Designing Networks with Good Equilibria

Ho-Lin Chen* Tim Roughgarden† Gregory Valiant‡

Abstract

In a network with selfish users, designing and deploying a protocol determines the rules of the game by which end users interact with each other and with the network. We study the problem of designing a protocol to optimize the equilibrium behavior of the induced network game. 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, different classes of cost-sharing methods, and different measures of the inefficiency of equilibria. One of our main technical tools is a complete characterization of the *uniform* cost-sharing protocols—protocols that are designed without foreknowledge of or assumptions on the network in which they will be deployed. We use this characterization result to identify the optimal uniform protocol in several scenarios: for example, the Shapley protocol is optimal in directed graphs, while the optimal protocol in undirected graphs, a simple priority scheme, has exponentially smaller worst-case price of anarchy than the Shapley protocol. We also provide several matching upper and lower bounds on the best-possible performance of non-uniform cost-sharing protocols.

*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: holin@stanford.edu.

†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.

1 Introduction

Designing Networks and Protocols to Minimize Inefficiency. The computer science view of networks has changed fundamentally over the last decade. In the past, networks were largely assumed to be planned in advance by a single designer, with users who were controllable, cooperative, or at least predictable. Today, many of the networks dear to computer science—from the Internet, to the Web, to peer-to-peer and social networks—are both 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 such applications is to quantify the magnitude of suboptimality inevitably caused by selfish resource allocation. This goal is analytic, not algorithmic. For example, 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 thus adopts inefficiency measures like the POA as *objective functions* to be minimized in novel network optimization problems. The optimal objective function value then quantifies the unavoidable loss in solution quality caused by selfish behavior, given the design constraints of the problem.

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. (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 cost function is $C(\emptyset) = 0$ and $C(T) = c_e$ for all non-empty T.) 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]. Note that the design decision of how to share the network cost determines the incentive structure and hence the set of Nash equilibria in 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 links [3]. Somewhat better bounds can be obtained by restricting attention to a subset of all Nash equilibria. Considering only the Nash equilibria reachable via best-response dynamics from the empty solution, as in [6, 7], the worst-case ratio drops to polylogarithmic in single-sink undirected networks [6], but this ratio 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 number $\mathcal{H}_k = \sum_{i=1}^{k} 1/i \approx \ln k$ [2]. (The worst-case POS of Shapley cost-sharing in undirected graphs is unknown [2, 14].)

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

Desiderata. To determine whether or not there are protocols superior to Shapley cost-sharing, we must precisely define a design space. Formulating such a design space requires a number of modeling choices that are inevitably subject to debate. We propose one such definition that we feel is natural and that leads to nontrivial problems and interesting results, but freely admit that there may be alternative, equally interesting design spaces to explore.

Throughout the entire paper, we retain the following properties of the Shapley cost-sharing protocol.

(1) *Budget-balance:* The cost of each edge in the formed network is fully passed on to its users.

(2) *Separability:* The cost shares of an edge are completely determined by the set of players that use it.

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

We call cost-sharing protocols that satisfy (1)–(3) *admissible*.

A case can be made for or against each of these requirements; we next discuss the most obvious pros and cons of and alternatives to these constraints. Budget-balance (1) is, of course, the *raison d'être* of a cost-sharing protocol and is the least contentious. One could also consider some version of approximate budget-balance; the present work does not.

Separability (2) precludes any explicit coordination or communication between the cost shares of different edges of a network. (Cost-sharing decisions of different edges still implicitly affect each other by influencing the set of equilibria.) Some important practical network protocols, such as TCP/IP congestion control with various packet dropping policies (e.g. [24, 33]), can be informally regarded as separable in this sense (since each edge makes independent packet dropping decisions based only on the local state, such as the current queue length). A diametrically opposite modeling decision would be to allow arbitrary communication between different edges, subject only to some computational complexity constraint (such as polynomial time); in this case, however, our protocol design problems essentially reduce to standard problems in (non-game-theoretic) polynomial-time approximation algorithm design. In particular, short of proving $P \neq NP$, no unconditional lower bounds are possible in this alternative model. Finding a natural generalization of separability that still permits unconditional lower bounds is an interesting research challenge.

