

**Dr. Ismail Ari**  
Assistant Professor at School of Engineering  
Ozyegin University



<http://www.soe.ucsc.edu/~ari>  
[http://ee.ozyegin.edu.tr/?page\\_id=332](http://ee.ozyegin.edu.tr/?page_id=332)

Work: +90-216-5592331  
[ismail.ari@ozyegin.edu.tr](mailto:ismail.ari@ozyegin.edu.tr)

Address: Kusbakisi Cad. No 2  
Altunizade, Uskudar, 34662  
Istanbul, Turkey

## Research Interests and Objective

I am a faculty member in Electrical and Electronics Engineering department at [Ozyegin University](http://www.ozyegin.edu.tr). Before joining Ozyegin, I worked as a senior research scientist at [Hewlett-Packard Laboratories](http://www.hp.com) in Silicon Valley, USA from 2004 to 2009. I have theoretical and practical experience in the design, implementation and performance characterization of file and storage systems, databases, and computer networks. My research focus areas can be listed as follows:

- Data stream and complex event processing
- Data mining models and model metadata management
- Business process management and business intelligence
- Cloud computing: Providing infrastructure, platform and software as services
- Web services, service-oriented architectures, session management
- System automation, integration and performance tuning
- Design of large-scale, distributed file and storage systems
- Developing novel caching and pre-fetching techniques
- Modeling, analysis and simulation of computer and telecomm systems
- Computer networks; Design of efficient, reliable data transfer protocols
- Providing Quality of Service (QoS) for networks

***My objective is to establish a world-class computer systems research center at Ozyegin University and share my experiences with students and other researchers world-wide.***

## Education

- 2001-2004**      **Ph.D. in Computer Science**, *Focus:* File and Storage Systems  
[University of California Santa Cruz](http://www.ucsc.edu) , *Advisor:* Prof. Ethan L. Miller
- 1998-2000**      **M.S. in Computer Science**, *Focus:* Computer Networks  
[University of Maryland Baltimore County](http://www.umaryland.edu)
- 1994-1998**      **B.S. in Electrical-Electronics Eng.**, *Focus:* Control Systems & Electronics  
[Bogazici University](http://www.bogazici.edu.tr), Istanbul, Turkey

## Work Experience

**2005-2009 Systems and Software Expert, Hewlett-Packard Laboratories**  
***Service-Oriented Retail Platform, Intelligent Information Management Lab***

I worked with the HP Labs team that developed a retail kiosk/web [platform](#), which was displayed in trade shows and demoed to numerous top retailers (e.g. [Safeway](#) and [Kroger](#)). This project took 2 years and had 6 team members. I contributed to the development of the 3 major components:

- [Session Management As a Service](#): This novel design solved multi-channel (web, store, mobile) integration problems for retail and banking applications. Local and remote session managers handled session timeout, state serialization, and reliable remote storage issues. I was the lead designer/developer. I used Visual Studio IDE (ASP.NET, C#) and the IIS web server.

- [Data Mining and Business Intelligence](#): For real-time personalized marketing in our retail platform we developed customer data mining models. As a team of 3, we built the system that integrated model scoring with the “coupon offer” process. I focused on the model quality management and other system integration issues. The code included C#, IIS, Biztalk, SQL Server components. We filed a patent application.
- [Data Stream and Complex Event Processing](#): We added an “event sequence detection” operator to a data stream processing engine called CAPE and integrated this with our retail platform. Complex events such as shoplifting or prospect customer arrival can be detected and managers can be alerted. I was the lead designer and I mentored an intern. Components included Java code, CORBA, Biztalk, ASP.NET, and Ajax.

**2004-2005 Systems and Software Specialist, Hewlett-Packard Laboratories**  
***Automated SAN-Level Caching, Emerging Technologies Lab***

HP Labs is a leading world-wide research center developing the emerging non-volatile memory technologies including [MRAM](#), [MemorySpot](#), and recently [Memristor](#). As systems researchers, our goal is to find the best ways to utilize these emerging technologies in existing computer and data center architectures.

- As a team of 5, we investigated the use of MRAM in Storage Area Networks (SAN) and other file & storage systems. I was the lead developer in the [SANCache](#) project. We automated the process of hot data migration from disk-based storage onto fast NVRAM and designed the caching algorithms. The benefits included on-demand storage performance improvement, selective workload isolation and differentiation. We filed two [patent](#) applications and published 2 related [papers](#). System included Windows storage driver development, storage virtualization, Fibre Channel (FC) configurations, multi-threaded C++ code handling 10K IO/sec, and distributed locks.
- In 2005, I mentored an intern to work on the decentralized SANCache implementation, which used a Distributed Hash Table (DHT) to provide fast consistency among nodes with a light-weight (network-efficient) protocol.

