**Proceedings Article**

# Improved Bounds on the Price of Stability in Network Cost Sharing Games

EUIWOONG LEE, Carnegie Mellon University
KATRINA LIGETT, California Institute of Technology

We study the price of stability in undirected network design games with fair cost sharing. Our work provides multiple new pieces of evidence that the true price of stability, at least for special subclasses of games, may be a constant.

We make progress on this long-outstanding problem, giving a bound of $O(\log\log\log n)$ on the price of stability of undirected broadcast games (where $n$ is the number of players). This is the first progress on the upper bound for this problem since the $O(\log\log n)$ bound of [Fiat et al. 2006] (despite much attention, the known lower bound remains at 1.818, from [Bilò et al. 2010]). Our proofs introduce several new techniques that may be useful in future work.

We provide further support for the conjectured constant price of stability in the form of a comprehensive analysis of an alternative solution concept that forces deviating players to bear the entire costs of building alternative paths. This solution concept includes all Nash equilibria and can be viewed as a relaxation thereof, but we show that it preserves many properties of Nash equilibria. We prove that the price of stability in multicast games for this relaxed solution concept is $\Theta(1)$, which may suggest that similar results should hold for Nash equilibria. This result also demonstrates that the existing techniques for lower bounds on the Nash price of stability in undirected network design games cannot be extended to be super-constant, as our relaxation concept encompasses all equilibria constructed in them.

Categories and Subject Descriptors: F.2.0 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Network Design, Price of Stability

## 1. INTRODUCTION

Network design games, which model cost-sharing of network links by selfish players, have received considerable attention in the algorithmic game theory literature. The price of stability of Nash equilibria has emerged as the most interesting question in such games, with a substantial line of work beginning with Anshelevich et al. [Anshelevich et al. 2004; Fiat et al. 2006; Li 2009; Bilò et al. 2010; Christodoulou et al. 2010; Bilo and Bove 2011; Kawase and Makino 2012] resulting in known bounds of $O(\log n)$ for all directed games and for undirected general games (wherein each player has a source and sink vertex she wishes to connect), $O(\log n/\log\log n)$ [Li 2009] for undirected multicast games (wherein all players share a single sink vertex), and $O(\log\log n)$ [Fiat et al. 2006] for undirected broadcast games (wherein there is a single sink vertex and a player at every node). While the upper bound is tight for directed games, the current lower bounds for undirected general games, multicast

Author's addresses: Euiwoong Lee, Computer Science Department, Carnegie Mellon University; Katrina Ligett, Department of Computing and Mathematical Sciences, California Institute of Technology.
EL is supported by the Samsung Foundation. KL gratefully acknowledges the generous support of the Charles Lee Powell Foundation.

games, and broadcast games are 2.245, 1.862, and 1.818, respectively (due to [Bilò et al. 2010]), leaving a huge gap between the upper and lower bounds.

It is remarkable that such a simple and elegant problem has defied analysis for so long, despite receiving considerable attention. This paper provides substantial new evidence for a constant price of stability for multicast games and broadcast games: first, we formally demonstrate the limitations of existing lower bound approaches (they cannot be used to obtain super-constant bounds), and second, we introduce new techniques and apply them to obtain a better upper bound.

Previous techniques for both the lower and upper bound results have limitations that make it difficult to close the gap by naturally extending them. For example, the games presented to establish lower bounds in [Bilò et al. 2010] and [Bilo and Bove 2011] are too "strong" in the sense that for every socially efficient equilibrium in these games, there exists a player who would benefit by deviating—even if she had to pay the entire cost of her new path. Any Nash equilibrium is robust to such "go-it-alone" deviations, but the converse is not true. Therefore, in this work, we seek to understand the set of strategy profiles robust to such deviations, and study their properties. We prove that the price of stability of this relaxed solution concept is bounded by a constant, implying that either the price of stability of Nash equilibria is also bounded by a constant, or we need new techniques to prove a super-constant lower bound.

With regard to upper bounds, almost all previous work (e.g., [Anshelevich et al. 2004], [Li 2009]) has focused on analyzing one particular equilibrium outcome, the global potential minimizer. Unfortunately, recent work of [Kawase and Makino 2012] demonstrates that in the worst case, the potential-minimizing equilibrium exceeds the cost of the social optimum by a super-constant factor (between $O(\sqrt{\log n})$ and $\Omega(\sqrt{\log\log n})$), and so potential-based techniques alone cannot be used to prove a constant upper bound on the price of stability. The upper bound of [Fiat et al. 2006], in a stark departure from the potential-based approach, carefully engineers a sequence of player movements to keep a certain structure of the strategy profile. We find that combining the machinery of Fiat et al. with potential minimization and additional new techniques allows us to make the first progress on upper bounds in over six years; we improve the upper bound on the price of stability in broadcast games from $O(\log\log n)$ to $O(\log\log\log n)$.

*Summary of our approach.* In Section 3, we propose and study an equilibrium relaxation, "go-it-alone equilibria," intended to capture the "restricted access" setting, wherein deviating players are unable to cost-share with compliant players. We provide a nearly complete characterization of the prices of anarchy and stability in directed and undirected graphs for this concept. In the case of directed graphs, the bounds parallel those known for Nash equilibria. In undirected graphs, the price of anarchy remains the same as its counterpart for Nash equilibria, but the price of stability is improved to $O(1)$ in multicast games and broadcast games (see Table I).

In Section 4, we leverage insights from our study of go-it-alone equilibria to present a new, improved upper bound of $O(\log\log\log n)$ for the Nash price of stability in broadcast games. Our two upper bound proofs ($O(1)$ for go-it-alone equilibria and $O(\log\log\log n)$ for Nash equilibria) each introduce new techniques. In both of them, starting from the socially optimum outcome, we construct a strategy profile iteratively, maintaining special structures. For go-it-alone equilibria, we require our strategy profile to form a chord-line graph, as has been used for lower bound constructions [Bilò et al. 2010]; for Nash equilibria, we maintain the structure proposed in [Fiat et al. 2006]. One technique we introduce is to consider a simultaneous move by multiple players (usually a set close to each other) to construct a desired equilibrium. This introduces an additional layer of complexity in the analysis, but it turns out to allow us to construct a more socially efficient equilibrium. Our second new technique, which is specific to our analysis of Nash equilibria, might be viewed as a combination of the two main upper bound techniques that appear in prior work: in each phase, we select the potential-minimizing strategy profile, subject to our structural requirements. This effectively allows us to exploit the benefits of each of those two existing approaches.

## 2. PRELIMINARIES

A network design game is defined by a tuple $(G, A, c)$. Here, $G = (V, E)$ is the underlying graph of the game, which, depending upon the scenario under consideration, may be either directed or undirected (we focus here on the undirected case). The function $c : E \mapsto \mathbb{R}^+$ defines the costs of the edges, and $A = \{1, ..., n\}$ is the set of players, where player $i$ is associated with a source-destination pair $(s_i, t_i)$ that she wishes to connect. A network design game is called a *multicast game* if all players share the same sink vertex ($t = t_1 = ... = t_n$), so that every player can be associated with her source in $V \setminus \{t\}$. A *broadcast game* is a special case of multicast games wherein every vertex in $V \setminus \{t\}$ is associated with a player. In multicast games and broadcast games, we sometimes abuse notation and identify a player and her source $v$.

The strategy space for player $i$ is the set $\mathcal{P}_i$ of paths from $s_i$ to $t_i$; we allow each player $i$ to use the same edge more than once. For an edge $e$ and a path $P_i$ for some player $i$, let $n_{e,i}$ be the number of times $e$ is used in $P_i$. Given a strategy profile $P = (P_1, ..., P_n)$ consisting of a path $P_i \in \mathcal{P}_i$ for each player, let $T(P) = \cup_i P_i$ be the set of the edges used in the current strategy profile. Let $c(T) = \sum_{e \in T} c(e)$ be the total cost of $T$, and $c(P) = c(T(P))$. The cost of player $i$, given a joint strategy profile $P$, is $c_i(P) = \sum_{e \in P_i} \frac{c(e)}{n_e} n_{e,i}$, where $n_e(P) = \sum_i n_{e,i}$ is the number of players using edge $e$, counting multiplicity. Note that $c(P) = \sum_i c_i(P)$ for any $P$. There are a number of basic economic motivations for equal sharing of costs among those that use an edge; this sharing rule can be derived from the Shapley value (hence the term "Shapley cost-sharing mechanism"), and it is the unique cost-sharing scheme satisfying a number of different sets of axioms [Anshelevich et al. 2004].

We often compare the cost of our strategy profile to that of a Steiner forest that satisfies the connectivity requirement for each player. Let $T^*$ be the min-cost Steiner forest in which $s_i, t_i$ are connected for every $i$. For any forest $T'$, let $c(T') = \sum_{e \in T'} c(e)$ be the cost of $T'$ and $d_{T'}(u, v)$ be the cost of the unique path from $u$ to $v$ in $T'$, assuming that $u$ and $v$ are connected.

A strategy profile $P = (P_1, ..., P_n)$ is a *Nash equilibrium* if no player $i$ can strictly decrease her cost by deviating to another strategy, i.e., $c_i(P_i, P_{-i}) \le c_i(P_i', P_{-i})$ for all $i, P_i' \in \mathcal{P}_i$. We introduce a new stability concept designed to model a situation where deviations can be disincentivized, for example by centralized control or by punishment of the sort that might emerge in a repeated game setting. We say a strategy profile $P = (P_1, ..., P_n)$ is a *go-it-alone equilibrium* if no player can strictly reduce her cost by constructing a path by herself, i.e., for any $i, P_i' \in \mathcal{P}_i$, $c_i(P) \le \sum_{e \in P_i'} c(e)$. Since $\sum_{e \in P_i'} c(e)$ is minimized when $P_i'$ is a shortest path from $s_i$ to $t_i$, this condition is equivalent to requiring that $c_i(P)$ does not exceed the cost of the shortest path from $s_i$ to $t_i$ for each $i$. Let $NE, GE$ be the sets of Nash equilibria and go-it-alone equilibria, respectively.

The existence of a Nash equilibrium is proved in [Anshelevich et al. 2004] via the potential function $\Phi(P) = \sum_{e \in E} c(e) H(n_e)$, where $H(k) = 1 + ... + (1/k)$. Note that in any Nash equilibrium no player pays more than the cost of her shortest path, so every Nash equilibrium is a go-it-alone equilibrium. It directly follows that a go-it-alone equilibrium exists in any network design game. Figure 1 shows that not every go-it-alone equilibrium of a network design game is a Nash equilibrium. The example is from [Epstein et al. 2007], who show that $P = ((e, c), (b, f))$ is the only Nash equilibrium and $c_1(P) = c_2(P) = 5$. Consider $P' = ((a, b, c), (b, c, d))$, where the two players cooperate. Note that $c_1(P') = c_2(P') = 4$ and the cost of the shortest path from $s_i$ to $t_i$ is 5 for each $i$, so $P'$ is a go-it-alone equilibrium.

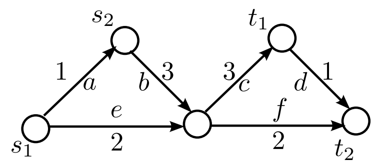


Fig. 1. A network design game where NE, GE, and the strong versions of these equilibrium concepts are pairwise distinct.

The price of anarchy (PoA) for a class of games is the ratio between the cost of the socially worst Nash equilibrium of a game instance and that of the social optimum of that instance, i.e., $\max_{P \in NE} \frac{c(P)}{c(T^*)}$. Similarly, the price of stability (PoS) for a class of games is the ratio between the cost of the socially *best* Nash equilibrium and the social optimum, $\min_{P \in NE} \frac{c(P)}{c(T^*)}$. We define analogues of these concepts with respect to go-it-alone equilibria; the go-it-alone price of anarchy (GPoA) is $\max_{P \in GE} \frac{c(P)}{c(T^*)}$ and the go-it-alone price of stability (GPoS) is $\min_{P \in GE} \frac{c(P)}{c(T^*)}$.

Since the set of go-it-alone equilibria contains all Nash equilibria, the lower bounds on the price of anarchy and the upper bounds on the price of stability directly carry over to go-it-alone equilibria; these results are proved by showing the existence of a good or bad equilibrium. However, the upper bounds on the price of anarchy and the lower bounds on the price of stability, whose proofs require considering all equilibria, do not carry over directly to go-it-alone equilibria. For the case of undirected graphs, Table I summarizes prior work, these immediate implications, and the results of this paper.

Table I. Best known results for upper and lower bounds on the price of anarchy (PoA) and price of stability (PoS) of go-it-alone equilibria (GE) and Nash equilibria (NE) in undirected graphs. Automatic implication of one result from another is denoted by an arrow (⇒). The first results from [Anshelevich et al. 2004] are marked by *, and new results from the present manuscript are marked by a †. The lower bound on the price of stability in undirected graphs has been studied extensively via constructing complicated examples, and we have not studied the implications for GE.

| | | GPoS(GE) | $\leq$ | PoS(NE) | $\leq$ | PoA(NE) | $\leq$ | GPoA(GE) |
|---|---|---|---|---|---|---|---|---|
| upper bounds | broadcast | $O(1)$† | | $O(\log\log\log n)$† | | $n$* | | $n$† |
| | multicast | $O(1)$† | | $O(\frac{\log n}{\log\log n})$ [Li 2009] | | | | |
| | general | $O(\log n)$ | ⇐ | $O(\log n)$* | | | | |
| lower bounds | broadcast | unstudied | | 1.818 [Bilò et al. 2010] | | $n$* | ⇒ | $n$ |
| | multicast | | | 1.862 [Bilò et al. 2010] | | | | |
| | general | | | 2.245 [Bilò et al. 2010] | | | | |

### 2.1. Related work

Motivated by the study of communication networks, network design games have received substantial attention over the past decade (e.g., [Anshelevich et al. 2003; Chawla et al. 2006; Hoefer 2006; Anshelevich and Caskurlu 2009a; 2009b]). Fair network design games, where an edge's cost is evenly divided among the players who use it, were first suggested by [Anshelevich et al. 2004] and have subsequently become one of the canonical areas of study of the quality of equilibrium outcomes. [Anshelevich et al. 2004] show that in both directed and undirected graphs, a Nash equilibrium always exists, the worst-case price of anarchy is exactly $n$, and the price of stability is $O(\log n)$.

