

RAJESH KRISHNA BALAN

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891

Office: (412) 268-3778
Home : (412) 441-6628
Cell : (412) 657-1353
Fax : (412) 268-5576
rajesh@cs.cmu.edu
<http://www.cs.cmu.edu/~rajesh/research/>
Citizenship: Singaporean, Visa status: F1 student

RESEARCH AREAS

Mobile and pervasive computing, operating systems; distributed computing; software engineering; usability of software systems; usability testing; networking

EDUCATION

Carnegie Mellon University, Pittsburgh, PA
Ph.D. in Computer Science, 2006 (expected)
Thesis Title: Simplifying Cyber Foraging
Advisor: Mahadev Satyanarayanan

National University of Singapore, Singapore
Masters in Computer Science, July 2000
Thesis Title: CHAMELEON – A System for Adaptive QoS Provisioning
Advisor: Akkihebbal L. Ananda

National University of Singapore, Singapore
B.Sc in Computer Science (Honours), Nov 1998
Thesis Title: Distributed Computing for the CAVE Virtual Reality Display Environment
Advisor: Akkihebbal L. Ananda

RESEARCH EXPERIENCE

My research has encompassed three main areas; 1) mobile computing / operating systems / distributed systems, 2) software engineering / usability testing, and 3) network protocols. I present a brief synopsis of this previous research, beginning with the most recent. Research items marked with a ★ are described in more detail in my research statement.

Improving Interactive Scientific Applications (Operating Systems / Usability Testing)

Carnegie Mellon University and University of Toronto, Fall 2005 - Present.

Submitted to EuroSys 2006

Interactive scientific applications typically have two phases; a compute phase that is memory, CPU, or network bound, and an interactive phase where a user interacts with the results of the compute phase. These phases can repeat multiple times. To improve the performance of such applications, I am working with researchers at Carnegie Mellon and the University of Toronto to develop a system, called AgentISR, that uses virtual machines to seamlessly and transparently migrate these applications automatically between a compute farm and the user's desktop machine. This results in the best possible performance during either application phase; the compute phase is able to use the resources available at the compute farm and users can crisply interact with the application locally during the interactive phase. A paper detailing the basic operation of AgentISR has been submitted to EuroSys 2006. For this paper, my key contribution was a user study demonstrating that a thin client solution was ineffective. In addition, I also assisted with the rest of the evaluation. We are currently extending AgentISR by testing its performance with many more applications and by developing policies for automatically migrating applications.

★ Developer Support for Cyber Foraging (Software Engineering / Usability Testing)

Ph.D. Thesis, Carnegie Mellon University, Spring 2003 - Present.

Cyber foraging refers to the opportunistic use of server resources in the environment by a mobile client. These resources allow clients to execute applications that they could not run locally. However, before cyber foraging can become a reality, developers must be able to quickly, easily, and effectively retarget existing applications to support cyber foraging. Previous methods of doing this used to take two to four weeks per application; an unacceptably long time. This problem is made more difficult as a good solution should be usable by even novice developers and should work with applications written in any language and style. As part of my thesis, I developed and rigorously validated a process, called RapidRe, that provides a solution. It allows even novice developers to retarget applications, written in arbitrary languages, to support cyber foraging in under 4 hours. In addition to dramatically reducing the retargeting time, my work is one of the first that rigorously validates the usability of a software system. This validation was performed via a software usability study involving 13 participants and 8 different applications (written in a variety of languages including C, C++, Java, Tcl/Tk, and Ada). This aspect of my research is described in more detail in my research statement.

★ Operating System Support for Cyber Foraging (Mobile Computing / Operating Systems)

Ph.D. Thesis, Carnegie Mellon University, Spring 2001 - Present.

Published at SIGOPS EW 2002 and MobiSys 2003.

RapidRe allows developers to quickly and easily retarget applications to support cyber foraging. To ensure that these retargeted applications also had excellent performance, I developed an adaptive dynamic runtime component called Chroma. This runtime is a key component of RapidRe and provides the common runtime components necessary for effectively benefiting from cyber foraging. For Chroma to be effective, I needed to figure out the optimal application settings (fidelity settings and runtime partitioning) that best matched

the available resources and user preferences. To do this, I integrated a number of resource measurement modules and a resource prediction module into Chroma. I then built an exhaustive yet efficient solver that picks the appropriate runtime application settings. I also developed an easy and effective way for Chroma to receive user preferences. I then showed that retargeted applications, using Chroma, can achieve excellent performance in a variety of mobile scenarios. A particular novel aspect of Chroma is that it can use extra server resources to automatically improve application performance. I describe this work in more detail in my research statement. The design and validation of Chroma was presented initially at SIGOPS EW 2002 and in more detail at MobiSys 2003.

