Designing Network Protocols for Good Equilibria*

Ho-Lin Chen† Tim Roughgarden‡ Gregory Valiant§

May 25, 2009

Abstract

Designing and deploying a network protocol determines the rules by which end users interact with each other and with the network. We consider the problem of designing a protocol to optimize the equilibrium behavior of a network with selfish users. We consider network cost-sharing games, where the set of Nash equilibria depends fundamentally on the choice of an edge cost-sharing protocol. Previous research focused on the Shapley protocol, in which the cost of each edge is shared equally among its users.

We systematically study the design of optimal cost-sharing protocols for undirected and directed graphs, single-sink and multicommodity networks, and different measures of the inefficiency of equilibria. Our primary technical tool is a precise characterization of the cost-sharing protocols that only induce network games with pure-strategy Nash equilibria. We use this characterization to prove, among other results, that the Shapley protocol is optimal in directed graphs, and that simple priority protocols are essentially optimal in undirected graphs.

1 Introduction

Most modern-day networks dear to computer science—from the Internet, to the Web, to peer-to-peer and social networks—are created and used by a vast number of autonomous individuals with diverse objectives. Research in the design and analysis of algorithms has responded in kind, with an increasing focus on optimization in networks with self-interested designers or users.

How do we model and analyze selfish behavior in networks? One important genre of problems posits that some aspect of network resource allocation—such as the routing of traffic, the balancing of jobs across machines, the division of bandwidth, or the available network infrastructure—is at least partially controlled by self-interested network users, rather than by the network designer or manager. Almost all work in this area studies applications in which resource allocation is *completely* controlled by selfish network users. The most common goal in these settings is to quantify the

*An extended abstract of this paper appeared in the Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2008.

†The Center for the Mathematics of Information, California Institute of Technology, Moore Laboratory, Pasadena, CA 91125. Research performed while at Stanford University and supported in part by NSF Award 0323766. Email: holinc@gmail.com.

‡Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by an NSF CAREER Award, an Alfred P. Sloan Fellowship, and an ONR Young Investigator Award. Email: tim@cs.stanford.edu.

§UC Berkeley, Computer Science Division, Soda Hall, Berkeley, CA 94720. Part of this work done while visiting Stanford University and supported in part by DARPA grant W911NF-05-1-0224. Email: gregory.valiant@gmail.com.

magnitude of suboptimality that is inevitably caused by selfish resource allocation. This goal is analytic, not algorithmic. One well-studied approximation measure used for this purpose is *the price of anarchy (POA)*—the ratio between the objective function values of a worst Nash equilibrium and that of an optimal solution.

But inefficiency measures like the POA are flexible enough to inform a broader question: how should we *design* networks and their protocols to minimize the efficiency loss caused by selfish behavior? A measure of inefficiency provides a comparative framework for rigorously answering this question—given a set of feasible solutions, the "optimal solution" is the one with the smallest-possible worst-case efficiency loss. This approach adopts inefficiency measures like the POA as *objective functions* to be minimized in novel network optimization problems. The optimal objective function value quantifies the unavoidable loss in solution quality caused by selfish behavior, given the design constraints of the problem.

1.1 Network Cost-Sharing Games

The question of how to design networks and network protocols to minimize the inefficiency of their equilibria can (and should) be asked in a range of models. In this paper, we focus on the conceptually simple but mathematically rich network cost-sharing games introduced by Anshelevich et al. [2, 3].

A *Shapley network design game* [2] is defined as follows. The game transpires in a graph, directed or undirected, with fixed edge costs; these might represent the cost of installing infrastructure between two vertices, or the cost of leasing a large amount of bandwidth on an existing link. Each player i is associated with a source-sink pair (s_i, t_i) and chooses an s_i-t_i path P_i to establish connectivity. Given a choice by each player, the network $H = \cup_i P_i$ is formed at cost $\sum_{e \in H} c_e$. The global objective function is to minimize this cost.

A key assumption in Shapley network design games is that the cost of the network formed is passed on to the players by sharing the cost of each edge $e \in H$ equally among the players that use it.[1] We assume that each player chooses a path to minimize the sum of its cost shares. Every such game admits at least one pure-strategy Nash equilibrium—a choice of a path for each player so that no player can strictly decrease its cost via a unilateral deviation [2]. Crucially, the design decision of how to share the network cost determines the incentive structure and hence the Nash equilibria of the network design game, but it does not affect the global optimization problem of connecting all players at minimum cost.

The inefficiency of equilibria in Shapley network design games is largely understood. The POA can be as large as the number k of players, even in a network of two parallel edges [3]. Better bounds can be obtained by restricting attention to a subset of all Nash equilibria (see also Section 2.2). Considering only the Nash equilibria reachable via best-response dynamics from the empty solution, as in [9, 10], the worst-case ratio — which we call the *reachable POA* — drops to polylogarithmic in k in single-sink undirected networks [9]. Unfortunately, the reachable POA remains polynomial in k in directed networks and multicommodity undirected networks (Seffi Naor, personal communication, May 2007). Considering only the best Nash equilibrium (the *price of stability (POS)*), as in [2, 3], the worst-case ratio in directed graphs is precisely the kth Harmonic

[1] This method of sharing the cost of a single edge e is the same as the Shapley value of the corresponding cooperative game, where the players S are the users of the edge and the characteristic function is $v(\emptyset) = 0$ and $v(T) = c_e$ for all $\emptyset \neq T \subseteq S$.

number $\mathcal{H}_k = \sum_{i=1}^{k} 1/i \approx \ln k$ [2].[2]

These lower bounds on the performance of the Shapley protocol motivate an obvious question: *can we design a better cost-sharing protocol?*

1.2 Our Results: Uniform Protocols

Which cost-sharing protocol minimizes the inefficiency of equilibria in network cost-sharing games? To answer this question, we must precisely define a protocol design space. Formulating such a design space requires a number of modeling choices that are inevitably subject to debate. Our basic model is characterized by the four requirements listed below and is defined more formally in Section 2. Naturally, a case can be made for or against each of these; the most obvious pros and cons of and alternatives to these requirements are discussed in Section 1.4, while Section 1.3 summarizes our results for alternative design spaces.

(1) *Budget-balance:* In each network design game induced by the cost-sharing protocol and in every outcome of the game, the cost of each edge in the formed network is fully passed on to its users.

(2) *Stability:* In each network design game induced by the cost-sharing protocol, there is at least one (pure-strategy) Nash equilibrium.

(3) *Separability:* In each network design game induced by the cost-sharing protocol, the cost shares of each edge are completely determined by the set of players that use it.

(4) *Uniformity:* Across all network design games induced by the cost-sharing protocol, the cost shares of an edge (for each potential set of users) depend only on the edge cost, and not on the network itself.

The first two constraints are self-explanatory. The third constraint insists that the cost shares of an edge in a given outcome depend only on the users of that edge, and are independent of which players use the other edges. The fourth constraint ensures that the cost shares of an edge are not tailored to a particular network. The Shapley protocol satisfies all four constraints: the cost shares of an edge depend only on its cost and its number of users, and are independent of all other properties of the network and the outcome. We call protocols satisfying (1)–(4) *uniform*.

Our main technical result is a *complete characterization* of linear uniform protocols — uniform protocols with cost shares that are a linear function of the edge cost (as in the Shapley protocol). The difficulty of this result stems from the complex stability requirement (2). More precisely, we prove a one-to-one correspondence between such protocols and *"direct products" of certain weighted potential functions*. Potential functions are a standard sufficient condition for the existence of pure-strategy Nash equilibria [47, 54], but many games with pure equilibria admit no potential function. The content of our characterization result is that the *only way* to obtain pure-strategy Nash equilibria via a cost-sharing protocol across all possible networks is via a generalized potential function approach.

We leverage our characterization of linear uniform protocols to identify optimal (not necessarily linear) uniform protocols in both undirected and directed graphs. In undirected networks, it is easy

[2]The worst-case POS of Shapley cost-sharing in undirected graphs is at most $\mathcal{H}_k$, but its exact value is unknown [2, 16] and could be $O(1)$.

to show that simple priority-based uniform protocols dramatically decrease the POA relative to the Shapley protocol — from polynomial in k to polylogarithmic in k. We prove a complementary logarithmic lower bound on the best-possible POA, which applies even in single-sink networks, demonstrating that *priority-based protocols are essentially optimal* in undirected networks. The proof idea is to use our characterization of uniform protocols to associate "weights" with the players, and then exhibit a (weight-dependent) family of networks for which the protocol induces games with large POA.

Our characterization quickly resolves the optimal uniform protocol design problem for directed graphs: for all of the measures of inefficiency we consider, *the Shapley protocol is optimal*. Thus the Shapley protocol, typically motivated by its simplicity and fairness properties, is equally well justified on efficiency grounds in directed graphs — fairness arises "for free" when optimizing for performance.

1.3 Our Results: Extensions

We also study optimal protocol design for the more powerful class of *non-uniform* protocols— those that satisfy only requirements (1)–(3). More formally, while a uniform protocol is defined as a mapping from every possible edge cost and player set to cost shares for these players, a non-uniform protocol is a mapping from edge costs, player sets, *and networks* to cost shares for the players. As we will see in Example 2.7 and thereafter, a simple but powerful way to leverage non-uniformity is to order the players according to some static property of the network, such as shortest-path distances.

For non-uniform protocols, we cannot rely on our characterization theorem and instead establish lower bounds via explicit constructions. For single-sink undirected networks, we show matching upper and lower bounds of 2 on the best-possible POA. For multicommodity networks, we prove a (nearly tight) logarithmic lower bound on the best-possible POA via a novel graph construction. This construction has several additional implications, most notably an information-theoretic $\Omega(\sqrt{\log k})$ lower bound for oblivious network design [20, 24] in k-commodity networks.

In directed graphs, we show that even non-uniform protocols do not admit non-trivial positive results for the POA or reachable POA, and thus focus on minimizing the POS. We show that a POS of 1 is always achievable in single-sink networks and is not always achievable in multicommodity networks. Unlike all of the other settings we study, there remains a non-trivial gap between our best upper and lower bounds in the latter case.

Table 1 summarizes our quantitative results for minimizing the POA and POS. Our upper bounds on the POA trivially carry over to the reachable POA. Minor modifications of our proofs extend most of our lower bounds on the POA to the reachable POA as well.

Finally, we generalize our results to a model that includes an "outside option" for each player, which allows it to opt out of the game at some fixed opportunity cost.

1.4 Discussion

We next discuss the four requirements (1)–(4) in detail. While our protocol design space is natural and leads to non-trivial problems and interesting results, we freely admit that there may be alternative, equally interesting design spaces to explore.

Budget-balance (1) is, of course, the *raison d'être* of a cost-sharing protocol and is the least contentious. It could be interesting to consider some version of approximate budget-balance; we

Network	Measure	Uniform	Non-Uniform
U-SS	POA	$\Theta(\log k)$	2
U-MC	POA	$\Theta(\text{polylog}(k))$	$\Theta(\text{polylog}(k))$
D-SS	POS	$\mathcal{H}_k$	1
D-MC	POS	$\mathcal{H}_k$	$[3/2, \mathcal{H}_k]$
D-SS	POA	k	k

Table 1: Summary of results. Quantities denote the smallest-possible worst-case inefficiency of equilibria, for the given class of networks, approximation measure, and class of cost-sharing protocols. The abbreviations "U", "D", "SS", and "MC" stand for undirected, directed, single-sink, and multicommodity networks, respectively. The $\mathcal{H}_k$ upper bound in directed networks follows from [2].

leave this challenging direction open for future research.

For the stability constraint (2), one line of criticism would argue that it is too strong: by Nash's theorem [50], every protocol always induces a game that has at least one mixed-strategy Nash equilibrium, by which we can measure the protocol's performance. However, the mixed-strategy Nash equilibrium is a notoriously problematic solution concept (see e.g. [52, §3.2]), and is adopted primarily when there is no alternative, in games that possess no pure-strategy Nash equilibria. *When designing the game being played*, as in protocol design, there is due cause for avoiding mixed-strategy Nash equilibria. (A similar argument applies for the "sink equilibria" of [21].) Analogously, algorithmic mechanism design [51] is a field that designs games (largely auctions) that have good equilibria, and almost all work in the area has sought games with *dominant-strategy* (pure) Nash equilibria, a much stronger requirement than stability (2). A second parallel is provided by work in the networking community on the BGP interdomain routing protocol [53], which can be viewed naturally as a game (e.g. [15, 19, 23, 44, 62]) — while mixed-strategy Nash equilibria always exist in the induced path selection game, research has focused entirely on the existence of pure-strategy Nash equilibria.

One could also criticize the stability constraint (2) for being too weak: pure-strategy equilibria should not only exist, but also be *easy to find*. Arguably the most natural strengthening of (2) is to insist that best-response dynamics always converges to a pure-strategy Nash equilibrium from an arbitrary initial state. (This has also been the focus of the literature on BGP, where this property is called "safety" [19, 23, 62].) While our lower bounds assume only the weaker stability requirement (2), all of our upper bounds are achieved using protocols that also satisfy this stronger convergence property. Indeed, a surprising corollary of our characterization of linear uniform protocols is that such a protocol always induces a game with pure-strategy Nash equilibria *if and only if* it always induces a game in which best-response dynamics is guaranteed to converge (Theorem 4.16).

The separability (3) requirement precludes any explicit coordination between different edges of a network. This assumption is restrictive but is satisfied by many important practical network protocols. For example, TCP/IP congestion control with various packet dropping policies (e.g. [31, 45]), can be informally regarded as separable in this sense — each edge makes independent packet dropping decisions based only on the local state, such as the current queue length. Finding a natural generalization of separability that still permits interesting protocol design optimality results is an important research challenge.

Whatever the merits of the uniformity constraint (4), we thoroughly study the optimal protocol design problem both with and without it. That said, a protocol often must be designed without foreknowledge of or assumptions about the network in which it will be deployed. Uniformity is an appropriate requirement in such cases. Moreover, uniformity ensures that a cost-sharing protocol remains well defined as the surrounding network evolves over time. Again, TCP/IP congestion control can be thought of as "uniform" in this high-level sense.

Remark 1.1 We do not allow cost-sharing protocols that can affect the underlying optimization problem of connecting all players at minimum cost. This rules out obtaining near-optimal equilibria for the "wrong reasons" — by increasing the optimal cost rather than by improving the quality of the Nash equilibria.

Remark 1.2 We restrict attention to network cost-sharing games in which players have control only over their connecting path. In particular, the cost-sharing protocol of Anshelevich et al. [3], which allows players to choose endogenously their own cost shares, falls outside of our design space. However, the network games induced by this protocol need not have pure-strategy Nash equilibria (except in single-sink networks), and can also have highly inefficient equilibria [3]. For these reasons, the cost-sharing protocol of [3] is not well suited to the design goals of this paper.

1.5 Further Related Work

Several previous papers studied network cost-sharing games [2, 3, 9, 10, 11, 14, 16, 48]. All of these considered only a fixed cost-sharing protocol and did not address the design questions studied here. The inefficiency of equilibria in other network design games has also been studied; see [61] for an overview. For other models of network formation and design with self-interested participants, see [6, 30, 33] and the references therein.

