Equilibria in Topology Control Games for Ad Hoc Networks*

Stephan Eidenbenz
eidenben@lanl.gov

V.S. Anil Kumar
anil@lanl.gov

Sibylle Zust
zust@lanl.gov

Los Alamos National Laboratory
Basic and Applied Simulation Science (CCS-5)
MS M997, P.O. Box 1663
Los Alamos, NM 87544, USA

ABSTRACT

We study topology control problems in ad hoc networks, where network nodes get to choose their power levels in order to ensure desired connectivity properties. Unlike most other work on this topic, we assume that the network nodes are owned by different entities, whose only goal is to maximize their own utility that they get out of the network without considering the overall performance of the network. Game theory is the appropriate tool to study such selfish nodes: we define several topology control games in which the nodes need to choose power levels in order to connect to other nodes in the network to reach their communication partners while at the same time minimizing their costs. We study Nash equilibria and show that – among the games we define – these can only be guaranteed to exist if all network nodes are required to be connected to all other nodes (we call this the STRONG CONNECTIVITY GAME). We give asymptotically tight bounds for the worst case quality of a Nash equilibrium in the STRONG CONNECTIVITY GAME and we improve these bounds for randomly distributed nodes. We then study the computational complexity of finding Nash equilibria and show that a polynomial-time algorithm finds Nash equilibria whose costs are at most a factor 2 off the minimum cost possible; for a variation called CONNECTIVITY GAME, where each node is only required to be connected (possibly via intermediate nodes) to a given set of nodes, we show that answering the question, if a Nash equilibrium exists, is NP-hard, if the network graph satisfies the triangle inequality. For a second game called REACHABILITY GAME, where each node tries to reach as many other nodes as possible, while minimizing its radius, we show that $1 + o(1)$-approximate Nash equilibria exist for randomly distributed nodes. Our work is a first step towards game-theoretic analyses of ad hoc networks.

*Los Alamos National Laboratory Report LA-UR-03-3269.

DialM–POMC'03, September 19, 2003, San Diego, California, USA.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network topology*

General Terms

Theory, Design, Algorithms.

Keywords

Ad Hoc Networks, Game Theory, Nash Equilibrium, Topology Control.

1. INTRODUCTION

Unlike traditional, fixed wireline networks, the next generation communication networks are likely to be ad hoc, or hybrid (i.e., a combination of ad hoc and wireline) networks. An ad hoc network consists of an arbitrary distribution of radios in some geographical location. One important feature of ad hoc networks is that nodes can move, and so the network changes dynamically. Earliest examples of ad hoc networks were in military applications (e.g. [8]). Recent advances in the commercialization of intelligent radio devices are likely to lead to the wide-spread emergence of ad hoc or hybrid networks [13].

Depending on its power level, on the nature of environmental interference, and on natural features, a node in an ad hoc network can reach all nodes in a certain range. To send a message to some far away node, the sending node must rely on intermediate nodes to forward the message. We consider scenarios where no fixed infrastructure is present and it is left to the nodes to choose their power levels to enable efficient communication. Once the radii (power levels) of the nodes are fixed, a digraph $G(V, E)$ can be abstracted out in the following manner: V is the set of nodes, and edge $e = (u, v)$ is present in E if node v is within the power range of node u; such a graph is called the *transmission graph* [14]. Efficient communication requires that the transmission graph satisfy certain properties such as connectivity, energy-efficiency and robustness. The area of *topology control* deals with choosing the radii such that the transmission graph has the desired properties; see [14] for a survey.

One of the key design criteria for any ad hoc network protocol is energy efficiency. In fact, the network nodes typically have limited battery power and energy consumption is the dominating cost component in an ad hoc network for an individual node [6]. The transmission range of a node u depends on the transmitting power P_u^{emit} of the node: the power $P_{u,v}^{rec}$ at which a node v at distance $d(u,v)$ to the transmitting node u receives the signal is [6]:

$$P_{u,v}^{rec} = \frac{K}{d(u,v)^\alpha} P_u^{emit}, \quad (1)$$

where K is a constant and α is the distance-power gradient varying between one and six depending on the environment conditions of the network. If this power exceeds a minimum level, a node v at this point can successfully receive the message and falls within the transmission range. Thus, energy requirements increase super-linearly in most cases (i.e., whenever $\alpha > 1$; under ideal conditions, we have $\alpha = 2$). Apart from high cost, large transmission radii also lead to higher interference during radio transmission, and thus to larger latencies.

Existing work on energy-efficient topology control [14] has focused on the problems of minimizing the sum of the radii (or the sum of some power of the radii) while ensuring that the transmission graph has the desired properties. Typically, such algorithms either require centralized control, or require that the nodes run a distributed algorithm while cooperating and trusting each other. While such an assumption on node behavior might hold for special networks, e.g. military or government applications, it is certainly unreasonable in commercial applications, which are strongly driven by economic incentives. More often than not, different nodes will be owned by different commercial entities, which would all like to communicate together, but – at the same time – individually want to incur as little cost as possible. Thus, in most scenarios, network nodes are selfish and each node's only goal is to maximize its own utility. This is a perfect scenario for studying as a non-cooperative game. The selfish nature of network nodes of course affects all layers of the protocol stack: a truly selfish node will try to exploit weaknesses of the protocols on any layer in order to improve its utility. We believe that the interactions of the network nodes, in particular on the data link and the network layer, should be studied as a game. On the data link layer, the protocol for fixing transmission radii and the protocol of allocating channel resources (on the MAC sublayer) could be exploited by selfish nodes; on the network layer, selfishness comes most obviously into play when a node is asked to forward data packets for another node, which only drains the battery of the forwarding node, thus bringing a negative utility to that node. A node can be made willing to forward packets by paying an appropriate amount of money; several schemes have been suggested that aim at solving the selfishness problem on the network layer [4, 1]. In this work, we will focus on the data link layer: we study different topology control problems as games by examining their equilibria, and by designing algorithms for reaching such equilibria. The ultimate goal of this line of research would be to combine the notion of selfishness such that it stretches across all protocol stack layers.

Computational Game Theory Game theory has been used as a tool to model and study different aspects of communication networks only in recent years. Transportation networks have been subject to game theoretic analyses (see e.g. [3, 5]) and policies on taxation and design of road networks have been influenced by game theoretic models [3]. Much of the classical game-theory work has been non-algorithmic in nature, without much focus on the computational complexity of finding good policies or designing good networks. The work by Roughgarden [15] represents recent attempts at addressing such algorithmic questions for traffic and wireline networks. Due to the intense interest in large networks, like the Internet, a lot of recent work in computational game theory focuses on *network design.* Roughgarden [16] considered the problem of designing networks that reduce the cost of selfish routing, and showed computational intractability of such problems. The work most closely related to the questions studied in this paper is [2, 7].

