On a Locally Minimum Cost Forwarding Game

Stephan Eidenbenz
Los Alamos National Laboratories
eidenben@lanl.gov

Gunes Ercal-Ozkaya
University of Kansas
gunes@ku.edu

Adam Meyerson
University of California at Los Angeles
awm@cs.ucla.edu

Allon Percus
Claremont Graduate University
Allon.Percus@cgu.edu

ABSTRACT

We consider the problem of all-to-one (reverse multicast) selfish routing in the absence of a payment scheme in wireless networks, where a natural model for cost is the power required to forward. Whereas each node requires a path to the destination, it does not care how long that path is, so long as its own individual or local forwarding cost is minimized. Thus, we refer to this setting as a Locally Minimum Cost Forwarding Game (LMCF). From a system-wide perspective, short paths are clearly desirable, yielding two related social objectives of finding topologies that minimize: (i) the maximum stretch factor, and (ii) the directed weighted diameter. We prove that Nash equilibria always exist for LMCF, in particular the directed MST always being one, and we analyze the ratio of the social cost of Nash equilibria to the global optimum. The worst (maximum) possible value of this ratio is called the *price of anarchy (PoA)*, and the best (minimum) possible value is called the *price of stability (PoS)*. For the maximum stretch factor we present a $\Omega(n)$ worst-case bound on PoA and PoS, and for the directed weighted diameter we present a $\omega(n^c)$ worst-case bound on PoA and PoS for all $c < 1$, even when restricted to Euclidean instances. We prove hardness of computing the optimal Nash equilibrium in three-dimensional Euclidean instances as well as approximation hardness in arbitrary instances. Finally, we propose a heuristic for finding Nash equilibria and analyze, via simulations and probabilistic arguments, the social costs given by the heuristic and by the MST. These results suggest that for random Euclidean power instances, the *expected PoA* is $\omega(1)$ while the *expected PoS* is $\Theta(1)$.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network Topology; G.m [**Mathematics of Computing**]: Miscellaneous

FOWANC'09, May 18, 2009, New Orleans, Louisiana, USA.

General Terms

Algorithms, Economics

Keywords

game-theory, price of stability, random Euclidean power graphs, reverse multicast

1. INTRODUCTION

Incentive-compatible topology control protocols play a central role in selfish wireless networking. These protocols determine which links are to be used for forwarding data packets from sources to destinations. Non-selfish topology control involves selecting a subset of possible communication edges such that the resulting induced subgraph satisfies a number of desirable properties, such as being a single component, having small maximum degree, and preserving shortest paths within a small factor (see [13, 16] for an overview of topology control results). Incentive-compatible topology control has only recently received attention from the research community, with the two traditional game-theoretic approaches being to characterize Nash equilibria [6, 5, 12] and to design VCG-based mechanisms for ad-hoc networks [14, 9, 6, 5, 1, 2].

We study selfish topology control where all participating nodes need to have a path to a single destination. This might be the access point (AP) that allows the nodes to connect to the Internet, or a base station in a hybrid cell phone network, or a central processing node in a sensor network. Our cost and utility model is as follows.

Individual nodes care only about minimizing their own power consumption, and adopt their strategies accordingly, whereas the global objective is to minimize the total energy used. The social goal is to minimize the expected amount of power consumed by transmitting a single message from source to destination. Assuming fixed routing tables (pure strategy) this will also be the cost-per-message in the long term. The goal of individual agents is to minimize their power consumed given that they must forward what has been given them in this reverse multi-cast scenario. The "social" goal is to minimize the total power consumed. We assume that messages are generated at random sources. For simplicity we will assume uniformly at random, but all our results will hold for *any* probability distribution which guarantees that each node has a non-zero probability of being a message source. Specifically, define:

- q_v = load generated at node v

- Q_v = total outgoing load at v from all data passing through v

- c_v = cost per-bit of forwarding from v to v's chosen neighbor

- $PATH_v$ = sum of c_i along the path from v to the destination

Player v wants to minimize $Q_v c_v$, but has no control over Q_v and can only minimize c_v. The global objective is to minimize the total induced load times the forwarding cost, $\sum_v Q_v c_v = \sum_v q_v PATH_v$. Given an aggregate load $\sum_v q_v$, the worst-case distribution would be concentrated along the costliest path, with exactly one nonzero q_v. The global objective is therefore to keep that longest directed path short.

