Mobile Networks and Applications 11, 143–159, 2006

DOI: 10.1007/s11036-005-4468-y

Equilibria in Topology Control Games for Ad Hoc Networks*

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

V.S. ANIL KUMAR
Virginia Bio-informatics Institute and Dept. of Computer Science, Virginia Tech, Blackburg, VA, 24061,

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

Published online: 31 March 2006

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 each network node is required to be connected to all other nodes (we call this the STRONG CONNECTIVITY GAME). 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 Nash equilibria do not necessarily exist. We further study how to find Nash equilibria with *incentive-compatible* algorithms and compare the cost of Nash equilibria to the cost of a social optimum, which is a radius assignment that minimizes the total cost in a network where nodes cooperate. We also study variations of the games; one where nodes not only have to be connected, but *k-connected*, and one that we call the REACHABILITY GAME, where nodes have to reach as many other nodes as possible, while keeping costs low. We extend our study of the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME to wireless networks with directional antennas and wireline networks, where nodes need to choose neighbors to which they will pay a link. Our work is a first step towards game-theoretic analyses of topology control in wireless and wireline networks.

Keywords: ad hoc networks, wireline networks, directional antenna networks, game theory, Nash equilibrium, topology control

1. Introduction

Unlike traditional, fixed wireline networks, 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 transceivers or radios in some geographical location that communicate with each other possibly via intermediate transceivers who forward data. Earliest examples of ad hoc networks were in military applications. Recent advances in the commercialization of intelligent radio devices are likely to lead to the wide-spread emergence of ad hoc or hybrid networks [16].

Topology control. Depending on its power level and on the nature of environmental interference, a node in an ad hoc network can reach all nodes in a certain range. 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 [7]:

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

where K is a constant and α is the distance-power gradient varying between one and six depending on the environmental conditions of the network. If this power exceeds a minimum level, a node v at this point can successfully receive the message from node u, and falls within the transmission range of u. Antennas are most often assumed to be omnidirectional, i.e., the power level depends only on the distance from the sender, and not on the direction; in this case, the radius is a proxy for the power level, and this gives us the *transmission graph* [17]. The transmission graph is a directed graph $G(V, E)$ that is defined as follows for a fixed set of radii: V is the set of nodes, and the directed edge $e = (u, v)$ is present in E if node v is within the power range of node u. We also consider graphs resulting from directional antennas and wireline networks, where this assumption of omnidirectionality does not hold.

*A preliminary version of this paper appeared in DIALM-POMC ’03 [8].

Efficient communication in an ad hoc network 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 [17] for a survey. Existing work on energy-efficient topology control [17] 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 to be studied 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. 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 [1,4]. 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. [3,6]) 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 [18] 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 [19] 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 [2,9].

For modeling communication networks as games, it is reasonable to think of each node as a player or an agent. Each player has a certain set of strategies: in the games we consider, a player needs to choose a radius or a set of neighbors, and such a choice 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 [15] for details). A choice of strategies $\bar{\sigma}$ for all players is said to be a Nash equilibrium, if no player has an incentive to deviate from $\bar{\sigma}$ 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 three types of scenarios on these networks:

(i) wireless networks with omnidirectional antennas
(ii) wireless networks with directional antennas, and
(iii) wireline networks.

In contrast to the transmission graphs resulting from omnidirectional antennas (as explained earlier), directional antennas allow specific links to be chosen.

As in the omnidirectional case, these links are not symmetric, i.e., a link from u to v does not imply a link in the reverse direction. It is predicted that directional antennas will be deployed extensively in future generation networks (see e.g. [20]). In the wireline case the links are symmetric, and so they need to be paid for only once.

