

My primary areas of expertise and interest are distributed systems, database systems, storage systems and high-performance data centers. I also have great interest in applying machine learning and queuing analysis to help simplify and automate management for such systems.

**EDUCATION****Carnegie Mellon University, Pittsburgh** **Fall 2002 – May 2007 (expected)**

- Ph.D. Candidate in Electrical and Computer Engineering
- Advisor: Prof. Gregory R. Ganger
- PhD Thesis: “Enabling what-if explorations in systems”
- Masters of Science (MS) degree obtained in 2003 with thesis title “A framework for implementing background storage applications using freeblock scheduling”

**Carnegie Mellon University, Pittsburgh****Fall 1998 – Spring 2002**

- Bachelors (BS): Electrical and Computer Engineering
- Second major: Computer Science. Minor in Mathematics
- GPA: 4.0/4.0

**HONORS AND AWARDS**

- FAST 2005 Best Paper Award
- FAST 2004 Best Student Paper Award
- Cisco Systems Achievement Program Award for individual and team contribution, July 2001
- Lockheed Martin Undergraduate Research Competition Award, May 2002
- University and College Honors at Carnegie Mellon University, 1998-2002

**RESEARCH SUMMARY**

**Enabling what-if explorations in systems (lead, 2004-present):** With a large percentage of total system cost going to system administration tasks, self-management remains a difficult and important goal in systems. As a step towards the self-management vision, I have developed a framework to enable systems to be self-predicting and answer “what-if” questions about their behavior with little or no administrator involvement. I have built a Resource Advisor inside two real systems: Microsoft’s SQL Server database and the Ursa Minor storage system at CMU. The Resource Advisor helps with upgrade and data placement decisions and provides what-if interfaces to external administrators (and internal tuning modules). The Resource Advisor is based on efficient system behavioral models that enable robust predictions in multi-tier systems. The models are robust in that they discover regions of operation where the prediction confidence is high and regions for which they choose not to predict. The models handle performance anomalies by pinpointing their likely cause (e.g., system misconfigurations) and continuously collect historical information and refine themselves to account for unforeseen (and hence not programmed-in) workload-system interactions.

**ABLE: Black-box learning of file usage characteristics (co-lead, 2003-2004):** Systems can specialize their policies when given knowledge of a file’s usage patterns, such as access pattern, maximum file size, and file lifetime. Since applications and users have proven unwilling (or unable) to provide explicit hints, we explored approaches to extracting them automatically using machine learning. Specifically, we found that the name and attributes assigned to a file when it is created often provide implicit hints as to its usage patterns. Predicting usage patterns can improve performance by guiding policies for how files should be laid out on disk, whether they should be cached and how they should be replicated.

**Freeblock scheduling: A framework for building unobtrusive disk maintenance applications. (2002-2003):** Most systems utilize many background, I/O bound services that balance load, re-encode data and repair corrupted data. I developed a clean API and framework for these background applications to tap into “free bandwidth” from busy disks. For background tasks, this framework uses idle time and freeblock scheduling, a new approach to utilizing more of disks’ potential media bandwidths. Specifically, by interleaving low priority disk activity with the normal workload, one can replace many foreground rotational latency delays with useful background media transfers. With this framework, maintenance applications such as backup, cache write backs and data migration can make steady forward progress with little-to-no impact on foreground application workloads.

## PUBLICATIONS

### *Refereed conference papers*

**Argon: performance insulation for shared storage servers.** Matthew Wachs, Michael Abd-El-Malek, Eno Thereska, Gregory R. Ganger. Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST’07), San Jose, CA. February 12-16 2007 (to appear).

**Stardust: Tracking activity in a distributed storage system.** Eno Thereska, Brandon Salmon, John Strunk, Matthew Wachs, Michael Abd-El-Malek, Julio Lopez, Gregory R. Ganger. Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS’06), Saint-Malo, France. June 26-30 2006.

**Informed data distribution selection in a self-predicting storage system.** Eno Thereska, Michael Abd-El-Malek, Jay J. Wylie, Dushyanth Narayanan, Gregory R. Ganger. Proceedings of the International Conference on Autonomic Computing (ICAC’06), Dublin, Ireland. June 12-16 2006.

**Ursa Minor: Versatile cluster-based storage.** Michael Abd-El-Malek, William V. Courtright, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. Proceedings of the 4th USENIX Conference on File and Storage Technologies (FAST’05), San Francisco, CA. December 13-16 2005.

