AVERY L. CHING

Yahoo! Search
Yahoo! Mission College – 2MC-08
2821 Mission College Boulevard
Santa Clara, CA 95054

Office: (408) 336-0639
Mobile: (847) 445-4846
Email: aching@ece.northwestern.edu
Web: www.ece.northwestern.edu/~aching

Research Interests

- High-performance and distributed computing
- Parallel I/O storage (middleware and file system optimizations)
- Large-scale and persistent services

Education

Northwestern University **Evanston, IL**

Ph.D. in Electrical Engineering and Computer Science, December 2007.
Advisor: Professor Alok Choudhary
Thesis: *Optimizing File System Techniques for Large-Scale Scientific Applications*
GPA: 3.9/4.0

B.S. in Computer Engineering, June 2002.
Graduated in Honors Program in Undergraduate Research
GPA: 3.7/4.0

Professional Experience

Yahoo! Search **Santa Clara, CA**

Senior Technical Yahoo
October 2007 to Current

Member of the content systems team. Deployed relevance changes to the selection rank algorithm for ordering over 100 billion web pages by the usefulness of their content. Current project is designing and implementing a gridified query service that is scalable, reliable, and persistent.

Sandia National Laboratories **Albuquerque, NM**

Graduate Student Internship
Summer 2006

Member of scalable I/O group in Organization 1423. Collaboratively developed a new, flexible two phase I/O implementation for ROMIO. Created a comprehensive MPI-IO test suite to ensure correctness. Implemented scalable locking techniques for parallel file systems.

Sandia National Laboratories **Albuquerque, NM**

Graduate Student Internship
July to October 2005

Member of scalable I/O group in Organization 1423. Implemented support in ROMIO for PVFS2 datatype I/O optimizations. Developed synthetic test suite *HPIO* for understanding problems with scientific data storage. All code given to PVFS2 development team for imminent integration.

Los Alamos National Laboratory **Los Alamos, NM**

Graduate Student Internship
April to July 2005

Member of RADIANT team. Designed and implemented a sequence similarity search algorithm simulator called "S3aSim." Work resulted in a better understanding of I/O access patterns for parallel sequence search algorithms and influenced the next generation design of mpiBLAST.

Argonne National Laboratory **Argonne, IL**

Givens Fellow
Summer 2004

Member of parallel tools group. Developed novel versioning based method for parallel file systems. Wrote initial test code for simulated evaluation and completed partial implementation in PVFS2. Presented simulated results.

Argonne National Laboratory **Argonne, IL**

Givens Fellow
Summer 2003

Member of parallel tools group. Designed and implemented datatype I/O processing for parallel file systems in PVFS1. Wrote corresponding PVFS1 driver code for ROMIO and compared against list I/O and POSIX I/O methods.

Argonne National Laboratory **Argonne, IL**

DOE ERULF Fellow
Summer 2002

Member of parallel tools group. Learned the ROMIO implementation of MPI-IO and optimized the PVFS1 driver for ROMIO to use the list I/O interface. Created new benchmarks based on tile reader code and the ASC FLASH code for testing list I/O through ROMIO.

Argonne National Laboratory **Argonne, IL**

DOE ERULF Fellow
Summer 2001

Member of parallel tools group. Implemented a list I/O interface for PVFS1. Developed native PVFS1 library benchmarks for testing typical block and cyclic I/O patterns. Wrote scripting tools for using Chiba City as a PVFS1 test platform.

Teaching Experience

Northwestern University, Department of EECS **Evanston, IL**

Co-Instructor
Winter 2007

ECE 362: Computer Architecture Projects
Instructor: Professor Alok Choudhary
Met with project groups on a bi-weekly basis to monitor progress. Duties included project management advice, progress evaluation, and assistance in fixing bugs. Maintained the course web page and provide support for relevant software. Evaluated completed projects and assigned final grades.

Teaching Associate
Fall 2005
Teaching Assistant
Fall 2004, 2006

ECE 361: Computer Architecture
Instructor: Professor Alok Choudhary
Presented numerous lectures and midterm review throughout the quarter. Held weekly office hours for students. Assisted in writing and grading homework, final exam problems, and class projects. Created and administered course web page.

