

P.O. BOX 704 • YORKTOWN HEIGHTS, NEW YORK 10598
PHONE (914) 784-7811 • E-MAIL DFB@WATSON.IBM.COM
URL WWW.RESEARCH.IBM.COM/PEOPLE/D/DFB

DAVID F. BACON

RESEARCH INTERESTS

Programming languages, run-time systems, garbage collection, compilers, real-time and embedded systems, concurrent and distributed systems, massively parallel supercomputers, and instruction set architecture. Raising the level of abstraction at which computer systems are described and programmed. Quantitative, scientific evaluation.

EDUCATION

Ph.D. in Computer Science, University of California, Berkeley, 1997

Dissertation: *Fast and Effective Optimization of Statically Typed Object-Oriented Languages*

Advisor: Susan L. Graham

A.B. in Computer Science, Columbia College, Columbia University, 1985

Advisor: Yechiam Yemini

H.S. Diploma, Hunter College High School, New York, 1981

PROFESSIONAL EXPERIENCE

IBM T.J. Watson Research Center

1997 – Present Research Staff Member

- Led projects in garbage collection, language design, virtual machine run-time systems, and compiler optimization.
- Led group in developing Metronome, a breakthrough technology in provably real-time garbage collection. Demonstrated that the prevailing real-time garbage collection methodology in use since 1978 was fundamentally flawed, and corrected by our algorithm. Second generation system is now being built in IBM's production virtual machine.
- Developed thin lock technology which sped up Java applications by an average of 25%. Now in use in virtually every Java virtual machine. Received ACM "PLDI Most Influential Paper" award.
- Successfully transferred runtime and compiler technologies into key IBM products.
- Filed 18 patent applications (5 issued to date).

1994 – 1997 Senior Staff Programmer

- Continued Ph.D. thesis work while at IBM as part of the Resident Study Fellowship.
- Implemented analysis algorithms from my thesis in an extensible C++ compiler.

- Optimized run-time system for IBM's C and C++ compilers to improve TLB mappings, leading to speedups of up to 20%.

University of California, Berkeley

1990-1997 *Doctoral Candidate*

- Developed Rapid Type Analysis, a core optimization algorithm for many compilers of object-oriented languages. Now the basis for two IBM products.
- Developed new networking kernel for CM-5 supercomputer; adopted by Thinking Machines.
- Worked with Astrophysics department team to obtain a factor of 40 speedup of an X-ray pulsar simulation.
- Developed algorithm for deterministic replay of parallel programs using bus-snooping.

IBM Santa Teresa Laboratory

1994 *Visiting Researcher*

- Developed graphical system for exploring application performance at the micro-architectural level with hardware performance monitor data from the POWER2 chip.
- Used tool to find several performance anomalies, including *cache jamming*, and developed compiler optimization techniques to prevent it.

IBM T.J. Watson Research Center

1985 – 1990 *Staff Programmer*

- Co-designer of the Hermes distributed programming language, which pioneered pointer-safe programming with an integrated lightweight process model and high-level data types.
- Led implementation of Hermes and distribution as an open source project.
- Developed and prototyped new algorithms for fault-tolerant and optimistic parallel computing.

DCC Laboratory, Columbia University Computer Science Department

1984-1985 *Research Associate*

- Created NEST (Network Simulation Testbed), widely used with a commercial version.
- Used NEST to prototype an incremental mobile position location system.

NON-VON Project, Columbia University Computer Science Department

1982-1983 *Research Associate*

- Designed and implemented NON-VON Pascal, a language for programming a custom VLSI tree-structured SIMD parallel supercomputer.

Columbia University Center for Computing Activities

1982-1985 *Manager of Student Consulting*

- Managed scheduling, hiring, and training of a staff of 40 student consultants.

Academic Directions Inc.

1980-1981 *Programmer*

- Startup. Implemented disk caching, forms processing, and recursion for FORTRAN.

Columbia Research Institute, Inc.

1979-1980 *Programmer*

- Maintained business systems on mini-computers; built Emacs-style editor for RT-11.

PUBLICATIONS

BOOK

B1. **Hermes: A Language for Distributed Computing**

Robert E. Strom, David F. Bacon, Andy Lowry, Arthur P. Goldberg, Daniel M. Yellin, and Shaula Yemini
Prentice-Hall, Series in Innovative Technology, 1991.

