc. Editor, *IEEE Trans. Knowledge and Data Engineering.*

P. Scheuermann*, Assoc. Editor, *J. Network and Computer Applications*.

P. Scheuermann*, Assoc. Editor, *J. Very Large Databases*.

Journal Papers

M. Kandemir, P. Banerjee*, A. N. Choudhary*, J. Ramanujam, and N. Shenoy*, “A global communication optimization technique based on data flow analysis and linear algebra,” *ACM Trans. Programming Languages and Systems*, vol. 21, no. 6, Nov. 1999.

A. Lain, D. Chakrabarti, and P. Banerjee*, “Compiler and run-time support for exploiting regularity within irregular applications,” *IEEE Trans. Parallel and Distributed Systems*, vol. 11, no. 2, Feb. 2000.

P. Prabhakaran and P. Banerjee*, “Parallel algorithms for force-directed scheduling of flattened and hierarchical signal flow graphs,” *IEEE Trans. Computers*, 1999.

A. N. 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, 2000, pp. 45–60.

A. N. 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 Trans. Aerospace and Electronic Systems*, vol. 36, no. 2, April 2000, pp. 528–548.

S.-G. Lee, L. J. Henschen*, J. Chun, and T. Lee, “Identifying relevant constraints for semantic query optimization,” *Information and Software Technology*, vol. 42, June 2000, pp. 899–914.

W.-S. V. Shih, W.-C. Lin*, and C.-T. Chen, “Volumetric morphologic deformation method for intersubject image registration,” *Int. J. of Imaging Systems and Technology*, vol. 11, 2000, pp. 117–124.

L. Biegler, J. Nocedal*, C. Schmidt and D. Terner, “Numerical experience with a reduced Hessian method for large-scale constrained optimization,” *Computational Optimization and Applications*, vol. 15, no. 1, 2000.

R. Byrd, M. E. Hribar and J. Nocedal*, “An interior point method for large scale nonlinear programming,” *SIAM J. Optimization*, vol. 9, no. 4, 1999, pp. 877–900.

J. L. Morales and J. Nocedal*, “Automatic preconditioning by limited memory quasi-Newton updating,” *SIAM J. Optimization*, vol. 10, no. 4, 2000, pp. 1079–1096.

P. Scheuermann *, L.-W. Lee, and R. Vingralek, “File assignment in parallel I/O systems with minimal variance of service time,” *IEEE Trans. Computers*, vol. 49, no. 2, Feb. 2000, pp. 127–140.

N. Shenoy*, Y. N. Srikant, and V. P. Bhatkar, “Automatic data partitioning by hierarchical genetic search, *J. Parallel Algorithms and Applications*, vol. 14, no. 2, 1999, pp. 119–147.

J. Chen and V. Taylor*, “ParaPART: Parallel mesh partitioning tool for distributed system,” *Concurrency: Practice and Experience*, 2000, vol. 12, pp. 111–123.

J. S. Gyorfi and C. H. Wu*, “Coordinated planning and control of automated assembly manufacturing,” *IEEE Trans. on Systems, Man, and Cybernetics*, *PART A: Systems and Humans*, vol. 30, no. 2, March 2000, pp. 173–180.

Symposium Sessions Organized / Chaired

P. Scheuermann*, Program Chair, *1999 Knowledge and Data Exchange Workshop* (KDEX’99), Chicago, IL, Nov. 1999.

P. Scheuermann*, Member Program Committee, *4th Pacific-Asia Conf. On Knowledge Discovery and Data Mining*, Kyoto, Japan, April 2000.

Invited Talks and Seminars

P. Banerjee*, “MATCH: A MATLAB Compilation Environment for Adaptive Computing Systems,” *Texas A&M University*, ECE Dept. Distinguished Lecture, March 2000.

P. Banerjee*, “PROPERCAD: Parallel Algorithms for VLSI CAD,” *Texas A&M University*, ECE Dept. Distinguished Lecture, March 2000.

