

---

## **Binny Gill**

Address: 5-1 Tuck Farm Road, Auburn, MA 01501.

Email: binnyg at us.ibm.com

Phone: (508) 936 6892

---

## **Research Interests and Current Role at IBM**

My research interests are in the field of adaptive cache management and in solving technical challenges and providing thought leadership in the broader field of distributed storage systems. In my current role at IBM Almaden Research Center I am entrusted with the responsibility to innovate and create intellectual property for IBM via patents and widely cited publications in esteemed conferences. My focus is on inventing and developing new autonomic algorithms for performance improvement of distributed storage systems.

---

## **Awards and Honors**

### **IBM Awards:**

#### **1. Outstanding Technical Achievement Award, Dec 2004**

The Outstanding Technical Achievement Award (OTAA) is the highest technical achievement award granted by IBM Research to an employee. It is granted by IBM to recognize outstanding achievements involving exceptional technical skill and insight. These accomplishments typically include a pioneer application of one or more principles or a major contribution to a technology. I received this award for my contribution to creating the Adaptive Replacement Cache algorithm. Below is the award citation:

"...for their contributions to creating Adaptive Replacement Cache and incorporating it in IBM's Enterprise Storage Server products.

Caching is used widely in storage systems, databases, web servers, middleware, processors, file systems, disk drives, RAID controllers, operating systems and in numerous other applications. For nearly four decades, the Least Recently Used (LRU) algorithm and its variants have remained the popular class of cache replacement policies. A long-standing question in cache management has been: Is it possible to improve on LRU across a wide range of workloads and cache sizes without incurring excess overhead or requiring workload-specific pre-tuning? Hundreds of attempts have been made, most significantly, FBR, LRU-2, 2Q, LRFU and MQ. However, until now, none has been universally successful. Adaptive Replacement Cache (ARC) dynamically adapts between recency (LRU) and frequency (Least Frequently Used (LFU)) to achieve higher cache hit rates, which imply better performance for

a server or application. ARC was successfully transferred to IBM's Enterprise Storage Server (ESS) products (DS 6000 and DS 8000) and was included in the product announcements as a significant innovation. On mixed random and sequential workloads, ARC was found to notably increase the cache hit rate of ESS on random workload (almost 2x in some cases) without impairing the hit rate on sequential workloads."

## **2. Invention Achievement Award**

**First Plateau : Oct 2004**

**Second Plateau: Apr 2006**

**Third Plateau: Feb 2007**

The Plateau award is awarded by IBM to recognize employees for being prolific submitters of patents within IBM and thus contributing significantly to the growth of intellectual property of IBM.

## **3. First Patent Application Invention Achievement Award, May 2004**

This is awarded when an employee's creative/innovative ideas result in a first patent of the employee for IBM. I received this award in May 2004.

## **Indian Institute of Technology, Kanpur, India:**

### **1. Award for the Best Bachelor of Technology Project (1999).**

This award is granted to the best undergraduate research project of the year by the Dept. of Computer Science and Engineering at the Indian Institute of Technology, Kanpur. I was the recipient of this award for the year 1999.

### **2. The Director's Award for Academic Excellence (1995-1996 & 1996-1997)**

This award is given to the top performing students in the Indian Institute of Technology, Kanpur, for their academic excellence.

## **Government of India:**

1. Awarded 9th position out of 150,000 nationwide (1995) in the Joint Entrance Examination for the Indian Institutes of Technology.
  2. The National Talent Search Scholarship (1993-1998) and Certificate of merit.
  3. All India Talent Search Award and scholarship (1992-1993).
  4. Certificate of Merit in Physics (1993)
-

---

## Education

August 1999 - January 2001

M.S. Computer Science (GPA 3.94/4.00)

University of Illinois at Urbana-Champaign, Urbana, IL

Advisor: Roy Campbell

August 1995 - May 1999

B.Tech. Computer Science and Engineering (GPA 9.7/10.0)

Indian Institute of Technology, Kanpur, India

---

## Experience

### January 2003 - present

Advisory Software Engineer (Research) at Almaden Research Center (IBM), San Jose, CA.

My work at IBM Research has the following dimensions:

1. Innovate and create intellectual property for IBM via patents and widely cited publications in esteemed conferences.

In the last 4 years I have filed 13 patents and 5 more are being filed. I have submitted three papers to USENIX Annual Technical Conference and USENIX File and Storage Technologies Conference.