**Continuous Resource Monitoring for Self-predicting DBMS.** Dushyanth Narayanan, Eno Thereska, Anastassia Ailamaki. Proceedings of the 13th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS’05), Atlanta, GA. September 25-27, 2005.

**File classification in self-\* systems.** Michael Mesnier, Eno Thereska, Daniel Ellard, Gregory R. Ganger, Margo Seltzer. Proceedings of the International Conference on Autonomic Computing (ICAC’04), New York, NY. May 17-18, 2004.

**A Framework for Building Unobtrusive Disk Maintenance Applications.** Eno Thereska, Jiri Schindler, John Bucy, Brandon Salmon, Christopher R. Lumb, Gregory R. Ganger. Proceedings of the 3rd USENIX Conference on File and Storage Technologies (FAST’04), San Francisco, CA. March 31st - April 2nd, 2004.

### *Invited journal papers*

**Early experiences on the journey towards self-\* storage.** Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. IEEE Data Engineering Bulletin 29(3). Special issue on self-managing database systems. September 2006.

**Challenges in building a DBMS Resource Advisor.** Dushyanth Narayanan, Eno Thereska, Anastassia Ailamaki. IEEE Data Engineering Bulletin 29(3). Special issue on self-managing database systems. September 2006.

*Workshop papers*

**Observer: Keeping System Models from Becoming Obsolete.** Eno Thereska, Dushyanth Narayanan, Anastassia Ailamaki, Gregory R. Ganger. Second Workshop on Hot Topics in Autonomic Computing. Jacksonville, FL. June 15, 2007.

**Categorizing and Differencing System Behaviours.** Raja R. Sambasivan, Alice X. Zheng, Eno Thereska, Gregory R. Ganger. Second Workshop on Hot Topics in Autonomic Computing. Jacksonville, FL. June 15, 2007.

**Towards self-predicting systems: What if you could ask “what-if”?** Eno Thereska, Dushyanth Narayanan, Gregory R. Ganger. Proceedings of the 3rd International Workshop on Self-adaptive and Autonomic Computing Systems. Copenhagen, Denmark. August 22, 2005.

**A Two-Tiered Software Architecture for Automated Tuning of Disk Layouts.** Brandon Salmon, Eno Thereska, Craig A. N. Soules, Gregory R. Ganger. First Workshop on Algorithms and Architectures for Self-Managing Systems. In conjunction with Federated Computing Research Conference (FCRC). San Diego, CA. June 11, 2003.

*Technical reports (not overlapping with above)*

**Challenges in Building a Two-Tiered Learning Architecture for Disk Layout.** Brandon Salmon, Eno Thereska, Craig A.N. Soules, John Strunk, Gregory R. Ganger. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-04-109. August 2004.

**Attribute-Based Prediction of File Properties.** Daniel Ellard, Michael Mesnier, Eno Thereska, Gregory R. Ganger, Margo Seltzer. Harvard Computer Science Group Technical Report TR-14-03, December 2003.

**Design and Implementation of a Freeblock Subsystem.** Eno Thereska, Jiri Schindler, Christopher R. Lumb, John Bucy, Brandon Salmon, Gregory R. Ganger. Carnegie Mellon Technical Report CMU-PDL-03-107. December 2003.

SELECTED  
PRESENTATIONS

“Enabling what-if explorations in distributed systems.” TU Vienna, Austria. January 2007. Host: Schahram Dustdar.

“Enabling what-if explorations in systems.” Microsoft Research Cambridge. July 2006. Host: Dushyanth Narayanan.

“Stardust: tracking activity in a distributed storage system.” SIGMETRICS 2006. Saint Malo, France. June 2006.

“Enabling what-if explorations in a distributed storage system.” Laboratoire d’Informatique de Paris 6 (LIP6). Paris, France. June 2006. Host: Marc Shapiro.

“Informed data distribution selection in a self-predicting storage system.” ICAC 2006. Dublin, Ireland. June 2006.

“Ursa Minor: versatile, self-predicting, cluster-based storage.” EMC. Boston, MA, January 2006. Hosts: Georges Brun-Cottan, Jiri Schindler.

“A framework for informed tuning in a self-managing system.” 1st Eurosys doctoral workshop. Brighton, UK. October 2005.

“Continuous resource monitoring for self-predicting DBMS.” MASCOTS 2005 Conference. Atlanta, GA, September 2005.

“Continuous resource monitoring for self-predicting DBMS.” Duke University. NC, September 2005. Host: Landon Cox.

“File classification in self-\* storage systems.” NetOS Seminar. University of Cambridge, UK, August 2004. Host: Steven Hand.