Since a node's interest is limited to having a path to the destination, and it does not care how long that path is, so long as its *individual* or *local* forwarding cost is minimized, we refer to this setting as a *Locally Minimum Cost Forwarding Game* (LMCF). Of course, from a system-wide perspective, short paths are desirable, as explained above so that our *social* optimum objectives as the following: (i) minimize the maximum stretch factor in the resulting topology with respect to true shortest path distances, or (ii) minimize the cost of the longest path in the resulting topology. Our aim is to find and to characterize Nash equilibria that optimize the social objective. The ratio of the objective value achieved by the best Nash equilibrium and the (non-selfish) value of the social optimum is called the *price of stability*, whereas that ratio for the worst Nash equilibrium is called the *price of anarchy*[11, 15, 3].

The prices of stability and anarchy have been extensively investigated in other network settings under other objectives, particularly for congestion-based games and fair-allocation games (for example [3, 11, 15, 8]). The problem of finding good Nash equilibria in the context of topology-control for ad-hoc networks has also been investigated [14, 9, 1, 2, 6, 5]. We note that the LMCF game differs from previously considered games in its objectives, both individual and social. Our work is related to the non-game-theoretic results of [10] that construct spanning structures balancing edge costs (MSTs) and path costs (SPTs). However the game-theoretic aspect of this work, in particular the locality of individual preferences, makes a crucial difference as the algorithms of [10] do not directly relate to Nash equilibria for LMCF.

We prove that Nash equilibria always exist for LMCF, and that in fact a minimum spanning tree is always a Nash, albeit rarely a socially optimal one. We give examples showing that both the prices of anarchy and stability can be linear with respect to the stretch-based social cost objective and $\omega(n^c)$ for any $c < 1$ for the maximum-distance based cost function. We show NP-hardness and inapproximability results for the problem of finding the socially optimal Nash equilibrium. We observe that there is hope for positive *average* results in various random graph models, such as Euclidean power cost functions which are a common model for communication costs in ad-hoc networks (Sect. 3.2), in view of previous work indicating that many nodes are involved in mutual nearest neighbor pairs in the relevant random models [17, 18]. We propose a greedy heuristic that we test in simulation, and find that the quality of the Nash equilibria found appear independent of the instance size (this is not true for a straight-forward MST heuristic). Our experiments suggest a plausible $\omega(1)$ average price of anarchy and $\Theta(1)$ average price of stability, and supports the use of our heuristic as a topology-control protocol for selfish all-to-one routing in ad-hoc networks.

2. MODEL AND PRELIMINARIES

Definition 2.1 (LMCF Game). *Given a connected, undirected, edge-weighted graph $G = (V, E)$ (with $V = (1, 2, \ldots, n)$), weight function $w : E \to R$ and designated destination node t, $LMCF(G, w, t)$ consists of the following: Players are nodes $v \in V \setminus \{t\}$, each player v with strategy set $N(v) = \{$ one-hop neighbors of $v\}$. Given a pure strategy-tuple[1] $S = (s_1, s_2, \cdots, s_{n-1})$ refer to G_S, the graph induced by S, as the directed graph formed by the set of directed edges of the form (u, s_u). Finally, the cost $c_S(v)$ of strategy-tuple S to player v is $c_S(v) = w(v, s_v)$ if G_S contains a path from v to t and ∞ otherwise.*

For any node v and any strategy-tuple S, at most one path may exist from v to t in G_S. Denote by $dist_S(v)$ the total weight of that path if such exists and ∞ otherwise. Clearly, this distance is minimum in a shortest path tree (SPT) rooted at t. Denote the shortest path distance simply by $dist(v)$. Now, we present two alternative formulations for the *Social Cost* of a strategy-tuple S for the LMCF Game. The first is based on the stretch factor of node-destination paths in G_S, the second based directly on the maximum distance of any node to the destination in G_S.

$$SC_{stretch}(S) = max_{v \in V \setminus \{t\}} \frac{dist_S(v)}{dist(v)} \quad (1)$$

$$SC_{md}(S) = max_{v \in V \setminus \{t\}} dist_S(v) \quad (2)$$

We investigate the *price of anarchy* and *price of stability*, as well as their computability, for *Nash equilibria* of the Locally Minimum Cost Forwarding (LMCF) Game. A **Nash equilibrium** is a fixed-point best-response strategy profile: a strategy-tuple from which no agent has a unilateral incentive to deviate, i.e., any such deviation would not improve the cost to the agent. The **price of anarchy**, PoA, with respect to a social cost function on a given instance is the maximum (worst) ratio of the social cost of a Nash equilibrium to the best possible social cost (that for the SPT). The **price of stability**, PoS, is the minimum (best) such ratio. We study these quantities in both the worst-case and average-case. Since by definition $SC_{stretch}(S) = 1$ for the SPT, PoA and PoS over $SC_{stretch}$ are precisely the maximum and minimum social cost over all Nash equilibria.

