

Capacity Allocation Games for Network-Coded Multicast Streaming*

Elliot Anshelevich¹, Bugra Caskurlu², Koushik Kar³, and Hang Zhang³

¹ Department of Computer Science,
Rensselaer Polytechnic Institute, Troy NY 12180, USA
eanshel@cs.rpi.edu

² Lane Department of Computer Science & Electrical Engineering,
West Virginia University, Morgantown WV 26506, USA
caskurlu@gmail.com

³ Department of Electrical, Computer & Systems Engineering,
Rensselaer Polytechnic Institute, Troy NY 12180, USA
{kark,zhangh10}@rpi.edu

Abstract. In this paper we formulate and study a capacity allocation game between a set of receivers (players) that are interested in receiving multicast data (video/multimedia) being streamed from a server through a multihop network. We consider *fractional* multicast streaming, where the multicast stream from the source (origin-server) to any particular receiver (end-user) can be split over multiple paths. The receivers are selfish and non-cooperative, but must collaboratively purchase capacities of links in the network, as necessary for delivery of the multicast stream from the source to the individual receivers, assuming that the multicast stream is *network coded*. For this multicast capacity allocation (network formation) game, we show that the Nash equilibrium is guaranteed to exist in general. For a 2-tier network model where the receivers must obtain the multicast data from the source through a set of relay nodes, we show that the price-of-stability is at most 2, and provide a polynomial-time algorithm that computes a Nash equilibrium whose social cost is within a factor of 2 of the socially optimum solution. For more general network models, we give a polynomial time algorithm that computes a 2-approximate Nash equilibrium whose cost is at most 2 times the social optimum. Simulation studies show that our algorithms generate efficient Nash equilibrium allocation solutions for a vast majority of randomly generated network topologies.

Keywords: capacity allocation games, the price of stability, Nash equilibrium, multicast networks.

1 Introduction

The last decade has witnessed an explosive growth in the number of *streaming* video (multimedia) applications. Some of these involve live video streaming,

* This work is supported in part by NSF grants CCF-0914782, CNS-1017932, and CNS-1018398.

while others stream video that is already available in stored format but too large for download-and-play. Stored video streaming applications like YouTube [4] are already contributing to a large fraction of the Internet traffic today¹, and IPTV and similar other efforts are likely to boost live video streaming through the Internet in the coming years [3]. Streaming stored video may be unicast or multicast, depending on the application: while IPTV [2] video streaming may mostly be multicast (broadcast), receiver-driven video streaming (like video streaming from YouTube [4]) will typically be unicast. Streaming live video will typically be multicast to possibly many receivers.

For multicast data delivery, use of *network coding* allows individual receivers to simultaneously attain data rates that equal their maxflow capacities [6], which in general may not be achievable through a routing-only approach. Naturally, this makes network coding ideally suited for multicast data delivery over a multi-hop network.

In this paper, we consider a capacity allocation game that end-users will play in buying resources for multicast streaming data delivery. More specifically, receivers (users) buy capacities on the links of the distribution network at fixed (possibly different for different links) per-unit cost, so as to ensure delivery of a multicast stream (with a given source rate) from its source to the receiver. The receivers are selfish and non-cooperative and are only interested in minimizing their individual costs, but must collaboratively pay for capacities bought on network links, as necessary for network coded multicast data delivery from the data source to the individual receivers. We consider *fractional* multicast streaming, i.e., the multicast data between the source and any particular receiver can be split across multiple paths that exist between the source-receiver pair. The problem we consider is a *network formation* game where the amount of capacity collaboratively bought on the different links in the network must be such that the maxflow from the source to each receiver is no less than the desired multicast data rate. Using network coding [6], this ensures that all receivers are able to obtain their full data rate from the source. This paper focuses on the questions of the existence, efficiency, and computation of the equilibria of this game. Initially, we focus on the 2-tier network model where the receivers must obtain the multicast data from the source through a set of relay nodes, and derive some strong results by exploiting structural properties of such topologies. Later we consider arbitrary topology networks for multicast data distribution, and study the existence and efficiency of approximate equilibria for that case.

To measure efficiency of equilibrium, we use the common measures of the *price of anarchy* and the *price of stability* [22] — the supremum of the ratios between the costs of the worst and best pure Nash equilibrium, respectively, and that of the globally optimal solution over all instances of the game.

Our Contributions. The specific technical contributions of this paper are as follows. For our fractional multicast network formation game, we show that pure

¹ It was estimated that in 2007 YouTube consumed as much bandwidth as the entire Internet in 2000 [1].

Nash equilibrium is guaranteed to exist in general. For the 2-tier model, defined in Section 4, we show a tight bound of 2 on the price-of-stability, and provide a polynomial-time algorithm that returns a Nash equilibrium whose social cost is within a factor of 2 of the socially optimum solution. The 2-tier model is essentially equivalent to the case where all nodes in the network are receivers, which is itself an important special case (e.g. [17]). For more general network topologies, we give a polynomial time algorithm that computes a 2-approximate Nash equilibrium whose cost is at most 2 times the social optimum. Simulation studies show that our algorithms generate efficient Nash equilibrium allocation solutions for a vast majority of randomly generated networks. Due to lack of space, the proofs of Theorem 1, and Theorem 5 appear in the full version of this paper. (see <http://www.ecse.rpi.edu/~koushik/>)

While network formation games have been studied in other contexts (see Section 2), the questions we consider are new for the context of network-coded fractional multicast streaming. Unlike integral multicast, data distribution networks for socially optimal or Nash equilibrium solutions for network-coded multicast need not be trees (Figure 1), and techniques for integral multicast do not extend to this context (see Section 2). Interestingly, however, we show that there exist solutions based on tree topologies that are at Nash equilibrium (exactly or approximately) and attain a near-optimal social cost. For the 2-tier network model, the solution is based on the minimum spanning tree; for more general network models, it is based on the minimum Steiner tree or polynomial-time approximations of it. Despite the complexity of the problem, our results show that there exist easily-computable exact or approximate distribution networks where receivers have no motivation to deviate from it unilaterally, and yet results in the set of receivers paying near-minimal cost as a group for multicast data delivery.

