Resume

Table of Contents

CURRICULUM VITAE

NAME: Guang R. Gao

OFFICE ADDRESS:

Department of Electrical Engineering
104 Evans Hall
University of Delaware
Newark, DE 19716
Tel: 302-831-8218
Fax: 302-831-4316
ggao@eecis.udel.edu

EDUCATION

Ph.D Degree in *Electrical Engineering and Computer Science*
Massachusetts Institutes of Technology, August 1986.
Member of Computational Structures Group at Laboratory of Computer Science, MIT, June 1982 to August 1986.

Master Degree in *Electrical Engineering and Computer Science*
Massachusetts Institutes of Technology, June 1982.

BS in *Electrical Engineering*
Tsinghua University, Beijing.

PROFESSIONAL EXPERIENCE

University of Delaware
Newark, DE.
Associate Professor, Department of Electrical and Computer Engineering, Sept. 96-present
Associate Professor, Department of Electrical and Computer Engineering, Sept. 96-present
Founder and a leader of the Computer Architectures and Parallel Systems Laboratory (CAPSL) .

McGill University
Montreal, Canada
Associate Professor, School of Computer Science, June'92-August,1996
Assistant Professor, School of Computer Science, Aug.'87-June'92
Founder and a leader of the Advanced Compilers, Architectures and Parallel Systems Group (ACAPS) at McGILL since 1988.

Philips Research Laboratories
Sept. 1986 – June 1987
Briarcliff Manor, NY, USA
Senior member of research staff of the Computer Architecture and Programming Systems Group. Played a major role in founding a multiprocessor system project, and research in parallelizing compilers.

Massachusetts Institutes of Technology
June 1980 - Aug. 1986
Member of the Computational Structures Group at the Laboratory of Computer Science, MIT.
Participated in the MIT Static Dataflow Architecture Project and other projects.
Proposed a novel methodology of organizing array operations to exploit the fine-grain parallelism of dataflow computation models. Developed a unique pipelined code mapping scheme for dataflow machines (later known as dataflow software pipelining).

Center Of Advanced Studies, IBM Toronto Lab
Aug 1993 – June 1994
Visiting scientist with a NSERC Senior Industrial Fellowship.

CURRENT RESEARCH AREAS:

Computer Architecture and Systems
Parallel and Distributed Systems
Optimizing and Parallelizing Compilers, Parallel Programming
VLSI and Application-Specific System Design

PROFESSIONAL MEMBERSHIP

I am a Senior Member of IEEE, Member of ACM, ACM-SIGARCH, ACM-SIGPLAN.
I am currently a Distinguished Visitor of IEEE Computer Society.

NATIONAL RECOGNITION:

- **IEEE Computer Society Distringuished Visitor, 1998-2001**

- **IEEE, Senior Member**

- **Program Committee Members of Recognized International Conferences**

  - IEEE International Symposium on Computer Architecture (HPCA-95, HPCA-99, HPCA-00)
  - ACM Symposium on Programming Language Design and Implementation (PLDI-98)
  - ACM International Conference on Supercomputing (ICS-95)
  - ACM/IEEE International Symposium on Microarchitectures (MICRO-95, 96, 97)
  - International Parallel Processing Symposium (IPPS-95)
  - IFIP and ACM SIGARCH International Conference on Parallel Architectures and Compilation Techniques (PACT-94,95,96,97,98,99)
  - International Conference on Algorithms And Architectures for Parallel Processing (ICAPP-95)
  - Parallel Architecture and Language Europe (PARLE-91,92,93,94,95)
  - International Conference on Parallel Processing (EURO-PAR-95,96)
  - Working Conference on Massively Parallel Programming Models (MPPM-93,95,97, 99)
  - High Performance Computing Symposium (HPCS-95, 96, 98), Canada.
  - International Conference on Compiler Construction (CC-98,99,00), Europe.
  - International Symposium on High Performance Computing (ISHPC99), Japan.

- **Conference Committee Chairmanship**

- Program Chairman of the 1994 ACM SIGARCH, International Conference on parallel Architectures and Compilation Techniques (PACT '94), Aug/. 1984. Montreal, Canada, co-sponsored by IFIP and in association with ACM SIGPLAN, IEEE TCCA (Technical Committee on Computer Architecture) and IEEE TCPP (Technical Committee on Parallel Processing).

- General Co-Chair of the 1998 International Conference on Parallel Architectures and Compilation Techniques (PACT '98), Oct. 1998, Paris, France., co-sponsored by IFIP and IEEE Computer Society

- Chair of the Third Workshop on Petaflop Computing, Feb. 1999. Annapolis, MD.

- Co-Chair of the Multithreaded Architecture Workshop, in Conjunction to HPCA99, Orlando, Florida, Jan. 1999.

- Co-Chair of the Compiler and Architecture Support for Embedded Systems (CASES98), Washington D.C., Oct. 1998.

- **Journal Editorialship**

- I am elected to the Editorial Board of IEEE Transactions on Computers (1998 -)
- I am elected to the Editorial Board of IEEE Concurrency Journal (1997 -)
- I joined the Editorial Board of the Journal on Programming Languages in Jan. 1996, and subsequently became one of the two Co-Editors of the journal.
- I have been a Guest Editor for the Special Issue on Dataflow and Multithreaded Computers, Journal of Parallel and Distributed Computing, Academic Press, June, 1993.

- **Invited Seminars and Distinguished Seminars**

I have given seminars in many industrial and academic organizations: IBMT.J. Watson Research Center, IBM Toronto Lab, AT&T Bell Laboratories, BNR, HP Labs, SGI, DEC, NRL(Navy Research Lab.), MIT, Stanford, UC Berkeley, NYU, Cornell U., University of Victoria are just named a few.

Section A: Teaching and Research Supervision

A.1 Teaching

A.1.a Teaching at University of Delaware

- **New courses introduced and taught:**

CPEG 324 Computer System Design

This is is a new undergraduate CE design course which I developed and taught in Spring 97. Now, it will be offered again under the title CPEG-422 this fall.

This course stresses the principal design concepts which are embodied in modern computer architectures, and emphasizes ideas which we believe will continue to apply into the future, in spite of a rapidly changing technological environment. The primary objective of the course is to show how the design and evaluation of architectural features, based on both qualitative and quantitative studies, can be used to achieve balanced, efficient systems, well-matched to the class of problems they overcome.