P. Banerjee*, “A MATLAB Compilation Environment for Adaptive Computing Systems,” *Illinois Institute of Technology*, Chicago, IL, ECE Dept. Outstanding Lecture, April 2000.

P. Banerjee*, “A MATLAB Compilation Environment for Adaptive Computing Systems,” *University of California–Irvine*, Computer and Information Science Dept. Distinguished Lecture, June 2000.

P. Banerjee*, “A MATLAB Compilation Environment for Adaptive Computing Systems,” *University of Toronto*, ECE Dept. Distinguished Lecture, Aug. 2000.

A. N. Choudhary*, “High-Performance Data Management for Large-Scale Scientific Computing,” Lawrence-Livermore National Laboratory, Livermore, CA, Jan. 2000

J. Nocedal*, “On the Failure of Newton Iterations,” *Optimization Technology Center*, *Northwestern University*, Nov. 1999.

J. Nocedal*, “Advances in Automatic Preconditioners,” *Meteo-France*, Toulouse, France, March 20, 2000.

J. Nocedal*, “Feasibility Control in Nonlinear Optimization,” *SIAM National Meeting,* San Juan, Puerto Rico, July 12, 2000.

J. Nocedal*, “Development of the NITRO Software Package,” *Int. Symp. on Mathematical Programming*, Atlanta, Georgia, Aug. 12, 2000.

P. Scheuermann*, “A Scalable Architecture for Replication on the Web,” *8th Int. Conf. on Computer Communications and Networks*, Boston-Natick, MA, Oct. 1999.

P. Scheuermann*, “Web++ — A System for User-Transparent Replication of Resources,” *Swiss Federal Institute of Technology*, Zurich, Switzerland, July 7, 2000.

P. Scheuermann*, “Web++ — A System for User-Transparent Replication of Resources,” *Swiss Federal Institute of Technology*, Lausanne, Switzerland, July 13, 2000.

Symposium Papers

P. Banerjee*, N. Shenoy*, A. N. 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,” *Proc. Int. Symp. on FPGA Custom Computing Machines* (FCCM-2000), Napa Valley, CA, April 2000.

P. Joisha and P. Banerjee*, “Efficient computation of ownership sets in HPF,” *Proc. Languages and Compilers for Parallel Computing* (LCPC-2000), Yorktown Heights, NY, Aug. 2000.

P. Joisha, A. Kanhere, P. Banerjee*, N. Shenoy*, and A. N. Choudhary*, “Handling context-sensitive syntactic issues in the design of a front-end for a MATLAB compiler,” *Proc. ACM Array Programming Languages Conf.* (APL-Berlin-2000), Berlin, Germany, July 24–27, 2000.

V. Kim, P. Banerjee*, and K. De, “Fine-grained parallel VLSI synthesis for commercial CAD on a network of workstations,” *Proc. Int. Conf. on Parallel Processing* (ICPP-2000), Toronto, Canada, Aug. 2000.

A. Nayak, P. Banerjee*, C. Chen, and M. Sarrafzadeh*, “Power optimization issues in dual voltage design,” *Proc. Int. Conf. on Design Automation* (ICDA 2000), Beijing, China, Aug. 21–25, 2000.

Z. Ye, P. Banerjee*, S. Hauck, and A. Moshovos*, “CHIMAERA: A high-performance architecture with a tightly-coupled reconfigurable functional unit,” *Proc. 27th Int. Symp. on Computer Architecture*, Vancouver, Canada, June 10–14, 2000.

Y. Yuan and P. Banerjee*, “A parallel 3-D capacitance extraction program,” *Proc. 6th Int. Conf. on High Performance Computing* (HiPC’99), Calcutta, India, Dec. 1999.

Y. Yuan and P. Banerjee*, “A parallel implementation of a fast multipole based 3-D capacitance extraction program on distributed memory multicomputers,” *Proc. 14th Int. Parallel and Distributed Processing Symp.* (IPDPS 2000), Cancun, Mexico, May 1–5, 2000 (Best Paper Award).

