# Summary Review Documentation for
# "Boosting Market Liquidity of Peer-to-Peer Systems Through Cyclic Trading"

Authors: Raphael Eidenbenz, Thomas Locher, Stefan Schmid, Roger Wattenhofer

## REVIEWER #1

*Strengths:* The strengths are (i) nice idea, (ii) clear evaluation, (iii) real implementation.

*Weaknesses:* The weaknesses are (i) lack of empirical results from implementation, (ii) lack of discussion of potential gaming, (iii) relatively small delta to previous work (e.g. [11] & [14]).

*Comments to Authors:* One problem with this idea is that it sort of assumes that different swarms are substitutable goods. In reality, I can use a torrent to download legitimate content (e.g. open source software distribution, legal computer game world data updates) and copyright breaching content—token exchange between these different systems might cause interesting problems. Second problem is potential for complex gaming across cycles. Finally, there's potential for re-identification of users across multiple torrents by clever adversaries.

## REVIEWER #2

*Strengths:* The strengths are (i) well structured and written, (ii) well presented motivation, simple and great idea, (iii) valid verification and convincing results.

*Weaknesses:* The weaknesses are (i) avoid so many buzzwords and internaly-used and defined terminology. Unnecessarily complicates reading and hides the value. That was the only thing that is really annoying while reading, (ii) is bitTorrent still relevant? Show the impact of your work in reality.

*Comments to Authors:* I really liked the paper, but W1 and W2 comments needs to be addressed.

## REVIEWER #3

*Strengths:* The strengths are (i) realistic improvement to a real-world system, (ii) well presented, with very informative graphs on the effect of k-cycles, (iii) a distributed version that appears to be sound and efficient (but not shown to be!).

*Weaknesses:* The weaknesses are (i) the distributed version is not evaluated. This is very important for the system to work, (ii) figure 5 should be augmented by the results of the distributed implementation. Maybe two figures side-by-side.

*Comments to Authors:* A real implementation and deployment would allow for a better assessment of the benefits earned, as the effect of TCP and the effect of timing in transitions from uploading to choking a neighbor could make a difference. One issue I see is that tit-for-tat is known to allow for some elasticity (necessary to bootstrap new nodes), which has been exploited by systems such as BitThief to free-ride on it. Cycles of length $> 2$ might be susceptible to more "elasticity", and they might be more prone to free-riding exploitation. A real implementation would allow to shed some light on that.

## REVIEWER #4

*Strengths:* The strengths are (i) it proposes a mechanism, cyclic trading, that has a lot promise for improving the performance of P2P file download, (ii) it is well-written and easy to understand, (iii) it describes not only a general design but also a plausible strategy for implementation.

*Weaknesses:* The weaknesses are (i) the utility of the mechanism depends on the prevalence of opportunities for cyclic sharing. But, since all evaluation is on synthetic workloads, the paper does not make a convincing case that cyclic sharing opportunities really exist in the real world, (ii) the mechanism they propose for preserving privacy does not actually do so, (iii) the system was not actually implemented, but is only evaluated in simulation.

*Comments to Authors:* The main problem I have with the paper is that its claims are not convincingly evaluated. The paper makes the argument that cyclic trading opportunities exist in sufficient number to provide utility. However, the evaluation of this claim is made only on highly artificial models of BitTorrent streams. There seems to be one data point that anchors the model, namely the distribution of the cardinality of each swarm. However, this is not nearly sufficient information from which to derive an entire population distribution of block requirements and availability. You devise a simple model, and in V-E a more complex model, but you never validate either model with any real data. I'd suggest obtaining a detailed trace of BitTorrent swarms, and using that to evaluate the presence of cyclic trading opportunities. Only then will you have good evidence that your approach is reasonable. Furthermore, your privacy-preserving mechanism for learning cycles does not really seem to preserve privacy. First of all, it exposes the participants in any 3-cycle. Thus, it clearly reveals "more private information than with an intra-swarm protocol", contradicting the last sentence of the first paragraph of section VI. Second, since it uses the same CID any time a particular edge occurs, even if that edge appears in multiple cycles, it is open to frequency analysis. Cryptographic schemes are only secure if things are encrypted differently each time they're used, but the fact that you use a deterministic and unsalted hash function means that it is not secure. I'm not sure why you even bother with this privacy-preserving scheme since your paper repeatedly makes the point that cycles with k¿3 are unworthwhile, and they only seem to provide a possible benefit for such long cycles.

## REVIEWER #5

*Strengths:* The strengths are (i) simple mechanism that gives good performance improvement, (ii) in depth understanding of the mechanism through simulations of various cases and optimizations, (iii) well-written and easy to understand.

*Weaknesses:* The weaknesses are (i) only simulations and the simulation model is not very realistic, (ii) the idea is simple and indirect bartering in BT has been explored before, (iii) no implementation or real experiments, only simulations.

*Comments to Authors:* The paper makes a good case for cycle trading. My biggest worry is whether the simulation model is accurate enough? Instead of preferential attachment, why not use actual traces or distributions that match real world traces? There are many variables that would benefit from that, e.g. upload/download bandwidth distribution over space and time, interest for different content and swarms, proximity information and distribution of latency/throughput of pairs of peers. This makes it very hard to know whether this actually would work so well. How can this system be gamed? Please discuss this. Certainly, it's easier to collude and lie in this system to create fake cycles and milk pieces out of it before the peers realize that the cycle is broken.

## RESPONSE FROM THE AUTHORS

We would like to thank the reviewers for their good feedback on our work.

We tried to address all issues that the reviewers raised. We explain our contributions better, reorganized the article and also put the contribution better into perspective. In particular, we discuss how our work relates to the BitTorrent protocol and provide better motivations and explanations for the simulated scenarios. We also emphasize that our evaluation is indeed based on a *distributed* simulation; this seems to have caused some confusion.

We conducted several additional simulations. For example, we now compare two different latency models: one with constant latencies of 60ms between all peers, and one where the latencies capture the distribution of peers over three continents (e.g., Asia, Europe, America), that is, we draw the latency of a connection from three different Gauss distributions depending on the number of continental hops (i.e., Normal distributions $\mathcal{N}(30, 10)$ for transmissions within a continent, $\mathcal{N}(60, 10)$ for one continental hop, and $\mathcal{N}(90, 10)$ for two continental hops). We find that both models yield a similar performance, also if the variance is increased further. This indicates that our results are relatively robust and general.

Furthermore, the new paper version attends to the incentive problem of cyclic trading and explains how under collusion, cycles may deteriorate the fairness of the protocol by one packet per cycle. Finally, the paper also explicitly states that user privacy aspects are beyond the scope of this paper, but they constitute an interesting direction for future research.