Furthermore, in directed graphs, the situation is simple and well-understood: the same work proves a matching lower bound on the price of stability of $\Omega(\log n)$. We study the quality of go-it-alone equilibria in the case of directed graphs in Appendix B, and show that the matching lower bound on the price of stability also holds for go-it-alone equilibria.

The price of stability of fair network design games in undirected graphs is still unknown, despite substantial work on both upper and lower bounds [Fiat et al. 2006; Li 2009; Bilò et al. 2010; Christodoulou et al. 2010; Bilo and Bove 2011]. Table I provides a summary. We prove an improved upper bound of $O(\log\log\log n)$ for the Nash price of stability in broadcast games and show a even better upper bound for go-it-alone equilibria; the price of go-it-alone stability is bounded by $O(1)$ for multicast (and thus also broadcast) games. Recently, [Kawase and Makino 2012] studied the quality of the Nash equilibrium with the minimum potential in network design games and showed that the potential-optimal price of stability is between $O(\sqrt{\log n})$ and $\Omega(\sqrt{\log\log n})$. Therefore, our work, in order to break the $\log\log n$ bound, first shows that the quality of the best Nash equilibrium can be asymptotically better than the quality of the potential-optimal Nash equilibrium.

[Charikar et al. 2008] consider an online version of multicast games where players arrive one by one and connect greedily to the sink. They show that the worst social cost after all players' arrival (not necessarily a Nash equilibrium), after best-response dynamics (a Nash equilibrium), and after interleaved arrivals and deviations are $O(\log^2 n), O(\log^3 n)$, $O(polylog(n)\sqrt{n})$, respectively. An online version of go-it-alone equilibria is rather straightforward, as each player has no incentive to deviate after her arrival, when she chose a path as efficient as the shortest path.

Strong equilibria, which resist deviations by coalitions of players, have also been studied in network design games. [Albers 2008] studied the price of anarchy of strong equilibria, giving an upper bound of $O(\log n)$ and a lower bound of $\Omega(\sqrt{\log n})$ in undirected graphs. While Nash strong equilibria need not exist, [Epstein et al. 2007] prove existence of a strong equilibrium under topological constraints. We discuss strong go-it-alone equilibria in Appendix C, where we study their existence and present a full characterization of the corresponding prices of anarchy and stability.

Arbitrary cost-sharing in network design games has also been considered (e.g., [Moulin 2009], [Hoefer 2010]).

Our notions of (strong) go-it-alone equilibria reflect some features of the *core* in cooperative game theory, where the goal is to centrally find a cost allocation to each player so that no coalition of players has incentive to deviate by satisfying their requirements by themselves. [Hoefer 2010] studied the relationship between strong Nash equilibria and the core in games with arbitrary cost sharing, and our model, somewhere in between, might be viewed as a bridge between competitive and cooperative game theory. Analogies can also be drawn to the literature on conjectural variations, as GE may be enforced by anticipation of punishment in future rounds of play.[1]

## 3. GO-IT-ALONE EQUILIBRIA

Since the set of go-it-alone equilibria contains all Nash equilibria, the GPoS can only be lower than the PoS, while the GPoA can only be greater than the PoA. It might be desirable for an equilibrium relaxation to have small such differences, i.e., that its price of stability and price of anarchy are not far from their Nash equilibrium counterparts. We show that for go-it-alone equilibria in network cost sharing games, this is indeed the case. For all such quantities for which there is a known tight asymptotic value for Nash equilibria, the same asymptotic value holds for go-it-alone equilibria. This correspondence even holds true for the strong equilibrium variants, which are robust to deviations of a group of players. Since the main case of interest is undirected graphs, we present the go-it-alone equilibrium results for directed graphs and strong equilibria in Appendices B and C.

In undirected graphs, we first see that considering go-it-alone equilibria does not push the price of anarchy beyond the (already quite high) bounds known for Nash equilibria. As

[1]Our thanks to Federico Echenique for suggesting this connection.

noted in Table I, lower bounds of $n$ on the Nash price of anarchy carry over immediately to go-it-alone equilibria; here we prove the matching upper bound.

THEOREM 3.1. *The go-it-alone price of anarchy can never exceed $n$.*

PROOF. Assume by way of contradiction that there is a go-it-alone equilibrium whose cost is more than $n$ times the cost of the social optimum. This implies the existence of a player who pays more than the cost of the social optimum, but that player could strictly reduce her cost by building the entire subgraph corresponding to the social optimum by herself. □

Before we turn to the study of the one remaining quantity—the price of stability in undirected cost sharing games—we briefly discuss the issue of equilibrium computation. [Anshelevich et al. 2004] proved existence of a Nash equilibrium in every network design game via a potential function argument, but unfortunately they show that best-response dynamics can take exponentially many steps to converge. First, we see that computing a go-it-alone equilibrium of minimum social cost is also $NP$-hard.[2] The proof, which is a reduction from set cover inspired by an exercise in [Vazirani 2001], appears in Appendix A.

THEOREM 3.2. *Computing the best go-it-alone equilibrium, even in single-source games, is $NP$-hard in both directed and undirected graphs. Furthermore, in directed graphs, it is $NP$-hard to approximate within an $O(\log n)$ factor.*

This is perhaps unsurprising, as it is well known that the Steiner tree problem is $APX$-hard in undirected graphs [Bern and Plassmann 1989], and a proof similar to that of the above theorem shows that it is unlikely to have a better approximation guarantee than $O(\log n)$ in directed graphs. (this was the original statement of the exercise in [Vazirani 2001]). However, these problems have been extensively studied and have known good approximations.[3] Therefore, it is natural to ask if it is possible to convert the socially efficient solutions found by these algorithms to go-it-alone equilibria, without losing too much efficiency. One intuitive way to accomplish this is via best-response dynamics: if there exists a player who pays more than the cost of the shortest path, let her deviate to build the shortest path. Conveniently, for go-it-alone equilibria, each player will make such a deviation at most once.

In Theorem 3.4, we use this idea to show that the go-it-alone price of stability is $O(1)$ in multicast and broadcast games, giving a polynomial-time algorithm that computes such a go-it-alone equilibrium. Note that this is the only case where we achieve a better result for go-it-alone equilibria than is currently known for Nash equilibria. While it is difficult to bound the resulting cost in terms of the social optimum when players are allowed to deviate in an arbitrary order, our approach works by (1) identifying a simple class of graphs on which it is sufficient to prove our theorem, (2) scheduling a chain of beneficial deviations to follow each deviating player and recursively treating the remaining subgraphs, and (3) introducing a charging scheme wherein the cost of each deviation is charged to the edges in the social optimum. The equilibrium we find does not satisfy the criterion that the edges which are used at least once should form a tree, but it does satisfy a slightly relaxed criterion that it only uses edges from two trees—the social optimum and the shortest path tree.

We first show that a slightly more strict result on a certain class of graphs actually suffices to prove the theorem. We call a graph $G = (V, E)$ as *line-shortcut* if $V = \{t = s_0, s_1, ..., s_n\}$,

[2]Note that this is not a direct corollary of the $NP$-hardness of the Steiner tree problem, because we want to compute not the globally best Steiner tree, but the best Steiner tree among go-it-alone equilibria.
[3]There is a 1.55-approximation algorithm in undirected graphs [Robins and Zelikovsky 2000], and a $O(n^{\epsilon})$-approximation algorithm in directed graphs for any $\epsilon > 0$ [Charikar et al. 1998]. Even in the Steiner Forest problem, which finds the social optimum of general multicommodity network design games, there is a 2-approximation algorithm in undirected graphs [Goemans and Williamson 1992] and a $O(n^{1/2+\epsilon})$-approximation algorithm for any $\epsilon > 0$ [Chekuri et al. 2008].

where $t$ is the common sink and each $s_i$ with $i > 0$ is the source of one player, and $E = \{(s_0, s_i) \mid 1 \leq i \leq n\} \cup \{(s_{i-1}, s_i) \mid 2 \leq i \leq n\}$. We also require that the edge costs must satisfy $c(s_{i-1}, s_i) \leq c(s_0, s_i) \leq \sum_{j=1}^{i} c(s_{j-1}, s_j)$ to ensure that the *line* $(s_0, s_1, ..., s_n)$ is the shortest spanning tree but each vertex has a direct *shortcut* to the sink (for $s_1$, the shortcut and the path to the sink on the line coincide). This class of graphs is also used in the lower bound constructions in [Fiat et al. 2006] and [Bilò et al. 2010], so our positive result on these graphs suggests that current Nash equilibrium lower bound techniques cannot easily be extended. Two examples are illustrated in Figure 2. Although line-shortcut graphs are quite simple, for go-it-alone equilibria a slightly more strict result on line-shortcut graphs will imply the theorem on general graphs.

LEMMA 3.3. *Suppose that for any undirected broadcast game on a line-shortcut graph, there is a polynomial time algorithm to compute a go-it-alone equilibrium $P$ with $c(P) = O(1)c(T^*)$ such that each player uses each edge at most once. We can use this to efficiently compute a go-it-alone equilibrium that is within a constant factor of optimal in a general graph (where in the resulting GE each player uses each edge at most three times).*

The full proof appears in Appendix D.1. Intuitively, given an instance of an undirected multicast game, we construct the corresponding line-shortcut graph as follows: the line corresponds to the preorder traversal of the (approximate) optimal Steiner tree, and each shortcut corresponds to the shortest path from the corresponding player to the sink. Then, the total cost of the line is at most two times the cost of the Steiner tree and the cost of each shortcut is preserved. Each edge in the original graph corresponds to at most three edges (two for the preorder traversal, one for the shortest path tree) in the line-shortcut graph, explaining the slightly increased number of times that each edge can be used.

Therefore, we can concentrate on line-shortcut graphs. Algorithm 1 finds a GE with total cost only a constant worse than a given strategy, on line-shortcut graphs. In the initial strategy $P$ where each player uses the path on the main line to get to the sink, each edge $(s_{k-1}, s_k)$ is used by exactly $n - k + 1$ players ($k$, ..., $n$). We find the *unsatisfied* (i.e., paying more than the cost of her shortcut) player $i$ with the largest index and make her deviate using her shortcut. Before $i$'s deviation, $c(s_0, s_i) < c_i(P) = \sum_{k=1}^{i} \frac{c(s_{k-1}, s_k)}{n-k+1}$ and $c(s_0, s_j) \geq c_j(P) = \sum_{k=1}^{j} \frac{c(s_{k-1}, s_k)}{n-k+1}$ for all $j > i$. All players $j > i$ then *follow* $i$'s deviation (i.e. $P_j \leftarrow (s_j, ..., s_i, s_0)$). They will still be satisfied, since following $i$'s shortcut further reduces their costs to $[\frac{c(s_0, s_i)}{n-i+1} + \sum_{k=i+1}^{j} \frac{c(s_{k-1}, s_k)}{n-k+1}] < c_j(P)$.

If $i \leq \lfloor n/2 \rfloor + 1$, the remaining graph consisting of nodes that have not deviated will have at most $\lfloor n/2 \rfloor$ vertices; we process it recursively. Otherwise, we divide the remaining nodes into two smaller line-shortcut graphs, with vertices $s_0, ..., s_{\lfloor n/2 \rfloor}$ and $s_0, s_i, s_{i-1}, ..., s_{\lfloor n/2 \rfloor + 1}$. Note that the second graph is *reversed*, going from right to left, even though it is a legitimate line-shortcut graph. We process these two graphs recursively. Figure 2 illustrates the two cases. Clearly, the procedure will terminate with all players satisfied. Furthermore, every player $j$'s strategy consists of (a path on the main line from $s_j$ to $s_i$ for some $i$) + $(s_i, s_0)$, so no player uses any edge more than once.

Whenever a player deviates to her shortest path, the cost of her shortest path is less than the amount she pays in the current strategy profile, which is simply a linear combination of the costs of the edges on her current path. Each coefficient in the linear combination can be thought of as debt that the corresponding edge has to pay in order to add a new shortest path. If we accumulate the debts of the edges in the (approximate) social optimum as we add shortest paths, the total cost of the added shortest paths is bounded by the cost of the social optimum times the maximal debt of each edge. The recursive splits with reversals mentioned above are carefully orchestrated to prevent the debt on any one connection to the root from increasing too quickly. The full proof appears in in Appendix D.2.

**ALGORITHM 1:** Divide-by-half

**Data**: Undirected broadcast game $(G, A, c)$ where $G$ is a line-shortcut graph with main line $(t = s_0, s_1, ..., s_n)$ and associated initial strategy $P$ s.t. $\forall i, P_i = (s_i, s_{i-1}, ..., s_0)$
**Result**: Strategy $P$ at go-it-alone equilibrium
Let $\alpha = \lfloor |A|/2 \rfloor$.
**if** *all players in* $A$ *are satisfied given* $P$ **then return**
**else**
    Let $i$ be the unsatisfied player in $A$ with largest index. `// Note` $c_i(P) > c(s_i, s_0)$
    $P_i \leftarrow (s_i, s_0)$
    **for** $j = i+1, ..., n$ **do** $P_j \leftarrow (s_j, s_{j-1}, ..., s_i) + P_i$
    **if** $i \leq \alpha + 1$ **then**
        Divide-by-half$(G, \{1, ..., i-1\}, c)$
    **end**
    **else**
        **for** $j = \alpha+1, ..., i-1$ **do** $P_j \leftarrow (s_j, s_{j+1}, ..., s_i) + P_i$
        Divide-by-half$(G, \{1, ..., \alpha\}, c)$
        Divide-by-half$(G, \{i, i-1, ..., \alpha+1\}, c)$ `// direction of the line is reversed`
    **end**
**end**

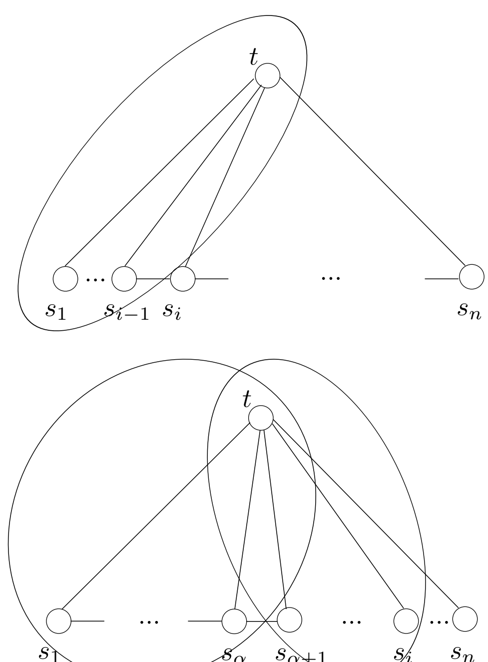


Fig. 2. Two situations depending on $i$. A component in each circle represents a line-shortcut graph recursively processed.

