



OPERATING SYSTEMS REVIEW

A Publication of the
Association for Computing Machinery
Special Interest Group on Operating Systems

Volume 41, Number 5

October 2007

SPECIAL TOPIC
GUEST EDITORS

Gossip-Based Computer Networking
Anne-Marie Kermarrec and Maarten van Steen

Gossip-Based Computer Networking	1
Anne-Marie Kermarrec (INRIA), Maarten van Steen (Vrije Universiteit Amsterdam)	
Gossiping in Distributed Systems	2
Anne-Marie Kermarrec (INRIA), Maarten van Steen (Vrije Universiteit Amsterdam)	
The Promise, and Limitations, of Gossip Protocols	8
Ken Birman (Cornell University)	
How Robust Are Gossip-Based Communication Protocols?	14
Lorenzo Alvisi (University of Texas at Austin), Jeroen Doumen (University of Twente), Rachid Guerraoui (Ecole Polytechnique Fédérale de Lausanne), Boris Koldehofe (University of Stuttgart), Harry Li (University of Texas at Austin), Robbert van Renesse (Cornell University), Gilles Tredan (IRISA/INRIA)	
A Generic Theoretical Framework for Modeling Gossip-Based Algorithms	19
Yaacov Fernandess (Hebrew University of Jerusalem), Antonio Fernández (Universidad Rey Juan Carlos), Maxime Monod (Ecole Polytechnique Fédérale de Lausanne)	
Formal Analysis Techniques for Gossiping Protocols	28
Rena Bakhshi (Vrije Universiteit Amsterdam), Francois Bonnet (ENS Cachan/IRISA), Wan Fokkink (Vrije Universiteit Amsterdam), Boudewijn Haverkort (University of Twente)	
The “Art” of Programming Gossip-based Systems	37
Patrick Eugster (Purdue University), Pascal Felber (University of Neuchâtel), Fabrice Le Fessant (INRIA)	
Compositional Gossip: A Conceptual Architecture for Designing Gossip-Based Applications	43
Étienne Rivière (IRISA), Roberto Baldoni (Roma University), Harry Li (University of Texas at Austin), José Pereira (University of Minho)	
Exploring the Interdisciplinary Connections of Gossip-Based Systems	51
Paolo Costa (Vrije Universiteit Amsterdam), Vincent Gramoli (INRIA), Márk Jelasity (University of Szeged), Gian Paolo Jesi (University of Bologna), Erwan Le Merrer (FTR&D/IRISA), Alberto Montresor (University of Trento), Leonardo Querzoni (University of Rome)	
Exploiting the Synergy Between Gossiping and Structured Overlays	61
Ali Ghodsi (Swedish Institute of Computer Science), Seif Haridi (KTH), Hakim Weatherspoon (Cornell University)	
Gossiping on MANETs: The Beauty and the Beast	67
Roy Friedman (Technion – Israel Institute of Technology), Daniela Gavidia (Vrije Universiteit Amsterdam), Luis Rodrigues (Universidade de Lisboa), Aline Carneiro Viana (INRIA), Spyros Voulgaris (ETH Zurich)	
Gossiping Over Storage Systems Is Practical	75
Hakim Weatherspoon (Cornell University), Hugo Miranda (Universidade de Lisboa), Konrad Iwanicki (Vrije Universiteit Amsterdam), Ali Ghodsi (Swedish Institute of Computer Science), Yann Busnel (IRISA)	
Reliable On-Demand Management Operations for Large-scale Distributed Applications	82
Jin Liang, Indranil Gupta, Klara Nahrstedt (University of Illinois at Urbana-Champaign)	

OPERATING SYSTEMS REVIEW

A Publication of the ACM Special Interest Group on Operating Systems

Chair

Keith Marzullo
Computer Science & Engineering
U.C. San Diego
La Jolla, CA 92093-0114
USA
+1-858-534-3729
marzullo@cs.ucsd.edu

Information Director

Geoffrey M. Voelker
Computer Science & Engineering
U. C. San Diego
9500 Gilman Dr., MC 0114
La Jolla, CA 92093-0114 USA
+1-858-822-3323
voelker@cs.ucsd.edu

Vice Chair

Gilles Muller
École des Mines de Nantes
4, rue Alfred Kastler – BP 20722
44307 Nantes Cedex 3
FRANCE
+33-2-51-85-82-04
gilles.muller@emn.fr