INVITED ARTICLE

I1. **Retrospective: Thin Locks**

David F. Bacon, Ravi Konuru, Chet Murthy, and Mauricio Serrano
Twenty Years of the ACM SIGPLAN Conference on Programming Language Design and Implementation: A Selection (2003), to appear.

JOURNAL ARTICLES

J1. **A Pure Reference Counting Garbage Collector**

David F. Bacon, V.T. Rajan, Clement R. Attanasio, Stephen E. Smith, and Han B. Lee
ACM Transactions on Programming Languages and Systems, to appear.

J2. **Kava: A Java Dialect with a Uniform Object Model for Lightweight Classes**

David F. Bacon
Concurrency—Practice and Experience, to volume 15, numbers 3-5, March-April 2003, pp. 185-206.

J3. **Compiler Transformations for High-Performance Computing**

David F. Bacon, Susan L. Graham, and Oliver Sharp
ACM Computing Surveys, volume 26 number 4, December 1994.

J4. **NEST: A Network Simulation and Prototyping Tool**

Alexander Dupuy, Jed Schwartz, Yechiam Yemini, and David F. Bacon
Communications of the ACM, volume 33 number 10, October 1990.

REFEREED ARTICLES IN CONFERENCE PROCEEDINGS

- C1. **A Unified Theory of Garbage Collection**
David F. Bacon, Perry Cheng, and V.T. Rajan
OOPSLA'04 Conference Proceedings: Object-Oriented Systems, Languages, and Applications (Vancouver, British Columbia, October 2004).
- C2. **Dynamic Selection of Application-Specific Garbage Collectors**
Sunil Soman, Chandra Krintz, and David F. Bacon
Proceedings of the International Symposium on Memory Management (Vancouver, British Columbia, October 2004).
- C3. **Write Barrier Elision for Concurrent Garbage Collectors**
Martin Vechev and David F. Bacon
Proceedings of the International Symposium on Memory Management (Vancouver, British Columbia, October 2004).
- C4. **Garbage Collection for Embedded Systems**
David F. Bacon, Perry Cheng, and David Grove
Proceedings of the Fourth ACM International Conference on Embedded Software (Pisa, Italy, September 2004).
- C5. **A Real-time Garbage Collector with Low Overhead and Consistent Utilization**
David F. Bacon, Perry Cheng, and V.T. Rajan
Conference Record of the Thirtieth ACM Symposium on Principles of Programming Languages (New Orleans, Louisiana, January 2003), pp. 285-298
- C6. **Controlling Fragmentation and Space Consumption in the Metronome, a Real-time Garbage Collector for Java**
David F. Bacon, Perry Cheng, and V.T. Rajan
Proceedings of the Conference on Languages, Compilers, and Tools for Embedded Systems (San Diego, California, June 2003). ACM SIGPLAN Notices Volume 38, number 7, July 2003, pp. 81-92.
- C7. **MJ: A Rational Module System for Java and its Applications**
John Corwin, David F. Bacon, David Grove, and Chet Murthy
OOPSLA'03 Conference Proceedings: Object-Oriented Systems, Languages, and Applications, ACM SIGPLAN Notices, volume 38, number 11, October 2003 (Anaheim, California), pp. 241-254.
- C8. **Space- and Time-Efficient Implementation of the Java Object Model**
David F. Bacon, Stephen J. Fink, and David Grove
Proceedings of the Sixteenth European Conference on Object-Oriented Programming (Málaga, Spain, June 2002), B. Magnusson, editor, Lecture Notes in Computer Science vol. 2374, pp. 111-132.