Finally, we explain why the stability constraint (3) is appropriate. One line of criticism would argue that it is too strong: by Nash's theorem [38], 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. [40, §3.2]), and is adopted primarily in games where there are no pure-strategy Nash equilibria. There may be no alternative when confronted with an arbitrary game. When *designing the game being played*, as in protocol design, there is no justification for settling for mixed-strategy equilibria. (A similar argument applies for the "sink equilibria" of [18].) For example, algorithmic mechanism design [39] is also concerned with designing 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 (3). A second parallel is provided by work in the networking community on the BGP interdomain routing protocol [41], which can be naturally viewed as a game (e.g. [13, 16, 19, 46]): while mixed-strategy Nash equilibria always exist in the induced path selection game, previous work has focused entirely on the existence of pure-strategy Nash equilibria.

One could also criticize the stability constraint (3) for being too weak: pure-strategy equilibria should not only exist, but also be *easy to find*. Arguably the most natural strengthening of (3) is to insist that best-response dynamics always converges to a pure-strategy equilibrium. (This has also been the focus of the literature on BGP discussed above, where this property is called "safety" [16, 19, 46].) While our lower bounds assume only the weaker stability requirement (3), all of our upper bounds are achieved using protocols that also satisfy this stronger convergence property. In fact, for an important subclass of protocols discussed next, a non-obvious consequence of our results is the equivalence of the two assumptions: a protocol always induces a game with pure-strategy equilibria *if and only if* it always induces a game in which best-response dynamics is guaranteed to converge.

We next explain a fourth constraint which, intuitively, requires that the cost-sharing decisions on an edge are independent of the network context. For example, in the Shapley protocol, the cost shares of an edge depend only on the edge cost and the number of users, and are independent of all other network properties (network size, location of sources and sinks, etc.).

(4) *Uniformity:* Consider two networks G_1 and G_2, each with the same player set, and two outcomes so that the users of edge $e_1 \in G_1$ and edge $e_2 \in G_2$ are the same subset S in both outcomes. If e_1 and e_2 have equal cost, then the players of S are charged the same cost shares in both outcomes.

We first emphasize that, whatever its merits, we thoroughly study the optimal protocol design problem both with and without this uniformity constraint. That said, a protocol often must be designed without foreknowledge of or assumptions about the network in which it will be deployed. Uniformity is natural 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.

More formally, a uniform cost-sharing protocol is defined as a mapping from every possible edge cost and player set to cost shares for these players, while 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 2.3 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.

Having defined a design space (either uniform or non-uniform admissible protocols) and an equilibrium concept (pure-strategy Nash equilibria), the final modeling decision required to rigorously define an "optimal protocol" is an objective function—a measure of the inefficiency of equilibria. As noted above, three such measures, ordered from most to least stringent, are the POA (determined by the worst equilibrium), the reachable POA of [6, 7] (determined by the worst equilibrium reachable via best-response dynamics from an empty state), and the POS (determined by the best equilibrium). Obviously, the strongest types of results are upper bounds on the POA and lower bounds on the POS. We seek out bounds on the POA where possible (in undirected graphs), and resort to bounds on the POS only when there are no other options (in directed graphs). While we do not emphasize the reachable POA in this extended abstract, our upper and lower bounds for the POA in undirected networks all have comparable analogues for the reachable POA, as we show in the full version.

Our Results. We systematically address the following question: *which admissible cost-sharing protocol minimizes the inefficiency of equilibria in network cost-sharing games?* We study this question for uniform

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 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].

and non-uniform cost-sharing protocols, for undirected and directed graphs, for single-sink and multicommodity networks, and for different measures of the inefficiency of equilibria. We give essentially matching upper and lower bounds in almost all cases. Our quantitative results are summarized in Table 1. (Upper bounds on the POA trivially carry over to the reachable POA. Minor modifications of our proofs extend all of our lower bounds on the POA to the reachable POA as well.)