Teaching Assistant Winter 2002	ECE 362: Computer Architecture Projects Instructor: Professor Alok Choudhary Assisted students in the technical details of class software and maintained course web page. Graded class projects and assisted in final grade assignment.

Professional Activities

Program Committee Member	Workshop on High Performance I/O Systems and Data Intensive Computing (HiperIO)
Reviewer	IEEE Transactions on Parallel and Distributed Systems (TPDS) Journal of Parallel and Distributed Computing (JPDC) Int'l Symposium on High-Performance Computer Architecture (HPCA) Supercomputing Conference (SC) Int'l Parallel & Distributed Processing Symposium (IPDPS) Int'l Conference on Cluster Computing (Cluster) Information Processing in Sensor Networks (IPSN)
Professional Travel	HPDC travel grant, 2006 Krell Institute SC travel grant, 2004 Northwestern University, Office of Research travel grant, 2003
Membership	IEEE Member, since 2002

Selected Honors and Awards

- DOE High-Performance Computer Science Fellowship, 2004 to present
- Argonne National Laboratory Givens Associate Fellowship, Summer 2003 & Summer 2004
- NASA Graduate Student Researchers Program Fellowship, 2003 - 2004
- University Scholar Tuition Award, 2003
- DOE Undergraduate Laboratory Fellowship Program (ERULF), Summer 2001 & Summer 2002
- Completed Honors Program in Undergraduate Research, 2002
- Member of Tau Beta Pi Honors Society, inducted 2002
- Walter P. Murphy Fellowship, 2001 - 2002
- Member of Eta Kappa Nu Engineering Honor Society, inducted 2001
- Eagle Scout, 1998
- National MathCounts Finalist, Washington D.C., 1994

Skills

- Parallel programming with MPI-1/2, HPF, and POSIX threads
- Parallel I/O expert in NetCDF, HDF5, MPI-IO, and parallel file systems (PVFS, Lustre, GPFS, etc.)
- Proficient in C/C++, Perl, Java, VHDL, Mentor Graphics, Matlab
- Knowledge in operating systems, distributed systems, computer hardware design and architecture and VLSI

Grant Proposals

P1. "The scientific data management center for enabling technologies"
Funding Agency: Office of Science, Office of Advanced Scientific Computing Research
Principal Investigator: Alok Choudhary
Status & Duration: Funded for $975,000 from 2006 - 2011
Contribution: Wrote several major sections and assisted in organizing proposal.

P2. "Scalable I/O middleware and file system optimizations for high-performance computing"
Funding Agency: National Science Foundation (NSF)
Principal Investigator: Alok Choudhary
Status & Duration: Funded for $540,000 from 2006 - 2009
Contribution: Wrote part of the motivation and several major sections.

P3. "Ultra-scalable system software tools for data-intensive computing"
Funding Agency: National Science Foundation (NSF)
Principal Investigator: Alok Choudhary
Status & Duration: Funded for $445,000 from 2005 - 2008
Contribution: Participated in overall proposal organization and wrote several sections.

Publications

Journal Publications

[1] Avery Ching, Alok Choudhary, Wei-keng Liao, Robert Ross, and William Gropp. Evaluating structured I/O methods for parallel file systems. In *International Journal of High Performance Computing and Networking*, 2:133-145, 2004.

Book Chapters

[2] Avery Ching, Kenin Coloma, Arifa Nisar, Wei-keng Liao, and Alok Choudhary. Distributed file systems. In *Wiley Encyclopedia of Computer Science and Engineering*. John Wiley & Sons, Inc., 2007.

[3] 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.

[4] Avery Ching, Kenin Coloma, and Alok Choudhary. Challenges for parallel I/O in GRID computing. In *Engineering the Grid: Status and Perspective*. American Scientific Publishers, 2006.

Conference and Workshop Publications

[5] Avery Ching, Robert Ross, Wei-keng Liao, Lee Ward, and Alok Choudhary. Noncontiguous locking techniques for parallel file systems. In *Proceedings of Supercomputing*, November 2007.

[6] Wei-keng Liao, Avery Ching, Kenin Coloma, Arifa Nisar, Alok Choudhary, Jackie Chen, Ramanan Sankaran, and Scott Klasky. Using MPI file caching to improve parallel write performance for large-scale scientific applications. In *Proceedings of Supercomputing*, November 2007.