2. Invent and develop new algorithms for performance improvement of distributed storage systems.

I have invented, developed and delivered in product, the Sequential Prefetching in Adaptive Replacement Cache algorithm. This algorithm is seen as a key differentiator for IBM's flagship storage products. It elegantly solves a long-standing question in cache management. It improves the cache hit rate and hence the performance of the system significantly.

I have also invented the Wise Ordering for Writes (WOW) algorithm. This algorithm solves the age old problem of combining temporal and spatial locality via a remarkably simple solution. The preliminary implementations have shown unprecedented performance improvements for the SPC-1 benchmark. I have implemented WOW for IBM's DS8300 storage servers.

Sequential prefetching is an important aspect of computer systems. I introduced a useful classification of sequential prefetching algorithms and developed a new algorithm called

Adaptive Multi-stream Prefetching (AMP) that shares a cache between multiple sequential streams in a provably optimal way. I have implemented a prototype for the DS8300 storage server.

3. Create new autonomic algorithms for distributed systems and solve difficult technical problems in storage systems.

I have filed multiple patents and publications covering autonomic algorithms for distributed storage systems.

4. Study, design, enhance and develop distributed storage systems. Add new functionality (autonomic, performance, security, networking) to IBM storage systems.

I have prototyped dozens of new caching algorithms to verify their feasibility during my research at IBM. I have also worked on the design of the Distributed Storage Tank product which extends the Storage Tank system's application for heterogeneous and distributed file sharing to include file sharing in wide-area networks. I also designed a security infrastructure for Storage Tank that would be scalable, applicable and powerful enough to solve the security needs of heterogeneous data grids and other distributed data-sharing systems.

5. I am also an elected early career member of the Storage Technology Council (STC), an IBM Academy of Technology Affiliate. The STC acts as a technical advisory resource for the IBM business line and IBM senior executives. It influences product direction and performs problem resolution. At any point of time IBM has only 30 early career members in this council.

### **January 2001 - January 2003**

Software Engineer and Project Lead at Zambeel Inc, Fremont, CA.

Worked on various aspects of design and innovation in highly scalable, manageable and fault-tolerant, cluster-based storage systems.

1. Unified Allocation

Lead a 5 member team to provide optimal resource allocation in the distributed storage cluster and reduce cost of unit storage, while maintaining reliability, availability and scalability requirements, by improving resource utilization and allowing multiple file system services on a single device.

2. Filesystem Snapshots

Designed and developed the management software to coordinate in parallel, the creation and deletion of point-in-time instantaneous snapshots in Zambeel's fluid file system.

3. Cluster Management

Designed, developed and maintained the Zambeel cluster management software. This involves dynamic resource discovery and management, SLA translation and implementation, failure recovery and health monitoring within the system.

4. NFS security

Developed authorization service for NFS v3 in the Aztera NAS system. Also improved performance of NFS v2 authorization by 35%.

5. NFS exports

Developed mechanism to centrally manage NFS exports by broadcasting configuration changes via RPC to participating file system gateways within the distributed system.

**August 1999 - January 2001**

Research Assistant, Systems Software Research, Univ. of Illinois at Urbana Champaign, IL

My research in the “GAIA active spaces” and the “2k distributed OS”:

1. Extending the reach of traditional computing systems to encompass the devices and the physical space surrounding them.
2. Architecting security infrastructure in GAIA, a middleware operating system for ubiquitous computing. The design involved policy specification and enforcement mechanisms using GSS API and ORB security interceptors.
3. Researching and inventing policy-driven Role-Based Access Control (RBAC) for the GAIA system. This involved securing the distributed services like the name service and the distributed file system using secure interceptors.
4. Designing security interceptors for TAO (The Ace ORB). Devising a modified ARBAC representation using the secure file system as a policy repository.

**May 2000 - August 2000**

Intern, Sun Microsystems, Palo Alto, CA

Worked on integrating iPlanet eCommerce solutions for the configuration engine on SunStore.

**May 1998 - July 1998**

Intern, Sun Microsystems, Palo Alto, CA

Worked on benchmarking automation of JaWS (the Java Web Server).

---

**Selected Product Impact**

IBM TotalStorage Servers DS8000 and DS6000 (aka Shark) are IBM's flagship storage controllers with a revenue of a few billion dollars a year. An essential element of these

products is the cache. My research in the field of caching algorithms led to a novel, self-tuning, low overhead, autonomic algorithm that continuously adapts to evolving workloads improving the performance of these systems dramatically. The algorithm, Sequential Prefetching in Adaptive Replacement Cache, is widely available since Oct 2004 as an important product differentiator in these IBM storage servers.

