

## YUTAO ZHONG

Department of Computer Science  
University of Rochester  
Rochester, NY 14627  
<http://www.cs.rochester.edu/~ytzhong>

Phone: (585) 275-5414  
Fax: (585) 273-4556  
[ytzhong@cs.rochester.edu](mailto:ytzhong@cs.rochester.edu)

### RESEARCH INTERESTS

---

Program analysis and optimization, compilers, and programming languages

### EDUCATION

---

University of Rochester, Rochester, NY

Ph.D. in Computer Science, July 2005 (expected), GPA 4.0/4.0

Thesis: Whole-Program Data Locality Hierarchy

Advisor: Dr. Chen DING

M.S. in Computer Science, March 2002, GPA 3.8/4.0

Nanjing University, Nanjing, China

M.E. in Computer Science and Technology, June 2000, GPA 3.7/4.0

Thesis: Agent-based Distributed Internet Information Retrieval

Advisor: Prof. Yongsen XU

B.S. in Computer Science and Technology, June 1997, GPA 3.8/4.0

### RESEARCH EXPERIENCE

---

#### Research Assistant, University of Rochester

Fall 2000 - present

Thesis work focuses on program locality behavior measurement, analysis, and optimization. The central concept is reuse distance, which measures the volume of intervening data between two consecutive references to the same data element. It is an architecture-independent program-specific locality metric. Long reuse distances reveal recurrence patterns in data references. My research starts by constructing a profiling system, which performs program instrumentation and collects and efficiently measures reuse distances for the whole program over its complete execution. I show that many benchmark programs exhibit consistent locality patterns which can be captured by the whole-program locality model based on reuse distance [5]. This further leads to the accurate prediction of cache miss rates for a wide range of programs, for all inputs on all cache sizes [4]. By attributing reuses to individual data elements, I also propose an approximation of data grouping according to reference affinity. Research on this topic includes a theoretical reference affinity model and practical data reorganization techniques that consistently improve performance [2]. Other research activities include locality-based phase detection and phase behavior prediction [1]. The tool we built for reuse distance measurement has been distributed to researchers in Michigan Tech., Rutgers University and University of Texas at Austin.

**Research Intern, Intel Corporation, Beijing, China** Summer 2004

Microprocessor Technology Lab (MTL), under the supervision of Dr. Youfeng Wu. Designed and implemented a system analyzing reference affinity for program instruction code. Evaluated the effectiveness of reference affinity in region-based compilation.

**Research Intern, Microsoft Research, Redmond, WA** Summer 2002

Program Analysis and Monitoring Group, Programmer Productivity Research Center (PPRC), under the supervision of Dr. Trishul Chilimbi. Extended a system that detects repeated sequences of data accesses to enable collecting and analyzing such sequences on shared data among multiple threads. The pattern discovered from these shared hot data streams can be used as hints to reduce the coherence and communication overhead for multi-threaded programs.

**Research Assistant, Nanjing University, China** 1997-2002

Joint project on Web-based information retrieval, Department of Computer Science and Technology and Network Information Center, Nanjing University. Served as the technical lead and developed a personal information retrieval agent model and the multi-agent network environment[6,7].

**TEACHING EXPERIENCE****Teaching Assistant, University of Rochester**

CSC 255/455 Advanced Programming Systems Fall 2002

*Advanced course cross listed for senior undergraduate students and graduate students.*

CSC 200 Undergraduate Problem Seminar Spring 2002

*Core course required for the honors B.S. degree.*

CSC 171 Computer Programming Fall 2001

*The first programming course and one of pre-major requirements in the CS curriculum.*

*Held office hours. Graded homeworks, exams and projects. Developed and maintained software and instructions for course projects.*

**Guest Lecturing, University of Rochester**

Delivered guest lectures in Computer Organization (CSC252) and Advanced Programming Systems (CSC 255/455).

