

Open Peer-Review Experiment in the Decentralized Coordination Workshop

MARIUS SILAGHI, SONG QIN
FLORIDA TECH
MARIUS.SILAGHI@FIT.EDU

WILLIAM CHEUNG
HONG KONG BAPTIST UNIVERSITY
WILLIAM@COMP.HKBU.EDU.HK

Given known game-theoretical analyses of peer-reviewing for conferences, the plan is to experiment with organizing an event based on techniques that in simulations yield the best incentives for high quality reviewing.

I. GAME THEORY FOR PEER-REVIEWING

Game theory is a powerful tool that can be used to model, analyze, and compare complex social interactions. If one can formalize and quantify motivations in terms of concrete utility values, then one can predict the behavior of the society under the assumption of rational participants. Moreover, one can simulate the impact of new regulations.

Peer-reviewing is an important social activity, whose quality directly impacts the advancement of science and economy. The traditional expectation for reviewers is to be altruistic, just as for politicians. It is nevertheless educational to see what would happen in case they behave like rational players. In general, exact utilities of human players are hard to quantify; in peer-reviewing as in auctions, negotiations and war. Nevertheless, one knows that often reviewers' revenues come from promotions and funding that depend on quantifiable metrics based on material facts such as the number of citations that they get, the number of articles that they publish, the number of reviewing boards on which they are invited, etc.

It is possible to use the aforementioned metrics to create approximate models linking the behavior of a reviewer in the peer-reviewing process to the net impact on the funding (utility) that he will get. Let us take an example where we assume funding offered to a researcher is related to the total

count of her publications, and conferences limit the number of accepted papers based on a fixed threshold. A conference organized by a community of n researchers has these researchers simultaneously submitting papers and reviewing submissions of their peers. The researchers are considered at the same level of expertise, and the papers are considered equally worthy. In this example it is assumed that a single blinded review is written for each paper and that the review can take two values: $\{low, high\}$. Each researcher submits one paper and reviews one paper. The conference only accepts a fraction $\frac{1}{k}$ of the n submitted papers (i.e. accepts $m = \frac{n}{k}$ submissions). The revenue of an author for publishing a paper is 1. The expected gain from rejecting a paper is $\frac{m}{n} - \frac{m}{n-1} \approx \frac{1}{kn}$. The pair-wise payoff matrix for researchers A and B blindly reviewing each-other's papers is given in Fig. 1. It reveals an equilibrium consisting in scoring each other's article *low*.

	high A's	low A's
high B's	0,0	0, $\frac{1}{kn}$
low B's	$\frac{1}{kn}, 0$	$\frac{1}{kn}, \frac{1}{kn}$

Fig. 1. Pair-wise payoff matrix in reviewer-author game with n researchers in conferences with a threshold on the number of accepted papers and funding based on the count of published papers.

Other models exist for funding based on citation influence, where reviewers have strategies to increase current and future citations (see [Peterson et.al.]).

While errors induced by approximations in such models should be subject to further investigation, current results suggest that certain versions of open peer-review schemes have better equilibria for truthful reviewing than common blind review procedures. The mentioned open peer-review schemes, further explained below, are scheduled to be experimented within a workshop on decentralized co-

ordination planned for Spring 2013.

II. OPEN PEER-REVIEW IN WORKSHOPS

With open peer-review, the reviews and the identity of the reviewers are published along with the endorsed and rejected submissions, as an incentive for improving their quality. The publication effectively creates a new link between reviews and the utility of the reviewer, since reviews can be cited, making it possible to create incentives for truthful reviewing. Even before a game-theoretic study provided any objective support for it, open-review has been advocated by various researchers. While strong calls for a shift towards open review have been issued in authoritative venues, such as the January 2009 IEEE Spectrum, the open peer-review is not yet common in computer science symposiums. However, several highly rated journals in natural sciences are currently employing open-review procedures (*Atmospheric Chemistry and Physics*, *Biology Direct*, *Journal of Medical Internet Research*, etc.). Some open-review schemes reveal only either the reviews or the name of the reviewers of accepted papers, while other venues publish entirely the name, reviews, and answers from authors. Sometimes the community can see submitted papers and researchers can propose themselves as reviewers of journal submissions. The obvious problem is that researchers may be reluctant to write negative reviews if they are going to be published. The way *Biology Direct* addresses this problem is by accepting only articles which receive at least three reviews [Koonin et.al.]. The corresponding reviewer-author graph that we obtained by parsing the publicly available data is shown in Fig. 2. The node size

is proportional to the number of reviews written. As seen in this image, such data about communities around research publication venues can help detect close knit sub-communities and highly influential reviewers (large nodes in the picture). Some sub-communities, such as the cluster on the bottom right, can be completely separated in the review process from the rest of the researchers. Other quite large communities can be linked via as few as 2 or 3 researchers.