THEOREM 3.4. *For any undirected multicast game, there exists a go-it-alone equilibrium* $P$ *where* $c(P) = O(1)c(T^*)$ *and each player uses the same edge no more than three times. This equilibrium can be computed in polynomial time.*

## 4. $O(\text{LOG LOG LOG } N)$ UPPER BOUND ON NASH POS

Motivated in part by our result for go-it-alone equilibria, we present an upper bound of $O(\log\log\log n)$ on the price of stability in undirected broadcast games, improving the $O(\log\log n)$ bound of [Fiat et al. 2006]. Their approach is based on the idea of carefully scheduling a sequence of player deviations. They classify each edge $e_v$ outside the social optimum (and the player $v$ who initially uses it) into one of two categories: *crowded* if there are more than $\log n$ players in a circle around $v$ of radius proportional to the cost of $e_v$, or *light* otherwise. In the final Nash equilibrium, the total cost of the crowded edges can be bounded by $O(1)c(T^*)$ using a relatively simple potential argument, supporting the intuition that the crowded edges must be sparse—especially when they are expensive. Then, one can draw a circle around each light vertex $v$ of radius proportional to $c(e_v)$. They observe that if one circle contains the center of another, the radius of the latter is at most half of the radius of the former. We make an improvement to this key fact, showing that such a containment reduces the radius by factor of $O(\frac{\log h}{h})$, where $h$ indicates that the number of players in the latter circle. This proves that the worst case for [Fiat et al. 2006], that one large circle might contain many other large circles recursively, is actually impossible. This observation, with careful induction, gives us a better upper bound.

The key improvement combines two techniques to find a Nash equilibrium with low cost: taking the global potential minimizer [Anshelevich et al. 2004; Li 2009; Bilo and Bove 2011] and carefully scheduling movements [Fiat et al. 2006]. Both techniques have their limitations: the price of stability of the potential minimizer was recently proved to be $\Omega(\sqrt{\log\log n})$ even in broadcast games [Kawase and Makino 2012], and the scheduler of [Fiat et al. 2006] resorts to arbitrary improving deviations at the end of each phase. Taking the strategy profile that minimizes the potential function and satisfies our structural requirement in each phase allows us to exploit having a low strategy profile as well as maintaining the structure.

Our algorithm also has an additional feature, which is to consider a sequence of *group moves* (wherein groups of players deviate simultaneously) that reduce the potential. Group moves have previously only been considered in the case of three players [Bilo and Bove 2011], where the analysis was already quite complex. Our result suggests that this additional sophistication, combined with the ability to select a strategy profile having desired properties, might be a promising approach to finding the true price of stability.

The proof is divided into two parts. In what follows, we prove the key observation—limiting the number of circles surrounding light vertices that can appear within the circle of any light vertex—using the two techniques mentioned above. The second part, where we use nontrivial induction to prove the result based on the key observation, appears in Appendix F.

*Limiting the crowding of light vertices.* Given an instance of the undirected broadcast game $(G, A, c)$, we work with the graph $\overline{G}$ as defined by [Fiat et al. 2006], where each edge in the optimal spanning tree $T^*$ is replaced by the two edges, colored red and blue, each with the same cost as the original edge. As in [Fiat et al. 2006], given a strategy profile $P$, an improving move (deviation) of a player $u$ to a new strategy $P'_u$ is called

— **EE** (Existing Edges) : $P'_u \subseteq T(P)$ and all edges in $P'_u$ are used in the same direction as in $P$.
— **OPT** : $P'_u \subseteq T(P) \cup T^*$ but $P'_u \not\subseteq T(P)$ and all edges in $P'_u \cap T(P)$ are used in the same direction as in $P$.
— $\overline{\textbf{OPT}}$: The first edge $(u, v) \notin T^* \cup T(P)$ but the other edges are in $T(P)$ and are used in the same direction as in $P$.

Lemma 2 of [Fiat et al. 2006] shows that $P$ is a Nash equilibrium if no **EE**, **OPT**, and $\overline{\textbf{OPT}}$ moves are possible. Their scheduler to find a desired Nash equilibrium consists of 4 steps per phase, with the last step consisting of arbitrary **EE** and **OPT** moves until no more are possible. We reintroduce the definitions of the first three steps here, as well as some additional notation. At the beginning of each phase, only one of the copies of an edge in $T^*$ is used; we assume that it is red. Let $z_u = c(u,v)$, and $D_u = \left\{w \mid d_{T^*}(u,w) \leq \frac{z_u}{4}\right\}$.

— $\overline{\textbf{OPT}}$**-move**: This step starts with some player $u$ changing her strategy by an improving $\overline{\textbf{OPT}}$ move. Let $e_u = (u,v) \notin T^* \cup T(P)$ be the first edge. As in [Fiat et al. 2006], we say $u$ is crowded if $|D_u| \geq \log n$, and light otherwise.
— **OPT-loop** : Start a breadth first search of $T^*$ from $u$ reaching each player $w$ in increasing order of $d_{T^*}(u,w)$. If (path from $w$ to $u$ in $T^*$ using blue edges) + (strategy of $u$) is improving, make $w$ deviate to follow $u$. Otherwise, truncate the breadth first search. Denote by $D'_u$ the set of players who changed their strategy in this step. [Fiat et al. 2006] prove that $D_u \subseteq D'_u$.
— **EE-loop** : For each player $w \in D_u$, let $M_w$ be the subset of descendants of $w$ in the tree $T(P)$ rooted at $t$, such that $x \in M_w$ if and only if $x \notin D_u$ and $w$ is the first player in $D_u$ along the path from $x$ to $t$ in $T(P)$. We traverse the vertices in $\cup_{w \in D_u} M_w$. For each player $x \in M_w$, make her deviate using (path from $x$ to $w$ in $T(P)$) + (strategy of $w$ following $u$). [Fiat et al. 2006] prove that this move is improving for each $x$.

We define one more move, **GroupOpt**, which stands for Group OPT move and refers to the strategy profile after simultaneous deviations of a group of players. Formally, **GroupOpt**$(P,x,w)$ is the strategy profile where each player $y$ whose strategy in $P$ contained $x$ changes her strategy to (path from $y$ to $x$ in $P$) + (path from $x$ to $w$ using blue edges in $T^*$) + (path from $w$ to $t$ in $P$). Note that every edge $e \in T(\textbf{GroupOpt}(P,x,w)) - T^*$ is used in the same direction as it was in $P$. In the algorithm, **GroupOpt** is executed only when it reduces the potential $\Phi(P) = \sum_{e \in E} c(e)H(n_e)$.

**ALGORITHM 2:** Find-NE

**Data**: Strategy profile $P$ where no **EE** or **OPT** move is possible
**if** *no* $\overline{\textbf{OPT}}$ *move is possible* **then** output $P$.
**else**
  Perform $\overline{\textbf{OPT}}$**-move**, **OPT-loop**, **EE-loop** steps defined above, resulting in $P'$.
  Let $u'_1, ..., u'_m$ be the vertices such that $u'_i$ is associated with $(u'_i, v'_i)$, $u \neq u'_i$, $d_{T^*}(u, u'_i) \leq \frac{z_{u'_i}}{8}$, and $u'_i$ is not contained in $P_v$.
  **for** $j = 1, ..., m$ **do**
    **if GroupOpt**$(P', u'_j, u)$ *decreases the potential* **then** $P' \leftarrow$**GroupOpt**$(P', u'_j, u)$.
  **end**
  Update $P$ to be the strategy profile `// no` **EE** `or` **OPT** `move is possible in` $P$
  of minimum potential subject to $T(P) \subseteq T(P'_i) \cup T^*$, with every $e \in T(P) \setminus T^*$ used in the same direction as in $P'$.
  Find-NE$(P)$
**end**

Let **GOPT-loop** refer to the loop in the above algorithm where each move is defined by **GroupOpt**. To prevent the propagation of constants in formulas, let $\alpha = 8$, so that for each $i$, $u$ is contained in the circle of radius $\frac{z_{u'_i}}{\alpha}$ around $u'_i$. After **OPT-loop** and **EE-loop**, we can classify every vertex $x$ into one of three categories:

— $x \in D'_u$: $x$'s strategy consists of (the path from $x$ to $u$ using blue edges) + $e_u$ + $P_v$.
— $x \in M_w$ for some $w \in D'_u$: $v$'s strategy consists of (the prefix of $P_x$ from $x$ to $w$) $+P'_w$.

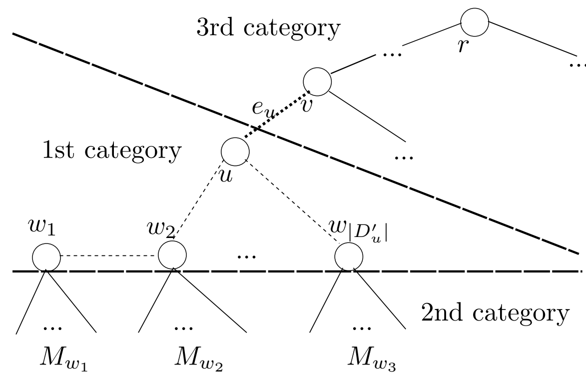


Fig. 3. Situation after **OPT-loop** and **EE-loop**. The solid lines represent edges of $T(P)$. In the first category, the only edges used (represented as thin dashed lines) are blue edges in $T^*$.

— otherwise: $P'_x = P_x$ since $x$ did not change her strategy during **OPT-loop** and **EE-loop**.

Figure 3 illustrates the situation after **OPT-loop** and **EE-loop**. Each $w \in D'_u$ has $M_w$ as descendants in solid lines as all of them will follow her during **EE-loop**. Note that every vertex contained in $P_v$ is in the third category; if $v'$ were an ancestor of $v$ in the tree $T(P)$ rooted at $t$ but a descendant of $w$ for some $w \in D'_u$, it would mean that $w$ is also an ancestor of $u$, and using $(u, w) \in T^*$ could have been another possible improving move for $u$, which contradicts the assumption that no **EE** or **OPT** move was possible in $P$.

We show that if there is another light vertex $u'$ such that $d_{T^*}(u, u') \leq \frac{z_{u'}}{\alpha}$, $z_u$ is significantly less than $z_{u'}$, and furthermore the ratio is proportional to $O(\frac{\log |D_u|}{|D_u|})$. This fact ensures that a player with many *neighbors* tends to use a cheaper edge to deviate, which fits common sense and helps bound the social cost of the final equilibrium. Such $u'$ are the ones considered in the **GOPT-loop** $(u'_1, .., u'_m)$ or the ones contained in $P_v$ (ancestors of $v$ in $T(P)$). In the latter case, this fact can be shown relatively easily since $u'$ is strictly closer to $t$ than $v$. Consider the situation where $u$ and the players who followed her stop using $e_u$ and start to follow $u'$ via edges in $T^*$. This group move might reduce the potential if $z_u$ is comparable to $d_{T^*}(u, u') \leq \frac{z_{u'}}{\alpha}$. However, in the previous phase we took $P$ to be the potential minimizer using $T(P) \cup T^*$, so the outcome of this new group move, which does not use $e_u$ and is contained in $T(P)$, cannot have a lower potential than $P$. This shows that $z_u$ must be significantly less than $d_{T^*}(u, u')$ and $z_{u'}$.

The proof of these lemmas can be found in Appendix E.

LEMMA 4.1. *Let $u'$ be a light vertex associated with $(u', v')$, $u \neq u'$, $d_{T^*}(u, u') \leq \frac{z_{u'}}{\alpha}$, and $u'_i$ contained in $P_v$ (thus, $u'$ is not among the $u'_1, ..., u'_m$ considered in the algorithm). If $|D_u| \geq 20$, $z_u \leq z_{u'} \frac{3H(|D_u|)}{\alpha |D_u|}$*

When $u'$ is not an ancestor of $v$ in $T(P)$, none of them can be said to be strictly closer to $t$, and $u' \in \{u'_1, ..., u'_m\}$ was considered in **GroupOpt-loop**. In addition to considering the group move from $u$ to $u'$ as in the previous lemma, we also consider the group move from $u'$ to $u$ and execute it if it reduces the potential. If this is possible, $e_{u'}$ will be completely unused. Otherwise, we have two seemingly contradictory situations happen (both group moves cannot reduce the potential), and the only possibility is to have $z_u$ far less than $z_{u'}$, as required.

LEMMA 4.2. *For each $u' \in \{u'_1, ..., u'_m\}$, suppose $P'$ was not replaced by* **GroupOpt**$(P', u', u)$. *If $|D_u| \geq 20$, $z_u \leq \frac{6H(|D_u|)}{|D_u|} z_{u'}$.*

These two lemmas together show the first part of the theorem.

LEMMA 4.3. *Let $N$ be the Nash equilibrium obtained by this algorithm. We can conclude that if two light vertices $u$ and $u'$ such that $d_{T^*}(u, u') \leq \frac{z_{u'}}{\alpha}$, $z_u \leq \min(\frac{4}{\alpha}, \frac{6H(|D_u|)}{|D_u|})z_{u'}$.*

The second part of the theorem, where we utilize this key improvement to prove the theorem, appears in Appendix F. Roughly, let $A[n']$ be the maximal sum of the radii of $n'$ circles, given that all of their centers are contained in the big circle of radius 1. By efficiently partitioning the big circle and assigning each portion to a subcircle that is not contained in any of its peers, we obtain a recurrence relation on $A$ that will lead us to show $A[n'] = O(\log \log n')$. Since the circle around a light vertex can have at most $\log n$ circles in it, the total cost of the light edges is $O(\log \log \log n)c(T^*)$. Together with the upper bound of $O(1)c(T^*)$ on the cost of the crowded edges achieved in [Fiat et al. 2006], this proves our main theorem regarding the Nash price of anarchy on broadcast games:

THEOREM 4.4. *The price of stability in undirected broadcast games is $O(\log \log \log n)$.*

## 5. CONCLUSION

The Nash price of stability for undirected graphs remains tantalizingly unresolved. For multicast and broadcast games, our results give two hints that the true Nash price of stability may be $\Theta(1)$. First, we have shown that the current lower bounds, which rely on "go it alone" deviations, cannot be extended beyond a constant, because the go-it-alone price of stability is bounded by $O(1)$. Furthermore, we suspect that our $O(\log \log \log n)$ bound on the Nash price of stability in broadcast games will not be the last word. Our approach is only the second centralized algorithm in the literature and is the first to show that the quality of a deliberately computed equilibrium is asymptotically better than that of the potential-minimizer. There may yet be more leverage to gain from centralized algorithms, both in the broadcast setting and in more general multicommodity games.