[7] Wei-keng Liao, Avery Ching, Kenin Coloma, and Alok Choudhary. Improving MPI independent write performance using a two-stage write-behind buffering method. In *the NSF Next Generation Software Workshop, held in conjunction with the International Parallel and Distributed Processing Symposium*, March 2007.

[8] Wu-chun Feng, Avery Ching, and Chung-hsing Hsu. Green supercomputing for the desktop. In *the 3rd Workshop on High-Performance, Power-Aware Computing*, *held in conjunction with the International Parallel and Distributed Processing Symposium*, March 2007.

[9] 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.

[10] Kenin Coloma, Avery Ching, Alok Choudhary, Wei-keng Liao, Robert Ross, Rajeev Thakur, and Lee Ward. A new flexible MPI collective I/O implementation. In *Proceedings of the IEEE International Conference on Cluster Computing*, September 2006.

[11] 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.

[12] Peter Aarestad, Avery Ching, George Thiruvathukal, and Alok Choudhary. Scalable approaches for supporting MPI-IO atomicity. In *Proceedings of the IEEE/ACM International Symposium on Cluster Computing and the Grid*, May 2006.

[13] Avery Ching, Alok Choudhary, Wei-keng Liao, Lee Ward, and Neil Pundit. Evaluating I/O characteristics and methods for storing structured scientific data. In *Proceedings of the International Parallel and Distributed Processing Symposium*, April 2006.

[14] Kenin Coloma, Alok Choudhary, Avery Ching, Wei-keng Liao, Seung Woo Son, Mahmut Kandemir, and Lee Ward. Power and performance in I/O for scientific applications. In *the NSF Next Generation Software Workshop*, *held in conjunction with the International Parallel and Distributed Processing Symposium*, April 2005.

[15] Avery Ching, Alok Choudhary, Wei-keng Liao, Robert Ross, and William Gropp. Efficient structured data access in parallel file systems. In *Proceedings of the IEEE International Conference on Cluster Computing*, December 2003.

[16] Avery Ching, Alok Choudhary, Kenin Coloma, Wei-keng Liao, Robert Ross, and William Gropp. Noncontiguous access through MPI-IO. In *Proceedings of the IEEE/ACM International Symposium on Cluster Computing and the Grid*, May 2003.

[17] Avery Ching, Alok Choudhary, Wei-keng Liao, Robert Ross, and William Gropp. Noncontiguous I/O through PVFS. In *Proceedings of the IEEE International Conference on Cluster Computing*, September 2002.

References

Professor Alok Choudhary, Department Chair
Department of EECS
Northwestern University
2145 Sheridan Road
Evanston, IL 60208-3118
Phone: (847) 491-4129 ; fax: (847) 491-4144
Email: choudhar@ece.northwestern.edu

Dr. Robert Ross - Computer Scientist
MCS Division
Argonne National Laboratory
9700 South Cass Avenue - Building 221
Argonne, IL 60439-4844
Phone: (630) 252-4588 ; fax: (630) 252-5986
Email: rross@mcs.anl.gov

Associate Professor Wu-chun Feng
Department of Computer Science
Virginia Tech
660 McBryde Hall (0106) - Virginia Tech
Blacksburg, VA 24061
Phone: (540) 231-1192 ; fax: (540) 231-9218
Email: feng@cs.vt.edu

Dr. Rajeev Thakur – Computer Scientist
MCS Division
Argonne National Laboratory
9700 South Cass Avenue - Building 221
Argonne, IL 60439-4844
Phone: (630) 252-7847 ; fax: (630) 252-5986
Email: thakur@mcs.anl.gov

Research Assistant Professor Wei-keng Liao
Department of EECS
Northwestern University
2145 Sheridan Road
Evanston, IL 60208-3118
Phone: (847) 491-2916 ; fax: (847) 491-4455
Email: wkliao@ece.northwestern.edu

Assistant Professor Gokhan Memik
Department of EECS
Northwestern University
2145 Sheridan Road
Evanston, IL 60208-3118
Phone: (847) 467-1168 ; fax: (847) 491-4144
Email: memik@ece.northwestern.edu