Forward

—

Welcome to the 5th Workshop on Hot Topics in Computer Networks!

We received 114 submissions. The ten members of the program committee submitted 421 reviews, and external reviewers submitted an additional six reviews. HotNets papers are shorter than a normal conference submission, and were generally enjoyable reading, but the load on the program committee was still non-trivial, and they all deserve our thanks. Additional thanks are due to our external reviewers: Kevin Fall (Intel Research Berkeley), Philip Levis (Stanford), Nikitas Liogkas (UCLA), Richard Mortier (Microsoft Research), Suman Nath (Microsoft Research), and Scott Shenker (UC Berkeley and ICSI).

Of course most thanks are due to the authors of submitted papers, both accepted and rejected. Thank you for sharing your work with us.

We would also like to thank our general chair, Xiaowei Yang, for her work in organizing the venue, proceedings, student travel grants, and so forth for the workshop.

We hope you find the papers in this workshop stimulating, and we expect two days of interesting conversation.

Paper analysis Now to the data. Figure 1 lists a number of discriminators for HotNets papers, ranging from review scores (1 was low and 5 was high) to selected topics and assigned reviewers. For each discriminator, we list how many papers fit that discriminator, and what percentage of those papers were accepted. Future authors may want to scrutinize this list to increase their chances at the next HotNets. Consider writing a theoretical paper that claims to be about no other topic (neither systems, routing, transport, security, robustness, nor link-level or wireless issues), and make sure it interests Cristian Estan. And above all, keep it short.

The program committee agreed more on accepted papers than rejected ones, and more on novelty and discussability than overall merit. Define the *spread* as the difference between a paper's highest score and its lowest score in a given score category. (There were four categories: overall merit, technical merit, novelty, and discussability.) Then the set of rejected papers had in every category higher average spreads than the set of accepted papers. This was dramatically true for novelty and discussability, where the average spread for rejected papers was roughly 60% higher than for accepted papers. Did we generally prefer lack of controversy (unfortunate for HotNets)? Were reviewers with outlier scores more successful at convincing the program committee to reject a paper than they were at convincing the committee to accept one? Or are good papers obviously good to most everybody?

Discriminator	**# Papers**	**Accepted**
Low "overall merit" ≥ 4	5	100.0%
High "novelty" $= 5$	4	100.0%
High "overall merit" $= 5$	9	66.7%
PDF submission < 100000 bytes	24	41.7%
Topic "None of the others"	5	40.0%
Topic "Theory"	6	33.3%
Reviewed by Cristian Estan	34	32.4%
Topic "Security & robustness"	27	29.6%
Mainly Computer Modern fonts	7	28.6%
Reviewed by Vern Paxson	34	26.5%
Reviewed by Sylvia Ratnasamy	34	26.5%
Non-anonymous submission	85	24.7%
Topic "Controversy"	21	23.8%
Reviewed by Frank Dabek	34	20.6%
Reviewed by Greg Minshall	113	20.4%
Any paper	*114*	*20.2%*
Reviewed by Eddie Kohler	35	20.0%
Reviewed by Ratul Mahajan	36	19.4%
Reviewed by Anja Feldmann	34	17.6%
Topic "Systems"	46	17.4%
Topic "Applications"	21	14.3%
PDF submission ≥ 500000 bytes	14	14.3%
At least one TrueType font	31	12.9%
Reviewed by Vivek Pai	33	12.1%
Reviewed by Rebecca Isaacs	34	11.8%
Topic "Routing & transport"	37	8.1%
Anonymous submission	29	6.9%
Topic "Link & wireless"	26	3.8%
High "overall merit" ≤ 3	54	0.0%
High "discussability" ≤ 2	23	0.0%

Figure 1: *How to get your paper into HotNets V, or not.*

We were more willing to give low scores for overall merit than for technical merit or novelty. For several papers all reviewers agreed to reject, but no paper had all reviewers agreeing on "no technical merit".