**2001-2004: Research Assistant, University of California Santa Cruz**  
***Storage Systems Research Center (SSRC)***

SSRC focuses on large-scale data and information management issues and is sponsored by biggest the IT companies as well as the US Government. Large-scale systems need to adapt to varying load conditions.

- I was the lead researcher for the [ACME](#) adaptive caching project. I applied machine learning techniques for selection and adaptation of the best cache replacement algorithm called “[expert](#)” used by the large-scale, distributed systems. I proposed adoption of ACME [via data snooping routers](#) for the next-generation Internet infrastructure as a remedy to ineffective multi-level caching and [flash-crowd problems](#) as a part of my [PhD thesis](#). ACME provides coordination among distributed network caches without explicit communication.
- During my PhD, I’ve gathered several real web-proxy, file and storage workload traces, analyzed them with Perl scripts, modeled LAN to WAN network topologies, implemented and used multi-cache [simulator](#) in C++, replayed traced traffic over simulated topologies and flash crowd scenarios. I demonstrated adverse effects of workload mixing and the effectiveness of adaptive caching solution to multiple Internet-scale traffic problems. I also worked with the machine learning team to further study [pre-fetching](#) techniques to support caching. My PhD-related work received more than 100 [citations](#) over the last 7 years.
- During my PhD course studies, I prepared a software engineering project [report](#) on analysis of two open source software projects by parsing through their email correspondence using Perl scripts. This work took 4 months and was later cited in a related [book](#). For the computer security-cryptography class, we prepared designed a method for fast and fully-encrypted file storage in Unix mixing HMAC using SHA-1 codes to assure authenticity. It took 3 of us 3 months and we used C to implement the system. I also had studies related to computational [complexity](#) and [statistics](#).

**2002-2003 Summer: Research Intern, HP Labs**  
***SANCache setup and configuration***

In 2002, I worked on the initial setup and configuration of the HP disk arrays, HP CASA storage virtualization platform, selection and installation of the FC-NVRAM device for the SANCache project. In 2003, I did the preliminary implementation of the caching module that automated the hot data migration process and enabled distributed operation via distributed locks and messaging. Most of the code was written in C++.

## 2000-2001 Summer: Research Intern, HP Labs

### *Internet SCSI (ISCSI) Driver Implementation and WAN Performance Studies*

ISCSI was an emerging storage inter-connect protocol in 2000 that provided remote access to block-storage devices over cheap TCP/IP/Ethernet networks. I was among the first to implement an ISCSI driver (initiator and target) for Unix FreeBSD operating system. This work required kernel-level development in C language. In 2001, I focused on solving ISCSI performance [problems](#) for Wide-Area Networks including head-of-queue blocking and bandwidth shrink issues. I combined the famous simulators [NS2](#) (written in C++/TCL) and [DiskSim](#) (written in C) to show the effects of the mentioned problems and suggested remedies on performance. I've written my Masters thesis on these problems and the modeling/implementation details for the simulation environment.

## Publications

### Book Chapters

- 1- Kemal Akkaya, Ismail. Ari, *In-network Data Aggregation in Wireless Sensor Networks*, The Handbook of Computer Networks, Vol.2, John Wiley & Sons.

### Refereed Journal Publications

- 1- Ismail Ari, Ahmed Amer, Robert Gramacy, Ethan L. Miller, Scott Brandt and Darrell D. E. Long, *Adaptive Caching using Multiple Experts*, Carleton Scientific, Proceedings in Informatics, volume 14, pages 143-158. 2002 (59 [Citations](#) in Google Scholar)