Our main technical tool for analyzing uniform cost-sharing protocols is a *complete characterization* of such protocols. In particular, we completely characterize the uniform protocols that always induce a network game with at least one pure-strategy equilibrium. Our proof approach is to show that every such protocol is induced by a direct product of weighted potential functions. While the existence of a potential function is a standard sufficient condition for the existence of pure-strategy Nash equilibria [35, 42], it is generally far from a necessary condition. The content of our characterization result is, therefore, showing that the *only way* to obtain pure-strategy Nash equilibria via a cost-sharing protocol across all possible networks is via the potential function approach. We know of no analogous result in the economics literature. There is also a close connection between the potential functions in our characterization and the "weighted Shapley values" defined by Kalai and Samet [28].

In undirected networks, simple uniform protocols dramatically reduce the POA (and, in multicommodity networks, the reachable POA) compared to that of Shapley cost-sharing (from polynomial in k to polylogarithmic in k). We provide a complete analysis of the best-possible worst-case POA in single-sink and multicommodity undirected networks, for uniform and non-uniform schemes. For uniform protocols, we prove a (nearly tight) logarithmic lower bound on the best-possible POA, even in single-sink networks. The proof idea to leverage our characterization of uniform protocols to associate weights with the players, prove that there is either a cluster of players with sufficiently similar weights or a large group of players with sufficiently different weights, and then exhibit a family of hard instances for each of the two cases. For undirected networks and non-uniform protocols, we cannot rely on our characterization theorem and establish lower bounds via explicit constructions. For single-sink 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 achievable by non-uniform protocols via a construction based on high-girth graphs. This construction has additional implications, most notably an $\Omega(\sqrt{\log k})$ lower bound for oblivious network design [17, 20] in k-commodity networks.

For directed graphs, our characterization theorem quickly resolves the uniform protocol case: the Shapley protocol is the optimal uniform protocol in directed graphs (for the POA, reachable POA, and POS). Thus, while the Shapley protocol is typically motivated by its simplicity and fairness properties, it can be equally well justified on efficiency grounds. Lastly, for non-uniform protocols in directed graphs, we show that a POS of 1 is always achievable in single-sink networks and is not always achievable in multicommodity networks. In the latter scenario, we also give a linear programming-based characterization of the *enforceable* outcomes of a network game—the outcomes that can arise as a Nash equilibrium with respect to some non-uniform cost-sharing protocol.

Related Work. Several previous papers have studied network cost-sharing games [2, 3, 6, 7, 8, 11, 14, 36]. All of these papers studied a fixed cost-sharing method; none considered the design questions addressed here. The inefficiency of equilibria in other network design games was studied in [1, 12]. For other models of network formation and design with self-interested participants, see [5, 23, 25] 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. First, Christodoulou, Koutsoupias, and Nanavati [9] and Immorlica et al. [22] design machine scheduling policies to minimize the worst-case POA in variants of the scheduling game proposed by Koutsoupias and Papadimitriou [31]. Second, Johari and Tsitsiklis [27] 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 [30] minimizes the worst-case efficiency loss. Third, in a mechanism design context, Moulin and Shenker [37] 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 [10, 15, 29] and Stackelberg routing [32, 44, 45]. The present paper differs from these previous works in that we 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 games with good equilibria bears some resemblance to that of algorithmic mechanism design [39]. 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 Preliminaries

Network Cost-Sharing Games. In a *network cost-sharing game*, we are given a graph $G = (V, E)$, which can be directed or undirected, where each edge $e \in E$ has a nonnegative cost $c_e \geq 0$. There is a set $\{1, \dots, k\}$ of k *players*, where player i is associated with a source s_i and a sink t_i. 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 $(P_1, \dots, P_k)$ is defined to be $C(P_1, \dots, P_k) = \sum_{e \in \cup_i P_i} c_e$.