- C9. **Java without the Coffee Breaks: A Non-intrusive Multiprocessor Garbage Collector**
David F. Bacon, Clement R. Attanasio, Han B. Lee, V.T. Rajan, and Stephen E. Smith
Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (Snowbird, Utah, June 2001), ACM SIGPLAN Notices, volume 36, number 5, May 2001, pp. 92-103.
- C10. **Concurrent Cycle Collection in Reference Counted Systems**
David F. Bacon and V.T. Rajan
Proceedings of the Fifteenth European Conference on Object-Oriented Programming (University Eötvös Loránd, Budapest, Hungary, June 2001), J.L. Knudsen, editor, Lecture Notes in Computer Science vol. 2072, pp. 207-235.
- C11. **Kava: A Java Dialect with a Uniform Object Model for Lightweight Classes**
David F. Bacon
Proceedings of the Joint ACM Java Grande/ISCOPE Conference, pp. 68-77, Stanford, California, June 2001
- C12. **Guava: A Dialect of Java without Data Races**
David F. Bacon, Robert E. Strom, and Ashis Tarafdar
OOPSLA'00 Conference Proceedings: Object-Oriented Systems, Languages, and Applications (Minneapolis, Minnesota, October 2000), ACM SIGPLAN Notices, volume 35, number 10, pp. 382-400.
- C13. **Thin Locks: Featherweight Synchronization for Java**
David F. Bacon, Ravi Konuru, Chet Murthy, and Mauricio Serrano
Proceedings of the ACM Conference on Programming Language Design and Implementation (Montreal, Canada, June 1998), ACM SIGPLAN Notices volume 33 number 6, pp. 258-268.
- C14. **Fast Static Analysis of C++ Virtual Function Calls**
David F. Bacon and Peter Sweeney
OOPSLA'96 Conference Proceedings: Object-Oriented Programming Systems, Languages, and Applications (San Jose, California, October 1996), ACM SIGPLAN Notices volume 31 number 10, pp. 324-341
- C15. **A Compiler Framework for Restructuring Data Declarations to Enhance Cache and TLB Effectiveness**
David F. Bacon, Jyh-Herng Chow, Dz-ching R. Ju, Kalyan Muthukumar, and Vivek Sarkar
CASCON '94, Toronto, Canada.
- C16. **Optimistic Parallelization of Communicating Sequential Processes**
David F. Bacon and Robert E. Strom
Proceedings of the Third ACM Symposium on Principles and Practices of Parallel Programming (Williamsburg, Virginia, April 1991), SIGPLAN Notices volume 26, number 7, July 1991, pp. 155-166.

- C17. **File System Measurements and their Application to the Design of Efficient Operation Logging Algorithms**
David F. Bacon
Proceedings of the Tenth IEEE Symposium on Reliable Distributed Systems (Pisa, Italy, September 1991), pp. 21-30.
- C18. **High-Level Language Support for Programming Distributed Systems**
Joshua Auerbach, David F. Bacon, Arthur P. Goldberg, G. Goldszmidt, Ajei Gopal, Mark Kennedy, Andy Lowry, James R. Russell, William Silverman, Robert E. Strom, Daniel M. Yellin and Shaula Yemini
Proceedings of the International Conference on Computer Languages (Oakland, California, April 1992), pp. 320-330.
- C19. **A Portable Run-Time System for the Hermes Distributed Programming Language**
David F. Bacon and Andy Lowry
Usenix Conference Proceedings, Summer 1990, Anaheim, California, pp. 39-50.
- C20. **Volatile Logging in n-Fault-Tolerant Distributed Systems**
Robert E. Strom, David F. Bacon, and Shaula Yemini
Proceedings of the Eighteenth International Symposium on Fault Tolerant Computing, IEEE Computer Society, 1988
- C21. **A Recoverable Object Store**
Robert E. Strom, Shaula Yemini, and David F. Bacon
Proceedings of the Hawaii International Conference on Systems Sciences, IEEE Computer Society, 1988.
- C22. **Toward Self-Recovering Operating Systems**
Robert E. Strom, Shaula Yemini, and David F. Bacon
Proceedings of the International Conference on Parallel Processing, North-Holland, 1987.

REFEREED ARTICLES IN WORKSHOP PROCEEDINGS

- W1. **Braids and Fibers: Language Constructs with Architectural Support for Adaptive Response to Hardware Latencies**
David F. Bacon and Xiaowei Shen
The First Watson Conference on Interaction between Architecture, Circuits, and Compilers, Yorktown Heights, New York, October 2004
- W2. **The Metronome: A Simpler Approach to Garbage Collection in Real-time Systems**
David F. Bacon, Perry Cheng, and V.T. Rajan
Workshop on Java Technologies for Real-Time and Embedded Systems (Catania, Sicily, November 2003), Proceedings of the OTM Workshops, R. Meersman and Z. Tari, eds., Lecture Notes in Computer Science vol. 2889, pp. 466-478.

