

MUTHIAN SIVATHANU

University of Wisconsin Madison
1210 West Dayton Street
7358 Computer Sciences,
Madison, WI 53706

Email: muthian@cs.wisc.edu
Phone: (608) 262 6623 (work)
Fax: (608) 262 9777
<http://www.cs.wisc.edu/~muthian>

RESEARCH INTERESTS

File and storage systems, distributed systems, operating systems.

EDUCATION

COMPUTER SCIENCES DEPARTMENT, UNIVERSITY OF WISCONSIN MADISON

Ph.D. (Computer Science)

Jun 2005

Advisors: Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau.

M.S (Computer Science)

Dec 2001

COLLEGE OF ENGINEERING GUINDY, ANNA UNIVERSITY, INDIA

B.E (Computer Science and Engineering)

May 2000

AWARDS AND DISTINCTIONS

- Outstanding Graduate Student Research Award, 2005, by UW Computer Sciences Department.
- Best student paper award in FAST 2004
- Awarded University Fellowship by the University of Wisconsin Madison during the year 2000 - 2001
- Awarded the Colgate Gold Medal Endowment Award, 2000, for being the top ranker in the B.E degree program in Anna University.
- Awarded the University Gold Medal, 2000 by Anna University for being the first rank holder in the CS department.
- Awarded the Arvind Mehta Memorial Award, 1997 by College of Engineering , Anna University for outstanding proficiency in English
- Secured state II Rank in science in X std. Matriculation Examination, 1994

PUBLICATIONS

Journal Publications

[*TOS05*] Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Improving Storage System Availability with D-GRAID", To Appear in ACM Transactions on Storage, March 2005.

Conference Publications

[*OSDI04*] Muthian Sivathanu, Lakshmi N. Bairavasundaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Life or Death at Block Level", Proceedings of the Sixth Symposium on Operating Systems Design and Implementation (OSDI '04), December 2004;

- [ISCA04] Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs", Proceedings of the 31st International Symposium on Computer Architecture (ISCA '04), June 2004;
- [FAST04] Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Improving Storage System Availability with D-GRAID", Proceedings of the Third USENIX Conference on File and Storage Technologies (FAST '04), March 2004;
Best Student Paper Award.
- [FAST03] Muthian Sivathanu, Vijayan Prabhakaran, Florentina I. Popovici, Timothy E. Denehy, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau; "Semantically-smart Disk Systems", Proceedings of the Second USENIX Conference on File and Storage Technologies (FAST '03), March 2003
- [ASPLOS02] Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau; "Evolving RPC for Active Storage", Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X), October 2002.
- [HIP01] Muthian Sivathanu, Venkateshwaran Venkataramani, and Remzi H. Arpaci-Dusseau; "Block Asynchronous I/O - An Infrastructure for High-Performance User-Level File systems", Proceedings of the 8th International Conference on High Performance Computing, December 2001 ; LNCS 2228 Springer 2001, p. 249 - 261.

Technical Reports

- [ANON02] Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau; "Adaptive Control for Anonymous Network-Attached Storage", February 2002.
- [MS01] Muthian Sivathanu; "The WiND Filesystem (WFS)", Masters' Project Report, Computer Sciences Department, University of Wisconsin Madison, December 2001.

DISSERTATION RESEARCH

Title: *Semantically-smart Disk Systems*

My thesis research proposes and evaluates a new class of storage systems that are aware of how the file system or DBMS is using them, and exploits this knowledge to implement new optimizations and functionality that are currently impossible to implement. A straightforward way to making storage systems semantically-aware is to change the interface to storage to convey richer information; however, modifying an interface as basic as SCSI is not very pragmatic, since it requires broad industry consensus and raises various legacy issues. In contrast, semantically-smart storage systems automatically track higher level semantic information about the file system or DBMS, by careful observation of block level reads and writes from underneath an unmodified SCSI interface. By requiring no changes to the existing interface, semantically-smart disks enable seamless deployment and adoption of new functionality.

I have demonstrated the utility of semantic knowledge within a storage system by design and prototyping of a variety of new functionality aimed at improving the availability, security, and performance of storage systems. In each of these cases, semantic knowledge was shown to enable functionality that is otherwise impossible to provide in the storage hierarchy. More details of my thesis research are available in my research statement.

EXPERIENCE

Graduate student researcher, UW ADSL group

A Logic of File systems

In order to facilitate reasoning about the extent of knowledge that can be tracked within a semantic disk underneath a modern file system, I formulated a logic and proof system for reasoning about the on-disk interaction of a file system. In addition to enabling reasoning about semantic disks, the logic was shown to be useful to file systems in general. Various file system consistency properties can be easily expressed using the logic and systematically proved from a basic set of postulates [LOGIC04].