Figures 2 and 3 rank the papers in increasing order by overall merit, technical merit, and novelty, then plot novelty rank vs. overall merit rank (Figure 2) and novelty rank vs. technical merit rank (Figure 3). Novelty

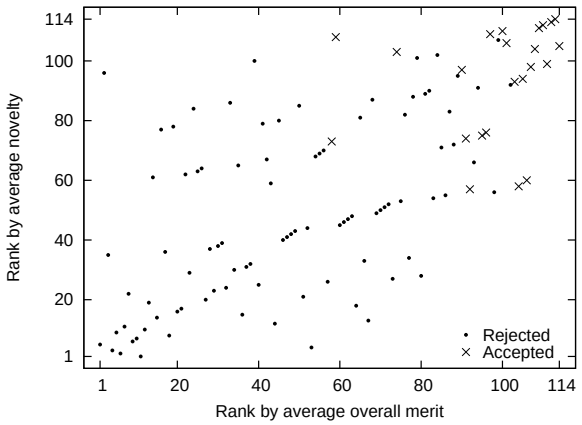

Figure 2: *Novelty vs. overall merit. We largely accepted papers that had either high overall merit or high novelty.*

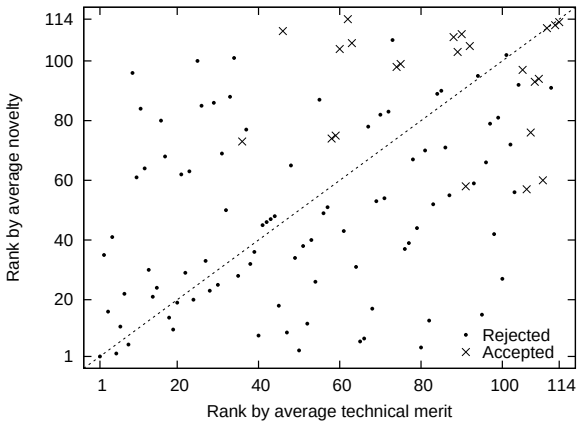

Figure 3: *Novelty vs. technical merit. We preferred "exciting, but flawed" papers, with higher novelty than technical merit (13 acceptances above the line), to "boring, but correct" papers, with higher technical merit than novelty (7 below it).*

rank is a better predictor of acceptance than technical merit rank, as it should be here. Rank plots of both novelty and technical merit vs. overall merit cluster near the diagonal, implying, as one would hope, that overall merit is correlated with both factors. Figure 3 shows less diagonal clustering, particularly at the high end. It's hard to be both new and right.

Anonymity HotNets V authors could choose whether or not to submit their papers anonymously, and just over one-fourth of authors chose to do so—fewer than we expected. Anonymous submissions were much less likely to be accepted than non-anonymous submissions. Of the discriminators in Figure 1 that are under the author's control, only the topic "Link & wireless" is correlated with a lower acceptance rate. This is also reflected in the overall merit scores, where anonymous submissions did substantially worse. Rejected anonymous submissions came from top worldwide industrial research labs and top US universities, among other places. Were our referees less likely to give anonymous submissions the benefit of a doubt? Or, perhaps, were submitters less likely to put their names on a less-than-top-quality submission?

HotNets V *reviewers* could also choose whether or not to submit their *reviews* anonymously. Both co-chairs planned to submit their reviews non-blind, allowing authors to see who wrote their reviews, and both co-chairs still believe philosophically in open reviews. In the end, however, only two people—Greg Minshall and one of our external reviewers—chose to submit their reviews non-blind. The other co-chair, Eddie Kohler, had wanted to use open review as an enforcement mechanism to ensure his reviews were careful, very high quality, and avoided unnecessary provocation, but wasn't sure they all met that standard in the end. If anything, this supports the argument for open reviews, but when even workshops like HotNets receive 114 submissions it's hard to know where all reviewers would get the time to complete their reviews that carefully.

Onward And with that, enjoy these papers: 2006's glimpse into the future of networking research!

—Eddie Kohler and Greg Minshall