For each player i, a *cost function* $c_i : \mathcal{P}_1 \times \cdots \times \mathcal{P}_k \to \mathcal{R}_+$ describes the cost incurred by player i in each outcome of the game. Our separability constraint (2) from the Introduction dictates that each cost function c_i can be expressed as a sum over *edge cost shares*. More formally, we assume that each edge e of the network G is endowed with a *cost-sharing method* $\xi_e : 2^{\{1,\dots,k\}} \to \mathcal{R}_+^k$. A cost-sharing method ξ_e 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. Separability also implies that $\xi_e(i, S) = 0$ for all players $i \notin S$.

The budget-balance constraint (1) from the Introduction imposes the following condition on each cost-sharing method ξ_e: $\sum_{i \in S} \xi_e(i, S) = c_e$. As a consequence, the cost of every outcome $(P_1, \dots, P_k)$ is partitioned among the players: $C(P_1, \dots, P_k) = \sum_{i=1}^k c_i(P_1, \dots, P_k)$.

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. More formally, the outcome $(P_1, \dots, P_k)$ is a PNE if for every player i and every strategy $P_i' \in \mathcal{P}_i$, $c_i(P_1, \dots, P_i, \dots, P_k) \leq c_i(P_1, \dots, P_i', \dots, P_k)$. By the stability constraint (3) from the Introduction, we are only interested in network cost-sharing games that have at least one PNE. Recall that an *admissible* scheme is separable, budget-balanced, and stable.

We now formalize cost-sharing schemes.

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

For example, Shapley cost shares define a cost-sharing scheme, with the method ξ_e of an edge e given by $\xi_e(i, S) = c_e/|S|$ for every set $S \subseteq \{1, \dots, k\}$ and player $i \in S$. More generally, a cost-sharing scheme can define the method ξ_e in a way that depends on additional information, including the identities of the players in S, the topology of G, and the costs of other edges of the network.

We next formalize the uniformity constraint (4) from the Introduction.

Definition 2.2 A cost-sharing scheme is *uniform* if the cost-sharing method ξ_e assigned to an edge e is a function only of the edge cost c_e and the player set $\{1, 2, \dots, k\}$.

Ordered cost shares are a simple and important example. Such cost shares are defined with respect to an ordering of the players. The first player in the ordering pays the full cost of all edges in its path; the second player pays the full cost of all edges in its path not already paid for by the first player; and so on. Ordered cost-sharing schemes can be radically better than Shapley cost-sharing in undirected networks; the next example demonstrates this with an ordered non-uniform scheme for single-sink undirected networks.

Example 2.3 (Prim Cost-Sharing Scheme) Consider an undirected single-sink network defined by an undirected network G with edge costs c, a sink vertex t, and source vertices $s_1, \dots, s_k$. Define a non-uniform ordered cost-sharing scheme by ordering the players as follows. The first player is the one with source s_i closest to the sink t; the second player is the one with source closest to the set $\{t, s_i\}$; and so on. We call this the *Prim cost-sharing scheme*.

We claim that every PNE of the network cost-sharing game defined by the Prim cost-sharing scheme has cost at most twice that of optimal. To see this, first note that the cost incurred by the first player—the player with source closest to the sink t—depends only on its strategy and is independent of the strategies chosen by the other players. Thus in every PNE, this player will choose a shortest path P_i between its source s_i and the sink t. By the same reasoning, in every PNE, the second player will choose a shortest path between its source and the path P_i, and will then follow P_i to the sink t. The cost incurred by the second player is thus at most the length of a shortest path between its source and the set $\{s_i, t\}$.

More generally, in every PNE, each player selects a shortest path from its source to the union of the paths chosen by earlier players in the ordering. The resulting outcome thus corresponds to 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. [47]), which is a minimum-cost outcome in the network game. Thus the POA in every network cost-sharing game defined by the Prim cost-sharing scheme is at most 2. Recall that the POA in Shapley network design games can be as large as the number k of players, even in undirected networks of parallel links [3].

Remark 2.4 Standard examples (e.g. [47, Example 3.4]) give a matching lower bound on the worst-case POA of every non-uniform admissible cost-sharing scheme in single-sink undirected networks.

