

School Address:
35 Olden St.
Princeton, NJ 08544
yshuf@cs.princeton.edu

YEFIM SHUF

<http://www.cs.princeton.edu/~yshuf>

Permanent Address:
100-25 Queens Blvd., #4GG
Forest Hills, NY 11375
(917) 476-7632

Objective: A permanent position starting in Fall 2002

Education

- 09/98-present **Ph.D. Candidate, Princeton University, Computer Science** (Degree expected: Fall 2002)
- 09/96-09/98 **M.A., Princeton University, Computer Science, GPA: 4.00**
- 01/93-05/95 **Hunter College, Summa Cum Laude, Honors in Computer Science, Major GPA: 4.00, Overall: 3.95**
- 09/88-07/92 Tashkent State Technical University, Major: **Computer Systems and Networks/Electrical Engineering**
- 09/86-06/88 High School #178, **Graduation with Honors and a Silver Medal**, Tashkent, Uzbekistan (former USSR)

Employment History

- 07/99-present *Visiting Researcher*, **IBM T. J. Watson Research Center, Exploratory Server Systems**
Designed, implemented, and evaluated techniques for improving the data locality and performance of Java workloads
- 03/96-09/96 *Lead Systems Engineer/New Product Development*, **Giles Scientific Inc., New York, NY**
Developed systems and application software for medical laboratory information systems; implemented confidential data transfer; improved the quality of existing software
- 12/94-03/95 *UNIX System Administrator's Assistant*, **SPARS/UNIX Laboratory, Hunter College, CUNY**
Installed and maintained a cluster of SUN workstations, SUNPics printers, network hubs, transceivers; modified shell scripts; performed backup on file servers; solved various technical problems
- 06/94-08/94 *Database Developer*, **Office of the Mayor, Midtown Enforcement, New York, NY**
Designed and implemented new features lacking in the existing IS using RDBMS; improved reports; optimized queries; integrated applications into the system; maintained consistency and integrity of DB

Research Positions

- 07/99-present *Visiting Researcher*, **IBM T. J. Watson Research Center, Exploratory Server Systems**
Project title: "Characterizing the data locality and improving the memory performance of Java applications"
- 02/99-07/99 *Research Assistant*, **Princeton University, Computer Science Department**
Project title: "Exploiting inherent spatial locality in applications"
- 06/98-09/98 *Research Intern*, **IBM T. J. Watson Research Center, Scalable Parallel Systems**
Project title: "Predicting Behavior of Parallel Applications in Large-Scale Shared-Memory Systems"
- 06/97-05/98 *Research Assistant*, **Princeton University, Computer Science Department**
Project title: "Fast Software Encryption"
- 06/96-10/96 *Research Fellow, Ronald E. McNair Research Program*, **Center for Advanced Study, CUNY**
Project title: "A Cyclic Garbage Collection Algorithm for a Distributed Object-Oriented System"
- 06/95-10/95 *Research Fellow, Ronald E. McNair Research Program*, **Center for Advanced Study, CUNY**
Project title: "Garbage Collection for a Shared Memory Multiprocessor"
- 06/91-12/91 *Senior Laboratory Assistant*, **Research Institute of Energetics, The Academy of Science, Uzbekistan**
Designed software for classified scientific projects related to laser systems; optimized and ported various libraries

Teaching Positions

- 09/96-05/98 *Assistant in Instruction*, **Princeton University, Computer Science Department**
- 02/96-05/96 *Master Tutor in Computer Science*, **Hunter College/Special Services, CUNY Research Foundation**

Awards and Honors

- *IBM T.J. Watson Research Center, Yorktown Heights, NY*
 - **First IBM Invention Plateau, 2002**
 - **Four IBM Invention Achievement Awards for four patent applications, 2000-2002**
- *Princeton University, Princeton, NJ*
 - **ACM SIGPLAN 2002 Scholarship**
 - **ACM SIGMETRICS 2001 Scholarship**
 - **2000 and 2001 HIAS Scholarship Awards**
 - **2001 Association of Princeton Graduate Alumni Scholarship**

- *Post-Graduate, Center for Advanced Study, CUNY, New York, NY*
 - **The Ronald E. McNair Post-Baccalaureate Achievement Program, '95 and '96**
 - **The Grove Foundation Scholarship for Graduate Study, '96**
- *Hunter College, New York, NY*
 - **Summa Cum Laude, '95**
 - **Graduation with Computer Science and General Honors, '95**
 - National Collegiate Computer Science Award, '95
 - US Achievement Academy, '94
 - All-American Scholar, '94
 - Golden Key National Honor Society, '93
 - Rose Biller Scholarship, 95
 - Class 1933 Scholarship, '94
 - Lillian Fisher Alpert Scholarship, '94
 - Grace and Marie Schut Scholarship, '93
 - Dean's List, '93-'95
- *Tashkent State Technical University, Tashkent, Uzbekistan (former USSR)*
 - A stipend for excellence in undergraduate studies, '88-'92
 - First place in mathematics and physics contests, '88
- *High School #178, Tashkent, Uzbekistan (former USSR)*
 - Graduation with Honors and a Silver Medal, '88