A few previous papers study how to design protocols to minimize the worst-case inefficiency of equilibria in models unrelated to ours. Christodoulou, Koutsoupias, and Nanavati [12] and several follow-up papers [5, 29, 40] design machine scheduling policies to minimize the worst-case POA in variants of the scheduling game proposed by Koutsoupias and Papadimitriou [42]. Johari and Tsitsiklis [35] design protocols for allocating a single divisible resource among heterogeneous players and show that, among all protocols that meet certain desirable "scalability" constraints, the Kelly protocol [38] minimizes the worst-case efficiency loss. In a mechanism design context, Moulin and Shenker [49] identify groupstrategyproof and budget-balanced mechanisms that minimize worst-case additive efficiency loss over all possible valuation profiles.

To a lesser extent, the goals of this paper are similar to previous approaches for improving the price of anarchy of a given game; see, for example, previous work on pricing selfish routing networks [13, 17, 37] and Stackelberg routing [43, 56, 60]. Our work here differs in its aim to design a single distributed protocol to minimize the worst-case inefficiency of equilibria over an entire *family of games*, rather than a centralized algorithm for improving the POA in a given game.

Finally, the goal of designing a game with good equilibria bears some resemblance to that of algorithmic mechanism design [51]. In mechanism design problems, however, there is generally some crucial data, such as players' valuations for different goods or resources, which are unknown to the mechanism designer. There is no private information in the games studied here; instead, the designer lacks full control over the *allocation* of resources. For this reason, the problems studied in this paper are technically very different from those in algorithmic mechanism design.

2 Definitions and Examples

2.1 Network Cost-Sharing Games

A *network* is an undirected or directed graph $G = (V, E)$ and a nonnegative cost c_e for each edge $e \in E$. A *network cost-sharing game* is additionally specified by a set $\{1, \ldots, k\}$ of k *players* that we identify with source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$ and a *cost-sharing method* $\xi_e : 2^{\{1,\ldots,k\}} \to \mathcal{R}_+^k$ for each edge e. A cost-sharing method assigns nonnegative cost shares to the players, as a function of the set of players that choose a path that contains the edge e. We abuse notation and write $\xi_e(i, S)$ for the cost share of player i for the edge e, given that S is the set of players using e.

The *strategy set* of player i is the set $\mathcal{P}_i$ of s_i-t_i paths. In an *outcome* of the game, each player i chooses a single path $P_i \in \mathcal{P}_i$. The *cost* of an outcome $(P_1, \ldots, P_k)$ is defined to be $C(P_1, \ldots, P_k) = \sum_{e \in \cup_i P_i} c_e$.

A cost-sharing method ξ_e is *budget-balanced* if for every set $S \subseteq \{1, 2, \ldots, k\}$:

(1) $\xi_e(i, S) = 0$ for all players $i \notin S$;

(2) $\sum_{i \in S} \xi_e(i, S) = c_e$.

A cost-sharing method is automatically separable in the sense of (3) in Section 1.2 in that its domain is simply the possible sets of users — if the users of an edge are the same in two different outcomes of the game, the cost shares of these users for this edge are also the same.

The cost-sharing methods determine the incentives in a network cost-sharing game by inducing a *cost function* $c_i : \mathcal{P}_1 \times \cdots \times \mathcal{P}_k \to \mathcal{R}_+$ for each player i, defined as

$$c_i(P_1, \ldots, P_k) = \sum_{e \in P_i} \xi_e(i, S_e),$$

where $S_e = \{j : e \in P_j\}$ denotes the set of players choosing a path that contains the edge e. If all of the cost-sharing methods of a network cost-sharing game are budget-balanced, then the cost of each outcome $(P_1, \ldots, P_k)$ is partitioned among the players: $C(P_1, \ldots, P_k) = \sum_{i=1}^k c_i(P_1, \ldots, P_k)$.

An outcome of a network cost-sharing game is a *pure-strategy Nash equilibrium (PNE)* if no player can decrease its cost by changing its strategy. Formally, the outcome $(P_1, \ldots, P_k)$ is a PNE if for every player i and every strategy $P_i' \in \mathcal{P}_i$,

$$c_i(P_1, \ldots, P_i, \ldots, P_k) \leq c_i(P_1, \ldots, P_i', \ldots, P_k).$$

2.2 Quantifying Inefficiency of Equilibria

We aspire toward network cost-sharing games with relatively efficient PNE. A standard and conservative measure of the inefficiency of the equilibria of a game is the *price of anarchy (POA)*, the largest ratio between the cost of a PNE and that of a minimum-cost outcome. When there are no interesting upper bounds on the cost of all equilibria, a common weaker goal is to bound the cost of a subset of equilibria. An extreme approach is to bound the *price of stability (POS)*, defined as the smallest ratio between the cost of a PNE and that of an optimal outcome. An intermediate measure is what we call the *reachable POA*, a quantity defined in [10]. The numerator of this ratio is the largest cost of an equilibrium reachable via the following process: players enter the game one-by-one in an arbitrary order, and each picks a path of minimum cost, given the choices of previous

players; after all players have entered, an arbitrary player is selected to re-optimize its path, given the current strategies of all other players (a "best response"); when the process reaches a PNE (as it must [2]), it stops. In general, not all PNE of a network cost-sharing game can be obtained via this process [10]. See [57] for further discussion and interpretations of these and related concepts.

2.3 Cost-Sharing Protocols

Formally, a cost-sharing protocol takes as input a network and player set, and outputs a collection of edge cost-sharing methods, thereby providing the final ingredient for a network cost-sharing game.

Definition 2.1 (Cost-Sharing Protocol) A *cost-sharing protocol* assigns, for every network $G = (V, E)$ with edge costs c, for every player set $\{1, 2, \ldots, k\}$, and every set $(s_1, t_1), \ldots, (s_k, t_k)$ of source-sink pairs, a cost-sharing method ξ_e to every edge $e \in E$.

Example 2.2 (The Shapley Protocol [2]) The *Shapley protocol* always assigns an edge e of cost c_e the cost-sharing method ξ_e given by $\xi_e(i, S) = c_e/|S|$ for every subset S of players and $i \in S$.

More generally, a cost-sharing protocol can assign cost-sharing methods in a way that depends on additional information, including the topology of G, the locations of player sources and sinks, and the costs of other edges of the network.

A cost-sharing protocol is *stable* if it only induces network cost-sharing games that possess at least one PNE. An *admissible* protocol is stable and only assigns budget-balanced cost-sharing methods; such a protocol meets the first three requirements from Section 1.2. *Uniform protocols* additionally meet the uniformity constraint (4) of that section.

Definition 2.3 (Uniform Protocols) An admissible cost-sharing protocol is *uniform* if the cost-sharing method ξ_e it assigns to an edge e of a network G is a function only of the edge cost c_e and the player set $\{1, 2, \ldots, k\}$.

Remark 2.4 Definition 2.3 allows a uniform protocol to assign cost-sharing methods $\{\xi_e\}_{e \in E}$ in a way that depends on the number k of players in the game. A natural extra requirement would be to insist that for every edge e, $\xi_e(i, S)$ is a function only of c_e, S, and i (and is independent of k). All of our positive results for uniform protocols remain valid under this extra constraint, and all of our negative results apply to all uniform protocols.

Are there admissible protocols other than the Shapley protocol? *Ordered* protocols form simple, important examples. Such cost-sharing protocols are defined with respect to an ordering of the players, and can be either uniform (independent of the network) or non-uniform (defined in a network-dependent way). In either case, the first player in the ordering pays the full cost of every edge in its path; the second player pays the full cost of every edge in its path not already paid for by the first player; and so on.

Definition 2.5 (Ordered Protocols) A cost-sharing method ξ_e in a network cost-sharing game with player set $\{1, 2, \ldots, k\}$ is *ordered according to* σ, where σ is a permutation of the players, if for every $S \subseteq \{1, 2, \ldots, k\}$ and $i \in S$, $\xi_e(i, S) = c_e$ if $\sigma(i) \leq \sigma(j)$ for all $j \in S$ and $\xi_e(i, S) = 0$

otherwise. A cost-sharing protocol is *ordered* if, for every network and player set, it assigns cost-sharing methods that are all ordered according to a common ordering.

Proposition 2.6 (Stability of Ordered Protocols) *Every ordered cost-sharing protocol is admissible.*

Proof: Budget-balance is clear. For stability, fix a network and a player set, and suppose the protocol assigns cost-sharing methods that are all ordered according to σ. Renaming the players, we can assume that σ is the identity. Define P_1 to be a shortest s_1-t_1 path with respect to c. For each $i > 1$ in turn, define P_i to be a shortest s_i-t_i path after zeroing out the cost of all edges of $P_1 \cup \cdots \cup P_{i-1}$. The outcome $(P_1, \ldots, P_k)$ is a PNE. ■

2.4 Example: The Prim Protocol

Ordered protocols can be radically better than the Shapley protocol in undirected networks; our first demonstration is for a non-uniform protocol for single-sink networks.

Example 2.7 (Prim Cost-Sharing Protocol) Consider an undirected single-sink network G with edge costs c, a sink vertex t, and source vertices $s_1, \ldots, s_k$. The *Prim cost-sharing protocol* is the (non-uniform) ordered cost-sharing protocol that orders the players as follows. The first player is the one with source s_i closest to the sink t with respect to the edge costs c; the second player is the one with source closest to the set $\{t, s_i\}$; and so on.

Proposition 2.8 (POA of Prim Protocol) *For every single-sink undirected network and player set, the Prim protocol induces a network cost-sharing game with POA at most 2.*

Proof (sketch): Fix a single-sink undirected network and a player set. Renaming the players, we can assume that the Prim protocol assigns cost-sharing methods that are ordered by the identity. In every PNE of the induced network cost-sharing game, every player i chooses a shortest path P_i between its source s_i and $P_1 \cup \cdots \cup P_{i-1}$. The cost incurred by this player is at most the shortest-path distance (w.r.t. the original network edge costs) between s_i and the set $\{t, s_1, \ldots, s_{i-1}\}$. Thus, every PNE has cost bounded above by a possible output of the MST heuristic for the Steiner tree problem, when implemented using Prim's MST algorithm. Every such output has cost at most twice that of a minimum-cost Steiner tree (see e.g. [63]), a minimum-cost outcome in the network game. ■

Recall that the worst-case POA of the Shapley protocol, even in undirected networks of parallel edges, is the number k of players [3].

Remark 2.9 (Optimality of the Prim Protocol) Standard examples (e.g. [63, Example 3.4]) give a matching lower bound on the worst-case POA and reachable POA of every (possibly non-uniform) admissible cost-sharing protocol in single-sink undirected networks.

Remark 2.10 (Forthcoming Lower Bounds) Both the non-uniformity of the Prim protocol and the restriction to single-sink undirected networks are necessary to obtain a constant worst-case POA. We prove, by completely different methods, that every uniform protocol has a worst-case POA of $\Omega(\log k)$, even in single-sink undirected networks (Theorem 4.3); and that every (possibly non-uniform) admissible protocol has a worst-case POA of $\Omega(\log k)$ in multicommodity undirected networks (Theorem 6.1).

3 Characterization of Linear Uniform Protocols: Overview

This section describes a complete characterization of the linear uniform protocols as those induced by the "concatenation" of "weighted potential functions". The stability constraint (2) from Section 1.2—a complex global constraint on all network games that might be induced by a protocol—makes this result highly non-trivial, and we defer the proof to Section 5. Section 4 gives several applications of this characterization, including lower bounds on the worst-case POA and POS achievable by (not necessarily linear) uniform protocols in undirected and directed networks, respectively.

3.1 Potential-Based Protocols

We begin with two definitions.

Definition 3.1 (Linear Protocols) A uniform protocol is *linear* if, for all $c_e \geq 0$, the cost-sharing method it assigns to an edge of cost c_e is $c_e \cdot \xi_1$, where ξ_1 is the cost-sharing method it assigns to an edge of unit cost.

We often abuse notation and refer to a linear uniform protocol by the cost-sharing method it assigns to a unit-cost edge.

Definition 3.2 (Positive Methods and Protocols) A cost-sharing method of a non-zero cost edge is *positive* if it always assigns strictly positive cost shares to its users. A cost-sharing protocol is *positive* if it assigns only positive cost-sharing methods to edges with non-zero cost.

For example, the Shapley protocol is positive, but ordered protocols are not.

Our plan is to use the two known methods of proving stability as a compass to map the terrain of linear and uniform protocols. Proposition 2.6 proves that ordered protocols are stable by explicitly exhibiting a PNE in every induced network cost-sharing game. The stability of the Shapley protocol is a more subtle "potential function argument" [2, 47, 55]: one exhibits a potential function for each induced network game such that local minima of the potential function are in bijective correspondence with the PNE of the game. To begin our development, do any other protocols admit analogous potential functions? This question motivates the following definition.

Definition 3.3 (Edge Potential) Let $\{1, 2, \ldots, k\}$ be a player set. A strictly positive function $f : 2^{\{1,\ldots,k\}} \to \mathcal{R}^+$ is an *edge potential* if it is strictly increasing ($f(T) < f(S)$ whenever $T \subset S$) and if

$$\sum_{i \in S} \frac{f(S) - f(S \setminus \{i\})}{f(\{i\})} = 1 \tag{1}$$

for every $S \subseteq \{1, \ldots, k\}$.

Every edge potential induces a positive, linear, and uniform cost-sharing protocol; we call these *potential-based protocols*.

Proposition 3.4 (Properties of Potential-Based Protocols) *Let f be an edge potential for the player set $\{1, 2, \ldots, k\}$. Define a cost-sharing protocol by assigning an edge of cost c the cost-sharing method ξ, where*

$$\xi(i, S) = c \cdot \frac{f(S) - f(S \setminus \{i\})}{f(\{i\})} \tag{2}$$

for every $S \subseteq \{1, \ldots, k\}$ and $i \in S$. This protocol is positive, linear, and uniform.

Proof: Linearity is clear. The method ξ for an edge of cost $c > 0$ is positive because the edge potential is, by definition, strictly positive and strictly increasing. The method is budget-balanced by (1). To prove that this cost-sharing protocol is stable and hence uniform, consider the following potential function for a (directed or undirected) network $G = (V, E)$ with edge costs c:

$$\Phi(P_1, \ldots, P_k) = \sum_{e \in E} c_e \cdot f(S_e), \tag{3}$$

where $S_e = \{i \in \{1, \ldots, k\} : e \in P_i\}$. By the definitions (2) and (3), whenever a player i changes its strategy from P_i to Q_i, the identity

$$\begin{aligned} \Delta c_i &= \sum_{e \in Q_i \setminus P_i} \xi(S_e \cup \{i\}) - \sum_{e \in P_i \setminus Q_i} \xi(S_e) \\ &= \sum_{e \in Q_i \setminus P_i} c_e \cdot \frac{f(S_e \cup \{i\}) - f(S_e)}{f(\{i\})} - \sum_{e \in P_i \setminus Q_i} c_e \cdot \frac{f(S_e) - f(S_e \setminus \{i\})}{f(\{i\})} \\ &= \frac{\Delta \Phi}{f(\{i\})} \end{aligned} \tag{4}$$

holds. Hence, every local minimum of Φ is a PNE of the network cost-sharing game — and there is at least one, the global minimum of Φ. ■

The Shapley protocol is potential-based with the edge potential $f(S) = \mathcal{H}_{|S|}$ for every subset S of players, where as usual $\mathcal{H}_k = \sum_{i=1}^{k} 1/i$ denotes the kth Harmonic number. Are there any others? Because of the budget-balance constraint in Definition 3.3, the answer is not clear. But our proof of Theorem 3.8 implicitly shows that there are a plethora of potential-based protocols, in bijective correspondence with the open unit cube $(0, 1)^{k-1}$ in $(k - 1)$ dimensions.