ELEGG-652: Topics in High-Performance Architecture

This is a graduate core course and I taught it first time in the fall of 1997.

This course examines the basic principles and methodology used in the design and evaluation of high-performance computer architectures, and its relation with the underlying program execution and architecture models. Topics include pipelining and vector processing, instruction level parallelism (ILP) architectures, multiprocessor architectures and high-speed networking, memory consistency models and cache-coherence issues, fine-grain parallelism and multithreaded architectures, and the roles of optimizing and parallelizing compilers.

ELEG 867-14: Topics in Hardware/Software Codesign

This new course introduces the concepts, principles and methods of digital system design from both a hardware and a software viewpoint. In the context of general purpose computer systems, the principles studied in this course include the close interaction between compiler technology and architecture design. In the context of special-purpose systems, such as embedded systems, the course will deal with the close interaction between software synthesis and hardware system design.

Topics to be discussed include the fundamentals of analysis, generation, synthesis, and optimization of computer code. Specific topics in this area include dependency analysis, code motion, scheduling, register and resource allocation. Among the hardware micro-architecture topics studied are pipeline co-design and memory models. Important case studies that illustrate the basic principles of software/hardware co-design will be introduced. Topics in the new emerging field of adaptive computing system design will be discussed.

- **Activities**

- New hardware/software tools introduced or developed for teaching laboratories:

Modern computer architecture and system design involve both intensive software and hardware design activities. In the new courses introduced, the students are exposed to both software/hardware tools

and methodology for computer architecture design (e.g. software simulation toolset) as well as hardware design tools and methodology (e.g. VHDL tools and environments) on digital systems.
Students are expected to learn modern design tools and related skills through lab assignments and course projects. To this end, we have invested extensive effort to develop the laboratory and introduce the VHDL design environment in the course.

- The SEMi instruction set architecture simulator which provide accurate timing simulation for RISC-like architectures and its cache memory.

- The EARTH architecture emulation testbedwith a 20-node multiprocessor hardware engine which provide tools to study parallel and multithreaded programming paradigms and architecture models. The EARTH-MANNA platform and PC based EARTH-Beowulf platform are developed and made available in the teaching.

- A series of VHDL based hardware design and simulation tools has been introduced and established, which include VHDL behaviour simulation tool, the VLSI synthesis tools, FPGA place and routing tools, and the FPGA based hardware experimental test boards. This has considerably enhanced the teaching capacity for undergraduate design course and courses in computer architectures.

- Various benchmark suits for architecture/compiler studies: SPEC, LINPACK, Whetstone, Drystone, Livermore Loops, NAS, Spice, GCC etc. have been introduced;

- The CAPSL laboratory seminar series.

I have established the Computer Architecture and Parallel System Laboratory since I joined UDel. In addition to perform research, one important objective of this laboratory has been to facilitate the teaching of the computer architecture and digital system courses, and training of graduate research and teaching assistants. The new courses and software tools described above depend directly on this laboratory. The laboratory is now equipped with various workstations, We have a wide variety of research and teaching software installed, and a number of my best graduate students have been actively participated and contributed to teaching. Activities organized include:

- organization of the CAPSL research seminar series;

- invitation of a number of distinguished speakers of international reputation to give such seminars;

A.1.b Other Teaching Experience

At McGill Uniersity, I have introduced and developed a set of new courses (308-505,308-605,308- 622) on high-performance computer architectures, parallel systems and parallelizing compilers. These courses have been consolidated and improved over the period of time, forming a core for students who are interested in the related subject areas. I have also taught a number of graduate seminar courses. (Details can be provided by request). The excellence of my teaching have been recognized through the following outstanding teaching award nominations:

- nomination for the McGill Engineering Class of 51' Award for Outstanding Teaching (1988); - nomination reconsideration for the Engineering Class of 51' Award for Outstanding Teaching (1989);

- nomination for the McGill Engineering Class of 51' Award for Outstanding Teaching (1990); - nomination reconsideration for the Engineering Class of 51' Award for Outstanding Teaching (1991);

A.2. Research Supervision

Current, graduate students under my supervision include:

Gieger, Thomas	(processing in memory and multithreading)
Marquez, Andres	(multithreaded architectures)
Ryan, Sean	(optimizing compilers)
Stouchinin, Artour	(instruction-lelvel parallelism, software pipelining)
Tang, Xi-Nan	(compiler for multitheading)
Thulasiraman, Parimala,	(parallel algorithmsand applications)
Yang, Hongbo	(instruction-level parallelism)
Douillet, Alban	(compiling for multithreading)

Current Postodoc fellows under my supervision include:

Amaral, Nelson	(system software, compilers)
Kevin, Theobald	(computer architecture, parallel systems)
Rupak, Thulasiram	(parallel applications)

Already Completed:

The applicant has completed the supervision of 7 Ph.D. and 18 M.Sc. students, and 5 postdoctoral fellows in the proposed research areas of high-performance computing.

Post-Doctor (4 Completed)	**Ph.D. Level (7 Graduaged)**		**M.Sc Level (18 Graduated)**			
G. Liao	E. Altman	(1991 - 1996)	H. Cai	(1995 - 1997)	R. Shanker	(1991 - 1993)
(1991-1993)	H. Hum	(1998 - 1992)	N. Elmasri	(1992 - 1995)	N. Shiri	(1990 - 1992)
O. Maquelin	S. Nemawarka	(1989 - 1996)	A. Emtage	(1988 - 1991)	R. Silvera	(1996 - 1997)
(1994 - 1998)	Q. Ning	(1990 - 1993)	S.H. Han	(1996 - 1997)	A. Stouchinin	(1994 - 1996)
G. Ramaswamy	V. C. Sreedhar	(1990 - 1995)	A. Jimenez	(1993 - 1996)	R. Wen	(1993 - 1995)
(1990 – 1994)	G. Tremblay	(1988 - 1994)	L. Lozano	(1992 - 1994)	Y-B Wong	(1989 - 1991)
X. Tian	R. Yates	(1988 - 1992)	S. Merali	(1993 - 1996)		
(1993 – 1996)			C. Moura	(1991 - 1993)		
J. Wang			C. Mukerji	(1991 - 1994)		
(1995 – 1997)			R. Olsen	(1989 - 1992)		
			Z. Paraskevas	(1987 - 1989)		
			H. Petry	(1995 - 1997)		

Those who have graduated are highly trained in the field of parallel architectures and compilers, as evidenced by the fact that they have been working (or worked) as tenure-track university professors (Ramaswamy, Tremblay); as engineers in key industrial sectors, e.g., Intel(Hum), Nortel (Wang), IBM (Altman,Nemawarkar, Sreedhar), BNR (Liao, Wen), HP (Lozano), Convex (Ning), NCUBE (Olsen), CAE(Nassur), AT&T (Petry); and as researchers in government labs, e.g., LLNL (Yates), or assuming other professional jobs.