The paper is structured as follows. In Section 2, we outline related work on this topic. In Section 3 we describe the model and problem formulation. In Section 4 we state and prove our main results on the existence, efficiency, and properties of Nash equilibrium distribution topology solutions for the 2-tier network model. We extend these results to arbitrary topology networks in Section 5. In Section 6 we describe the results of experiments conducted on randomly generated network topologies.




Fig. 1. Example: Fractional multicast streaming is better than integral multicast streaming. Source S sends multicast data at rate 1 to receivers A and B. All link capacities are 1 unit in each direction; numbers across links are per-unit capacity (bidirectional) purchase costs. Optimal fractional solution involves purchase of 1 unit capacities on links CA and DB, and 0.5 unit capacities on links SC, CD, SD; total cost = 25. Optimal integral solution involves purchase of 1 unit capacities along SC, SD, CA, DB; total cost = 30.

2 Related Work

In contrast to the work presented in this paper, the models in most network formation game literature do not allow players to reserve an arbitrary amount of bandwidth on the links; rather, a link can either be constructed and be utilized to full extent or won't be constructed at all. Our game represents a more realistic scenario by allowing players to buy certain amount of capacity (bandwidth) on that link. This becomes particularly relevant for network coded multicast streaming [6,18,21,23] that we consider in this paper, where the capacity that needs to be bought on a link for successful multicast streaming is often less than the link capacity, as well as the source data rate.

There have been several variants of “integral” network formation games where instead of allocating capacity, an edge can be either fully present or non-existent. One of the most important decisions when modeling network design involving strategic agents is to determine how the total cost of the solution is going to be split among the players. Among various alternatives [13], the most popular one in the literature is the “fair sharing” mechanism [8,11,12,17]. In this cost sharing mechanism, the cost of each edge of the constructed network is shared equally by the players using that edge. Since in our model, each player is allowed to purchase any amount of capacity on an edge, the “arbitrary sharing” model of network formation [9,7,15,19,20] is closer to being an “integral” version of our game. In this model, players contribute to the cost of an edge, and an edge is present in the network if the player contributions are larger than its cost. This model has many differences from the “fair sharing” model: e.g., the game is not a congestion game, but the price of stability is much better than with fair sharing, etc.

While many interesting results have been proven for network formation games with arbitrary sharing, most do not extend to our “fractional” context where players (multicast data users/receivers in our case) are allowed to reserve an arbitrary amount of bandwidth on the links. For example, [9] proved that the minimum-cost Steiner tree is always a Nash equilibrium in this integral network formation game with arbitrary sharing, but the same does not hold for the fractional version. Consider, for example, a graph with n receiver nodes, node s , and an extra node v . All receivers wish to receive a rate of 1 from node s . For any receiver node u , the cost of allocating x capacity on an edge (u, v) is x , and the same is true for edge (v, s) . The cost of allocating x capacity on an edge (u, s) is $x(2 - \epsilon)$. No other edges exist in the graph. Then, the minimum-cost Steiner tree has cost $n + 1$ (for this example, it is also the optimum fractional solution, although Figure 1 shows that this is not always the case). However, the min-cost Steiner tree in this example is not a Nash equilibrium: there must be some receiver u who is paying for 1 capacity of edge (u, v) , and at least for $1/n$ capacity of edge (v, s) . This player could reduce its payments to both edges by $1/n$, and instead pay for $1/n$ capacity on edge (u, s) , costing it $(2 - \epsilon)/n$, and thus strictly decreasing its cost.

Our game assumes that there is no central authority that can dictate the network cost-sharing mechanism, and thus players simply purchase edge capacities

directly. Other directions in multicast games include cooperative games and mechanism design (see, e.g. [16] and [22, Chapters 14-15] and references therein), where the goal is to come up with a cost-sharing mechanism with good properties that could be implemented by a central authority.

3 System Model and Formulation

We consider a network modeled by an undirected graph $G = (V, E, b)$, where vertices V denote the set of nodes, and (undirected) edges E denote the set of (bidirectional) links in the network. For each edge $e \in E$, $b(e) \geq 1$ denotes the capacity (bandwidth)² of the corresponding link in each direction. In other words, at full capacity $b(e)$ units of traffic can travel on link e in one direction, and simultaneously $b(e)$ units of traffic can travel on e in the opposite direction. This network is to be used for delivery of a given traffic stream of rate 1. One of the nodes, $s \in V$ is designated as the source of the traffic stream, and a subset of the nodes, $R \subseteq V$ are receivers (users) of the traffic stream. In our model, all receivers must collaboratively pay for capacities of the links that are used for carrying the traffic stream, for shared use by all receivers in R . We assume that an edge e is associated with a cost of $c(e)$ that buys the corresponding link, i.e., $c(e)$ is the cost of buying capacity of 1 of the link in each direction. If each receiver $i \in R$ pays $p_i(e)$ for edge e , then the total *purchased capacity* on the corresponding link is $\theta(e) = (\sum_{i \in R} p_i(e))/c(e)$. In other words, if the players R pay a fraction of edge e 's cost in total, then the same fraction of e 's desired capacity becomes available in both directions.

To formally define the game, we note that R is the set of players, and $\mathbf{p} = (p_i(e), i \in R, e \in E)$, or the prices paid by the receivers for the links, constitute the player strategies. (Formally, a strategy of player $i \in R$ is a function $p_i : E \rightarrow R_{\geq 0}$ that determines how much i is offering to pay for each edge.) Once the strategies are given, the network used for multicast traffic streaming, or the *distribution network*, is (V, E, θ) , where the purchased capacities $\theta(\cdot)$ depend on the strategy vector \mathbf{p} . We assume that each node in the network is capable of network-coding. A necessary and sufficient condition for delivery of the given traffic stream to all receivers is that the maxflow from the source s to each receiver $i \in R$ in distribution network (V, E, θ) is no less than the stream rate of 1 [6]. In other words, a feasible strategy vector \mathbf{p} is one that ensures that

$$f_i(\mathbf{p}) \geq 1, \tag{1}$$

holds for each receiver $i \in R$, where $f_i(\mathbf{p})$ denotes the maxflow from source node s to receiver node i in distribution network (V, E, θ) . The goal of each receiver i is to receive the full rate of 1, but pay as little as possible. Formally, the cost of player i with strategy vector \mathbf{p} is $\sum_e p_i(e)$ if $f_i(\mathbf{p}) \geq 1$, and is very large otherwise.

² We assume that the links are loss-free, and the terms ‘capacity’ and ‘bandwidth’ are used synonymously.

Among all strategies that form feasible networks, we are specifically interested in those that are Nash equilibria. A solution is a Nash equilibrium if each receiver does not have an incentive to unilaterally deviate from it. To state this formally, consider a strategy vector \mathbf{p}^* that is feasible for all players, and let $\mathbf{p}^*_{-i} = (p_j^*(e), j \in R \setminus i, e \in E)$, denote the strategies of (purchase prices paid by) all receivers other than i . Then the strategy vector \mathbf{p}^* is said to be a Nash equilibrium if for any $\mathbf{p}_i = (p_i(e), e \in E)$ such that $(\mathbf{p}_i, \mathbf{p}^*_{-i})$ satisfies the feasibility condition (1) for receiver i , we have that $\sum_{e \in E} p_i^*(e) \leq \sum_{e \in E} p_i(e)$. In other words, given the strategies of (prices paid by) the other receivers, and subject to maintaining the feasibility condition $f_i(\mathbf{p}) \geq 1$ that is necessary for receiver i to receive the streamed data at the full rate, the total price paid by receiver i is minimized at Nash equilibrium. A solution that is not feasible for all players will never be a Nash equilibrium, since the infeasible players will have very large cost, and will have incentive to purchase more capacity in the network in order to become feasible.

To study the efficiency of the Nash equilibrium, we next define the *social optimum* against which the Nash equilibrium solution will be compared in terms of the total capacity purchase cost. A price vector \mathbf{p} is said to be a social optimum if it minimizes $\sum_{i \in R, e \in E} p_i(e) = \sum_e \theta(e)c(e)$, subject to satisfying the feasibility constraints (1) for all receivers $i \in R$. We will denote this solution as *OPT*. The supremum of the ratio of the overall cost of the *worst* Nash equilibrium to that of the social optimum over all instances of a game is defined as the *price of anarchy* of the game. Similarly, the supremum of the ratio of the overall cost of the *best* Nash equilibrium to that of the social optimum over all instances of a game is defined as the *price of stability* of the game. In this paper, we use these two notions to characterize the efficiency of the worst and best Nash equilibria for the game.

4 Capacity Allocation Games on the 2-Tier Network Model

It has been envisioned that in the near future, streaming of video over the Internet will be done through the use of multiple, dedicated servers that would relay video from the source server to the end-users [5]. This network of video relay servers is expected to play a role similar to that of content distribution networks (CDN) for delivery of various kinds of non-streaming and non-real time data (content). This relay network will deliver video from the source to the receiver through multiple relay-hops, possibly performing network coding at the intermediate (relay) nodes/servers. This is illustrated in Figures 2-3.

In this section, we restrict our attention to the 2-tier network model. In this model, $V = \{s\} \cup L \cup R$, i.e., the vertex set of the graph is composed of the source vertex s , the set of receivers $R = \{r_1, r_2, \dots, r_n\}$, and the set of relay nodes $L = \{l_1, l_2, \dots, l_k\}$. Each receiver node r_i of G has *exactly* one incident edge and is adjacent to a relay node of G , i.e., for each $r_i \in R$ there exists $l_j \in L$ such that $(r_i, l_j) \in E$, and r_i has no other incident edges in G . Each relay node




Fig. 2. Streaming video delivery in the 2-tier network model




Fig. 3. Graph representation of the 2-tier network model in Figure 2

is adjacent to one or more receiver nodes, i.e., for each $l_j \in L$ there exists $r_i \in R$ such that $(r_i, l_j) \in E$. These assumptions imply that each receiver is directly connected to exactly one relay node, from which it must receive the data being streamed (this relay node can obtain the data from the source through other relay nodes, however). In addition, each relay node serves at least one receiver. This assumption can be interpreted in the following way: if a relay node has no receiver to serve, then it does not participate in the multicast data distribution.

Note that the 2-Tier model is equivalent to the model where all nodes are receivers, in the following sense. If G is an instance of our game in the 2-Tier model, let G' be a game obtained by contracting all the edges incident to a receiver node in G . Since every relay node in G must be adjacent to a receiver, then all nodes in G' except s are receivers, although there may be several receiver players located at the same node. Since every Nash equilibrium in G must have each receiver r_i purchasing 1 capacity on the edge (r_i, l_j) , then it is easy to see that there is a one-to-one correspondence between the Nash equilibria in G and in G' . Below we argue about the price of stability for 2-Tier networks, but all the arguments can easily be extended to the model where the network has arbitrary topology, but every node except s is a receiver/player.

We prove below that if G has a 2-tier topology as described above, then there exists a Nash equilibrium solution that does not cost much more than the cost of the socially optimal solution. We first show, however, that there are examples where all Nash equilibria cost a factor of 2 more than the social optimum.

Theorem 1. *The price of stability of the capacity allocation game with the 2-Tier topology is at least 2.*

In the proof of Theorem 1, we show that for any $\epsilon > 0$ we can construct an instance of the capacity allocation game with the 2-Tier topology such that the social cost of all Nash equilibria are at least $(2 - \epsilon)$ times the social cost of OPT .

Consider the cheapest solution satisfying Condition (1) for our game, which we denoted by OPT . Notice that OPT is exactly the optimum fractional solution to the LP-relaxation of the classic Steiner tree problem (see, e.g. [24]), with the terminal nodes being $R \cup \{s\}$ and edge costs being $c(e)$. Since the integrality gap

of this LP is at most 2 [24], this implies that a minimum-cost Steiner tree has cost at most twice the cost of OPT . More precisely, if we take a Steiner tree T of graph G with terminals $R \cup \{s\}$ and edge costs $c(e)$, and set the capacity of each edge in T to be 1, and each edge not in T to be 0, then the cost of this tree is at most twice the cost of OPT . We show in Section 5 that we can always form a 2-approximate Nash equilibrium on a Steiner tree, and there are simple examples where there is no Nash equilibrium that buys the minimum-cost Steiner tree. For the 2-Tier network topology, however, notice that the minimum-cost Steiner tree is simply the minimum spanning tree (MST) of G . This is because every node r_i of R has *exactly* one incident edge and is adjacent to a relay node of G , and each relay node is adjacent to one or more receiver nodes. We next prove that there always exists a Nash equilibrium solution that buys the minimum spanning tree, which gives the result stated by Theorem 2, since the cost of the minimum spanning tree is at most twice the cost of the socially optimal solution as argued above.

The proof of Theorem 2 gives a polynomial-time algorithm that returns a strategy profile $\mathbf{p} = (p_i(e), i \in R, e \in E)$ that is a Nash equilibrium, and reserves 1 unit of bandwidth on the edges of the MST. Since the cost of MST is at most 2 times the cost of OPT , the price of stability is at most 2. Theorem 1 and Theorem 2 together imply that the price of stability is 2.

Theorem 2. *There is a polynomial-time algorithm that returns a Nash equilibrium of the capacity allocation game for the 2-tier topology whose social cost is within a factor of 2 of the cost of OPT , and thus the price of stability is at most 2.*

Proof. We prove the result by showing that there always exists a Nash equilibrium solution that reserves 1 unit of bandwidth on the edges of the minimum spanning tree and 0 units of bandwidth on the remaining edges. Our proof is constructive, i.e., we explicitly form payments that purchase the minimum spanning tree. In our payment scheme, for each edge e of minimum-cost spanning tree T , there is a corresponding receiver r_i that reserves 1 unit of bandwidth on it. Notice that even though the 'arbitrary-sharing' cost-sharing scheme allows receivers to share the cost of reserving 1 unit of bandwidth on the edges of T , our payment scheme does not use this property. Therefore, in order to fully specify the payment scheme all we need to do is to assign one receiver for each edge e of T .

Without loss of generality, for an edge $e = (i, j)$, we will assume that j is the node that is closer to s on T than i . If $e = (i, j)$ is an edge incident to a receiver, i.e., $i \in R$, then the receiver i makes a payment on e that is sufficient to reserve 1 units of bandwidth on it. Let $e = (i, j)$ be an edge between two relay nodes, or a relay node and the source, i.e., $i \in L$. Let r_k be an arbitrary receiver that has a direct link to the relay node i , i.e., (r_k, i) is an edge of T . Then, r_k makes a payment of $c(e)$ on e .

Since we have fully specified the payment scheme on all the edges of T , let us now prove that this payment scheme is indeed a Nash equilibrium. Notice that

in our payment scheme each receiver is reserving 1 units of bandwidth on either one or two edges of T . More precisely, each receiver r_i paying for the cost of the edge (r_i, l_j) incident to her node, and possibly she is also paying for the cost of the other incident edge of l_j .

Recall that any solution that is feasible for r_i must satisfy Inequality (1) for r_i . If a receiver r_i is only paying for the cost of her incident edge e in T , she trivially does not have an incentive of unilateral deviation since r_i does not have any other incident edges in G , and therefore any solution where r_i satisfies Equation 1 reserves 1 unit of bandwidth on e . Suppose r_i is paying for the cost of both $e = (r_i, l_j)$ incident to her node, and the cost of another incident edge f of l_j . Since the payments of the receivers $R - \{r_i\}$ reserve 1 unit of bandwidth on $T - \{e, f\}$, the best response of r_i must include enough capacity for a flow of size 1 between r_i and $T - \{e, f\}$. The cheapest way to obtain this capacity is to reserve 1 unit of bandwidth along a single path between r_i and $T - \{e, f\}$, and the cost of this path cannot be more than the total cost of e and f , since otherwise T would not be a minimum spanning tree. Thus, any deviation of r_i that results in a feasible solution for r_i is at least as expensive as $c(e) + c(f)$. Since no receiver has an incentive for unilateral deviation, the resulting payment scheme is a Nash equilibrium. ■

5 Generalizations for Arbitrary Network Models

In this section, we consider our general game, with the graph G having arbitrary topology, and an arbitrary subset of receiver nodes. The capacity allocation game is guaranteed to have a Nash equilibrium by Theorem 3, however, the cost of some Nash equilibria can be prohibitive by Theorem 5.

Theorem 3. *Nash equilibrium in pure strategies is guaranteed to exist in the capacity allocation game.*

Proof. For each player $i \in R$, let S_i denote the strategy space of player i . Player i selects a strategy $s_i \in S_i$ when she plays the game. Let $S = \prod_{i \in R} S_i$ denote the strategy space of the game. Notice that S is the product space of the strategy spaces of the players. A *strategy profile* $s \in S$ is an n -tuple $s = (s_1, \dots, s_n)$ such that each entry s_i of s is a strategy of player i . We use the common notational convenience and write a strategy profile as $s = (s_i, s_{-i})$ where $s_{-i} \in \prod_{j \in R - \{i\}} S_j$. Notice that a strategy $s_i \in S_i$ of player i is a vector of size m (with m being the number of edges in the graph), since a strategy for a player consists of a nonnegative payment for each edge e of G . Without loss of generality, we will assume $s_i(e) \leq \max_e c(e)$. Notice that S is a nonempty, convex, and compact set since S is a cube in $R^{n \times m}$: specifically it is just the cross product of the closed interval $[0, \max_e c(e)]$ taken nm times.

In order to prove the result, we use the technique used in Nash’s proof for showing existence of mixed Nash equilibrium in finite games, that uses Kakutani’s fixed point theorem. Our proof uses standard techniques, except for the

part showing that the graph $\Gamma(F)$ is closed, which requires somewhat different arguments due to the fact that our cost functions are not continuous.

Recall that Kakutani's fixed point theorem is defined as follows:

Theorem 4 (Kakutani's Fixed Point Theorem). *Let S be a non-empty, compact and convex subset of some Euclidean space R^n . Let $F : S \rightarrow 2^S$ be a set-valued function on S with a closed graph $\Gamma(F)$, and the property that $F(s)$ is nonempty and convex for all $s \in S$. Then F has a fixed point.*

A set-valued function $F : S \rightarrow 2^S$ is some rule that maps each element $s \in S$ to a subset of S , i.e., $F(s) \subset S$. Notice that each element of $F(s)$ is a strategy profile of the capacity allocation game. Since S is nonempty, compact, and convex, Kakutani's fixed point theorem states that if the function graph $\Gamma(F) = \{(s, t) | s \in S, t \in F(s)\}$ (which is a subset of the product space $S \times S$) is a closed set, and $F(s)$ is a nonempty and convex set for all $s \in S$, then there exists $s \in S$ such that $s \in F(s)$, i.e., a fixed point.

For a strategy profile $s = (s_i, s_{-i})$, let $\chi_i(s_{-i})$ denote the set of best responses of player i to the strategies s_{-i} of other players. Given the strategies s_{-i} of other players (which correspond to some capacity reservation on the edges of G), each element $s'_i \in \chi_i(s_{-i})$ is a minimum cost strategy of player i that will ensure that a flow of size 1 can be send from the source to i in the distribution network purchased by (s'_i, s_{-i}) . It is easy to see that for each s_{-i} we can express $\chi_i(s_{-i})$ as the set of optimal solutions of a linear program, with only non-strict inequalities. Therefore, $\chi_i(s_{-i})$ is a closed and convex subset of S_i for all $s_{-i} \in \prod_{j \in R - \{i\}} S_j$. Moreover, $\chi_i(s_{-i})$ is non-empty, since player i always has at least one best response.

We define the mapping $F : S \rightarrow 2^S$ as follows. Given a strategy profile s , $t \in F(s)$ if t is a strategy profile that can be obtained if each player $i \in R$ deviates from her strategy s_i to one of her best responses, i.e., to an element of $\chi_i(s_{-i})$. Formally, we define F as $F(s_1, \dots, s_n) = \{(t_1, \dots, t_n) | t_i \in \chi_i(s_{-i})\}$. In other words, $F(s_1, \dots, s_n) = \prod_{i \in R} \chi_i(s_{-i})$. Since $F(s)$ is the product space of nonempty, closed, and convex sets, then $F(s)$ is nonempty, closed, and convex for all $s \in S$. Therefore, if the graph $\Gamma(F) = \{(s, t) | s \in S, t \in F(s)\}$ is a closed set, then by Kakutani's fixed point theorem there exists $s \in S$ such that $s \in F(s)$. Notice that $s \in F(s)$ if and only if $s_i \in \chi_i(s_{-i})$ for all players i , i.e., the strategy of all the players is a best response of them to the strategies of the other players. Hence, a fixed point of F is a Nash equilibrium of the capacity allocation game. Therefore, in order to complete the proof all we need to show is that $\Gamma(F)$ is a closed set.

Let $(x^1, y^1), (x^2, y^2), \dots$ be an arbitrary convergent sequence of points in $\Gamma(F)$, and denote its limit by (x^*, y^*) . To show that $\Gamma(F)$ is closed, all we need to show is that $(x^*, y^*) \in \Gamma(F)$. Recall that $\Gamma(F) \subset S \times S$, and $S \times S$ is closed, so $(x^*, y^*) \in S \times S$. Therefore, all we need to show is that $y_i^* \in \chi_i(x_{-i}^*)$ for all $i \in R$. In order to do this, fix some arbitrary player i .

Let $C_i(s)$ denote the cost of player $i \in R$ for strategy profile s , and let $\hat{C}_i(s)$ be the cost of i 's best response to s , i.e., $\hat{C}_i(s) = \min_{s'_i \in S_i} \{C_i(s'_i, s_{-i})\}$. Since

$C_i(s'_i, s_{-i})$ is minimized if and only if $s'_i \in \chi_i(s_{-i})$, then we can equivalently define \hat{C}_i as $\hat{C}_i(s) = C_i(s'_i, s_{-i})$ for some $s'_i \in \chi_i(s_{-i})$.

We next define a function $\overline{C}_i : S \rightarrow \mathcal{R}_{\geq 0}$ for each $i \in R$ as follows: $\overline{C}_i(s) = C_i(s) - \hat{C}_i(s)$. Notice that $\overline{C}_i(s) \geq 0$ for all $s \in S$ and $\overline{C}_i(s) = 0$ if and only if $s_i \in \chi_i(s_{-i})$. In other words, $\overline{C}_i(s) = 0$ if and only if s is a stable strategy profile for player i , i.e., player i does not have an incentive of unilateral deviation from s . Notice that for any point $(s, t) \in \Gamma(F)$, we have that $\overline{C}_i(t_i, s_{-i}) = 0$ since $t_i \in \chi_i(s_{-i})$ by definition of F . Therefore, $\overline{C}_i(y_i^k, x_{-i}^k) = 0$ for all $k > 0$. In the usual argument about the existence of Nash equilibrium, \overline{C}_i is continuous over S and therefore this completes the proof, since this implies that $\overline{C}_i(y_i^*, x_{-i}^*) = 0$, and thus that $y_i^* \in \chi_i(x_{-i}^*)$. In our game, however, \overline{C}_i is not continuous.

Recall that \overline{C}_i is defined as the difference of two functions, i.e., $\overline{C}_i(s) = C_i(s) - \hat{C}_i(s)$. The function $\hat{C}_i(s)$ is continuous on S , since for any strategy profile s' such that $\|s - s'\| \leq \epsilon$, we have that $|\hat{C}_i(s) - \hat{C}_i(s')| \leq \epsilon$. However, $C_i(s)$ is not necessarily continuous on S , since when the mincut between the source and i becomes less than 1 in the distribution network, then the cost for player i suddenly becomes unbounded. Let Δ_i be the set of strategy profiles where player i is feasible, i.e., $\Delta_i = \{s | f_i(s) \geq 1\}$. For any $s \in \Delta_i$, the cost $C_i(s)$ is simply equal to $|s_i| = \sum_e s_i(e)$. Thus C_i is clearly continuous on the domain Δ_i . Notice that Δ_i can be formulated as a set of linear constraints with non-strict inequalities and therefore, Δ_i is a closed (and convex) set. Since both \hat{C}_i and C_i are continuous on Δ_i , then \overline{C}_i is also continuous on Δ_i .

Notice that for any $k > 0$, we have that $(y_i^k, x_{-i}^k) \in \Delta_i$, since y_i^k is a best response of player i to x_{-i}^k , and thus results in a solution feasible for player i . Since Δ_i is closed, then (y_i^*, x_{-i}^*) is also in Δ_i . Thus, since $\overline{C}_i(y_i^k, x_{-i}^k) = 0$ for all k , and \overline{C}_i is continuous on Δ_i , then $\overline{C}_i(y_i^*, x_{-i}^*) = 0$. This implies that $y_i^* \in \chi_i(x_{-i}^*)$. Since this is true for all i , then $(x^*, y^*) \in \Gamma(F)$, as desired. ■

Theorem 5. *The price of anarchy for the capacity allocation game is N and this bound is tight.*

In network formation games on undirected graphs, bounding the price of stability is usually a lot more challenging than bounding the price of anarchy [8,10]. The price of stability is known to be 1 for arbitrary sharing in the discrete model, i.e., there exists a Nash equilibrium that buys the Steiner Tree [9]. However, the analysis for the discrete model does not carry over for capacity allocation games (i.e., the fractional model), for which the price of stability is shown to be at least 2 in Section 4. In fact, the example in Section 2 shows that the Steiner tree is not necessarily a Nash equilibrium for our game, and so it is not the case that this factor of 2 arises simply because of the gap between integral and fractional solutions. Thus, even though there always exists a Nash equilibrium that buys the cheapest integral solution in any 2-tier topology (which is the minimum spanning tree in this case), this is not true for general undirected networks.

Though we do not have an upper bound for the price of stability in general undirected networks and therefore cannot guarantee the existence and efficient computation of cheap Nash equilibrium, we prove that there always exists a

cheap *approximate* Nash equilibrium that can be efficiently computed. By an α -approximate Nash equilibrium, we mean a solution where no player can reduce her cost by a factor of α by unilateral deviation [9].

Similar to the 2-tier topology in spirit, we start with a cheap integral feasible solution T and form payments on the edges of T . The cheapest integral solution T is the Steiner tree that connects the source and the receiver nodes; however, the Steiner Tree is not efficiently computable. We therefore use an approximation to Steiner Tree that does not cost more than twice the cost of OPT. We obtain the integral solution by using the primal-dual approximation algorithm for minimum-cost Steiner Forest problems (see, e.g. [24]). We use T to denote the tree returned by the primal-dual algorithm with $R \cup \{s\}$ being terminal nodes, and root it at s . Since the primal-dual algorithm uses the fractional optimal solution OPT as its lower bound, the cost of the integral solution returned by the primal-dual algorithm is not only within a factor of 2 of the cost of the minimum-cost Steiner Tree, but also the cost of OPT. The proof of Theorem 6 gives a polynomial-time algorithm that returns a strategy profile $\mathbf{p} = (p_i(e), i \in R, e \in E)$ that is a 2-approximate Nash equilibrium, and reserves 1 unit of bandwidth on the edges of T .

The rest of this section is devoted to proving Theorem 6.

Theorem 6. *There is a polynomial-time algorithm that returns a 2-approximate Nash equilibrium whose social cost is at most twice the cost of OPT.*

We use the term *edge block* to refer to maximal length paths of T whose interior nodes are degree 2 nonreceiver nodes. Notice that each edge of T is part of an edge block. It is easy to show that there are at most $2n - k$ edge blocks in T , where n is the number of receiver nodes, and k is the number of non-leaf receiver nodes. Notice that the removal of an edge block \mathbf{e} from T will divide T into 2 connected components $T_1(e)$ and $T_2(e)$. Without loss of generality we will assume that \mathbf{e} constitutes the cheapest path between these connected components in the rest of the text, since otherwise we can obtain a cheaper integral solution with this property by simply replacing the cheapest path between these 2 connected components for \mathbf{e} .

Our payment algorithm assigns each edge block \mathbf{e} to a receiver r_i and asks r_i to reserve 1 unit of bandwidth on *all* the edges of \mathbf{e} . We use the notation $p_i = \{\mathbf{e}\}$ if edge block \mathbf{e} is assigned to receiver r_i and all other edge blocks of T are assigned to the other receivers. Similarly, we say $p_i = \{\mathbf{e}, \mathbf{f}\}$ if edge blocks \mathbf{e} and \mathbf{f} are assigned to receiver r_i and all other edge blocks of T are assigned to the other receivers. Let $\chi_i\{\mathbf{e}\}$ and $\chi_i\{\mathbf{e}, \mathbf{f}\}$ denote the cheapest deviations of receiver r_i under the two possible strategies of her described above. Let $|\chi_i\{\mathbf{e}\}|$ and $|\chi_i\{\mathbf{e}, \mathbf{f}\}|$ denote the cost of these deviations to player i . We can now show the following lemmas.

Lemma 1. *For an edge block \mathbf{e} between s and r_i on T , we have that $\chi_i\{\mathbf{e}\} = \{\mathbf{e}\}$.*

Proof. For an edge block \mathbf{e} between s and r_i on T , the cheapest deviation $\chi_i\{\mathbf{e}\}$ of receiver r_i to strategy $p_i = \{\mathbf{e}\}$, where other players are reserving 1 unit of

bandwidth on all other edge blocks of T , must reserve enough capacity to send 1 unit of traffic between $T_1(e)$ and $T_2(e)$. The cheapest way to do this is to reserve 1 unit of bandwidth along the cheapest path between $T_1(e)$ and $T_2(e)$, and thus $|\chi_i\{e\}|$ is at least the cost of this path. Since e constitutes the cheapest path between $T_1(e)$ and $T_2(e)$, we have that $\chi_i\{e\} = \{e\}$. ■

Lemma 2. *For any two edge blocks e and f between s and r_i on T , we have $|\chi_i\{e, f\}| \geq \max\{|\chi_i\{e\}|, |\chi_i\{f\}|\}$.*

Proof. By Lemma 1, we know that $\chi_i\{e\}$ is the cheapest path between s and r_i if all the edge blocks of T other than e are bought by other players, and similarly for $\chi_i\{f\}$. By the same argument as in the previous lemma, we know that $\chi_i\{e, f\}$ costs at least as much as the cheapest path between s and r_i if all the edge blocks of T other than e and f are bought by other players. The lemma holds since r_i can find a cheaper path between r_i and s if other players buy more edges. ■

Lemma 1 and Lemma 2 imply that we will form a 2-approximate Nash equilibrium if for each receiver r_i , we assign at most 2 edge blocks between r_i and s to it. This is due to the fact that if r_i is assigned edge blocks e and f , then $|\chi_i\{e, f\}| \geq \max\{|\chi_i\{e\}|, |\chi_i\{f\}|\}$ (by Lemma 2), which is exactly the maximum of the cost of e and f by Lemma 1. Thus, the cost of r_i 's best deviation $|\chi_i\{e, f\}|$ to its current strategy costs at least half of what r_i is currently paying. Since this is true for every player, this forms a 2-approximate Nash equilibrium. All that is left to show is that we can form such an assignment.

Recall that there are at most $2n$ edge blocks of T and we can make an assignment where each receiver is assigned at most 2 edge blocks. However, we cannot make an arbitrary assignment of edge blocks to receivers, since Lemma 1 and Lemma 2 hold only if a receiver r_i is assigned edge blocks that are between r_i and s . In order to make an assignment with the desired property, we will root T at s , and loop through the edge blocks of T in the reverse BFS order from s . For each edge block e in this order, we select an arbitrary receiver r_i under e that is not assigned 2 edge blocks yet, and assign e to r_i . It is easy to show by induction that at the time the algorithm decides the assignment of an edge block e , there is always a receiver r_i that is not assigned 2 edge blocks yet. By the above argument, this assignment creates a 2-approximate Nash equilibrium that purchases 1 unit of bandwidth on edges of T , and thus has cost at most twice that of OPT .

6 Experimental Results

In this section, we experimentally evaluate, using randomly generated networks, the performance of the algorithm used in the proof of Theorem 2 for the 2-tier network model, and the algorithm used in the proof of Theorem 6 for general undirected networks. Experimental studies indicate that both of the algorithms perform much better than the theoretical guarantees in the corresponding theorems.

We first present the experimental results for the 2-tier network model. Let β denotes the ratio of the cost of the Nash equilibrium (as computed by our MST-based algorithm used in the proof on Theorem 2) to the cost of the socially optimal solution. Recall that $\beta \leq 2$ by Theorem 2. In the representative results shown in Table 1, the total number of non-receiver nodes, i.e., the relay nodes plus the source, is varied from 5 to 25. For each value of the number of nodes, we compute the maximum and average values of β , namely β_{avg} and β_{max} , over 200 random runs (network samples). Edges are drawn between the non-receiver nodes randomly, in the following manner: the nodes are picked one by one, and the node picked at any step is connected to each of the nodes already included (in previous steps) with probability $p = 0.5$. If the chosen node remains unconnected at the end of the step, to maintain connectivity, an edge is drawn between this node and a randomly chosen node that is already included. The cost of these edges follows a uniform distribution between 1 and 100. Finally receivers are assigned to the relay nodes randomly, such that each relay node is associated with at least one user: first we assign one receiver to each relay node, and then assign the remaining receiver nodes to the relay nodes randomly. The cost of edges between the receivers and their peers follows a uniform distribution between 1 and 5. Table 1(a) and (b) shows the results for two different numbers of (receiver nodes/non-receiver nodes) ratios. From the results, we observe that β value of the solution computed by our algorithm is very close to 1 on the average, and less than 1.5 in the worst case.

Table 1. Cost-approximation values for randomly generated 2-tier networks

#non-receivers	5	10	15	20	25	#non-receivers	5	10	15	20	25
β_{avg}	1.093	1.199	1.201	1.195	1.194	β_{avg}	1.079	1.174	1.162	1.147	1.130
β_{max}	1.49	1.48	1.37	1.35	1.36	β_{max}	1.37	1.42	1.34	1.29	1.28

(a) #receiver nodes = $2 \times$ #non-receiver nodes (b) #receiver nodes = $4 \times$ #non-receiver nodes

Next we experimentally evaluate the performance of our algorithm used in the proof of Theorem 6 for general undirected networks (which buys a primal-dual approximation of the Steiner Tree), to compute α -approximate Nash equilibria that attains a cost that is within a factor of β of the cost of OPT . Recall that both α and β are at most 2 by Theorem 6. Our experimental study on randomly generated networks show that the observations on β are similar to those observed for the 2-tier model, so we only show the α values in Table 2. In these experiments, the total number of nodes (receivers, relays, and the source) is varied from 20 to 100. For each value of the number of nodes, we compute the maximum and average values of α , namely α_{avg} and α_{max} , over 500 random runs (network samples). Edges are drawn randomly as in the 2-tier model, but across all nodes instead of only the non-receiver nodes. The cost of these edges follows a uniform distribution between 1 and 100, as before. From Table 2, we observe that α value of the solution computed by our algorithm is very close to 1 on average,

and equals 1 for a large fraction of the networks. Therefore, for general networks, our algorithm generates a solution that is an exact equilibrium or extremely close to an equilibrium, and has low cost (typically within 1.5 times the socially optimal cost). For (receiver/non-receiver) node number ratio of 2 (not shown in the table), α was observed to be 1 in all 500 runs (network samples). Therefore, from these results we conclude that a larger (receiver/non-receiver) node number ratio creates better (smaller) α values.

Table 2. Equilibrium-approximation values for randomly generated general networks

#nodes	20	40	60	80	100	#nodes	20	40	60	80	100
α_{avg}	1.0003	1.0003	1.0002	1.0009	1.0007	α_{avg}	1.0000	1.0009	1.0000	1.0000	1.0000
α_{max}	1.1250	1.1429	1.0833	1.2000	1.3333	α_{max}	1.0000	1.1667	1.0000	1.0000	1.0000

(a) #receiver nodes = $0.5 \times$ #non-receiver nodes (b) #receiver nodes = #non-receiver nodes

References

1. Web could collapse as video demand soars. Daily Telegraph (April 7, 2008)
2. ITU IPTV Focus Group, <http://www.itu.int/ITU-T/IPTV/>
3. Future Looks Bright For IPTV. Satellite Today (May 2, 2005)
4. <http://www.youtube.com>
5. Cisco Virtual Video Infrastructure Managing Complexity and Scale in a Next-Generation Video Network. White Paper, Cisco Systems (2008)
6. Ahlswede, R., Cai, N., Li, S., Yeung, R.: Network information flow. *IEEE Trans. Inform. Theory* 46(4), 1204–1216 (2000)
7. Anshelevich, E., Caskurlu, B.: Price of Stability in Survivable Network Design. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 208–219. Springer, Heidelberg (2009)
8. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing* 38(4), 1602–1623 (2008)
9. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-Optimal Network Design with Selfish Agents. *Theory of Computing* 4, 77–109 (2008)
10. Bild, V., Caragiannis, I., Fanelli, A., Monaco, G.: Improved Lower Bounds on the Price of Stability of Undirected Network Design Games. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 90–101. Springer, Heidelberg (2010)
11. Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., Orda, A.: Non-cooperative multicast and facility location games. In: Proceedings of the 7th ACM Conference on Electronic Commerce (EC), Ann Arbor, Michigan, pp. 72–81 (2006)
12. Chen, H., Roughgarden, T.: Network Design with Weighted Players. In: SPAA (2006)
13. Chen, H., Roughgarden, T., Valiant, G.: Designing Networks with Good Equilibria. In: SODA 2008 (2008)
14. Christodoulou, G., Koutsoupias, E.: On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)

15. Epstein, A., Feldman, M., Mansour, Y.: Strong Equilibrium in Cost-Sharing Connection Games. In: EC 2007 (2007)
16. Feigenbaum, J., Papadimitriou, C., Shenker, S.: Sharing the Cost of Multicast Transmissions. *Journal of Computer and System Sciences* 63, 21–41 (2001)
17. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the Price of Stability for Designing Undirected Networks with Fair Cost Allocations. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006, Part I. LNCS, vol. 4051, pp. 608–618. Springer, Heidelberg (2006)
18. Ho, T., Medard, M., Shi, J., Effros, M., Karger, D.R.: On randomized network coding. In: Proc. 41st Annual Allerton Conf. Comm., Control, & Computing, Monticello, IL (October 2003)
19. Hoefer, M.: Non-cooperative Facility Location and Covering Games. In: Asano, T. (ed.) ISAAC 2006. LNCS, vol. 4288, pp. 369–378. Springer, Heidelberg (2006)
20. Hoefer, M., Krysta, P.: Geometric Network Design with Selfish Agents. In: Wang, L. (ed.) COCOON 2005. LNCS, vol. 3595, pp. 167–178. Springer, Heidelberg (2005)
21. Li, S.-Y.R., Yeung, R.W., Cai, N.: Linear network coding. *IEEE Trans. Inform. Theory* 49(2), 317–381 (2003)
22. Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V.(eds.): *Algorithmic Game Theory*. Cambridge University Press
23. Sanders, P., Egner, S., Tolhuizen, L.: Polynomial time algorithms for network information flow. In: Proc.of the Fifteenth Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 286–294 (2003)
24. Vazirani, V.V.: *Approximation algorithms*. Springer, Berlin (2001)