On these networks, we study the CONNECTIVITY GAME: we are given a set of source-destination pairs, and each pair needs to be connected (see Section 2 for details). Each node chooses the smallest possible radius (in the omnidirectional wireless case) and cheapest set of links (in the directional wireless and wireline cases) to achieve this goal. We study the following questions: do Nash Equilibria exist, and if they do, what is their quality compared to the optimal choice? Also, how do local algorithms for finding such equilibria perform? A special case of this game is the STRONG CONNECTIVITY GAME, where each node needs to reach every other node. We also extend these games to the cases where k node-disjoint paths are needed between each source and destination pair; these are denoted by k-CONNECTIVITY GAME and the STRONG k-CONNECTIVITY GAME; this captures the requirement of robustness in such networks, so that the network operation is not disrupted by some node deletions.

Our main results on the various connectivity games are summarized below.

1. For the omnidirectional antenna network, 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. The first of these results extends to the k-CONNECTIVITY GAME: for any k, the game does not necessarily have a pure Nash equilibrium. We further show that the CONNECTIVITY GAME need not always have a pure Nash equilibrium in neither the directional antenna network nor the wireline network.
2. For Euclidean instances, the STRONG CONNECTIVITY GAME always has a pure Nash equilibrium, in all the three networks we consider. In fact, there are multiple Nash equilibria, whose costs can vary widely. Also, any local optimum to the total energy function is a Nash equilibrium in the omnidirectional antenna network. This also holds for the STRONG k-CONNECTIVITY GAME.
3. For Euclidean instances, there is a simple local improvement algorithm that yields a Nash equilibrium for the STRONG CONNECTIVITY GAME for the wireless cases. Using an observation from [13], this yields an algorithm to find a Nash equilibrium of cost at most twice the optimum. For the k-connected version, a constant factor can be achieved for the case of $\alpha = 1$, but for larger α, the factor depends on k.
4. For Euclidean instances, the cost of any Nash equilibrium for the STRONG CONNECTIVITY GAME on omnidirectional wireless instances 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 even holds for the STRONG k-CONNECTIVITY GAME, where the lack of dependence on k is notable. This bound is 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.

We consider a related game, called the REACHABILITY GAME, in which 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 (and each node chooses a radius in order to maximize its utility). Informally, each node wishes to maximize the number of nodes it connects to while keeping its cost low. We show that there are omnidirectional wireless (Euclidean) instances of this game with no pure Nash equilibrium, even when the points are located on a line for the case of $\alpha = 1$. 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.

Related results. The CONNECTIVITY GAME can be viewed as a power level version of Anshelevich et al. [2]. Their work involves players on a network, with edges having costs, and each player has a set of terminals that need to be connected - so edges have to be put in to achieve the desired connectivity, but an edge can be used only if it is paid for by the players; therefore, strategies for the players involve choosing payments for the edges so that their connectivity requirements are met. Anshelevich et al. show that if each player has to connect a single terminal to a common source, a Nash equilibrium exists, and they describe an algorithm to compute a $(1+\epsilon)$-approximate equilibrium. In contrast, if the players want to connect multiple terminals, no Nash equilibrium might exist, but they show how to construct an approximate equilibrium. Our model can be viewed as a restriction of their model to an ad hoc network setting, where a node can only reach all the 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 our wireline and directional wireless versions of the game, a node can only pay for incident edges, and the payment cannot be divided among nodes. The game defined in [9] is similar to the wireline version of our games and the authors also study quality of Nash equilibria, but the utility function of their game is drastically different from our game. A good survey on topology control is [17], and [12] considers the problem of finding a radius assignment that minimizes the power while maintaining k-connectivity.

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 STRONG CONNECTIVITY GAME and the STRONG k-CONNECTIVITY GAME, Section 4 describes the results on the CONNECTIVITY GAME and the k-CONNECTIVITY GAME, and Section 5 describes the results on the REACHABILITY GAME. In Section 6 we describe the extensions of the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME to wireless networks with directional antennas, and in Section 7 to wireline networks. We conclude in Section 8. Some proofs are presented in the Appendix.