Section B: Scholarship

B.1 Research Activity and Interests

1. Computer Architecture and Systems.

> *One main question facing modern computer architects is: Is it ever possible to build a high-performance parallel architecture combining the power of hundreds, or even thousands, of processors to solve real world applications (regular or irregular) with scalable performance?*

My research interests have been seeking an answer to this challenge. In particular, our primary work has been concentrated on multithreaded program execution models and architectures. To this end, I have initiated/led or played a major roles in a number of research projects in this area.

- In the EARTH (Efficient Architecture for Running THreads) project, our focus has been, given the conventional off-the-shell processor technology, how can a multithreaded program execution model and architecture be developed which can exploit fine-grain parallelism and deliver scalable performance with affordable cost. Our current activities include: refinement of the EARTH program execution model and shared-memory architecture support (partially supported via a NSF-MIPS grant joint with USC), study and implementation of EARTH model on a cluster of SMP workstations linked with high-speed networks (via a NSF-CISE infrastructure grant), the study and implementation of a real world large irregular application (the crack propagation) on EARTH platforms (partially supported via a NSF-CISE grant joint with Cornell), and the investigation of compiling techniques for multithreaded architectures (partially supported via a NSF-CISE grant).

- In the HTMT (Hybrid Technology Multithreaded Architecture Project), our focus is on very high-end parallel supercomputer architectures based on advanced technology and beyond the off-the-shell processor architectures. Our recent and current activities include an initial "point design" study of the HTMT architecture model (funded partially via a NSF grant), and subsequently a feasibility study of the HTMT program execution and architecture model for a petaflop-scale architecture which employs and integrates the combined capabilities of semiconductor, superconductor, and optical technologies, as well as the PIM (processing-in-memory) technology (funded via a grant from DARPA/NSA/ NASA through JPL/Caltech).

- In the Data-IntensiVe Architecture Project (DIVA), our goal is to exploit alternative multithreaded execution model to fully utilize the processing power and memory bandwidth provided by the DIVA PIM chip for large-scale high-performance data base applications (funded through a grant from DARPA via Caltech/JPL).

2. Optimizing compiler technology.

- **Modulo scheduling and software pipelining**. My interest in modulo scheduling and software pipelining stemmed from work on register allocation for loops on dataflow machines. This work culminated in a mathematical formulation of the problem in a linear periodic form. It was soon discovered that this formulation can also be applied to software pipelining for conventional architectures. This formulation was then used to prove an interesting theoretical result: the minimum storage assignment problem for rate-optimal software pipelined schedules can be solved using an efficient polynomial-time method provided the target machine has enough functional units so resource constraints can be ignored. At the same time, I and my colleague and students have proposed the use of an ‘‘interval graph” based register allocation algorithms which appear to provide a good representation to study combined instruction scheduling and register allocation. Subsequently, we extended our framework to handle resource constraints, resulting in a unified integer linear programming formulation for the problem of simple pipelined architectures. The work was subsequently generalized to more complex architectures. This work was implemented in MOST --- the Modulo Scheduling Toolset -- developed in my group. Recently, I and co-workers have proposed "co-scheduling", a FSA (Finite State Automata) based framework for simultaneous design of hardware pipelines structures and software-pipelined schedules (partially funded through a NSF-CCR grant).

- **Program analysis techniques**. I have been interested in program analysis techniques for compiler optimization. V.C. Sreedhar (my Ph.D student) and myself have proposed a novel program representation, called the DJ graph. Based on DJ-graph, we have developed a surprisingly simple algorithm for computing Phi-nodes for arbitrary flowgraphs (reducible or irreducible) that runs in linear time. Based on DJ graphs, we have developed other novel and efficient algorithms to a series of problems in flowgraph analysis such as multiple-node immediate dominator analysis, identification of reducible and irreducible graphs, incremental algorithm for maintaining dominator trees, and exhaustive and incremental dataflow analysis based on DJ-graphs.

- **Program parallelization**: I have been interested in methodology of collective loop optimization. We have developed a methodology which has been applied to a collection of loops to perform a novel optimization called array contraction, that saves space and time by converting an array variable into a scalar variable or a buffer containing a small number of scalar variables. We have shown that the array contraction problem can be solved efficiently for a class of loops.

- **Thread partitioning**. I have been interested in the automatic thread partitioning and the threaded-code generation problem. We have developed a new heuristic algorithms based on an extension of the classical list scheduling algorithm. Based on a cost model, our algorithm groups instructions into threads by considering the trade-offs of the following characteristics: exploitation of parallelism, latency tolerance, minimizing thread switching costs and sequential execution efficiency. The proposed algorithm has been implemented, and a quantitative performance study of our algorithm has been conducted. Currently, we are extending our method and studying new thread partitioning algorithms which can integrate the scheduling and register allocation under the same framework (partially funded via a NSF-CCR grant).

3. <u>Other areas</u>

- Memory consistency models. I am interested in the problem of defining a memory model that does not rely on the memory coherence assumption, and also the problem of designing a cache consistency protocol based on such a memory model. I and my colleague have defined a new memory consistency model, called Location Consistency (LC), in which the state of a memory location is modeled as a partially ordered multiset (pomset) of write and synchronization operations. We have proved that LC is strictly weaker than existing memory models, but is still equivalent to stronger models for parallel programs that have no data races. We also introduced a new multiprocessor cache consistency protocol based on the LC memory model.

B.2: List of Research Contributions

Refereed Journal Publications

X. Tang and Guang R. Gao, Automatic partitioning threads for multithreaded architectures, Special Issues on Compilation and Architectural Support for Parallel Applications, Accepted for Application, June, 99.