Sep '04 - Nov '04

Semantic Disks for Database Systems

This work involved exploring the extraction of semantic information underneath a DBMS. With a prototype implementation underneath the Predator/Shore DBMS, it was shown that DBMS specific information could be reliably tracked from within the storage system. Some of the same case studies implemented in the context of file systems were studied in the DBMS context [DB04].

Feb '04 - May '04

Semantically-smart disks for file systems

My first piece of work on semantic disks involved exploration of the basic techniques for tracking dynamic file system specific information [FAST03]. Although I demonstrated in this work that functionality such as secure delete and journaling could be implemented within a semantic disk, it was under the assumption of a restricted (synchronous) file system behavior. In subsequent work, these restrictions were relaxed.

Subsequent to this, I designed, implemented and evaluated a series of case studies that demonstrate the utility of semantic knowledge within a storage system. In D-GRAID, semantic intelligence was utilized to greatly improve the availability of the storage system under catastrophic failures [FAST04]. In more recent work, information on semantic liveness was shown to enable a variety of optimizations within storage, specifically the ability to delete data securely [OSDI04]. I was also involved in X-RAY, a storage system that utilizes semantic information to perform exclusive caching [ISCA04]. Implementing these case studies involved discovery of fundamental limits to the accuracy of tracking semantic information, and novel techniques to circumvent such uncertainty. Since this exploration was performed underneath a variety of file systems, it also led to the identification of key dynamic file system properties that are relevant to semantic disks [OSDI04].

Apr '02 - Jan '04

Distributed file systems: WFS and Scriptable RPC

In the context of the Wisconsin Network Disks (WiND) project, I designed and implemented a distributed file system called WFS targeted at network attached disks [MS01]. This file system was then used as the platform to explore and evaluate a new communication paradigm for distributed systems, called Scriptable RPC (SRPC), a non-intrusive addition to traditional RPC [ASPLOS02]. I also explored distributed file system security in WFS, proposing a mechanism for access anonymity in a distributed file system, that automatically adapts its message traffic in order to reduce network traffic overhead [ANON02].

Jan '01 - Mar '02

Intern, Microsoft Inc, Redmond (Advanced Operating Systems group)

Was part of a small core team that is working on a new operating system kernel from scratch. I designed and prototyped a key component of the kernel's IPC subsystem.

Jun '04 - Aug '04

Intern, Google Inc. Mountain View (Infrastructure group)

Designed and implemented various techniques for performance isolation on Google's production servers running Linux. The techniques extended across all major resources such as CPU, memory, disk I/O, network, and the OS file cache.

May '03 - Aug '03

Intern, IBM Almaden Research Center, San Jose (Storage Systems Group)

Designed an architecture for scalable, decentralized resource management in wide-area storage. The architecture, targeted at wide-area storage service providers, enables efficient replication of competing client virtual disks, in terms of the number of replicas and their geographic placement.

May '02 - Aug '02

Intern, Hewlett Packard Labs, Palo Alto (Storage Systems program)

Designed and implemented a protocol for scalable, fault-tolerant wide area replication that permits concurrent conflicting updates from all replicas in a geographically distributed storage system. A prototype system was built with clients using iSCSI to communicate with disk nodes, which themselves co-ordinate using RPC.

May '01 - Aug '01

TECHNICAL TALKS

Life or Death at Block-Level,
OSDI '04, Dec 2004.

Improving Storage System Availability with D-GRAID,
FAST '04 (Mar 2004), HP Labs (Aug 2003), IBM Almaden(Aug 2003).

Performance Isolation on Google Servers,
Google Inc, Aug 2003.

Evolving RPC for Active Storage,
ASPLOS '02 (Oct 2002), HP Labs (Aug 2002).

Scalable Resource Management for Wide-Area Storage,
IBM Almaden, Aug 2002.

Scalable Fault-tolerant Replication in Wide-Area Storage,
HP Labs, Aug 2001.

The WiND Filesystem,
HP Labs, June 2001.

PROFESSIONAL ACTIVITIES

Reviewed papers for:

FAST '03

USENIX '03

SOSP '03

ICDCS '04

USENIX '04

TOCS '04

VLDB '04

TOS '04

USENIX '05

Member of ACM and USENIX.

PERSONAL INFORMATION

Visa Status: F1

Country of Citizenship : India

REFERENCES

Available on request.