★ Distributed Infrastructure for Massively Multiplayer Games (Distributed Systems)

IBM Research Watson and Carnegie Mellon University, Summer 2003 - Spring 2005.

Published at Middleware 2005. Submitted to IEEE TC Journal

During my internship at IBM Research Watson, I worked with Maria Ebling, Paul Castro, and Archan Misra to investigate the problem of building a distributed infrastructure for massively multiplayer online games. The scope of the problem was immense as these games could potentially have up to 1,000,000 players in the same game. The key to allowing such a large number of players is to distribute the game environment among many (10,000 or more) different servers; each serving a subset of the total player population. Building and managing such an infrastructure requires overcoming many different challenges. For my research, I concentrated on the problems of a) ensuring that each game server was as consistent as it needed to be, b) developing mechanisms to automatically handle load spikes on any game server (caused by hotspots or other events), and c) ensuring that the infrastructure is attractive to and easy to use by real game developers. We called our infrastructure Matrix. I built a prototype implementation of Matrix and showed that a) Matrix could automatically add and remove servers to handle even very dense load hotspots, b) Matrix has reasonable overheads, and c) Matrix is able to satisfy the demands of real game players (verified with a user study explained below). My research statement contains more details of this work. This work was presented at Middleware 2005 and has been submitted to the IEEE TC Journal.

Investigating the Effect of Matrix on Game Players (Usability Testing)

IBM Research Watson and Carnegie Mellon University, Summer 2004

Published at Middleware 2005. Submitted to IEEE TC Journal

A key design consideration for Matrix was that it should not add any perceivable latency to a game player's experience. To validate this, I conducted a user study, with the help of Maria Ebling, in the summer of 2004. We asked 5 participants to play a game called Bzflag that had been modified to use Matrix. During the user study, load was introduced and removed from the game to trigger Matrix's load balancing algorithms. Overall, our results were very promising and the participants perceived almost no additional Matrix-induced latency. The full user study details appear in the IEEE TC journal submission.

Multi-Modal Network Protocols (Network Protocols)

Carnegie Mellon University, Spring 2001 - Fall 2001.

Presented as a poster at SIGCOMM 2001 and SOSP 2001

In this work, I, together with Srinivasan Seshan and Aditya Akella from Carnegie Mellon University, looked at developing network protocols that could change behaviour based on the current situation. We looked at two main problems. First, we looked at ways to modify TCP's behaviour in order to save power. We showed that by changing TCP's delayed-ack threshold, you could actually reduce the total amount of energy needed to transfer data, at the loss of some throughput. Second, we looked at developing an overlay routing infrastructure that would pick the appropriate overlay routing algorithm for the current overlay. For example, for small overlays, you would pick a flat routing algorithm whereas for large overlays, a hierarchical routing algorithm would be better. This work was presented as a poster at SIGCOMM 2001 and SOSP 2001.

Porting Linux to the iPAQ (Operating Systems)

Carnegie Mellon University, Fall 2000 - Spring 2001.

In Fall 2000, I was asked to port Linux to the iPAQ for a DARPA demo. I used the ARM Linux distribution and kernel created by Compaq developers as a starting point. To complete the port, I needed to a) fix bugs in the infrared network driver, and b) write a new driver to access flash RAM on the iPAQ. The infrared fixes and new driver were subsequently merged into the Linux-iPAQ source tree.

Improving TCP's performance over wireless links (Network Protocols)

Masters Thesis, National University of Singapore, Spring 1999 - Summer 2000.

Published at INFOCOM 2001 and as an extended journal paper (2002 Elsevier's Computer Networks)

As part of my masters thesis, I developed an extension to TCP that improved its performance over wireless links. The key idea was to add an extra option to TCP that contained the checksum of the header of the TCP packet. When a packet got corrupted by the wireless medium, this extra checksum was checked. If the checksum passed, the information in the header was used to tell the sender of the packet that corruption of the packet, and not congestion, had occurred. The sender would then re-send the packet without reducing its sending rate or timing out. This extension to TCP was implemented in the Linux kernel 2.2.10 and its performance was tested under a variety of wireless conditions. The results showed that the extension significantly improved TCP's performance in the presence of frequent packet corruption. These results were presented at INFOCOMM 2001.