Remark 3.5 (Weighted Potentials and Weighted Shapley Values) Due to the player-specific scaling factor in (4), the function in (3) is more properly called a *weighted* potential function [47], as opposed to the *exact* potential function for the Shapley protocol (where $f(\{i\}) = 1$ for all i). It is tempting but incorrect to regard each value $f(\{i\})$ of an edge potential as a "weight" for player i in the weight-proportional sharing sense of [2, 11]. The appropriate interpretation is in terms of the weighted Shapley value [36, 59]; the value $f(\{i\})$ corresponds to a player weight of $1/f(\{i\})$ in the sense of Kalai and Samet [36]. See also Hart and Mas-Colell [27] and Monderer and Shapley [47] for similar connections between weighted potential functions and weighted Shapley values in other contexts.

3.2 Concatenation

Ordered protocols demonstrate that not all linear uniform protocols are positive. This motivates an operation that combines two cost-sharing protocols into a single (non-positive) one. In the following definition, we refer to a linear protocol in terms of the cost-sharing method it assigns a unit-cost edge.

Definition 3.6 (Concatenation) Let ξ_1 and ξ_2 be linear, uniform cost-sharing protocols for disjoint player sets A_1 and A_2, respectively. The *concatenation of* ξ_1 *and* ξ_2 is the cost-sharing protocol $\xi_1 \oplus \xi_2$ for the player set $A_1 \cup A_2$ defined by

$$(\xi_1 \oplus \xi_2)(i, S) = \begin{cases} \xi_1(i, S \cap A_1) & \text{if } i \in A_1 \\ \xi_2(i, S) & \text{if } S \subseteq A_2 \\ 0 & \text{otherwise.} \end{cases}$$

In words, players of A_1 share the cost of an edge as if no players of A_2 were present (according to ξ_1); if only players of A_2 use an edge, they share its cost according to ξ_2. Ordered protocols are precisely the k-fold concatenations of trivial one-player protocols.

Proposition 3.7 (Concatenation Preserves Linearity and Uniformity) *The concatenation of two linear uniform cost-sharing protocols is a linear uniform protocol.*

Proof (sketch): Arguing as in the proof of Proposition 2.6 shows that concatenation preserves stability; all other properties are obvious. ■

We can now state our characterization result: *every* linear uniform cost-sharing protocol arises as the concatenation of potential-based protocols.

Theorem 3.8 (Characterization of Linear Uniform Protocols) *A cost-sharing protocol is linear and uniform if and only if it is the concatenation of potential-based protocols.*

Remark 3.9 (Interpretation via Weighted Shapley Values) As alluded to in Remark 3.5, a positive protocol induced by a potential f coincides with the weighted Shapley value for the weight vector $1/f(\{1\}), \ldots, 1/f(\{k\})$ (see [27, 36, 47]). Theorem 3.8 can therefore be interpreted as: *the linear uniform cost-sharing protocols are precisely the concatenations of weighted Shapley values.*

We provide a full proof of Theorem 3.8 in Section 5, after exploring several applications in Section 4. We record here one of the major milestones of the proof, which is also directly useful in the applications in the next section.

Lemma 3.10 (Monotonicity of Linear Uniform Protocols) *Every linear and uniform protocol ξ for a player set $\{1, \ldots, k\}$ is* monotone*: for every $S \subseteq T \subseteq \{1, \ldots, k\}$ and $i \in S$, $\xi(i, S) \geq \xi(i, T)$.*

4 Characterization of Uniform Protocols: Applications

This section presents three results that build on Theorem 3.8: ordered protocols have near-optimal worst-case price of anarchy in undirected networks; the Shapley protocol has optimal worst-case price of stability in directed networks; and PNE exist in all network cost-sharing games induced by a protocol if and only if better-response dynamics always converges (to a PNE) in these games.

4.1 Undirected Graphs: Near-Optimality of Ordered Protocols

Which uniform protocol minimizes the worst-case POA in undirected graphs? Ordered protocols offer some simple but interesting upper bounds.

Proposition 4.1 (POA of Ordered Protocols) *Every uniform ordered protocol has a worst-case POA of $O(\log k)$ in single-sink undirected networks and a worst-case POA of $O(\log^2 k)$ in multicommodity undirected networks.*

Proof: Consider a uniform ordered protocol for single-sink undirected networks with the player set $\{1, 2, \ldots, k\}$. Renaming the players, we can assume that the corresponding ordering is the identity. Consider a single-sink network G and define paths $P_1, \ldots, P_k$ as in the proof of Proposition 2.8, breaking ties among equal-cost paths in a worst-case manner. The POA of the corresponding network cost-sharing game is the ratio between the cost of these paths $\cup_{i=1}^k P_i$ and that of a minimum-cost Steiner tree spanning $\{s_1, \ldots, s_k, t\}$.

An alternative interpretation of the network $\cup_{i=1}^k P_i$ is as the output of a natural greedy algorithm for the online Steiner tree problem: player 1 arrives first and is immediately connected to t via a shortest path; then player 2 arrives and is connected to the network-so-far via a shortest path; and so on. This process generates the exact same sequence of paths $P_1, \ldots, P_k$. Thus, the worst-case POA of this (arbitrary) ordered uniform protocol is precisely the worst-case competitive ratio of this online Steiner tree algorithm. Imase and Waxman [28] proved that the latter quantity is $O(\log k)$, and the assertion for single-sink networks follows immediately.

For multicommodity networks, the worst-case POA of an ordered uniform protocol is precisely the worst-case competitive ratio of the natural greedy algorithm for the online *generalized* Steiner tree problem. Here, pairs (s_i, t_i) arrive online, and the algorithm connects s_i to t_i using a shortest path relative to the network already constructed (with previously constructed edges treated as zero-cost). Awerbuch, Azar, and Bartal [4] proved that this competitive ratio is $O(\log^2 k)$; the claim for multicommodity networks follows. ■

Remark 4.2 (Lower Bounds for Ordered Protocols) Theorem 4.3 below shows that this $O(\log k)$ upper bound for single-sink networks cannot be improved by *any* uniform protocol, ordered or otherwise. For multicommodity networks, there is a $\Theta(\log k)$ factor gap between our upper and lower bounds on the worst-case POA of uniform protocols, both in general and for the special case of ordered protocols. Obtaining a tight analysis of the greedy online algorithm of [4] remains an open problem, and its competitive ratio might well be $O(\log k)$. A proof of such an upper bound would, of course, close the remaining gap between our upper and lower bounds for multicommodity networks. We note that Berman and Coulston [8] devised a non-greedy $O(\log k)$-competitive online algorithm for the generalized Steiner tree problem, but it does not seem to have any implications for our game-theoretic protocol design problems.

Next, we leverage our main characterization result (Theorem 3.8) to prove that ordered protocols are almost optimal uniform protocols.

Theorem 4.3 (Near-Optimality of Ordered Protocols) *Every uniform cost-sharing protocol has a worst-case POA of $\Omega(\log k)$, even in single-sink undirected networks.*

We develop the proof of Theorem 4.3 in several steps. The high-level proof plan is, given a uniform protocol, to first apply Theorem 3.8 to classify players i according to their $f(\{i\})$-values

in the associated edge potential f. We then make precise the intuition that players with similar f-values should have similar cost shares. This enables a dichotomy lemma showing that there are either $\log k$ players that receive near-identical cost shares, or $k/\text{polylog}(k)$ players for which the protocol is close to an ordered protocol. In each case, we construct a family of networks in which the protocol induces games with $\Omega(\log k)$ POA.

We first dispense with a preliminary result: a reduction of Theorem 4.3 to the special case of linear uniform protocols. Every uniform protocol induces a *linear extension*, the linear protocol that assigns the cost-sharing method $c_e \cdot \xi$ to an edge of cost c_e, where ξ is the cost-sharing method that the given uniform protocol assigns to a unit-cost edge.

Lemma 4.4 *The linear extension of a uniform protocol is uniform.*

Proof (sketch): We prove stability of the linear extension; all other properties are obvious. Suppose there is a counterexample network in which the linear extension fails to induce a game with at least one PNE. Appealing to linearity and the density of the rationals, there is also such a counterexample with rational edge costs. Scaling the edge costs and invoking linearity, there is a counterexample network with integral edge costs. Again by linearity, subdividing edges yields a counterexample in which all edges have cost 0 or 1. But the linear extension and the original protocol coincide on such a network, contradicting the stability of the latter. ■

Arguing as in the proof of Lemma 4.4 also shows that the worst-case POA of linear uniform protocols is determined by *0-1 networks* — those in which all edges have cost 0 or 1.

Lemma 4.5 *For every linear uniform protocol and every $k \geq 1$, its worst-case POA in k-player single-sink undirected networks is attained, up to an arbitrarily small additive constant, in a 0-1 network.*

We can now deduce that the minimum worst-case POA of uniform protocols is the same as that of linear uniform protocols.

Lemma 4.6 (Reduction to Linear Protocols) *For every $k \geq 1$, the worst-case POA of a uniform protocol in k-player single-sink undirected networks is no smaller than that of its linear extension.*

Proof: Lemma 4.5 implies that, up to an arbitrarily small error, there is a worst-case network G for the linear extension which is also 0-1. Since the original protocol and its linear extension coincide on 0-1 networks, they induce the same game (with the same POA) in G. ■

Conceptually, Lemma 4.6 implies that adding linearity to the requirements (1)–(4) of Section 1.2 would not affect the optimal solutions to our protocol design problems.

Remark 4.7 In preparation for Theorem 4.15 below, we note that Lemmas 4.4–4.6 also hold, with the same proofs, in directed networks. Moreover, Lemmas 4.5 and 4.6 also hold with the POA replaced by the POS.

We next prove four lemmas that relate proximity in edge potential values to proximity in cost shares. Henceforth, we write $f(i)$ for $f(\{i\})$, which we call the *f-value* of player i. We first use (1)

to obtain the following useful expression for the value $f(S)$ of an edge potential in terms of the f-values of the players of S:

$$f(S) \;=\; \left(1+\sum_{i\in S}\frac{f(S\setminus\{i\})}{f(i)}\right)\cdot\left(\sum_{i\in S}\frac{1}{f(i)}\right)^{-1}. \tag{5}$$

Lemma 4.8 *Let f be an edge potential for the players $\{1,2,\ldots,k\}$, $\alpha \geq 1$ a constant, and S a subset of players with $f(i) \leq \alpha \cdot f(j)$ for every $i,j \in S$. For every subset $T \subseteq S$ and player $j \in S$, $f(T) \leq \alpha \mathcal{H}_{|T|} f(j)$.*

Proof: By induction on $|T|$. For the inductive step, fix $T \subseteq S$. Use (5) to obtain

$$f(T) \;=\; \left(\sum_{i\in T}\frac{f(T\setminus\{i\})}{f(i)}\right)\cdot\left(\sum_{i\in T}\frac{1}{f(i)}\right)^{-1} + \left(\sum_{i\in T}\frac{1}{f(i)}\right)^{-1}.$$

The first term on the right-hand side is a weighted average of the $f(T\setminus\{i\})$'s, which is at most $\alpha\mathcal{H}_{|T|-1}f(j)$ by the inductive hypothesis. The second term is at most $\alpha f(j)/|T|$ by the definition of α, completing the inductive step. ■

Lemma 4.9 *Let f, S, and α be as in Lemma 4.8. For every two (not necessarily disjoint) subsets $T_1, T_2 \subseteq S$ with ℓ players each, $f(T_1) \leq \alpha^{2\ell-1} f(T_2)$.*

Proof: By induction on ℓ. Lemma 4.8 provides the base case. For the inductive step, let $T_1, T_2 \subseteq S$ have size ℓ. By (5), the definition of α, and the inductive hypothesis, we have

$$\begin{aligned} f(T_1) &= \left(1+\sum_{i\in T_1}\frac{f(T_1\setminus\{i\})}{f(i)}\right)\cdot\left(\sum_{i\in T_1}\frac{1}{f(i)}\right)^{-1} \\ &\leq \left(1+\sum_{i\in T_2}\frac{\alpha^{2\ell-3}\cdot f(T_2\setminus\{i\})}{(f(i)/\alpha)}\right)\cdot\left(\sum_{i\in T_2}\frac{1}{\alpha\cdot f(i)}\right)^{-1} \\ &\leq \alpha^{2\ell-1}\cdot f(T_2). \end{aligned}$$

■

We can now show that players with similar f-values receive similar cost shares. For the rest of this section, all logarithms are base 2 (say).

Lemma 4.10 (Proximity Lemma) *Let f, S, and α be as in Lemma 4.8. Assume further that $|S| \leq \log k$ and that α satisfies*

$$\alpha^{2\log k} \leq 1 + \log^{-3} k. \tag{6}$$

Let ξ be the potential-based cost-sharing method induced by f for a unit-cost edge. Then,

$$\xi(i,S) \leq \alpha\left(\xi(j,S) + 2\log^{-2} k\right)$$

for every pair i,j of players of S.

Proof: Lemma 4.9 and (6) imply that for every two subsets $T_1, T_2 \subseteq S$, $f(T_1)$ and $f(T_2)$ differ by at most a factor of $1 + \log^{-3} k$. Lemma 4.8 implies that $f(T_1)$ and $f(T_2)$ are both at most $f(i) \cdot 2 \log k$ for any player $i \in S$. Hence, $f(T_1)$ and $f(T_2)$ differ by at most an additive $2 \log^{-2} k \cdot f(i)$ amount for any player $i \in S$.

Therefore, for every two players $i, j \in S$,

$$\begin{aligned} \xi(i,S) &= \frac{f(S) - f(S \setminus \{i\})}{f(i)} \\ &\geq \frac{f(S) - f(S \setminus \{j\}) - 2\log^{-2} k \cdot f(i)}{f(i)} \\ &\geq \frac{\xi(j,S)}{\alpha} - 2\log^{-2} k, \end{aligned}$$

which proves the lemma. ■

Our final lemma prior to the proof of Theorem 4.3 is a weak converse of Lemma 4.10.