1. Vugraman C. Sreedhar, Guang R. Gao and Yong-Fong Lee, A New Framework for Elimination0Based Dataflow Analysis Using DJ Graphs, ACM Transaction on Programming Languages and Systems, Vol 20, No. 2, PP 388-433, march 1998.

2. Erik Altman, Guang R. Gao, Optimal Modulo Scheduling Through Enumeration, International Journal on Parallel Programming, Accepted for publication, 1998.

3. Erik Altman, Guang R. Gao, A Unified Framework for Instruction Scheduling and Mapping for Function Units with Structural Hazards, Journal of Parallel and Distributed Computing, No. 39, pp 259-293, 1998.

4. Vugranam C. Sreedhar, Guang R. Gao, and Yong-fong Lee. Incremental computation of dominator trees. ACM Transactions on Programming Languages and Systems, 1996. Vol 19, No. 2, pp239-252, March 1997.

5. Vugranam C. Sreedhar, Guang R. Gao, and Yong fong Lee. A quadratic time algorithm for computing multiple node immediate dominators. Journal of Programming Languages, 1996. Accepted for Publication.

6. R. Govindarajan, Erik R. Altman, and Guang R. Gao. A framework for resource-constrained rate-optimal software pipelining. IEEE Transactions on Parallel and Distributed Systems, pages 1133-1149, November 1996.

7. Herbert H. J. Hum, Olivier Maquelin, Kevin B. Theobald, Xinmin Tian, Guang R. Gao, and Laurie J. Hendren. A study of the EARTH-MANNA multithreaded system. International Journal of Parallel Programming, 24(4):319-347, August 1996.

8. Vugranam Sreedhar, Guang R. Gao, and Yong fong Lee. Identifying loops using dj graphs. ACM Transactions on Programming Languages and Systems, 1996. Accepted.

9. Vugranam C. Sreedhar and Guang R. Gao. A linear time algorithm for placing OE-nodes. Journal of Programming Languages, 1995. Accepted.

10. Ning Qi, Vincent V. Dongen, and Guang R. Gao. Automatic data and computation decomposition for distributed memory machines. Parallel Processing Letters, 5(4):539-550, April 1995.

11. Vugranam Sreedhar and Guang Gao. Computing OE-nodes in linear time using DJ-graphs. Journal of Programming Languages, April 1995. 3(1995), page 191-213.

12. E. Arjomandi, W. O'Farrell, I. Kalas, G. Koblents, F. Ch. Eigler, and G. R. Gao. ABC++: Concurrency by inheritance in C++. IBM Systems Journal, 34(1):120-137, 1995.

13. R. Govindarajan and Guang R. Gao. Rate-optimal schedule for multi-rate DSP computations. Journal of VLSI Signal Processing, 9(3), April 1995. page 211-232.

14. G. R. Gao. An efficient hybrid dataflow architecture model. Journal of Parallel and Distributed Computing, 19(4):293-307, December 1993.

15. Laurie J. Hendren, Guang R. Gao, Erik R. Altman, and Chandrika Mukerji. A register allocation framework based on hierarchical cyclic interval graphs. The Journal of Programming Languages, 1(3):155-185, 1993.

16. Qi Ning and Guang R. Gao. Optimal loop storage allocation for argument-fetching dataflow machines. International Journal of Parallel Programming, 21(6):421-448, December 1992.

17. H. H. J. Hum and G. R. Gao. A high-speed memory organization for hybrid dataflow/von Neumann computing. Future Generation Computer Systems, 8:287-301, 1992.

18. G. R. Gao, H. H. J. Hum, and Y-B Wong. Toward efficient fine-grain software pipelining and the limited balancing techniques. International Journal of Mini and Microcomputers, 13(2):57-68, 1991.

19. Guang R. Gao. Exploiting fine-grain parallelism on dataflow architectures. Parallel Computing, 13(3):309-320, March 1990.

Publications in Refereed Conference Proceedings (Last SixYears Only)

I have more than 80 publications in refereed conferences. Due to space limitations, only those in the last 6 years are listed. The rest can be provided by request.

1. G. Heber, R. Biswas, and Guang R. Gao, Self-Adative Walks over Adaptive Unstructured Grids, In the Proceedings of Irregular'99 in conjuction to the International Parallel Processing Symposium (IPPS/SPDP), pp 969-977, San Juan, Puerto Rico, April 12-16, 1999.

2. G. Heber, R. Biswas, P. Thulasiram and Guang R. Gao, Using Multithreading for Automatic Load Balancing of Adaptive Finite Element Meshes, In the Proceedings of Irregular'99 in conjuction to the International Parallel Processing Symposium (IPPS/SPDP), pp 969-977, San Juan, Puerto Rico, April 12-16, 1999.

3. A. Khokhar, G. Heber, Parimala Thulasiraman and Guang R. Gao, Load Adaptive Algorithms and Implementation for the 2D Discret e Wavelet Transform on Fine-Grain Multihtreaded Architctures, In the Proceedings of the International Parallel Processing Symposium (IPPS/SPDP), pp 360-364, San Juan, Puerto Rico, April 12-16, 1999.

4. G. Heber, R. Biswas, and Guang R. Gao, Self-Avoiding Walks over Adaptive Triangular Grids, In Proceedings of SIAM Parallel Processing Conference for Scientific Computing, San Antonio, Texas, April, 1999.

5. Chihong Zhang, R. Govindarajan, and Guang R. Gao, Efficient State-Diagram Construction Methods for Software Pipelining, In Proceedings of the International Conference on Compiler Construction, CC'99, held as part of ETAPS'99, Amsterdam, The Netherland, March 22 – 26, 1999.

6. K. Theobald, Guang R. Gao and T. Sterling, Superconducting Processors for HTMT: Issues and Challenges, In Proceedings of The Seventh Symposium on The Frontiers of Massively Parallel Computation (Frontiers'99), pp 260-267, Annopolis, Maryland, February 21-25, 1999.

7. H. Cai, O. Maquelin, P. Kakulavarapu and Guang R. Gao, Design and Evaluation of Dynamic Load Balancing Schemes under a Fine-Grain Multithreaded Execution Model, In Proceedings of Workshop on Multithreaded Execution, Architecture and Compilation (MTEAC), in conjunction to the 1999 IEEE Symposium on High-Performance Computer Architecture (HPCA99), Orlando, Florida, Janurary, 1999.