For modeling communication networks as games, it is reasonable to think of each node as a player or agent. Each player has a certain set of strategies: in the games we consider, a player needs to choose a radius, and a choice of the radius is a strategy. Each player is endowed with a local utility function. A lot of work in game theory has been devoted to stable operating points in non-cooperative games, and the most popular notion is that of a Nash equilibrium (see [12] for details). A choice of strategies σ for all players is said to be a Nash equilibrium, if no player has an incentive to deviate from σ in order to improve its utility. A Nash equilibrium can be pure or mixed: a mixed equilibrium is relevant if players randomize on their strategies. In this work we will generally not consider mixed strategies, as they do not seem to be practical in the context of such design problems (see also [2] for a similar argument). If the game has a Nash equilibrium, game theorists believe that such a game – played repeatedly – would tend to end up in a Nash equilibrium. Therefore, questions of existence of Nash equilibria and algorithms for finding them are of crucial importance.

Our Results We consider topology control problems in ad hoc networks, and model them as non-cooperative games. Ad hoc networks are characteristically different from other infrastructure networks, e.g. transportation systems, in many ways, and very little game theoretic analysis has been done so far for ad hoc networks. We consider two topology control games for static ad hoc networks, the CONNECTIVITY GAME and the REACHABILITY GAME, in this paper. The cost of a radius vector $\bar{r}$ is $C(\bar{r}) = \sum_v r_v^\alpha$ (see also Section 2).

In the CONNECTIVITY GAME, each node has to communicate with another node. Therefore, we need directed paths connecting each source node to the corresponding destination node. Each node aims to minimize its radius. The version of this game where each node needs to connect with everyone else is called the STRONG CONNECTIVITY GAME. In our model of selfishness, a node only cares about reaching its destination; a node will not make its radius larger to help other nodes reach their destinations. Once the radii are fixed, however, a node is willing to forward data packets even if these packets are from a third node (see [4] for a method of achieving this willingness).

The CONNECTIVITY GAME can be viewed as a wireless version of [2]. The work of [2] involves players on a network, with edges having costs. An edge can be used if it is paid for by the players, and strategies for the players involve choosing payments for the edges so that their connectivity requirements are met. In the ad hoc setting, a node can reach all nodes within its transmission range. Also, in an ad hoc network, a node only has control on its power, in contrast to [2], where a node can pay for far away edges.

In the REACHABILITY GAME, the utility function for a node v is defined as the difference between the number of nodes reached from v and r_v^α, where r_v is the radius chosen by v and α is a constant. Each node chooses a radius in order to maximize its utility.

Our results are summarized below.

1. The CONNECTIVITY GAME need not always have a pure Nash equilibrium, not even a β-approximate Nash equilibrium, for any $\beta > 0$. Deciding whether an instance of this game has a pure Nash equilibrium is NP-complete if the underlying graph satisfies the triangle inequality.

2. The STRONG CONNECTIVITY GAME always has a pure Nash equilibrium. In fact, there are multiple Nash equilibria, whose cost can vary widely. Also, any local optimum is a Nash equilibrium.

3. There is a simple local improvement algorithm that yields a Nash equilibrium for the STRONG CONNECTIVITY GAME. Using an observation from [10], this yields an algorithm to find a Nash equilibrium of cost at most twice the optimum.

4. The cost of any Nash equilibrium for the STRONG CONNECTIVITY GAME is bounded by $O(n^\alpha)$ times the optimal cost. This is interesting, because it is independent of the distances between points, and only depends on their number. This bound is also tight: there is an instance which has a Nash equilibrium of cost $\Theta(n^\alpha)$ times the cost of the optimum. This tight instance has a special structure, and typical instances have a much better ratio. Indeed, for a random distribution of n points in a $\sqrt{n} \times \sqrt{n}$ plane region, the ratio of the cost of the worst Nash equilibrium to the optimal cost is bounded by $O(n^{\alpha/2} \log^\alpha n)$, with high probability. Also, we show that a local improvement algorithm results in Nash equilibria of cost $O(\log^{O(1)} n)$ times the optimal, with high probability.

5. There are instances of the REACHABILITY GAME with no pure Nash equilibrium, even when the points are located on a line for the case of $\alpha = 1$.

6. For a random distribution of n points in a $\sqrt{n} \times \sqrt{n}$ plane region, the REACHABILITY GAME has a $1 + o(1)$-approximate Nash equilibrium, with high probability.

Organization Section 2 defines all the basic graph theoretic and game theoretic concepts, and the models we study. Section 3 describes the results on the CONNECTIVITY GAME and Section 4 describes the results on the REACHABILITY GAME. We conclude in Section 5.

2. PRELIMINARIES

Our input is always a graph $H(V, E', \bar{w})$, with $|V| = n$ and with $\bar{w}$ being the weight vector on edges (i.e., w_e is the weight of edge $e \in E'$). A radius vector $\bar{r} \in \mathbb{R}^n$ (r_v being the radius of $v \in V$), induces a directed graph $G(V, E)$ in the following manner: $e = (u, v) \in E'$ is present in E if $r_u \geq w_e$. The graphs we consider here will not be arbitrary – they are either Euclidean or the weight vector $\bar{w}$ satisfies the triangle inequality. H is Euclidean if there is an embedding of V in $\mathbb{R}^k$ (k will usually be 2 or 3) such that $w_e = d(u, v)$, where $d()$ denotes the Euclidean distance function. In the Euclidean case, we will denote the graph induced by radius vector $\bar{r}$ by $G(V, \bar{r})$, since the weights are defined by the points themselves. Most of our discussion will be restricted to the Euclidean case, except in Section 3.3, where the graphs will not be Euclidean but would satisfy the triangle inequality. The cost of a radius vector $\bar{r}$ is defined as $C(\bar{r}) = \sum_v r_v^\alpha$, where α is a constant known as the distance power gradient, usually being 2.

Now we define the game theory notation we need; see [12] for more details. Formally, a game in its normal form is defined as the tuple $(I, \{S_v\}, \{U_v()\})$, where I is the set of players, S_v is the set of strategies for player $v \in I$ and $U_v : \Pi_v S_v \rightarrow \mathbb{R}$ is the utility function for player $v \in I$. In our models, each node v is associated with an independent, selfish agent; so $I = V$. We will identify the agent with the point in the description below. Each point v has to choose a radius (power level) and so the set of strategies S_v for $v \in V$ is $\mathbb{R}$, the set of all possible radii (note that it is sufficient to consider the finite set $\{d(v, w),\ w \in V\}$ for the set of possible radii for point v, instead of $\mathbb{R}$). A choice of strategies for all points is just a radius vector $\bar{r}$, with r_v being the radius chosen by v. The game is fully specified once we define the utility functions. In this paper we consider the following games.

The Connectivity Game In the CONNECTIVITY GAME, we are given pairs $(s_1, t_1), \ldots, (s_k, t_k)$, and each s_i needs to connect to t_i. Each s_i has to choose a radius so that it gets connected to t_i (possibly over several intermediate nodes), while keeping the radius as small as possible. For a radius vector $\bar{r}$, the utility U_{s_i} of point s_i is defined as $-M$ if s_i does not connect to t_i, M being some very large number, and is $-r_{s_i}^\alpha$ if s_i connects to t_i. The utilities of all other points are 0. The social optimum for such a game is a radius vector $\bar{r}$ such that s_i reaches t_i, for each i in $G(V, \bar{r})$, and $C(\bar{r})$ is minimized.