The idea that a limited centralized authority might prevent deviating players from free riding on compliant players is not specific to network design games, nor to fair cost sharing. We hope that the "go-it-alone equilibrium" approach we propose here may also serve as an interesting relaxation concept in other domains where studying Nash equilibria directly seems intractable.

## ACKNOWLEDGMENTS

The authors would like to thank Federico Echenique, Amos Fiat, Bobby Kleinberg, and Rob van Stee for their useful suggestions and comments.

## REFERENCES

ALBERS, S. 2008. On the value of coordination in network design. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '08. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 294–303.

ANSHELEVICH, E. AND CASKURLU, B. 2009a. Exact and approximate equilibria for optimal group network formation. In *Algorithms - ESA 2009*, A. Fiat and P. Sanders, Eds. Lecture Notes in Computer Science Series, vol. 5757. Springer Berlin Heidelberg, 239–250.

ANSHELEVICH, E. AND CASKURLU, B. 2009b. Price of stability in survivable network design. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*. SAGT '09. Springer-Verlag, Berlin, Heidelberg, 208–219.

ANSHELEVICH, E., DASGUPTA, A., KLEINBERG, J., TARDOS, E., WEXLER, T., AND ROUGHGARDEN, T. 2004. The price of stability for network design with fair cost allocation. In *Proceedings of 45th Annual IEEE Symposium on Foundations of Computer Science*. FOCS '04. 295 – 304.

ANSHELEVICH, E., DASGUPTA, A., TARDOS, E., AND WEXLER, T. 2003. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*. STOC '03. ACM, New York, NY, USA, 511–520.

Bern, M. and Plassmann, P. 1989. The steiner problem with edge lengths 1 and 2. *Information Processing Letters 32,* 4, 171 – 176.

Bilo, V. and Bove, R. 2011. Bounds on the price of stability of undirected network design games with three players. In *Journal of Interconnection Networks*. Vol. 12. 1–17.

Bilò, V., Caragiannis, I., Fanelli, A., and Monaco, G. 2010. Improved lower bounds on the price of stability of undirected network design games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*. SAGT '10. Springer-Verlag, Berlin, Heidelberg, 90–101.

Charikar, M., Chekuri, C., Cheung, T.-y., Dai, Z., Goel, A., Guha, S., and Li, M. 1998. Approximation algorithms for directed steiner problems. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '98. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 192–200.

Charikar, M., Karloff, H., Mathieu, C., Naor, J. S., and Saks, M. 2008. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th Annual Symposium on Parallelism in Algorithms and Architectures*. SPAA '08. ACM, New York, NY, USA, 70–76.

Chawla, S., Roughgarden, T., and Sundararajan, M. 2006. Optimal cost-sharing mechanisms for steiner forest problems. In *Proceedings of the 2nd International Conference on Internet and Network Economics*. WINE '06. Springer-Verlag, Berlin, Heidelberg, 112–123.

Chekuri, C., Even, G., Gupta, A., and Segev, D. 2008. Set connectivity problems in undirected graphs and the directed steiner network problem. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '08. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 532–541.

Christodoulou, G., Chung, C., Ligett, K., Pyrga, E., and van Stee, R. 2010. On the price of stability for undirected network design. In *Proceedings of the 7th International Conference on Approximation and Online Algorithms*. WAOA'09. Springer-Verlag, Berlin, Heidelberg, 86–97.

Epstein, A., Feldman, M., and Mansour, Y. 2007. Strong equilibrium in cost sharing connection games. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. EC '07. ACM, New York, NY, USA, 84–92.

Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., and Shabo, R. 2006. On the price of stability for designing undirected networks with fair cost allocations. In *Proceedings of the 33rd International Conference on Automata, Languages and Programming*. ICALP '06. Springer-Verlag, Berlin, Heidelberg, 608–618.

Goemans, M. X. and Williamson, D. P. 1992. A general approximation technique for constrained forest problems. In *Proceedings of the 3rd Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '92. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 307–316.

Hoefer, M. 2006. Non-cooperative tree creation. In *Proceedings of the 31st International Conference on Mathematical Foundations of Computer Science*. MFCS '06. Springer-Verlag, Berlin, Heidelberg, 517–527.

Hoefer, M. 2010. Strategic cooperation in cost sharing games. In *Proceedings of the 6th International Conference on Internet and Network Economics*. WINE '10. Springer-Verlag, Berlin, Heidelberg, 258–269.

Kawase, Y. and Makino, K. 2012. Nash equilibria with minimum potential in undirected broadcast games. In *Proceedings of the 6th International Conference on Algorithms and Computation*. WALCOM '12. Springer-Verlag, Berlin, Heidelberg, 217–228.

Li, J. 2009. An upper bound on the price of stability for undirected shapley network design games. *Information Processing Letters 109,* 15, 876 – 878.

Moulin, H. 2009. Pricing traffic in a spanning network. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. EC '09. New York, NY, USA, 21–30.

Robins, G. and Zelikovsky, A. 2000. Improved steiner tree approximation in graphs. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '00. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 770–779.

Vazirani, V. V. 2001. Approximation algorithms. Springer.

# Online Appendix to: Improved Bounds on the Price of Stability in Network Cost Sharing Games

EUIWOONG LEE, Carnegie Mellon University
KATRINA LIGETT, California Institute of Technology

## A. COMPUTATION OF GO-IT-ALONE EQUILIBRIA

THEOREM A.1. *Computing the best Go-it-alone equilibrium, even in single-source games, is NP-hard in both directed and undirected graphs. Furthermore, in directed graphs, it is NP-hard to approximate within an $O(\log n)$ factor.*

PROOF.
We reduce an instance of the set cover problem to a single-source network design game. Given universe of elements $U = \{u_1, ..., u_n\}$ and a collection of subsets $C = \{S_1, ..., S_m\}$, our underlying graph for the network design game is $G = (U \cup C \cup \{t\}, E)$. We include $(u_i, S_j) \in E$ if and only if $u_i \in S_j$, and $(S_j, t) \in E$ for all $1 \leq j \leq m$. We have $n$ players $1, ..., n$, where player $i$'s source is $u_i$ and destination is $t$. The construction of $G$ is shown in Figure 4.

If $G$ is directed, the cost of $(u_i, S_j) \in E$ is 0 for each $i$ and $j$, and the cost of $(S_j, t) \in E$ is 1 for each $j$. Any strategy profile where each player $i$ connects $u_i$ to $t$, which corresponds to a feasible set cover, is a Go-it-alone equilibrium, because every shortest path from $u_i$ to $t$ costs 1, and each player does not pay more than that in any strategy profile since every path from $u_i$ to $t$ consists of only two edges $(u_i, S_j), (S_j, t)$ for some $S_j$. The cost of any strategy profile is exactly the number of $(S_j, t)$'s used in this profile, so there is a set cover of cost $k$ if and only if there is a Go-it-alone equilibrium of cost $k$. Therefore, we have a gap-preserving reduction from the set cover problem, and the hardness of $O(\log n)$-approximation follows.

If $G$ is undirected, the cost of $(u_i, S_j) \in E$ is $1 + \epsilon$ for each $i$ and $j$, and the cost of $(S_j, t) \in E$ is 1 for each $j$, where $\epsilon > 0$ is a small constant. For each set cover, its corresponding strategy profile, where each player is only connected to one set and each used set is connected to $t$, is a Go-it-alone equilibrium as in directed graphs. Note that not every Go-it-alone equilibrium is guaranteed to correspond to a feasible set cover, because some $(S_j, t)$ may not be in the strategy profile even though $S_j$ is visited by some player; $(u_{i_1}, S_{j_1}), (S_{j_1}, u_{i_2}), (u_{i_2}, S_{j_2}), (S_{j_2}, t)$ is a possible path from $u_{i_1}$ to $t$, but it does not guarantee that $u_{i_1}$ is covered by some $S_j$. However, we want to compute the best Go-it-alone equilibrium, and we can show that it corresponds to the best set cover. Let $P^*$ be the socially best Go-it-alone equilibrium and suppose $T(P^*)$ contains $m'$ $S_j$'s as Steiner vertices. Note that $m'$ should be at least the size of the smallest set cover, since each player $i$ should move to some set containing $e_i$ via her first edge. Since the number of edges in $T(P^*)$ is at least $(n + m' + 1) - 1 = n + m'$, $C(P^*) \geq n + m'$, so $P$ cannot be the best Go-it-alone equilibrium unless $m'$ is the size of the smallest set cover. Once we have a fixed $m'$, having all the chosen $S_j$'s connected to $t$ is the best strategy, since the $c(S_j, t)$ is slightly cheaper than $c(e_i, S_j)$. Therefore, $P^*$ is the strategy profile where each $e_i$ is connected exactly one $S_j$ and $(S_j, t)$ is used for each $S_j \in V(T(P^*))$, which corresponds to the best set cover. □

## B. GO-IT-ALONE EQUILIBRIA IN DIRECTED GRAPHS

Theorem 3.1 also serves to settle the Go-it-alone Price of Anarchy in Directed Graphs.

For the Go-it-alone Price of Stability, we have the following:


DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

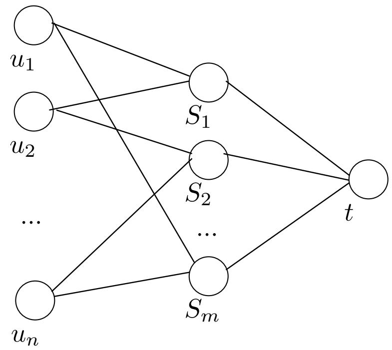


Fig. 4. A reduction from an instance of the set cover problem to a single-source network design game. This partial information shows that $\{u_1, u_2\} \subseteq S_1, \{u_2, u_n\} \subseteq S_2, \{u_1, u_n\} \subseteq S_m$.

THEOREM B.1. *The Go-it-alone price of stability in directed graphs is* $\Omega(\log n)$.

PROOF. The example for Nash equilibria introduced in [Anshelevich et al. 2004] also works for Go-it-alone equilibria; see Figure 5. It is easy to see that $((s_1, t), (s_2, t), ..., (s_n, t))$ is the only Go-it-alone equilibrium, while the social optimum is $(s_i, v, t)$ for every player $i$. The cost of the Go-it-alone equilibrium is $1 + 1/2 + ... + 1/n = H(n) = \Omega(\log n)$, where $H(n)$ is the $n$th harmonic number. The cost of optimal solution is $1 + \epsilon$. $\square$

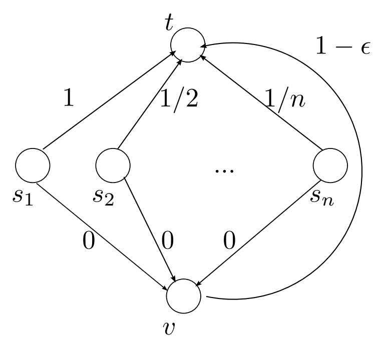


Fig. 5. An instance with price of stability $H(k)$.

## C. STRONG GO-IT-ALONE EQUILIBRIA

### C.1. Definition and Existence

Let $A' = \{i_1, ..., i_{n'}\}$ be a nonempty coalition of players. Given a strategy profile $P$, let $P_{A'}$ be the strategy profile of players $i \in A'$. $T(P_{A'}), c(P_{A'})$, and $c_i(P_{A'})$ are naturally defined from the earlier definitions, as if there were no players outside $A'$. $P$ is a *strong equilibrium*, if, for no nonempty coalition $A'$, there exists a strategy change $P'_{A'}$ such that $c_i(P'_{A'}, P_{-A}) < c_i(P)$ for all players $i \in A'$. Similarly, $P$ is a *strong Go-it-alone equilibrium*, if, for no nonempty coalition $A'$, there exists a strategy change $P'_{A'}$ on $A'$ such that $c_i(P'_{A'}) < c_i(P)$ for all players $i \in A'$. Note that for any $P$ and $P'_{A'}$, $c_i(P'_{A'}) \geq c_i(P'_{A'}, P_{-A'})$ since the former ignores possible cost sharing with players outside $A'$.

Assuming the existence of a strong equilibrium or a strong Go-it-alone equilibrium, the strong price of anarchy, the strong price of stability, the strong Go-it-alone price of anarchy, the strong Go-it-alone price of stability can be defined similarly. Let $NE, SE, GE, SGE$ be the set of Nash, strong, Go-it-alone, strong Go-it-alone equilibria respectively. It is straightforward from the definitions that every Nash equilibrium is a Go-it-alone equilibrium, and every strong equilibrium is a strong Go-it-alone equilibrium, so $NE \subseteq GE$ and $SE \subseteq SGE$.

LEMMA C.1. *There is a network design game in which the sets of equilibria corresponding to the four equilibrium concepts are pairwise distinct.*

PROOF. Figure 6 is from [Epstein et al. 2007], who showed that $P = ((e, c), (b, f))$ is the only Nash equilibrium (with $c_1(P) = c_2(P) = 5$), and that there is no strong equilibrium. Consider $P' = ((a, b, c), (b, c, d))$ where the two players cooperate. $c_1(P') = c_2(P') = 4 \leq e_1 = e_2 = 5$, so $P'$ is a Go-it-alone equilibrium. Furthermore, $P'$ is the social optimum and the unique strong Go-it-alone equilibrium. $P$ is not a SGE (or a SE), since simultaneous deviation to $P'$ benefits both players. Therefore, GE = $\{P, P'\}$, NE = $\{P\}$, SGE = $\{P'\}$, SE = $\emptyset$, which shows that all 4 notions are different. □

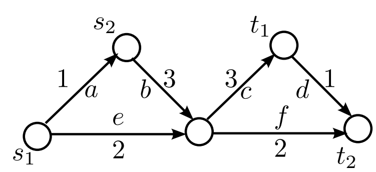


Fig. 6. A network design game where all 4 notions of equilibria are pairwise different

The existence of a Nash equilibrium is proved in [Anshelevich et al. 2004] using a potential function. It directly follows that a Go-it-alone equilibrium exists in any network design game. However, as shown in Figure 6, a strong equilibrium is not guaranteed to exist even in a game with two players where a strong Go-it-alone equilibrium exists. The next theorem shows that this was not a coincidence.