3 A Characterization of Linear, Uniform, Admissible Cost-Sharing Schemes

This section provides a complete characterization of the linear and uniform cost-sharing schemes that are admissible. The stability constraint (3) from the Introduction—a complex "global" constraint on all network games that might be induced by a scheme—makes this result highly non-trivial. This characterization has several consequences, including lower bounds in the next two sections on the worst-case POA and POS achievable by (not necessarily linear) uniform admissible schemes in undirected and directed networks, respectively.

Formally, a uniform cost-sharing scheme for a player set $\{1, 2, \ldots, k\}$ is *linear* if, for all $c_e \geq 0$, the cost-sharing method ξ 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 sometimes abuse notation and refer to a linear and uniform scheme (for a fixed player set) by the cost-sharing method it assigns to a unit-cost edge. Also, a cost-sharing method is *positive* if it always assigns strictly positive cost shares to all of the players using an edge of non-zero cost.

To map the terrain of linear, uniform, admissible schemes, we begin with the Shapley scheme. The stability of the Shapley scheme follows a "potential function argument" [2, 35, 43]: one exhibits a potential function for each network game induced by the scheme such that local minima of the potential function are in bijective correspondence with the PNE of the game. Do any other schemes admit potential functions? This question motivates the following definition.

Definition 3.1 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 ($f(S) < f(T)$ whenever $S \subset T$) and if

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

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

It is straightforward to show that every edge potential induces a positive, linear, uniform, and admissible cost-sharing scheme.

Proposition 3.2 *Let f be an edge potential for the player set $\{1, 2, \ldots, k\}$. Define a positive, linear, and uniform cost-sharing scheme by assigning a unit-cost edge the cost-sharing method ξ, where $\xi(i, S) = (f(S) - f(S \setminus \{i\}))/f(\{i\})$ for every $S \subseteq \{1, \ldots, k\}$ and $i \in S$. Then, ξ is admissible.*

We call a cost-sharing scheme *potential-based* if it is induced by an edge potential as in Proposition 3.2.

The Shapley cost-sharing scheme corresponds to the edge potential f with $f(S) = \mathcal{H}_{|S|}$ for every subset S of players. Because of the budget-balance constraint in Definition 3.1, it is not immediately obvious that further edge potentials exist. But as the proof of our characterization theorem implicitly shows, there are a plethora of others, with a bijective correspondence between them and the open unit cube $(0, 1)^{k-1}$ in $(k-1)$ dimensions. In fact, each such edge potential can be interpreted as a sum of weighted Shapley cost shares in the sense of Kalai and Samet [28]. (A different notion of weighted Shapley cost shares has been explored recently in the computer science literature [2, 8].)

Not all linear, uniform, and admissible cost-sharing schemes are positive, as uniform variants of the Prim cost-sharing scheme (Example 2.3) show. (See also Proposition 4.1 below.) This motivates the following operation for combining two cost-sharing schemes into a single (non-positive) one.

Definition 3.3 Let ξ_1 and ξ_2 be linear, uniform cost-sharing schemes for disjoint player sets A_1 and A_2, respectively. The *concatenation of* ξ_1 *and* ξ_2 is the cost-sharing scheme $\xi_1 \oplus \xi_2$ 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 are using an edge, then they share its cost according to ξ_2. The concatenation operation preserves linearity and uniformity by definition; it also clearly preserves separability and budget-balance. Arguing as in Example 2.3, it also preserves stability.

Proposition 3.4 *Let ξ_1 and ξ_2 be linear, uniform, separable, and budget-balanced cost-sharing schemes. The concatenation $\xi_1 \oplus \xi_2$ is stable if and only if both ξ_1 and ξ_2 are stable.*

We can now state our characterization result: *every* linear, uniform, admissible cost-sharing scheme arises as the concatenation of potential-based schemes.

