Going Beyond Tit-for-Tat: Designing Peer-to-Peer Protocols for the Common Good

Ryan Peterson Emin Gün Sirer

Peer-to-peer systems, in which participants pool their resources to accomplish their goals, have become ubiquitous. Since rational peers engage in strategic behavior, much past work has examined the design of mechanisms that incentivize peers to provide resources. The predominant design paradigm to date has been tit-for-tat, in which each peer benefits from every interaction. In this paper, we discuss the limitations of the tit-for-tat design paradigm and make the case for a new approach based on seeking the globally optimal outcome known as the common good.

Peer-to-peer (P2P) systems, in which participants simultaneously act as both clients and servers, have become ubiquitous in recent years and dominate bandwidth consumed on the Internet. Compared to client-server systems that preceded them, peer-to-peer systems promise an unprecedented level of performance, fault-resilience and scale because they command a large pool of resources contributed by their participants. Yet getting peers to make their resources available to other peers poses a fundamental challenge: what incentives does a peer have to provide a valuable commodity, such as bandwidth, storage space, CPU time, peripherals, a network vantage point, access to devices, and so forth, to other peers?

Ever since initial studies uncovered strategic behavior at the network level, much past work on peer-to-peer systems has focused on the design of algorithms that provide incentives for peers to contribute their resources. Most of this work, however, has been limited to algorithms that ensure pair-wise parity in peer-to-peer interactions, known as the so-called *tit-for-tat* paradigm. BitTorrent [1] is perhaps the best-known embodiment of the tit-for-tat paradigm, where peers exchange blocks preferentially with other peers with whom they have successfully exchanged blocks in the past at high bandwidth. Samsara [2] applies the tit-for-tat paradigm to block storage, where every participant that stores a block on a distant peer provides a transferable guarantee (a claim) that it will provide storage for the distant peer when needed. And systems such as SHARP [3] generalize the tit-for-tat paradigm by, in effect, providing a multi-currency scheme where every resource used by a peer is compensated by a virtual claim that entitles the remote peer to local resources in the future. The tit-for-tat paradigm is so pervasive that it is instructive to further distinguish between its applications: *immediate fair-exchange* systems, like BitTorrent, ensure that both parties exchange physical resources at the time of the interaction, while *time-dilated tit-for-tat* systems ensure that a second compensatory action will take place to compensate the resource provider in the initial interaction. Similarly, *pair-wise* tit-for-tat systems ensure parity between pairs of participants, while *delegation-based* tit-for-tat systems enable resource exchange among peers whose interactions form a cycle.

The tit-for-tat paradigm has tremendous appeal from a system designer's point of view, as it greatly simplifies reasoning about incentives. If every interaction is arranged such that every participant expects to receive a desirable resource (either immediately or later through a promissory note), clearly free-riding behavior can be contained and the system as a whole will provide value. Further, tit-for-tat intertwines the desired outcome for the overall system with the enforcement activity; every node has incentives not only to provide resources directly, but also to spend resources checking that other peers keep up their end of the bargain. Consequently, the tit-for-tat approach has become the dominant paradigm in robust peer-to-peer systems.

Yet the tit-for-tat paradigm is undesirable for a large class of systems. First, tit-for-tat fails when resources are not interchangeable. A well-known problem with BitTorrent is that a node that has seeded a particular file gains no benefit from this altruistic behavior when downloading another file. Immediate exchange tit-for-tat fails when peers do not have resources to contribute at that instant in time. For instance, in BitTorrent, peers that do not initially have unique blocks are unable to participate in the protocol, leading to the well-known bootstrap problem where hosts cannot achieve full download bandwidth until they have amassed sufficiently many unique blocks from seeders [4]. Similarly, in a large swarm, it is quite possible for two nearby peers with very high mutual bandwidth to never discover each other because their existing peering relationships provide them with non-unique blocks. Second, time-delayed tit-for-tat carries inherent risks of bankruptcy of the underwriting peer, as well as risk of currency inflation induced by certain issuers and lack of liquidity in the market (for instance, due to churn) that can keep a peer from receiving its owed compensation. Mitigating these risks often requires such systems to be coupled with a reputation system, yet reputation systems themselves are inherently difficult to construct and prone to gaming. Finally, all tit-for-tat systems suffer from an inherent tension between discovery and exploitation that leads to shortcuts that open up vulnerabilities in the protocol. For instance, BitTorrent peers "optimistically unchoke" other peers, providing them with the benefit of the doubt and engaging in a transaction in the hopes of discovering a high-bandwidth peer despite lack of a prior history. This, in turn, enables parasitic behavior where a node can exploit all members of the swarm, and achieve download speeds far in excess of regular downloaders [5, 6]. Overall, tit-for-tat fails because each peer selfishly and locally seeks to maximize its profits, without concern for the global good and often without conferring with other nodes. From a game-theoretic perspective, the entire system can get stuck at one of many Nash equilibria, one in which the global outcome is not optimal, and ad hoc mechanisms not only provide no guarantee that the system will behave well, but themselves introduce opportunities for misbehavior.