M. Haldar, A. Nayak, A. Choudhary*, and P. Banerjee*, “Parallel algorithms for FPGA placement,” *Proc. Great Lakes Symp. on VLSI* (GVLSI 2000), Chicago, IL, March 2000.

M. Kandemir, A. N. Choudhary*, and J. Ramanujam, “I/O-conscious tiling for disk-resident data sets,” *Proc. 5th Int. Euro-Par Conf. on Parallel Processing* (Euro-Par’99), Toulouse, France, Aug.-Sept. 1999.

M. Kandemir, A. N. Choudhary*, and J. Ramanujam, “Compiler optimizations for I/O-intensive computations,” *Proc. 1999 Int. Conf. on Parallel Processing* (ICPP’99), Aizu, Japan, Sept. 1999, pp. 164–171.

M. Kandemir, A. N. Choudhary*, J. Ramanujam, and P. Banerjee*, “A framework for interprocedural locality optimization using both loop and data layout transformations,” *Proc. 1999 Int. Conf. on Parallel Processing* (ICPP’99), Aizu, Japan, Sept. 1999, pp. 95–102.

M. Kandemir, A. N. Choudhary*, J. Ramanujam, and P. Banerjee*, “On reducing false sharing while improving locality on shared memory multiprocessors,” *Proc. 1999 Int. Conf. on Parallel Architectures and Compilation Techniques* (PACT’99), Newport Beach, CA, Oct. 12-16, 1999.

G. Memik, M. Kandemir, and A. N. Choudhary*, “Design and evaluation of a compiler-directed I/O technique,” *Proc. European Conf. on Parallel Computing* (Euro-Par’2000), Munich, Germany, Aug. 2000.

G. Memik, M. Kandemir, and A. N. Choudhary*, “Design and evaluation of smart disk architectures for commercial workloads,” *Proc. Int. Conf. on Parallel Processing* (ICPP 2000), Toronto, Canada, Aug. 21–24, 2000.

G. Memik, M. Kandemir, A. N. Choudhary*, and V. Taylor*, “APRIL: A run-time library for tape resident data,” *Proc. 8th NASA Goddard Conf. on Mass Storage Systems and Technologies*, in cooperation with the *17th IEEE Symp. on Mass Storage Systems*, College Park, MD, March 27–30, 2000, pp. 61–75.

S. More and A. N. Choudhary*, “Tertiary storage organization for large multidimensional datasets,” *Proc. 8th NASA Goddard Space Flight Center Conf. on Mass Storage Systems and Technologies and 17th IEEE Symp. on Mass Storage Systems*, Baltimore, MD, March 2000.

H. Nagesh, S. Goil, and A. N. Choudhary*, “PMAFIA: A scalable parallel subspace clustering algorithm for massive datasets,” *Proc. Int. Conf. on Parallel Processing* (ICPP’2000), Toronto, Canada, Aug. 21–24, 2000.

X. Shen and A. N. Choudhary*, “A distributed multi-storage resource architecture and I/O performance prediction for scientific computing,” *Proc. High Performance Distributed Computing Conf.* (HPDC’2000), Pittsburgh, PA, Aug. 2000.

A. Srinilta and A. N. Choudhary*, “Multi-pool caching in continuous media server,” *Proc. Conf. on Multimedia Modeling* (MMM’99), Ottawa, Canada, Oct. 29, 1999, pp. 267–282.

W. Liao*, A. N. Choudhary*, D. Wiener, and P. Varshney, “I/O implementation and evaluation of parallel pipelined STAP on high performance computers,” *Proc. 6th Int. Conf. on High Performance Computing* (HIPC’99), Calcutta, India, Dec. 1999.

W. Liao*, A. N. Choudhary*, D. Wiener, and P. Varshney, “Design and evaluation of I/O strategies for parallel pipelined STAP applications,” *Proc. Int. Parallel and Distributed Processing Symp.*, May 2000, pp. 655–662.