We will focus on the prices of anarchy and stability on *Euclidean power graphs* and *random link* graphs. A Euclidean p-power graph in dimension d is a complete graph consisting of nodes embedded into d-dimensional Euclidean space with edge weights defined by $w(i, j) = d^p(i, j)$, the pth power of the distance.[2] *Random* Euclidean power graphs are in-

[1]We restrict ourselves to pure strategies when discussing strategies and later note why this restriction is *w.l.o.g.* for the questions considered.

[2]When the dimension is not specified, we may assume it is 2. When the power is not specified then assume that it is

duced by placing each node uniformly at random into the d-dimensional unit-cube. These are especially relevant models of ad-hoc networks due to the randomness of placement and the modeling of energy. *Random link graphs* are unstructured, formed simply by assigning i.i.d. weights from some distribution to the edges of the complete graph. Here, for simplicity, we will take weights to be distributed uniformly at random over some bounded interval.

3. RESULTS

3.1 Examples and Lower Bounds

We first present examples of Nash equilibria that provide some intuition on the nature of the problem, as well as lower bounds on PoA and PoS.

EXAMPLE 3.1 (MST). *Given a graph G and destination t, construct a minimum spanning tree T of G and direct its edges towards t. Note that this forms a Nash equilibrium: If a node u has an incentive to switch from its current forwarding choice v to a new node v', forming T', then doing so does not introduce a cycle and $w(u, v') < w(u, v)$. But then T' is also a spanning tree, with total cost less than that of T, contradicting that T is a MST.*

Now consider the MST for the Euclidean "Horseshoe" graph G_H of Figure 1 given in [10]. This example immediately gives a $\Omega(n)$ lower bound on PoA with respect to $SC_{stretch}$, since both a clock-wise and a counter-clockwise path to t are Nash equilibria. Further, it can inductively be checked that the best Nash equilibria for this case with respect to both $SC_{stretch}$ and SC_{md} is that of Figure 2, thus also giving a $\Omega(n)$ lower bound on PoS for $SC_{stretch}$. We note the contrast with [10]'s approximate solution for balancing MST cost and SPT cost which yields a constant bound for G_H (by actually connecting the dots as a horseshoe) but is not a Nash equilibrium. Thus, we have:

EXAMPLE 3.2 (HORSESHOE). *The Euclidean "Horseshoe" graph of Figure 1 given in [10] yields linear lower bounds on both PoA and PoS under $SC_{stretch}$, the optimal Nash being the counter-intuitive one of Figure 2 (unlike [10]'s constant stretch approximation for a non-game-theoretic scenario).*

Now consider a Euclidean Spiral graph G_{spiral} with t at center such as the nodes of Figure 3. One may imagine forming this graph just as one draws the spiral from inside out, where each new node gives an edge that was a little bit longer than the one before as well as a little bit shorter than the new node's distance to the closest neighbor in an inner layer of the spiral. It can be checked that the unique Nash in such a class of instances is that of directing the Spiral inward towards t as shown in the Figure. Moreover, precise parameters may be set such that the number of spiral layers is proportional to $\Omega(n^c)$ for any constant $0 < c < 1$, leading to the following (see [7] for details):

EXAMPLE 3.3 (SPIRAL). *The Euclidean "Spiral" graph of Figure 3 yields a $\omega(n^c)$ (for any constant $c < 1$) lower bound on PoA and PoS under SC_{md}.*

1. Note that for powers higher than 1 these graphs do not necessarily obey a metric though they are induced by such.

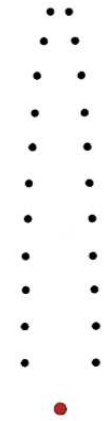

Figure 1: Euclidean Horseshoe

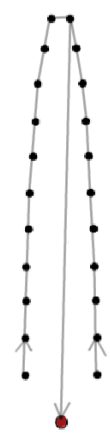

Figure 2: Optimal Nash for Horseshoe

3.2 Observations on the Structure of Nash Equilibria

It is not a coincidence that the Nash equilibria examples thus far have been trees. We briefly return to consideration of general Nash, including the mixed case:

REMARK 3.4. *In any connected graph G, Nash equilibria always exist (guaranteed via MST), and all Nash equilibria form an acyclic spanning graph with destination t as sink. In particular every pure Nash equilibria forms a spanning tree directed towards t.*

Note that mixed Nash can be viewed as flows in G and that any non-zero flow through a cycle will have infinite cost. Now, given any Nash-induced sub-graph G_S, denote as T_t the maximal acyclic flow in G_S including t. Due to connectedness, if there are any cycles in G_S, then there is

Figure 3: Euclidean Spiral

always some node on some cycle that has a neighbor in T_t, thus giving finite rather than infinite flow weighted cost (an infinite flow weighted cost is still infinite), and an incentive to switch. So, there can be no cycles and T_t spans G. And, due to the single out-degree nature of pure Nash, if G_S is a pure Nash then T_t is a spanning tree.