Using TCP Tunnels to Improve Gateway Performance (Network Protocols)

Masters Thesis, National University of Singapore, Spring 1999 - Summer 2000.

Published at LCN 2000 and as an extended journal paper (2002 Elsevier's Computer Networks)

The other half of my masters thesis involved the use of TCP tunnels to provide QoS guarantees at gateway routers. TCP tunneling is similar to IP tunneling except that TCP connections are being tunneled over another TCP connection instead of over an IP tunnel. The benefits of this approach were that the tunnels used TCP's congestion control mechanisms and were able to share bandwidth fairly with other TCP tunnels. I built a system, CHAMELEON, that used TCP tunnels between gateway routers to provide QoS guarantees

for application traffic using those routers. This work is detailed in my Masters thesis. I also showed, in a paper presented at LCN 2001, how TCP tunnels could be used to prevent congestion collapse from occurring at these gateway routers.

Distributed Computing for the CAVE Virtual Reality Environment (Distributed Systems)

Honours Thesis, National University of Singapore, Spring 1998 - Fall 1998.

Published at a local high performance computing conference in 1999

The CAVE is a virtual reality environment, developed by the University of Illinois at Chicago, that uses a room to completely immerse the user in the virtual environment. It requires powerful computers to handle the graphics requirements of the virtual environment. As such, the servers controlling the CAVE do not have many free CPU cycles to do computation on behalf of programs wanting to use the CAVE. For my honours thesis, I wrote software that allowed the computational aspect of an application to be distributed to other machines, thus leaving the main servers free to concentrate solely on the graphics requirements of the CAVE. This work was published in a high performance computing conference

Creation of 3D Meshes (Computer Vision)

Exchange Student, Berkeley, CA, Spring 1997.

I was an exchange student at the University of California at Berkeley during the Spring of 1997. While I was there, I worked with Dr. T. T. Blackmon who was part of the Intelligent Mechanisms Group at NASA Ames. Dr. Blackmon was responsible for creating a 3D virtual world from the images returned by the Pathfinder Probe that landed on Mars in July 1997. On my part, I wrote software to automatically create 3D meshes from the 2D images returned by the probe. These meshes were used to create the final 3D world.

PUBLICATIONS

Papers Under Submission

- Andrés Lagar-Cavilla, Niraj Tolia, **Rajesh Krishna Balan**, Eyal de Lara, Mahadev Satyanarayanan, and David O'Hallaron. "Varying Client Thickness for Interactive Scientific Applications". *Submitted to the 1st ACM SIGOPS European Systems Conference (EuroSys)*, April 2006, Leuven, Belgium.
- **Rajesh Krishna Balan**, Maria Ebling, Paul Castro, and Archan Misra. "Matrix: Adaptive Middleware for Distributed Multiplayer Games". *Submitted to IEEE Transactions on Computers*. Extended version of paper that appeared *In Proceedings of the ACM/IFIP/USENIX 6th International Middleware Conference (Middleware)*, Nov 2005, Grenoble, France.

Refereed Conferences/Workshops

- **Rajesh Krishna Balan**, Maria Ebling, Paul Castro, and Archan Misra. "Matrix: Adaptive Middleware for Distributed Multiplayer Games". *Proceedings of the ACM/IFIP/USENIX 6th International*

Middleware Conference (Middleware) (short paper), Nov 2005, Grenoble, France.