Theorem 3.5 *Let ξ be a linear, uniform, separable, and budget-balanced cost-sharing scheme. Then ξ is stable if and only if it is the concatenation of potential-based schemes.*

The proof of Theorem 3.5 is highly involved; for details, see the full version of this paper. Its four major steps are as follows. First, we show that every linear, uniform, and admissible scheme ξ must be *monotone* in the sense that $\xi(i, S) \geq \xi(i, T)$ whenever $i \in S \subseteq T \subseteq \{1, 2, \ldots, k\}$. Second, we prove that for every ξ as above, 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.) Third, for every scheme ξ as above that is also positive, all of its cost shares are uniquely determined by the $k-1$ pairwise cost shares $\xi(1, \{1, 2\})$, $\xi(1, \{1, 3\})$, ..., $\xi(1, \{1, k\})$. Finally, we prove that for every set of $k - 1$ pairwise cost shares as above, there exists a potential-based cost-sharing scheme with the prescribed cost shares.

4 Undirected Networks: Minimizing the POA

This section identifies the best-possible worst-case POA achievable in undirected graphs. All of our upper bounds follow from simple ordered cost-sharing methods. Our lower bound for uniform schemes builds on our characterization of uniform schemes (Theorem 3.5), while our lower bound for non-uniform schemes is an explicit construction derived from a family of high-girth graphs.

Uniform Cost-Sharing Schemes. We begin with some easy upper bounds using uniform ordered cost-sharing schemes, where the players are ordered in a fixed and arbitrary way (lexicographically, say). Reasoning as in Example 2.3, the worst-case POA of such a scheme is precisely the worst-case competitive ratio of the natural online greedy algorithms for online Steiner tree [21] (in the single-sink case) and generalized Steiner tree [4] (in the multicommodity case). We thus have the following positive consequences of the upper bounds in [4, 21] (where as usual, k denotes the number of players).

Proposition 4.1 *There is a uniform admissible cost-sharing scheme for single-sink undirected networks with worst-case POA $O(\log k)$, and for multicommodity undirected networks with worst-case POA $O(\log^2 k)$.*

We next leverage our characterization of uniform cost-sharing schemes (Theorem 3.5) to prove that the guarantees of Proposition 4.1 cannot be significantly improved by any uniform cost-sharing scheme, ordered or otherwise.

Theorem 4.2 *Every uniform admissible cost-sharing scheme has worst-case POA $\Omega(\log k)$, even in single-sink undirected networks.*

Given a uniform admissible cost-sharing scheme, our proof of Theorem 4.2 first separates the players into different classes and, for each player class, orders the players according to the value $f(\{i\})$ of the associated edge potential f (recall Theorem 3.5). Intuitively, players that have similar f-values should also have similar cost shares. We prove a dichotomy lemma that shows that there are either $\log k$ players that are sufficiently similar or $k/\text{polylog}(k)$ players that almost form an ordered cost-sharing scheme. In both of these two cases, we can construct an example with $\Omega(\log k)$ POA. The details are not trivial and are given in the full version.

Non-Uniform Cost-Sharing Schemes. Non-uniform cost-sharing schemes overcome the lower bound of Theorem 4.2 in single-sink networks. (Recall Example 2.3.) Is this also true in multicommodity networks?

For example, there is a natural analogue of the Prim cost-sharing scheme 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.

Since our characterization (Theorem 3.5) is only for uniform cost-sharing schemes, it offers no assistance for this question. Instead, we devise a different lower bound tailored to multicommodity networks that is robust to non-uniform cost-sharing schemes. (See below for additional applications of this result.)

Theorem 4.3 *For all sufficiently large k, the worst-case POA with respect to every (non-uniform) admissible cost-sharing scheme for k-player undirected multicommodity networks is $\Omega(\log k)$.*

The proof of Theorem 4.3 is based on the following combinatorial lemma, which we prove using a high-girth graph construction of Erdös and Sachs (see [34, Exercise 15.3.1]) and Hall's Marriage Theorem.