It is without loss of generality to consider pure Nash equilibria with respect to maxima and minima of our social cost functions, for the following reason. For any branching out (i.e. a node forwarding in a mixed manner to more than one neighbor) that occurs in a mixed Nash, the multiple options must have identical cost to the branching node (and given that, the branching node may distribute the probability flow in any manner). Any mixed flow achieving some relevant minima or maxima (with respect to PoS or PoA) can therefore be converted into a directed tree achieving the same such minima or maxima by continuously shifting the flow to a path that induces the extreme value. Thus, from now on we discuss Nash trees without loss of generality.

We may further state the following regarding the structure of Nash trees:

REMARK 3.5. *In a weighted graph G, if i is j's unique nearest neighbor and j is also i's unique nearest neighbor, then we refer to i and j as* mutual nearest neighbors. *Edges between mutual nearest neighbors (excluding t) are always used in some direction, in any Nash tree.*

To see this, consider a mutual nearest neighbor pair i, j and Nash tree T such that edge (i, j) is *not* used in either direction. Since the weight of this edge is minimal amongst all neighbors for both i and j, the only way it cannot be present in a Nash equilibrium is if it lies on a cycle. But if directing from i to j would create a cycle, then there must already be a path from j to i in T, and likewise for the opposite direction. So there must already be a cycle in T, namely from i to j and back to i, contradicting that T is a Nash tree. Thus, (i, j) must be used in some direction in every Nash tree.

As a corollary, we may also relate this to *generating* Nash equilibria. Due to the uniqueness condition in the above definition, any set of mutual nearest neighbor edges must be an independent set. Moreover, noting that in a complete graph we may always complete a spanning tree after fixing any independent set of edges as a subgraph, we have the following:

COROLLARY 3.6. *In any complete graph, for every directionality of the set of mutual nearest neighbor edges (excluding t) there exists a corresponding Nash equilibrium.*

Euclidean graphs are especially relevant cases for analysis of the LMCF Game. While we have already noted that a restriction to Euclidean graphs is rich enough to generate arbitrarily bad examples, this class also has some further structural properties:

REMARK 3.7. *For any Nash tree in any 2-dimensional Euclidean power graph of any power, the incoming node degree is at most 6.*

The reason is as follows: Consider a set of seven nodes incoming to a vertex v in some Nash tree T. By a regular hexagonal decomposition into 6 parts, it may be seen that at least one of these incoming neighbors u must be strictly closer to another incoming neighbor w than to v. Moreover, if by switching u's forwarding choice from v to w a cycle was created in the graph, then w and hence w's own forwarding choice v must have already had a path to u in T. But since u forwards to v in T, there must already be a cycle through u in T, contradicting that T is a Nash tree.

Of course, the smallest edge in any graph must consist of a mutual nearest neighbor pair. Exactly how many might we expect? To address this question, we may say something in the case of random instances based on results of [17, 18] on random Euclidean graphs of any dimension and on results of [18] on random link graphs.

REMARK 3.8. *For random Euclidean power instances of any dimension and any power*[3]*, and for random link graphs, at least half the nodes are expected to be involved in some mutual nearest neighbor relation.*

Recall that the "bad" Euclidean examples, Figure 1 and Figure 3, each had at most $O(1)$ nodes involved in a mutual nearest neighbor and a sparse set of possible Nash equilibria. As such situations are highly unlikely, it is reasonable to hope that greater optimism is warranted for random instances. We discuss this further in Sect. 3.4 and beyond.

3.3 Hardness of Optimal Nash Equilibria

We now provide hardness results for computing and approximating optimal Nash equilibria.

THEOREM 3.9. *The optimal Nash equilibrium for the LMCF Game is NP-Hard to approximate to any constant factor for both the SC_{md} and $SC_{stretch}$ social cost functions.*

PROOF. We start by showing that finding the optimal Nash equilibrium for SC_{md} is NP-hard. Following arguments from [10], a related construction then shows NP-hardness for $SC_{stretch}$ as well as hardness of approximation for both social cost functions.