The Strong Connectivity Game The STRONG CONNECTIVITY GAME is a special case of the CONNECTIVITY GAME, where each point needs to connect with every other point. Therefore, for a radius vector $\bar{r}$, the utility of point v is $-M$ if v does not reach some point, and is $-r_v^\alpha$ if v does reach all other points. The social optimum for such a game is a radius vector $\bar{r}$ such that each $v \in V$ reaches all other points in V and $C(\bar{r})$ is minimized.

The Reachability Game Given $\bar{r}$, let $f_{\bar{r}}(v)$ denote the number of vertices $w \in W$ reachable from v in $G(V, \bar{r})$. The utility of a player $v \in V$ is defined as $U(v) = f_{\bar{r}}(v) - r_v^\alpha$.

For the CONNECTIVITY GAME and the STRONG CONNECTIVITY GAME, we denote the social optimum by OPT. The exponent α in the utility functions (i.e., in $-r_v^\alpha$) for the two connectivity games models the emission energy level and thus the cost that a node v incurs when sending to node at distance r_v; our results would still hold if we defined the utility to be the negative of any strictly positive power of r_v, which could be different from α.

In all these games, we will be interested in the Nash equilibria. A choice of strategies $\bar{r}$ is said to be a Nash equilibrium if $U_v(r_v, r_{-v}) \geq U_v(r'_v, r_{-v}), \forall v \in V$, where r_{-v} is the vector denoting the radii of all points other than v. Informally, $\bar{r}$ is a Nash equilibrium, if no point v has incentive to locally change its radius (while others keep their choices fixed).

Remark The Nash equilibrium defined above is called a *pure Nash equilibrium*, because the players are not allowed to randomize on their strategies. In the case where players choose their strategies according to a probability distribution, the appropriate notion is that of a *mixed Nash equilibrium.* We will consider only pure strategies and pure Nash equilibria in this paper, as mixed strategies do not seem to be very reasonable in studies of network design, such as ours (see also [2]). Finding Nash equilibria is desirable, however, pure Nash equilibria need not necessarily exist in all games; the notion of a β-approximate Nash equilibrium is a possibility to deal with this: a choice of strategies σ for all players is said to be a β-approximate Nash equilibrium (for $\beta \geq 1$), if unilateral deviation from σ by an individual player will increase its utility by at most a factor β. Approximate Nash equilibria might be a more suitable notion when only partial information is available.

Random Points in the Plane We consider random distributions of n points within a $\sqrt{n} \times \sqrt{n}$ region of the plane, denoted by A. Each point is thrown into this region independently and uniformly at random. This experiment places points roughly uniformly in the region, as is shown in the following lemma, which will be used later.

LEMMA 1. *Partition the region A into $n/\log^2 n$ parts of dimensions $\log n \times \log n$. For any such part B in an instance $\mathcal{P}$ of random points, the number of points in B is in the interval $[(1-\epsilon)\log^2 n, (1+\epsilon)\log^2 n]$, with high probability, where ϵ is a small, strictly positive constant.*

PROOF. The proof is a simple application of the Chernoff bound [11]. Let Z_i be the event that the ith point lies in region B; $Pr[Z_i = 1] = area(B)/n = \log^2 n/n$. Let $Z = \sum_i Z_i$ be the random variable denoting the number of points in region B. Then, $E[Z] = \log^2 n$.

By the Chernoff bound, $Pr[|Z - E[Z]| \geq \epsilon E[Z]] \leq 1/n^2$. The number of parts B is $O(n)$, and therefore, with probability at least $1 - 1/n$, the statement holds for each such part B. $\square$

3. EQUILIBRIA IN THE CONNECTIVITY GAME

3.1 The Strong Connectivity Game

In this section, we show that any instance of the STRONG CONNECTIVITY GAME has a pure Nash equilibrium. We prove an upper bound on the cost of all Nash equilibria and show with an example, that this upper bound is tight. We further present an algorithm that finds a Nash equilibrium for any given instance. We then show how our local improvement algorithm combined with an algorithm given in [10] yields a Nash equilibrium with no more than twice the optimum cost. All results in this section also hold for non-Euclidean instances where the triangle inequality is satisfied.

LEMMA 2. *Any instance of the* STRONG CONNECTIVITY GAME *has a pure Nash equilibrium. In fact, any local optimum is a Nash equilibrium.*

PROOF. Consider the set of radius vectors $\{\bar{r} \mid G(V, \bar{r})$ *is strongly connected*$\}$. This set is clearly non-empty and has local optima. Every local optimum is a Nash equilibrium. This is due to the fact that in a local optimum, no node has an incentive to decrease its radius. If any node were to decrease its radius, it would thereafter not be able to reach at least one other node anymore, which would decrease its utility to $-M$ and therefore not be in its interest. $\square$

Unfortunately, this game can have multiple Nash equilibria. The costs of different equilibria can vary widely. It is, however, surprising that the ratio of the cost of any Nash equilibrium to the optimal cost depends only on n, and is independent of the actual interpoint distances. The following lemma bounds the maximum cost of any Nash equilibrium.

LEMMA 3. *Any Nash equilibrium for the* STRONG CONNECTIVITY GAME *has cost at most n^α times the optimal cost.*

PROOF. Assume the radius vector $\bar{r}$ constitutes a Nash equilibrium for the STRONG CONNECTIVITY GAME and $\bar{s}$ is a choice of radii such that $C(\bar{s})$ is minimal. Denote by K_V the undirected, complete graph over all nodes in V, with edge weights $\ell(e) = d(u,v)$ for any edge $e = \{u, v\}$.

Fix any vertex v_0. Since there is a path from any other vertex to v_0 in $G(V, \bar{s})$, we can construct a rooted intree T, rooted at v_0: all vertices v have a directed path towards v_0 of the smallest hop length. By construction, $C(\bar{s}) \geq \sum_{e \in T} \ell(e)^\alpha$.

Next, observe that $r_v \leq \sum_{e \in T} \ell(e)$, because $d(v,w) \leq \sum_{e \in T} \ell(e)$, for any point $w \in V$ (by triangle inequality). Therefore, $C(\bar{r}) \leq n(\sum_{e \in T} \ell(e))^\alpha = nC(T)^\alpha$, where $C(T) = \sum_{e \in T} \ell(e)$. The ratio of the cost of the Nash equilibrium to the optimal cost is therefore bounded by

$$\frac{C(\bar{r})}{C(\bar{s})} \leq \frac{nC(T)^\alpha}{\sum_{e \in T} \ell(e)^\alpha} \tag{2}$$

Since $C(T)$ is fixed, this ratio is maximized when $\sum_{e \in T} \ell(e)^\alpha$ is minimized, subject to $\sum_{e \in T} \ell(e) = C(T)$. The minimum value of $\sum_{e \in T} \ell(e)^\alpha$ is $n(C(T)/n)^\alpha$, which is achieved when $\ell(e) = C(T)/n$ for each $e \in T$. Therefore, the ratio $C(\bar{r})/C(\bar{s})$ is bounded by n^α. $\square$

The bound in the above lemma is tight: there is an instance where the cost of a Nash equilibrium is $\Theta(n^\alpha)$ times the optimal cost. The example given in Fig. 1 demonstrates this.