THEOREM C.2. *A strong Go-it-alone equilibrium exists in every network design game with two players.*

We present two proofs, each of which partially extends to different restricted versions of the case of $n$ players. We will state the first proof, introduce generalizations, and then give the second proof also with its generalizations.

PROOF. Let $f_i$ be the cost of the shortest path from $s_i$ to $t_i$. Let $P^* = \underset{P}{\operatorname{argmin}} \{c(P) \mid c_1(P) \leq f_1, c_2(P) \leq f_2\}$. The set $\{P \mid c_1(P) \leq f_1, c_2(P) \leq f_2\}$ is nonempty since the union of shortest paths is in this set, so $P^*$ is well defined. In $P^*$, neither player 1 nor player 2 wants to deviate by herself because $c_1(P^*) \leq f_1$ and $c_2(P^*) \leq f_2$. Furthermore, if they deviate simultaneously to $P'$ and both their costs are strictly reduced, $c(P') = c_1(P') + c_2(P') < c_1(P) + c_2(P) = c(P)$, which contradicts the choice of $P^*$. Therefore, $P^*$ is a strong Go-it-alone equilibrium. □

The definition of $P^*$ can easily be extended to the case where there are more than two players: with $n$ players, we can define $P^* = \underset{P}{\operatorname{argmin}} \{c(P) \mid c_i(P) \leq e_i, 1 \leq i \leq n\}$. In this case, $P^*$ prevents the deviation of any coalition of exactly size 1 or exactly size $n$, but it is not guaranteed to be a strong Go-it-alone equilibrium. One situation where the definition of $P^*$ can be applied to a more general model is when the collection of the possible deviating coalitions is restricted to correspond to a family tree. Formally, let $F = (V(F), E(F))$ be a rooted tree where each leaf corresponds to a player in the original game. Each possible deviating coalition $A'$ corresponds to a node $v$ in $F$: if $v$ is a leaf corresponding to player $i$, $A' = \{i\}$; otherwise, $A'$ is the set of the players that correspond to the descendants of $v$.

THEOREM C.3. *Given an instance of a network design game and tree $F$ that represents the family tree of possible deviating coalitions prevent, as above, there is a Go-it-alone equilibrium where no coalition that corresponds to a node in $F$ can deviate to strictly reduce the cost of each player.*

PROOF. Let $w$ be the root of $F$. While traversing $F$ in postorder, we can build a Go-it-alone equilibrium which does not allow the deviation of any coalition that corresponds to a node of $F$. If $v$ is the leaf that corresponds to player $i$, let $P^v$ be the shortest path from $s_i$ to $t_i$. If $v$ is an internal node with descendants $v_1, ..., v_q$ that correspond to $A_1, ..., A_q$ respectively, let

$$P^v = \operatorname*{argmin}_P \left\{ c(P) \mid \sum_{j \in A_i} c_j(P) \le c(P^{v_i}), 1 \le i \le q \right\}$$

$P^v$ is well defined because the union of subgraphs that correspond to its children gives one subgraph that satisfies the constraints.

We will argue that $P^w$ is the strategy profile that we want. Assume by way of contradiction that for some $v \in V(F)$ and the corresponding subset of players $A' \subseteq A$, they can strictly reduce each player's cost by deviating to $P_{A'}$; $\forall i \in A', c_i(P_{A'}) < c_i(P^w)$. This would imply $c(P_{A'}) < \sum_{i \in A'} c_i(P^w) \le c(P^v)$. By the choice of $P^v$, we can deduce that there is a descendant $r$ of $v$ that corresponds to $A'' \subseteq A'$ such that $\sum\limits_{j \in A''} c_j(P_{A'}) > c(P^r)$. This implies at least one player in $A''$ cannot strictly reduce her cost, which leads to contradiction. □

Although the number of coalitions whose deviations are prevented by this construction is only $|V(F)| = O(n)$ in this family tree model, one might argue that in reality deviating coalitions may be formed in a hierarchical manner such as that described here; a player or coalition might deviate if all of its *siblings* deviate to make a bigger coalition. The second proof of Theorem C.2 shows there is a Go-it-alone equilibrium that prevents the deviation of $O(n^2)$ coalitions.

PROOF. Let $f_i$ be the cost of the shortest path from $s_i$ to $t_i$. Let $P^1$ be a shortest path from $s_1$ to $t_1$. Then $c_1(P^1) = f_1$. Let $P^2 = \operatorname*{argmin}_P \{c_2(P) | c_1(P) \le c_1(P^1)\}$. The set $\{P | c_1(P) \le c_1(P^1)\}$ is nonempty since the union of shortest paths is in this set, so $P^2$ is well defined and $c_2(P^2) \le f_2$. In $P^2$, neither player 1 nor player 2 wants to deviate by herself because $c_1(P^2) \le e_1$ and $c_2(P^2) \le e_2$. Furthermore, if they deviate together to $P'$ where both of their costs are strictly reduced, $c_1(P') < c_1(P^1), c_2(P') < c_2(P^2)$, which contradicts the choice of $P^2$. Therefore, $P^2$ is a strong Go-it-alone equilibrium. □

The definition of $P^2$ can easily be extended to the case where there are more than two players. For $1 \le m \le n$, let $\Sigma_m$ be the space of strategy profiles of players $1, ..., m$. Define $P^m = \operatorname*{argmin}_{P \in \Sigma_m} \{c_m(P) | c_i(P) \le c_i(P^i), 1 \le i \le m-1\}$. Note that this definition yields the same $P^1$ and $P^2$ used in the above proof. The next theorem shows that $P^n$ prevents the deviation of any coalition which consists of *adjacent* players: $\{j, j+1, ..., k\}$ for some $j \le k$. Since there are $O(n^2)$ of such coalitions, it provides a Go-it-alone equilibrium which is more robust to possible deviations than the equilibrium defined using the previous proof. This fact is also useful in practice because in reality coalitions might easily be formed by adjacent players.

THEOREM C.4. *Given an instance of a network design game, there is a Go-it-alone equilibrium where no coalition that consists of adjacent players can deviate to strictly reduce the cost of each player.*

Proof. We claim that $P^n$ is such a Go-it-alone equilibrium. Assume by way of contradiction that $A' = \{j, j+1, ..., k\}$ is such a coalition, with strategy profile $P' = (P'_j, P'_{j+1}, ..., P'_k)$, i.e., $c_i(P') < c_i(P^n)$ for all $j \leq i \leq k$. Let $P'' \in \Sigma_k$ be the concatenation of $P^{j-1}$ and $P'$; the concatenation of a profile of players $1, ..., j-1$ and a profile of players $j, ..., k$ yields a profile of players $1, ..., k$. By the definition of $P^{j-1}$ and our assumption, $c_i(P'') \leq c_i(P^i)$ for all $1 \leq i \leq k$ with the inequality strict for $j \leq i \leq k$. However, this contradicts the choice of $P^k$ since $P^k$ minimizes $c_k(P^k)$ subject to $c_i(P^k) \leq c_i(P^i)$ over any $P \in \Sigma_k$. $P''$ satisfies all the constraints required for $P^k$, and has a lower $c_k$ value. □

These two proofs of Theorem C.2 and their simple extensions might lead one to conjecture that a strong Go-it-alone equilibrium might exist even in more complicated games. Unfortunately, this is not true, because there is a simple counterexample with three players that does not admit a strong Go-it-alone equilibrium.

Theorem C.5. *Strong Go-it-alone equilibria need not exist in games with three or more players.*

Proof. Figure 7 defines a simple game on a symmetric, planar graph. Each player has only two paths from $s_i$ to $t$, and two players may share costs if they meet at an intermediate vertex which corresponds to an edge of the triangle. Consider the strategy profile $((a, g), (b, g), (e, h))$; players 1 and 2 cooperate to share the cost of $g$, and player 3 chooses the shorter route by herself. This is not a strong Go-it-alone equilibrium since player 2 and 3 can deviate to $((c, i), (d, i))$: player 2 can reduce her cost by using a cheaper first edge, maintaining the cost sharing of the second edge. Player 3 can reduce its cost by creating new cooperation. Since this game is symmetric, every profile can be shown to not be a strong Go-it-alone equilibrium by a similar argument. Therefore, this game does not admit a strong Go-it-alone equilibrium. □

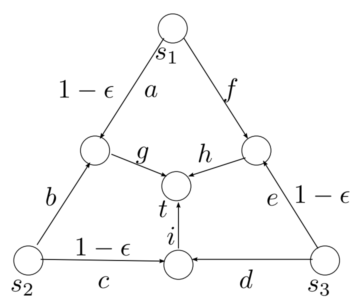


Fig. 7. An instance without a strong Go-it-alone equilibrium. The weight of each edge is 1 except $a, c, e$.

[Epstein et al. 2007] gave a topological classification of directed graphs according to the existence of a strong equilibrium. Any single-source game based on a series-parallel graph has at least one strong equilibrium, and there is a game based on a general graph that does not admit a strong equilibrium. In multicommodity cases, any game based on an extension-parallel graph (which restricts series operations to be performed only when one of the operands is a single edge) has a strong equilibrium, while a strong equilibrium may not exist even in a series-parallel graph.

Since all the existence results for strong equilibria directly imply existence of strong Go-it-alone equilibrium, it is only remains to check whether similar nonexistence results hold for strong Go-it-alone equilibria. Unfortunately, the counterexamples of [Epstein et al. 2007] that do not admit a strong equilibrium do indeed have a strong Go-it-alone equilibrium, so we need more sophisticated counterexamples. As noted above, figure 7 demonstrates nonexistence of a strong Go-it-alone equilibrium in a single-source game based on a non-series-parallel graph. However, the question of whether a strong Go-it-alone equilibrium exists in every game based on a series-parallel graph remains open. The proof techniques of [Epstein et al. 2007] cannot be directly applied to strong Go-it-alone equilibria. Their proof is based on the fact that if there are two players with the same source and the same destination, in every equilibrium these two players take the same path. However, this is not the case for strong Go-it-alone equilibria, as we show here.

THEOREM C.6. *There is an instance where two players with the same commodity take different paths in the unique strong Go-it-alone equilibrium.*

PROOF. Consider a game on the graph in Figure 8 with 6 players, where players 1 and 2 have the commodity $(s_1, t_1)$, players 3 and 4 have the commodity $(s_2, t_2)$, and players 5 and 6 have the commodity $(s_3, t_2)$. In a strong Go-it-alone equilibrium, it is obvious that players 3 and 4 have to take the same path, and players 5 and 6 have to take the same path, since otherwise they can cooperate to take $(s_2, t_2)$ or $(s_3, t_2)$ to strictly reduce their costs. Therefore, it is more convenient to think of players 3 and 4 as one player (player 3), and players 5 and 6 as one player (player 4), with their weights doubled.[4]

Consider strategy profiles where players 1 and 2 take the same path. By symmetry we need only consider the case where $P_1 = P_2 = ((s_1, a), (a, b), (b, t_1))$. It is also easy to see that in any strong Go-it-alone equilibrium either players 3 and 4 cooperate to take $(e, t_2)$ or they take their own shortcuts to $t_2$, since otherwise one of them has to pay 1000 for $(e, t_2)$ and has an incentive to take his shortcut. If players 3 and 4 cooperate, $c_1(P) = c_2(P) = 1.5, c_3(P) = 503, c_4(P) = 506$. Then player 4 can strictly reduce her cost by taking her own shortcut, so this is not a strong Go-it-alone equilibrium. If players 3 and 4 take their own shortcuts, $c_1(P) = c_2(P) = 3, c_3(P) = 505, c_4(P) = 505$. Then all players can simultaneously deviate to $((s_1, a), (a, b), (b, t_1)), ((s_1, c), (c, d), (d, t_1)), ((s_2, a), (a, b), (b, e), (e, t_2)), ((s_3, c), (c, d), (d, e), (e, t_2))$, where players 3 and 4 cooperate for the expensive edge, and players 1 and 2 take different paths to each cooperate with one of the players 3 and 4, which makes $c_1(P) = c_2(P) = 2, c_3(P) = 504, c_4(P) = 504$. This is the unique strong Go-it-alone equilibrium, which forces player 1 and 2 take different paths. □

The previous theorem supports our intuition that (strong) Go-it-alone equilibria reflect global fairness and cooperation, since the strong Go-it-alone equilibria force each player to cooperate not only with players who have the same interest (commodity), but with other players with whom cooperation will yield a socially efficient and fair solution. This unusual property makes it more difficult to prove or disprove the existence than for Nash equilibria and strong equilibria, and requires totally different proof techniques. However, as we saw in the body of the paper, Go-it-alone equilibria are more computationally tractable than their Nash counterparts, while preserving many desired properties.

We now turn to the study of the quality of strong Go-it-alone equilibria.

### C.2. Quality

The price of anarchy and the price of stability of strong equilibria have been less studied in the literature than those of Nash equilibria. We find tight but rather dull results in directed graphs; Figure 5 also serves an example to show the lower bound of $\Omega(\log n)$ on the strong price of anarchy and the strong price of stability in directed graphs. As far as we know, there is no result on the strong price of stability in undirected graphs. The only two nontrivial results on this topic are that of [Albers 2008], giving the lower bound on the strong price of

[4]See [Anshelevich et al. 2004] for the definition of a weighted network design game. Intuitively, the cost one player has to pay for an edge $e$ is now proportional to the ratio of the weight of the player to the sum of the weights of all players using $e$.