The proof of Theorem 3.9 is based on the 3-SAT reduction of [10] for the minimal-stretch MST problem, modified with appropriate "choice" gadgets between positive and negative literals of the same variable. For the purpose of showing NP-Hardness, the constructed graph G is 3-SAT represented as the union of the clause-literal bipartite graph with edges of length B, along with additional paths E between the positive and negative literals of each variable, as well as a destination t connected to every literal by edges of length $A \leq B$. The choice gadget for E is simply a symmetric path of edges with small (meaning even the heaviest edge has small cost) decreasing cost then increasing cost. This replaces the edges of path E in [10]'s construction. Note that for every positive and negative variable nodes, say x and $\bar{x}$, the choice gadget enforces that every Nash tree has either a path from x to $\bar{x}$ or vice versa. Since each literal is directly connected to t by an edge shorter than the edge to a clause, the literal with the incoming path from its corresponding choice gadget must necessarily then forward to t in every Nash tree as well. Moreover, each clause must choose one of its corresponding literals to forward to. Therefore, there are only two possible kinds of paths from a clause to the destination depending on the choice gadget's direction: Zig-zagging $B \to E \to A$ or bypassing E via $B \to A$. For every pair of literals, directing E from $\bar{x}$ to x if the corresponding 3-SAT variable assignment is true and from x to $\bar{x}$ if it is false, we see that

[3]Power does not change nearest neighbor relations.

a clause node that directs into its chosen literal does not zig-zag. Since every Nash equilibrium for G uniquely specifies the direction of all paths E between literals and vice versa, the reduction is clear: The 3-SAT instance is satisfiable iff there exists a choice of all path directions for E such that no clause zig-zags. For sufficiently long B, zig-zagging is the only way to increase the directed diameter, and so NP-hardness for optimal SC_{md} follows.

For demonstrating NP-hardness for $SC_{stretch}$ as well as hardness of approximation, the construction is augmented by an additional node R connected to t by a path of length D and to the clause nodes via edges of length W. By making W sufficiently long and D sufficiently short, we can ensure that zig-zagging is the only way to increase the maximum stretch for any Nash equilibrium, from which NP-hardness for optimal $SC_{stretch}$ follows. Finally, following the identical method as [10], we can set edge and path length to induce an arbitrary constant approximation gap, from which NP-hardness of approximation follows for both social cost functions. We refer the reader to [10] for details. □

Note that the theorem above also holds under the restriction to complete graphs, by ensuring that any edges added to this construction are large enough not to be used in any Nash equilibrium.

More specifically, we show that it is hard to compute optimal Nash equilibria even when we restrict ourselves to the 3-dimensional Euclidean case.

THEOREM 3.10. *The social cost of the optimal Nash equilibrium for the LMCF Game on 3-dimensional Euclidean graphs is NP-Hard to compute for both the SC_{md} and $SC_{stretch}$ social cost functions.*

PROOF. Here, it suffices to modify the 3-SAT construction G in the previous proof so that it can be embedded into 3-dimensional Euclidean space. Place the destination node t at the origin. For each variable x_i connect its positive and negative literal via a choice gadget E as in the previous proof, and place each such connected pair equidistant from neighboring pairs on a sufficiently large circle about t. We introduce a new gadget here as well, which we call the "directionality" gadget: a sequence of nodes $< v_0, v_1, \ldots, v_q >$ such that for each $0 \leq i < q$ the nearest neighbor of v_i is uniquely v_{i+1}. Namely, the distances of consecutive points is strictly decreasing and chosen small enough to guarantee that there are no closer points elsewhere in the remaining construction. Now, for each literal, draw a directionality gadget from the literal to t along the line connecting those two points, identical for every pair. These replace the A-edges in the previous proof, so let us refer to these as A as well. Note that the choice gadget E is simply two directionality gadgets mirroring each other about the central shortest edge. Whereas in a Nash equilibrium E has two choices of direction, A has only one choice of direction: all v_i except possibly v_0 direct towards v_{i+1}. We retain the caveat that the longest edge of E is still shorter than its neighboring edges outside of E. Now, for the placement of clauses, and the paths connecting clauses to literals: Place clause nodes on a sufficiently large sphere about the origin, sufficiently far apart from each other. For each clause node c_i place three idential directionality gadgets connecting c_i to its corresponding literals, replacing the edges B in the previous construction. Note that this can certainly be accomplished in 3-dimensional space without any two paths B coming too close by choosing the sufficiently large sphere. Moreover, these gadgets B can be made identical in total length as well by elongating short lines by a curve. This completes the specifications for the construction of the embedding: let us call it G_{3D}. What remains is this. Every Nash equilibrium for the LMCF game on G_{3D} is uniquely determined by the choice of literal to which c_i connects via a B-gadget and the choice of direction for each E gadget. Again note: The maximum possible stretch factor and weighted-hop-distance to t in a Nash equilibrium are achieved by a zig-zagging of $B \rightarrow E \rightarrow A$. Identically to the previous proof, such zig-zagging is only necessary for unsatisfiable 3-SAT instances. The NP-Hardness then follows. □

3.4 Heuristics

Given the hardness results, we look to find an intuitive heuristic to compute Nash trees for LMCF with low Social Cost. First, we present a meta-heuristic, **LMHeur**, to compute general Nash equilibria for the LMCF Game. The main idea behind the meta-heuristic is that since equilibria are directed trees, for any Nash tree there exists a forwarding order such that, maintaining a forest of directed edges for nodes already chosen, the next chosen node forwards to its nearest neighbor that does not introduce a cycle into the forest. The choice of forwarding order may be dictated by whichever global cost function we wish to optimize, or which kind of Nash we are looking for.