8. A. Marquez, K. Theobald, X. Tang and Guang R. Gao, The Superstrand Model, In Proceedings of Workshop on Multithreaded Execution, Architecture and Compilation (MTEAC), in conjunction to the 1999 IEEE Symposium on High-Performance Computer Architecture (HPCA99), Orlando, Florida, Janurary, 1999.

9. Sylvain Lelait, Guang R. Gao and Christine Eisenbeis, A New Fast Algorithm for Optimal Register Allocation in Modulo Scheduled Loops, In Proceedings of the International Conference on Compiler Construction, CC'98, held as part of ETAPS'98, 1998, Kai Koskimies, volume 1383, Lecture Notes in Computer Science, pp 204--218, Springer, Lisbon, Portugal, March 28 - April 4.

10. R. Govindrarajan, Narasimba Rao, E.R. Altman and Guang R. Gao, An Enhanced Co-Scheduling Method using Reduced MS-State Diagrams, In the Proceedings of the International Parallel Processing Symposium (IPPS/SPDP), pp 168-175, Orlando, Florida, April, 1998.

11. Maria-Dana Tarlescu, Kevin Theobald and Guang R. Gao, Elastic History Buffer: A Low Cost Method to Improce Branch Prediction Accuracy, In the Proceedings of the International Conference on Computer Design (ICCD'97), pp 82-87, Austin, TX., Oct. 1997.

12. Rauls Silvera, Jian Wang, Guang R. Gao and R. Govindarajan, A Register Pressure Sensitive Instruction Scheduler for Dynamic Issue Processors, In the Proceedings of the International Conference on Parallel Architecture and Compiler Techniques (PACT'97), San Francisco, CA, Nov. 1997.

13. X,N. Tang, Rakesh Ghiya, Laurie Hendren, Guang R. Gao, Heap Analysis and Optimizations for Threaded Programs, In the Proceedings of the International Conference on Parallel Architecture and Compiler Techniques (PACT'97), San Francisco, CA, Nov. 1997.

14. Xinan Tang, Guang R. Gao, How "Hard" is Thread Partitioning and How "Bad" is a List Scheduling Based Partitioning Algorithm, In Proceedings of Tenth Annual ACM Symposium on Parallel Algorithms and Architectures,Puerto Vallarta, Mexico, pp130--139, June,1998

15. Angela Sodan, Guang R. Gao, Olivier Maquelin, Jens-Uwe Schultz, and Xin-Min Tian. Experience with non-numeric applications on multithreaded architectures. *In Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Las Vegas, Nevada, pp124-135, June, 1997

16. X. N. Tang, J. Wang, K. Theobald, and Guang R. Gao. Thread Partition and Schedule Based on Cost Model. *In Proceedings of the 9th Annual Symposium on Parallel Algorithms and Architectures (SPAA),* Newport, Rhode Island, pp272-281, July 1997.

17. Shashank S. Nemawarkar and Guang R. Gao. Latency tolearance: A metric for performance analysis of multithreaded architecture. In Proceedings of the International Parallel Processing Symposium, April 1997.

18. Parimala Thulasiraman, Xin-Min Tian, and Guang R. Gao. Multithreading implementation of a distributed shortest path algorithm on earth multiprocessor. In Proc. of the Internatinal Conference on High Performance Computing, Trivandrum, India, pp336-341, December 1996.

19. Xin-Min Tian, Shashank S. Nemawarkar, Guang R. Gao, et al. Quantitive studies of data locality sensitivity on the EARTH multithreaded architecture: Preliminary results. In Proc. of the Internatinal Conference on High Performance Computing, Trivandrum, India, pp362-367December 1996.

20. Guang Gao, Konstantin K. Likharev, Paul C. Messina, and Thomas L. Sterling. Hybrid technology multi-threaded architecture. In Proceedings of Frontiers '96: The Sixth Symposium on the Frontiers of Massively Parallel Computation, pages 98-105, Annapolis, Maryland, October 1996.

21. Laurie J. Hendren, Xinan Tang, Yingchun Zhu, Guang R. Gao, Xun Xue, Haiying Cai, and Pierre Ouellet. Compiling C for the EARTH multithreaded architecture. In Proceedings of the 1996 Conference on Parallel Architectures and Compilation Techniques (PACT '96), pages 12-23, Boston, Massachusetts, October 1996. IEEE Computer Society Press.

22. Erik R. Altman and Guang R. Gao. Optimal software pipelining through enumeration of schedules. In Proceedings of Euro-Par'96, pages 833-840, Lyon, France, August 1996.

23. Vivek Sarkar, Guang R. Gao, and Shaohua Han. Data locality analysis for distributed shared memory multiprocessors. In Proceedings of the Ninth Workshop on Languages and Compilers for Parallel Computing, San Jose, California, August 1996.

24. Olivier Maquelin, Guang R. Gao, Herbert H. J. Hum, Kevin B. Theobald, and Xin-Min Tian.Polling Watchdog: Combining polling and interrupts for efficient message handling. In Proceedings of the 23rd Annual International Symposium on Computer Architecture, pages 178-188, Philadelphia, Pennsylvania, May 1996.

25. John Ruttenberg, G. R. Gao, A. Stouchinin, and W. Lichtenstein. Software pipelining showdown: Optimal vs. heuristic methods in a production compiler. In Proceedings of the ACM SIGPLAN '96 Conference on Programming Language Design and Implementation, pages 1-11, Philadelphia, Pennsylvania, May 1996.

26. Vugranam C. Sreedhar, Guang R. Gao, and Yong fong Lee. A new framework for exhaustive and incremental data flow analysis using DJ graphs. In Proceedings of the ACM SIGPLAN '96 Conference on Programming Language Design and Implementation, pages 278-290, Philadelphia, Pennsylvania, May 1996.

27. Jian Wang and Guang R. Gao. Pipelining-dovetailing: A transformation to enhance software pipelining for nested loops. In Proceedings of the 6th International Conference on Compiler Construction, Lecture Notes in Computer Science, Linkoping, Sweden, April 1996. Springer-Verlag.

28. R. Govindarajan, Erik R. Altman, and Guang R. Gao. Instruction scheduling in the presence of structureal hazards: An integer programming approach to software pipeline. In Proc. of the nternatinal Conference on High Performance Computing, Goa, India, December 1995.