OBSERVATION 1. *The instance given in Fig. 1 has a Nash equilibrium of cost $\Theta(n^\alpha)$ times the optimal cost.*

PROOF. Let n vertices be placed in the plane as illustrated in Fig. 1. There are $n/8$ nodes in each of the levels A to F, and the remaining $n/4$ nodes are situated in the set R along the right side of those levels. For n sufficiently big, this graph is Euclidean. The nodes in level A to F have a horizontal distance of 1 from each other and the vertical distances between the levels are $n/8$, 0.5, 0.4, 0.5, and 1 as given in the figure. The nodes in R all have distance 1 from each other.

Assume each node chooses radius $r_v = 1$. Then the induced graph $G(V, \bar{r})$ is strongly connected and the cost of the radius vector is $C(\bar{r}) = \sum_{v \in V} r_v^\alpha = n$. Now consider a second choice of radii $\bar{r}'$. We denote by r'_X the radius of any node v in the subset $X \subset V$.

$$r'_A = n/8,\ r'_B = 0.5,\ r'_C = 0.4,\ r'_D = 0.5,$$
$$r'_E = 1,\ r'_F = 1,\ r'_R = 1$$

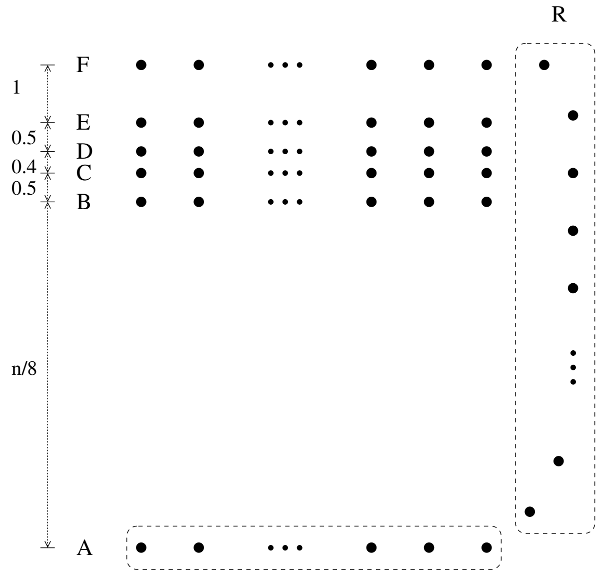

Figure 1: A graph where we can find a Nash equilibrium of cost $\Theta(n^\alpha)$ times the optimal cost.

The following proves that this choice of radii constitutes a Nash equilibrium. Certainly, in order for a node to have a utility function that is bigger than $-M$, that is, for the induced graph to be strongly connected, each node has to choose a positive (nonzero) radius. The nodes in F and R choose their smallest nonzero radius and therefore do not have an incentive to decrease it. Each of the nodes in level A to E chooses its radius such that it can reach the corresponding node in the level above. Observe that this induces a graph that can be traversed in various cycles in clockwise order, thus the induced graph is strongly connected. However, the graph cannot be traversed in counterclockwise direction, since there are no edges going from the nodes in level C to the nodes in level B. Hence, if any one of the nodes in levels A to E decreased its radius, it would not reach the node in the level above anymore and the cycle would be broken. The cost of this Nash equilibrium is $n/8 * (n/8)^\alpha + O(n) = O(n^{\alpha+1})$ and we have therefore found a Nash equilibrium of cost $\Theta(n^\alpha)$ times the optimal cost. □

Finding a Nash equilibrium via Local Improvement The following local improvement always leads to a Nash equilibrium.

1. Start with any choice of radii $\bar{r}^{(0)} = \bar{r}$ such that $G(V, \bar{r})$ is strongly connected.
2. Order the vertices arbitrarily as $v_1, \ldots, v_n$ and consider the vertices in this order.
3. In step i, the choice of radii is $\bar{r}^{(i-1)}$ initially. Vertex v_i chooses the smallest radius so that it can still reach every other vertex. Let $\bar{r}^{(i)}$ denote the choice of radii after v_i updates its radius.

LEMMA 4. *The above local improvement algorithm always leads to a Nash equilibrium.*

PROOF. Throughout the algorithm, the graph $G(V, \bar{r}^{(i)})$ stays strongly connected. Therefore, no vertex ever has an incentive to increase its radius. It remains to be proven that, for all i, vertex v_i does not have an incentive to further decrease its radius after any of the steps $i+1, \ldots, n$. Since in those steps no vertex increases its radius, no new paths are generated and v_i never develops an incentive to decrease its radius further. □

In order to find a choice of radii such that $G(V, \bar{r})$ is strongly connected, it is sufficient to go through all the nodes once in arbitrary order and let each of them choose the radius that maximizes its utility function. However, by determining the start vector for the local improvement algorithm in a more clever way, we can give an upper bound of the cost of the resulting Nash equilibrium. The following algorithm given in [10] constructs in $O(n)$ time a vector $\bar{r}^{(0)}$ such that $G(V, \bar{r}^{(0)})$ is strongly connected.

1. Construct the undirected, complete graph K_V with edge weights $d(u, v)^\alpha$ for all u and v.
2. Find a minimum weight spanning tree T of K_V.
3. For all $v \in V$ let $r_v = \max\{d(v, w) \mid \{v, w\} \in T\}$.

COROLLARY 1. *A Nash equilibrium of cost at most twice the optimum can be found in polynomial time.*

PROOF. Construct vector $\bar{r}^{(0)}$ with the algorithm given above. Lemma 4 implies that applying the local improvement algorithm with the start vector $\bar{r}^{(0)}$ yields a vector of radii $\bar{r}^{(n)} = \bar{r}$ which constitutes a Nash equilibrium of the STRONG CONNECTIVITY GAME. In [10] it is shown that $C(OPT) > C(T)$ and that

$$\begin{aligned} C(\bar{r}^{(0)}) &= \sum_{i=1}^{n} \max_{\{j \mid \{v_i, v_j\} \in T\}} d(v_i, v_j)^\alpha \\ &< \sum_{i=1}^{n} \sum_{\{j \mid \{v_i, v_j\} \in T\}} d(v_i, v_j)^\alpha \\ &= 2 * C(T) < 2 * C(OPT). \end{aligned}$$

Clearly $C(\bar{r}) \leq C(\bar{r}^{(0)})$. □

3.2 The Strong Connectivity Game for Random Points in the Plane

The bound on the ratio of the cost of the worst Nash equilibrium to the optimal cost in the previous section is tight, but the tight instance has a special structure. Most arrangements of points in the plane are likely to lack such a structure. Our results in this section show that this is indeed true: if the n points are distributed randomly in a square region of dimensions $\sqrt{n} \times \sqrt{n}$, the ratio is much smaller. As in Section 2, let $\mathcal{P}$ denote a random distribution of the n points. A denotes the region in which the points are thrown.