Individual Submissions Editors

John Chandy, Univ. of Connecticut (Coordinator)
Antônio Fröhlich, Federal Univ. of Santa Catarina
Ashvin Goel, Univ. of Toronto

Secretary-Treasurer

Editor of Operating System Review

Jeanna Matthews
Department of Computer Science
Clarkson University
8 Clarkson Avenue, MS 5815
Potsdam, NY 13699
USA
+1-315-268-6288
jnm@clarkson.edu

OPERATING SYSTEMS REVIEW is a publication of the ACM Special Interest Group on Operating Systems (SIGOPS), whose scope of interest includes: Computer operating systems and architecture for multiprogramming, multiprocessing, and time sharing; resource management; evaluation and simulation; reliability, integrity, and security of data; communications among computing processors; and computer system modeling analysis.

Membership in SIGOPS (at \$20 per annum) is open to ACM members and associate members, and student members (at \$8 per annum). Non-members of ACM may subscribe to OPERATING SYSTEMS REVIEW at \$30 per year. All SIGOPS members receive OPERATING SYSTEMS REVIEW. SIGOPS membership application forms are available from ACM Headquarters, 2 Penn Plaza, Suite 701, New York, NY 10121-0701, telephone +1-212-869-7440. Changes of address and other matters pertaining to the SIGOPS mailing list should be directed to ACM Headquarters, not to any of the SIGOPS officers.

Contributions to OSR Several types of contributions to OSR are solicited from the operating systems community. First, many issues of OSR will be organized around a special topic or theme. Guest editors are sought to solicit, review, select and coordinate articles related to the special topic. If you are interested in suggesting a special topic or in serving as a guest editor, contact the OSR editor. Second, contributions related to upcoming special topics are requested. Such contributions should be sent to the guest editor for the issue of interest. Calls for participation in upcoming special topics issues will be posted on the OSR web page (<http://www.acm.org/sigops/osr.html>) and announced via the SIGOPS mailing list. Finally, individual submissions not related to specific special topics are welcome. Proposals for individual submissions should be sent to the OSR editor for review and to discuss possible publication dates. Each proposal should clearly say how the submission would be of interest to the SIGOPS community. Traditionally, the SOSP proceedings will be mailed as the December issue in odd-numbered calendar years, and the ASPLOS proceedings will be mailed in even-numbered calendar years.

OPERATING SYSTEMS REVIEW (ISSN 0163-5980) is published five times a year (January, April, July, October, and December) by the Association for Computing Machinery, 2 Penn Plaza, Suite 701, New York, NY 10121-0701. Periodicals postage paid at New York, NY 10001, and at additional mailing offices. POSTMASTER: Send address changes to OPERATING SYSTEMS REVIEW, ACM, 2 Penn Plaza, Suite 701, New York, NY 10121-0701.

Gossip-Based Computer Networking

Anne-Marie Kermarrec
INRIA, Rennes, France
Anne-Marie.Kermarrec@irisa.fr

Maarten van Steen
Vrije Universiteit Amsterdam, Netherlands
steen@cs.vu.nl

In the past decade, the experienced shift in scale of networked systems has forced the research community to revisit many of its assumptions and solutions. In this context, we observed an increased use of gossip protocols in large-scale distributed systems.

This special issue consists of papers resulting from a workshop entitled “Gossip-Based Computer Networking.” The workshop was organized in December 2006 at the Lorentz Center in Leiden, The Netherlands. The main goal was to get an overview of the state-of-the-art in gossiping networks, and to identify the important topics that we need to address in the near future.

The format of the workshop was somewhat unusual in the sense that there were relatively few presentations, but many discussion slots. Roughly, each day started with a plenary presentation in which people were asked to be somewhat provocative. The total group of approximately 50 participants were divided into smaller groups, each concentrating on a specific theme, and which were asked to prepare a position paper. This led to discussions during the afternoon sessions, with plenary feedback at the end of the day.

The results of those discussions have been laid down in the position papers that form this special issue. Each paper has been anonymously reviewed by other participants (and interestingly enough, the reviewers were not particularly kind to each other), often leading to substantial improvements. What you now find is therefore a reflection of lots of in-depth discussions and hard work to clarify and substantiate points of views regarding gossip-based solutions.

By-and-large, the workshop itself was experienced as a great success. Although it did take quite some effort from everyone to finalize the discussions in the form of reviewed position papers, this special issue contains an interesting set of papers reflecting current and future trends in this specific area.