Lemma 4.11 (Separation Lemma) *Let f and ξ be as in Lemma 4.10, and i, j two players with $f(i) \geq \beta \cdot f(j)$. Then:*

(a) $\xi(i, \{i,j\}) \geq \frac{\beta}{\beta+1}$;

(b) for every $S \supseteq \{i,j\}$, $\xi(j,S) \leq \frac{1}{\beta+1}$.

Proof: Part (b) follows from (a) and the monotonicity of ξ (Lemma 3.10). To prove (a), we can assume without loss that $f(j) = 1$. First assume that $f(i)$ is exactly β. By (5), $f(\{i,j\}) = 1 + \frac{\beta^2}{\beta+1}$. By (1), $\xi(i, \{i,j\}) = \beta/(\beta+1)$. Since this is increasing in β, part (a) follows. ■

Proof of Theorem 4.3: Consider a uniform cost-sharing protocol for the player set $\{1, 2, \ldots, k\}$. We can assume that k is sufficiently large (at least 2^{32}, say). By Lemmas 4.4 and 4.6, we can assume that the protocol is a linear protocol ξ.

Theorem 3.8 implies that f is a concatenation of potential-based cost-sharing protocols $\xi_1, \ldots, \xi_m$ for disjoint player sets $A_1, \ldots, A_m$, where the protocol ξ_i is derived from an edge potential f_i. Scaling the f_i's, we can assume that $f_i(\{j\}) \geq 1$ for every i and j. We bucket the players of each A_i by f_i-values into *groups*, using intervals of the form $[\alpha^\ell, \alpha^{\ell+1})$ for nonnegative integers ℓ, where $\alpha > 1$ is chosen to satisfy (6) with equality.

The proof now breaks into two cases. First suppose that there are fewer than $k/\log k$ distinct non-empty groups across all of the A_i's. Then, there is a single group that contains a set S of $\log k$ players. Budget-balance implies that there exists a player j for which $\xi(j,S) \leq 1/\log k$. The Proximity Lemma (Lemma 4.10), our choice of α, and the fact that k is sufficiently large then imply that $\xi(j,S) \leq 2/\log k$ for every $j \in S$. Create a single-sink undirected network as follows (Figure 1(a)). Each player not in S has a zero-cost edge from its source vertex to the single sink t. The players of S share a common source vertex s, and there are two parallel edges from s to t, with costs 1 and $(\log k)/3$. The optimal solution clearly has cost 1. On the other hand, if all players of S share the more expensive s-t edge, then each incurs a cost share strictly less than 1 and will not unilaterally deviate to the unit-cost edge. This outcome is therefore a PNE, and the POA in this network game is $(\log k)/3$.

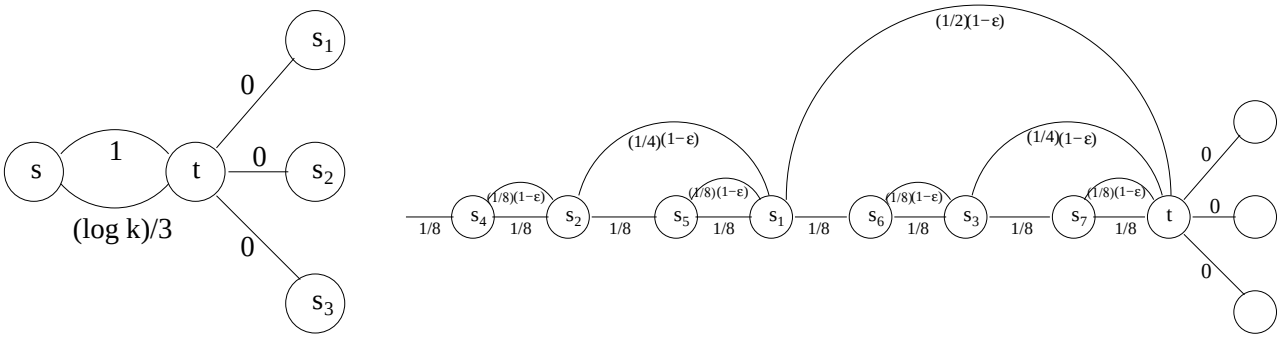

(a) Case 1: A Large Group of Similar Players

(b) Case 2: Many Distinct Groups

Figure 1: Proof of Theorem 4.3. Networks that induce games with large POA for uniform protocols with many similar players and for those with many different players, respectively. In (b), the precise value of ϵ is different for different edges (see text).

In the second case, there are at least $k/\log k$ distinct non-empty groups. Our choice of α ensures that $\alpha^{a \log^5 k} \geq k+1$ for a sufficiently large constant a that is independent of k. We can therefore pick one player out of every $\Theta(\log^5 k)$ groups to obtain a set S of $2^p - 1$ players such that: p is a positive integer; $|S| = \Theta(k/\log^6 k)$; and every pair j, h of distinct players of S either come from different A_i's or have f_i-values that differ by at least a $(k+1)$ factor. We rename players of S so that they are ordered according to the A_i's, and in decreasing order of f_i-values within an A_i.

We next construct a single-sink undirected network (Figure 1(b)). Each player not in S again has a direct zero-cost edge to t. The rest of the network is similar to a lower bound construction for the online Steiner tree problem [28] and is defined in rounds. In the 0th round, we add a unit-cost edge (the *main path*) incident to t. For $r = 1, \ldots, p$, in the rth round, we bisect the 2^{r-1} edges of the main path created in previous rounds with the sources $s_{2^{r-1}}, \ldots, s_{2^r-1}$ — that is, each such edge is replaced by two edges (in series) with half the cost, with a new source vertex between them. Additionally, we create a *shortcut* edge between each source added in the rth round and its neighbor on the main path that is closer to t, of cost $2^{-r}((k-1)/k)^{p-r+1}$. The cost of all of the shortcuts added in a round is $\Omega(1)$. The optimal outcome, with all players following the main path, has unit cost.

The outcome in which all players completely eschew the main path has cost $\Omega(p) = \Omega(\log k)$, and we complete the proof by showing that it is a PNE. Consider this outcome and a player j using the path P_j. The first edge of P_j is the shortcut (s_j, s_h). Every s_j-s_h path other than the shortcut comprises only edges of the main path (currently used by no players) and edges added in subsequent rounds (each currently used only by players with larger index). The claim is that player j's cost shares on such a path would total at least a $(k-1)/k$ fraction of the path cost, and this is no less than the cost incurred on its shortcut. To argue the claim, consider the players with which j would share edges on this s_j-s_h path: there are some from later groups A_i (who pay nothing in j's presence), and some players in the same group A_i as j but which have f_i-values (relative to $f_i(j)$) that are bounded above by a geometric sequence with ratio $1/(k+1)$. The Separation Lemma implies that the sum of the cost shares of all of these players is at most a $1/k$ fraction of the path cost, which implies the claim.

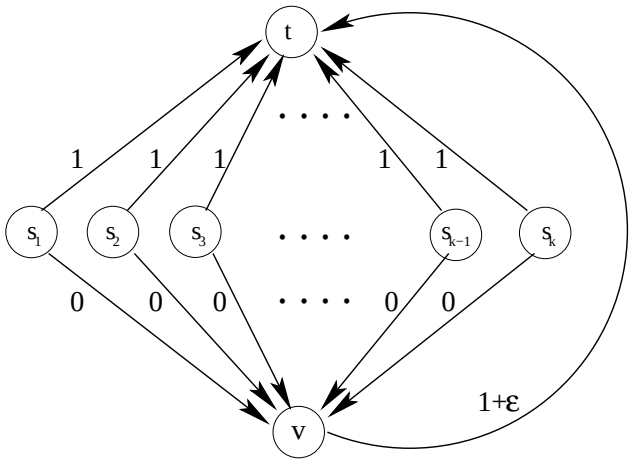

Figure 2: Proof of Proposition 4.12. Every uniform protocol has worst-case POA at least k in directed single-sink networks.

After player j reaches s_h, it shares each remaining edge of P_j with at least one player of lower index. By the Separation Lemma, the sum of its cost shares on these edges is strictly less than $1/k$, a lower bound on the cost it would incur for being the sole inhabitant of some edge on the main path. Thus P_j is a best response for player j, and this outcome is a PNE. ■

4.2 Directed Graphs: Optimality of the Shapley Protocol

We now turn to directed networks, and prove that the Shapley cost-sharing protocol is the optimal uniform protocol. We first note that the worst-case POA is a useless measure for differentiating between competing protocols in directed networks.

Proposition 4.12 (POA in Directed Networks) *For every $k \geq 1$, the worst-case POA of every uniform protocol in k-player directed networks is k. The lower bound holds even in single-sink directed networks.*

Proof: Fix k and a uniform protocol. For the upper bound, consider a directed network and the induced network cost-sharing game. Let $(P_1^*, \ldots, P_k^*)$ denote an optimal outcome, with cost C^*. Every player i can guarantee itself a cost of at most C^*, independent of the protocol and the other players' choices, by choosing the path P_i^*. Thus in a PNE, the cost of each player is at most C^*. By budget-balance, the cost of every PNE is at most $k \cdot C^*$.

The lower bound is provided by the single-sink network shown in Figure 2. The optimal outcome has cost $1 + \epsilon$, where $\epsilon > 0$ is arbitrarily small. The outcome in which each player i selects the direct $s_i \to t$ path is a PNE (for any protocol) and has cost k. ■

Remark 4.13 Proposition 4.12 and its proof hold even for *non-uniform* cost-sharing protocols (see Section 6.3).

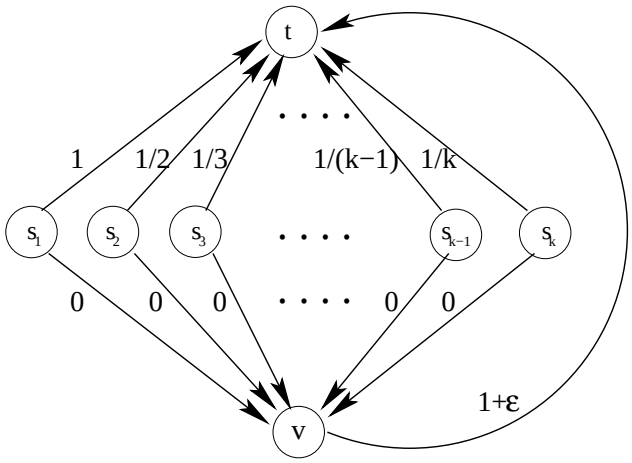

Figure 3: Proof of Theorem 4.15. Every linear monotone protocol has worst-case POS at least $\mathcal{H}_k$ in directed single-sink networks.

Proposition 4.12 also holds, with the same proof, for the worst-case reachable POA in (single-sink) directed networks. We therefore resort to our weakest inefficiency measure, the price of stability (POS). Anshelevich et al. [2] proved the following tight guarantee for the Shapley protocol.

Proposition 4.14 (POS of the Shapley Protocol [2]) *For every $k \geq 1$, the worst-case POS of the Shapley protocol in k-player directed networks is the kth Harmonic number $\mathcal{H}_k$: $1+\frac{1}{2}+\cdots+\frac{1}{k} = \ln k + O(1)$.*

Using the technical tools already developed in this and the previous section, we can quickly prove that the Shapley protocol is optimal.

Theorem 4.15 (Optimality of the Shapley Protocol) *For every $k \geq 1$, every uniform cost-sharing protocol has a worst-case POS of at least $\mathcal{H}_k$ in k-player single-sink directed networks.*

Proof: Fix $k \geq 1$ and a uniform cost-sharing protocol for the player set $\{1, 2, \ldots, k\}$. By Lemmas 4.4–4.6 and Remark 4.7, we can assume that the protocol ξ is linear. Lemma 3.10 ensures that ξ is *monotone* in the sense that $\xi(i, S) \geq \xi(i, T)$ whenever $i \in S \subseteq T \subseteq \{1, \ldots, k\}$. Since ξ is budget-balanced, there is a player i_k satisfying $\xi(i_k, \{1, 2, \ldots, k\}) \geq 1/k$. Similarly, for each $j = k-1, \ldots, 1$, there is a player i_j satisfying $\xi(i_j, \{1, \ldots, k\} \setminus \{i_{j+1}, \ldots, i_k\}) \geq 1/j$.

Now consider the single-sink directed network of Figure 3, taken from [2]. There is a sink t, source vertices $s_1, \ldots, s_k$, and an additional vertex v. Player i_j has source s_j and sink t. For each j, there is an edge of cost $1/j$ from s_j to t and an edge of zero cost from s_j to v. There is also an edge of cost $1+\epsilon$ from v to t. The optimal solution has cost $1+\epsilon$. On the other hand, we claim that in the network cost-sharing game induced by ξ, the only PNE has cost $\mathcal{H}_k$. To see this, consider an arbitrary outcome of the game, and let $S \subseteq \{1, 2, \ldots, k\}$ denote the set of players that share the edge (v, t), with the rest of the players choosing their one-hop paths to t. Suppose that $S \neq \emptyset$ and let $i_j \in S$ be the player of S with maximum index j. By construction, $\xi(i_j, \{i_1, \ldots, i_j\}) \geq 1/j$. Since ξ is monotone and linear, player i_j incurs cost at least $(1+\epsilon)/j$ in this outcome. Since player

i_j incurs cost only $1/j$ by choosing the path $s_j \to t$, this outcome cannot be a PNE. Thus $S = \emptyset$ in the unique PNE, which has cost $\mathcal{H}_k$. ■

4.3 Convergence of Better-Response Dynamics

Our final application of Theorem 3.8 is technically straightforward but conceptually interesting. Guarantees that equilibria (such as PNE) always exist in a family of games are important but lack predictive power: how do we know that rational players will successfully *reach* such an equilibrium? A guarantee that selfish participants converge to an equilibrium through repeated experimentation is much more compelling that a mere existence result.

Better-response dynamics is a simple and well-studied model of repeated experimentation: while the current outcome is not a PNE, choose an arbitrary player that could decrease its cost by switching paths, and update its path to a better one. If better-response dynamics is guaranteed to converge in a game, then the game obviously has at least one PNE; the converse generally fails (see e.g. [47]). But Theorem 3.8 implies a converse of sorts for network cost-sharing games: the *only way* to guarantee the existence of PNE in all such games is to guarantee the convergence of better-response dynamics.

Theorem 4.16 (Convergence of Better-Response Dynamics) *In every network cost-sharing game induced by a linear and uniform protocol, better-response dynamics always converges to a PNE.*

Proof: First consider a potential-based protocol ξ for a player set $\{1, 2, \ldots, k\}$. Every network cost-sharing game induced by ξ admits a weighted potential function Φ, given in (3), and every time a player changes its path in better-response dynamics, Φ strictly decreases (4). Since there are only finitely many outcomes, better-response dynamics must terminate.