LEMMA 5. *For an instance $\mathcal{P}$, let $\bar{s}$ be a radius vector that minimizes the cost $C(\bar{s})$. Then, $C(\bar{s}) \geq \Omega(n/\log^\alpha n)$, with high probability.*

Proof. Partition the region A into square grid regions of dimensions $\log n \times \log n$. By Lemma 1, the number of points in each grid cell B of A is very close to $\log^2 n$. Let $\bar{s}$ be the optimal radius vector for this random instance.

Consider any such grid cell B in A that is not a boundary cell. Let S be the set consisting of B and the 8 cells adjacent to B. We first show that $\sum_{B' \in S} \sum_{i \in B'} r_i \geq \log n$. Since $G(V, \bar{s})$ is strongly connected, points in B must connect to points in cells not adjacent to it. Thus, there must be a directed path in $G(V, \bar{s})$ from a point in B to some cell B'' that is distance 2 away from B (cells adjacent to cells in S are said to be distance 2 away from B). Since this path has length at least $\log n$, $\sum_{B' \in S} \sum_{i \in B'} r_i \geq \log n$.

Next, we show that $\sum_{B' \in S} \sum_{i \in B'} r_i^\alpha \geq 1/\log^{\alpha-1} n$. This follows directly from convexity. Since $\sum_{B' \in S} \sum_{i \in B'} r_i \geq \log n$, $\sum_{B' \in S} \sum_{i \in B'} r_i^\alpha$ is minimized when all the r_i are equal, and therefore,

$$\begin{aligned}\sum_{B' \in S} \sum_{i \in B'} r_i^\alpha &\geq n_S (\log n / n_S)^\alpha \\ &= \log^\alpha n / n_S^{\alpha-1} \\ &\geq \Omega(1/\log^{\alpha-2} n)\end{aligned}$$

where $n_S = \Theta(\log^2 n)$ denotes the number of points contained in cells in S.

Finally, partition A into $n/(9 \log^2 n)$ parts, each part consisting of 9 cells. By using the above bound on the sum of the powers of the radii of the points in it, the lemma follows. □

Lemma 6. *Let $\bar{r}$ be any Nash equilibrium for an instance $\mathcal{P}$ and let $\bar{s}$ denote an optimal assignment of radii. Then, $C(\bar{r})/C(\bar{s}) \leq n^{\alpha/2} \log^\alpha n$, with high probability.*

Proof. By construction, $d(u, v) = O(\sqrt{n})$ for any two points u, v. Therefore, $r_u = O(\sqrt{n})$ for any point u, and $C(\bar{r}) \leq O(n^{\alpha/2+1})$. From Lemma 5, $C(\bar{s}) = \Omega(n/\log^\alpha n)$ for the optimal radius vector $\bar{s}$, and the lemma now follows. □

For a random distribution of points in the plane, the local improvement algorithm described earlier tends to result in Nash equilibria of better quality. In fact, the next lemma shows for $\alpha = 2$ that if we start with $\bar{r}^{(0)}$ such that $r_v^{(0)}$ is the largest possible radius, the resulting Nash equilibrium is quite good.

Lemma 7. *Let $\bar{r}^{(0)}$ be a radius vector that satisfies $r_v^{(0)} \geq \max_w \{d(v, w)\} \, \forall v \in V$. Let $\bar{r}'$ be the Nash equilibrium resulting from the local improvement algorithm, for any order of updating the vertices, and let $\bar{s}$ be the optimal radius vector. Then, for $\alpha = 2$, $C(\bar{r}')/C(\bar{s}) \leq O(\log^{O(1)} n)$, with high probability.*

Proof. The proof basically improves the bound computed in the proof of Lemma 6. Partition the $\sqrt{n} \times \sqrt{n}$ region, A, into n/k^2 blocks of dimensions $k \times k$ each, where k will be defined later. Observe that at the end of the local improvement algorithm, all except possibly one vertex in each block have radius $O(k)$. Intuitively, the vertex within a block that got updated last might have a large radius (even $\Theta(\sqrt{n})$) but all other vertices that got updated earlier need to choose a radius sufficient to connect to this *leader*.

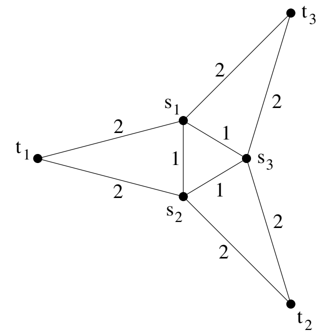

Figure 2: Connectivity game instance without Nash equilibrium.

Since there are n/k^2 blocks, there are at most this many leaders with large radius; the contribution of the remaining nodes to the cost of the Nash equilibrium is $O(nk^2)$. Each leader has radius $O(\sqrt{n})$, and the maximum contribution from the leaders is $O(n^2/k^2)$. Thus, the total cost is $O(nk^2 + n^2/k^2) = O(n\sqrt{n})$ for $k = n^{1/4}$. Using the bound from Lemma 5, this gives a bound of $O(\sqrt{n} \log^{O(1)} n)$ on the ratio of the cost of $\bar{r}'$ to the optimal.

We can improve this partitioning process repeatedly till the number of leaders becomes small. Consider the next step: partition the n/k^2 leaders from step 1 into n/k^4 blocks of k^2 elements each. One thing to note is that within the block defined in this step, the distances between two elements could be $O(k^2)$. Again, there is at most one leader in each block, who could possibly have radius larger than $O(k^2)$. This bounds the contribution of the non-leaders to the cost by $O(k^4 \frac{n}{k^2}) = O(nk^2)$. In general, if we repeat this process i times, the number of elements to consider at the ith step would be n/k^{2i-2}, and the radii of non leaders at the end of step i would be bounded by k^i. Therefore, the contribution of the non leaders to the cost is $O(nk^2)$. If we choose $k = O(\log n)$ and repeat this process for $i = O(\log n)$ steps, the number of elements in step i becomes $O(1)$, and the total cost over all steps is $O(nk^2 i) = O(n \log^{O(1)} n)$. The lemma now follows from the bound in Lemma 5. □

3.3 The Connectivity Game

In this section, we show that, unlike the Strong Connectivity Game, the general Connectivity Game need not have pure Nash equilibria, not even approximate ones. We also prove that determining whether a game instance has a pure Nash equilibrium is NP-complete by proposing a reduction from Monotone 2-in-3-Three-Satisfiability.