As we have proposed $SC_{stretch}$ and SC_{md} as reasonable social cost functions to consider for the LMCF Game, we propose the **DeltaHeur**, a member of the LMHeur class, to compute good Nash trees. The ordering priority for DeltaHeur is based on maximal progress towards shortest path. A priority queue is kept holding the nodes that have yet to forward, sorted by the difference between the candidate node's shortest path distance to the destination and its available nearest neighbor's shortest path distance to the destination ("available" means not introducing a cycle within the current forest). It is straightforward to show in the Euclidean case that the ordering induced by the DeltaHeur corresponds exactly to a "maximum projection" heuristic, which we call **ProjHeur**, where the projection in question is that of the vector from the candidate node to its available nearest neighbor projected onto the vector from the candidate node to the destination t.

3.5 Expected Social Costs on Random Graphs

While we have given *worst-case* lower bounds on PoA and PoS, it is also of practical interest to consider the *expected* values of these quantities for classes of random graphs. Here we discuss the case of random Euclidean instances, arguing that based on the structure of the MST, PoA is likely to be $\omega(1)$, whereas based on the behavior of DeltaHeur, PoS may be $\Theta(1)$.

The argument for PoA (at least under $SC_{stretch}$) is as follows. Consider the MST on nodes placed uniformly at random in a d-dimensional unit hypercube. Now imagine a $(d-1)$-dimensional hyperplane that cuts one of the edges incident on t in the spanning tree, but none of the others. Because of the uniform density of nodes, we expect that this (curved) hyperplane will pass between many pairs of nodes separated by distance $\Theta(n^{-1/d})$, i.e., the typical distance between a node and its close neighbors. The hyperplane continues to the boundaries of the hypercube, and so some

of these pairs of nodes are likely to be at constant distance from t. However, since the nodes in the pair are on opposite sides of the hyperplane, the shortest MST path between them must pass through the cut edge, and thus be of constant length. While we are concerned with the maximum stretch factor of a node to t rather than between two arbitrary nodes on the graph, note that in the random case, t is in fact equally likely to be *any* of the nodes in the MST. This suggests the maximum stretch factor could be $\Theta(n^{1/d})$, so PoA would then be $\Omega(n^{1/d})$ under $SC_{stretch}$. The argument is qualitatively similar for Euclidean power graphs.

Now consider the behavior of DeltaHeur in constructing Nash equilibria on the same random Euclidean instances. Not all directed edges in the Nash tree will point towards t, but we expect there to be a distinct positive bias in favor of this: it is easy to show, for instance, that the algorithm orients all mutual nearest neighbors in the direction of t. Given that mutual nearest neighbors pairs are a constant fraction of edges and their angles of orientation are uniform at random, at each step the expected progress towards destination (the criterion on which DeltaHeur ranks edge) is likely to be a constant fraction of the edge length. The typical distance on the Nash tree from a node to t is then a constant, giving an *average* (not maximum) social cost of $\Theta(1)$. Moreover, if it turns out that the random variables representing the progress towards destination along a path to t are only weakly correlated, the distribution of these distances will be close to a Gaussian with constant mean and variance $\Theta(n^{-1/d})$. The maximum of n such Gaussian variables has mean $\Theta(\alpha+\sqrt{\log n}\, n^{-1/d})$ [4] where α is a constant independent of d, suggesting that this too might be the expected *maximum* distance. A very similar argument holds for maximum stretch factor, leading to the possible scenario that PoS is $\Theta(1)$.