For an arbitrary linear and uniform protocol ξ, we can apply Theorem 3.8 to express ξ as the concatenation of potential-based protocols $\xi_1, \ldots, \xi_m$ for disjoint player sets $A_1, \ldots, A_m$. Players of A_1 are unaffected by the choices of other players, so the above argument can be applied to ξ_1 and A_1 to conclude that, in an arbitrary network cost-sharing game induced by ξ, better-response dynamics eventually reaches an outcome from which no player of A_1 has an incentive to switch paths. This property cannot be violated by subsequent moves by players from $A_2, \ldots, A_m$, so no player of A_1 will ever change paths again. Proceeding inductively on the A_i's, we conclude that better-response dynamics eventually terminates. ■

Theorem 4.16 implies that, for linear protocols, strengthening the stability constraint (2) in Section 1.2 to require the convergence of better-response dynamics has no effect on the protocol design space.

5 Characterization of Uniform Protocols: The Proof

This section presents a proof of Theorem 3.8. The "if" direction is immediate from Propositions 3.4 and 3.7. The four major milestones of the "only if" direction are as follows.

1. Every linear and uniform protocol ξ must be *monotone* in the sense that $\xi(i, S) \geq \xi(i, T)$ whenever $i \in S \subseteq T \subseteq \{1, 2, \ldots, k\}$. (Recall Lemma 3.10.)

2. For every linear and uniform protocol ξ, the players can be partitioned into ordered equivalence classes so that $\xi(i, S) > 0$ if and only if i belongs to the lowest-indexed class that intersects S. Different equivalence classes correspond to disjoint player sets that are combined via concatenation.

3. For every linear and uniform protocol ξ that is also positive, all of its cost shares are uniquely determined by the $k-1$ pairwise cost shares $\xi(1, \{1,2\})$, $\xi(1, \{1,3\})$, $\ldots$, $\xi(1, \{1,k\})$.

4. For every linear and uniform protocol ξ that is also positive, ξ is potential-based.

We work with undirected networks throughout the proof. Our proof can be modified trivially to use only directed networks.

Our first technical lemma builds toward the first milestone of the proof by showing that any failures of monotonicity in a linear uniform protocol must be "symmetric". The proof of this lemma also develops arguments crucial to the second milestone of the proof of Theorem 3.8 (see Lemma 5.4, below).

Lemma 5.1 *Let ξ be a linear and uniform cost-sharing protocol with player set $\{1, 2, \ldots, k\}$. Let S be a non-empty set of players and i and j distinct players not in S, and suppose that $\xi(i, S \cup \{i\}) < \xi(i, S \cup \{i,j\})$. Then $\xi(j, S \cup \{j\}) \leq \xi(j, S \cup \{i,j\})$.*

Proof: Fix ξ, i, j, and S, and assume that $\xi(i, S \cup \{i\}) < \xi(i, S \cup \{i,j\})$; in particular, $\xi(i, S \cup \{i,j\}) > 0$. The heart of the proof is the following claim: for all positive constants α and β, if

$$\alpha\left(\xi(i, S \cup \{i\}) - \xi(i, S \cup \{i,j\})\right) < \beta\left(\xi(i, \{i\}) - \xi(i, \{i,j\})\right), \tag{7}$$

then

$$\alpha\left(\xi(j, S \cup \{j\}) - \xi(j, S \cup \{i,j\})\right) \leq \beta\left(\xi(j, \{j\}) - \xi(j, \{i,j\})\right). \tag{8}$$

To see why the claim implies the lemma, set $\beta = 1$. Since $\xi(i, \{i\}) = 1$ and $\xi(i, \{i,j\}) \leq 1$, the right-hand side of (7) is non-negative. Since $\xi(i, S \cup \{i\}) < \xi(i, S \cup \{i,j\})$ by assumption, inequality (7) holds for all values of $\alpha > 0$. By the claim, inequality (8) holds for every $\alpha > 0$. Since the right-hand side of (8) is nonnegative, we conclude that $\xi(j, S \cup \{j\}) \leq \xi(j, S \cup \{i,j\})$.

To prove the claim, fix α and β and consider the network shown in Figure 4. All players of S have source s and sink t. Players outside $S \cup \{i,j\}$ are confined to a different, disjoint subnetwork. The parameter M is sufficiently large relative to α, β, and $1/\xi(i, S \cup \{i,j\})$. The parameter A is at least M^2; its precise value will be fixed later. We show that if the claim fails, then we can choose edge costs so that the game induced by ξ in this network has no PNE, thereby contradicting the stability of ξ.

We first argue that, at equilibrium, no player chooses a path containing both the edges (s_i, u) and (s_i, v). This is true for player i because it must choose a simple s_i-t_i path. For a contradiction, assume that at least one player of $S \cup \{j\}$ chooses a path containing both these edges, and let e be one of the two edges that is not used by player i. By budget-balance, one of the players using e must pay at least M^2/k for it. No player of S will incur such a large cost share at equilibrium because of the deviation $s \to u \to w \to x \to v \to t$ (assuming M is sufficiently large). Player j will not incur such a cost share at equilibrium, because the w-v subpath of its path that contains e can be replaced by the subpath $v \to x \to w$ (removing the resulting cycles if necessary) to decrease its cost. Thus no player will use edge e in a PNE.

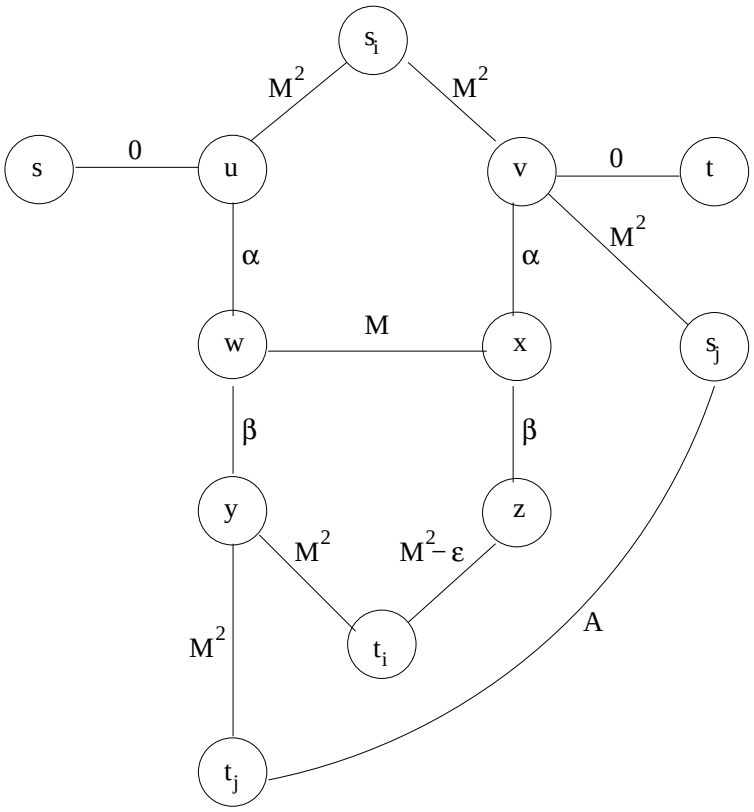

Figure 4: Network in the proofs of Lemmas 5.1 and 5.4.

A similar argument shows that no player will use the edges (y, t_i) and (z, t_i) in a PNE. Thus no player other than i uses an edge incident to s_i or t_i in a PNE. So, player j uses either the path $s_j \to v \to x \to w \to y \to t_j$ (denoted Q_1) or the one-hop path $s_j \to t_j$ (denoted Q_2) in a PNE.

Next we argue that no player of $S \cup \{i\}$ uses the edge (s_j, t_j) in a PNE. Player j uses at most two of the three edges (t_j, y), (s_j, t_j), and (v, s_j). If any players of $S \cup \{i\}$ use the edge (s_j, t_j), one of them incurs a cost share of at least M^2/k for one of these three edges. (Recall that A is at least M^2.) As above, no player of S will incur such a large cost share in a PNE. If player i incurs a cost share of at least M^2/k on these edges, then its overall cost is at least $M^2(2 + 1/k)$ (since no players other than i use edges incident to s_i or t_i). Since player i can guarantee itself a cost of at most $2M^2 + \alpha + \beta$, this cannot occur in a PNE provided M is sufficiently large.

Summarizing, player j will take path Q_1 or Q_2 in a PNE, and all players of S will follow the path $s \to u \to w \to x \to v \to t$. Let P_1 and P_2 denote the s_i-t_i paths $s_i \to u \to w \to y \to t_i$ and $s_i \to v \to x \to z \to t_i$, respectively. If player j chooses path Q_1, then our assumption that $\xi(i, S \cup \{i, j\}) > 0$ implies that player i will not use the edge (w, x) in a PNE (assuming M is sufficiently large) and hence uses P_1 or P_2. If player j chooses path Q_2, then the cost incurred by player i on path P_2 is no more than that on any other s_i-t_i path. Thus, if there is a PNE, there is one in which player i chooses either P_1 or P_2. We label the four candidates for a PNE according to the paths selected by players i and j, and proceed to choose parameters to rule them out.

We have $c_j(P_1, Q_2) = c_j(P_2, Q_2) = A$, and

$$\begin{aligned}
c_i(P_1, Q_1) &= 2M^2 + \alpha\xi(i, S \cup \{i\}) + \beta\xi(i, \{i, j\});\\
c_i(P_2, Q_1) &= 2M^2 + \alpha\xi(i, S \cup \{i, j\}) + \beta - \epsilon;\\
c_i(P_1, Q_2) &= 2M^2 + \alpha\xi(i, S \cup \{i\}) + \beta;\\
c_i(P_2, Q_2) &= 2M^2 + \alpha\xi(i, S \cup \{i\}) + \beta - \epsilon;\\
c_j(P_1, Q_1) &= 2M^2 + M\xi(j, \{S \cup \{j\}) + \alpha\xi(j, S \cup \{j\}) + \beta\xi(j, \{i, j\});\\
c_j(P_2, Q_1) &= 2M^2 + M\xi(j, \{S \cup \{j\}) + \alpha\xi(j, S \cup \{i, j\}) + \beta.
\end{aligned}$$

If (7) holds but (8) fails and ϵ is sufficiently small, then $c_i(P_1, Q_1) < c_i(P_2, Q_1)$. Setting the cost of edge (s_j, t_j) to be a number A satisfying $c_j(P_2, Q_1) < A < c_j(P_1, Q_1)$, we obtain a network game induced by ξ with no PNE: player i wants to deviate from (P_2, Q_1) and (P_1, Q_2), while player j wants to deviate from (P_1, Q_1) and (P_2, Q_2). (Note that $A \geq M^2$, as required.) This contradicts the stability of ξ, completing the proof. ■

The next lemma is a restatement of Lemma 3.10, and establishes the monotonicity of linear uniform protocols.

Lemma 5.2 (Monotonicity of Linear Uniform Protocols) *Every linear and uniform protocol ξ for a player set $\{1, \ldots, k\}$ is* monotone*: for every $S \subseteq T \subseteq \{1, \ldots, k\}$ and $i \in S$, $\xi(i, S) \geq \xi(i, T)$.*

Proof: We show if ξ is linear and uniform but not monotone, then it is not stable (a contradiction). By definition, if ξ is not monotone, there is a set $S \subseteq \{1, \ldots, k\}$ and a pair $i, i' \in S$ of players such that $\xi(i, S \backslash \{i'\}) < \xi(i, S)$. This can only occur if S contains a player other than i, i'. Among all such sets, choose one of minimum-possible cardinality. We assume henceforth that $\xi(i', S \setminus \{i\}) \leq \xi(i', S)$ as well; otherwise, by Lemma 5.1, the proof is complete.

Since the cost-sharing method ξ is budget-balanced, we have

$$\sum_{j \in S \backslash \{i\}} \xi(j, S \setminus \{i\}) \; + \sum_{j \in S \backslash \{i'\}} \xi(j, S \setminus \{i'\}) = 2 = \sum_{j \in S} \xi(j, S) \; + \sum_{j \in S \backslash \{i, i'\}} \xi(j, S \setminus \{i, i'\}).$$

Thus,

$$\begin{aligned}
\xi(i, S) + \xi(i', S) \;\; &+ \sum_{j \in S \backslash \{i, i'\}} \left[\xi(j, S) + \xi(j, S \setminus \{i, i'\})\right] =\\
\xi(i, S \setminus \{i'\}) + \xi(i', S \setminus \{i\}) \;\; &+ \sum_{j \in S \backslash \{i, i'\}} \left[\xi(j, S \setminus \{i\}) + \xi(j, S \setminus \{i'\})\right].
\end{aligned}$$

Since $\xi(i, S \setminus \{i'\}) < \xi(i, S)$ and $\xi(i', S \setminus \{i\}) \leq \xi(i', S)$, there is a player $j \in S \setminus \{i, i'\}$ with

$$\xi(j, S) + \xi(j, S \setminus \{i, i'\}) < \xi(j, S \setminus \{i\}) + \xi(j, S \setminus \{i'\}).$$

Now consider the network shown in Figure 5. All players of $S \setminus \{i, i', j\}$ have source s and sink t. As in the proof of Lemma 5.1, players outside of S are confined elsewhere. The parameter M is a

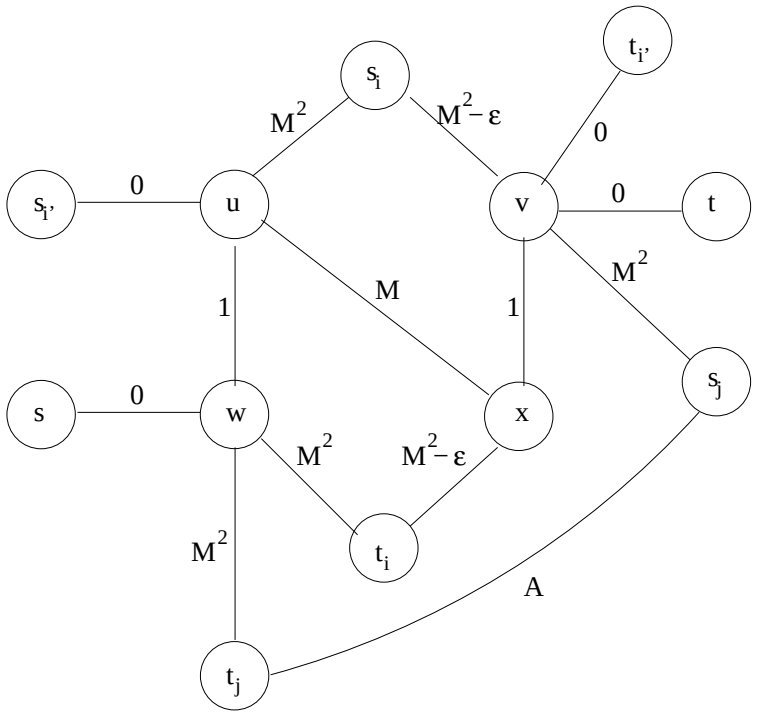

Figure 5: Network in the proof of Lemma 5.2.

sufficiently large function of $\xi(i, S)$. (By assumption, $\xi(i, S) > \xi(i, S \setminus \{i'\}) \geq 0$.) The parameter ϵ is small enough so that $\xi(i, S \setminus \{i'\}) + 2\epsilon < \xi(i, S)$. The parameter A is chosen to satisfy

$$2M^2+M\xi(j,S\setminus\{i\})+\xi(j,S)+\xi(j,S\setminus\{i,i'\}) < A < 2M^2+M\xi(j,S\setminus\{i\})+\xi(j,S\setminus\{i\})+\xi(j,S\setminus\{i'\}). \tag{9}$$

We claim that the network cost-sharing game induced by ξ in this network admits no PNE.