Figure 2 shows a game instance without pure Nash equilibrium. The instance consists of three sources s_1, s_2, s_3 and three sinks t_1, t_2, t_3; The sources form an equilateral triangle with edge length 1; vertices s_1, t_1, s_2 form an isosceles triangle with t_1 being at distance 2 from both s_1 and s_2; similarly vertices s_2, t_2, s_3 and vertices s_3, t_3, s_1 form isosceles triangles.

Observation 2. *No pure Nash equilibrium exists for the connectivity game instance given in Figure 2.*

Proof. Assume for the sake of contradiction that such an equilibrium exists with radii r_1, r_2, r_3 for the three source vertices s_1, s_2, s_3 respectively. We note immediately that $r_i \in \{1, 2\}$ as any radius $r_i < 1$ would mean that source s_i does not reach any other vertex in the graph and thus certainly will not reach its sink t_i, whereas any radius $r_i > 2$ cannot be part of a Nash equilibrium as reducing r_i to 2 would still allow source s_i to reach its sink t_i with a better utility. The following equivalences hold:

$$r_1 = 2 \iff r_3 = 1 \quad (3)$$
$$r_2 = 2 \iff r_1 = 1 \quad (4)$$
$$r_3 = 2 \iff r_2 = 1 \quad (5)$$

The "$\Longrightarrow$" direction holds because the source on the right-hand side of the implication would increase its utility by reducing its radius to 1; the "$\Longleftarrow$" direction holds because the source on the left-hand side would not reach its sink otherwise. Since we only have two possible values for each radius, Equations 3 – 5, imply the following equivalences:

$$r_1 = 1 \iff r_3 = 2 \quad (6)$$
$$r_2 = 1 \iff r_1 = 2 \quad (7)$$
$$r_3 = 1 \iff r_2 = 2 \quad (8)$$

Combining Equations 3, 8, and 4, we obtain the following contradiction:

$$r_1 = 2 \Longrightarrow r_3 = 1 \Longrightarrow r_2 = 2 \Longrightarrow r_1 = 1.$$

Thus, no pure Nash equilibrium exists for this instance. □

With respect to approximate Nash equilibria, a slight adaptation of the instance from Figure 2 yields the following negative result:

Corollary 2. *An instance of the connectivity game does not necessarily have an approximate Nash equilibrium.*

Proof. Consider the instance from Figure 2 and replace each edge of length 2 by an edge of length B, for an arbitrary $B > 1$. In geometric terms, this corresponds to making the three isosceles triangles longer. Each source node will now use a radius of either 1 or B. In any feasible combination of radii of the three sources as given in the proof of the previous lemma, a reduction from radius B to 1 will improve the utility of the corresponding source by a factor of B^α. Thus, this modified instance does not have an B^α-approximate Nash equilibrium. Since we can choose B arbitrarily large, the corollary follows. □

As a puzzling observation, we briefly look at the *mixed* Nash equilibrium for the instance in Figure 2. Recall that the utility function for a source is $-M$ if it cannot reach its sink, it is -2^α if it chooses radius 2, and it is -1 for radius 1. Straight-forward analysis shows that the only mixed Nash equilibrium is for all sources to choose their radius $r_i = 1$ with probability $\frac{1}{M-1}$, and $r_i = 2$ with probability $1 - \frac{1}{M-1}$. If we let $M \longrightarrow \infty$, the mixed Nash equilibrium lets all sources choose radius $r_i = 2$ with probability 1. Thus, the mixed Nash equilibrium appears to be identical to a pure-strategy solution that is clearly not a pure Nash equilibrium. In fact, if we set $M = \infty$, the game does not even have a *mixed* Nash equilibrium, which – intuitively – is due to the non-continuity of the utility function.

Knowing that Nash equilibria do not always exist does not necessarily prevent us from designing a polynomial-time algorithm that finds a pure Nash equilibrium if it exists. However, we now show that the simple question (dubbed Pure Nash Connectivity with Triangle Inequality) whether a given connectivity game has a pure Nash equilibrium is NP-hard to answer, if the triangle inequality holds on the input graph. The corresponding problem for purely geometric graphs (with embeddings in the plane) remains open. We show this hardness result by reducing Monotone 1-in-3 Three Satisfiability to Pure Nash Connectivity with Triangle Inequality.

Definition 1. *The problem* Monotone 1-in-3 Three Satisfiability *consists of finding a truth assignment to the variables of a given formula with three positive literals in each clause such that exactly one literal in each clause is true.*

Monotone 1-in-3 Three Satisfiability is NP-hard [9].

Lemma 8. Pure Nash Connectivity with Triangle Inequality *is NP-hard.*

The proof of Lemma 8 is given in the appendix.

4. THE REACHABILITY GAME

In this section, we show that the Reachability Game need not have a pure Nash equilibrium, even for a 1-dimensional instance. The simple example of Figure 3 is one such instance. Multiple vertices are located at the same point (e.g. points 1, 3, 4) in this figure. This is only for the purpose of keeping the example simple; the colocated points can be perturbed slightly to be located very close to each other.

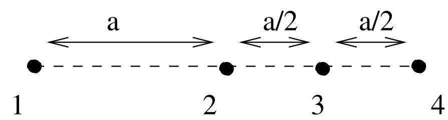

Figure 3: Instance of Reachability Game with no pure Nash equilibrium: there are $a/2 - 1$ vertices located together at points 1 and 3, three vertices at point 4 and a single vertex at point 2. The distances between the points are as shown. The value a is any number larger than 4.

Observation 3. *The* Reachability Game *instance in Figure 3 with $\alpha = 1$ does not have a pure Nash equilibrium.*

Proof. Note that in any Nash equilibrium, only one of a set of colocated vertices can have positive radius – all the other vertices can keep their radius 0 without affecting their utility. In what follows, we use r_3 (r_1, r_4, respectively) to denote the radius of the vertex located at point 3 (1, 4, respectively) with the largest radius, keeping in mind that the other vertices at point 3 (1, 4, respectively) have radius 0.

The total number of vertices in this instance, n is $a + 2$. Therefore, no vertex has radius more than $a + 1$. Also, $r_1 = 0$, since $U_1(0, \sigma_{-1}) = a/2 - 1$, $U_1(x, \sigma_{-1}) = a/2 - 1 - x$, for any $x < a$ and $U_1(a, \sigma_{-1}) \le 2$, for any choice σ_{-1} of radii by vertices at points $2, 3, 4$. Further, r_4 has no influence on the utilities of vertices at points 2 or 3, namely $U_2(), U_3()$: vertices at these points cannot reach any more vertices if $r_4 > 0$. Therefore, $U_2()$ and $U_3()$ depend only on r_2 and r_3, and are denoted by $U_2(r_2, r_3)$ and $U_3(r_2, r_3)$ in the discussion below.

The observation now follows from the following four implications.

1. $r_3 < a/2 \Rightarrow r_2 = a$: If $r_3 < a/2$, vertices at point 3 do not reach vertices at point 4. Therefore, $U_2(0, r_3) = 1, U_2(a/2, r_3) = 0$ and $U_2(a, r_3) = 2$, which implies $r_2 = a$.