Publications

- **Characterizing the Memory Behavior of Java Workloads: A Structured View and Opportunities for Optimizations**
Yefim Shuf, Mauricio J. Serrano, Manish Gupta, and Jaswinder Pal Singh (Princeton U., Intel, and IBM)
Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2001 / Performance 2001), Cambridge, MA, pages 194 - 205, June 2001
- **Exploiting Prolific Types for Memory Management and Optimizations**
Yefim Shuf, Manish Gupta, Rajesh Bordawekar, and Jaswinder Pal Singh (Princeton U. and IBM)
The 29th Annual ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages (POPL 2002), Portland, OR, pages 295 - 306, January 2002
- **A Study of Memory Behavior of Java Workloads**
Yefim Shuf, Mauricio J. Serrano, Manish Gupta, and Jaswinder Pal Singh (Princeton U., Intel, and IBM)
Java Microarchitectures (Book) / Editors: Vijaykrishnan Narayanan and Mario Wolczko
The Kluwer International Series in Engineering and Computer Science, Volume 679, Chapter 2
Kluwer Academic Publishers, ISBN 1-4020-7034-9, Boston, MA, April 2002
- **Multiple Page Size Support in the Linux Kernel**
Simon J. Winwood, Yefim Shuf, Hubertus Franke (U. of New South Wales, Princeton U., and IBM)
The 4th Annual Ottawa Linux Symposium (OLS 2002), Ottawa, Canada, June 2002 (to appear)
- **Creating and Preserving Locality of Java Applications at Allocation and Garbage Collection Times**
Yefim Shuf, Manish Gupta, Hubertus Franke, Andrew Appel, and Jaswinder Pal Singh (Princeton U. and IBM)
The 17th Annual ACM Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA 2002), Seattle, WA, November 2002 (to appear)

Patent activity (work done at IBM T.J. Watson Research Center)

- **Method and apparatus for efficient memory management (Filed)**
Yefim Shuf, Hubertus Franke, Manish Gupta, Marc Snir
- **Method and apparatus for efficient cache management (Filed)**
Yefim Shuf, Hubertus Franke, Manish Gupta
- **A method for efficient garbage collection based on object types (Filed)**
Yefim Shuf, Manish Gupta, Rajesh Bordawekar
- **A method for reducing write barrier overhead (Filed)**
Yefim Shuf, Manish Gupta, Rajesh Bordawekar
- **A method for reducing the memory requirements in object-oriented programs (in preparation)**
Yefim Shuf, Manish Gupta, Rajesh Bordawekar

- **A method for reducing the memory requirements and improving data locality in programs (in preparation)**
Yefim Shuf, Manish Gupta, Rajesh Bordawekar
- **Method and apparatus for efficient spilling and refilling the contents of registers (in preparation)**
Yefim Shuf, Hubertus Franke, Manish Gupta

Public Presentations

- *Conference Presentations*
 - **Creating and Preserving Locality of Java Applications at Allocation and Garbage Collection Times**
OOPSLA 2002, Seattle, WA (11/2002, invited)
 - **Exploiting Prolific Types for Memory Management and Optimizations**
POPL 2002, Portland, OR (01/2002)
 - **Characterizing the Memory Behavior of Java Workloads: A Structured View and Opportunities for Optimizations**
SIGMETRICS 2001, Cambridge, MA (06/2001)
- *Presentations to Industry Leaders*
 - **Distinguishing Between Prolific and Non-Prolific Types for Efficient Memory Management**
2001 Computer Science Department Affiliates Day, Princeton U., Princeton, NJ (05/2001)
 - **The Memory Behavior of Java Workloads**
2000 Computer Science Department Affiliates Day, Princeton U., Princeton, NJ (11/2000)
- *Poster Presentations*
 - **Using Prolific Types for Memory Management and Optimizations**
2001 IBM Student Research Day, Yorktown Heights, NY (07/2001)
 - **The Memory System Behavior of Java Programs**
2000 AT&T Labs-Research Student Research Symposium, Florham Park, NJ (10/2000)
 - **Characterizing the Memory Behavior of Java Workloads**
2000 IBM Student Research Day, Yorktown Heights, NY (08/2000)
- *Research Seminars*
 - **Improving the Memory Performance of Java Applications**
Computer Science Department, Princeton U., Princeton, NJ (04/2002)
 - **Exploiting Prolific Types for Memory Management and Optimizations**
Computer Science Department, Princeton U., Princeton, NJ (01/2002)
 - **Characterizing the Memory Behavior of Java Workloads**
Computer Science Department, Princeton U., Princeton, NJ (05/2001)
 - **A Case for Memory Management Based on Object Types**
IBM T.J. Watson Research Center, Yorktown Heights, NY (04/2001)
 - **Reducing the Memory Requirements of Java Applications**
IBM T.J. Watson Research Center, Yorktown Heights, NY (01/2001)
 - **A Study of the Memory Behavior of Java Workloads**
IBM T.J. Watson Research Center, Yorktown Heights, NY (09/2000)

