

Conferences – Bad in theory?

Jeff Mogul
HP Labs, Palo Alto

Conferences vs. Journals: The wrong question

- There are good conferences and bad conferences
 - There are also good journals and bad ones
 - Don't generalize from "some conferences are bad" to "conferences are inherently bad"
 - *Blanket* statements of the form "conferences do not do careful reviewing" are unsupported by evidence
- Setting up the debate as "conferences vs. journals" misses the point:
 - Both are necessary vehicles
 - Both are broken in their own ways
 - Perfection in peer-reviewing is a chimera

Who likes what & why?

- Theory people seem to like journals
 - Multiple review rounds with author feedback
 - Unhurried, in-depth reviews are easier
- Systems people seem to like conferences
 - Timeliness and exploration of alternatives
 - Highly selective conferences have impacts as high as good journals (Chen & Konstan, CACM 2010)
 - Good conferences take “wrong” papers
 - Deadlines can be useful to spur progress

Top 5 signs of a bad conference

Top 5 signs of a bad conference

1. They didn't accept *my* paper

Top 5 signs of a bad conference

1. They didn't accept *my* paper
2. More people on PC than at conference

Top 5 signs of a bad conference

1. They didn't accept *my* paper
2. More people on PC than at conference
3. Has a history of accepting papers generated by Markov processes

Top 5 signs of a bad conference

1. They didn't accept *my* paper
2. More people on PC than at conference
3. Has a history of accepting papers generated by Markov processes
4. Everybody but the speaker and session chair are on Waikiki beach

Top 5 signs of a bad conference

1. They didn't accept *my* paper
2. More people on PC than at conference
3. Has a history of accepting papers generated by Markov processes
4. Everybody but the speaker and session chair are on Waikiki beach
5. Submission deadline extended – 4 times

What makes a conference good?

It's a social process:

- Carefully chosen PC
 - Famous people aren't always the best
 - Use peer pressure to promote good behavior
- F2F PC meeting with lots of discussion
 - NEVER make PC-meeting decisions based on the scores
- PC chairs who aren't afraid of being wrong
 - Push PC to take interesting rather than safe papers
 - I *want* to hear: "Why did you accept that crap paper?"
- Shepherding of every accepted paper
- Strong community support
 - Only a few conferences can be "must-attend" – 500+ people
 - Get enough student travel funding

What can the community do?

- Avoid creating too many conferences
 - Dilutes paper quality and adds PC workload
 - Use workshops for narrowly focused topics, and be honest about the difference
 - Do we need a way to “accredit” the good conferences?
- Figure out a “CS arXiv” model
 - Perhaps this means giving up double-blind reviewing for conferences?
- Make “remote attendance” a *really* good option
 - Are we the cobbler’s children for telepresence?
- Teach grad students how to write really well
 - Well-organized papers/paragraphs enable good reviewing
- Accept that theory and systems people are different

Practices from the best systems conferences

- Multiple review rounds (sometimes 3)
- Detailed reviews (“journal quality”)
- Lots of online discussion before PC meeting – avoid score-focused decisions
- Pre-selected discussion leads
- Sometimes: author rebuttals accepted
- Relatively long papers (12-14 pages)
- Mandatory shepherding

Ideas for improving conference reviewing IT

- Global registry of CS papers
 - Avoid re-reviewing the same paper 10 times
 - Avoid simultaneous-submission problems
 - TurnItIn.com for CS papers?
- Global registry of reviewers
 - Recognize people who are doing the work
 - Help PC chairs choose good reviewers
 - But not a “black list” – legal/ethical issues here