- W3. **A Comparative Evaluation of Parallel Garbage Collectors**
Clement R. Attanasio, David F. Bacon, A. Cocchi, and Stephen E. Smith
Proceedings of the Fourteenth Annual Workshop on Languages and Compilers for Parallel Computing (Cumberland Falls, Kentucky, August 2001), Lecture Notes in Computer Science, Springer-Verlag.
- W4. **Hardware-Assisted Replay of Multiprocessor Programs**
David F. Bacon and Seth Copen Goldstein
Proceedings of the ACM SIGPLAN/SIGOPS Workshop on Parallel and Distributed Debugging (Santa Cruz, California, May 1991), SIGPLAN Notices volume 26, number 12, December 1991, pp. 194-206.
- W5. **Transparent Recovery in Distributed Systems**
David F. Bacon
Proceedings of the Fourth ACM SIGOPS European Workshop on Reliability in Distributed Systems (Bologna, Italy, September 1990), Operating Systems Review, volume 25, number 2, April 1991, pp. 1-4.
- W6. **Transparent Recovery of Mach Applications**
Arthur P. Goldberg, David F. Bacon, Ajei Gopal, Kong Li, and Robert E. Strom
Proceedings of the 1990 Usenix Mach Workshop.

MAGAZINE ARTICLES

- M1. **Measure for Measure**
Oliver Sharp and David F. Bacon
Byte, volume 19 number 10, October 1994.
- M2. **Cache Advantage**
David F. Bacon
Byte, volume 19 number 8, August 1994

PATENTS

PATENTS ISSUED

- P1. **Method to Provide Concurrency Control over Objects Without Atomic Operations on Non-Shared Objects**
D. Bacon and S. Fink
United States Patent number 6772153, issued August 3, 2004.
- P2. **Removal of Unreachable Methods in Object-oriented Applications Based on Program Interface Analysis**
David F. Bacon, Chris Laffra, Peter F. Sweeney, and Frank Tip
United States Patent number 6654951, issued 11 November 2003.

- P3. **Method for Determining Reachable Methods in Object-Oriented Applications that use Class Libraries**
David F. Bacon, Chris Laffra, Peter F. Sweeney, and Frank Tip
United States Patent number 6,463,581, issued 8 October 2002.
- P4. **Locking and Unlocking Mechanism for Controlling Concurrent Access to Objects**
David F. Bacon
United States Patent number 6,247,025, issued 12 June 2001.
- P5. **Object-Oriented Dispatch Optimization**
David F. Bacon, Mark Wegman, and F. Kenneth Zadeck
United States Patent number 6,041,179, issued 21 March 2000.

PATENTS PENDING

- PP1. **Method for Reducing Object Size in an Object-Oriented System**
David F. Bacon, Perry Cheng, and David Grove
Filed 25 September 2004.
- PP2. **Hardware-assisted Compression of Addresses and Data**
David F. Bacon, Perry Cheng, and David Grove
Filed 2 October 2004.
- PP3. **Cache Coloring**
David F. Bacon, Orran Krieger, and Xiaowei Shen
Filed 15 December 2003.
- PP4. **Split-Phase Instructions for Avoidance of Memory Latency**
David F. Bacon and Xiaowei Shen
Filed 15 February 2003.
- PP5. **Method for Optimizing Allocation and Access of Multi-level Objects**
Matthew Arnold, David F. Bacon, and Perry Cheng
Filed 1 March 2003.
- PP6. **Optimization of Software Read Barriers**
David F. Bacon and Perry Cheng
Filed 10 February 2003.
- PP7. **Incremental Defragmentation in a Mostly Non-moving Garbage Collector**
David F. Bacon, Perry Cheng, and V.T. Rajan
Filed 10 February 2003.
- PP8. **Space-Efficient Object Models for Object-Oriented Programming Languages**
David F. Bacon, Stephen J. Fink, and David Grove
Filed 10 December 2002.

- PP9. **Method for Detection of Unused Cyclic Data Structures in Linear Time**
David F. Bacon
Filed 11 December 2000.
- PP10. **Method for Concurrent Detection of Unused Cyclic Data Structures**
David F. Bacon and V.T. Rajan
Filed 10 December 2002.
- PP11. **Space-Efficient Object Models for Object-Oriented Programming Languages**
David F. Bacon, Stephen J. Fink, and David Grove
Filed 11 December 2000.
- PP12. **A Serially Reusable Virtual Machine**
David F. Bacon, Robert Berry, Rajesh Bordawekar, Donna Dillenberger, Beth Hutchison, S. Paice, D. Schmidt, A. Webb, and J. West
Filed 1999.
- PP13. **Method for Efficient Program Transformations for Object-Oriented Applications that use Class Libraries**
David F. Bacon, Chris Laffra, Peter F. Sweeney, and Frank Tip
Filed 14 December 1998.

