



Computer Science Publications

Jennifer Rexford

Princeton University

<http://www.cs.princeton.edu/~jrex>

Where I'm Coming From



- Chair of ACM SIGCOMM (2004-2007)
 - Explosive growth in the number of papers
 - Wider interest in data networks due to the Internet
 - Main conference with a reputation of cliquishness
- Issues SIGCOMM has been grappling with
 - Low acceptance rates at conferences
 - Encouraging publication of new and risky ideas
 - Improving the number and quality of reviews
 - Full-length versions of workshop papers
- Other issues that arise often
 - Non-anonymous paper reviews
 - Recycling of paper submissions



Providing New Kinds of Venues

- New and risky ideas
 - Hot topics in networking
- Active sub-areas in networking
 - Internet measurement, network software systems, network hardware architecture, sensor networks, ...
- Seeding new sub-areas
 - Workshops co-located with the main conference
 - Inter-disciplinary topics, and emerging research areas
- Regional networking workshops
 - Latin America
 - Asia



Changing the Main Conference

- Co-located workshops (2003--)
 - Three-day conference and two days of workshops
 - Broadens participation and builds community
- Accepting more papers (2005)
 - From high 20s to high 30s, while still single track
- New geographic rotation (2007--)
 - Old: twice North America, once Europe
 - New: North America, Europe, and wildcard
- Experimented with position papers (2003-2004)
 - Had them just for a couple of years
 - But, it proved to be a bit of a mixed bag



Prior Work at Workshops

- Workshop paper grows up into conference paper
 - Extending the work in a meaningful way
- How do you “count” the prior workshop paper?
 - If it counts, you limit the quality/excitement of workshops
 - If it doesn’t count, conferences may become repetitive
- We’re converging on a policy
 - If the conference paper makes a significant delta (e.g., adding implementation, evaluation, analysis, etc.)...
 - ... ignore workshop paper in judging the submission
- But, we can’t control what other conferences do...



Managing the Reviewer Load

- Early decision on some papers
 - E.g., “quick reject” with 1-2 reviews
- Two-tiered program committee (SIGCOMM’06)
 - “Lite” PC that reviews paper but doesn’t attend meeting
 - “Heavy” PC that reads most of the serious contenders
- External reviewers
 - Though, these are hard to manage, and calibrate
- Helping train future reviewers
 - Shadow program committees reviewing the papers
 - E.g., EU shadow PC, and university departments



Public Reviews

- “NY Times Book Review” of each accepted paper
 - Signed review, sometimes by a PC member
 - Different from paper-selection reviews
 - Provides context, and discusses pros/cons
- Started at HotNets’04
 - Lots of discussion and enthusiasm
 - Assigned as reading in some graduate courses
 - Now done in our SIG newsletter, and SIGCOMM’05
- Challenges and issues
 - Should authors be able to publish a rebuttal?
 - Risk of retaliation, brown-nosing, and conflict of interest
 - Why not publish the reviews of the submission?



Challenges

- **Cliquishness**
 - Limits the vibrancy, creativity, and breadth of a discipline
 - Need ways to have turn over, & convey value structure
- **Low acceptance rates**
 - Increases the noise in the process, and the role of bias
 - Need ways to publish risky ideas, & discourage recycling
- **Review quality**
 - Very heavy load of papers to review
 - Need to reduce load, spread load, provide incentives, ...
- **Prevalence of workshops**
 - Adds another stage to the publication pipeline
 - Need ways to address overlaps in paper contents