Fine-Grained and Coarse-Grained Parallelism CSC255/455, Fall 2004

Handling Control Flow CSC255/455, Fall 2004

Advanced Pipelining and Instruction Level Parallelism CSC252, Spring 2004

Data Layout Optimizations CSC252, Spring 2004

**Informal Undergraduate Independent Study Advising, University of Rochester**

Co-advised (with Prof. Chen Ding) two undergraduate students in their research projects.

Maksim Orlovich 2003

*Implementing a C program instrumentation and transformation tool. Currently a graduate student in Cornell University.*

Ahren Studer 2003-2004

*Implementing an interactive visualization tool for cache miss rate prediction across program inputs and cache sizes. Currently a graduate student in Carnegie Mellon University.*

**Teaching Assistant, Nanjing University**

Operating Systems

Fall 1997

Core course required for undergraduates.

*Led lab sessions for more than 50 undergraduate students. Graded homeworks, exams and projects.***PUBLICATIONS**

---

**Papers at Referred Conferences**

- [1] “Locality Phase Prediction,” by X. Shen, Y. Zhong, and C. Ding. *Eleventh International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’04)*, Boston, MA, October 2004.
- [2] “Array Regrouping and Structure Splitting Using Whole-Program Reference Affinity,” by Y. Zhong, M. Orlovich, X. Shen, and C. Ding. *ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation (PLDI’04)*, Washington, DC, June 2004.
- [3] “Regression-Based Multi-Model Prediction of Data Reuse Signature,” by X. Shen, Y. Zhong, and C. Ding. *Fourth Annual Symposium of the Los Alamos Computer Science Institute (LACSI’03)*, Santa Fe, NM, October 2003.
- [4] “Miss Rate Prediction across All Program Inputs,” by Y. Zhong, S. G. Dropsho, and C. Ding. *Twelfth International Conference on Parallel Architectures and Compilation Techniques (PACT’03)*, New Orleans, LA, September 2003.
- [5] “Predicting Whole-Program Locality through Reuse Distance Analysis,” by C. Ding and Y. Zhong. *ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation (PLDI’03)*, San Diego, CA, June 2003.
- [6] “Agent Society: A Personalized Intelligent Information Retrieval System through WWW,” by X. Du, Y. Zhong, Y. Xu, J. Chen. *IASTED International Conference on Applied Informatics*, Innsbruck, Austria, February, 2001.
- [7] “A Multi-Agent Cooperation Model in Internet Information Retrieval,” by Y. Zhong, W. Huang, X. Chen, Y. Xu, J. Chen. *Third National Conference on Object-Oriented Technology and Applications*, Nanjing, China, September 1999.

**Papers at Referred Workshops**

- [8] “Phase-Based Miss Rate Prediction Across Program Inputs,” by X. Shen, Y. Zhong, and C. Ding. *Seventeenth International Workshop on Languages and Compilers for Parallel Computing (LCPC’04)*, West Lafayette, IN, September 2004.
- [9] “Finding Reference Affinity Groups in Trace Using Sampling Method,” by C. Zhang, Y. Zhong, C. Ding and M. Ogihara. *Third Workshop on Mining Temporal and Sequential Data*, in conjunction with the ACM SIG-KDD’04, Seattle, WA, August 2004.
- [10] “A Hierarchical Model of Reference Affinity,” by Y. Zhong, X. Shen, and C. Ding. *Sixteenth International Workshop on Languages and Compilers for Parallel Computing (LCPC’03)*, College Station, TX, October 2003.
- [11] “Compiler-Directed Run-Time Monitoring of Program Data Access,” by C. Ding and Y. Zhong. *First ACM SIGPLAN Workshop on Memory System Performance (MSP’02)*, Berlin, Germany, June 2002.

- [12] “Reuse Distance Analysis for Scientific Programs,” by Y. Zhong, C. Ding, and K. Kennedy. *Sixth Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR’02)*, Washington, DC, March 2002.