First, note that player i can guarantee itself a cost share of at most $2M^2 + 1$; player j a cost share of at most $2M^2 + M + 2$; and all other players a collective cost share of at most $M + 2$. Hence, in every PNE, the total cost incurred by all players is strictly less than $5M^2 - 2\epsilon$ (provided M is sufficiently large). Every outcome in which both edges incident to s_i or both edges incident to t_i are used by some player has cost at least $5M^2 - 2\epsilon$—player i must use one edge incident to each of s_i and t_i, and player j must incur cost at least $2M^2$ on the edges (s_j, v), (s_j, t_j), and (w, t_j)—and thus none of these are PNE. By similar reasoning, no player other than j uses the edge (s_j, t_j) in a PNE.

We have established the following for every PNE: player i' must take the path $s_{i'} \to u \to x \to v \to t_{i'}$; players of $S \setminus \{i, i', j\}$ must take the path $s \to w \to u \to x \to v \to t$; and player j must take either the path Q_1, defined as $s_j \to v \to x \to u \to w \to t_j$, or the one-hop path Q_2 from s_j to t_j.

Let P_1 denote the path $s_i \to u \to w \to t_i$ and P_2 the path $s_i \to v \to x \to t_i$. The following four statements will complete the proof:

(1) In a PNE, if player j chooses the path Q_1, then player i chooses the path P_1.

(2) In a PNE, if player i chooses the path P_1, then player j chooses the path Q_2.

(3) In a PNE, if player j chooses the path Q_2, then player i chooses the path P_2.

(4) In a PNE, if player i chooses the path P_2, then player j chooses the path Q_1.

Statement (2) follows immediately from the second inequality in (9), and (4) follows from the first inequality in (9). (Recall that in a PNE, no player other than j uses an edge incident to s_j or t_j.) To prove (1), suppose that player j chooses path Q_1. Since $\xi(i,S) > 0$ and M is sufficiently large, player i will not use edge (u,x) in a PNE, and hence it chooses either P_1 or P_2. Since the edges incident to s_i and t_i are used by no other player, and since ϵ satisfies $\xi(i, S \setminus \{i'\}) < \xi(i,S) - 2\epsilon$, player i strictly prefers path P_1 to path P_2.

To establish (3), suppose player j chooses path Q_2. No players other than i use an edge incident to s_i or t_i, and the minimality of S implies that $\xi(i, S \setminus \{j\}) \leq \xi(i, S \setminus \{i', j\})$. These facts imply that P_2 is the unique best response of player i (provided $M \geq 1$), and the proof is complete. ■

Lemma 5.2 and budget-balance yield a useful corollary.

Corollary 5.3 *Let ξ be a linear uniform protocol for the player set $\{1,2,\ldots,k\}$. Suppose that $\xi(i,S) = 0$ for some $S \subseteq \{1,\ldots,k\}$ and $i \in S$. Then $\xi(j, S \setminus \{i\}) = \xi(j,S)$ for every $j \in S \setminus \{i\}$.*

The second milestone in the proof of Theorem 3.8 states that, for every linear uniform protocol ξ, the players can be partitioned into ordered equivalence classes so that $\xi(i,S) > 0$ if and only if i belongs to the lowest-indexed class that intersects S. This and Corollary 5.3 will imply that every linear uniform protocol is the concatenation of positive protocols in the sense of Definition 3.6. The next two lemmas work toward this goal by establishing further restrictions on how two players can influence each others cost shares.

Lemma 5.4 *Let ξ be a linear uniform protocol with player set $\{1,2,\ldots,k\}$. Let S be a non-empty set of players and i,j distinct players not in S, and suppose that $\xi(i, S \cup \{i,j\}) > 0$. Then,*

$$\Big(\xi(i,S\cup\{i\}) - \xi(i,S\cup\{i,j\})\Big)\Big(\xi(j,\{j\}) - \xi(j,\{i,j\})\Big) \quad = \\ \Big(\xi(j,S\cup\{j\}) - \xi(j,S\cup\{i,j\})\Big)\Big(\xi(i,\{i\}) - \xi(i,\{i,j\})\Big). \tag{10}$$

Proof: Fix ξ, i, j, and S. Denote the four differences in (10) by Δ_1, Δ_2, Δ_3, and Δ_4, respectively. By Lemma 5.2, these are all nonnegative numbers.

From the key claim in the proof of Lemma 5.1 (relating (7) and (8)), we know that for every pair α, β of positive constants,

$$\alpha\Delta_1 < \beta\Delta_4 \quad \Rightarrow \quad \alpha\Delta_3 \leq \beta\Delta_2. \tag{11}$$

(The proof of that key claim did not use the hypothesis that $\xi(i, S\cup\{i\}) < \xi(i, S\cup\{i,j\})$, and relied only on the fact that $\xi(i,S\cup\{i,j\}) > 0$.) Exchanging the costs of (y,t_i) and (z,t_i) in the network in Figure 4, an analogous argument proves a partial converse:

$$\alpha\Delta_3 < \beta\Delta_2 \quad \Rightarrow \quad \alpha\Delta_1 \leq \beta\Delta_4 \tag{12}$$

for each pair α, β of positive constants.

We conclude with a case analysis. If $\Delta_2 = \Delta_4 = 0$, then the lemma clearly holds. If $\Delta_2, \Delta_4 > 0$, then Δ_1/Δ_4 and Δ_3/Δ_2 are well-defined nonnegative numbers; the implications (11) and (12) imply that they are equal, and hence the lemma holds. The final two cases are symmetric; if $\Delta_2 > 0$ while $\Delta_4 = 0$ (say), then (12) implies that $\Delta_1 = 0$, and again the lemma holds. ■

Lemma 5.5 *Let ξ be a linear uniform protocol with player set $\{1, 2, \ldots, k\}$. Suppose there are distinct players i, j, ℓ such that $\xi(\ell, \{j, \ell\}) = 0$ and $\xi(i, \{i, j\}) > 0$. Then $\xi(\ell, \{i, \ell\}) = 0$.*

Proof: By monotonicity (Lemma 5.2), $\xi(\ell, \{i, j, \ell\}) = 0$. By Corollary 5.3, $\xi(i, \{i, j, \ell\}) = \xi(i, \{i, j\}) > 0$ and $\xi(j, \{i, j, \ell\}) = \xi(j, \{i, j\})$. We can apply Lemma 5.4 with $S = \{\ell\}$ to obtain

$$\Big(\xi(i, \{i, \ell\}) - \xi(i, \{i, j, \ell\})\Big)\Big(1 - \xi(j, \{i, j\})\Big) = \Big(\xi(j, \{j.\ell\}) - \xi(j, \{i, j, \ell\})\Big)\Big(1 - \xi(i, \{i, j\})\Big). \quad (13)$$

By hypothesis and budget-balance, $\xi(j, \{j, \ell\}) = 1 - \xi(\ell, \{j, \ell\}) = 1$. Rewriting (13) we obtain

$$\Big(\xi(i, \{i, \ell\}) - \xi(i, \{i, j\})\Big)\Big(1 - \xi(j, \{i, j\})\Big) = \Big(1 - \xi(j, \{i, j\})\Big)\Big(1 - \xi(i, \{i, j\})\Big).$$

Since $1 - \xi(j, \{i, j\}) = \xi(i, \{i, j\}) > 0$ by assumption, we can divide through to obtain $\xi(i, \{i, \ell\}) = 1$, which completes the proof. ■

The next lemma and its corollary completes the second milestone of the proof of Theorem 3.8.

Lemma 5.6 *Let ξ be a linear uniform protocol for the player set $\{1, 2, \ldots, k\}$. Then there exists a partition of the players into classes $\mathcal{C}_1, \ldots, \mathcal{C}_\ell$ with the following property: for every subset S of players and $i \in S$, $\xi(i, S) > 0$ if and only if i is contained in the lowest-indexed class that intersects S.*

Proof: Define a relation on players by $i \to j$ if and only if $\xi(i, \{j, i\}) > 0$. Lemma 5.5 implies that this relation is transitive. Interpret this relation as a directed graph H (where the relation $i \to j$ yields the arc (i, j)). By budget-balance, for every pair i, j of players, either $i \to j$ or $j \to i$ (or both); thus H includes a tournament as a subgraph. As a consequence, there is a unique topological ordering $\mathcal{C}_1, \ldots, \mathcal{C}_\ell$ of the strongly connected components of H. By transitivity of $\to$, each C_h is a complete directed graph.

Next, consider a subset T of players and a player $i \in T \setminus C_h$, where C_h denotes the lowest-indexed component that intersects T. Since $i \notin C_h$, there is a player $j \in T \cap C_h$ with $i \not\to j$ and hence $\xi(i, \{i, j\}) = 0$. By monotonicity (Lemma 5.2), $\xi(i, T) = 0$.

To complete the proof, assume for contradiction that there is a set $T \subseteq \{1, 2, \ldots, k\}$ and a player $i \in T \cap C_h$ with $\xi(i, T) = 0$, where C_h denotes the lowest-indexed component that intersects T. Among all sets T with this property, choose one of minimum cardinality. By the previous paragraph and Corollary 5.3, $T \subseteq C_h$. (If $T \not\subseteq C_h$, then deleting a player of $T \setminus C_h$ yields a smaller set with the same cost shares, contradicting the minimality of T.) Since C_h is a complete directed graph, the definition of $\to$ implies that $|T| \geq 3$.

Pick $j \in T \setminus \{i\}$ arbitrarily. By Corollary 5.3, $\xi(j, T \setminus \{i\}) = \xi(j, T)$. By the minimality of T, $\xi(i, T \setminus \{j\}) > 0$. Applying Lemma 5.4 with $S = T \setminus \{i, j\}$, we have $\xi(j, \{i, j\}) = \xi(j, \{j\}) = 1$ and hence $i \not\to j$. This contradicts the fact that T is a complete directed graph. ■

Corollary 5.7 (Linear Uniform Protocols are Concatenations of Positive Protocols) *A cost-sharing protocol is linear and uniform only if it is the concatenation of positive protocols.*

Proof: Let ξ be a linear uniform protocol for the player set $\{1, 2, \ldots, k\}$. Partition the player set into classes $A_1, \ldots, A_m$ according to Lemma 5.6. For a subset T lying entirely in some A_h and a subset $S \subseteq \{1, 2, \ldots, k\}$, we say that S *induces* T if S contains no players of $A_1, \ldots, A_{h-1}$ and if $S \cap A_h = T$. Lemma 5.6 implies that the players i with $\xi(i, S) > 0$ are precisely those in the induced subset T. Corollary 5.3 implies that the values of these cost shares depend only on T and not on S's intersection with $A_{h+1}, \ldots, A_m$. For each h, we can therefore define a linear uniform protocol ξ_h on A_h by setting $\xi_h(i, T)$ equal to i's cost share $\xi(i, S)$ in every subset S that induces T. Then $\xi = \xi_1 \oplus \xi_2 \oplus \cdots \oplus \xi_m$, as desired. ■

To complete the proof of Theorem 3.8, we need to show that every linear uniform cost-sharing protocol that is also positive is potential-based in the sense of Proposition 3.4. The third milestone in the proof states that every such protocol is uniquely determined by its *pairwise cost shares* for a given player (player 1, say): the cost shares of the player in the $k - 1$ player pairs to which it might belong. This fact is proved over the next three lemmas. The first is a calculation.

Lemma 5.8 *Let ξ be a linear uniform protocol for the player set $\{1, 2, \ldots, k\}$ that is also positive. For every three players i, j, ℓ,*

$$\xi(i, \{i, j\}) \times \xi(j, \{j, \ell\}) \times \xi(\ell, \{\ell, i\}) = \xi(j, \{i, j\}) \times \xi(\ell, \{j, \ell\}) \times \xi(i, \{\ell, i\}). \tag{14}$$

Proof: Denote the three cost shares on the left-hand side of (14) by x, y, and z, respectively; the lemma is then equivalent to the identity $xyz = (1 - x)(1 - y)(1 - z)$. Denote the cost shares $\xi(i, \{i, j, \ell\})$, $\xi(j, \{i, j, \ell\})$, and $\xi(\ell, \{i, j, \ell\})$ by a, b, and c, respectively. The plan of the proof is to use Lemma 5.4 to express x, y, z in terms of a, b, c, apply the budget-balance constraint $a+b+c = 1$, and then solve for x, y, z. By positivity of ξ, x, y, and z are strictly between 0 and 1.

An application of Lemma 5.4 (with $S = \{\ell\}$ and i, j as themselves) gives

$$[(1 - z) - a]x = (y - b)(1 - x); \tag{15}$$

another application, after permuting the roles of the three players, yields

$$[(1 - y) - c]z = (x - a)(1 - z).$$

Solving for b and c gives

$$b = \frac{xa - x(1 - z) + y(1 - x)}{1 - x}$$

and

$$c = \frac{(1 - z)(a - x) + z(1 - y)}{z}.$$

Applying the budget-balance constraint $a + b + c = 1$ we can solve for a in terms of x, y, z:

$$a = \frac{(1 - z)(x - x^2 + xz)}{1 - x + xz}.$$

A symmetric argument (with the substitutions $x \to y$, $y \to z$, and $z \to x$) yields a solution for b:

$$b = \frac{(1 - x)(y - y^2 + yx)}{1 - y + yx}.$$

Substituting into (15) and clearing denominators gives

$$x(1-y+yx)(1-z)(1-x)^2 = (1-x)(1-x+xz)yx^2;$$

cancelling through and subtracting $xy(1-x)(1-z)$ from both sides yields $(1-x)(1-y)(1-z) = xyz$, as desired. ■

Lemma 5.9 *Let ξ_1, ξ_2 be linear uniform protocols for the player set $\{1, 2, \ldots, k\}$ that are also positive. If ξ_1 and ξ_2 have identical pairwise cost shares for some player i, then they have identical pairwise cost shares for all players.*

Proof: For every pair j, ℓ of other players, the pairwise cost shares for player i and budget-balance determine the first and third terms on both sides of (14). Equation (14), positivity, and budget-balance then uniquely determine the second terms of both sides. ■

The next lemma completes the third step of the proof of Theorem 3.8.

Lemma 5.10 (Pairwise Cost Shares Determine a Positive Protocol) *Let ξ_1, ξ_2 be linear uniform protocols for the player set $\{1, 2, \ldots, k\}$ that are also positive. If ξ_1 and ξ_2 have identical pairwise cost shares for some player i, then they are identical cost-sharing protocols.*

Proof: We show that for every $m \geq 3$, the cost shares of a positive, linear uniform protocol for subsets of at most $m-1$ players uniquely determine those for each subset of m players; the lemma then follows from Lemma 5.9 and induction.