---

---

### **Press mentions of my work in caching**

IBM's announcement of two new products DS6000/DS8000,  
BusinessWire, October 12, 2004.

[Picked up by many major media outlets.]

“New caching technology from IBM Research called Adaptive Replacement Cache (ARC) is designed to help clients achieve dramatically greater throughput and faster response times than previous IBM TotalStorage Enterprise Storage Server 800 systems. ARC incorporates autonomic, self-optimizing technology and a more efficient and effective method for the widely used process of replacing data pages in computer cache memories. The breakthrough technology, available in both the DS6000 and DS8000 series, dynamically optimizes the storage system's performance for both sequential and randomly accessed workloads.”

EE Times,  
Nicolas Mokhoff, October 12, 2004.

“The new caching technology, adaptive replacement cache (ARC), was developed by IBM's Almaden Labs (San Jose, Calif.) and is available in both DS6000 and DS8000 series. ARC incorporates autonomic, self-optimizing algorithms and a more efficient method for the widely-used process of replacing data pages in cache memories.”

Information Week,  
Martin J. Garvey, October 12, 2004.

“For both the DS6000 and the DS8000, IBM Labs is offering the Adaptive Replacement Cache, which is designed to help customers increase system performance for different forms of data access while taking up less capacity with cache.”

EnterpriseITPlanet,  
Pedro Hernandez, October 12, 2004.

“The system also debuts with Adaptive Replacement Cache, a technology from IBM Research that considerably boosts the hardware's data retrieval rates over that of the ESS800.”

CXOtoday.com,

October 13, 2004.

“The systems also feature a new caching technology called Adaptive Replacement Cache (ARC), which delivers a faster response time than previous TotalStorage Enterprise Storage Server 800 systems.”

InternetNews.com Storage News,  
Clint Boulton, October 12, 2004.

“...the DS8000 utilizes new caching technology from the company's research lab that helps users realize greater throughput and response times than the ESS800. The Adaptive Replacement Cache (ARC) uses autonomic technology to provide a more intuitive caching experience.”

---

## Publications

"AMP: Adaptive Multi-stream Prefetching in a Shared Cache"

Binny Gill and Luis Angel D. Bathen,  
USENIX Conference on File and Storage Technologies (FAST 07), San Jose, CA,  
February 13-16, 2007.

“WOW: Wise Ordering for Writes—Combining Spatial and Temporal Locality in Non-Volatile Caches”

Binny Gill and Dharmendra S. Modha  
Proceedings of the 2005 USENIX Conference on File and Storage Technologies,  
San Francisco, CA, December 14-16, 2005.

“SARC: Sequential Prefetching in Adaptive Replacement Cache.”

Binny Gill and Dharmendra S. Modha  
Proceedings of the 2005 USENIX Annual Technical Conference,  
Anaheim, CA, April 10–15, 2005.

“Secure Dynamic Reconfiguration of Scalable CORBA Systems with Mobile Agents.”

Fabio Kon, Binny Gill, Manish Anand, Roy Campbell and Dennis Mickunas  
Proceedings of the IEEE Joint Symposium on Agent Systems and Applications /  
Mobile Agents (ASA/MA'2000), Zurich, September, 2000.

“Looking beyond the mouse: Gesture Based Internet Games.”

Amitabha Mukerjee, Debabrata Dash, Amit, Binny Gill and Mukesh P. Singh  
Proceedings of the South-East Asian Regional Computer Conference SEARCC-97,  
New Delhi, India, December 3-5, 1997.

“Security Architecture in GAIA.”

Prashant Viswanathan, Binny Gill and Roy Campbell.

Technical Report UIUCDCS-R-2001-2215 UILU-ENG-2001-1720,  
University of Illinois at Urbana-Champaign, 2001.

Master's Thesis

"Dynamic Policy-Driven Role-Based Access Control for Active Spaces."  
University of Illinois at Urbana-Champaign, IL, 2001

Bachelor's Thesis

"Design and Implementation of a Push Protocol for the Internet",  
Indian Institute of Technology, Kanpur, India, 1998

---

---

**Technical Competencies**

OS: Linux, Unix (Solaris, AIX), Windows 2000  
Programming Languages: C, C++, Java, Perl, shell, HTML, Expect, Tcl/TK.  
Protocols and standards: NFS, TCP/IP, RPC, CORBA, HTTP, XML  
PKI, Kerberos, Cryptographic methods.

---

**Professional Affiliations**

1. IEEE member
  2. USENIX member
  3. ACM member
-