In this paper, we argue that a new, principled paradigm is necessary for high-performance peer-to-peer services. We call this new design methodology the *common good* paradigm. Whereas tit-for-tat systems are designed to yield a locally good outcome for each participant in each interaction, common good systems are designed to yield a globally good outcome for their participants. That is, they relax the design constraint that each participant benefit from each interaction, greatly expanding the space of possible behaviors open to peers. This freedom comes at the cost of requiring participants to engage in behaviors, such as contributing resources, that might temporarily go against their local best interests. We provide policing mechanisms to ensure that peers will engage in such behavior when called upon, and claim that seeking the globally optimal outcome is not at odds with seeking the locally optimal outcome. To the contrary, a set of peers that seek the globally optimal strategy can explore more of the state space compared to peers that only engage in locally beneficial behaviors, and arrive at an outcome where the peers' local utility functions are optimally satisfied, whereas it is quite difficult to ensure that choosing purely the locally-beneficial interactions will yield a globally good outcome.

The key challenges in designing peer-to-peer systems that seek the common good are to define the globally optimal outcome, to disseminate sufficient information to each peer such that it can work towards this outcome, and to build enforcement mechanisms that will disincentivize peers from making selfish decisions against the interests of the collective. While building systems that seek the common good against peers' self-interests may seem utopian, a suitably designed system will make good behavior the rational choice. We next outline a file download system, akin to BitTorrent but based instead on the common good paradigm, to illustrate the new paradigm and show how it differs from the traditional designs based on tit-for-tat.

High-performance file downloads showcase the contrast between the tit-for-tat and the common good paradigms. Recall that in BitTorrent, each peer exchanges blocks tit-for-tat with each peer as it discovers opportunities to do so; in this endeavor, it is limited by the unique blocks it has available, its peers' behavioral histories, and the result of ad hoc mechanisms such as randomized unchoking; consequently, the resulting set of block exchanges may be far from optimal. In contrast, an alternative file download system based on the common good paradigm can achieve far higher performance. Suppose, for a moment, that there exists an omniscient entity that can determine a network-wide block delivery schedule, see the result of every interaction and hold internal state. Such an entity can iteratively, and efficiently, determine a block distribution graph in which nearby nodes are preferentially placed together, new nodes are immediately accommodated, and misbehaving nodes are cut out of the network. It can converge on a globally desirable configuration efficiently because it can iteratively try, and learn from, all past configurations. It relies entirely on principled decisions instead of ad hoc mechanisms. And it does not allow a misbehaving node to exploit every single peer separately, because it can hold state. Such an omniscient coordinator for a peer-to-peer system can enable the network to determine, top-down, an outcome that is far better than the bottom-up outcome of the uncoordinated activities of the peers.

Practically implementing such an omniscient entity is not impossible. One consideration is providing each peer with enough information about the current state of the system to determine how it should behave to benefit the common good. The two straightforward methods, namely, to rely on a centralized trusted entity or to send all information to all peers, both suffer from scalability problems. In between these two points of operation, there exists a design space in which peers are organized into a hierarchy that channels aggregate system information until each peer has an approximate global view of the system. Another consideration is enforcement; misbehaving peers need to be identified and excised, and doing so is not difficult in the presence of global information channels.

Overall, we believe that designing systems that seek the common good will yield better, more robust outcomes than systems in which participants solely seek their own narrowly-defined self-interest.

References

[1] B. Cohen. Incentives Build Robustness in BitTorrent. In *P2PECON Workshop*, Berkeley, CA, May 2003.

[2] L. P. Cox and B. D. Noble. Samsara: Honor Among Thieves in Peer-to-Peer Storage. In *SOSP*, Bolton Landing, NY, October 2003.

[3] Y. Fu, J. Chase, B. Chun, S. Schwab and A. Vahdat. SHARP: An Architecture for Secure Resource Peering. In *SOSP*, Bolton Landing, NY, October 2003.

[4] C. Gkantsidis and P. R. Rodriguez. Network Coding for Large Scale Content Distribution. In *INFOCOM*, Miami, FL, March 2005.

[5] T. Locher, P. Moor, S. Schmid and R. Wattenhofer. Free Riding in BitTorrent is Cheap. In *HotNets Workshop*, Irvine, CA, November 2006.

[6] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy and A. Venkataramani. Do Incentives Build Robustness in BitTorrent?. In *NSDI*, Cambridge, MA, April 2007.