2. Preliminaries

Our input is always an undirected graph $H(V, E', w)$, with $|V| = n$ and with $\bar{w} \in \mathbb{R}_+^{|E|}$ being the weight vector on edges (i.e., $w_e \geq 0$ is the weight of edge $e \in E'$). In the wireless networks with omnidirectional antennas, 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 $r_u \geq w_e$. We will denote the graph induced by radius vector $\bar{r}$ with $G(V, \bar{r})$. In Sections 6 and 7 we will introduce the equivalent definitions for a wireless network with directional antennas and a wireline network.

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.

We say two nodes u and v are *k-connected*, if they are connected via k internally node-disjoint paths. If u and v are *k-connected*, we can delete any arbitrary set $S \subset V \backslash \{u, v\}$, $|S| = k - 1$, without disconnecting u and v. Note that if there is a direct link between u and v, the second statement holds, but the nodes are not *k-connected* according to our definition. In particular, if there is a direct link between u and v, then u and v are *k-connected* only iff we can in addition delete an arbitrary set $S \subset V\backslash\{u,v\}$, $|S| = k - 2$, without disconnecting u and v.

Now we define the game theory notation we need; see [15] 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: $\Pi_v S_v \to \mathbb{R}$ is the utility function for player $v \in I$. In our models, each node v is associated with an independent, selfish agent; so $I = V$. Each point v has to choose a radius (power level). Therefore 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 $\{w_e, e = \{v, v'\} \in E'\}$ for the set of possible radii for point v, instead of $\mathbb{R}$). A choice of strategies $\bar{\sigma}$ for all points is a radius vector $\bar{r}$, with r_v being the radius chosen by v. The cost a node v incurs is $C(v) = r_v^\alpha$ for strategy r_v, where α is a constant known as the distance power gradient, usually being 2. We can then define the cost of a strategy vector as $C\bar{r} = \sum_v r_v^\alpha$. The game is fully specified once we define the utility functions.

Nash equilibria. In all the described games, we will be interested in Nash equilibria. A choice of strategies $\bar{\sigma}$ is said to be a *Nash equilibrium* if $U_v(\sigma_v, \sigma_{-v}) \geq U_v(\sigma_v', \sigma_{-v}) \forall v \varepsilon V$, where σ_{-v} is the vector denoting the strategies of all points other than v. Informally, $\bar{\sigma}$ is a Nash equilibrium, if no point v has incentive to locally change its strategy (while others keep their choices fixed).

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 $\bar{\sigma}$ for all players is said to be a β-approximate Nash equilibrium, if unilateral deviation from $\bar{\sigma}$ by an individual player will increase its utility by at most a factor $\beta \geq 1$ for positive utility functions, and by a factor of at least $1/\beta \leq 1$ for negative utility functions. Approximate Nash equilibria might be a more suitable notion when only partial information is available.

In this paper we consider the following games.

The connectivity games. In the CONNECTIVITY GAME we are given j source-sink pairs $(s_1, t_1), \ldots, (s_j, t_j)$; each source s_i needs to connect to its target or sink t_i. The sources and targets are located on the vertices of the input graph. We denote by $S(v)$ (resp. $T(v)$) the set of sources (resp. targets) that are located on vertex v. Each vertex v has to choose a strategy σ_v, such that every source $s \in S(v)$ gets connected to its corresponding target (possibly over several intermediate nodes) in the resulting transmission graph, while keeping its cost $C(v)$ as small as possible. We assume for the CONNECTIVITY GAME that there is a direct link between each source and its sink in the input graph H.

For a strategy vector $\bar{\sigma}$, the utility $U_v(\bar{\sigma})$ of vertex v is defined as $-M$ if at least one $s_i \in S(v)$ does not connect to t_i, M being some very large number, and is $-C(v)$, if all $s_i \in S(v)$ connect to their t_i. The utilities of all points for which $S(v)$ is empty is 0. The social optimum for such a game is a strategy vector $\bar{\sigma}$ such that each s_i reaches t_i in the transmission graph $G(V, \bar{\sigma})$ and $C(\bar{\sigma})$ is minimized.