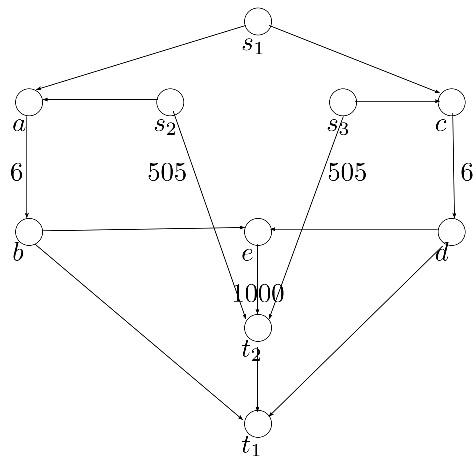


Fig. 8. An instance where two players with the same commodity take different paths in the unique strong Go-it-alone equilibrium. The weight of each unspecified edge is 0.

Table II. The results on PoA and PoS of strong equilibria and strong Go-it-alone equilibria. An asterisk indicates results proved in the present manuscript, that do not follow directly from prior work.

| | strong NE | strong GE |
|---|---|---|
| PoA(Undirected) | $O(\log n)$[Epstein et al. 2007] | $O(\log n)^*$ |
| | $\Omega(\sqrt{\log n})$[Albers 2008] | $\Omega(\sqrt{\log n})$ |
| PoA(Directed) | $O(\log n)$[Epstein et al. 2007] | $O(\log n)^*$ |
| | $\Omega(\log n)$[Anshelevich et al. 2004] | $\Omega(\log n)$ |
| PoS(Directed) | $O(\log n)$[Anshelevich et al. 2004] | $O(\log n)^*$ |
| | $\Omega(\log n)$[Anshelevich et al. 2004] | $\Omega(\log n)^*$ |

anarchy in undirected graphs, and that of [Albers 2008] and [Epstein et al. 2007], giving an upper bound on the strong price of anarchy which holds in both directed and undirected graphs.

Table II summarizes the results on strong Nash equilibria and strong Go-it-alone equilibria. Since a strong equilibrium is not guaranteed to exist in every network design game, we need to be careful when translating the results of strong equilibria to strong Go-it-alone equilibria. As before, lower bounds on the strong price of anarchy hold for strong Go-it-alone equilibria. However, upper bounds on the strong price of stability are not guaranteed to hold for strong Go-it-alone equilibria, because strong Go-it-alone equilibria may exist in an instance which does not admit a strong equilibrium, and the ratio of the cost of best equilibrium to the cost of the social optimum in this instance could be larger than the price of stability of strong equilibria.

However, this does not in fact occur In directed graphs, the upper bound on the strong price of anarchy matches the lower bound on the strong price of stability. In undirected graphs, we have no result on the price of stability other than the trivial upper bound coming from the upper bound on the price of anarchy.

Therefore, there is only one nontrivial result that needs to be translated to strong Go-it-alone equilibria, which is the upper bound on the price of anarchy. [Albers 2008] and [Epstein et al. 2007] independently proved this result, and we reproduce the proof of [Epstein et al. 2007], highlighting why it also holds for strong Go-it-alone equilibria.

THEOREM C.7. *[Epstein et al. 2007] The strong Go-it-alone price of anarchy of a network design game with $n$ players is at most $H(n)$.*

PROOF. Let $P = (P_1, ..., P_n)$ be a strategy profile. Given a subset $A'$ of the set of players, recall that $P_{A'}$ are the strategies of players $i \in A'$ and $n_e(P_{A'})$ is the number of players in $A'$ using $e$ (counting multiplicity). We denote by $\Phi(P_{A'})$ the potential function of the profile $P_{A'}$, where

$$\Phi(P_{A'}) = \sum_{e \in E} c(e) H(n_e(P_{A'}))$$

Let $P$ be a strong Go-it-alone equilibrium, and let $P^*$ be the profile of the optimal solution. We define an order on the players as follows. For each $k = n, ..., 1$, since $P$ is a strong Go-it-alone equilibrium, there exists a player in $A_k = \{1, ..., k\}$, w.l.o.g. call it player $k$, such that,

$$c_k(P) \leq c_k(P^*_{A_k})$$

where $c_k(P^*_{A_k})$ denotes the cost of player $k$ in the game with the set of players $A_k$ and the strategy profile $P^*_{A_k}$. In this way, $A_k$ is defined recursively, such that for every $k = n, ..., 2$ it holds that $A_{k-1} = A_k - \{k\}$, i.e., after the renaming, $A_k = \{1, ..., k\}$. It is easy to see that $c_k(P^*_{A_k}) \leq \Phi(P^*_{A_k}) - \Phi(P^*_{A_{k-1}})$, even considering the fact that each edge can be used multiple times. Therefore,

$$c_k(P) \leq c_k(P^*_{A_k}) \leq \Phi(P^*_{A_k}) - \Phi(P^*_{A_{k-1}})$$

The above equation is where strong Go-it-alone equilibria perfectly replace strong equilibria. The original proof says that $c_k(P) \leq c_k(P_{-A_k}, P^*_{A_k}) \leq c_k(P^*_{A_k})$, which simply ignores the contribution of the players outside $A_k$. Together with the upper bound on the price of stability of Go-it-alone equilibria, it shows that these proof techniques implicitly assumed that deviating coalitions will independently create the network that serves them, which provides additional justification of our new equilibrium concept.

Summing over all players, we obtain

$$\begin{aligned}
\sum_{i \in A} c_i(P) &\leq \Phi(P^*_A) \\
&= \sum_{e \in P^*} c_e H(n_e(P^*)) \\
&\leq \sum_{e \in P^*} c_e H(n) \\
&= H(n) \cdot OPT
\end{aligned}$$

□

While the strong Go-it-alone price of anarchy and stability are exactly the same as the strong price of anarchy and stability in directed graphs, in undirected graphs, where a strong Go-it-alone equilibrium might be strictly better than a strong equilibrium (as a Go-it-alone equilibrium is better than a Nash equilibrium), the strong price of stability has never been studied even. It would be interesting to study this problem in terms of strong (Go-it-alone) equilibria and see how they improve the price of stability.

## D. FULL PROOF OF THEOREM 3.4

### D.1. Proof of Lemma 3.3

Given an instance of an arbitrary multicast game, compute a constant approximation Steiner tree as suggested in [Robins and Zelikovsky 2000]. Traverse the tree in preorder with the sink $t$ as the root. Note that some vertices in this traversal are sources of the players while some are not, and the same vertex can appear more than once in the traversal. Name the sink and the sources $t = s_0, s_1, ..., s_n$ in the order they first appear in the traversal. The line-shortcut graph has $s_0, ..., s_n$ as the vertices, where $c(s_{i-1}, s_i)$ is the cost of the path along the traversal between the first appearances of $s_{i-1}$ and $s_i$ ($(s_{i-1}, s_i)$ in the line-shortcut graph is the path on the Steiner tree traversal, contracted to one edge). The total cost of the line is at most the two times the cost of the computed Steiner tree, which is a constant times the social optimum. For $i \geq 1$, let $c(s_0, s_i)$ in the line-shortcut graph be the cost of the shortest path from $s_i$ to $s_0$ in the original graph ($(s_0, s_i)$ in the line-shortcut graph is player $i$'s shortest path in the original graph, contracted to one edge).

Suppose we find a Go-it-alone equilibrium in the line-shortcut graph, where each player pays no more than the cost of her shortcut and the cost of the equilibrium is bounded by a constant times the cost of the social optimum in the line-shortcut graph. We can reconstruct the corresponding Go-it-alone equilibrium in the original graph, expanding each edge to a path if necessary, while respecting the direction in which it was used. For every edge in the line-shortcut graph, the number of players using the corresponding path in the original graph is at least the number of players using that edge in the line-shortcut graph, so in the original graph each player pays no more than the cost she paid in the line-shortcut graph, which is no more than the cost of her shortest path (in both line-short and original graph). Even if no edge is used by any player more than once in the line-shortcut graph, some edges can be used up to three times by the same player in the original graph, since the preorder traversal can use the same edge twice, and it might also overlap with the shortest path tree. However, these overlaps only reduce the cost of some players, as well as the total cost of the equilibrium, giving a desired Go-it-alone equilibrium in the original graph. The conversion to and from the line-shortcut graph can be done in polynomial time.

### D.2. Proof of Theorem 3.4

We wish to show that Algorithm 1 finds a Go-it-alone equilibrium in a line-shortcut graph with cost at most a constant times the cost of the social optimum.

At each recursive call with players $\{s_0, ..., s_n\}$ (indices and the number of players can be different from original graph), we bound the cost of strategy profile $P$ by $\sum_{i=1}^{m} C[i]c(s_{i-1}, s_i)$, by recursively computing $C$. Initially, we set $C[i] = 1$ for all $1 \leq i \leq n$. The cost of the initial strategy profile, the line, is then precisely $\sum_{i=1}^{n} C[i]c(s_{i-1}, s_i)$. When $i$ is the unsatisfied player with the largest index, we add $(s_0, s_i)$ to the strategy profile. However, $c(s_0, s_i) < c_i(P) = \sum_{k=1}^{i} \frac{c(s_{k-1}, s_k)}{n-k+1}$, so the cost of the added edge can be bounded by the costs of the existing edges. If $i \leq \alpha + 1$ ($\alpha = \lfloor n/2 \rfloor$), only one recursion with players $\{s_0, ..., s_{i-1}\}$ is called. Let $C'[1], ..., C'[i-1]$ be the recursively computed $C$ for that subgraph. Then, $C[k] \leftarrow C'[k] + \frac{1}{n-k+1}$ satisfies the invariant ($1 \leq k < i$). We do not need to consider $C[i]$ since $(s_{i-1}, s_i)$ gets completely unused.

If $i > \alpha + 1$, the added shortcut $(s_0, s_i)$ is used as the first edge in a reversed subgraph with players $\{s_0, s_i, s_{i-1}, ..., s_{\alpha+1}\}$, accumulating its own debt in subsequent calls in that subgraph. In that case, the edges responsible for $(s_0, s_i)$, which are $(s_{k-1}, s_k)$ for $1 \leq k \leq i$, need to pay the debt for $(s_0, s_i)$. In addition to $C'[1], ..., C'[\alpha]$ for the first subgraph, we continue to compute $C''[1], ..., C''[i - \alpha]$ of the reversed subgraph recursively, where the debt for $(s_0, s_i)$ is stored in $C''[1]$. Then $C[k] \leftarrow C'[k] + \frac{C''[1]}{n-k+1}$ for $1 \leq k \leq \alpha$, and $C[k] \leftarrow C''[i - k + 1] + \frac{C''[1]}{n-k+1}$ for $\alpha < k \leq i$ satisfies the invariant. Table III summarizes the rule to compute $C$.

If we continue to do the same *accounting* in all recursive calls, in the original line-shortcut graph after the termination of the algorithm, $c(P) \leq \sum_{i=1}^{n}(C[i]c(s_{i-1}, s_i)) \leq (\max_i C[i]) \sum_{i=1}^{n} c(s_{i-1}, s_i) \leq (\max_i C[i])c(T^*)$, so showing $C[i]$ is bounded by a constant suffices.

Table III. Update rule to compute $C$ when $(s_0, s_i)$ is added.

| | Range | Update rule |
|---|---|---|
| $i \leq \alpha + 1$ | $1 \leq k < i$ | $C[k] \leftarrow C'[k] + \frac{1}{n-k+1}$ |
| $i > \alpha + 1$ | $1 \leq k \leq \alpha$ | $C[k] \leftarrow C'[k] + \frac{C''[1]}{n-k+1}$ |
| | $\alpha < k \leq i$ | $C[k] \leftarrow C''[i-k+1] + \frac{C''[1]}{n-k+1}$ |

These updates depend on the size of a graph and the costs of edges, but we can show that $C$ is bounded by a constant for any number of players and any costs of edges. Let $B[n, j] (1 \leq j \leq n)$ be the supremum of all possible $C[j]$'s on any line-shortcut graph with $n$ vertices (over any costs of edges satisfying the definition of a line-shortcut graph). For example, $B[1, 1] = 1$, $B[2, 1] = 1.5$, $B[2, 2] = 1$ (when $n = 2$, player 2 can deviate when $c(s_0, s_2) < \frac{c(s_0,s_1)}{2} + c(s_1, s_2)$). $B[n, 1]$ can be easily bounded since the first edge always remains to be the first edge in a reduced graph with at most $\lfloor n/2 \rfloor$ vertices.

Lemma D.1. *For any $n$, $B[n, 1] < 5$*

Proof. We prove the following claim which implies the lemma: if $2^{l-1} < n \leq 2^l$,

$$B[n, 1] \leq \prod_{k=0}^{l-1} \left(1 + \frac{1}{2^k}\right)$$

$B[1, 1] = 1$ and $B[2, 1] = 1.5$, so it holds for $n = 1, 2$. For $n \geq 3$ (so $l \geq 2$), by the update rules in Table III ($C[1] \leq C'[1] + \frac{C''[1]}{n}$),

$$\begin{aligned}
B[n, 1] &\leq B[\lfloor n/2 \rfloor, 1] + \frac{1}{n} \max_{\lfloor n/2 \rfloor < i \leq n} B[i - \lfloor n/2 \rfloor, 1] \\
&\leq \prod_{k=0}^{l-2} \left(1 + \frac{1}{2^k}\right) + \frac{1}{n} \prod_{k=0}^{l-2} \left(1 + \frac{1}{2^k}\right) \\
&\leq \left(1 + \frac{1}{n}\right) \prod_{k=0}^{l-2} \left(1 + \frac{1}{2^k}\right) \\
&\leq \left(1 + \frac{1}{2^{l-1}}\right) \prod_{k=0}^{l-2} \left(1 + \frac{1}{2^k}\right) \\
&\leq \prod_{k=0}^{l-1} \left(1 + \frac{1}{2^k}\right)
\end{aligned}$$

The lemma follows from the fact that $\prod_{k=0}^{\infty} \left(1 + \frac{1}{2^k}\right) < 5$.

□

When $j > 1$, it is harder to bound $B[n, 1]$. This is partially because if $\lfloor n/2 \rfloor < j < i$, where $i$ deviated, $j$ belongs to the second, reversed component, and this reversal can occur multiple times with both $n$ and $j$ changed. However, by a more involved induction on both $n$ and $j$, we can prove that $B[n, j]$ for any $n, j$ is bounded by a constant.