Consider a positive, linear uniform protocol ξ for $\{1, 2, \ldots, k\}$ and a set S of $m \geq 3$ players. For every distinct $i, j \in S$, we can use Lemma 5.4 (with $S \setminus \{i, j\}$ playing the role of S) and positivity to write

$$\xi(j, S) = \xi(j, S \setminus \{i\}) + \frac{1 - \xi(j, \{i, j\})}{1 - \xi(i, \{i, j\})}(\xi(i, S) - \xi(i, S \setminus \{j\})).$$

This equation shows that the given cost shares for all subsets with at most $m-1$ players and a choice of a cost share $\xi(i, S)$ uniquely determine the cost share $\xi(j, S)$ for every other player j of S. Moreover, the cost shares $\{\xi(j, S)\}_{j \neq i}$ are strictly increasing with the choice of $\xi(i, S)$. There can be only one choice of $\xi(i, S)$ that satisfies the budget-balance constraint $\sum_{j \in S} \xi(j, S) = 1$, completing the proof. ■

The fourth and final main step in the proof of Theorem 3.8 is to show that the only positive, linear uniform protocols are potential-based. The idea is that every set of positive pairwise cost shares for a player can be extended to a potential-based protocol; and by Lemma 5.10, this is the only possible extension.

Lemma 5.11 (Positive Protocols Are Potential-Based) *If ξ is a positive linear uniform protocol for the player set $\{1, 2, \ldots, k\}$, then ξ is potential-based.*

Proof: Given ξ, set $a_i = \xi(1, \{1, i\})$. Define $f(\{1\}) = 1$ and, for $i \in \{2, \ldots, k\}$, assign $f(\{i\})$ so that $1/[1 + f(\{i\})] = a_i$. Extend f inductively to all of $2^{\{1,\ldots,k\}}$ using (5) and let $\hat{\xi}$ denote the protocol induced by f via equation (2). By construction and equations (1) and (2), ξ and $\hat{\xi}$ have identical pairwise cost shares for player 1. Lemma 5.10 and its inductive proof show that $\xi = \hat{\xi}$ everywhere. Since $\hat{\xi}$ is potential-based by construction, ξ is also potential-based. ■

Proof of Theorem 3.8: The "if" direction follows from Propositions 3.4 and 3.7. For the "only if" direction, let ξ be a linear uniform cost-sharing protocol. By Corollary 5.7, we can write $\xi = \xi_1 \oplus \cdots \oplus \xi_p$ for protocols $\xi_1, \ldots, \xi_p$ that are linear, uniform, and positive. Lemma 5.11 implies that each ξ_i is potential-based, completing the proof. ■

6 Extensions

This section studies two extensions to the basic protocol design problem: relaxing the design space to include non-uniform protocols (Section 6.1 and 6.3), and allowing a player to withdraw from the game in favor of an "outside option" (Subsection 6.4). We also discuss the applications to approximation algorithms of one of our constructions (Section 6.2).

6.1 Non-Uniform Protocols in Undirected Graphs

Recall from Definition 2.1 that a (non-uniform) cost-sharing protocol assigns a cost-sharing method ξ_e to every edge of a network G. Unlike uniform protocols, the definition of ξ_e can depend on the network G and the edge e. We continue to require admissibility (i.e., budget-balance and stability).

Example 2.7 and Theorem 4.3 rigorously separate the power of uniform and non-uniform protocols: in single-sink networks with k players, there is a non-uniform protocol with worst-case POA at most 2, while every uniform protocol has worst-case POA $\Omega(\log k)$. We now systematically study the power and limitations of non-uniform protocols.

This section considers undirected networks. Example 2.7 and Remark 2.9 resolve the best-possible worst-case POA in single-sink networks, so we turn to multicommodity networks. Is a constant (independent of k) worst-case POA achievable by a non-uniform protocol? For example, there is a natural analogue of the Prim cost-sharing protocol in such networks: the first player i is the one minimizing the distance between its source s_i and sink t_i; the second player is the one minimizing the distance between its source and sink, after all edges in the shortest s_i-t_i path have been reset to zero; and so on. Our characterization result (Theorem 3.8) pertains only to uniform protocols and offers no assistance. Nevertheless, we can devise an explicit multicommodity example that provides a logarithmic lower bound on the best-possible worst-case POA of non-uniform protocols.

Theorem 6.1 *For all sufficiently large k, the worst-case POA of every (non-uniform) admissible cost-sharing protocol for k-player undirected multicommodity networks is $\Omega(\log k)$.*

Theorem 6.1 and Proposition 4.1 imply that, unlike in single-sink networks, non-uniform cost-sharing protocols provided little advantage over uniform ones.

The proof of Theorem 6.1 is based on the following combinatorial lemma.

Lemma 6.2 *For all sufficiently large n, there exists a 3-regular graph $G = (V, E)$ with $2n$ vertices and a perfect matching M in G satisfying the following two properties. First, deleting all of the edges of M yields a graph with $O(n/\log n)$ connected components. Second, contracting all of the edges of M yields a graph with girth $\Omega(\log n)$.*

Proof: For sufficiently large n, we begin with a 4-regular graph $G_1 = (V_1, E_1)$ that has n vertices and girth $\Omega(\log n)$. Such a graph exists due to, for example, a construction of Erdös and Sachs (see [46, Exercise 15.3.1]). First, we pick a cycle cover (a 2-regular subgraph, not necessarily connected)

of G_1; an application of Hall's Theorem (e.g. [58, Theorem 22.1]) shows that one exists. Let C_1 denote the edges in this cover. Since G is 4-regular, the set C_2 of the rest of the edges of G_1 is also a cycle cover. Since G has girth $\Omega(\log n)$, C_1 and C_2 each contain $O(n/\log n)$ cycles.

Letting $\{v_1, v_2, \ldots, v_n\}$ denote the vertices of G_1, we construct the following 3-regular graph G. There are $2n$ vertices $v_{11}, v_{12}, \ldots, v_{1n}, v_{21}, v_{22}, \ldots, v_{2n}$. For every edge e between v_i and v_j in the original graph, we add an edge between v_{1i} and v_{1j} if e is in cycle cover C_1 and add an edge between v_{2i} and v_{2j} if e is in the cycle cover C_2. We also add an edge between v_{1i} and v_{2i} for every i and call the resulting graph G.

First, G certainly has $2n$ vertices and is 3-regular. It also contains a perfect matching M consisting of the edges (v_{1i}, v_{2i}) for every i. If we contract all of the edges of M, then we recover the graph G_1, which has girth $\Omega(\log n)$. If we delete every edge of M, then each connected component that remains corresponds to a cycle in the graph G_1 and hence contains $\Omega(\log n)$ vertices. Since G has $2n$ nodes, there are at most $O(n/\log n)$ connected components after we delete the edges of M. ■

Proof of Theorem 6.1: Let $G = (V, E)$ be the graph described in Lemma 6.2. Let G_1 be the graph obtained from G by contracting the perfect matching M. G_1 has girth at least $2c \log n$ for some constant $c > 0$. Assign cost $c \log n$ to all edges of M in G. All other edges in G have cost 1.

We add an additional vertex v and, for every connected component C of $G \setminus M$, we add an edge with cost $2c \log n$ between v and an arbitrary vertex of C. There are n players in the game, with one for each edge e of M. The endpoints of e are the source and sink vertices of the corresponding player.

Now fix arbitrary cost-sharing methods for the edges of G. We claim that the outcome in which every player chooses its one-hop path is a PNE with respect to these cost-sharing methods. First, every deviation from this outcome must use either an edge incident to the extra vertex v or all of the edges of a cycle in the graph G_1. Since all such edges are currently unused by all of the players, the budget-balance constraint ensures that the deviating player must pay their full cost. Since each edge incident to v has cost $2c \log n$ and G_1 has girth $2c \log n$, every deviation by every player incurs cost at least $2c \log n$. Since every player's one-hop path has cost only $c \log n$, it follows that this outcome is indeed a PNE with respect to an arbitrary set of cost-sharing methods. The cost of this outcome is $cn \log n$.

To finish the proof, it suffices to exhibit a connected subgraph with cost $O(n)$. Consider taking all of the edges incident to v and all of the edges in G but not M. By construction, this subnetwork is connected. All of the edges in G but not M have cost 1 and there are $2n$ of them. All of the edges incident to v have cost $2c \log n$ and there are $O(n/\log n)$ of them, for a total cost of $O(n)$. ■

6.2 Applications of Theorem 6.1 to Approximation Algorithms

The construction in Theorem 6.1 is relevant to a number of well-studied NP-hard network design problems. The most interesting application is to *oblivious network design* [20, 24], where the goal is to simultaneously route one unit of flow between source-sink pairs in an undirected network at minimum cost. The key assumption is that the cost of routing a given amount of flow on an edge is governed by a concave function that is *unknown* to the algorithm. Can the flow be routed in a way that is competitive with an optimal solution that is privy to this cost function? Precisely, the input is specified by an undirected graph G with a cost c_e for each edge, and known source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$. Given this information, the design problem is to select a

multicommodity flow x that routes one unit for each source-sink pair. An adversary then chooses a concave function f, and the cost incurred is $C_f(x) = \sum_e c_e \cdot f(x_e)$, where x_e is the total flow routed across edge e. A flow x is α-*competitive* if, for *every* concave f, $C_f(x) \leq \alpha \cdot C_f(x_f^*)$, where x_f^* denotes a minimum-cost multicommodity flow for the function f.[3]

Amazingly, for every undirected network, there is a multicommodity flow with a polylogarithmic competitive ratio [24]. The construction in the proof of Theorem 6.1 implies an unconditional lower bound that rules out a constant competitive ratio.

Corollary 6.3 *For all sufficiently large k, there are k-commodity oblivious network design instances such that no routing of the demands is $o(\sqrt{\log k})$-competitive with respect to every concave cost function.*

Proof (sketch): We require only two different concave cost functions: a linear function and a constant function. (A constant competitive ratio is possible for this special case in single-sink networks [20, 39].) Begin with the network G and corresponding source-sink pairs in the proof of Theorem 6.1, and set the scaling factor c_e to $\sqrt{\log n}$ for all edges of M and to 1 for all other edges. Consider the two cost functions $f(x) = x$ and $f(x) = \sqrt{\log n}$, where x denotes the amount of flow on an edge. For each cost function, there is a routing of the traffic in G with cost $\Theta(n\sqrt{\log n})$. On the other hand, every fixed routing of the demands has cost $\Omega(n \log n)$ with respect to one of the two functions. ■

The other implications of the proof of Theorem 6.1 concern the generalized Steiner tree (GST) problem [1]: given an undirected graph with fixed edge costs and source-sink vertex pairs, the objective is to compute a minimum-cost subgraph that includes a path between each source-sink pair. This is the underlying optimization problem of network cost-sharing games in multicommodity networks. This problem and numerous variants have been extensively studied from an approximation algorithms viewpoint.

The best currently known approximation algorithms for the GST problem are fairly sophisticated primal-dual or LP-rounding algorithms that obtain an approximation ratio of 2 [1, 22, 32, 41]. Unlike the simpler Steiner tree problem, no constant-factor greedy approximation algorithm for GST is known. An obvious candidate heuristic is: iterate through the source-sink pairs in some order, always connecting the current pair via a shortest path (with already built edges viewed as zero-cost). The most natural order is to select unconnected pairs to greedily minimize the additional cost incurred at each iteration; this yields a 2-approximation algorithm when there is a common sink vertex (cf., Example 2.7). If this or any other ordering gave a constant-approximation algorithm for GST, it would also lead directly to the simplest-known constant-factor approximation algorithm for the multicommodity rent-or-buy problem (see [7, 18, 25]) and the first constant-factor approximation algorithm for the Stochastic Steiner Forest problem (see [26]).

Unfortunately, the network in the proof of Theorem 6.1 shatters all of these hopes: for *every* ordering of its source-sink pairs, the greedy heuristic above outputs a network (the matching M) that has cost $\Omega(\log k)$ times that of an optimal solution.

[3] In [24], this version of the problem is called "function-oblivious"; polylogarithmic lower bounds on the best-possible competitive ratio were already known for the "demand-oblivious" version [24, 34], in which the sources themselves are unknown.

6.3 Non-Uniform Protocols in Directed Graphs

This section studies non-uniform protocols in directed networks. Recall from Remark 4.13 that the proof of Proposition 4.12 carries over without change to non-uniform protocols, and shows that the worst-case POA and reachable POA remain linear for all such protocols.

Proposition 6.4 (POA in Directed Networks (Non-Uniform)) *For every $k \geq 1$, the worst-case POA of every admissible protocol in k-player directed networks is k. The lower bound holds even in single-sink directed networks.*

Proposition 6.4 justifies adopting the POS as an inefficiency measure. Are there non-uniform protocols more powerful than the optimal uniform one (the Shapley protocol)? To answer this question, call an outcome $(P_1, \ldots, P_k)$ of a network *enforceable* if there exists a cost-sharing method ξ_e for each edge e such that the outcome $(P_1, \ldots, P_k)$ is a PNE in the resulting network cost-sharing game. For example, a POS of 1 is achievable in a network (via a non-uniform protocol) if and only if some optimal outcome is enforceable.

Proposition 6.5 *Every single-sink directed network admits an enforceable optimal solution.*

Proof: Anshelevich et al. [3], motivated by network design games with endogenous cost shares, prove the following result: there always exists a minimum-cost solution $P_1, \ldots, P_k$ of a single-sink directed network and nonnegative payments $\{\pi_e^{(i)}\}_{e \in E, i \in \{1,\ldots,k\}}$ with the following properties:

(a) $\sum_{i=1}^{k} \pi_e^{(i)} = c_e$ for every edge e;

(b) $\pi_e^{(i)} = 0$ for every player i and edge $e \notin P_i$;

(c) for every player i and s_i-t path $\hat{P}_i$, $\sum_{e \in P_i} c_e^{(i)} \leq \sum_{e \in \hat{P}_i} c_e^{(i)}$, where $c_e^{(i)} = \pi_e^{(i)}$ for $e \in \cup_i P_i$ and $c_e^{(i)} = c_e$ otherwise.

Part (c) states that every player simultaneously chooses a minimum-cost path given the current payments π, with the understanding that the player would have to pay the full cost of an edge not currently in use.

Given this result, the proof of the proposition is easy. Consider a single-sink network, let $(P_1, \ldots, P_k)$ denote the optimal solution guaranteed above, and define $S_e = \{i : e \in P_i\}$. Let π denote payments satisfying properties (a)-(c) above. Define $\xi_e(i, S_e) = \pi_e^{(i)}$ for every edge e and player $i \in S_e$. Define $\xi_e(i, S_e \cup \{i\}) = c_e$ for every e and $i \notin S_e$; this ensures that the penalty of deviating from $(P_1, \ldots, P_k)$ is at least as severe as in property (c) above. Define other cost shares arbitrarily, subject to budget-balance. Property (c) implies that $(P_1, \ldots, P_k)$ is a PNE with respect to these cost-sharing methods, completing the proof. ■

In multicommodity networks the optimal POS can be strictly larger than 1.