29. R. Govindarajan, Erik R. Altman, and Guang R. Gao. Co-scheduling hardware and software pipelines. In Second International Symposium on High-Performance Computer Architecture, San Jose, California, February 1996.

30. Shashank S. Nemawarkar and Guang R. Gao. Measurement and modeling of EARTH-MANNA multithreaded architecture. In Proceedings of the Fourth International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, pages 109-114, San Jose, California, February 1996. IEEE Computer Society TCCA and TCS.

31. Luis A. Lozano C. and Guang R. Gao. Exploiting short-lived variables in superscalar processors. In Proceedings of the 28th Annual International Symposium on Microarchitecture, pages 292-302, Ann Arbor, Michigan, November-December 1995.

32. J. B. Dennis and G.R. Gao. On memory models and cache management for shared-memory multi-processors. In Proceedings of Seventh IEEE International Symposium on Parallel and Distributed Proccesing. IEEE, October 1995.

33. Olivier C. Maquelin, Herbert H. J. Hum, and Guang R. Gao. Costs and benefits of multithreading with off-the-shelf RISC processors. In Proceedings of the First International EURO-PAR Conference, number 966 in Lecture Notes in Computer Science, pages 117-128, Stockholm, Sweden, August 1995. Springer-Verlag.

34. R. Wen, Guang R. Gao, and Vincent V. Dongen. The design and implementation of the accurate array data-flow analysis in the HPC compiler. In Proceedings of High Performance Computing Symposium '95, Canada's Ninth Annual International High Performance Computing Conference and Exhibition, pages 144-155, Montr'eal, Qu'ebec, July 1995. Centre de recherche informatique de Montr'eal.

35. Nasser Elmasri, Herbert H. J. Hum, and Guang R. Gao. The Threaded Communication Library: Preliminary experiences on a multiprocessor with dual-processor nodes. In Conference Proceedings, 1995 International Conference on Supercomputing, pages 195-199, Barcelona, Spain, July 1995.

36. Erik R. Altman, R. Govindarajan, and Guang R. Gao. An experimental study of an ILP-based exact solution method for software pipelining. In Proceedings of the 8th International Workshop on Languages and Compilers for Parallel Computing, Lecture Notes in Computer Science, pages 2.1 - 2.15, Columbus, Ohio, August 1995. Springer-Verlag.

37. Guang R. Gao and Vivek Sarkar. Location consistency: Stepping beyond the memory coherence barrier. In 24th International Conference on Parallel Processing, pages II-73-II-76, University Park, Pennsylvania, August 1995.

38. Herbert H. J. Hum, Olivier Maquelin, Kevin B. Theobald, Xinmin Tian, Xinan Tang, Guang R. Gao, Phil Cupryk, Nasser Elmasri, Laurie J. Hendren, Alberto Jimenez, Shoba Krishnan, Andres Marquez, Shamir Merali, Shashank S. Nemawarkar, Prakash Panangaden, Xun Xue, and Yingchun Zhu. A design study of the EARTH multiprocessor. In Proceedings of the IFIP WG 10.3 Working Conference on Parallel Architectures and Compilation Techniques, PACT '95, pages 59-68, Limassol, Cyprus, June 1995. ACM Press.

39. E. R. Altman, R. Govindarajan, and G. R. Gao. Scheduling and mapping: Software pipelining in the presence of structural hazards. In ACM SIGPLAN Symposium on Programming Language Design and Implementation, June 1995. page 139-150.

40. G. Tremblay and G. R. Gao. The impact of laziness on parallelism and the limits of strictness analysis. In Proceedings of the High Performance Functional Computing Conference, pages 119- 133, Denver, Colorado, April 1995. Lawrence Livermore National Laboratory. CONF-9504126.

41. Vugranam C. Sreedhar and Guang R. Gao. A linear time algorithm for placing OE-nodes. In Conference Record of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pages 62-73, San Francisco, California, January 1995.

42. Vugranam C. Sreedhar, Guang R. Gao, and Yong fong Lee. Incremental computation of dominator trees. In Proceedings of the ACM SIGPLAN Workshop on Intermediate Representations (IR'95), pages 1-12, San Francisco, California, January 22, 1995. SIGPLAN Notices, 30(3), March 1995.

43. Kevin B. Theobald, Herbert H. J. Hum, and Guang R. Gao. A design framework for hybrid-access caches. In Proceedings of the First International Symposium on High-Performance Computer Architecture, pages 144-153, Raleigh, North Carolina, January 1995.

44. R. Govindarajan, Erik R. Altman, and Guang R. Gao. Minimizing register requirements under resource-constrained rate-optimal software pipelining. In Proceedings of the 27th Annual International Symposium on Microarchitecture, pages 85-94, San Jose, California, November-December 1994.

45. R. Govindarajan, Erik R. Altman, and Guang R. Gao. A framework for resource-constrained rate optimal software pipelining. In Proceedings of the Third Joint International Conference on Vector and Parallel Processing (CONPAR 94 - VAPP VI), number 854 in Lecture Notes in Computer Science, pages 640-651, Linz, Austria, September 1994. Springer-Verlag.

46. R. Govindarajan, Guang R. Gao, and Palash Desai. Minimizing memory requirements in rate optimal schedules. In Proceedings of the 1994 International Conference on Application Specific Array Processors, pages 75-86, San Francisco, California, August 1994. IEEE Computer Society.

47. S. S. Nemawarkar, R. Govindarajan, G. R. Gao, and V. K. Agarwal. Performance of interconnection network in multithreaded architectures. In Proceedings of PARLE '94 - Parallel Architectures and Languages Europe, number 817 in Lecture Notes in Computer Science, pages 823-826, Athens, Greece, July 1994. Springer-Verlag.

48. V. Van Dongen, C. Bonello, and Guang R. Gao. Data parallelism with High Performance C. In Proceedings of Supercomputing Symosium '94, Canada's Eighth Annual High Performance Computing Conference, pp 128-135, Toronto, Ontario, June 1994. University of Toronto.

49. Herbert H. J. Hum, Kevin B. Theobald, and Guang R. Gao. Building multithreaded architectures with off-the-shelf microprocessors. In Proceedings of the 8th International Parallel Processing Symposium, pp 288-294, Cancun, Mexico, April 1994. IEEE Computer Society.