### Refereed Conference Publications

- 1- Mo Liu, Elke A. Rundensteiner, Kara B. Greenfield, Chetan Gupta, Song Wang, Ismail Ari, Abhay Mehta, E-Cube: Multi-Dimensional Event Sequence Analysis Using Hierarchical Pattern Query Sharing, **(Submitted to)** 26<sup>th</sup> IEEE International Conference on Data Engineering (ICDE) 2010.
- 2- Chetan Gupta, Song Wang, Ismail Ari, Ming Hao, Umeshwar Dayal, Abhay Mehta, Manish Marwah, Ratnesh Sharma, CHAOS: A Data Stream Analysis Architecture for Enterprise Applications, 11th IEEE Conference on Commerce and Enterprise Computing (CEC.09)
- 3- Jun Li, Ismail Ari, Jhilmil Jain, Alan H. Karp, Mohamed Dekhil, Mobile In-Store Personalized Services, In 7th IEEE International Conference on Web Services (ICWS) July 2009, Los Angeles
- 4- Jhilmil Jain, Ismail Ari, Jun Li, Understanding the Challenges Faced During the Management of Data Mining Models, In 2nd ACM Symposium on Computer Human Interfaces for Management of IT (CHIMIT) 2008, November 2008.
- 5- Jhilmil Jain, Ismail Ari, Jun Li, Using Dashboards for Managing Model Lifecycle, In 2nd ACM Symposium on Computer Human Interfaces for Management of IT (CHIMIT) 2008, November 2008.
- 6- Ismail Ari, Jun Li, Alex Kozlov, Mohamed Dekhil, *Data Mining Model Management to Support Real-Time Business Intelligence in Service-Oriented Architectures*, HP Software University Association Workshop, Morocco, June 2008.
- 7- Ismail Ari, Jun Li, Riddhiman Ghosh, Mohamed Dekhil, *Providing Session Management As a Core Business Service*, In Proceedings of WWW 2007 Conference.
- 8- Ismail Ari, Melanie Gottwals, Dick Henze, *Performance boosting and workload isolation in Storage Area Networks using SANCache*, In Proceedings of the 23rd IEEE / 14th NASA Goddard Conference on Mass Storage Systems and Technologies 2006.
- 9- Ismail Ari, Ethan L. Miller, *Caching support for push-pull data dissemination using data snooping routers*, In Proceedings of the 10<sup>th</sup> IEEE International Conference on Parallel and Distributed Systems (ICPADS'04).
- 10- Ismail Ari, Melanie Gottwals, Dick Henze, *SANBoost: Automated SAN-Level Caching in Storage Area Networks*, In Proceedings of the First IEEE International Conference on Autonomic Computing (ICAC'04) (7 [citations](#))
- 11- Ismail Ari, Bo Hong, Ethan L. Miller, Scott Brandt and Darrell D. E. Long, *Managing Flash Crowds on the Internet*, In Proceedings of the 11<sup>th</sup> IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'03) pp. 246-249. (31 [citations](#))
- 12- Robert Gramacy, *Manfred Warmuth*, Scott Brandt and Ismail Ari, *Adaptive Caching By Refetching*, In Proceedings of the 2002 Neural Information Processing Systems (NIPS'02) pp. 1465-1472. (25 [citations](#)).
- 13- Ismail Ari, Ahmed Amer, Robert Gramacy, Ethan L. Miller, Scott Brandt and Darrell D. E. Long, *Adaptive Caching using Multiple Experts*, In Proceedings of the 2002 Workshop on Distributed Data and Structures (WDAS'02) pp. 143-158.

## Patents Granted and Published

- 1- Ismail Ari, *System and method to access storage inside an inactive computing device*, [U.S. Patent 7,418,538](#), Granted on August, 26, 2008.
- 2- US20090037363, Alex Kozlov, Jun Li, Ismail Ari, *Methods and systems for managing a data mining model*, Feb, 2009
- 3- US20080270231, Jun Li, Ismail Ari, Elsa Durante, *Methods and systems for tracking customer response to a coupon*, Nov, 2008
- 4- Ismail Ari, Melanie Gottwals, Dick Henze, *Providing an alternative caching scheme at the storage area network level*, U.S. Patent [US2006/0026229\(A1\)](#), Published Feb 2006.
- 5- Ismail Ari, Melanie Gottwals, Dick Henze, *A method for managing data storage*, Pending Patent Application, # US 200500771.

(I have 3 other pending patent applications with USPTO.)

## Thesis

- 1- Ismail Ari, *Design and Management of Globally Distributed Network Caches*, PhD Thesis, University of California Santa Cruz, September 2004. (2 citations)

## Selected Technical Reports

- 1- Mingzhu Wei, Ismail Ari, Jun Li, Mohamed Dekhil, *ReCEPTor: Sensing Complex Events in Data Streams for Service-Oriented Architectures*, HPL-2007-176
- 2- Ismail Ari, *Storage Embedded Networks (SEN) and Adaptive Caching using Multiple Experts (ACME)*, UCSC-CRL-03-01, 2003.
- 3- Ismail Ari, *Multicache Simulation Environment, Version 1.0 Reference Guide*, UCSC-CRL-03-02, 2003.
- 4- Ismail Ari, Bo Hong, Ethan L. Miller, Scott A. Brandt, Darrell D. E. Long, *Modeling, Analysis and Simulation of Flash Crowds on the Internet*, UCSC-CRL-03-15, 2004.
- 5- Ismail Ari, *Using Statistical Correlation For Dependency Analysis Of Cache Replacement Policies*, UCSC-CRL-03-16, 2004.