2. $r_3 \geq a/2 \Rightarrow r_2 = a/2$: In this case, $U_2(a/2, a/2) = a/2 + 3 - a/2 = 3 > U_2(a, a/2)$ and so $r_2 = a/2$.

3. $r_2 = a \Rightarrow r_3 = a/2$: In this case, $U_3(a, a/2) = a + 2 - a/2 = a/2 + 2 > U_3(a, 0) = a/2 - 1$, and so $r_3 = a/2$.

4. $r_2 < a \Rightarrow r_3 = 0$: In this case, the vertex at point 2 does not reach the vertices at point 1. As a result, $U_3(r_2, a/2) = a/2 + 3 - a/2 < U_3(r_2, 0) = a/2 - 1$, and so $r_3 = 0$.

Suppose, $r_3 < a/2$. Then implications (1) and (3) lead to a contradiction. Suppose $r_3 \geq a/2$. Then implications (2) and (4) lead to a contradiction. □

4.1 The Reachability Game for Random Points in the Plane

In this section, we show that an approximate Nash equilibrium always exists for a random distribution of points in the plane, as described in section 2.

LEMMA 9. *For an instance $\mathcal{P}$ of random points in the plane, the* REACHABILITY GAME *with any α has a $1 + o(1)$-approximate Nash equilibrium, with high probability.*

PROOF. The proof is by constructing a radius vector $\bar{r}$ that is an approximate Nash equilibrium. The graph $G(V, \bar{r})$ will actually be strongly connected.

Partition the region A into square regions of dimensions $\log n \times \log n$: there are $n/\log^2 n$ such regions. Within each such region, choose one node arbitrarily as a leader for that region. The leader of each region chooses a radius of $4 \log n$, so that it is connected to the leaders of the regions immediately adjacent to it. All the other points in each region choose a radius in the range $[0, \log n]$ so that they get connected to the leader of that region. Let $\bar{r}$ be the resulting radius vector. It is easy to check that $G(V, \bar{r})$ is strongly connected, and each point has utility at least $n - (4 \log n)^\alpha$. The maximum utility of any point is n, and therefore this choice is a $\frac{n}{n-(4\log n)^\alpha} = 1 + o(1)$-approximate Nash equilibrium. □

5. CONCLUSIONS AND OPEN PROBLEMS

We consider two topology control games arising in ad hoc networks in the presence of selfish, non-cooperative agents in this paper and study the existence of Nash equilibria, their quality and algorithms for computing them. Our work motivates further game theoretic study of protocols for ad hoc networks. Some of the interesting open questions are the following.

1. What is the *average* cost of a Nash equilibrium for the STRONG CONNECTIVITY GAME? Design mechanisms or pricing schemes that reduce their cost.

2. Is it NP-hard to decide whether an instance of the REACHABILITY GAME has a Nash equilibrium?

3. Intermediate nodes in the games we study have to be paid by the source for each message. Augment the games to include this price.

6. ACKNOWLEDGMENTS

We are very grateful to Madhav Marathe for his constant encouragement and valuable suggestions.

7. REFERENCES

[1] L. Anderegg, S. Eidenbenz. Ad hoc-VCG: A Truthful and Cost-Efficient Routing Protocol for Mobile Ad Hoc Networks with Selfish Agents. *Proceedings of Mobicom*, 2003.

[2] E. Ansheleich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Proceedings of the Symposium on Theory of Computing*, 2003.

[3] M. Beckmann, C.B. McGuire and C.B. Winsten. *Studies in the Economics of transportation*. Yale University Press, 1956.

[4] L. Buttyan and J. Hubaux. Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks; *Accepted for publication in ACM/Kluwer Mobile Networks and Applications MONET)*, October 2003, Vol. 8 No. 5.

[5] S.C. Dafermos and F.T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards*, Series B, 73 B(2), pp 91-118, 1969.

[6] S. Doshi, S. Bhandare, T. Brown. An On-demand minimum energy routing protocol for a wireless ad hoc network; in: *ACM Mobile Computing and Communications Review*, vol. 6, no. 3, July 2002.

[7] A. Fabrikant, A. Luthra, E. Maneva, S. Papadimitriou, and S. Shenker. On a network creation game.

[8] J. Freebersyer and B. Leiner. A DoD perspective on mobile ad hoc networks. In C. Perkins, ed. *Ad Hoc Networking*, pp 314-318, 2001.

[9] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of* NP*-completeness*; Freeman, San Fransisco, 1979.

[10] L. M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks, *Theoretical Computer Science* 243 (2000) pp 289-305.

[11] R. Motwani and P. Raghavan. *Randomized Algorithms*, Cambridge University Press, 1995.

[12] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1994.

[13] C. Perkins. Ad hoc networking: an introduction. In C. Perkins, ed. *Ad Hoc Networking*, pp 314-318, 2001.

[14] R. Rajaraman. Topology control and routing in ad hoc networks: a survey. *SIGACT News* 33:60-73, June 2002.

[15] T. Roughgarden. *Selfish Routing*. Ph.D. Thesis, Cornell University, 2002.

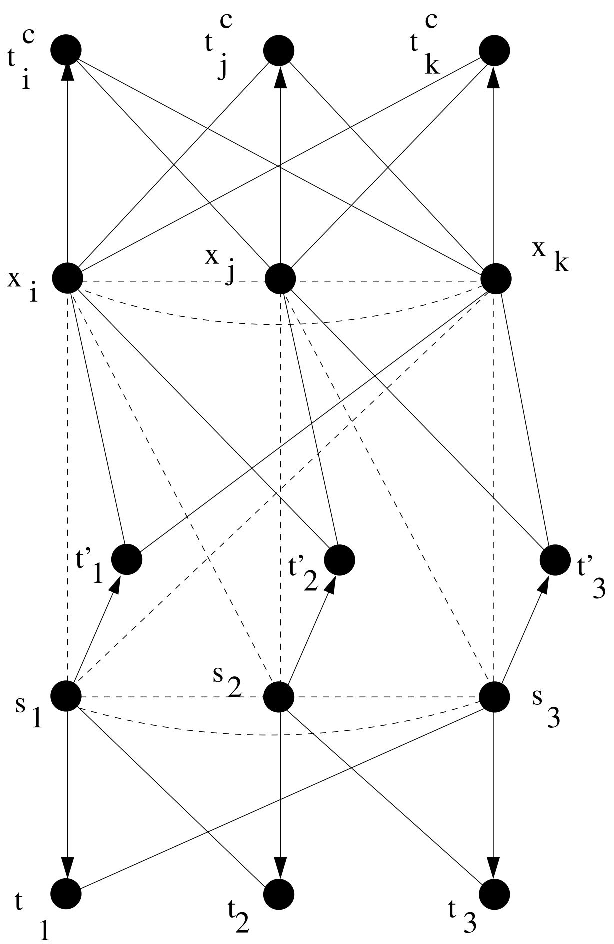

Figure 4: Clause gadget: source-sink relationships are indicated by arrows; dashed lines denote edges with weight 1, solid lines denote edges with weight 2.

