

Animesh Nandi

Max Planck Institute for Software Systems (MPI-SWS)
Campus E1 4
66123 Saarbrücken, Germany

Phone: +49 152 0148 6023
animesh@mpi-sws.org
<http://www.mpi-sws.org/~animesh>

RESEARCH INTERESTS

Internet-scale networked and distributed systems, middleware for communication systems, content distribution networks, Web2.0 applications (Internet television, virtual classroom, online content recommendation systems, cloud computing, and online social networks), value-added services in converged networks.

EDUCATION

Ph.D. in Computer Science May 2009
Rice University, Houston, Texas, USA
Thesis: *Exploring the design space of cooperative streaming multicast*
Advisor: Dr. Peter Druschel

M.S. in Computer Science May 2004
Rice University, Houston, Texas, USA
GPA: 4.0 / 4.0
Thesis: *LALA: Location aware load aware overlay anycast*
Advisors: Dr. Peter Druschel and Dr. Ion Stoica

B.Tech. in Computer Science May 2001
Indian Institute of Technology, Kharagpur, India
GPA: 9.18 / 10.0
Thesis: *Designing a network monitoring system*
Advisors: Dr. Arobinda Gupta and Dr. Anupam Basu

RESEARCH PROJECTS

Understanding the design tradeoffs of cooperative streaming multicast: Cooperative end-system multicast (CEM) is emerging as a promising paradigm for video distribution since the bandwidth overhead of disseminating content is shared among the participants themselves. Prior research in CEMs has focused either on the design of new protocols or on comparisons of complete systems. However, the inherent tradeoffs between different CEM designs, particularly under a wide range of operating conditions, are not well understood. In this work [1], we empirically and systematically evaluate CEM design choices. We synthesize our results into a model that captures the inherent tradeoffs in CEMs, which in particular, conjectures that no CEM design can simultaneously achieve low overhead, low lag, and high streaming quality.

SAAR: A shared control plane for cooperative streaming multicast: Common to all cooperative end-system multicast (CEM) systems is the control task of selecting overlay neighbors. We demonstrate that the control task of neighbor acquisition can be factored out to a separate control overlay that provides a single primitive: a configurable anycast for peer selection. Different data overlay organizations (single-tree, multi-tree, mesh, and hybrids) can be supported simply and efficiently by expressing appropriate anycast policies. We designed and implemented SAAR [2], a shared control plane for CEMs. An experimental evaluation demonstrates that SAAR's control efficiency enables fast channel switching and high streaming quality even under high churn rates in an application like IPTV.

A realistic Internet latency generator to support large-scale overlay simulations: Understanding the characteristics of the Internet delay space (i.e. the set of all-pair static propagation delays amongst edge networks) is important for the design and evaluation of large scale distributed systems such as overlay networks. In this work [3], we measure and analyze delays among thousands of edge networks and build a model of the Internet delay space. The model preserves relevant metrics of distributed system design, allows for compact representation, and can be used to synthesize data at a scale where direct measurements and storage are impractical.

Scrivener : Providing incentives in cooperative content distribution networks: Participants of cooperative peer-to-peer applications do not necessarily have an incentive to donate resources to the system if they can get the system's resources for free. We design and implement Scrivener [4,7], a decentralized system that ensures fair sharing of bandwidth among the participating nodes, and provides service to a user proportional to the user's contribution.

Resource discovery in decentralized systems: Anycast is a powerful network service that can be used to efficiently manage and locate resources in decentralized distributed systems. Ideally, an anycast system must be scalable, location aware, and load aware. We motivate, design, and implement LALA, a system that realizes the following anycast functionality: given a client request, our goal is to select the closest anycast server that has enough resources (i.e. it is not overloaded) to satisfy the client's request.

SplitStream : High bandwidth content distribution in a cooperative environment: In tree-based multicast systems, only a small number of interior nodes carry the load of forwarding multicast messages. This poses a problem in cooperative end-system multicast where participants are expected to contribute resources proportional to the benefit they expect in return. SplitStream [5,8] is designed to distribute the forwarding load among all participants and to accommodate nodes with different bandwidth capacities.

FreePastry: I contributed to the FreePastry code, which is an open source implementation of Pastry, a generic, scalable and efficient substrate for peer-to-peer applications (<http://freepastry.org/FreePastry/>). FreePastry applications can be deployed in the general Internet, executed in a wide-area testbed like Planetlab, or executed on top of an emulator. As part of the exploration of the design space of cooperative video streaming [1,2], I built from scratch (approx. 40,000 lines of code) different peer-to-peer video streaming multicast approaches (single-tree, multi-tree, mesh-based and hybrids) on the FreePastry substrate.

**PROFESSIONAL
SERVICE**

Program Committee member of the First IRIS Student Workshop on Peer-to-Peer Systems, Aug 10, 2003 in Cambridge, MA.

