

DecentPeeR: A Self-Incentivised & Inclusive Decentralized Peer Review System

Johannes Gruendler¹ Darya Melnyk¹ Arash Pourdamghani^{1,2} Stefan Schmid^{1,2}
¹TU Berlin, Germany ²Weizenbaum Institute, Germany

Abstract—Peer review, as a widely used practice to ensure the quality and integrity of publications, lacks a well-defined and common mechanism to self-incentivize virtuous behavior across all the conferences and journals. This is because information about reviewer efforts and author feedback typically remains local to a single venue, while the same group of authors and reviewers participate in the publication process across many venues. Previous attempts to incentivize the reviewing process assume that the quality of reviews and papers authored correlate for the same person, or they assume that the reviewers can receive physical rewards for their work. In this paper, we aim to keep track of reviewing and authoring efforts by users (who review and author) across different venues while ensuring self-incentivization. We show that our system, *Owly*, incentivizes reviewers to behave according to the rules, i.e., it has a unique Nash equilibrium in which virtuous behavior is rewarded.

I. INTRODUCTION

Peer review systems have an extensive impact in today’s academia: they have become the backbone of the scientific publication process [1], [2], the distribution of academic funds [3], [4], as well as of the open source software development [5], [6]. Since the introduction of peer review systems, various methods have been proposed to make the peer review procedure more inclusive. The goal of inclusivity is to give authors a chance to publish their work based on its quality rather than on other aspects. Ensuring inclusivity is challenging in the academic peer review process: submissions on different research topics may not be comparable; reviewers may have personal preferences depending on the topic of the submission; due to large amounts of published papers, evaluations from only few reviewers can be used to decide on the quality of a submission. Previous efforts to ensure inclusivity range from enforcing prior-announcement of conflict-of-interest [7], [8], double-blindness [9], [10], and actions from the editor to promote quality, integrity, and fairness [11]. Most traditional solutions are focused on how to make a *single* conference more inclusive. However, authors and reviewers likely take part at multiple venues over their careers. Thus, a cross-venue measure would be more viable today, given the advancements in decentralized technologies.

Initial proposals for a blockchain-based peer review systems [12], [13] focus on providing an alternative coin instead of Bitcoin. These attempts to alter a financial system for peer review, however, are inherently flawed: wealthy participants can game the system to their benefit.

This work is supported by the European Research Council (ERC) under grant agreement No. 864228 (AdjustNet), 2020-2025.

More recent systems [14], [15] aimed to tackle the challenge of a collaborative system. In doing so, they showed what are possible ways to provide self-invitation based on a game-theoretic perspective. In a nutshell, they showed that the allowed rules of the game are in the best interest of all users. However, they did not show what happens when a failure happens and how the system could recover. Furthermore, they did not consider the fact that not only a single venue exists, and a system should not be restarted whenever a new request for a venue appears.

In this work, we keep track of the actions of users *over time* using a *reputation system* to ensure inclusivity. We build a decentralized system where users tend to follow the rules based on their best interests. While incentivizing users who behave rationally, the system should not punish academic work of good quality and thus violate inclusivity. To this end, we develop a self-incentivized system based on a game theoretical approach, showing that achieving the unique Nash equilibrium is only possible by adhering to the rules of the system.

Our method incorporates three mechanisms that lead to a high level of inclusivity:

- Our system only considers the reputation score in borderline cases: if the high quality of a paper is already agreed upon, we consider that paper as accepted.
- Our system uses a randomness mechanism to form a program committee and a reviewing team: the weighted randomness ensures that selected users can be trusted while giving chance to every user to participate.
- A reputation score function has been implemented with the goal of ensuring fast recovery for users who have limited misbehavior.

II. SYSTEM DESIGN

In the following, we detail how our system benefits from the decentralized structure of blockchain systems. In particular, how we keep track of the venues’ data and allow users to interact with the data. We then detail the steps taken to realize a peer-review process.

A. Blockchain-Based Implementation

We now describe the blockchain setup that can be used to implement the *Owly*.

Storage on blockchain. We benefit from the blockchain storage capabilities in two ways: Firstly, by storing the papers’ and the users’ metadata (e.g., name) on a public ledger, so they can be retrieved and used easily in the future. Secondly, we keep

the confidential information about papers (e.g., their content) and users (e.g., their reputation) encrypted in a second-layer storage system like IPFS [16].