- **Rajesh Krishna Balan**, Mahadev Satyanarayanan, SoYoung Park, and Tadashi Okoshi. “Tactics-Based Remote Execution for Mobile Computing”. *Proceedings of the 1st USENIX International Conference on Mobile Systems, Applications, and Services (MobiSys)*, May 2003, San Francisco, California.
- **Rajesh Krishna Balan**, Jason Flinn, Mahadev Satyanarayanan, Shafeeq Sinnamohideen, and Hen-i Yang. “The Case for Cyber Foraging”. *In Proceedings of the 10th ACM SIGOPS European workshop*, September 2001, Saint-Emilion, France.
- **Rajesh Krishna Balan**, B. P. Lee, K. R. R. Kumar, L. Jacob, W. K. G. Seah, and A. L. Ananda. “TCP HACK: TCP Header Checksum Option to Improve Performance Over Lossy Links”. *In Proceedings of 20th IEEE Conference on Computer Communications (INFOCOM)*, April 2001, Anchorage, Alaska.
- B. P. Lee, **Rajesh Krishna Balan**, L. Jacob, W. K. G. Seah, and A. L. Ananda. “TCP Tunnels: Avoiding Congestion Collapse”. *In Proceedings of 25th IEEE Conference on Local Computer Networks (LCN)*, November 2000, Tampa, Florida.
- **Rajesh Krishna Balan**, A. L. Ananda, and H. Y. Wong. “Distributed Computing for the CAVE Virtual Reality Display Environment” *In Proceedings of the Conference on Applications of High Performance Computing*, May 1999, National University of Singapore, Singapore.

Journal Articles

- **Rajesh Krishna Balan**, Aditya Akella, and Srinivasan Seshan. “Multi-Modal Network Protocols”. *Computer Communications Review*, Vol. 32, No. 1, Pg. 60, ACM Press, 2002.
- **Rajesh Krishna Balan**, B. P. Lee, K. R. R. Kumar, L. Jacob, W. K. G. Seah, and A. L. Ananda. “TCP HACK: TCP Header Checksum Option to Improve Performance Over Lossy Links”. *Computer Networks*, Vol. 39, No. 4, Pg. 347–361, Elsevier Publications, 2002. Extended version of paper that appeared *In Proceedings of 20th IEEE Conference on Computer Communications (INFOCOM)*, April 2001, Anchorage, Alaska.
- B. P. Lee, **Rajesh Krishna Balan**, L. Jacob, W. K. G. Seah, and A. L. Ananda. “TCP Tunnels: Avoiding Congestion Collapse”. *Computer Networks*, Vol. 39, No. 2, Pg. 207–219, Elsevier Publications, 2002. Extended version of paper that appeared *In Proceedings of 25th IEEE Conference on Local Computer Networks (LCN)*, November 2000, Tampa, Florida.

Patents

- **Rajesh Krishna Balan**, Paul Castro, Maria R. Ebling, and Archan Misra. “Method And System For A Generalized On-Demand Service Architecture For Interactive Applications” Under Review by the US Patent Office.

Thesis

- **Rajesh Krishna Balan** “Simplifying Cyber Foraging”, *Ph.D. Thesis*, expected Spring 2006, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania.
- **Rajesh Krishna Balan** “CHAMELEON - A System for Adaptive QoS Provisioning” *Masters Thesis*, August 2000, School of Computing, National University of Singapore, Singapore.
- **Rajesh Krishna Balan** “Distributed Computing for the CAVE Virtual Reality Display Environment” *Honours Thesis*, November 1998, School of Computing, National University of Singapore, Singapore.

Technical Reports and Other Articles

- Jesse Cheng, **Rajesh Krishna Balan**, and Mahadev Satyanarayanan. “Exploiting Rich Mobile Environments”. *Technical Report*, CMU-CS-05-199, Carnegie Mellon University, Dec 2005.
- João Pedro Sousa, **Rajesh Krishna Balan**, Vahe Poladian, Mahadev Satyanarayanan, and David Garlan. “Giving Users the Steering Wheel for Guiding Resource-Adaptive Systems”. *Technical Report*, CMU-CS-05-198, Carnegie Mellon University, Dec 2005.
- **Rajesh Krishna Balan**, Darren Gergle, Mahadev Satyanarayanan, and Jim Herbsleb. “Simplifying Cyber Foraging for Mobile Devices”. *Technical Report*, CMU-CS-05-157R, Carnegie Mellon University, Aug 2005.
- **Rajesh Krishna Balan**. “Powerful Change Part 2: Reducing the Power Demands of Mobile Devices”. *IEEE Pervasive Computing*, Vol. 3, No. 2, Pg. 71–73, IEEE Press, 2004.
- **Rajesh Krishna Balan**, João Pedro Sousa, and Mahadev Satyanarayanan. “Meeting the Software Engineering Challenges of Adaptive Mobile Applications”. *Technical Report*, CMU-CS-03-111, Carnegie Mellon University, Feb 2003.
- **Rajesh Krishna Balan**, K. K. Sung, and T. T. Blackmon. “Creation of 3D models for the Mars Pathfinder Project”. In *Proceedings of the 3rd National Undergraduate Research Congress*, September 1997, Nanyang Technological University, Singapore.
- **Rajesh Krishna Balan** and E. W. Chionh. “Linux Internals”. In *Proceedings of the 2nd National Undergraduate Research Congress*, September 1996, National University of Singapore, Singapore.

