Stevenson, Washington, USA
October 14–17, 2007

Association for
Computing Machinery

Advancing Computing as a Science & Profession

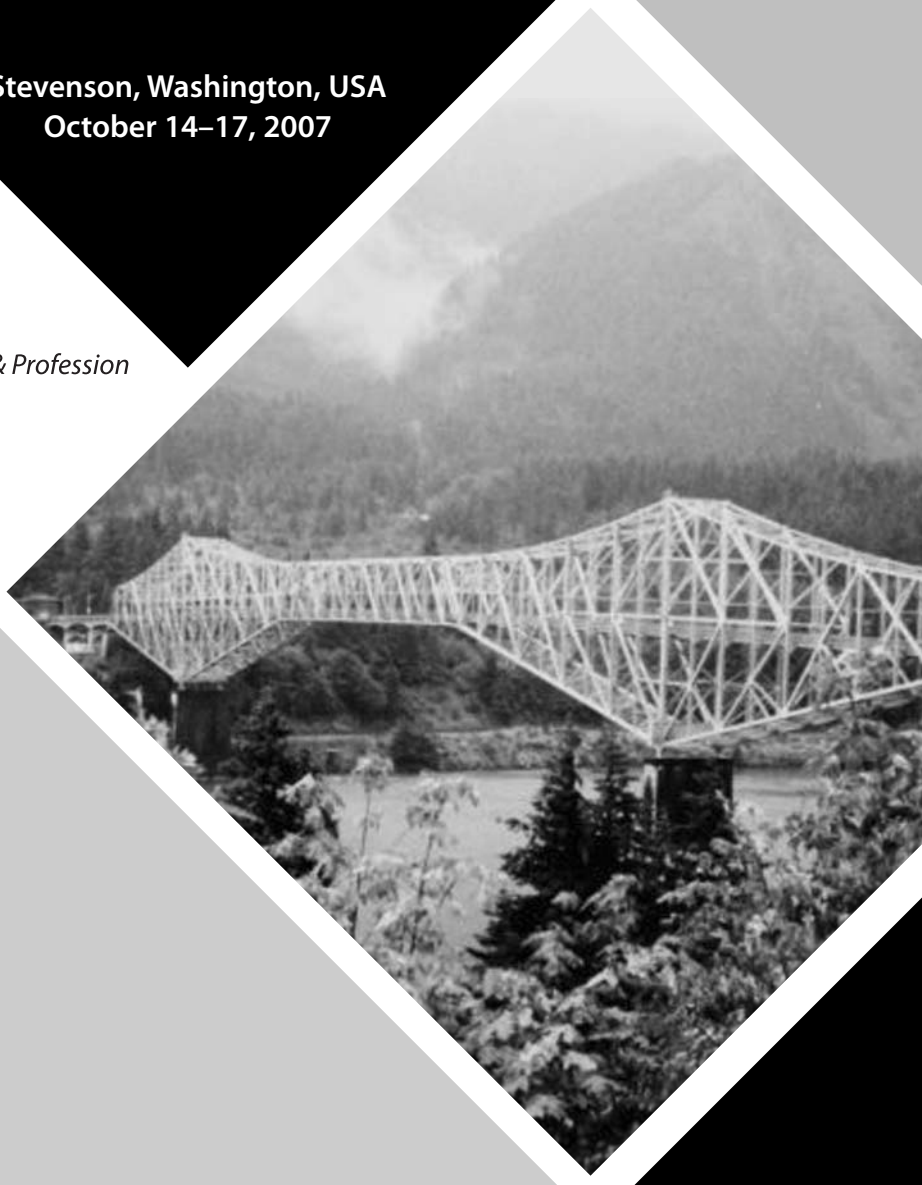

SOSP'07

Proceedings of the 21st ACM Symposium on
Operating Systems Principles

Sponsored by:

ACM SIGOPS

Association for Computing Machinery

Advancing Computing as a Science & Profession

The Association for Computing Machinery
2 Penn Plaza, Suite 701
New York, New York 10121-0701

ISBN: 978-1-59593-591-5

Additional copies may be ordered prepaid from:

ACM Order Department
PO Box 11405
New York, NY 10286-1405

Phone: 1-800-342-6626
(US and Canada)
+1-212-626-0500
(all other countries)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org

ACM Order Number 83300706

Printed in the USA

Foreword

Please enjoy the proceedings of the 21st ACM Symposium on Operating Systems Principles–SOSP'07. In the SOSP tradition, the 25 papers herein explore a wide range of computer systems topics, including traditional ones such as concurrency as well as new ones such as "hardening" Web browsers. Collectively these papers report on some of the most creative and thought-provoking ideas in computer systems today and how they work out in practice. The 25 papers were shepherded by PC members to ensure that they are easy to read. We hope you will enjoy learning from these papers.