50. G. Liao, E.R. Altman, V.K. Agarwal, and Guang R. Gao. A comparative study of DSP multiprocessor list sheduling heuristics. In Proceedings of the 27th Annual Hawaii International Conference on System Sciences, Kihei, Hawaii, 1994.

51. S. S. Nemawarkar, R. Govindarajan, Guang R. Gao, and V. K. Agarwal. Analysis of multithreaded multiprocessors with distributed shared memory. In Proceedings of the Fifth IEEE Symposium on Parallel and Distributed Processing, pp 114-121, Dallas, Texas, December 1993.

52. R. Govindarajan and Guang R. Gao. A novel framework for multi-rate sheduling in DSP applications. In Proceedings of the 1993 International Conference on Application Specific Array Processors, pp 77-88, Venice, Italy, October 1993. IEEE Computer Society.

53. Guang R. Gao, Vivek Sarkar, and Lelia A. Vazquez. Beyond the data parallel paradigm: Issues and options. In W.K. Giloi, S. Jahnichen, and B.D. Shriver, Editors Proceedings – 1993 Programming Models for Massively Parallel Computers, pp 191-197, Berlin, Germany, September 20-23, 1993. IEEE Computer Society Press.

54. Guang R. Gao, Qi Ning, and Vincent Van Dongen. Extending software pipelining techniques for scheduling nested loops. In Proceedings of the 6th International Workshop on Languages and Compilers for Parallel Computing, number 768 in Lecture Notes in Computer Science, pp 340-357, Portland, Oregon, August 1993. Springer-Verlag.

55. Erik R. Altman, Vonod K. Agarwal, and Guang R Gao. A novel methodology using genetic algorithms for the design of caches and cache replacement policy. In Stephanie Forrest, editor, Proceedings of the 5th International Conference on Genetic Algorithms, pp 392-399. Morgan Kaufmann Publishers, Inc., July 1993. University of Illinois at Urbana-Champaign.

56. Kevin B. Theobald, Guang , and Laurei J. Hendren. Speculative execution and branch prediction on parallel machines. In Conference Proceedings, 1993 ACM International Conference on Supercomputing, pp 77-86, Tokyo, Japan, July 1993.

57. Robert Kim Yates and Guang R. Gao, A Kahn principle for networks of nonmonotonic real-time processes. In Proceedings of PARLE '93 – Parallel Architectures and Languages Europe, number 694 in Lecture Notes in Computer Science, pp 209-227, Munich, Germany, June 1993. Springer-Verlag.

58. Herbert H. J. Hum and Guang R. Gao. Supporting a dynamic PMD model in a multi-threaded architecture. In Digest of Papers, 38th IEEE Computer Society International Conference, COMPCON Spring '93, pp 165-174, San Francisco, California, February 1993.

59. Qi Ning and Guang R. Gao, A novel framework of register allocation for software pipelining. In Conference Record of the Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp 29-42, Charleston, South Carolina, January 1993.

60. Kevin B. Theobald, Guang R. Gao, and Laurie J. Hendren. On the limits of program parallelism and its smoothability. In Proceedings of the 25th Annual International Symposium on Microarchitecture, pp 10-19, Portland, Oregon, December 1992.

61. V. Van Dongen, Guang R. Gao, and Q. Ning. A polynomial time method for optimal software pipelining. In Proceedings of the Conference on Vector and Parallel Processing, CONPAR-92, number 634 in Lecture Notes in Computer Science, pp 613-624, Lyon, France, September 1-4, 1992. Springer-Verlag.

62. J. M. Monti and Guang R Gao. Efficient interprocessor synchronization and communication on a dataflow multiprocessor architecture. In the Proceedings of 1992 International Conference on Parallel Processing, pp I-220-224, St. Charles, IL, August 1992.

63. Guang R Gao, R. Olsen, V. Sarkar, and R. Thekkath. Collective loop fusion for array contraction. In Proceedings fo the 5th International Workshop on Languages and Compilers for Parallel Computing, number 757 in Lecture Notes in Computer Science, pp 281-295, New Haven, Connecticut, August 1992. Springer-Verlag.

64. L. Hendren, C. Donawa, M. Emami, Guang R Gao, Justiani, and B. Sridharan. Designing the McCAT compiler based on a family of structured intermadiate representations. In Proceedings of the 5th International Workshop on Languages and Compilers for Parallel Computing, number 757 in Lecture Notes in Computer Science, pp 406-420, New Haven, Connecticut, August 1992. Springer-Verlag.

65. Herbert H. J. Hum, Kevin B. Theobald, and Guang R. Gao. Building multithreaded architectures with off-the-shelf microprocessors. In Proceedings of the 8th International Parallel Processing Symposium, pages 288-294, Canc'un, Mexico, April 1994. IEEE Computer Society.

Monographs, Books and Book Chapters

1. G. R. Gao., J-L. Gaudiot, and L. Bic, editors. Advanced Topics in Dataflow and Multithreaded Computers. IEEE Computer Society Press, 1995.

2. Jack B. Dennis and Guang R. Gao. Multithreaded architectures: Principles, projects, and issues. In Robert A. Iannucci, Guang R. Gao, Robert H. Halstead, Jr., and Burton Smith, editors, Multithreaded Computer Architecture: A Summary of the State of the Art, chapter 1, pages 1-72. Kluwer Academic Publishers, Norwell, Massachusetts, 1994.

3. Robert A. Iannucci, Guang R. Gao, Robert H. Halstead, Jr., and Burton Smith, editors. Multi-threaded Computer Architecture: A Summary of the State of the Art. Kluwer Academic Publishers, Norwell, Massachusetts, 1994. Book contains papers presented at the Workshop on Multithreaded Computers, Albuquerque, New Mexico, November 1991.

4. G. R. Gao. A Code Mapping Scheme for Dataflow Software Pipelining. Kluwer Academic Publishers, Boston, Massachusetts, December 1990.

B.3 Research Significance

The theme of my research in computer architecture and systems, compiler technology, and memory models not only enriches the field of parallel computing and encompass a host new techniques for high-performance architectures and compiling technology but also provides a new horizon for mapping applications, both regular or irregular, onto these architectures. Furthermore, the research activities are not only themselves intellectually stimulating, interesting and competitive, but also exposes students with a dynamic new field with excellent prospect of employment and a productive career.