Organization via smart contracts. Our peer review system benefits from smart contracts [17], [18] that are used to implement its core functionalities. Actions such as registration, random generation for reviewer assignment (using tools such as [19], [20]), and paper submissions are handled via the smart contract. It plays a vital role in safely decentralized tracking of the behavior of the users, reputation scores, and topic tags. In addition, it allows us to implement the system across venues. The possible required payments for maintaining the contract is considered to be included in the conference fee.

B. Conference Processes

Our system consists of a set of *users* that have two roles: *authoring* and *reviewing* papers. We use N to denote the set of all n users in the system. Every user in this mechanism has a *reputation score* that is initially set to $R_i^t = \frac{1}{2}$ for a user i and time t_i . As we later show, the reputation score can influence the status of the paper and the users themselves. A schematic design of *Owly's* is shown in Figure 1

Venue & reviewer pool creation. To initiate a new venue (or add a new iteration of the venue), program chairs can create a venue-specific instance. Based on the topics that they mentioned in the definition of the conference, the system suggests a pool of reviewers with high reputation who have indicated similar topics to the conference description.

Paper submission & reviewer assignment. When an author wants to submit a paper, we consider the case that a part of their reputation is stored as a deposit to avoid spam submissions. After the successful submission of a paper, a reviewer is chosen uniformly at random from the set of reviewers (excluding the authors and conflicts of interest) for this paper. Then reviewers get a call to review individually. If they accept this call, they should complete the review in the given time frame.

To balance the reviewers based on their expertise in the respective area, a *confidence score* is calculated for each reviewer. Such a score can be calculated with any of the similarity detection techniques. One can compare the degree of similarity between the tags provided by a reviewer with the tags of the paper that need to be reviewed. Let us consider the output of such a similarity detection as $\sigma(T_R, T_P) \in [0, 1]$, where T_R and T_P are the reviewer's and the paper's tags. Hence, the total competence of a reviewer j reviewing paper p is $C_j^p := \sigma(T_R, T_P)$.

User fault detection. If users do not adhere to the rules, this behavior should be detected by a peer review system. One example is the submission of identical reviews for a new version of the paper. In our cross-venue system, comparison to older review versions becomes possible.

Weighted score of a paper. Consider a paper which is scored by r reviewers with respective scores $S_j^p \in [1, 5]$ for reviewers $j \in \{1, \dots, r\}$. We say that a paper is honest if the average score is above a preset threshold.




Fig. 1: *Owly* processes and their division into an off/on-chain as well as venue-specific/cross-venue.

A borderline paper receives a weighted score W^p . This score is based on the reputation of an author R_i^t and the score for the paper S_j^p by reviewer j . The weighted score is calculated as $W^p := R_i^t \cdot \frac{\sum_{j=1}^r C_j^p \cdot R_j \cdot S_j^p}{r \cdot \sum_{j=1}^r C_j^p \cdot R_j}$. The link between the score of a paper and the reputation of a user is an incentive to gain a high reputation score.

Paper selection. Papers that receive a score above a predefined threshold for acceptance are accepted directly, and similarly for papers with very low score. Besides these two extremes, there might exist a set of borderline papers. If there is still a possibility to accept more papers, our method chooses those based on the average reputation score of authors. This motivates the reviewers to submit high-quality reviews, as a high reputation leads to acceptance of their own borderline-scored papers. Furthermore, it is still possible for all authors to be part of the scientific community if their papers are good enough, despite their poor review reputation.

Reputation score update. A user's reputation score is updated at the end of the reviewing process of any given venue. The honest behavior of a user leads to a higher score, and faulty behavior leads to a score reduction.

Our system has a *normal form game* in its heart, modeling the change in the users' reputation based on their behavior. We designed the parameters of the game in a way that new users can obtain the reputation score quickly. Also, the loss of reputation for faulty behavior is set to be high, but the lost reputation can be recovered over time. Details on how the reputation score is defined and updated can be found in the full version of the paper.