3.6 Experimental Analysis

We have performed a number of simulations in order to test these predictions, and more generally to assess the quality of the Delta Heuristic in comparison to directed MSTs for LMCF. We have run these experiments for 2-dimensional Euclidean power graphs of powers $1, 2, 3$, and random link graphs. The resulting plots of *prices* (ratios of the social cost for the computed Nash to the optimal SPT cost) are shown in Figures 4 through 11. Each Euclidean power graph of size n was formed by picking n integer-coordinate points uniformly at random from a 560x560 grid, with edge weights induced by the given power of the distance between nodes. Each random link graph of size n was formed by assigning edge weights uniformly at random from the interval $(0, 10001)$ independently to each link of the complete graph. For each random graph instance with size n, ranging from 30 nodes to 350 nodes, we ran $\sqrt{n}$ experiments computing the prices of the DeltaHeur Nash solutions and the MST Nash solutions and plotted the average price obtained for each node. The results are summarized in Figures 4, 5, 6, 7, 8, 9, 10, and 11. In all figures note that light blue plot refers to the directed MST while the dark red plot refers to the DeltaHeur. If there is a number i in the figure caption, it indicates that the random graph model considered is a 2-dimensional Euclidean power graph with power i. Otherwise if capital R is in the figure caption, the random link model is presented. Of course, "stretch" in the caption refers to $SC_{stretch}$ price, whereas MD in the caption refers to SC_{md} price.

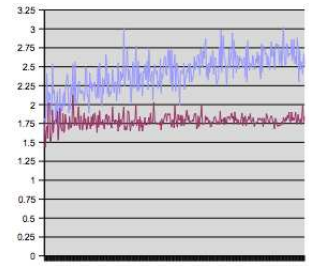

Figure 4: MD1

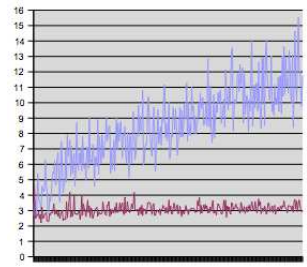

Figure 5: Stretch1

In all experiments, the concentration of the blue plot can be seen above the concentration of the red plot, though at varying angles and to varying degrees of concentration, thus confirming the intuition that DeltaHeur yields heuristic improvement. This improvement is perhaps most pronounced in the standard Euclidean case under the stretch-based social cost function, as seen in Figure 5, which is especially of interest as the DeltaHeur is identically **ProjHeur** in that case (and was, in fact originally conceived from such). We can see that whereas the stretch-based price of DeltaHeur appears highly concentrated about a small constant, 3, that of the directed MST is increasing asymptotically. For the higher powers considered, aside from sparse but extreme outliers for DeltaHeur (likely due to the increase in variance with power and limited sample size), the same observation still holds except that the rate of increase of the directed MST plot slows with power, and the small constant about which the DeltaHeur price is concentrated also seems to decrease with power. That both prices appear to decrease on average with power is intuitive: As the power of the distance grows, SPT preference is given to smaller edges. Both expected DeltaHeur and MST prices stay near small constants for the maximum distance based (md-based) social cost function. Even here, MST prices appear growing slowly whereas DeltaHeur prices demonstrate stability at constants close to 1. For the random link models under both social cost functions, both DeltaHeur and MST prices appear growing, though for the stretch-based pricing, again MST appears growing asymptotically faster, and there is too much variance in the DeltaHeur plot to make further inferences. Here also, DeltaHeur is still an improvement over MST (whether by a constant or asymptotically growing factor), and the md-prices are still quite small.

Finally, it is interesting to note that these results for DeltaHeur may be explained by our Euclidean predictions together with an interpretation of the random link model as the limit of a Euclidean model in high dimensions. If the expected DeltaHeur social cost is $\Theta(\alpha + \sqrt{\log n}\, n^{-1/d})$,

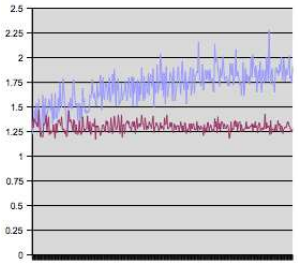

Figure 6: MD2

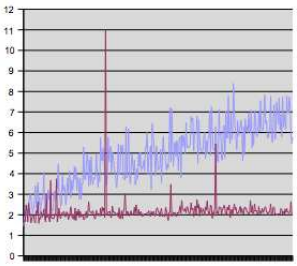

Figure 7: Stretch2

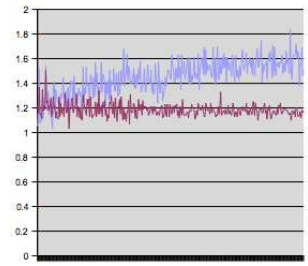

Figure 8: MD3

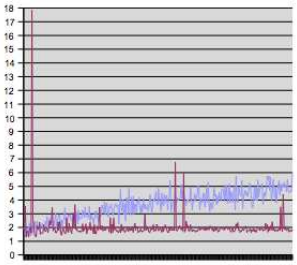

Figure 9: Stretch3

then unlike for finite d where this is constant, in the large d limit it grows logarithmically. This is consistent with our simulations.