X. Shen, W. Liao*, A. N. Choudhary*, G. Memik, M. Kandemir, S. More, G. Thiruvathukal, and A. Singh, “A novel application development environment for large-scale scientific computations,” *Proc. Int. Conf. on Supercomputing*, Santa Fe, NM, May 2000.

X. Shen, G. Thiruvathukal, W. Liao*, A. N. Choudhary*, and A. Singh, “A JAVA graphical user interface for large-scale scientific computations in heterogeneous systems,” *Proc. 4th Int. Conf. on High Performance Computing –ASIA*, Beijing, China, May 2000.

C.-C. Lin and W.-C. Lin*, “Interactions of center-on and center-off receptive field operators: The potential role in holistic perception from computational perspective,” *Proc. 5th Joint Conf. on Information Sciences*, Atlantic City, NJ, Feb. 27 – March 3, 2000, vol. 2, pp. 297–304.

P. Scheuermann* and M. Sayal, “A distributed clustering algorithm for web-based access patterns,” *Proc. 2nd ACM-SIGKDD Workshop on Distributed and Parallel Knowledge Discovery*, Boston, MA, Aug. 2000.

N. Shenoy*, A. Choudhary*, and P. Banerjee*, “A system-level synthesis algorithm with guaranteed solution quality,” *Proc. Design Automation and Test in Europe* (DATE 2000), Paris, France, March 27–30, 2000.

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,” *Proc. Int. Conf. on Parallel Processing* (ICPP-2000), Toronto, Canada, Aug. 2000.

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,” *Proc. ISCA 13th Int. Conf. on Parallel and Distributed Computing Systems* (PDCS-2000), Las Vegas, NV, Aug. 8–10, 2000.

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. 11th IASTED Parallel and Distributed Computing and Systems Conf.* (PDCS’99), Cambridge, MA, Nov. 1999.

Z. Ye, N. Shenoy*, and P. Banerjee*, “A C compiler for a processor with a reconfigurable functional unit,” *Proc. ACM/SIGDA Symp. on Field Programmable Gate Arrays*, Monterey, CA, Feb. 2000.

V. Taylor*, J. Chen, T. Canfield and J. Richard, “Efficient use of distributed systems for scientific applications,” *Proc. CAS 2000 Workshop*, NASA Ames Research Center, Feb. 2000.

W. Smith, I. Foster and V. Taylor*, “Scheduling with advanced reservations,” *Proc. IPDPS*, May 2000, Cancun, Mexico.

S. Mok, C. H. Wu*, and D. T. Lee, “A hierarchical workcell model for intelligent assembly and disassembly,” *Proc. 1999 IEEE Int. Symp. on Computational Intelligence in Robotics and Automation* (CIRA’99), Monterey, CA, Nov. 8–9, 1999, pp. 125–130.

S. Mok, C. H. Wu*, and D. T. Lee, “A system for automatic assembly and disassembly operations,” *Proc. 2000 IEEE Int. Conf. on Robotics and Automation*, San Francisco, CA, April 2000.

Ph.D. Dissertations

Dhruva Chakrabarti, *Compiler Support for Parallel Irregular Applications* (2000, advisor: P. Banerjee*)

Yanhong Yuan, *Parallel Algorithms for 3D Resistance and Capacitance Extraction* (2000, advisor: P. Banerjee*)

Jung J. Kim, *Disciplined Design Capture and Evolution of Object-Oriented Systems* (1999, advisor: L. J. Henschen*)

Jie Feng, *Fractional Fractal Geometry for Image Processing* (2000, advisor: L. J. Henschen*)

Lisa O. Singh, *Generating Constrained Association Rules from Semi-Structured Data* (1999, advisor: P. Scheuermann*)

Mehmet O. Sayal, *A Scalable and Adaptive Web Client-Server Architecture* (2000, advisor: P. Scheuermann*)