Fig. 2. Reviewer-author relation at *Biology Direct*.

Another kind of information offered by open review is illustrated by the reviewer-paper graph. The reviewer-paper relations for *Biology Direct* is shown in Fig. 3. It reveals that many papers are reviewed only by researchers not involved in reviewing anything else for this journal. Under the working assumptions, this raises questions about whether the given paper is relevant to the core community. Meanwhile, a few authors review a significant number of papers, yielding an unmatched influence on what is being published.




Fig. 3. Reviewer-paper relation at *Biology Direct*. Darker nodes show reviewers and red nodes show papers.

Using open review for workshops and conferences is complicated by the limits on the time available for writing publishable, good quality reviews. A rare example is a 2007 ConnectED workshop in Design. The next section details the reviewing mechanism planned for the 2013 workshop on decentralized coordination.

III. DECENTRALIZED COORDINATION WORKSHOP (DCW)

Decentralized coordination is a challenging problem in multi-agent system as well as in human societies in general. With byzantine behavior from agents, even the seemingly simple problem of agreeing on a bit is not trivial. A robust deterministic agreement protocol was proved impossible even in the case of a single failure. The area of distributed computing has seen significant work on the problem of byzantine consensus. Typically a limit is assumed on the number of supported incorrect participants. Similarly, the area of distributed CSPs has a significant impossibility result concerning self-stabilization when all participants have equal priority.

With decentralized coordination, the focus is on techniques and applications where the decisions are construed via a distributed process by multiple participants. The participants are assumed to have a fair say in the final decision. A challenge is to make the coordination process robust to attempts of manipulation by a subset of the players.

It is a common practice that the articles accepted in a venue are not so much decided based on a global merit but based on their merit with respect to the interests and expertise of the reviewers in the community around that venue. This explains the decision mechanism used by the *Biology Direct* journal where any article receiving three open peer-reviews is published.

In DCW, submitted papers will be posted such that workshop committee members can bid on reviewing the ones they find interesting and where they feel they can write a meaningful review. Each article is allocated to some reviewers that bade on it. Remaining reviewing assignments are randomly allocated to reviewers that did not get the papers for which they did bid.

Authors will get an opportunity to write a response to the received reviews,

and the answer will be published together with the reviews. A reviewer can withdraw her review after seeing the answer to it. After reading the articles, reviewers assigned to a submission can decide to not actually submit a review. If the reviewer did bid for that paper, her name will still be officially marked as an assigned reviewer of the corresponding article. Articles that receive reviews will be sorted and accepted for either oral or poster presentation. If the authors do not withdraw them after seeing the reviews, submissions will be published together with the reviews and author answers, in the peer-reviewed section of the proceedings. Submissions for which nobody bids and that nobody reviews will only be made available as non-reviewed technical reports in the final proceedings.

A given program committee has a limited capacity in terms of number of quality reviews that it can provide for a workshop. This capacity can be dynamically extended by inviting new reviewers after submissions. By opening the organization process, the obtained experiment is useful not only as reference for future organizers, future studies of peer-reviewing processes, and research on community detection, but can also facilitate the community formation. The workshop call for papers is available at: <http://cs.fit.edu/~msilaghi/WDC>.

BIBLIOGRAPHY

[Peterson et al.]; “Game Theoretical Modeling and Studies of Peer-Reviewing Methods”, ISPR, Vol I, pp 267-272; 2009.

[Koonin et al.]; “A community experiment with fully open and published peer review”, *Biology Direct*, 1:1; 2006

Marius Silaghi is assistant professor in Computer Sciences at Florida Tech. His main research interests lie in the application of artificial intelligence, game theory and cryptology to support distributed decision making.

Song Qin is a PhD candidate in Computer Science at Florida Tech. He is interested in decentralized coordination and complex networks.

William Cheung is associate professor in Computer Science at Hong Kong Baptist University. His work focuses on data mining and intelligent systems.