My work on EARTH model and architecture has important relevance to the design and development of future generation of parallel computer architectures. The research results have been published widely in a range of recognized international professional conferences and journals. It has attracted a considerable level of research support from NSF through 4 NSF research grants encompassing the areas from architecture and memory support, the efficient implementation of multithreaded execution models on SMP workstation cluster based parallel systems, the application of EARTH model to large irregular applications such as the crack propagation problem, and the compilation technology for multithreading. It has also attracted industry interests and funding such as the DRP grant we received with support from ACORN Inc. An extension of our work on fine-grain multithreading and EARTH to be applied to high-end Supercomputing has become an important component to the HTMT project, one of the few nation's on-going petaflow architecture project funded by DARP, NSA and NASA.

My work on modulo scheduling and software pipelining also have immediate relevance to the computer industry in their effort to exploit high performance with instruction level parallelism. The research results have been published widely in a range of recognized international professional conferences and journals. The technology developed in our group has been used in the evaluation of the software pipelining techniques in the SGI production compiler, and to foster the future collaboration, we have received the donation of two SGI workstations with special SGI software.

The co-scheduling technique has been funded by NSF through a research grant. The co-scheduling technology developed by me and my colleagues have also attracted strong industry attention, and Rockwell Semiconductor Systems has already committed funding to this research and a DRP grant on retargetable compiler for DSP architectures with Rockwell funding and university matching has just been awarded.

The significance and novelty of my work on program analysis and memory models have also been recognized by the research community. Three papers out of the work on program Analysis have been accepted for publications on the most prestigious journal -- the ACM Transactions on Program Languages and Systems.

B.4 Research Support

Agency	Grant Number	Title	Amount	Period	Status
NSF	CCR 9808522	Compiling Irregular Applications on a Multithreaded Architecture	$319,156	08/97-07/00	co-PI
NSF	MIPS 9707125	A New Generation Multithreaded Processors	$400,000	07/97-06/00	co-I
NSF	CDA 9703088	Parallel and Distributed Computing: Systems and Applications Development	$633,513	07/97-06/02	co-PI
	(Infrastructure Grant)				
DARPA/ NSA/ NASA	ASC 9612105	Hybrid Technology Multithreaded Architecture for Petaflops	$800,000	06/97-05/99	co-I
NSF	CCR 9711477	A Framework of Modulo Scheduling Based on Finite Automaton	$139,263	06/97-05/99	PI
		(with REO	$ 6,250	06/97-05/98	
NSF	CISE 9726388	Challenges in CISE-Crack Propagation Project	$264,952	01/98-12/00	co-I
DRP	Approved	Retargetable Compilers for Embedded DSP Processors	$75,000	98-00	PI
	(with Rockwell Semiconductor Systems Inc.)				

Section C: Services

C.1 University Activities and Services

- Special Activities:
  - attended recruiting activities of new faculty members
  - the tenure review of Prof. Dan Van Weide, Prof. Paul Berger
  - participating faculty retreat meeting (1998)
  - Dean's ad hoc group for supercomputing (1998)
  - Participate Engieering Outreach program
  - An advisor in the university Undergraduate Research Opportunity program

- Departmental and College Committees

  - chairing the departmental Committee on Promotion & Tenure (1998)
  - College Election Committee (1998)

- University Committees

  - the ICRSS committee (Instructional, Computing and Research Support Services Committee)

C.2 Profession Services

- **IEEE Computer Society Distringuished Visitor, 1998-2001**

- **IEEE, Senior Member (since 1997)**

- **Program Committee Members of Recognized International Conferences**

  - IEEE International Symposium on Computer Architecture (HPCA-95, HPCA-99)
  - ACM Symposium on Programming Language Design and Implementation (PLDI'98)
  - ACM International Conference on Supercomputing (ICS-95)
  - ACM/IEEE International Symposium on Microarchitectures (MICRO-95, 96, 97)
  - International Parallel Processing Symposium (IPPS'95)
  - IFIP and ACM SIGARCH International Conference on Parallel Architectures and Compilation Techniques (PACT'94,95,96,97,98)
  - International Conference on Algorithms And Architectures for Parallel Processing (ICAPP-95)
  - Parallel Architecture and Language Europe (PARLE-91,92,93,94,95)
  - International Conference on Parallel Processing (EURO-PAR-95,96)
  - Working Conference on Massively Parallel Programming Models (MPPM-93,95,97,99)
  - High Performance Computing Symposium (HPCS-95, 96, 98).

- **Program Committee Chairmanship**

- I have been elected as the Program Chairman of the 1994 ACM SIGARCH, International Conference on parallel Architectures and Compilation Techniques (PACT '94), Aug/. 1984. Montreal, Canada, co-sponsored by IFIP and in association with ACM SIGPLAN, IEEE TCCA (Technical Committee on Computer Architecture) and IEEE TCPP (Technical Committee on Parallel Processing).

- I am elected as the General Co-Chair of the 1998 International Conference on Parallel Architectures and Compilation Techniques (PACT '98), Oct. 1998, Paris, France., co-sponsored by IFIP and IEEE Computer Society

- I am elected as the Chair of the Third Workshop on Petaflop Computing, Feb., 1999. Annapolis, MD.

- **Other Activities in Recognized Professional Conferences**

I have served as a workshop chair, a session chair, an organizing committee or steering committee member of many international conferences.

- **Journal Editorialship**

- I am elected to the Editorial Board of IEEE Transactions on Computers (1998 -)
- I am elected to the Editorial Board of IEEE Concurrency Journal (1997 -)
  - I joined the Editorial Board of the Journal on Programming Languages in Jan. 1996, and subsequently became one of the two Co-Editors of the journal.
  - I am a Guest Editor for the Special Issue on Dataflow and Multithreaded Computers, Journal of Parallel and Distributed Computing, Academic Press, June, 1993.

- **Invited Seminars and Distinguished Seminars**

  I have given seminars in many industrial and academic organizations: IBMT.J. Watson Research Center, IBM Toronto Lab, AT&T Bell Laboratories, BNR, HP Labs, SGI, DEC, NRL(Navy Research Lab.), MIT, Stanford, UC Berkeley, University of Victoria are just named a few.

- **Others**: A panelist, session chair, organization/steering committee member, advisory board member for many recognized professional conferences. (detail to be provided upon request)