4. CONCLUSION

The primary observation throughout the plots is as follows: For both random link graphs and Euclidean power graphs, DeltaHeur is consistently cheaper than the MST for LMCF under both SC_{md} and $SC_{stretch}$. In particular: With respect to both social cost functions considered, while the expected MST price **grows asymptotically**, the expected price of DeltaHeur stays **concentrated at small constants** throughout the 2-dimensional Euclidean power instances for powers $1, 2, 3$, validating the usefulness of the DeltaHeur for all-to-one topology-control in the presence of selfish node behavior.

5. ACKNOWLEDGEMENTS

We would like to thank Bhaskar Krishnamachari, Hua Liu, and Ho-Lin Chen for useful discussions and the latter's helpful observations on hardness of the Euclidean case.

6. REFERENCES

[1] Ambühl, C., Clementi, A. E. F., Penna, P., Rossi, G., and Silvestri, R. On the approximability of the range assignment problem on radio networks in presence of selfish agents. *Theor. Comput. Sci. 343*, 1-2 (2005), 27–41.

[2] Anderegg, L., and Eidenbenz, S. Ad hoc-VCG: a truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In *MOBICOM 03: Mobile Computing and Networking* (2003), ACM, pp. 245–259.

[3] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., and Roughgarden, T. The price of stability for network design with fair cost allocation. In *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04)* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 295–304.

[4] Berman, S. A law of large numbers for the maximum in a stationary gaussian sequence. *Ann. Math. Stat. 33* (1962), 93–97.

[5] Crescenzi, P., Ianni, M. D., Lazzoni, A., Penna, P., Rossi, G., and Vocca, P. Equilibria for broadcast range assignment games in ad-hoc networks. In *ADHOC-NOW* (2005), pp. 4–17.

[6] Eidenbenz, S., Kumar, V. S. A., and Zust, S. Equilibria in topology control games for ad hoc networks. *Mob. Netw. Appl. 11*, 2 (2006), 143–159.

[7] Eppstein, D. Beta-skeletons have unbounded dilation. *Comput. Geom. Theory Appl. 23*, 1 (2002), 43–52.

[8] Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C. H., and Shenker, S. On a network creation game. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing* (New York, NY, USA, 2003), ACM Press, pp. 347–351.

[9] Kannan, R., Sarangi, S., and Iyengar, S. S. Sensor-centric energy-constrained reliable query routing for wireless sensor networks. *J. Parallel Distrib. Comput. 64*, 7 (2004), 839–852.

[10] Khuller, S., Raghavachari, B., and Young, N. E. Balancing minimum spanning trees and shortest-path trees. *Algorithmica 14*, 4 (1995), 305–321.

[11] Koutsoupias, E., and Papadimitriou, C. Worst-case equilibria? In *STACS* (1999).

[12] Laoutaris, N., Poplawski, L. J., Rajaraman, R., Sundaram, R., and Teng, S.-H. Bounded budget connection (bbc) games or how to make friends and influence people, on a budget. In *PODC '08: Proceedings of the twenty-seventh ACM symposium on*

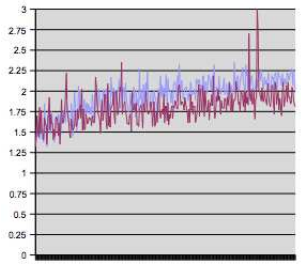

Figure 10: MDR

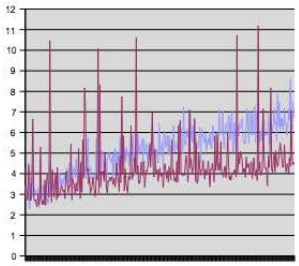

Figure 11: StretchR

Principles of distributed computing (New York, NY, USA, 2008), ACM, pp. 165–174.

[13] Li, X.-Y., Song, W.-Z., and Wang, Y. Efficient topology control for ad-hoc wireless networks with non-uniform transmission ranges. *Wirel. Netw. 11*, 3 (2005), 255–264.

[14] Liu, H., and Krishnamachari, B. A price-based reliable routing game in wireless networks. In *Proceedings of the First Workshop on Game Theory for Networks, GAMENETS-06* (2006).

[15] Roughgarden, T. *Selfish Routing and the Price of Anarchy*. MIT Press, May 2005.

[16] Santi, P. *Topology Control in Wireless Ad Hoc and Sensor Networks*. John Wiley and Sons Inc, 2005.

[17] Schwarz, G., and Tversky, A. On the reciprocity of proximity relations. *Mathematical Psychology 22* (1980), 157–175.

[18] Terariol, C. A. S., de Moura Kiipper, F., and Martinez, A. S. An analytical calculation of neighbourhood order probabilities for high dimensional poissonian processes and mean field models. *J. Phys. A: Math. Theor 40* (2007), 1981–1989.