PROFESSIONAL ACTIVITIES

CONFERENCE ORGANIZATION

- CO1. Steering Committee Chair, International Symposium on Memory Management (ISMM).
- CO2. General Chair, International Symposium on Memory Management (ISMM), Vancouver, British Columbia, October 2004.
- CO3. Organizing Committee, Workshop on the Future of the Java Language, October 2004.

PROGRAM COMMITTEES

- PC1. Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA), San Diego, California, October 2005.
- PC2. Third ACM SIGPLAN Workshop on Memory Systems Performance (MSP), Chicago, Illinois, June 2005.
- PC3. Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA), Vancouver, British Columbia, October 2004.
- PC4. The First Watson Conference on Interaction between Architecture, Circuits, and Compilers (PAC2), Yorktown Heights, New York, October 2004.
- PC5. Workshop on Concurrency and Synchronization in Java Programs (CSJP), St. John's, Newfoundland, July 2004.
- PC6. Symposium on Principles of Programming Languages (POPL), Venice, Italy, January 2004.
- PC7. International Workshop on Types in Programming Language Design and Implementation (TLDI), New Orleans, Louisiana, January 2003.

- PC8. Java Grande/ISCOPE Conference, Seattle, Washington, October 2002.
- PC9. Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA), Seattle, Washington, October 2002.
- PC10. International Symposium on Memory Management (ISMM), Berlin, Germany, June 2002.

PROGRAM CHAIR

- CH1. Second Annual Programming Languages Day, IBM T.J. Watson Research Center, Hawthorne, New York, April 2001.

GRANT REVIEW PANELS

- GR1. IBM Eclipse Innovation Grants, November 2004.
- GR2. NSF Information Technology Research program, May 2002.

TECHNOLOGY TRANSFER

- TT1. Thin Locks (also known as “Bacon bits”) implemented in IBM’s AIX Java virtual machine and adopted by product team for all IBM Java products. Also used by Intel Open Runtime Platform, Gnu GCJ 3.1, Mono .NET runtime, Stanford joeq, Excelsior LLC JET, McGill JVM, and IBM Jikes Research VM.
- TT2. Key components of the Metronome garbage collection technology implemented in and adopted for IBM’s Java 2 Micro Edition (J2ME) virtual machine, achieving major improvements in space consumption, processing time, and code density.
- TT3. Space- and time-efficient object model implemented in and adopted for IBM’s next-generation virtual machine product, leading to significant improvements in SPEC benchmark performance.
- TT4. Serially Reusable Virtual Machine shipped in the IBM mainframe versions of the Java Virtual Machine.
- TT5. Rapid Type Analysis and related patents used in the Smartlinker product from Object Technology International, an IBM subsidiary.
- TT6. Rapid Type Analysis used to develop the Java Application Extractor (JAX), for IBM alphaWorks.
- TT7. The NEST Network Simulation Testbed used at over 100 universities and corporations, with a commercialized version distributed by Netlabs Inc.

AWARDS

- A1. *Thin Locks: Featherweight Synchronization for Java* selected as one of the 50 most influential papers published in the twenty years of the ACM Conference on Programming Language Design and Implementation, 2003.

- A2. IBM Second Invention Plateau Award, 2002
- A3. IBM Technical Group Award for contributions to the Jikes Research VM, 2002
- A4. IBM First Invention Plateau Award, 1999
- A5. IBM Research Division Award for High Performance Monitors for Java, 1998.
- A6. IBM Resident Study Fellowship, University of California, Berkeley, 1990-1994
- A7. IBM Research Division Award for public domain software, 1988

PROFESSIONAL MEMBERSHIPS

Association for Computing Machinery (SIGPLAN, SIGARCH, SIGOPS)
IEEE and IEEE Computer Society

LANGUAGES

Fluent in German; conversational in Spanish.

PERSONAL

Born 24 February 1963, New Haven, Connecticut.
U.S. Citizen.