External reviewer of research papers in conferences including SOSP, OSDI, NSDI, SIGCOMM, IPTPS, EuroSys and journals including TOCS.

**RESEARCH
EXPERIENCE**

Max Planck Institute for Software Systems (MPI-SWS), Germany: Visiting research student in the Distributed and Networked Systems group, under the guidance of Dr. Peter Druschel who is a pioneer in the area of structured peer-to-peer overlay networks, May 2006 - present.

Rice University, Houston, USA: M.S. and Ph.D. student in the Dept. of Computer Science, Distributed and Networked Systems group, under the guidance of Dr. Peter Druschel who is a pioneer in the area of structured peer-to-peer overlay networks, Aug 2001 - May 2006.

University of California, Berkeley, USA: Summer intern in the Networked Systems group under the guidance of Dr. Ion Stoica, for the following 2 summers: (i) May-August 2003, and (ii) June-August 2004.

**REFEREED
SHORT PAPERS**

[1] **Animesh Nandi**, Bobby Bhattacharjee, Peter Druschel. What a mesh: Understanding the design tradeoffs for streaming multicast. In Proceedings of ACM SIGMETRICS Performance Evaluation Review, special issue on poster session of SIGMETRICS 2009, as a short 2-page paper.

**REFEREED
CONFERENCE
PUBLICATIONS**

[2] **Animesh Nandi**, Aditya Ganjam, Peter Druschel, T.S. Eugene Ng, Ion Stoica, Hui Zhang, Bobby Bhattacharjee. SAAR: A shared control plane for overlay multicast. In Proceedings of 4th USENIX Symposium of Networked Systems Design and Implementation (*NSDI 2007*), Cambridge, MA, USA.

[3] Bo Zhang, T.S. Eugene Ng, **Animesh Nandi**, Rudolf Riedi, Peter Druschel, Guohui Wang. Measurement based analysis, modeling and synthesis of the Internet delay space. In Proceedings of ACM/USENIX Internet Measurement Conference (*IMC 2006*), Rio de Janeiro, Brazil.

[4] **Animesh Nandi**, Tsuen-wan “Johnny” Ngan, Atul Singh, Dan Wallach, Peter Druschel. Scrivener: Providing incentives in cooperative content distribution systems. In Proceedings of USENIX 6th International Middleware Conference (*Middleware 2005*), Grenoble, France.

[5] Miguel Castro, Peter Druschel, A. M. Kermarrec, **Animesh Nandi**, Antony Rowstron, Atul Singh. SplitStream: High bandwidth multicast in cooperative environments. In Proceedings of 19th ACM Symposium on Operating Systems Principles (*SOSP 2003*), Bolton Landing, New York, USA.

[6] **Animesh Nandi**, Arobinda Gupta. Estimating link characteristics of Internet paths with FPathchar. In Proceedings of 4th International Conference on Internet Computing (*IC 2003*), Las Vegas, Nevada.

**REFEREED
WORKSHOP
PUBLICATIONS**

[7] Tsuen-wan “Johnny” Ngan, **Animesh Nandi**, Atul Singh, Dan Wallach, Peter Druschel. On designing incentives-compatible peer-to-peer systems. In Proceedings of 2nd Workshop on Future Directions in Distributed Computing (*Fu-DiCo II: S.O.S 2004*), Bertinoro, Italy.

[8] Miguel Castro, Peter Druschel, A. M. Kermarrec, **Animesh Nandi**, Antony Rowstron, Atul Singh. SplitStream: High bandwidth content distribution in cooperative environments. In Proceedings of 2nd International Workshop on Peer-to-Peer Systems (*IPTPS 2003*), Berkeley, CA, USA.

TEACHING

Teaching assistant in the following courses in graduate school at Rice University:

COMP 413, Distributed Program Construction	Spring 2002
COMP 212, Intermediate Programming	Fall 2002
COMP 421, Operating Systems and Concurrent Programming	Spring 2003
COMP 526, High Performance Computer Architecture	Fall 2003
COMP 280, The Mathematics of Computation	Fall 2004
COMP 620, Graduate Seminar in Distributed Computing	Spring 2005

HONORS AND AWARDS

Rice University Graduate Fellowship, 2001

All India Rank of 139 out of approx. 150,000 applicants in the All India IIT Joint Entrance Examination, 1997.

All India Rank of 25 out of approx. 100,000 applicants in the All India Roorkee Entrance Examination, 1997.

Topped the state level examination conducted by the Association of Mathematics Teachers of India and qualified to participate in the Indian National Mathematics Olympiad (INMO-1995).

Recipient of the National Talent Search Examination (NTSE) scholarship for 4 consecutive years (1994-1997)

TECHNICAL SKILLS

Extensive experience in Unix-based environments, Java, Perl, and C. Ability to pickup familiarity with other languages, scripting tools, and software, as and when required.

STATUS

Citizenship: India
Visa: U.S. F-1 Student Visa, German Residence Permit

REFERENCES

Available upon request.