Lemma D.2. *For any $n$ and $1 \leq j \leq n$, $B[n, j] < O(1)$*

Proof. Let $B' = \max(5, \max_{1\le j\le n\le 4} B[n,j])$. We prove the following claim which proves the lemma: if $2^{l-1} < n \le 2^l$, $2^{p-1} < j \le 2^p$,

$$B[n,j] \le 2\left(\sum_{k=0}^{l-3}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B'$$

By the definition of $B'$, the claim holds when $n \le 4$ or $j = 1$ (assume $\sum_{k=a}^{b} = 1$ when $a > b$). For $n > 4$ and $j > 1$ ($l \ge 3$ and $p \ge 1$), consider a line-shortcut graph with $n$ vertices. Let $i$ be the unsatisfied player with the largest index. If $j \le \lfloor n/2 \rfloor$, $j$ belongs to the first component which contains at most $\lfloor n/2 \rfloor$ vertices. If $j > \lfloor n/2 \rfloor$, $j$ belongs to the second component with $i - \lfloor n/2 \rfloor$ vertices. In the latter case, also note that the direction is reversed, so $j$ is changed to $i - j + 1$. By the update rules in Table III,

$$\begin{aligned}
&j \le \lfloor n/2 \rfloor : \\
&B[n,j] \le B[\min(\lfloor n/2 \rfloor, i-1), j] + \frac{1}{n-j+1} B[i - \lfloor n/2 \rfloor, 1] \\
&j > \lfloor n/2 \rfloor : \\
&B[n,j] \le B[i - \lfloor n/2 \rfloor, i-j+1] + \frac{1}{n-j+1} B[i - \lfloor n/2 \rfloor, 1]
\end{aligned}$$

In the first case, $j \le \lfloor n/2 \rfloor$ and $n - j + 1 \ge 2^{l-2}$,

$$\begin{aligned}
B[n,j] &\le B[\min(\lfloor n/2 \rfloor, i-1), j] + \frac{1}{n-j+1} B[i - \lfloor n/2 \rfloor, 1] \\
&\le 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{1}{n-j+1} B' \\
&\le 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{1}{2^{l-2}} B' \\
&\le 2\left(\sum_{k=0}^{l-3}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B'
\end{aligned}$$

We further divide the second case into two. When $\lfloor n/2 \rfloor < j < \lfloor \frac{3n}{4} \rfloor$, $n - j + 1 \ge \frac{n}{4} + 1 \ge 2^{l-3}$ and $n - j + 1 \le j$,

$$\begin{aligned}
B[n,j] &\le B[i - \lfloor n/2 \rfloor, i-j+1] + \frac{1}{n-j+1} B[i - \lfloor n/2 \rfloor, 1] \\
&\le 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{1}{n-j+1} B' \\
&\le 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{1}{2^{l-3}} B' \\
&\le 2\left(\sum_{k=0}^{l-3}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B'
\end{aligned}$$

Finally, if $\lfloor \frac{3n}{4} \rfloor < j$, $n-j+1 \leq \frac{n}{4} \leq \frac{j}{2}$. Let $2^{p'-1} < n-j+1 \leq 2^{p'}$ where $p' \leq p-1$.

$$\begin{aligned} B[n,j] &\leq B[i-\lfloor n/2 \rfloor, i-j+1] + \frac{1}{n-j+1} B[i-\lfloor n/2 \rfloor, 1] \\ &\leq 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p'-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{1}{n-j+1} B' \\ &\leq 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p'-1}\left(\frac{1}{2^k}\right)\right) B' + \frac{2}{2^{p'}} B' \\ &\leq 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p'}\left(\frac{1}{2^k}\right)\right) B' \\ &\leq 2\left(\sum_{k=0}^{l-4}\left(\frac{1}{2^k}\right) + \sum_{k=0}^{p-1}\left(\frac{1}{2^k}\right)\right) B' \end{aligned}$$

Therefore, the claim is proved and the lemma easily follows from the fact that $\sum_{k=0}^{\infty} \frac{1}{2^k} = 2$.

□

Therefore, for any line-shortcut graph, Algorithm 1 finds a desired Go-it-alone equilibrium. This proves the theorem.

## E. PROOF OF LEMMAS IN SECTION 4

### E.1. Proof of Lemma 4.1

Let $Q$ be the strategy profile during **OPT-loop** immediately after vertices in $D_u$ change their strategies to follow $u$. (Note that this is possible since we process $w_i$ in the increasing order of $d_{T^*}(u, w_i)$, and by the definition of $D_u$.) By Lemma 7 of [Fiat et al. 2006], the potential of $Q$ is less than the potential of $P$ by at least $\sum_{i=1}^{|D_u|} z_u[(\frac{\lfloor D_u \rfloor}{2} - (H(|D_u|+1)-1))] \geq \frac{\lfloor D_u \rfloor}{3} z_u$; (note $H(21)-1 \approx 2.645$ and $10-2.645 \geq 20/3$). The vertices in $D_u$ use $e_u$ in their strategies. Consider the group move where each vertex $x \in D_u$ changes her strategy to (the current path from $x$ to $u$) + (the path from $u$ to $u'$ using blue($T^*$) edges) + (the path from $u'$ to $t$ in $P$). After this deviation, some vertices might have the same blue edge more than once, but correcting these situations will only decrease the potential. After this move, $e_u$ will be completely unused again, so the potential of $Q$ will decrease by at least $z_u H(|D_u|)$ from $e_u$. The only edges which are used more are the blue edges on the path from $u$ to $u'$, since $u'$ was already in the strategy of every vertex in $D_u$. The increase of the potential from these edges will be at most $H(|D_u|) d_{T^*}(u, u') \leq H(|D_u|) \frac{z_{u'}}{\alpha}$. However, the strategy profile after this group move will not use $e_u$ and will use the edges in $T(P) - T^*$ in the same direction as they did in $P$. Thus, the potential of $P$ must be no more than the potential of this strategy profile, since we took $P$ to be the strategy profile achieving the minimum potential using the edges in $T(P) \cup T^*$ and obeying the directions for $T(P) - T^*$. Therefore, $H(|D_u|) \frac{z_{u'}}{\alpha} \geq H(|D_u|) z_u + \frac{\lfloor D_u \rfloor}{3} z_u \geq \frac{\lfloor D_u \rfloor}{3} z_u$ and $z_u \leq z_{u'} \frac{3H(|D_u|)}{\alpha |D_u|}$.

### E.2. Proof of Lemma 4.2

As in the previous lemma, let $Q$ be the strategy profile during **OPT-loop** immediately after exactly the vertices in $D_u$ changed their strategies to follow $u$. The potential of $Q$ is less than the potential of $P$ by at least $\frac{|D_u|}{3}z_u$. Let $e_u = (u,v)$, $e_{u'} = (u',v')$ and $r$ be the least common ancestor of $v$ and $v'$ in $T(Q)$. Consider the same group move as before, where each vertex $x \in D_u$ changes her strategy to (the current path from $x$ to $u$) + (the path from $u$ to $u'$ using blue($T^*$) edges) + (the path from $u'$ to $t$ in $P$), followed by improving movements to avoid using same the same edge more than once. Let $Q^r_v$ and $Q^r_{v'}$ be the paths from $v$ to $r$ and $v'$ to $r$ in $Q$, respectively. After the group move, $e_u$ will be completely unused, and the number of players using edges in $Q^r_v$ will be decreased. Since exactly $|D_u|$ vertices stop using these edges, the amount of the potential decreased from these edges is $H(|D_u|)z_u + \sum_{e\in Q^r_v}(H(n_e) - H(n_e - |D_u|))c_e$. On the other hand, the potential is increased in the blue edges on the path from $u$ to $u'$, $e_{u'}$, and the edges in $Q^r_{v'}$. Since $|D_u|$ vertices start using these edges, the amount of the potential increase is at most $H(|D_u|)\frac{z_{u'}}{\alpha} + H(|D_u|)z_{u'} + \sum_{e\in Q^r_{v'}}(H(n_e + |D_u|) - H(n_e))c_e$. Since the potential of $Q$ is already less that the potential of $P$ by at least $\frac{|D_u|}{3}z_u$, and this group move leads to another strategy profile which does not use $e_u$ and use the edges in $T(P) - T^*$ in the same direction as they did in $P$, by the same argument as before, we have

$$\begin{aligned} &H(|D_u|)\frac{z_{u'}}{\alpha} + H(|D_u|)z_{u'} + \sum_{e\in Q^r_{v'}}(H(n_e + |D_u|) - H(n_e))c_e \\ &> H(|D_u|)z_u + \frac{|D_u|}{3}z_u + \sum_{e\in Q^r_v}(H(n_e) - H(n_e - |D_u|))c_e. \end{aligned}$$

If $z_u > \frac{6H(|D_u|)}{|D_u|}z_{u'}$, then $H(|D_u|)z_u + \frac{|D_u|}{3}z_u > H(|D_u|)\frac{z_{u'}}{\alpha} + H(|D_u|)z_{u'}$ and we have

$$\sum_{e\in Q^r_{v'}}(H(n_e + |D_u|) - H(n_e))c_e > \sum_{e\in Q^r_v}(H(n_e) - H(n_e - |D_u|))c_e \tag{1}$$

Now, let $P'$ be the strategy profile right before we process $u'$ as a candidate for **GroupOpt**, after **OPT-loop** and **EE-loop**. Let $M$ be the set of players whose strategy in $P'$ contains $u'$. Note that $u'$ can be neither in $D_{u'}$ (after **OPT-loop** and **EE-loop** $e_{u'}$ will be completely unused), nor in $Q^r_v$ (we required those who did not change strategy between $P$ and $Q$ to not be an ancestor of $v$ in the tree $T(P)$). It follows that $u'$ will be either in the second ($u' \in M_w$ for some $w \in D'_u$) or the third category. Note that if $u'$ is in the second category, **GroupOpt**$(u', u, P')$ is guaranteed to execute because it makes $z_{u'}$ completely unused (thus decreasing the potential by at least $H(|M|)z_{u'}$) while increasing the potential from the blue edges on the path from $u'$ to $u$ by at most $H(|M|)\frac{z_{u'}}{\alpha}$.

If $u'$ in the third category, it means that her ancestors in $P$, including herself, did not change strategy at all during **OPT-loop**, **EE-loop** and previous executions in **GOPT-loop** (if one of her ancestors $u''$ had changed her strategy according one of these moves, it would mean that $(P')_{u''}$ contained $u$, so $u''$ would be either in the first or in the second category, and $u'$ would be in the second category). Since the ancestors of $v$ in $P$ did not change the strategy either, $Q^r_v$ and $Q^r_{v'}$ in (1) refer to the same path in each of $P$, $Q$ and $P'$. The number of players using those edges, however, will be different. Let $n_e$ denote $n_e(Q)$ while using $n'_e$ for $n_e(P')$. Note that during **EE-loop** and **GOPT-loop** the number of players using $Q^r_v$ will increase, as some players will change their strategy to follow $u$. On the other hand, the number of players using $Q^r_{v'}$ will decrease, as some of its descendants might change their strategy to follow $u$. Therefore, $n_e \leq n'_e$ for $e \in Q^r_v$ and $n_e \geq n'_e$ for

$e \in Q^r_{v'}$. Therefore, equation (1) implies that

$$\sum_{e \in Q^r_{v'}} (H(n'_e + |D_u|) - H(n'_e))c_e > \sum_{e \in Q^r_v} (H(n'_e) - H(n'_e - |D_u|))c_e, \quad (2)$$

since $H(x) - H(y) \geq H(x+z) - H(y+z)$ for $x > y \geq 0$ and $z \geq 0$. Let us consider the amount of the potential change caused by $\mathbf{GroupOpt}(u', u, P')$. $e_{u'}$ will be completely unused and the edges in $Q^r_{v'}$ will have less players, so the amount of the potential decrease in these edges is at least $H(|M|)z_{u'} + \sum_{e \in Q^r_{v'}} (H(n'_e) - H(n'_e - |M|))c_e$. On the other hand, the potential will be increased by the blue edges on the path from $u'$ to $u$, $e_u$ and by the edges in $Q^r_{v'}$. The amount of the potential increase is at most $H(|M|)\frac{z_{u'}}{\alpha} + H(|M|)z_u + \sum_{e \in Q^r_{v'}} (H(n'_e + |M|) - H(n'_e))c_e$. Since $\frac{H(x+y+z) - H(x+y)}{z} \leq \frac{1}{x+y} \leq \frac{H(x+y) - H(x)}{y}$ for any $x, y, z > 0$,

$$\begin{aligned} & \sum_{e \in Q^r_{v'}} (H(n'_e) - H(n'_e - |M|))c_e \\ \geq & \sum_{e \in Q^r_{v'}} (H(n'_e + |D_u|) - H(n'_e))c_e \\ \geq & \sum_{e \in Q^r_v} (H(n'_e) - H(n'_e - |D_u|))c_e \text{ by (2)} \\ \geq & \sum_{e \in Q^r_v} (H(n'_e + |M|) - H(n'_e))c_e \end{aligned}$$

Furthermore, note that $\frac{z_u}{4} \leq d_{T^*}(u, u') \leq \frac{z_{u'}}{\alpha}$, which implies $z_u \leq \frac{4}{\alpha} z_{u'}$. For $\alpha \geq 5$, $\frac{z_{u'}}{\alpha} + z_u \geq z'_u$. Finally we have

$$\begin{aligned} & H(|M|)z_{u'} + \sum_{e \in Q^r_{v'}} (H(n'_e) - H(n'_e - |M|)) \\ > & H(|M|)(\frac{z_{u'}}{\alpha} + z_u) + \sum_{e \in Q^r_v} (H(n'_e + |M|) - H(n'_e)) \end{aligned}$$

which indicates that $P'$ will be replaced by $\mathbf{GroupOpt}(u', u, P')$ if $z_u > \frac{6H(|D_u|)}{|D_u|} z_{u'}$.

### E.3. Proof of Lemma 4.3

Note that $\frac{z_u}{4} \leq d_{T^*}(u, u')$ since otherwise $u' \in D_u$ and $e_{u'}$ will be completely unused. Therefore, $z_u \leq \frac{z_{u'}}{2} = \frac{4}{\alpha} z_{u'}$ regardless of $|D_u|$.