Proposition 6.6 *For every $\epsilon > 0$, there is a directed network in which every enforceable outcome costs at least $(3/2 - \epsilon)$ times that of an optimal solution.*

Proof: To illustrate the main ideas, consider the network depicted in Figure 6. The unique optimal outcome has both players sharing the middle route with cost 2. For this outcome to arise as a PNE

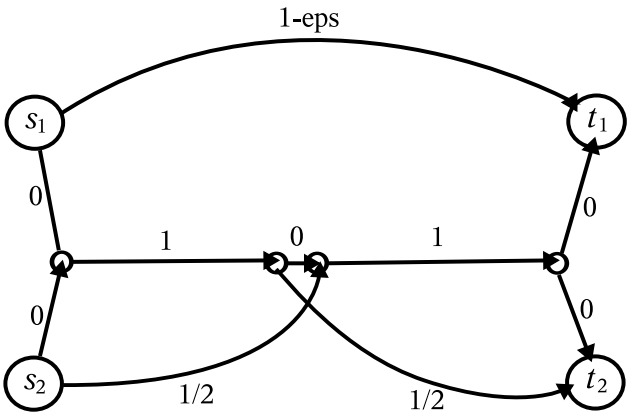

Figure 6: A network with POS strictly greater than 1 with respect to every (non-uniform) cost-sharing method.

with respect to a cost-sharing protocol, the sum of the cost shares charged to the first player must not exceed $1 - \epsilon$ (otherwise the first player would deviate to its 1-hop path). Similarly, the sum of the cost shares charged to the second player cannot exceed 1, as otherwise it would substitute one of the two edges of cost $1/2$ in its path. Budget-balance now implies that this optimal outcome is unenforceable, and the POS for this example is $5/4 - \epsilon$.

We build on this example to show for every $\epsilon > 0$, there is a network in which every enforceable outcome has cost at least $3/2 - \epsilon$ times that of the optimal solution. The base case of our iterative construction is the network of Figure 6. The second network in the sequence is shown in Figure 7, and the general process is as follows. To obtain the ith network from the $(i-1)$th network in the sequence, we first "split" the last player i of the latter network into two identical players, with respective sources and sinks s_i, t_i and s_{i+1}, t_{i+1}. Each of these two players has the same links incident to its respective source and sink as player i had in the $(i-1)$th network, except with all costs divided in half. Next, we add a direct link from s_i to t_i of cost $(i+1)/2^i - \epsilon$. We subdivide the middle route so that it has 2^i sections of cost $1/2^{i-1}$ each, with a zero-cost link separating each such segment. We then add $2^i - 1$ pairs of new edges. Each pair has one edge emanating from s_{i+1} and one entering t_{i+1}, and the other ends of these edges are connected to the tail and head (respectively) of one of the newly added zero-cost segments of the middle route. Finally, the new edges' costs are set so that every s_{i+1}-t_{i+1} path using exactly one positive-cost segment of the middle route has total cost $(i+2)/2^i$.

We now argue that if the first $i-1$ players of the $(i-1)$th network use their direct source-sink paths in every PNE (with respect to every cost-sharing protocol), then the same property holds for the first i players of the ith network. To see this, first note that in the latter network, players i and $i+1$ are collectively willing to contribute at most $1/(2^{i-2})$ toward paying for the middle route (the optimal shared path), and this is the same maximum amount that player i is willing to contribute to the middle route in the $(i-1)$th network. This implies that players $1, \ldots, i-1$ will continue to take their direct source-sink paths in the ith network in every PNE. As for player i, it uses some part of the middle route — at least two consecutive positive-cost segments — only if it pays at most $1/2^{i-1} - \epsilon$ for these segments. Player $i+1$ deviates to a path with only one positive-cost segment unless it pays strictly less than $1/2^{i-1}$ for these multiple segments. Thus, player i must

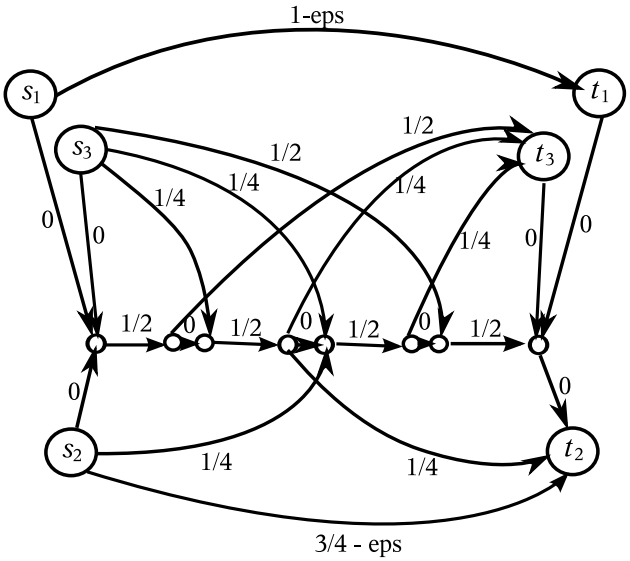

Figure 7: A network with POS $\approx 11/8$ with respect to every (non-uniform) cost-sharing method.

use its direct source-sink path in a PNE. Since the sum of the costs of the direct source-sink paths approaches 3 as $i \to \infty$, the proof is complete. ■

We leave open the challenging question as to whether or not every multicommodity network admits an enforceable near-optimal outcome.

6.4 Outside Options

We have assumed so far that each player in a network cost-sharing game is required to choose a path connecting its source and sink. This section considers briefly the obvious generalization in which each player has the "outside option" of not participating in the game, thereby suffering a constant opportunity cost. Every PNE then satisfies a "voluntary participation" constraint, in that a player agrees to select a path in the network only if its incurred cost does not exceed that of its outside option.

Formally, a *network cost-sharing game with outside options* is specified by the same data as a standard network cost-sharing game, but with an additional constant opportunity cost a_i for each player i. The strategy set of each player i is its s_i-t_i paths, and also its own outside option. (No player can select an outside option belonging to a different player.) A PNE is defined in the obvious way with respect to this enlarged strategy set. The cost of an outcome is the sum of the costs of the edges used by players that select paths, plus the sum of the opportunity costs of the players that choose their outside option.

An outside option for player i is essentially equivalent to an additional directed s_i-t_i edge that cannot be used by any other player. For this reason, the best-possible worst-case POA and POS in such games (in both undirected and directed networks) is the same as that in directed network cost-sharing games without outside options. Precisely, we have the following.

Theorem 6.7 (Optimality of the Shapley Protocol with Outside Options)

(a) For every $k \geq 1$, the worst-case POA of every cost-sharing protocol in k-player networks with outside options is k. The lower bound holds even in single-edge networks.

(b) For every $k \geq 1$, the worst-case POS of the Shapley protocol in k-player networks with outside options is the kth Harmonic number $\mathcal{H}_k$: $1 + \frac{1}{2} + \cdots + \frac{1}{k} = \ln k + O(1)$.

(c) For every $k \geq 1$, the worst-case POS of every uniform protocol in k-player networks with outside options is $\mathcal{H}_k$. The lower bound holds even in single-edge networks.

Proof (sketch): Part (a) follows from the proof of Proposition 4.12. The lower bound in that proof can be simulated with a network with one edge (s, t) with cost $1 + \epsilon$, and k players all with $a_i = 1$.

Part (b) follows from the potential function argument of Anshelevich et al. [2]. In more detail, define the function Φ as in (3), with $f(S)$ defined as $\mathcal{H}_{|S|}$ for every set S of players. This function is defined for all strategy profiles in which each player i selects either an s_i-t_i path or its own outside option; in the definition of Φ, the latter can be treated as an edge with cost a_i. The function Φ can only overestimate the cost of a strategy profile, and only by an $\mathcal{H}_k$ factor. It follows that the global minimizer of the potential function is a PNE with cost at most $\mathcal{H}_k$ times that of an optimal solution.

Finally, part (c) follows from the proof of Theorem 4.15; that example can be simulated by a single edge of cost $1 + \epsilon$ and opportunity costs $a_i = 1/i$ for all $i \in \{1, 2, \ldots, k\}$. ■

7 Conclusions and Open Questions

Our work suggests a number of promising directions for future research. One concrete question is to narrow the gap between our upper and lower bounds for non-uniform cost-sharing protocols in directed multicommodity networks. More broadly, different trade-offs between the constraints (1)–(4) discussed in the Introduction should be studied. For example, can the separability assumption be weakened in a useful way? Can the trade-off curve between efficiency and budget-balance be rigorously quantified? What about quantitative trade-offs between fairness and efficiency in undirected networks?

Finally, as the techniques for analyzing the inefficiency of equilibria in static games mature rapidly, the more algorithmic and challenging worst-case design questions should assume a central role in algorithmic game theory research. This paper presented one study of such problems; we expect similar pursuits in other settings—from routing, to facility location, to other models of network formation—to be equally fruitful.

Acknowledgements

We thank Anupam Gupta for helpful discussions during the early stages of this research, and two anonymous referees for their useful comments.

References

[1] A. Agrawal, P. Klein, and R. Ravi. When trees collide: an approximation algorithm for the generalized Steiner problem on networks. *SIAM Journal on Computing*, 24(3):440–456, 1995.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

[3] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.

[4] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized Steiner problem. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 68–74, 1996.

[5] Y. Azar, K. Jain, and V. S. Mirrokni. On Nash equilibria for a network creation game. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 323–332, 2008.

[6] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.

[7] L. Becchetti, J. Könemann, S. Leonardi, and M. Pál. Sharing the cost more efficiently: Improved approximation for multicommodity rent-or-buy. *ACM Transactions on Algorithms*, 3(2), 2007. Article 23.

[8] P. Berman and C. Coulston. On-line algorithms for Steiner tree problems. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 344–353, 1997.

[9] M. Charikar, H. Karloff, C. Mathieu, J. Naor, and M. E. Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 70–76, 2008.

[10] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.

[11] H. Chen and T. Roughgarden. Network design with weighted players. *Theory of Computing Systems*, 2009. In press.

[12] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. In *Proceedings of the 31st Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3142 of *Lecture Notes in Computer Science*, pages 345–357, 2004.

[13] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 521–530, 2003.

[14] J. Corbo and D. C. Parkes. The price of selfish behavior in bilateral network formation games. In *Proceedings of the 24th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 99–107, 2005.

[15] A. Fabrikant and C. H. Papadimitriou. The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 844–853, 2008.

[16] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *Proceedings of the 33rd Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 4051 of *Lecture Notes in Computer Science*, pages 608–618, 2006.

[17] L. K. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 277–285, 2004.

[18] L. K. Fleischer, J. Könemann, S. Leonardi, and G. Schäfer. Simple cost-sharing schemes for multicommodity rent-or-buy and stochastic steiner tree. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 663–670, 2006.

[19] L. Gao and J. Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, 2001.

[20] A. Goel and D. Estrin. Simultaneous optimization for concave costs: Single sink aggregation or single source buy-at-bulk. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 499–505, 2003.

[21] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[22] M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.

[23] T. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(1):232–243, 2002.

[24] A. Gupta, M. T. Hajiaghayi, and H. Räcke. Oblivious network design. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 970–979, 2006.

[25] A. Gupta, A. Kumar, M. Pál, and T. Roughgarden. Approximation via cost-sharing: Simpler and better approximation algorithms for network design. *Journal of the ACM*, 54(3):Article 11, 2007.

[26] A. Gupta, M. Pál, R. Ravi, and A. Sinha. Boosted sampling: Approximation algorithms for stochastic optimization. In *Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 417–426, 2004.

[27] S. Hart and A. Mas-Colell. Potential, value, and consistency. *Econometrica*, 57(3):589–614, 1989.

[28] M. Imase and B. M. Waxman. Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.

[29] N. Immorlica, L. Li, V. S. Mirrokni, and A. S. Schulz. Coordination mechanisms for selfish scheduling. *Theoretical Computer Science*, 410(17):1589–1598, 2009.

[30] M. O. Jackson. A survey of models of network formation: Stability and efficiency. In G. Demange and M. Wooders, editors, *Group Formation in Economics; Networks, Clubs, and Coalitions*, chapter 1, pages 11–57. Cambridge University Press, 2005.

[31] V. Jacobson. Congestion avoidance and control. In *Proceedings of SIGCOMM*, pages 314–329, 1988.

[32] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[33] K. Jain and V. V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–372, 2001.

[34] L. Jia, G. Lin, G. Noubir, R. Rajaraman, and R. Sundaram. Universal approximations for TSP, Steiner tree, and set cover. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 386–395, 2005.

[35] R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 2009. In press.

[36] E. Kalai and D. Samet. On weighted Shapley values. *International Journal of Game Theory*, 16(3):205–222, 1987.

[37] G. Karakostas and S. G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 268–276, 2004.

[38] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.

[39] S. Khuller, B. Raghavachari, and N. E. Young. Balancing minimum spanning and shortest path trees. *Algorithmica*, 14(4):305–322, 1995.

[40] K. Kollias. Non-preemptive coordination mechanisms for identical machine scheduling games. In *Proceedings of the 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 197–208, 2008.

[41] J. Könemann, S. Leonardi, G. Schäfer, and S. van Zwam. A group-strategyproof mechanism for the Steiner forest game. *SIAM Journal on Computing*, 37(5):1319–1341, 2008.

[42] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[43] V. S. Anil Kumar and M. V. Marathe. Improved results for Stackelberg scheduling strategies. In *Proceedings of the 29th Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2380 of *Lecture Notes in Computer Science*, pages 776–787, 2002.

[44] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2008.

[45] S. H. Low and D. E. Lapsley. Optimization flow control I: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, 1999.

[46] J. Matoušek. *Lectures on Discrete Geometry*. Springer, 2002.

[47] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[48] T. Moscibroda, S. Schmid, and R. Wattenhofer. On the topologies formed by selfish peers. In *Proceedings of the 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 133–142, 2006.

[49] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[50] J. F. Nash. Equilibrium points in N-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

[51] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, 2001.

[52] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[53] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (BGP-4). Network Working Group Request for Comments 4271, January 2006.

[54] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[55] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.

[56] T. Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33(2):332–350, 2004.

[57] T. Roughgarden and É. Tardos. Introduction to the inefficiency of equilibria. In *Algorithmic Game Theory*, chapter 17, pages 443–459. Cambridge University Press, 2007.

[58] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.

[59] L. S. Shapley. *Additive and Non-Additive Set Functions*. PhD thesis, Department of Mathematics, Princeton University, 1953.

[60] C. Swamy. The effectiveness of Stackelberg strategies and tolls for network congestion games. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1133–1142, 2007.

[61] É. Tardos and T. Wexler. Network formation games and the potential function method. In *Algorithmic Game Theory*, chapter 19, pages 487–516. Cambridge University Press, 2007.

[62] K. Varadhan, R. Govindan, and D. Estrin. Persistent route oscillations in inter-domain routing. *Computer Networks*, 32(1):1–16, 2000.

[63] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.