## Teaching and Advising

### 2005-2008 Mentoring PhD Students

I have mentored the following PhD students during their HP Labs internships. I helped shape their summer projects and PhD thesis research. Together, we have published internal and external reports or invention disclosures. Part of my teaching philosophy is to expand their breadth of knowledge, help them scope their work and prioritize tasks.

- **2008-2009 Mo Liu**, *Data Stream Warehousing*, Worcester Polytechnic Institute (WPI),  
PhD Advisor: Prof. Elke Rundensteiner (co-mentored with **Dr. Gupta and Dr. Wang**)
- **2007 Mingzhu Wei**, *Complex Event Processing as a Service*, Worcester Polytechnic Institute (WPI),  
PhD Advisor: Prof. Elke Rundensteiner
- **2006 Winnie Cheng** Information flow control and privacy-preserving access control, MIT  
PhD Advisor: Prof. Barbara Liskov (co-mentored with **Dr. Jun Li**)
- **2005 Qiaoqin Ma**, *Decentralized SAN-Level Caching using DHT*, Northeastern University  
PhD Advisor: Prof. Gene Coopernan

### Selected Invited Talks

- **2008**: Data Stream and Complex Event Processing, Istanbul Bilgi, Isik and Teditepe Universities, October.
- **2007**: Data Stream and Complex Event Processing, SSRC [Invited Talk](#), UC Santa Cruz, November
- **2006**: [SANCache](#): Performance boosting and workload isolation, IEEE Mass Storage Conference, May
- **2003**: Storage Embedded Networks and Adaptive Caching, SSRC [Hot Seat](#), UC Santa Cruz, Feb
- **2001**: Hilbert's 10<sup>th</sup> Problem is Unsolvable, Computational Complexity Project [Presentation](#), March

### 2000 Fall: Instructor, University of Maryland Baltimore County (UMBC)

For one semester, I taught the object-oriented C++ programming class at UMBC that focused on the implementation of data structures and basic search/sort algorithms. I was given a Teaching Assistant who helped

students with the projects and homeworks I assigned and with the grading of exams.

### **1998 – 2000: Research and Teaching Assistantships, UMBC Maryland Center for Telecommunications Research**

- As an RA, I studied all the networking protocols listed in the first page (especially TCP/IP) during this time and investigated methods to provide Quality of Service (QoS) within IP networks ([DiffServ](#), [IntServ](#)). I learned network programming using sockets. I implemented a module in the NIST simulator for [SSCOP](#) reliable connection-oriented transport protocol for ATM switches.
- As a TA, I helped teach the Computer Architecture CS411 class to senior computer science students: arithmetic, pipelining, CPU, registers, cache, I/O.
- Selected MS courses: I implemented a distributed, load balancing scheduler in Unix using Java for OS class. I compared [ad-hoc](#) routing protocol performances for the wireless networking class.

## **Awards and Honors**

- April 2009 IBM Smarter Planet University Jam “Top Faculty Contributor Award” \$20,000.**
- 2002-2009** My current scientific impact (Harzing) index is 5.
- 2002-2003** My [Erdős number](#) is 3 due to the publications with Prof. Darrell Long and Prof. Manfred Warmuth.
- 2004-2001** My PhD research was sponsored by HP Labs over 3 years with a research grant of **\$100,000**.
- 1998-1994** I finished B.S.E.E with an **honors degree** and graduated from high school with a **2<sup>nd</sup> rank**
- 1994** **Ranked 11<sup>th</sup>** among 1.2 million students in Turkish National University Entrance Exam. I won several awards and scholarships from governmental, educational, and private institutions.

## **Professional Activities**

### **Organizing Program Committee (PC) Member for:**

- Web Information Systems Engineering (WISE) [2009](#)
- IEEE Modeling Analysis and Simulation of Computer and Telecomm Networks ([MASCOTS](#)) [2008](#), [2009](#)
- IEEE International Conference on Parallel Processing ([ICPP](#)) [2007](#), Data Intensive and I/O Computing
- ACM/USENIX Symposium on Computer Human Interaction for Management of IT ([CHIMIT](#)) [2007](#).

### **Served as an External Reviewer for the following Journals and Conferences**

- Turkish Journal of Electrical Engineering and Computer Sciences
- IEEE Transactions on Computing (TOC)
- ACM Transactions on Database Systems (TODS)
- ACM Transactions on Computer Systems (TOCS)
- The Computer Journal (Oxford Press)
- Turkish Journal of Electrical Engineering and Computer Sciences
- **2008-2007** IEEE GlobeCom, MASCOTS, ICPP, CHIMIT
- **2004-2005** ACM SIGMETRICS 2004, IEEE MASCOTS 2004-2005, IEEE/NASA MSST 2004-2005.
- **2003** IEEE ICDCS
- **2002** USENIX FAST 2002 and ;Login Magazine

### **Professional Memberships**

- ACM, IEEE, USENIX
- [TASSA](#) (Turkish American Scientists and Scholars Association)

### **Professional Courses**

- Supercomputing and Parallel Programming Workshop: Istanbul Technical University, June 15-26, 2009
- Project Management: One week long applied class covering fundamentals of PM (counts towards PMI)
- Star@Work: One week long applied class teaching leadership, team dynamics, time management, etc.

## Programming Skills, Platforms, and Tools

*I can teach or build systems using these programming languages, tools, protocols and platforms.*

**Programming Languages:** C#, C++, C, Java, Perl, TCL, ASP.NET, SQL

**Web site and Web service** development using IIS Server with ASP.NET and C#

**DB: SQL Server (SS)**, ADO.NET, SS Data Analysis/Mining, OLAP, Notification Services

**Middleware:** Microsoft Biztalk, WWF, Business Rule Engine, SS Integration Services

**Packages:** MS Office (Word, Excel, PowerPoint, Visio, Access). Matlab.

**Software Engineering & Collaboration Suites:** Visual Studio.NET, Eclipse, CVS, NUnit/Testing, SRS docs, twiki, SharePoint, Moodle CMS

**Enterprise design patterns:** 3-Layered Architecture, MVC, Factory, Observer, Singleton, SOA.

**OS Kernel & Driver Dev:** Windows *driver*, *Kernel-level dev.& drivers* in Linux, FreeBSD

**Mobile Platforms:** Windows Mobile Smart Device and *Forms-Desktop* apps in VS.NET

**Network Protocols:** HTTP, SOAP/XML, NFS, TCP/IP. *Routing:* OSPF, RIP, BGP, SSCOP, *Ad-hoc.* MAC: 802.11, 802.3. *A/V:* SIP, RTP/RTCP. *Storage:* ISCSI, FCIP.

**Distributed:** Web Services, CORBA, Java RMI, Socket-level/Threads/Parallel programs, MPI, openmp

## Extracurricular Activities

- Helping children in underprivileged neighborhoods learn basic math and science skills through fun activities (e.g. preparing for [Lego Mindstorm](#) competition)
- Engaging in various sports activities, traveling, and cooking, reading books, watching movies, social networking and other family activities.
- **Personal Information:** **Military Service** completed, **Marital Status:** Married, **Residency** Turkish citizen and US Permanent Resident  
**Qualifications** Strong and diverse technical background, analytical thinking, attention to detail, excellent communication skills, willing to learn and teach, initiator, team player, can prioritize tasks.

## References

### Advisor:

Prof. Ethan L. Miller, UC Santa Cruz, [elm@cs.ucsc.edu](mailto:elm@cs.ucsc.edu) (831) 459-1222  
1156 High St Santa Cruz, CA, USA

### Advisor:

Prof. Darrell D E Long, UC Santa Cruz, [darrell@cs.ucsc.edu](mailto:darrell@cs.ucsc.edu)  
1156 High St Santa Cruz, CA, USA

### Manager:

Dick Henze, HP Labs, [dick.henze@hp.com](mailto:dick.henze@hp.com)  
1501 Page Mill Rd, Palo Alto, CA, 94304, USA

### Vice Rector for Research and Technology @ Ozyegin University

Prof. Sirin Tekinay, [sirin.tekinay@ozyegin.edu.tr](mailto:sirin.tekinay@ozyegin.edu.tr), +90-216-559-2380  
Kusbakisi Cad, No 2, Altunizade 34662, Uskudar, Istanbul, Turkey

**Other references are available upon request.**