

Amitanand S. Aiyer

Department of Computer Sciences
C00500, University of Texas at Austin,
Austin, TX-78712
anand@cs.utexas.edu
<http://www.cs.utexas.edu/users/anand/>

3106 Duval Street, #302
Austin, TX-78705
(512) 297-0180

- OBJECTIVE Obtain a full-time position involving *research* and *development*.
- INTERESTS **Distributed Systems, Algorithms, and Security.**
- EDUCATION Ph.D. Exp. *Summer 2010*
Department of Computer Sciences,
University of Texas at Austin.
- M.S.C.S. *May 2006*
Department of Computer Sciences,
University of Texas at Austin.
- B. Tech. *Aug 2003*
Department of Computer Science and Engineering,
Indian Institute of Technology - Madras.
- PUBLICATION **Referred conference publications**
- ◇ Amitanand S. Aiyer, Lorenzo Alvisi, Rida A. Bazzi and Allen Clement.
Matrix Signatures: From MACs to Digital Signatures. 22nd International Symposium on Distributed Computing (DISC 2008), September 22-24, 2008, Arcachon, France.
 - ◇ Amitanand S. Aiyer, Lorenzo Alvisi and Rida A. Bazzi.
Bounded Wait-Free Implementation of Optimally resilient Byzantine Storage without (Unproven) Cryptographic assumptions. 21st International Symposium on Distributed Computing (DISC 2007), September 24-26, 2007, Lemesos, Cyprus.
 - ◇ Harry C. Li, Allen Clement, Amitanand S. Aiyer and Lorenzo Alvisi.
The Paxos Register. 26th IEEE International Symposium on Reliable Distributed Systems (SRDS 2007), October 10-12, 2007, Beijing, China.
 - ◇ Amitanand S. Aiyer, Lorenzo Alvisi and Mohamed G. Gouda.
Key Grids: A Protocol Family for Assigning Symmetric Keys. IEEE International Conference on Network Protocols (ICNP 2006), November 12-15, 2006, Santa Barbara, California, USA.
 - ◇ Amitanand S. Aiyer, Lorenzo Alvisi and Rida A. Bazzi.
Byzantine and Multi-writer K-quorums. 20th International Symposium on Distributed Computing (DISC 2006), September 18-20, 2006, Stockholm, Sweden.
 - ◇ Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Micheal Dahlin, Jean-Philippe Martin and Carl Porth.
BAR Fault Tolerance for Cooperative Services. 20th Symposium on Operating Systems Principles (SOSP 2005), October 23-26, 2005, Brighton, UK. (*Award Paper*)
 - ◇ Amitanand S. Aiyer, Lorenzo Alvisi and Rida A. Bazzi.
On the availability of Non-strict Quorum Systems. 19th International Symposium on Distributed Computing (DISC 2005), September 26-29, 2005, Cracow, Poland.

- ◇ Amitanand S. Aiyer, Sanketh Indarapu, Srinathan Kannan, Vinod Vaikuntanathan and C. Pandu Rangan.

Distributed Consensus in the presence of Sectional Faults. 22nd ACM Symposium on the Principles of Distributed Computing (PODC 2003), July 13-16, 2003, Boston, Massachusetts, USA.

Referred short articles

- ◇ Amitanand S. Aiyer, Eric Anderson, Xiaozhou Li, Mehul A. Shah and Jay J. Wylie.
Consistability: Describing Usually Consistent Systems. Fourth Workshop on Hot Topics in System Dependability (HotDep 2008), December 7, 2008, San Diego, California, USA.
- ◇ Amitanand S. Aiyer, Lorenzo Alvisi and Rida A. Bazzi.
Bounded Wait-Free Implementation of Optimally resilient Byzantine Storage without (Unproven) Cryptographic assumptions. 26th ACM Symposium on the Principles of Distributed Computing (PODC 2007), August 12-15, 2007, Portland, Oregon, USA.

Patents

- ◇ *Multimodal object de-duplication.*
Jin Li, Li-wei He, Sudipta Sengupta and Amitanand Aiyer.
Microsoft Corporation – Seattle, WA.

