

Incentive-Compatible Distributed Greedy Protocols

Noam Nisan^{*} Michael Schapira[†] Gregory Valiant[‡] Aviv Zohar[§]

Abstract

Under many distributed protocols, the prescribed behavior for participants is to behave greedily, that is, to repeatedly “best respond” to the others’ actions. We ask the following general question: “*When is it in the best interest of a long-sighted participant to adhere to a distributed greedy protocol?*”. We take a game-theoretic approach and exhibit a class of games where greedy behavior (*i.e.*, repeated best-response) is incentive compatible for all players. We identify several environments of interest that fall within this class, thus establishing the incentive compatibility of the natural distributed greedy protocol for each. These environments include models of the Border Gateway Protocol (BGP) [4], that handles routing on the Internet, and of the Transmission Control Protocol (TCP) [3], and also stable-roommates assignments [2] and cost-sharing [5, 6], that have been extensively studied in economic theory.

1 Motivation: When is it Best to be Greedy?

The Internet has transformed computation from a largely local endeavor to one that frequently involves diverse collections of self-interested individuals and organizations. While in traditional cooperative networks non-faulty nodes are trusted to follow the protocol specification, in Internet environments we can no longer take for granted that nodes blindly and unquestioningly follow a prescribed behavior. To guarantee correct execution of a protocol we must ensure that the self-concerned nodes have incentive to follow the protocol. Over the past decade, much research has been devoted to the design and analysis of incentive compatible computational protocols in both centralized [7] and distributed [1] settings.

Given the considerable size and volatility of the Internet, protocols that run on the Internet are often distributed across multiple computational nodes, whose actions rely on very limited processing power and memory, and on local information only. Often, under such protocols, nodes’ prescribed behavior is simply to repeatedly “best respond” to each others’ actions, *i.e.*, to behave in a greedy manner (*e.g.*, under the Border Gateway Protocol, that handles routing on the Internet, every router repeatedly selects the “best” route made available to it by its neighbors). We wish to understand when following such a distributed greedy protocol is *rational* from a computational node’s perspective. We ask the following general question: “*When is it best, in the long run, for a participant to repeatedly best-respond to others’ actions?*”, that is, “*When can’t a long-sighted participant improve over this repeated myopic optimization?*”.

We take a game-theoretic approach to distributed greedy protocols. We now briefly and informally present our main contributions. We refer the reader to [8] for a more thorough exposition.

^{*}School of Eng. and Computer Science, The Hebrew University of Jerusalem. noam@cs.huji.ac.il

[†]Dept. of Computer Science, Princeton University. ms7@cs.princeton.edu

[‡]Dept. of Computer Science, UC Berkeley. gvaliant@eecs.berkeley.edu

[§]Microsoft Research, Silicon Valley. avivz@microsoft.com

2 Our Framework and General Result

Our first contribution is a formal framework for reasoning about incentives for repeated best-response. We model the computational nodes as players in a partial-information game and their possible actions as strategies in this game. Each player aims to maximize a *private* utility function. We consider dynamic interactions where players take turns selecting strategies; at each (discrete) time step, some player (possibly more than one) selects and announces a strategy. The prescribed behavior for every player is to repeatedly choose his best-response (given his private utility function) to the most recently announced strategies of the others. We point out that there is nothing preventing a player from not following the prescribed behavior and from selecting strategies based on more complex computation that takes into account the entire history of interaction.

Intuitively, repeated best-response is incentive-compatible if, when all other players are repeatedly best-responding, a player has incentive to do the same, that is, no player can gain from unilaterally deviating from repeated best-response. In distributed protocol contexts, this translates to the desideratum that as long as other computational nodes execute the protocol, the best course of action for a node be to also adhere to the protocol. Defining incentive compatibility in our setting involves many intricacies. We opt to focus on a very general notion of incentive compatibility that, we believe, captures essentially any variant that the reader may desire.