Selecting 25 papers out of 131 submissions was difficult because so many of the submissions were of high quality. To make the selection process as fair and as consistent as possible the program committee employed a different process than used by previous SOSPs (but used successfully by other conferences such as SIGCOMM). The program committee consisted of 13 "heavy"-load and 13 "light"-load members. The heavy-load members reviewed about 34 submissions each and attended the face-to-face PC meeting in Cambridge, MA USA. The light-load members reviewed about 24 papers each and did not attend the PC meeting. In contrast, recent SOSPs used a small number of PC members (12-15) who read a large fraction of all submissions, sometimes assisted by external reviewers. SOSPs before that required all PC members to read all submissions.

The goal of the new process was to resolve the tension between having high-quality, consistent reviews, a large number of submissions (it has been steadily growing over the years), and a productive face-to-face meeting. With more PC members the PC did not have to rely on external reviews, which can be inconsistent because the external reviewers see only a small sample of the submissions, yet the workload for the individual PC members was manageable, allowing thorough reviewing. By having a subset of the PC members meet in person, the PC was able to have in-depth discussion and reach consensus through discussion (rather than voting). The larger overall PC also allowed a broader group of people to participate in the decisions.

Paper selection was a three round process, with multiple reviews by the PC generated in each round and with reviewers targeted by subject expertise. The first two rounds reduced the pool of considered papers by 50%. The 62 remaining papers produced another two reviews apiece and all 705 reviews were assessed in preparation for the PC meeting. At the PC meeting, the 62 papers were ranked by review scores for discussion order and each assigned a champion to summarize content and strengths and to lead the discussion on individual papers. The PC discussion for each paper followed until consensus was reached. Throughout the process anonymity was maintained and conflicts of interest precluded by removing authors or those with direct association with an author from the discussion. In the final selection, 3 papers were co-authored by heavy-load PC members, and 6 were co-authored by light-load PC members.

Did the PC make good decisions? This question is probably best answered by you after reading the papers! It is interesting to note, however, that a shadow PC chaired and organized by Rebecca Isaacs (Microsoft Research, Cambridge, England) reviewed 101 of the 131 submissions (which included 18 of the 25 papers accepted by the real PC) and accepted 16 papers. Of the 18 papers accepted by the real PC, 9 were accepted by the shadow PC, 4 were discussed by the shadow PC, and 5 didn't make it to the discussion at the shadow PC meeting (the shadow PC discussed 40 submissions). An informal review suggests that the variations in decisions were partially due to the fact that the shadow PC's goals were different from the real PC's. The shadow PC's main goal was to educate participants about how a PC works, how to review papers, etc. and members volunteered to participate; the real PC members were carefully chosen to provide both depth and breadth across a

wide range of topics. This difference in focus resulted in a few important modifications to the decision process: the shadow PC members produced 4 reviews per submission and saw fewer submissions, had less time to absorb the reviews before the meeting, and had less expertise in certain areas. A full report will be submitted to SIGOPS Operating Systems Review.

A successful conference goes beyond the accepted papers, building and supporting its community. At SOSP this year, and in celebration of the 20th anniversary of the SYSTERS group, we have introduced two special programs. First, we recognize the importance of increasing the participation of women and underrepresented minorities in systems research. And to be successful, this participation has to reach to undergraduates and show them the excitement and interesting problems in systems. Toward this end, we established an additional scholarship opportunity, supported by industry contributors, that has supported these targeted groups to attend SOSP. Second, the participation by women works best when undergraduates and graduates are shown the way by women already participating in the field. For this, we created a special one-day workshop for women to develop this community as a prelude to the beginning of the SOSP conference. Support from our industry contributors and from NSF and CRA-W has been outstanding. Equally impressive has been the support from the organizing team and all who have made this possible. These initiatives have been embraced enthusiastically.

SOSP is a great conference mostly because it attracts so many high-quality submissions, and we would like to thank all the authors who submitted. We would also like to thank the PC members for the tremendous amount of work they did: reviewing the submissions, providing constructive feedback, and shepherding the accepted submissions. Organizing a conference is a team effort, and we would like to thank the team: Jon Walpole (Local arrangements), Jay Lorch (Industry contributions), Jason Flinn (Publicity), Robbert van Renesse and Hakim Weatherspoon (Scholarships), Mike Kozuch (Registration), Eddie Kohler (HotCRP), Mema Roussopoulos (Posters), David Mazières (WIPs), and Carla Ellis, Sharon Perl, and Barbara Liskov (Women's workshop). All of these people have shown great dedication and have worked very hard to make this conference a success. Finally we would like to thank the sponsors—without their financial support the conference could not have happened.