“A framework for building unobtrusive disk maintenance applications.” Microsoft Research Cambridge. June 2004. Host: Dushyanth Narayanan.

“A framework for building unobtrusive disk maintenance applications.” FAST 2004 Conference. San Francisco, CA, March 2004.

#### PATENTS

**Predicting database system performance.** Dushyanth Narayanan, Eno Thereska. Pending, filed October 2005.

#### PROFESSIONAL EXPERIENCE

##### **Carnegie Mellon University** **2002 - present**

- Conducted research on various aspects of operating systems and distributed systems, with a focus on system management problems.
- Designed and deployed Stardust, an instrumentation and querying framework for monitoring a large-scale storage system (Ursa Minor).
- Designed and implemented analytical and simulation-based what-if modules inside Ursa Minor that can answer various hypothetical questions on data distribution changes, resource upgrades, and service placement.
- Designed and deployed the data migration module inside Ursa Minor.
- Co-lead the Attribute-Based Learning Environments (ABLE) project and assisted in the Continuous Reorganization project (see [www.pdl.cmu.edu](http://www.pdl.cmu.edu) for details).
- Designed and implemented a novel disk scheduling algorithm (freeblock scheduling) inside the FreeBSD operating system.

##### **Microsoft Research, Cambridge, UK** **May 31 - August 20, 2004**

- Augmented the next-generation SQL Server (version 2005 then) with a resource planning advisor. The advisor monitors the resources consumed by a workload and helps database administrators determine the performance impact when upgrading those resources.

##### **Cisco Systems** **May 13 - July 27, 2002**

- Implemented a complete IOC Client simulator that allowed software and micro-code developers to test their IOS Client code and IXP2800 micro-code respectively, before the hardware (IXP2800 routers) was available.

##### **Cisco Systems** **January 8 - July 31, 2001**

- Implemented extra features regarding security for an edge-router, wrote IOS Port Adapter driver code and worked closely with other developers in lab to test the software.

##### **Oracle Corp.** **June 1 - August 6, 2000**

- Worked on Self-Service Web Applications and their back-end databases.

#### TEACHING EXPERIENCE

Guest lecturer for Performance Modeling (15857) during Fall '06

Guest lecturer for Operating Systems (15410) during Spring '04, Fall '04, Fall '05, Spring '06  
Teaching Assistant for Storage Systems (18746), Fall '05  
Teaching Assistant for Distributed Systems (18842), Spring '05  
Teaching Assistant for Computer Architecture (18347), Spring '02  
Undergraduate Laboratory Assistant for Intro to EE (18220), Fall '00

SERVICE

Reviewer for several conferences, including Eurosys, FAST, OSDI, SIGMOD.  
Reviewer for ACM Computing Reviews.  
Shepherding committee for the First Eurosys Authoring Workshop'06.  
SOSP 2007 Shadow PC.

PROFESSIONAL  
AFFILIATIONS

- Parallel Data Lab (PDL), Carnegie Mellon University
- Storage Network Industry Association (SNIA)
- Institute of Electrical and Electronics Engineers (IEEE)
- Association for Computing Machinery (ACM)
- SIGOPS Eurosys Chapter Member (and Volunteer)
- Phi Kappa Phi Honors Society, Carnegie Mellon Chapter
- The National Society of Collegiate Scholars (NSCS)

OTHER  
AFFILIATIONS

- Animal Friends of Pittsburgh: volunteer
- Pittsburgh Filmmakers: photographer

LANGUAGES

- English, German, Italian (studied), French (studied), Albanian (mother tongue)

REFERENCES

**Dr. Greg Ganger**  
Professor of ECE & CS  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA, USA  
ganger@ece.cmu.edu  
tel: +1-412-268-1297

**Dr. Anastassia Ailamaki**  
Associate Professor of ECE & CS  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA, USA  
natassa@cs.cmu.edu  
tel: +1-412-268-7848

**Dr. Mor Harchol-Balter**  
Associate Professor of CS  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA, USA  
harchol@cs.cmu.edu  
tel: +1-412-268-7893

**Dr. Dushyanth Narayanan**  
Researcher  
Microsoft Research  
7 J J Thomson Avenue  
Cambridge, CB3 0FB, UK  
dnarayan@microsoft.com  
tel: +44-1223-479808

**Dr. John Wilkes**  
Hewlett-Packard Fellow  
Hewlett Packard Laboratories  
1501 Page Mill Road, MS 1134  
Palo Alto, CA 94304-1100  
john.wilkes@hp.com  
tel: +1-650-857-3568