Our main results are identifying a class of games for which repeated best-response is incentive compatible, and exhibiting several interesting environments that fall within this class. We formally present our class of games in [8]. Here, we settle for an intuitive exposition of this class of games. Games that belong to our class have the following property. When each player i considers the game after the other players have already iteratively eliminated “dominated strategies”, *i.e.*, strategies that are inferior to other strategies and that can be eliminated regardless of what i does, he can already tell that he can do no better than the outcome that is reached via repeated best-response.

Our main, and quite easy to prove, general theorem is that in this class of games, for every starting state of the system and every (finite or infinite) order of player activations with “sufficiently many” rounds (where a round is a sequence of consecutive time steps in which each player gets to act at least once), repeated best-response is incentive compatible for every player. This result holds even when players are activated in an asynchronous manner and even when players’ strategy announcements can be arbitrarily delayed.

3 Applications: Routing, Congestion Control, and More

We prove that each of the four environments below can be formulated as a game that falls within our class of games. Thus, our general result above implies that repeated best-response is incentive compatible in all the contexts below.

Internet routing. The Border Gateway Protocol (BGP) establishes routes between the smaller networks that make up the Internet. We abstract the results in [4] and prove that BGP is incentive compatible in realistic environments.

Congestion control. The Transmission Control Protocol (TCP) handles congestion on the Internet. Building upon [3], that models key aspects of TCP, we consider behavior that is somewhat similar to TCP: increase your attempted transmission rate until encountering congestion, and then decrease the transmission rate. We show that such behavior is in equilibrium, in the sense that no participant has incentive to deviate from it.

Stable-roommates. In this classic setting [2], students must be paired for the purpose of sharing dorm rooms, and each student has a private full order over possible roommates. The objective is

to find a “stable matching” where no two students prefer each other to their assigned roommates. We show that a natural distributed protocol, in which a student repeatedly proposes to his most preferred roommate among those that would not immediately reject him, and immediately rejects all proposers except for his most preferred proposer, is incentive compatible and converges to a stable matching in well-studied environments (interns-hospitals, correlated markets).

Cost-sharing. Cost-sharing arises in situations in which the cost of some public service (*e.g.*, building a bridge) must be distributed among self-interested users that can benefit from this service to different extents. We present a distributed protocol that achieves this goal in an incentive-compatible manner. Our protocol implements the outcome of the Moulin mechanism [5, 6].

4 Future Research

We view this work as a first step towards a more general research agenda. Natural myopic dynamics, *e.g.*, repeated best-response, fictitious play and regret minimization, have been extensively studied in game theory and economics. Yet, little attention has been given to the question of when such “myopic optimization” dynamics are also rational to follow in the long run. We have tackled this question in the context of repeated best-response. However, we believe that the examination of other dynamics from game theory/economics literature is an interesting direction for future research. Positive and negative results along these lines can help shed new light on the incentive structure of existing protocols (see our results for BGP and TCP and the results in [3, 4]), and provide new insights into the design of new incentive-compatible protocols.

In addition, our results establish *sufficient* conditions for distributed greedy protocols to be incentive compatible. We still lack a *characterization* of environments where distributed greedy protocols are incentive compatible.

References

- [1] Joan Feigenbaum, Michael Schapira, and Scott Shenker. *Distributed Algorithmic Mechanism Design*. “*Algorithmic Game Theory*”, Chapter 14. Cambridge University Press, 2007.
- [2] David Gale and Lloyd S. Shapley. College admissions and stability of marriage. *Amer. Math. Monthly*, (69):9–15, 1962.
- [3] P. Brighten Godfrey, Michael Schapira, Aviv Zohar, and Scott Shenker. Incentive compatibility and dynamics of congestion control. *SIGMETRICS Perform. Eval. Rev.*, 38(1):95–106, 2010.
- [4] Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. In *STOC*, pages 57–66, 2008.
- [5] Herve Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
- [6] Herve Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- [7] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1):166–196, 2001.
- [8] Noam Nisan, Michael Schapira, Gregory Valiant, and Aviv Zohar. Best-response mechanisms. In *ICS '11: Proceedings of the 2nd symposium on Innovations in Computer Science*, 2011.