### Technical Reports

- [13] “An Interactive 3-D Visualization Tool for Accurate Estimation of Miss Rates,” by A. Studer, Y. Zhong, and C. Ding. Computer Science Department, University of Rochester, Rochester, *UR-CS-TR-826*, December 2003.
- [14] “Reuse Distance Analysis,” by C. Ding, and Y. Zhong. Computer Science Department, University of Rochester, *UR-CS-TR-741*, February 2001.

### Book Chapters

- [15] “Java Programming and Development Environment,” by Y. Xu, J. Chen, W. Huang, Y. Zhong and C. Sun. Nanjing University Publishing House, 1998, in Chinese.

### Journal Submissions

- [16] “Distance-Based Locality Analysis and Prediction,” by C. Ding, Y. Zhong, and X. Shen. Submitted to *ACM Transactions on Programming Languages and Systems*, September, 2004.
- [17] “Miss Rate Prediction across Program Inputs and Cache Configurations,” by Y. Zhong, S. G. Dropsho, X. Shen, A. Studer, and C. Ding. Submitted to *IEEE Transactions on Computers*, October, 2004.
- [18] “Phase Behavior Prediction across Program Inputs,” by X. Shen, Y. Zhong, and C. Ding. Submitted to *ACM Transactions on Architecture and Code Optimization*, December, 2004.

### PATENT

---

Temporal Affinity Analysis Using Reuse Signatures, C. Ding and Y. Zhong, provisional filed, 2003.

### PRESENTATIONS

---

- “Array Regrouping and Structure Splitting Using Whole-Program Reference Affinity”, ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation, Washington, DC, June 2004.
- “Miss Rate Prediction across All Program Inputs”, Twelfth International Conference on Parallel Architectures and Compilation Techniques, New Orleans, LA, October 2003.
- “A Hierarchical Model of Reference Affinity”, Sixteenth International Workshop on Languages and Compilers for Parallel Computing, College Station, TX, October 2003.
- “Regression-based Program Locality Prediction”(poster presentation), ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation, San Diego, CA, June 2003.
- “Reuse Distance Analysis for Scientific Programs”, Sixth Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers, Washington, DC, March 2002.
- “A Multi-Agent Cooperation Model in Internet Information Retrieval”, Third National Conference on Object-Oriented Technology and Applications (China), Nanjing, China, September 1999.

---

**ACTIVITIES AND SERVICE**

---

- Student member of the ACM, SIGPLAN, and SIGMETRICS.
- Reviewed papers for the following conferences and journals: ICPP'03, IPDPS'04, NPC'04, PLDI'05, Journal of Parallel and Distributed Computing.
- Helped to organize the review of language/compiler track of ICPP'04 and NPC'04.
- Volunteer and summary writer for the Fifth Symposium on Operating Systems Design and Implementation, Boston, MA, 2002.
- Comprehensive Exams Committee Member, University of Rochester, Computer Science Department, 2004 - 2005.
- Graduate Admission Committee Member, University of Rochester, Computer Science Department, 2002 - 2003.
- Graduate Recruiting Committee Member, University of Rochester, Computer Science Department, 2001 - 2002.

---

**REFERENCES**

---

Prof. Chen Ding  
Department of Computer Science  
University of Rochester  
Rochester, NY 14627  
cding@cs.rochester.edu  
+1-585-275-1373

Prof. Michael L. Scott  
Department of Computer Science  
University of Rochester  
Rochester, NY 14627  
scott@cs.rochester.edu  
+1-585-275-7745

Prof. Sandhya Dwarkadas  
Department of Computer Science  
University of Rochester  
Rochester, NY 14627  
sandhya@cs.rochester.edu  
+1-585-275-5647

Dr. Youfeng Wu  
Principal Researcher  
Programming Systems Lab  
Intel Corporation  
2200 Mission College Blvd.  
Santa Clara, CA 95052  
youfeng.wu@intel.com  
+1-408-765-9415