HONORS

2003 - 2005	IBM Ph.D. Fellowship Recipient for 2003/2004 and 2004/2005 This is a highly prestigious fellowship awarded annually to the top Ph.D. candidates worldwide
April 2002	Agency for Science and Technology Research International Fellowship Recipient This is the Singaporean equivalent of the NSF that awards fellowships to the best Singaporean graduate students
Jan 2002 - Jan 2003	Awarded a \$40,000 personal research grant from USENIX for research on Cyber Foraging

TEACHING EXPERIENCE

- Masters Thesis Research Mentor, Carnegie Mellon University, Fall 2004. Guided a masters student, Jesse Cheng, from the Information Networking Institute for his 6 month long masters thesis. The goal of Jesse's thesis was to develop mechanisms to automatically use extra server resources in the environment. I met with Jesse at least once a week and helped him to understand the issues relevant to this thesis, to identify solutions to the problems posed by this thesis, to understand the Chroma codebase and various test applications, and finally, to implement the solution into the codebase and test applications. Jesse successfully defended his thesis, entitled "Data Decomposition " in April 2005. A paper describing his work is currently under submission at Pervasive 2006.
- Undergraduate Senior Thesis Mentor. Guided a computer science undergraduate student, Maksim Ioffe, for his year long senior thesis. This thesis attempted to create a networking service that would allow mobile users to retain their network connections even when their IP addresses and networking interfaces changed. Maksim defended his thesis, entitled "PluriNet: Dynamic, Environment-Aware Routing", successfully in May 2002.
- Teaching Assistant, Carnegie Mellon University, Fall 2002, *15-213: Introduction to Computer Systems* (lower-level class with weekly recitations). This class was taught by Prof. Dave O'Halloran and Prof. Randy Bryant.
- Teaching Assistant, Carnegie Mellon University, Fall 2001, *15-441: Computer Networks* (senior-level class). This class was taught by Prof. Srinivasan Seshan.
- Teaching Assistant, National University of Singapore, Spring 1999, Fall 1999, Spring 2000, *CS3103 - Computer Networks II* (senior-level class). This class was taught by Prof. A. L. Ananda.

PROFESSIONAL AND COMMUNITY SERVICE

- Served on the Programme Committee for the 2004 System Support for Ubiquitous Computing Workshop (UbiSys). This was a workshop held in conjunction with UbiComp 2004.
- Reviewed submissions to INFOCOMM, SIGCOMM, MobiCom, MobiSys, OSDI, SOSP, UbiComp, Pervasive, IEEE Transactions on Mobile Computing, and the IEEE Pervasive Computing Magazine.
- Served on the Admissions Committee of Carnegie Mellon University's Computer Science Department for two years (2002 and 2003).
- Helped organize the Spring 2002 Open House for prospective Computer Science graduate students.
- Student Coordinator for Fall 2001 Immigration Course for new Computer Science graduate students. This is 3 week programme that helps incoming students both identify prospective advisors and integrate into the department.
- Helped organize the Spring 2001 Open House for prospective Computer Science graduate students.

REFERENCES

Professor Mahadev Satyanarayanan
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
Office: (412) 268-3743 Fax: (412) 268-4136
satya@cs.cmu.edu

Professor David Garlan
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
Office: (412) 268-5056 Fax: (412)-268-3455
garlan@cs.cmu.edu

Professor Srinivasan Seshan
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
Office: (412) 268-8734 Fax: (412) 268-5576
srini@cmu.edu

Doctor Maria R. Ebling
Research Staff Member
IBM T. J. Watson Research Center
P. O. Box 704
Yorktown Heights, NY 10598-0704
Office: (914) 784-7949 Fax: (914) 784-7455
ebling@us.ibm.com

Professor James D. Herbsleb
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
Office: (412) 268-8933 Fax: (412) 268-7287
jd@cs.cmu.edu

Additional references available on request.