[16] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42nd Symposium on the Foundations of Computer Science*, pp 472-481, 2001.

APPENDIX

Proof of Lemma 8 Given a Monotone 1-in-3 Three Satisfiability instance I consisting of variables $x_1, \dots, x_n$ and m clauses with each clause being a 3-tuple of positive literals, we construct a Pure Nash Connectivity with Triangle Inequality instance I' as follows: For each variable x_i, we create a source node x_i in the graph that we call a variable node. We insert an edge of weight 1 between two nodes x_i and x_j, if there exists a clause in which both variables appear as positive literals.

Figure 4 shows a clause gadget: for each clause $c = (x_i, x_j, x_k)$, we create three nodes t_i^c, t_j^c, t_k^c, where t_i^c is a sink node that source node x_i must reach, accordingly for t_j^c and t_k^c. Edges of weight 2 are inserted between the three source nodes x_i, x_j, x_k and the three sink nodes t_i^c, t_j^c, t_k^c. We call the part of the clause gadget containing these six nodes the upper part.

In contrast, the lower part of the clause gadget consists of nine nodes that are created individually for each clause. The six nodes $s_1, s_2, s_3, t_1, t_2, t_3$ (see Figure 4) form exactly the same graph as the one given as an example of a graph without Nash equilibrium in Figure 2 with source-sink pairs (s_1, t_1), (s_2, t_2), and (s_3, t_3). In addition, each of the sources s_1, s_2, s_3 needs to reach a second sink node t'_1, t'_2, t'_3. These additional sink nodes are connected to their corresponding source nodes by edges of length 2. The upper and the lower part of the clause gadget are connected through edges (x_i, s_1), (x_i, s_2), (x_j, s_2), (x_j, s_3), (x_k, s_3), (x_k, s_1) of length 1 and through edges (x_i, t'_1), (x_i, t'_2), (x_j, t'_2), (x_j, t'_3), (x_k, t'_3), (x_k, t'_1) of length 2.[1]

This completes the description of the Pure Nash Connectivity with Triangle Inequality instance I'. As we only have edge weights 1 and 2, our graph satisfies the triangle inequality. We have created one node for each variable and 12 nodes for each clause, giving a total number of $n + 12m$ nodes, thus the reduction is polynomial. The key idea of the construction is that the lower part of each clause gadget will only have a Nash equilibrium if exactly one of the variable nodes in the upper part sets its radius to 2 and the other two variable nodes set their radii to 1.

To be more precise, if a "1-in-3" satisfying truth assignment exists for the variables of the Monotone 1-in-3 Three Satisfiability instance I, we obtain a radius vector for the nodes of the Pure Nash Connectivity with Triangle Inequality instance I' that constitutes a Nash equilibrium by setting the radii of exactly those variable nodes x_i in I' to 2, of which the corresponding variables x_i in I are set to true in the truth assignment. All other variable nodes set their radius to 1. For better illustration, assume w.l.o.g. (due to the symmetry of the construction) that variables x_i and x_j are set to false, while variable x_k is set to true in the assignment, thus the clause $c = (x_i, x_j, x_k)$ is "1-in3" satisfied, and thus the radii of nodes x_i and x_j are 1 and the radius of node x_k is 2. Thus, variable node x_k reaches its sink t_k^c directly and nodes x_i and x_k reach their sinks t_i^c and t_j^c via node x_k. This radii assignment also forces a radius assignment for the sources on the lower part of the clause gadget: source s_2 has to set its radius to 2 in order to reach sink t'_2 as nodes x_i and x_j have both set their radii to 1 and thus do not reach t'_2; this makes it sufficient for source s_3 to set its radius to 1 as it can reach sink t_3 via s_2 and sink t'_2 via upper part node x_k; this in turn forces s_1 to set its radius to 2 as it cannot reach t_1 otherwise. To see that this radius vector constitutes a Nash equilibrium, first note that all sources reach their sinks and thus have no incentive to increase their radii. Similarly, each source with radius set to 2 would lose the connection to at least on of its sinks if it reduced its radius to 1. Thus, we have found a radius vector that constitutes a Nash equilibrium.

We also need to show that any Nash equilibrium of I' induces a "1-in-3" satisfying truth assignment of the variables of I. Assume we are given a radii assignment for all sources in I that constitutes a Nash equilibrium. We first note that no source will choose a radius larger than 2 in any Nash equilibrium as it will directly reach all its sinks with a radius of 2, neither will a source set its radius to less than 1, as it will not reach any other node with such a small ra-

[1] The nine nodes of the lower part of the clause gadget would be more aptly named $s_1^c, s_2^c, s_3^c, t_1^c, t_2^c, t_3^c, t_1'^c, t_2'^c, t_3'^c$, as they are individual to clause c, but for ease of presentation, we have chosen to drop the c-index.

dius. Let us consider the clause gadget representing clause $c = (x_i, x_j, x_k)$. We distinguish four cases of radii assignment for the three source nodes s_1, s_2, s_3 as they are in the lower part of the clause gadget. For simplicity, let $(2, 1, 1)$ denote the radius of source s_1 set to 2 and the radii of the other two sources s_2 and s_3 set to 1; accordingly for other radii choices:

- Radii vector $(1, 1, 1)$:
  In this case, none of the sinks t_1, t_2, t_3 is reached by its source, thus the radius assignment cannot be a Nash equilibrium.
- Radii vector $(2, 1, 1)$:
  In this case, sink t_3 is not reached by its source, thus this cannot be a Nash equilibrium. The radii vectors $(1, 2, 1)$ and $(1, 1, 2)$ are equivalent due to symmetry.
- Radii vector $(2, 2, 2)$:
  In this case, all variable nodes x_i, x_j, x_k of the clause must have set their radius to 1 as at least two of the sources s_1, s_2, s_3 would have an incentive to reduce their radius otherwise. However, with the radii of x_i, x_j, x_k all set to 1, sinks t_1^c, t_2^c, t_3^c in the upper part of the clause will not be reached by their sources. Thus, this cannot be a Nash equilibrium.
- Radii vector $(2, 2, 1)$:
  In this case, either x_j or x_k must have radius 2 as sink t_3' would not be reached otherwise. If x_j had its radius set to 2, then source s_2 would have an incentive to reduce its radius to 1, independent of the radii of x_i and x_k. Similarly, if x_i had its radius set to 2, then source s_2 would have an incentive to reduce its radius to 1, independent of the radii of x_j and x_k. However, if only x_k has its radius set to 2 and x_i and x_j set to 1, then we have a valid Nash equilibrium, from which we can easily read off a truth assignment for the variables: x_i and x_j are false, x_k is true. We can argue for the radii vectors $(2, 1, 2)$ and $(1, 2, 2)$ similarly.

Thus, the a radii vector can only be a Nash equilibrium, if it has exactly one variable node in each clause set to radius 2 and the other two variable nodes set to radius 1. From this, we can assign a "1-in-3" satisfying truth assignment to the variables of I immediately. This completes our proof. □