Lemma 4.4 *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 of Theorem 4.3: Let $G = (V, E)$ be the graph described in Lemma 4.4. 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 $O(n)$. ■

Applications of Theorem 4.3. In *oblivious network design* [17, 20], the goal is to simultaneously route one unit of flow between source-sink pairs in an undirected network at minimum cost. However, 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? (In [20], this version of the problem is called "function-oblivious"; lower bounds were already known for the "demand-oblivious" version [20, 26], in which the sources themselves are unknown.) Amazingly, a polylogarithmic competitive ratio is possible [20]. Our next result is an unconditional lower bound that rules out constant competitive ratios for this problem.

Corollary 4.5 *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.*

The idea is to set the lengths of the one-hop paths in the network in the proof of Theorem 4.3 to $\sqrt{\log n}$, and to consider the two cost functions $c(x) = x$ and $c(x) = \sqrt{\log n}$, where x denotes the amount of flow on an edge. For each function, there is a routing of the traffic 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.

In the full version, we discuss implications of the construction in Theorem 4.3 for several NP-hard network design problems.

5 Directed Networks: Minimizing the POS

In directed networks, a simple example (see the full version) proves that every admissible (non-uniform) cost-sharing scheme has worst-case POA equal to the number k of players. We therefore study only the POS in directed networks.

Uniform Cost-Sharing Schemes. Our main result in this section is that the Shapley cost-sharing scheme is the optimal uniform scheme. This result justifies Shapley cost-sharing in a novel way: in addition to being simple and fair, it minimizes worst-case efficiency loss over all uniform schemes in directed networks. Given our complete characterization of uniform schemes (Theorem 3.5), the proof is not difficult (see the full version).

Theorem 5.1 *For all $k \geq 1$, among all uniform cost-sharing schemes, the Shapley cost-sharing scheme minimizes the worst-case POS in k-player directed networks.*

Non-Uniform Cost-Sharing Schemes. Are there non-uniform schemes more powerful than the optimal uniform one (i.e., Shapley)? We answer this question in the affirmative for single-sink networks, and give partial results for multicommodity networks.

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 (with a non-uniform scheme) if and only if some optimal outcome is enforceable.

Proposition 5.2 *In single-sink directed networks, optimal solutions are enforceable.*

We prove Proposition 5.2 via a reduction to an analogous result in [3] about network cost-sharing games with endogenous cost shares. In multicommodity networks the optimal POS can be strictly larger than 1 (see the full version). An obvious question is whether or not every multicommodity network admits an enforceable near-optimal outcome. We leave this as a challenging open problem. We contribute in the full version a linear-programming-based characterization of the enforceable outcomes of a network game that we expect will prove useful in resolving this question.

References

[1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 89–98, 2006.

[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. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

[3] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 511–520, 2003.

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

[6] M. Charikar, H. Karloff, C. Mathieu, J. Naor, and M. Saks. Best response dynamics in multicast cost sharing. Manuscript, 2007.

[7] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative multicast and facility location games. In *7th Annual ACM Conference on Electronic Commerce (EC)*, pages 72–81, 2006.

[8] H. Chen and T. Roughgarden. Network design with weighted players. In *Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architextures (SPAA)*, pages 28–37, 2006.

[9] 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.

[10] 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.

[11] 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.

[12] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. J. Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[13] 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)*, 2008. To appear.

[14] 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.

[15] 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 Com-

puter Science (FOCS), pages 277–285, 2004.

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

[17] 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.

[18] 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.

[19] 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.

[20] 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.

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

[22] N. Immorlica, L. Li, V. S. Mirrokni, and A. S. Schulz. Coordination mechanisms for selfish scheduling. In *Proceedings of the First Annual International Workshop on Internet and Network Economics (WINE)*, volume 3828 of *Lecture Notes in Computer Science*, pages 55–69, 2005.

[23] 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. Cambridge University Press, 2005.

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

[25] 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.

[26] 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.

[27] R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. Submitted, 2007.

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

[29] 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.

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

[31] 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.

[32] 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.

[33] 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.

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

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

[36] 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.

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

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

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

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

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

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

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

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

[45] 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.

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

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