SELECTED PROJECTS

- ◇ **Relaxed consistency for higher availability**
Large scale storage systems aim to provide good consistency and good availability. However, in presence of network partitions achieving both these objectives simultaneously is not possible. Existing approaches have either relied on strict quorum systems that enforce good consistency giving up on availability; or have relied on probabilistic quorum systems that provide high availability but may return incorrect values during adversarial conditions. We propose a solution based on K-quorums that achieves a middle ground between strict quorum systems and probabilistic quorum systems. K-quorums provides a strict bound on the staleness of a returned value, while at the same time achieving higher availability whenever possible.
- ◇ **Atomic wait-free storage without digital signatures**
Atomic wait-free storage systems require readers to update the state at the servers. When readers may be Byzantine, servers cannot trust the values from the readers; thus, existing protocols have relied on digital signatures, based on unproven cryptographic assumptions, to tolerate malicious readers. We have proposed a storage protocol that tolerates Byzantine readers without relying on any unproven cryptographic assumptions, by leveraging secret sharing techniques instead. Our solution also reduces the number of replicas required to tolerate f Byzantine faults from $4f + 1$ to $3f + 1$. This protocol is optimal in the number of replicas and uses a bounded amount of storage and communication.
- ◇ **Chunk based de-duplication**
Large storage systems may contain duplicate data across different files and their versions. Chunk based de-duplication mechanisms save storage by detecting common chunks across various files. De-duplication using a smaller chunk sizes is preferable for efficiency, but increases the amount of metadata to be stored. We have designed and evaluated a multimodal chunker that combines different strategies for chunking, based on the size and content of the file, to efficiently save space while maintaining a small amount of metadata.
- ◇ **Formally verifying a large scale storage system**
Large scale storage systems are designed to work under a variety of conditions and may be required to continue operating under conditions involving node failures and network partitions. Verifying the guarantees provided by such storage protocols for different operating conditions can be cumbersome and error-prone. We have specified the functioning of a large scale key-blog archive, that is spread over multiple data centers, using the +CAL language and verified the consistency semantics provided by the protocol for different failure conditions.

◇ **Fault Tolerance for Cooperative Services**

Byzantine Altruistic and Rational (BAR) fault model captures the self-interest prevalent in systems spread across multiple administrative domains. In addition to Byzantine faults that occur due to software bugs and virus attacks, rational participants may also deviate from the protocol to maximize their own self-interest. Thus, protocols spanning multiple administrative domains need to tolerate Byzantine nodes as well as rational nodes. Relying on game-theoretic notion of an equilibrium, we have developed a replicated state machine protocol for the BAR model and used it to implemented a cooperative backup system.

◇ **Replacing digital signatures with MACs**

Digital signatures provide stronger guarantees than message authentication codes (MACs). However, digital signatures are 2-3 orders of magnitude slower to compute. Thus, many distributed systems strive to avoid the use of digital signatures, replacing them with MACs whenever possible. Matrix-signatures is a construct that captures the collective knowledge among a set of servers using pair-wise message authentication codes (MACs). This construction simplifies the conversion process by providing all the digital signatures' properties using only message authentication codes.

◇ **Key distribution protocols for sensor networks**

Secure communication between n processes requires each process to stores $n - 1$ keys, one for every other process. However, if processes do not collude with one another secure communication can be achieved using far fewer keys. We propose a family of key distribution protocols, known as Key-Grids, that drastically reduce the required number of keys for a collection of n sensor nodes to communicate securely with each other. Key-Grids trades-off the number of keys that a node needs to be store against the maximum size of colluding nodes that the protocol can tolerate. In the absence of any collusion, we show that nodes can securely communicate with each other with just $O(\log^2 n)$ keys.

WORK EXPERIENCE	<p>Research Intern <i>Summer 2008</i> HP Labs – Palo Alto.</p> <p>Research Intern <i>Summer 2007</i> Microsoft Research – Redmond.</p> <p>Engineering Intern <i>Summer 2006</i> Google Inc – New York.</p> <p>Software Development Intern <i>Summer 2005</i> Microsoft IDC – Hyderabad.</p> <p>Research/Teaching Assistant <i>Fall 2003 - Present</i> The University of Texas at Austin.</p>
-----------------	---

AWARDS	<p>◇ Obtained <i>All India 18th Rank</i> among over 110,000 candidates in IIT-JEE 1999 (2nd in South India).</p> <p>◇ Awarded <i>Prathibha Scholarship for Academic Excellence</i> from The State Government of Andhra Pradesh for the years 1999-2003.</p> <p>◇ Secured 13th position in <i>Algorithmic Intensive Online Programming Contest: BitWise 2k+2</i> conducted by IIT-Kharagpur.</p> <p>◇ Secured 1st position in <i>Shaastra Programming Contest 2001</i>, conducted by IIT-Madras.</p>
--------	---

COURSES	<p><i>Graduate Courses:</i> Distributed Computing, Algorithms: Techniques and Theory, Advanced Operating Systems, Advanced Security, (audited) Design and Implementation of Trusted Services, Information Theory, Combinatorics and Graph Theory, Artificial Intelligence, Computer Graphics.</p> <p><i>Undergraduate Courses:</i> Design and Analysis of Algorithms, Data Structures, Operating Systems, Computer Networks, Compilers, Database Management Systems.</p>
---------	--

ACTIVITIES Playing Tabla and Sitar. Listening to Indian classical music.
Tennis, Badminton and Climbing.

Prof. Lorenzo Alvisi
Department of Computer Sciences
University of Texas at Austin
lorenzo@cs.utexas.edu

Prof. Rida A. Bazzi
School of Computing and Informatics
Arizona State University
bazzi@asu.edu

REFERENCE

Prof. Mohamed G. Gouda Department of Computer Sciences University of Texas at Austin gouda@cs.utexas.edu	Dr. Jay J. Wylie HP Research Labs Palo Alto jay.wylie@hp.com
---	---

Dr. Sudipta Sengupta
Researcher
Microsoft Research
sudipta@microsoft.com

Li-Wei He
Principal Architect
Microsoft Research
lhe@microsoft.com