Ph. D. Thesis

- **Thesis Title: Improving the Memory Performance of Java Applications**
- *Thesis Committee*
 - **Professor Jaswinder Pal Singh**, Princeton U.
 - **Dr. Manish Gupta**, IBM T.J. Watson Research Center
 - **Professor Andrew W. Appel**, Princeton U.
 - **Professor Brian Kernighan**, Princeton U.
 - **Professor David Walker**, Princeton U.
- *Preliminary Thesis Defense*: 04/15/2002
- *Final Thesis Defense*: Fall 2002 (expected)

Main Areas of Interests

- computer architecture, operating systems, compilers, run-time systems, Java
- memory management, garbage collection, techniques for improving the locality of code and data
- workload characterization and analysis, performance measurements and optimizations

- understanding interactions between applications and various system components

Other Areas of Interests

- databases, data analysis and visualization, pattern discovery, data mining, knowledge management and discovery
- parallel systems and applications, scalability issues in large systems
- computer security, viruses and anti-viruses, intrusion detection, reliability and stability issues
- embedded and low-power systems, design and performance issues in resource-constrained environments

Selected Projects

- Demonstrated the benefit of multiple page size support for modern C/C++ workloads and Java workloads in particular
- Designed and implemented a new type-based memory management and garbage collection scheme
- Designed and implemented a new locality enhancing object co-allocation scheme based on object types
- Designed and implemented a new locality-friendly traversal algorithm for garbage collection; the algorithm has good scalability properties for multi-processors (e.g., low synchronization and false sharing overheads)
- Designed and implemented a hybrid type-based memory allocation technique for embedded environments; the technique reduces the memory requirements of Java applications
- Implemented a tool for tracing the memory accesses of Java programs; implemented software for analyzing collected traces and simulating the memory system behavior; created data visualization software
- Implemented a parallel MPEG-2 encoder for shared memory multi-processors; demonstrated good scalability of that application
- Demonstrated that on fast superscalar processors memory copy with encryption can be virtually as fast as memory copy by itself
- Reverse engineered a single-threaded version of a device driver for a video/image capture card; created a multi-threaded version of that driver
- Optimized and improved image analysis and expert system software
- Restructured several database applications; optimized database queries; demonstrated significant performance improvements
- Reverse engineered a LAN communication software package; replaced its authentication protocol with a new one which prevented users from running pirated versions of that software; created a new installation utility
- Reverse engineered numerous (boot sector, COM, EXE, and other) viruses; created anti-virus and inoculation software
- Reverse engineered various system utilities; documented their usage of system services
- Reverse engineered several copy-protection schemes; created improved copy-protection utilities

Languages: C/C++, Java, Assembly x86/Z80/MC6502/VAX-11/IBM-370/MIPS, Ada, Pascal, Fortran, PL/M, SQL, ML

Special Skills: Excellent software testing skills, **reverse engineering**

Professional Characteristics and Qualities

- Proven team player
- Excellent presentation, communication, and interpersonal skills
- Creative, energetic, self-motivated
- Good understanding of complex systems and application environments
- Receptive to new ideas and new areas

Public Service

- Reviewer for PLDI 2002, SC 2001, SIGMETRICS 2000, IPDPS 2000, ICPADS 2000, LCR 2000, PACT 1999
- Princeton U. School of Engineering and Applied Science graduate student committee member
- Princeton U. Computer Science Department graduate student committee member
- Princeton U. Hillel graduate student representative

Society Memberships: Association for Computing Machinery (ACM), Golden Key National Honor Society

Hobbies: Table Tennis / Ping-Pong

Language Skills: English (fluent), Russian (fluent)

Security Clearance: The Academy of Science, Uzbekistan (former USSR), 1991

Citizenship: US

References

- **Professor Jaswinder Pal Singh**
Princeton University, Department of Computer Science, 35 Olden Street, Princeton, NJ 08544
jps@cs.princeton.edu, (609) 258-5329
- **Dr. Manish Gupta, Senior Manager, Emerging System Software**
IBM T.J. Watson Research Center, Route 134, P.O. Box 218, Yorktown Heights, NY 10598
mgupta@us.ibm.com, (914) 945-2494
- **Professor Andrew W. Appel**
Princeton University, Department of Computer Science, 35 Olden Street, Princeton, NJ 08544
appel@cs.princeton.edu, (609) 258-4627
- **Professor Brian Kernighan**
Princeton University, Department of Computer Science, 35 Olden Street, Princeton, NJ 08544
bwk@cs.princeton.edu, (609) 258-2089
- **Dr. Hubertus Franke, Manager, Enterprise Linux Group**
IBM T.J. Watson Research Center, Route 134, P.O. Box 218, Yorktown Heights, NY 10598
frankeh@us.ibm.com, (914) 945-2003

Last updated on: 05/19/2002 20:55