

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, compilers, run-time systems, garbage collection, real-time systems, concurrent systems, instruction set architecture, massively parallel supercomputers, distributed computing, and fault-tolerance. 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, JVM runtime optimization, and compile-time optimization.
- Led group in developing breakthrough technology in real-time garbage collection.
- Developed thin lock technology which sped up Java applications by an average of 25%.
- Successfully transferred runtime and compiler technologies into key IBM products.
- Filed 13 patent applications (3 issued to date).

1994 – 1997 Senior Staff Programmer

- Implemented analysis algorithms from my thesis in Montana, an extensible C++ compiler.
- Filed for patent on the core Ph.D. thesis algorithm, Rapid Type Analysis.

1985 – 1990 *Staff Programmer*

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

IBM Santa Teresa Laboratory

1994 *Visiting Researcher*

- Developed graphical toolkit for evaluation of data from hardware performance monitors in the IBM POWER2 chip.
- Used toolkit to find several performance anomalies, including *cache jamming*, a previously undiscovered phenomenon.

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.
- 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 logging mechanism.

Columbia University Center for Computing Activities

1983-1985 *Manager of Student Consulting*

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

1982-1983 *Student Consultant*

- Staffed public terminal rooms to answer user queries. Created new TOPS-20 help system.

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 for ARPA.

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.

Academic Directions Inc.

1980-1981 *Programmer/Analyst*

- Employee #3 at a startup providing scholarship data to high school students.
- Implemented recursion for FORTRAN, disk caching subsystem, forms processing.

Columbia Research Institute, Inc.

1979-1980 *Programmer*

- Maintained small business systems on various mini-computers.
- Implemented Emacs-style editor in RT-11 TECO.

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 appear.
- 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 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
- C2. **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.
- C3. **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.
- C4. **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.
- C5. **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.
- C6. **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.
- C7. **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.
- C8. **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.

- C9. **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.
- C10. **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.
- C11. **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.
- C12. **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.
- C13. **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.
- C14. **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.
- C15. **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.
- C16. **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.
- C17. **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.

- C18. **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.

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 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.
- P2. **Locking and Unlocking Mechanism for Controlling Concurrent Access to Objects**
David F. Bacon
United States Patent number 6,247,025, issued 12 June 2001.
- P3. **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. **Split-Phase Load Instructions for Avoidance of Memory Latency**
David F. Bacon and Xiaowei Shen
Filed 15 February 2003.
- PP2. **Method for Optimizing Allocation and Access of Multi-level Objects**
Matthew Arnold, David F. Bacon, and Perry Cheng
Filed 1 March 2003.
- PP3. **Optimization of Software Read Barriers**
David F. Bacon and Perry Cheng
Filed 10 February 2003.

- PP4. **Incremental Defragmentation in a Mostly Non-moving Garbage Collector**
David F. Bacon, Perry Cheng, and V.T. Rajan
Filed 10 February 2003.
- PP5. **Space-Efficient Object Models for Object-Oriented Programming Languages**
David F. Bacon, Stephen J. Fink, and David Grove
Filed 10 December 2002.
- PP6. **Method for Detection of Unused Cyclic Data Structures in Linear Time**
David F. Bacon
Filed 11 December 2000.
- PP7. **Method for Concurrent Detection of Unused Cyclic Data Structures**
David F. Bacon and V.T. Rajan
Filed 10 December 2002.
- PP8. **Space-Efficient Object Models for Object-Oriented Programming Languages**
David F. Bacon, Stephen J. Fink, and David Grove
Filed 11 December 2000.
- PP9. **Method to Provide Concurrency Control over Objects Without Atomic Operations on Non-Shared Objects**
David F. Bacon and Stephen J. Fink
Filed 7 August 2000.
- PP10. **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.
- PP11. **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

PROGRAM COMMITTEES

- PC1. Symposium on Principles of Programming Languages (POPL), Venice, Italy, January 2004.
- PC2. International Workshop on Types in Programming Language Design and Implementation (TLDI), New Orleans, Louisiana, January 2003.
- PC3. Java Grande/ISCOPE Conference, Seattle, Washington, October 2002.
- PC4. Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA), Seattle, Washington, October 2002.
- PC5. 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 PANEL

GR1. National Science Foundation Information Technology Research program, May 2002.

TECHNOLOGY TRANSFER

TT1. Thin Locks (also known as “Bacon bits”) implemented in IBM’s product JVM and adopted by product team for all IBM Java products. Also used by Stanford joeq, Excelsior LLC JET, Gnu GCJ 3.1, Mono open source .NET runtime, Intel Open Runtime Platform (ORP), McGill University JVM, and IBM Jikes Research VM.

TT2. Serially Reusable Virtual Machine shipped in the IBM Mainframe version of the Java Virtual Machine.

TT3. Rapid Type Analysis used in the Smartlinker product from Object Technology International, an IBM subsidiary, and the Java Application Extractor (JAX), available from IBM alphaWorks.

TT4. The NEST Network Simulation Testbed is in use at over 100 universities and corporations, with a commercialized version distributed by Netlabs Inc.

AWARDS

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.

IBM Second Invention Plateau Award, 2002

IBM Technical Group Award for contributions to the Jikes Research VM, 2002

IBM First Invention Plateau Award, 1999

IBM Research Division Award for High Performance Monitors for Java, 1998.

IBM Resident Study Fellowship, University of California, Berkeley, 1990-1994

IBM Research Division Award for public domain software procedures, 1988

PROFESSIONAL MEMBERSHIPS

Association for Computing Machinery
IEEE and IEEE Computer Society

LANGUAGES

Fluent in German; conversational in Spanish.

REFERENCES

Furnished upon request.