$z_u \leq \frac{6H(|D_u|)}{|D_u|} z_{u'}$ is proved when $|D_u| \geq 20$. For $|D_u| \leq 20$, $\frac{4}{\alpha} \leq \frac{6H(i))}{i}$ for any $i \leq 20$.

## F. SECOND PART OF PROOF OF THEOREM 4.4

Let $N$ be the Nash equilibrium finally obtained by the algorithm. We want to bound the sum of $z_u$, where $u$ is a light vertex associated with an edge not in $T^*$. As in [Fiat et al. 2006], we draw a circle of radius $\frac{z_u}{\alpha}$ around each light vertex $u$, and show that the sum of the radii can be bounded in terms of the cost of $T^*$. However, we do not draw circles directly on $T^*$. As in Lemma 3.3, let $T' = (v_1, ..., v_{n'})$ be a sequence of light vertices ordered by a preorder traversal of $T^*$ rooted at $t$. Note that we choose not to include the sink $t$ or any crowded vertices. We define a similar (line) metric on $T'$ by $d_{T'}(v_i, v_{i+1}) = d_{T^*}(v_i, v_{i+1})$ and $d_{T'}(v_i, v_j) = \sum_{i \leq k < j} d_{T'}(v_k, v_{k+1})$ for $i < j$. By the properties of the preorder traversal,

$c(T') \le 2c(T^*)$. Therefore, we could instead show that the sum of the radii can be bounded in terms of $c(T')$ to prove the desired result. From now on, we will often identify $T'$ with a segment of the real line where $v_1, ..., v_{n'}$ correspond to real numbers $x_1 < ... < x_n$ such that $d_{T'}(v_i, v_j) = |x_i - x_j|$. This will allow us to use simple geometric arguments.

Suppose that, on $T'$, we draw a circle of radius $r_u = \frac{z_u}{\alpha}$ around each light vertex $u$. If the circle around $u'$ contains the center $u$ of another circle, $d_{T^*}(u, u') \le d_{T'}(u, u') \le \frac{z_{u'}}{\alpha}$, so by Lemma 4.3, $z_u \le \min(\frac{4}{\alpha}, \frac{6H(|D_u|)}{|D_u|})z_{u'}$, and thus also $r_u \le \min(\frac{4}{\alpha}, \frac{6H(|D_u|)}{|D_u|})r_{u'}$. Furthermore, the fact that $d_{T^*}(u, u') \le r_{u'} \le \frac{z_{u'}}{4}$ but $u$ and $u'$ are both light vertices implies that $e_{u'}$ was added earlier than $e_u$ was added; therefore, there must be no sequence $u_0, ..., u_m = u_0$ such that the circle around $u_i$ contains $u_{i+1}$.

Consider a directed graph $G' = (V', E')$ where $V' \subseteq V$ is the set of the light vertices and $(u, u') \in E'$ if and only if the circle around $u'$ contains $u$ or there exists $u''$ such that $(u'', u') \in E'$ and the circle around $u''$ contains $u$. Then $G'$ is acyclic. Let $B_u = \{w \mid (w, u) \in E'\}$. For each vertex $u \in G'$ without outgoing edges, we will show that $\sum_{w \in B_u} r_w$ is $O(\log\log|B_u|)r_u$. $(w, u) \in E'$ means that $d_{T'}(w, u) \le r_u$ or $((w', u) \in E'$ and $d_{T'}(w, w') \le r_{w'})$ for some $w'$. Since the $r_w$ is decreased multiplicatively by at least $\frac{4}{\alpha}$ for each containment, $d_{T'}(w, u) \le (1 + \frac{4}{\alpha} + \frac{4^2}{\alpha^2} + ...)r_u = (1 + \frac{4}{\alpha} + \frac{4^2}{\alpha^2} + ...)\frac{z_u}{\alpha} \le \frac{z_u}{4}$. This implies $B_u \subseteq D_u$ and $|B_u| \le |D_u| \le \log n$. Therefore, $\sum_{w \in B_u} r_w$ is $O(\log\log\log n)r_u$, and since sum of those $r_u$'s can be at most $4c(T')$ (if no circle contains the center of another, each point in $T'$ can be covered only twice and there is no circle with a radius larger than $T'$), $\sum_w z_w = O(\log\log\log n)c(T')$, which is the main theorem of this paper.

Consider a light vertex $u \in V'$ without outgoing edges (i.e., it is not contained in other any circle), and the line segment from the leftmost point to the rightmost point covered by the circles around the vertices in $B_u$ (these intervals do not need to be contained in $T'$). The length $j_u$ is at least $r_u$, and since $(1 + \frac{4}{\alpha} + \frac{4^2}{\alpha^2} + ...) < 2$, any point on the circle around some $w \in B_u$ cannot be farther than $2r_u$ from $u$, so $j_u \le 4r_u$. Therefore, bounding the sum of radii of the circles around the vertices in $B_u$ in terms of $j_u$ will be enough for our goal. This interval has several properties that we will use inductively, even though partitioning one interval into several subintervals is more involved than forming these first-level intervals:

— Every $w' \in V'$ on this segment is in $B_u$. Let $B'_u \subseteq B_u$ be the set of the vertices on this segment.
— Any segment covered by the circle $w \in B'_u$ cannot go beyond the left end of the segment.
— $\frac{1}{4}j_u \le r_u \le j_u$.

Let $A[n']$ be the maximal sum of the radii $r_1, ..., r_{n'}$ of the $n'$ circles with centers $w_1, ..., w_{n'}$ on the segment of length $j_u = 1$, satisfying the above three requirements and the previous lemma that if $w_i$ is contained in the circle around $w_j$, $r_i \le \min(\frac{4}{\alpha}, \frac{6H(|B_i|)}{|B_i|})r_j$ where $B_i$ is computed as above but only for $n'$ circles on this segment. For example, $A[0] = 0$ and $A[1] = \frac{4}{\alpha}$ (since any $w$ is in $B_u$ and its radius cannot be larger than $\frac{4}{\alpha}r_u \le \frac{4}{\alpha}j_u = \frac{4}{\alpha}$). We prove the following recurrence relation, which shows that $A$ is very slowly growing. As opposed to the selection of the intervals in the first level (those around $u$'s without outgoing edges), the essence of the lemma is to divide a segment into several subsegments such that (1) all of them are contained in the big segment, and (2) subsegments are pairwise disjoint.

LEMMA F.1.

$$A[n'] \le \left( \max_{1 \le k \le n'} \max_{\substack{i_1 + ... + i_k = n' \\ j_1 + ... + j_k \le 1 \\ j_l \le 24H(i_l)/i_l}} \sum_{l=1}^{k} A[i_l - 1]j_l \right) + 3$$

PROOF. Let $w_1, ..., w_{n'}$ and $r_1, ..., r_{n'}$ be the centers and the radii of $n'$ circles achieving $A[n']$. Without loss of generality, the big segment is $[0, 1]$ and assume that $0 < w_1 < ... < w_k < 1$ are the centers that are not contained in any other circle. Note that $w_1 \geq r_1$ since no circle can extend beyond 0. We will find $k$ disjoint line segments $g_l$ of length $j_l$, contained in $[0, 1]$ such that together with $k$ circles, each segment-circle pair satisfies the same requirement for $A$.

First set $g_l = [w_l - r_l, w_l]$; this interval is from the left end of the circle to the center of the circle. This ensures that the length of the interval will be at least the radius of the corresponding circle, since the subsequent operations only enlarge the intervals. Now, consider the circles in $B_l$. Some circles will be on the left of $w_l$ while the others will be on the right. Let

$$L_l = \min_{w_q \in B_l, w_q < w_l} w_q - r_q$$

$$R_l = \max_{\substack{w_q \in B_l, w_q > w_l \\ w_q \notin B_{l'}, \forall l' \in [k]-\{l\}}} w_q$$

Let $g_l \leftarrow g_l \cup [L_l, R_l]$ and $j_l$ be the length of $g_l$. Intuitively, $g_l$ contains all the circles in $B_l$ which are on the left of $w_l$; it extends to the leftmost point, not just the centers. On the right, it just extends to cover the centers which were not previously covered by the aggressive left extension. Figure 9 illustrates an example.

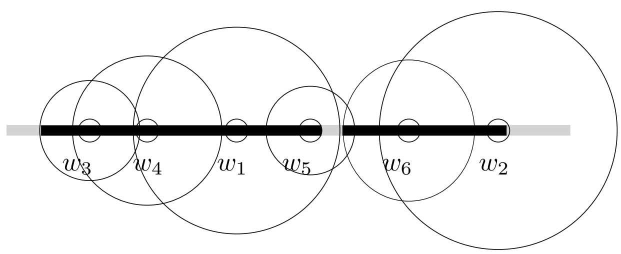


Fig. 9. An example of the division of the big interval to $g_1$ and $g_2$

Since each point can be contained by at most two circles, one on the left and one on the right, it is clear that every $w_q$ will be covered by some $g_l$, $(1 \leq l \leq k)$. Let $B'_l$ be the centers covered by $g_l$. Then $\sum_{l=1}^{k} |B'_l| = n' - k$. Assume that $g_l$ and $g_{l+1}$ intersect. It implies $w_q \geq w_{q'} - r_{q'}$ for some $q \in B_l \cup \{l\} - B_{l+1}, q' \in B_{l+1} \cup \{l+1\}$. However, then $w_q$ is contained in the circle around $w_{q'}$ and $w_q \in B_{l+1}$, which contradicts the definitions of $L_l$ and $R_l$. Therefore, the $g_l$'s are disjoint and $\sum j_l \leq 1$. As previously mentioned, $r_l \leq j_l$ by the definition of $g_l$. Also, $j_l \leq 4r_l$ since the line segment covered by the circles in $B_l$ cannot be more than $2(1 + \frac{4}{\alpha} + \frac{4^2}{\alpha^2} + ...)r_l < 4r_l$. Therefore, $j_l \leq 4r_l \leq \frac{24H(|B_l|+1)}{|B_l|+1} r_u \leq \frac{24H(|B_l|+1)}{|B_l|+1}$.

Note that $g_l$ satisfies all 3 requirements for being bounded by $A[|B'_l|]$: every center on $g_l$ belongs to $B_l$, no circle can extend beyond the left end of $g_l$ (by the definition of $L_l$), and $\frac{1}{4} j_l \leq r \leq j_l$. Finally, the sum of the radii of the circles with the centers in $B'_l$ is bounded by $A[|B_l|]j_l$. The sum of $r_1, ..., r_k$ will be at most 3, since each point $[0, 1]$ can be covered by at most 2 circles, each point in $[1, 2]$ can be covered by only the rightmost circle, and no circle can extend beyond either 0 or 2.

Therefore,

$$\begin{aligned}
A[n] &\leq \left(\sum_{l=1}^{k} A[|B'_l|]j_l\right) + 3 \\
&\quad \text{where } \sum |B'_l| = n' - k, \sum j_l \leq 1, j_l \leq \frac{24H(|B_l|+1)}{|B_l|+1} \\
&= \left(\sum_{l=1}^{k} A[i_l - 1]j_l\right) + 3 \\
&\quad \text{where } i_l = |B'_l| + 1, \sum j_l \leq 1, j_l \leq \frac{24H(i_l)}{i_l} \\
&\leq \left(\max_{1\leq k\leq n'} \max_{\substack{i_1+...+i_k=n' \\ j_1+...+j_k\leq 1 \\ j_l\leq 24H(i_l)/i_l}} \sum_{l=1}^{k} A[i_l-1]j_l\right) + 3
\end{aligned}$$

□

Given the recurrence relation, the main result is then proved by introducing an inverse function. Let $D[m] = \min_{n'} \{A[n'] \geq m\}$.

LEMMA F.2. *For $m \geq 7$,*

$$D[m] \geq \frac{D[m-6]^2}{48H(D[m-6])}$$

PROOF. For given $m$, let $n' = D[m]$. Then, there exists $i_1, ..., i_k, j_1, ..., j_k$ for some $k$ such that

— $i_1 + ... + i_k = n'$
— $j_1 + ... + j_k \leq 1$
— $j_l \leq \frac{24H(i_l)}{i_l}$ for each $1 \leq l \leq k$
— $(\sum_{l=1}^{k} A[i_l - 1]j_l) + 3 \geq A[n'] \geq m$

Without loss of generality, assume $i_1 \geq ... \geq i_k$, and let $k'$ be such that

$$A[i_1 - 1] \geq ... \geq A[i_{k'} - 1] \geq (m-6) > A[i_{k'+1} - 1] \geq ... \geq A[i_k - 1],$$

which means

$$i_1 \geq ... \geq i_{k'} \geq D[m-6] + 1 > i_{k'+1} \geq ... \geq i_k.$$

$(\sum_{l=1}^{k} A[i_l - 1]j_l)$ can be thought of the weighted average of $A[i_l - 1]$'s (if $\sum j_l < 1$, assume that there is an element 0 of weight $1 - \sum j_l$). By the choice of $n'$ and $k'$, $A[i_1 - 1] \leq m$ and $A[i_{k'+1} - 1] \leq m - 6$. Therefore, to make $(\sum_{l=1}^{k} A[i_l - 1]j_l)$ at least $m - 3$, at least half of the weight should be given to $A[i_1 - 1], ..., \geq A[i_{k'} - 1]$, which means $j_1 + ... + j_{k'} \geq 0.5$. Since $j_l \leq \frac{24H(i_l)}{i_l} \leq \frac{24H(D[m-6])}{D[m-6]}$ for each $1 \leq l \leq k'$, $k' \geq \frac{D[m-6]}{48H(D[m-6])}$. Furthermore, $i_1 \geq ... \geq i_{k'} \geq D[m-6]$ implies that $n' \geq \sum_{l=1}^{k'} i_l \geq \frac{D[m-6]^2}{48H(D[m-6])}$, as desired.

□

LEMMA F.3.

$$A[n'] = O(\log \log n')$$

PROOF. Note that $D$ is indeed a very fast increasing function diverging to infinity, so there exists $m'$ such that $\sqrt{D[m'-6]} > 48H(D[m'-6])$. After $m'$, $D[m] \geq \frac{D[m-6]^2}{48H(D[m-6])}$ implies $D[m] \geq D[m-6]^{1.5}$. This implies $D[m] = 2^{1.5^{\Omega(m)}}$ and $A[D[m]] = O(\log\log D[m])$. The monotonicity of $A$ proves the lemma. □