III. CONCLUSION & NEXT STEPS

This paper introduced a self-incentivized, cross-venue, and inclusive peer review system. We designed a mechanism that provably motivates users to follow the rules while submitting and reviewing a paper, ensuring a fair chance for every user to participate in our mechanism. Our method relies on a reputation score, which in turn can be utilized across venues, assisting program chairs and the choices they have to make.

In the next step, we aim to augment our model as a smart contract, showing the effectiveness of our approach via simulations on public peer-review data sets. Our goal is to make this system accessible to the public, so it can be tested beyond our simulations.

REFERENCES

- [1] T. Ross-Hellauer, A. Deppe, and B. Schmidt, "Survey on open peer review: Attitudes and experience amongst editors, authors and reviewers," *PLoS one*, 2017.
- [2] R. Walker and P. Rocha da Silva, "Emerging trends in peer review—a survey," *Frontiers in neuroscience*, 2015.
- [3] S. Sato, P. M. Gyax, J. Randall, and M. Schmid Mast, "The leaky pipeline in research grant peer review and funding decisions: challenges and future directions," *Higher Education*, 2021.
- [4] E. Reale and A. Zinilli, "Evaluation for the allocation of university research project funding: Can rules improve the peer review?" *Research Evaluation*, 2017.
- [5] J. Wang, P. C. Shih, and J. M. Carroll, "Revisiting linus's law: Benefits and challenges of open source software peer review," *Int. J. Hum. Comput. Stud.*, 2015.
- [6] P. C. Rigby, B. Cleary, F. Painchaud, M. D. Storey, and D. M. Germán, "Contemporary peer review in action: Lessons from open source development," *IEEE Softw.*, 2012.
- [7] I. Radun, "Nonfinancial conflict of interest in peer-review: Some notes for discussion," *Accountability in Research*, 2021.
- [8] J. S. Ancker and A. Flanagin, "A comparison of conflict of interest policies at peer-reviewed journals in different scientific disciplines," *Sci. Eng. Ethics.*, 2007.
- [9] A. Tomkins, M. Zhang, and W. D. Heavlin, "Reviewer bias in single-versus double-blind peer review," *Proceedings of the National Academy of Sciences*, 2017.
- [10] M. Sun, J. B. Danfa, and M. Teplitskiy, "Does double-blind peer review reduce bias? evidence from a top computer science conference," *J. Assoc. Inf. Sci. Technol.*, 2022.
- [11] D. B. Resnik and S. A. Elmore, "Ensuring the quality, fairness, and integrity of journal peer review: A possible role of editors," *Sci. Eng. Ethics*, 2016.
- [12] M. Sharples and J. Domingue, "The blockchain and kudos: A distributed system for educational record, reputation and reward," in *EC-TEL*, 2016.
- [13] T. Wang, S. C. Liew, and S. Zhang, "Pubchain: A decentralized open-access publication platform with participants incentivized by blockchain technology," in *IEEE ISNCC*, 2020.
- [14] A. Ayyukt, G. S. Ramachandran, and B. Krishnamachari, "A decentralized review system for data marketplaces," in *IEEE ICBC*, 2021.
- [15] J. X. Lim, B. Monnot, and G. Piliouras, "Blockchain-based mechanism design for collaborative mathematical research," in *IEEE ICBC*, 2022.
- [16] E. Daniel and F. Tschorsch, "IPFS and friends: A qualitative comparison of next generation peer-to-peer data networks," *IEEE Commun. Surv. Tutorials*, 2022.
- [17] V. Y. Kemmoe, W. Stone, J. Kim, D. Kim, and J. Son, "Recent advances in smart contracts: A technical overview and state of the art," *IEEE Access*, 2020.
- [18] T. M. Hewa, Y. Hu, M. Liyanage, S. S. Kanhare, and M. Ylianttila, "Survey on blockchain-based smart contracts: Technical aspects and future research," *IEEE Access*, 2021.
- [19] K. Chatterjee, A. K. Goharshady, and A. Pourdamghani, "Probabilistic smart contracts: Secure randomness on the blockchain," in *IEEE ICBC*, 2019.
- [20] P. Qian, J. He, L. Lu, S. Wu, Z. Lu, L. Wu, Y. Zhou, and Q. He, "Demystifying random number in ethereum smart contract: Taxonomy, vulnerability identification, and attack detection," *IEEE Trans. Softw.*, 2023.