We hope that you will find the SOSP program interesting and that the symposium will provide you with a valuable opportunity to share ideas with other researchers and practitioners from institutions around the world.

Tom Bressoud
General Chair
Denison University

Frans Kaashoek
Program Chair
MIT

Table of Contents

21st ACM Symposium on Operating Systems Principles Organization

General Chair & Treasurer: Thomas C. Bressoud *(Denison University, USA)*

Program Chair: M. Frans Kaashoek *(Massachusetts Institute of Technology, USA)*

Financial Support Chair: Jacob R. Lorch *(Microsoft Research, USA)*

Local Arrangements Chair: Jonathan Walpole *(Portland State University, USA)*

Registration Chair: Michael Kozuch *(Intel Research, USA)*

Publicity Chair: Jason Flinn *(University of Michigan, USA)*

SOSP Women's Workshop: Carla Ellis *(Duke University, USA)*
Sharon Perl *(Google, USA)*
Barbara Liskov *(Massachusetts Institute of Technology, USA)*

Student Scholarship Committee: Robbert van Renesse *(Cornell University, USA)*
Hakim Weatherspoon *(Cornell University, USA)*

Poster Session Chair: Mema Roussopoulos *(Harvard University, USA)*

Work-in-progress Session Chair: David Mazières *(Stanford University, USA)*

HotCRP Czar: Eddie Kohler *(University of California Los Angles, USA)*

Shadow PC Chair: Rebecca Isaacs *(Microsoft Research, England)*

Hall of Fame Award Committee: Ken Birman *(Cornell University, USA)*
Peter Chen *(University of Michigan, USA)*
Peter Druschel *(MPI-SWS, Germany)*
Mike Jones *(Microsoft Research, USA)*
Frans Kaashoek *(Massachusetts Institute of Technology, USA)*
Jeff Mogul *(Chair, HP Labs, USA)*
Larry Peterson *(Princeton University, USA)*
John Wilkes *(HP Labs, USA)*

Program Committee:

Andrea Arpaci-Dusseau *(University of Wisconsin, USA)*
Paul Barham *(Microsoft Research, England)*
Peter Chen *(University of Michigan, USA)*
Jeff Dean *(Google, USA)*
Richard Draves *(Microsoft Research, USA)*
Peter Druschel *(Max Planck Institute for Software Systems, Germany)*
Carla Ellis *(Duke University, USA)*
Hermann Härtig *(Dresden University of Technology, Germany)*
M. Frans Kaashoek *(Massachusetts Institute of Technology, USA)*
Orran Krieger *(IBM, USA)*
Hank Levy *(University of Washington, USA)*
David Mazières *(Stanford University, USA)*
Andrew Myers *(Cornell University, USA)*
Jeff Mogul *(HP Labs, USA)*
Vivek Pai *(Princeton University, USA)*
Adrian Perrig *(Carnegie Mellon University, USA)*
Sharon Perl *(Google, USA)*
Sylvia Ratnasamy *(Intel Research, USA)*
Mendel Rosenblum *(Stanford University, USA)*
Mema Roussopoulos *(Harvard University, USA)*
Stefan Savage *(University of California San Diego, USA)*
Michael Schroeder *(Microsoft Research, USA)*
Ion Stoica *(University of California Berkeley, USA)*
Yuanyuan Zhou *(University of Illinois, USA)*
WeiMin Zheng *(Tsinghua University, China)*

Additional Reviewers:

Krste Asanovic
Hari Balakrishnan
Cullen Bash
Lujo Bauer
Micah Brodsky
David Brumley
Tanya Bragin
Mike Burrows
Russ Cox
Rodrigo Fonseca
Bryan Ford
Armando Fox
Greg Ganger
Steven Gribble
Andreas Haeberlen
Max Krohn
Chris Lesniewski-Laas
Dave Levin
Z. Morley Mao
Petros Maniatis
Jon McCune
Alan Mislove
Robert Morris
Robert O'Callahan
Bryan Parno
Ansley Post
Charles Reis
Mike Reiter
Partha Ranganathan
Andrey Rybalchenko
Arvind Seshadri
Atul Singh
Emil Sit
Dawn Song
Jacob Strauss
Jeremey Stribling
K. Vikram
Benjamin Wester
Alex Yip
Lantian Zheng
Xin Zheng

Sponsor:

Supporters: