Online Multicast with Egalitarian Cost Sharing

Moses Charikar
Dept. of Computer Science
Princeton University
Princeton, NJ 08540
moses@cs.princeton.edu

Howard Karloff
AT&T Labs
180 Park Avenue
Florham Park, NJ 07932
howard@research.att.com

Claire Mathieu
Computer Science Dept.
Brown University
Providence, RI 02912
claire@cs.brown.edu

Joseph (Seffi) Naor*
Computer Science Dept.
Technion
Haifa 32000, Israel
naor@cs.technion.ac.il

Michael Saks†
Dept. of Mathematics
Rutgers University
Piscataway, NJ 08854
saks@math.rutgers.edu

ABSTRACT

We consider a multicast game played by a set of selfish noncooperative players (i.e., nodes) on a rooted undirected graph. Players arrive one by one and each connects to the root by greedily choosing a path minimizing its cost; the cost of using an edge is split equally among all users using the edge. How large can the sum of the players' costs be, compared to the cost of a "socially optimal" solution, defined to be a minimum Steiner tree connecting the players to the root? We show that the ratio is $O(\log^2 n)$ and $\Omega(\log n)$, when there are n players. One can view this multicast game as a variant of ONLINE STEINER TREE with a different cost sharing mechanism.

Furthermore, we consider what happens if the players, in a second phase, are allowed to change their paths in order to decrease their costs. Thus, in the second phase players play best response dynamics until eventually a Nash equilibrium is reached. We show that the price of anarchy is $O(\log^3 n)$ and $\Omega(\log n)$.

We also make progress towards understanding the challenging case where arrivals and path changes by existing terminals are interleaved. In particular, we analyze the interesting special case where the terminals fire in random order and prove that the cost of the solution produced (with arbitrary interleaving of actions) is at most $O(\text{polylog}(n)\sqrt{n})$ times the optimum.

Categories and Subject Descriptors

F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General

*Supported in part by ISF grant 1366/07 and BSF grant 2002276. Part of this work was done while visiting Microsoft Research.

†This work was supported in part by NSF grant CCF-0515201. Part of this work was done while visiting Microsoft Research.

General Terms

Algorithms, Theory

Keywords

Price of Anarchy, Shapley Value, Best Response, Nash Equilibrium.

1. INTRODUCTION

Given an undirected connected rooted graph $G = (V, E)$ with nonnegative edge costs $c(e)$ for $e \in E$, consider a sequence $X = t_1, \ldots, t_n$ of *players*, or *terminal* vertices, who arrive in an online fashion; each new terminal greedily chooses a path connecting it to the root r so as to minimize its payment, as determined by the *egalitarian* (or Shapley) model: the cost of an edge is split evenly among the players currently using it in their paths to the root. Thus, if k players use edge e, then each player pays $c(e)/k$ for using e. Such egalitarian cost-sharing was recently introduced in an algorithmic game-theoretic context by Anshelevich et al. [2] and further studied by [3, 4, 5]. (It is also the outcome of the Shapley value; see [14] for details.)

The offline structure minimizing the total cost to all the players (the so-called "social optimum") is the minimum Steiner tree on $X \cup \{r\}$. How much larger is the cost of the structure resulting from a sequence of selfish choices of cost-minimizing paths? The answer turns out to be polylogarithmic: here, we prove that its cost is larger by a factor which is no larger than $O(\log^2 n)$ (Theorem 1), and that there are instances for which it is larger by a factor of $\Omega(\log n)$. This dramatically improves upon the previously known upper and lower bounds of $O(\sqrt{n}\log n)$ and $\Omega(\log n/\log\log n)$ [3].

Our upper bound is proven by a gap-revealing linear program, an argument which we find appealingly simple. Our lower bound is proven by relating the problem to the ONLINE STEINER TREE problem, whose competitive ratio [7] is known to be $\Theta(\log n)$. (Recall that the *competitive ratio* of an online algorithm A is the worst case, over all input sequences X, of the value of A for input X, divided by the value of the optimal offline solution for X.)

Naturally, later arriving terminals may render certain paths more attractive and produce some regret for terminals who have previously chosen paths which are no longer cheapest for them. In order to deal with such regret, Chekuri et al. [3], following a first phase in which the players arrive one by one, proposed adding a second phase during which any existing player may change its path to a different, currently cheapest path (i.e., use *best response dynamics*): that is, once all players have arrived, players are allowed to "refire" or "replay" an arbitrary number of times and in arbitrary order, until a Nash equilibrium is reached, in the sense that no player has any incentive to deviate unilaterally from its currently-chosen path. This defines a two-phase multicast game. The inefficiency resulting from such a non-cooperative game is quantified by the *price of anarchy* [8], the worst-case ratio between the cost of a Nash equilibrium and the cost of the optimum Steiner tree. Here, as a corollary of Theorem 1, we prove that the ratio is $O(\log^3 n)$ when all initial firings precede the first refiring; and our lower bound extends to prove that the worst-case ratio is also $\Omega(\log n)$. Again, this dramatically improves upon the previous bounds of $O(\sqrt{n}\log^2 n)$ and $\Omega(\log n/\log\log n)$ [3], also proven in the case in which all initial firings precede the first refiring.

We emphasize that in our game all players must arrive before any player refires. Thus, we do not talk about a Nash equilibrium until all nodes have arrived, each firing once. Only then does the second phase, of refirings, begin, and then we discuss a Nash equilibrium. In the last section, we pose as an open problem the case in which arbitrary intermingling of initial firings and refirings is allowed.

For the two-phase multicast game with egalitarian cost sharing, we have thus proven that the Nash equilibria that are *reachable* by our best response dynamics (all initial firings before any refirings) have cost no more than a polylogarithmic factor times the optimum. In contrast, for this game there do exist Nash equilibria of the traditional kind, without any issue of reachability, with cost larger than the optimum by a factor of n: consider an instance [2] with two nodes, s and r, and two parallel edges between s and r, one of cost n, one of cost $1+\epsilon$. Now suppose that n players use the heavy edge, each shouldering one nth of the burden. This configuration has cost n, whereas the optimal cost is $1+\epsilon$. Yet, no player will move to the light edge, and so this is a Nash equilibrium; hence the price of anarchy is (at least) n. In a dynamic environment as modeled by our game, *this equilibrium simply cannot be reached*, since we start from an initial empty configuration.

The latter example motivated Anshelevich et al. to introduce the notion of *price of stability*: the ratio of the cost of the *cheapest* Nash equilibrium to the cost of the optimal solution. They showed that this ratio is $O(\log n)$ for the multicast game on undirected graphs, a bound slightly improved by Agarwal and Charikar [1] to $O(\log n/\log\log n)$, and by Fiat et al. [5] to $O(\log\log n)$ in the special case in which every node in the graph is a terminal. However, the proof of the logarithmic upper bound on the price of stability of [2] heavily relies on the assumption that the starting point is an optimal (or near optimal) Steiner tree on the terminals.

In light of this prior work, one can view our result either as a result about ONLINE STEINER TREE with egalitarian cost sharing (Phase 1 alone), or as a result about the modest price of anarchy in a multicast game in which price of anarchy is defined only over configurations which are reachable from the initial empty configuration, albeit in a specific way (all initial firings preceding all refirings).

Extending our analysis to the more general case in which terminals may refire before all terminals have arrived, seems to be surprisingly challenging. In this case, the online sequence is an arbitrary sequence of "firings" of the terminals, with repetitions, where the first firing corresponds to the arrival of the terminal. We make some progress towards understanding this case by studying the version where the firing sequence is random. For this version we prove that the cost of the solution produced (with arbitrary interleaving of actions) is at most $O(\text{polylog}(n)\sqrt{n})$ times the optimum.

2. ONLINE MULTICAST WITH EGALITARIAN COST SHARING: UPPER BOUND

Theorem 1 *The greedy algorithm for online multicast with egalitarian cost sharing has competitive ratio $O(\log^2 n)$.*

Let $H(n) = 1+1/2+1/3+\cdots+1/n$, be the nth harmonic number. Assume the terminals are labeled so that their arrival order is $< t_1, t_2, ..., t_n >$. Denote by $d(\cdot,\cdot)$ the distance function on the graph G as defined by the edge cost $c(\cdot)$.

Lemma 1 *The total cost of all edges used by the terminals' paths is bounded above by the optimal value of the following "gap-revealing" linear program with variables $s(1),...,s(n),b(1),...,b(n)$:*

$$\max \quad \sum_{i=1}^{n} b(i) \quad \textit{s.t.}$$

$$\begin{aligned} s(j)-s(i)+b(i)/2 &\leq d(i,j) \quad \forall\, 1\leq i<j\leq n \\ \sum_i s(i) - \sum_i b(i)H(n) &\leq 0 \\ s(0)=b(0) &= 0 \quad \textit{(for the root)} \\ s(i), b(i) &\geq 0 \quad \forall\, 1 \leq i \leq n \end{aligned}$$

PROOF. It suffices to construct a feasible solution to the above LP whose objective function value equals the total cost of the edges used by the terminals. For $1\leq i\leq n$, let P_i denote the path selected by terminal t_i. Let $s(i)$ be the cost share of P_i upon arrival of t_i, specifically,

$$s(i) = \sum_{e\in P_i} c(e)/|\{j : (e\in P_j)\wedge(1\leq j\leq i)\}|$$

and let $b(i)$ to be the sum of $c(e)$ over those e that were used for the first time by P_i.

We claim that this choice of $s(i), b(i)$ is feasible for the LP. The first set of constraints is equivalent to an observation of Chekuri et al. [3]: For any two terminals i and j, $1\leq i<j\leq n$,

$$s(j)\leq d(i,j)+s(i)-b(i)/2.$$

To see this, note that when terminal j selects his path, one option is to select the shortest path to i (costing $d(i,j)$), and then follow the same path selected by i, which costs at most $s(i)-b(i)/2$ since the cost share to j of using edges that were first used by i is at most $b(i)/2$.

For the second set of constraints, observe that the total of all cost shares for the selected path by each terminal i upon arrival is at most:

$$\Phi = \sum_i (b(i)+s(i)) \leq \sum_{e\in\cup_i P_i} c(e)H(n) = \sum_i b(i)H(n).$$

Finally, the total cost of the selected tree is given by $\sum_i b(i)$, which is the value of the objective function. ■

To get an upper bound on the optimal value of this LP, we consider a relaxation of this LP, and construct a solution to the dual of the relaxation.

In what follows, we view the root r as a terminal t_0 that arrived before any other terminal and define $W = X\cup\{t_0\} =$

$\{t_0, \ldots, t_n\}$. Let T be a tree on vertex set W rooted at t_0 (not necessarily a subgraph of the network) such that for each $j \in X$, its parent $p(j)$ in T arrived before j. For a node j, let $C(j)$ denote the set of children in T. Consider the relaxation of the LP in Lemma 1, denoted LP_T, where for the first set of constraints we only keep those in which $i = p(j)$. The dual of this linear program, DLP_T, has dual variables $\{z_j : j \in X\} \cup \{y\}$ and has the form:

$$\min \quad \sum_{j \in X} d(p(j), j) z_j \quad \text{s.t.}$$

$$\begin{aligned} -H(n)y + \sum_{j \in C(i)} z_j/2 &\geq 1 \ \forall i \in W \\ y + z_i - \sum_{j \in C(i)} z_j &\geq 0 \ \forall i \in W \\ z_i &\geq 0 \ \forall i \in W \\ y &\geq 0 \end{aligned}$$

For any choice of T and any feasible solution to DLP_T we get an upper bound on the optimal value of the LP in Lemma 1 and therefore an upper bound on the cost of tree produced by the algorithm.

The next lemma constructs a specific rooted tree T' with some nice properties. We then show how to construct a low cost feasible solution to $DLP_{T'}$ that yields the upper bound of Theorem 1.

When we write $d(i, j)$, we mean the distance in the underlying network (not in the tree T' being constructed). The *level* of a node is one more than its distance from the root in the tree T' (so the root is at level 1). For $j \in X$ we write $p(j)$ for the parent of j in $'$.

Lemma 2 *Let τ be a positive integer. There exists a rooted tree T' (not necessarily a subgraph of the input graph) on the set $W = \{t_0, \ldots, t_n\}$ with root t_0 such that:*

1. *For every terminal $t \neq t_0$, the parent of t arrived earlier than t.*
2. *Every node has at most two children.*
3. *Every node u with two children is at level divisible by τ.*
4. $\sum_{j \in X} d(j, p(j)) \leq 2(\tau + 1)\lceil \log n \rceil \cdot$ *OPT.*

PROOF. Select a minimum Steiner tree R on the network. Let $\pi = (\pi(0), \pi(1), \ldots, \pi(n))$ be the ordering of W that lists the terminals in order of first appearance along an Eulerian tour of R that starts from t_0. We have $\pi(0) = t_0$ and $\sum_i d(\pi(i), \pi(i+1)) \leq 2 \cdot \text{OPT}$.

The construction of T' is recursive. Let S be the set of τ terminals, $t_0, \ldots, t_{\tau-1}$ with earliest arrival time (listed in increasing order of arrival time). Let π' be the sequence obtained from π by deleting S. Split π' into two sequences π_1, π_2 of nearly equal size, so that $|\pi_1| \leq |\pi_2| \leq |\pi_1|+1$. For $i \in \{1, 2\}$, let r_i be the terminal in π_i that arrived first. Recursively build trees T_1 for π_1 with root r_1 and T_2 for π_2 with root r_2. The output tree T' consists of the union of the path $t_0, \ldots, t_{\tau-1}$, the edges $(t_{\tau-1}, r_1)$ and $(t_{\tau-1}, r_2)$ and the trees T_1, T_2.

By construction, T' satisfies the first three properties. Furthermore, a simple induction on n shows that the number of levels of T' is at most $\tau\lceil \log n \rceil$.

For the final property claimed in the lemma we want to prove an upper bound on $\sum_{j \in X} d(j, p(j))$. For each terminal $j \neq \pi(n)$ let $f(j)$ be the terminal that immediately follows it in π order and let $D(j)$ be the set of terminals that appear between j and $p(j)$ in π order together with the member of $\{j, p(j)\}$ that appears first in π order. By repeated use of the triangle inequality $d(j, p(j)) \leq \sum_{i \in D(j)} d(i, f(i))$ and therefore

$$\sum_{j \in X} d(j, p(j)) \leq \sum_{i \in X} d(i, f(i)) |\{j : i \in D(j)\}|.$$

Claim 3 *For any two terminals j, j', if $D(j) \cap D(j') \neq \emptyset$ then $p(j)$ is an ancestor of $p(j')$ or $p(j')$ is an ancestor of $p(j)$.*

To prove the claim, suppose, for contradiction that the least common ancestor k of $p(j)$ and $p(j')$ is distinct from both, and without loss of generality $p(j)$ is in the left subtree of k and $p(j')$ is in the right subtree. Then j and $p(j)$ both precede j' and $p(j')$ in π order and so $D(j) \cap D(j')$ is empty, a contradiction that proves the claim.

From the claim it follows that for any terminal i, there is a single root to leaf path that contains the parents of all terminals j such that $i \in D(j)$. Such a path contains at most $\tau\lceil \log n \rceil$ terminals, of which at most $\lceil \log n \rceil$ have two children, so $|\{j : i \in D(j)\|| \leq (\tau + 1)\lceil \log n \rceil$. We conclude that

$$\begin{aligned} \sum_{j \in X} d(j, p(j)) &\leq (\tau + 1)\lceil \log n \rceil \sum_{i \in X} d(i, f(i)) \\ &\leq 2(\tau + 1)\lceil \log n \rceil OPT, \end{aligned}$$

as required to complete the proof of the lemma. ■

Note that in the previous lemma, τ is a free parameter. We now complete the proof of Theorem 1 by fixing τ and defining a specific feasible solution to $DLP_{T'}$, for the tree T' defined in Lemma 2. The value of that solution will be the required upper bound.

For the moment we leave τ unspecified and also introduce another positive parameter a that will be fixed shortly. Let $y = \frac{a}{2(\tau-1)}$. For each integer q satisfying $1 \leq q \leq \tau$, for every terminal j whose level is congruent to $q \bmod \tau$ set $z_j = \frac{a}{2} + (q - 1)y$. Observe that for every j, $z_j \in [a/2, a]$.

For any choice of a and τ, the second set of dual constraints is satisfied. To see this, note that if the level of j is not a multiple of τ then j has one child, so $y + z_i - \sum_{j \in C(i)} z_j = 0$. If j is at level a multiple of τ then $z_j = a$ and j has two children each having z value $a/2$, so $y + z_j - \sum_{i \in C(j)} z_i = y \geq 0$.

For the first set of constraints, it is enough that $yH(n) \leq a/4 - 1$ and this is true if we take $\tau = \lceil 4H(n) \rceil + 1$ and $a = 8$.

By the third property of Lemma 2, the value of the solution is

$$\begin{aligned} \sum_j d(p(j), j) z_{p(j),j} &\leq \sum_j d(p(j), j) a \\ &\leq 16\lceil \log n \rceil (\lceil 4H(n) \rceil + 1) \cdot \text{OPT}, \end{aligned}$$

completing the proof of Theorem 1.

3. TWO-PHASE MULTICAST GAME: PRICE OF ANARCHY

Recall that we consider a *two-phase game* for connecting the terminals to the root. In Phase 1 the players play in the order $t_1, \ldots, t_n$, and each selects a greedy (best response) path relative to the selection of paths by the previous players. In Phase 2, the following step is repeated: take an arbitrary player among $t_1, \ldots, t_n$, whose current path is not a greedy path relative to the other paths, and replace its current path by a greedy path. When each player's path is the greedy path relative to the other paths, we have reached a Nash equilibrium and the game ends. Our goal is to analyze the cost of the resulting solution. (As the game proceeds, the union of the paths chosen by the players need not be a tree, but it can be shown [3] that any Nash equilibrium induces a tree.)

The multicast game belongs to the class, first defined by Rosenthal [12] and widely investigated [6, 9, 11, 14, 15], of *congestion games*. Rosenthal [12] showed that a potential function can be defined for each congestion game such that the potential decreases if a player makes a move that improves its selfish cost. It follows that every congestion game has a pure Nash equilibrium. Moreover, there is a one-to-one correspondence between Nash equilibria and the solutions defining a local minimum of Rosenthal's potential function. For the multicast game, Rosenthal's potential function Φ reduces to the following:

$$\Phi = \sum_{e \in E} \left(c(e) \cdot \sum_{i=1}^{n(e)} \frac{1}{i} \right),$$

where $n(e)$, the *usage of* e, denotes the number of players using edge e at the specified time. If a player has previously fired and changes its connection to the root from path P to path Q, then the potential function precisely captures the change in cost (to the player) from P to Q. Since best response dynamics can only decrease Rosenthal's potential function, it follows that this process must terminate in a Nash equilibrium. From Theorem 1, it is easy to deduce a bound on the price of anarchy of reachable Nash equilibria. (Chekuri et al. [3] used the same technique.)

Corollary 4 *The Nash equilibrium reached by the two-phase Multicast Cost Sharing game with best response dynamics has cost $O(\log^3 n)$OPT.*

PROOF. Consider Rosenthal's potential function Φ defined above. Let Φ_1 and Φ_2 be its values at the end of Phases 1 and 2, respectively. Since Φ cannot increase in Phase 2, $\Phi_2 \le \Phi_1$. Now the cost at the end of Phase 1 is at most $O(\log^2 n) \cdot$ OPT by Theorem 1. Hence, Φ_1 is at most $H(n)$ times this value, i.e., $O(\log^3 n) \cdot OPT$. But the cost of the graph at the end of Phase 2 is at most $\Phi_2 \le \Phi_1$. ∎

4. LOWER BOUNDS

Using the lower bound for the ONLINE STEINER TREE greedy algorithm, we prove a lower bound of $\Omega(\log n)$ for the greedy online Steiner problem with egalitarian cost sharing (i.e., Phase 1 alone), and for the two-phase multicast cost sharing game, improving upon the previous lower bound proof of $\Omega(\log n/\log\log n)$ by [3]. We note that our proof is simpler compared to the proof of [3].

In the ONLINE STEINER TREE problem, given are a graph $G = (V, E)$ and a root $r \in V$. The algorithm will maintain a connected subgraph T of G; initially $T = \{r\}$. In each step, given a new terminal t of V, the greedy algorithm selects a cheapest path from t to r and adds it to T. Its *competitive ratio* is the worst-case ratio between the cost of the tree it constructs and that of the cheapest Steiner tree on the union of $\{r\}$ and the set of the given terminals.

Theorem 5 *[7] The competitive ratio of the greedy algorithm for* ONLINE STEINER TREE *on n-vertex unweighted graphs is $\Omega(\log n)$.*

The proof of Theorem 5 given by Imase and Waxman uses the ℓ-*level diamond graph* with 4^ℓ edges. However, we will use the lower bound as a black box.

Theorem 6 *The competitive ratio of the greedy algorithm for online multicast with cost sharing and the price of anarchy of the two-phase multicast cost sharing game on $n(n+1)$-vertex graphs are both at least half of the competitive ratio of the greedy algorithm for* ONLINE STEINER TREE *on n-vertex unweighted graphs.*

PROOF. Consider an instance for ONLINE STEINER TREE on a graph $G = (V, E)$ with n vertices. We define an instance of the multicast game as follows. Let $N = n + 1$ and $\epsilon = 1/n^2$. Replace each vertex $v \in V$ by a star S_v of N vertices by adding $N - 1$ new vertices $v_1, v_2, ..., v_{N-1}$ each at distance ϵ from v. This defines a new graph H. Replace each request to a vertex $v \in V(G)$ by a "v-batch," i.e., a request to $v_0 = v$ of H followed by a sequence of requests to all the leaves $v_1, ..., v_{N-1}$ in $V(H)$ in the star S_v of H. This defines an instance of the multicast game.

An edge of length ϵ will be called a *short* edge; others (of length 1) will be called *long* edges. We will prove that during Phase 1 of the multicast cost sharing algorithm, the following invariant is maintained: *At any time, the union of the request paths forms a tree T rooted at r.* The invariant implies that at the beginning of every batch of requests, for every long edge e, the number of requests which use e is an integer multiple of N. It also implies that for every vertex $v_i \in S_v$, if v is in T then the path from v_i to r is the concatenation of the short edge from v_i to v and the *same* path from v to r. The invariant is valid initially. Assume that it is true so far, and consider a request to some vertex $v_i \in S_v$.

If this is the beginning of a new batch ($i = 0$), the path serving v traverses some number ℓ, $\ell \ge 1$, of long edges to first reach a vertex which is already in T. From that point onward, following the edges of T, each of which is used at least N times, costs at most n/N in total. Following unnecessarily even one long edge not in T would cost $1 > n/N$. It follows that the algorithm will simply minimize the number ℓ of non-T edges it needs to first hit T (just as the greedy ONLINE STEINER TREE algorithm does); from that point on, it will lazily follow T to r.

Now suppose $1 \le i \le N - 1$. When v_i is requested, it must take the short edge to v followed by some path from v to r. The main observation is that whatever path P that v_{i-1} chose to take to go from v to r is still the best path for v_i. (Compare P to some other path P'. Edges in the intersection I of P and P' contribute the same cost to P and P'. After v_{i-1} has chosen path P, edges in $P \setminus P'$ are cheaper than before, and edges in $P' \setminus P$ are the same price as before. Since P was of no greater cost than P' before path P was added to the collection, it is certainly still of no greater cost than P' afterward.)

Thus the invariant is maintained. The structure at the end of Phase 1 in the two-phase multicast cost sharing game is thus a tree corresponding to the tree produced by the ONLINE STEINER TREE algorithm on G. Moreover, this tree is a Nash equilibrium since no vertex can wish for a different path (even using one long edge not in T is a disaster), so Phase 2 is empty of events.

The output tree has total length exactly equal to $n(N-1)\epsilon + \ell(T)$, $\ell(T)$ denoting the number of long edges in T. The optimal multicast structure on H must be a tree, which has length $n(N-1)\epsilon + \ell(T^*)$, where T^* is the corresponding tree in G and $\ell(T^*)$ is the number of its long edges (all of them). By optimality, T^* is the optimal Steiner tree on G. Thus the competitive ratio of the optimal Steiner tree problem (on this instance) is $\ell(T)/\ell(T^*)$ and the price of anarchy of the multicast game (on this instance) is

$$\frac{n(N-1)\epsilon + \ell(T)}{n(N-1)\epsilon + \ell(T^*)} \ge \frac{\ell(T)}{n(N-1)\epsilon + \ell(T^*)} \ge (1/2)\frac{\ell(T)}{\ell(T^*)},$$

since $\ell(T^*) \ge 1 \ge n(N-1)\epsilon$ by the definition of ϵ. Finally, the competitive ratio of the optimal Steiner tree problem is simply the worst-case value of $\ell(T)/\ell(T^*)$. ■

Corollary 7 *The price of anarchy of the two-phase Multicast Cost Sharing game is* $\Omega(\log n)$.

5. MULTICAST WITH RANDOM ARRIVALS

In this section, we consider the setting in which arrivals and replays can be mixed in arbitrary order. The general setting seems quite challenging to analyze and giving a guarantee in this case is an open problem. We make some progress towards understanding this problem in the interesting special case of *random arrivals*. Here, the order of arrivals is a random permutation of the terminals, refirings are (adversarially) intermingled with arrivals. In this case, we show that the expected cost of the solution produced is $O(\text{polylog}(n)\sqrt{n})\cdot$ OPT.

More precisely, the model we consider can be described as a semi-random adversary model as follows. At each time step, an adversary decides either to refire a specific terminal that has previously joined, or decides that a new terminal should arrive. In the latter case, the new terminal is selected uniformly at random from among the terminals that have not yet arrived.

To analyze the cost of the solution produced, we analyze the evolution of the potential function.

5.1 Bounding the potential change

Let $\Psi(k)$ denote the value of the potential function immediately prior to the kth arrival and let $\Phi(k)$ denote the value of the potential function immediately following the kth arrival. Upon arrival of the $(k+1)$th new terminal v, the potential change is $\Phi(k+1)-\Psi(k+1)$, which is at most

$$\min\{d(v,u)+\sum_{e\in P(u)} c(e)/(n(e)+1): \ \ u \text{ arrived before } v\ \},$$

where $n(e)$ is the number of paths using edge e immediately prior to the arrival of v and $P(u)$ is the path currently used by vertex u when v arrives.

We can bound this minimum from above by considering an average over u selected according to some distribution. We now describe the distribution we will use. As in the previous section, let π be the permutation of the terminals giving their order of first appearance along an Eulerian tour of the minimum Steiner tree on $X\cup\{r\}$. Here we view π as a cyclic permutation. Below we will fix a positive integer s. Let $S(v)$ (resp. $S^-(v)$) denote the first s vertices following (resp. preceding) v along π. Let $R(k)$ denote the first k terminals of X that arrived (in time) and $B(k)$ denote the remaining terminals. We pick u uniformly at random among $S(v)\cap R(k)$.

Let $A_v(k+1)$ be the indicator of the event that v arrives at time $k+1$. Since replays can only decrease the potential function, $\Psi(k+1)\le\Phi(k)$, and so $\Phi(k+1)-\Phi(k)\le\Phi(k+1)-\Psi(k+1)$; then

$$\Phi(k+1)-\Phi(k)\ \le\ \sum_v A_v(k+1)\Big[\max_{u\in S(v)} d(v,u)$$
$$+\sum_{u\in R(k)\cap S(v)} 1(v\text{ picks }u)\sum_{e\in P(u)}\frac{c(e)}{n(e)+1}\Big]. \tag{1}$$

5.2 Chernoff Bound

Each fixed vertex belongs to $R(k)$ with probability k/n so the expected size of $R(k)\cap S(v)$ is sk/n and the expected size of $B(k)$ is $s(n-k)/n$. For the analysis we will need to show that with probability close to 1, for all v $|R(k)\cap S(v)|$ is not much smaller than its expectation (and $|B(k)\cap S(v)|$ is not much bigger than its expectation.) For fixed v, $R(k)\cap S(v)$ is a sum of s indicator random variables each corresponding to a vertex $w\in S(v)$ and indicating whether $w\in R(k)$. If these were independent we could use standard tail bounds for sums of independent random variables, but they are not independent. The problem of deriving tail bounds in similar cases has been considered extensively, but we don't know a result that it is in a form that is convenient for our purposes, so we prove it here.

Lemma 8 *Let* $\epsilon\in(0,1/2)$. *Let* $k_0=48\ln n/\epsilon^2$. *Let* $k\in[k_0,n-k_0]$ *and* $s=k_0n/\min(k,n-k)$. *Define the event* $E_1(k)$*:*

for every v,

$$|S(v)\cap R(k)|\ge\frac{ks}{n}(1-\epsilon)$$

and

$$|S(v)\cap B(k)|\ge\frac{(n-k)s}{n}(1+\epsilon).$$

Then, event $E_1(k)$ *has probability at least* $1-O(1/\epsilon n^2)$.

First observe that the definition of k_0 ensures that $s<n$. We will bound the two cardinalities separately and use the union bound for $1-\Pr\{E_1(k)\}$. By symmetry, $|S(v)\cap B(k)|$ has the same distribution as $|S(v)\cap R(n-k)|$, so we will just analyze $|S(v)\cap R(k)|$ and bound its tail distribution on both sides of the mean.

To prove the lemma, fix k, fix v, and let $Z=|S(v)\cap R(k)|$. As noted above, Z has expectation $(k/n)s$. Using a standard technique, we will obtain an upper bound on the probability that $Z<(1-\epsilon)E[Z]$, by approximating Z by a sum of independent random variables. Let $R'=X_1+\ldots+X_n$ and $Z'=X_1+\cdots+X_s$, where the X_i's are i.i.d. indicator random variables each having the same (as yet unspecified) expectation. Observe that when Z' is conditioned on the event $R'=k$, has the same as the distribution of Z, and when Z' is conditioned on the event $R'<k$, it is stochastically dominated by Z. Thus:

$$\begin{aligned}\Pr\{Z\le(1-\epsilon)(k/n)s\} &\le \Pr\{Z'\le(1-\epsilon)(k/n)s|R'\le k\}\\ &\le \frac{\Pr\{Z'\le(1-\epsilon)(k/n)s\}}{\Pr\{R'\le k\}}.\end{aligned}$$

Fixing the common expectations of the X_i to be $(k/n)/(1+\epsilon/2)$ we have $E[R']=k/(1+\epsilon/2)$, and by Markov's inequality, we obtain the (crude but adequate) bound:

$$\Pr\{R'\le k\}=1-\Pr\{R'>k\}>1-1/(1+\epsilon/2)>\epsilon/3.$$

Therefore we have:

$$\begin{aligned}&\Pr\ \ \{Z\le(1-\epsilon)(k/n)s\}\\ &\le\ \frac{3}{\epsilon}\Pr\{Z'\le(1-\epsilon)(k/n)s\}\\ &\le\ \frac{3}{\epsilon}\Pr\{Z'\le(1-\epsilon)(1+\frac{\epsilon}{2})E[Z']\}\\ &\le\ \frac{3}{\epsilon}e^{-\epsilon^2(k/n)s/(12(1+\epsilon/2))},\end{aligned}\tag{2}$$

where the last inequality is obtained by a standard Chernoff-type bound.

Similarly, since Z is stochastically dominated by Z' if $R'\ge k$,

$$\begin{aligned}&\Pr\ \ \{Z\ge(1+\epsilon)(k/n)s\}\\ &\le\ \Pr\{Z'\ge(1+\epsilon)(k/n)s|R'\ge k\}\\ &\le\ \Pr\{Z'\ge(1+\epsilon)(k/n)s\}/\Pr\{R'\ge k\}.\end{aligned}$$

Assume now that X_i satisfies[1] $E[X_i] = (k/n)/(1-\epsilon/2)$. Then $E[R'] = k/(1-\epsilon/2)$, and (given that the maximum possible value is n), by Markov's inequality, the probability that R' is less than k is at most $1 - \epsilon k/(2(n-k))$.

Therefore we have:

$$\begin{aligned}
\Pr\quad & \{Z \geq (1+\epsilon)(k/n)s\} \\
\leq\quad & \frac{2(n-k)}{\epsilon k} \Pr\{Z' \geq (1+\epsilon)(k/n)s\} \\
\leq\quad & \frac{2(n-k)}{\epsilon k} \Pr\{Z' \geq (1+\epsilon)(1-\epsilon/2)E[Z']\} \\
\leq\quad & \frac{2(n-k)}{\epsilon k} e^{-\epsilon^2 (k/n)s/(12(1-\epsilon/2))}, \qquad (3)
\end{aligned}$$

where the last inequality is obtained by standard Chernoff bounds. Summing the bounds in (2) and (3) and plugging in the value of $s = 4 \ln n \cdot 12n/(\min(k, n-k)\epsilon^2)$ proves the lemma.

5.3 Concluding the Analysis

We start from Equation (1). Writing $\sigma = rv_1 v_2 \ldots v_n$, the first part averages to

$$\begin{aligned}
E[D_1] &= \sum_i \Pr\{v_i \text{ arrives at } k+1\} \sum_{i \leq j < i+s} d(v_j, v_{j+1}) \\
&\leq (1/n)s\ 2\mathrm{OPT}. \qquad (4)
\end{aligned}$$

Rewrite the second part as

$$D_2 = \sum_{u \in R(k)} \sum_{e \in P(u)} \frac{c(e)}{n(e)+1} \sum_{v \in S^-(u) \cap B(k)} A_v(k+1) 1(v \text{ picks } u).$$

To bound the expectation of D_2, we condition on the event $E_1(k)$ from Lemma 8 with $\epsilon = 1/(c \ln(n))$. We have:

$$\begin{aligned}
E[D_2] &\leq E[D_2 | E_1(k)] + (1 - \Pr\{E_1(k)\}) \cdot \mathrm{OPT} \\
&\leq E[D_2 | E_1(k)] + \frac{O(\ln n)}{n^2} \cdot \mathrm{OPT}. \qquad (5)
\end{aligned}$$

So now we need to bound $E[D_2 | E_1(k)]$. The event $E_1(k)$ depends only on the history up to time k. Fix a history up to time k, such that $E_1(k)$ holds, and take expectations over v, being careful to do things in the correct order: the probability that v picks u is $1/|S(v) \cap R(k)|$, which can be bounded by the definition of the event $E_1(k)$. $E[A_v|\text{history}]$ equals $E[A_v | v \in B(k)] = 1/(n-k)$. The number of non-zero terms is $|S^-(u) \cap B(k)|$ and can also be bounded by the definition of event $E_1(k)$. Thus:

$$\begin{aligned}
E\quad & [D_2 | \text{history}] \\
\leq\quad & \sum_{u \in R(k), e \in P(u)} \frac{c(e)}{n(e)+1} \sum_{v \in S^-(u) \cap B(k)} \frac{1}{n-k} \frac{n}{(1-\epsilon)ks} \\
\leq\quad & \sum_{u \in R(k), e \in P(u)} \frac{c(e)}{n(e)+1} \frac{(1+\epsilon)s(n-k)}{n} \frac{1}{n-k} \frac{n}{(1-\epsilon)ks} \\
=\quad & \frac{1+\epsilon}{(1-\epsilon)k} \sum_e c(e) \frac{n(e)}{n(e)+1}.
\end{aligned}$$

Now, note that for any $i \geq 1$, we have

$$i/(i+1) \leq (1/2)(1 + 1/2 + 1/3 + \cdots + 1/i),$$

and so

$$\sum_e c(e)n(e)/(n(e)+1) \leq (1/2)\Psi(k+1).$$

[1] If this quantity is greater than n then the probability that Z exceeds it is 0.

Thus, we can substitute and average over histories such that $E_1(k)$ holds:

$$\begin{aligned}
E[D_2 | E_1(k)] &\leq \left(\frac{1+3\epsilon}{2k}\right) \cdot E[\Psi(k+1) | E_1(k)] \\
&\leq \left(\frac{1+3\epsilon}{2k}\right) \cdot E[\Phi(k) | E_1(k)],
\end{aligned}$$

since $\Psi(k+1) \leq \Phi(k)$ as replays can only decrease the potential function. Continuing we have:

$$\begin{aligned}
E[D_2 | E_1(k)] &\leq \left(\frac{1+3\epsilon}{2k}\right) \cdot \frac{E[\Phi(k)]}{\Pr\{E_1(k)\}} \\
&\leq \left(\frac{1+O(1/\ln n)}{2k}\right) \cdot E[\Phi(k)], \qquad (6)
\end{aligned}$$

where the last inequality uses the chosen value of ϵ and Lemma 8. Therefore, combining (4), (5) and (6) we get:

$$\begin{aligned}
E[\Phi(k+1) - \Phi(k)] &= E[D_1 + D_2] \\
&\leq \left(\frac{s}{n} + \frac{O(\ln n)}{n^2}\right) OPT \\
&\quad + \left(\frac{1+O(1/\ln n)}{2k}\right) \cdot E[\Phi(k)].
\end{aligned}$$

Rewriting this, we get that for $k \in [k_0, n-k_0]$:

$$\begin{aligned}
E[\Phi(k+1)] &\leq \frac{O(1)\ \ln^3(n)}{\min(k, n-k)} \cdot \mathrm{OPT} \\
&\quad + \left(1 + \frac{1+O(1)/\ln(n))}{2k}\right) \cdot E[\Phi(k)].
\end{aligned}$$

Moreover, the initial value of the recurrence is

$$\Phi(k_0) \leq k_0 \mathrm{OPT}.$$

Replacing k_0 by its value, and doing repeated back substition from $n - k_0$ down to k_0 and using a coarse upper bound, yields

$$E[\Phi(n-k_0)] \leq O(\ln^4(n)) \cdot \mathrm{OPT} \cdot \prod_{k_0}^{n-k_0} \left(1 + \frac{1+O(1)/\ln(n))}{2k}\right).$$

$$E[\Phi(n-k_0)] \leq O(\ln^4(n)) \cdot \mathrm{OPT} \cdot\ e^{(1+O(1)/\ln(n))\ln(n)/2}.$$

$$E[\Phi(n-k_0)] \leq \mathrm{OPT} \cdot\ O(\ln^4(n))\sqrt{n}.$$

Finally, $\Phi(n) - \Phi(n-k_0) \leq k_0 \mathrm{OPT}$. Recall that the value of the potential function is an upper bound on the cost of the solution. This gives the claimed bound of $O(\text{polylog}(n)\sqrt{n} \cdot \mathrm{OPT})$.

6. DISCUSSION

The main problem that remains open is analyzing the model where we are allowed to mix arrivals and replays. We made some progress towards understanding this model in the special case of random arrivals. We conjecture that the upper bound remains polylogarithmic even in the adversarial model, but leave it as a tantalizing (and difficult) open problem.

Another interesting direction is the multi-source case where each request (player) is a pair of terminals that need to be connected. We assume egalitarian cost sharing between the players. Unfortunately, the following is an easy example that shows that in this case the greedy algorithm (for phase 1) has competitive ratio $\Omega(n)$. Take the complete graph in which all edges e have distinct costs c_e which are approximately 1. The terminals pairs consist of all possible pairs of vertices, and the requests $\{i, j\}$ appear in order of decreasing cost of $c_{\{i,j\}}$. Then, upon arrival of request $\{i, j\}$, even if all other more expensive edges already have a user, buying edge $\{i, j\}$ to serve the request is cheaper than using a path $\{i, k\}, \{k, j\}$. Thus, every request will be served by a new edge and the total cost of the solution is approximately $n(n-1)/2$, whereas the optimal solution has cost about $n - 1$, giving a competitive ratio of at least $n/2$.

7. REFERENCES

[1] A. Agarwal and M. Charikar. On the price of stability for undirected multicast, manuscript, 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. *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science*, 2004, pp. 295-304.

[3] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Noncooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, Vol. 25 (2007), pp. 1193-1206.

[4] H.-L. Chen and T. Roughgarden. Network design with weighted players. *Proceedings of the 18th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, 2006, pp. 29–38.

[5] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, 2006, pp. 608-618.

[6] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, Vol. 21 (1997), pp. 85–101.

[7] M. Imase and B. Waxman. Dynamic steiner tree problem. *SIAM J. on Discrete Mathematics*, Vol. 4 (1991), 369–384.

[8] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *16th Annual Symposium on Theoretical Aspects of Computer Science*, 1999, pp. 404–413.

[9] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, Vol. 13 (1996), 111–124.

[10] V. Mirrokni and A. Vetta. Convergence issues in competitive games. *7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2004, pp. 183–194.

[11] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, Vol. 14 (1996), pp. 124–143.

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

[13] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, Vol. 49 (2002), pp. 236–259.

[14] L. S. Shapley. A value for n-person games. Contribution to the Theory of Games, pp. 31-40, Princeton University Press, Princeton, 1953.

[15] M. Voorneveld, P. Borm, F. van Megen, S. Tijs and G. Faccini. Congestion games and potentials reconsidered. *International Game Theory Review*, Vol. 1 (1999), pp. 283–299.