The k-CONNECTIVITY GAME is a generalization of the CONNECTIVITY GAME, where each node needs to reach its target via k internally disjoint paths. In this game the input graph H needs to be k-connected.

The strong connectivity games. The STRONG CONNECTIVITY GAME is a special case of the CONNECTIVITY GAME, in which each point needs to connect to *every* other point. The STRONG k-CONNECTIVITY GAME is a generalization of the STRONG CONNECTIVITY GAME, where each node needs to reach every other node via k internally disjoint paths. We also assume that there must be a direct link between each source and its sink in the input graph H, which implies in this case that H is required to be a complete graph.

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

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. The lemma follows from the standard Chernoff bound (e.g. [5]); "with high probability" means the probability is at least $1-1/n^{O(1)}$.

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 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.*

3. Equilibria in the Strong Connectivity Games

In this section we deal with the STRONG CONNECTIVITY GAME on omnidirectional wireless instances, and will not state this explicitly in all the results. Many of the results are stated for the STRONG k-CONNECTIVITY GAME, where each node has to reach every other node over k internally disjoint paths.

Existence of Nash equlibria

Lemma 2. *Any instance of the* STRONG k-CONNECTIVITY GAME *has a pure Nash equilibrium for $k \leq n-1$. In fact, any radius vector $\bar{r}$ that is a local optimum[1] for the cost function $C(\bar{r})$ is a Nash equilibrium.*

Proof: Consider the set of radius vectors $\{\bar{r} \mid G(V, \bar{r})$ *is strongly k-connected*$\}$. This set is non-empty for any $k \leq n-1$ and has local optima with respect to the cost function $C(\bar{r})$. We will prove by contradiction that every local optimum is a Nash equilibrium. Let $\bar{s}$ be a local optimum. Suppose that there is a node v that has an incentive to decrease its radius from s_v to s'_v, that is, it still has k internally disjoint paths to every other node in $G' := G(V, \bar{s}')$, even after at least one of its links (v, v') gets deleted. We claim that every node u still has k internally disjoint paths to every other node in G', which then contradicts that $\bar{s}$ is a local optimum. Let us first prove that u is still *k-connected* to v'. If this were not the case, there must be a separator $S \subset V \backslash \{u, v'\}$ of size $k-1$ (or $k-2$ if u has a direct link to v') that separates u and v'. First, observe that $v \notin S$; if this were not the case, i.e., $v \in S$, then u and v' would not be *k-connected* even in G, since the only difference between G and G' is the presence of extra edges going out of v. Therefore, $v \notin S$. Assume u reaches v in $G' \backslash S$. Since v reaches v' even in $G' \backslash S$, u must reach v' in $G' \backslash S$. Now assume u and v are disconnected in $G' \backslash S$. But this would imply that S is a separator for u and v in the graph $G(V, \bar{s})$ already, since we deleted only links going out of v, and therefore this contradicts the fact that $G(V, \bar{s})$ is strongly k-connected. The same arguments hold for any node $w \neq v'$. □

[1] Let $R = \{\bar{r} \mid G(V, \bar{r})$ *is strongly* k-connected$\}$. $\bar{r}' \in R$ is a *neighbor* of $\bar{r} \in R$, if there is a $w \in V$ such that $r_w \neq r'_w$ and $r_w = r'_w$ for all $v \in V$, $v \neq w$. $\bar{r}$ $\bar{r} \in R$ is a *local optimum* of R, if $C(\bar{r}) \geq C(\bar{r}'$ for all radius vectors $\bar{r}'$ that are neighbors of $\bar{r}$.

Quality of Nash equlibria. Unfortunately, this game can have multiple Nash equilibria of widely varying costs. 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 as well as of the required connectivity k of the transmission graph. The following lemma bounds the maximum cost of any Nash equilibrium.

Lemma 3. *Any Nash equilibrium for the* STRONG k-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 k-CONNECTIVITY GAME and $\bar{s}$ is a choice of radii such that $C(\bar{s})$ is minimal and $G(V, \bar{s})$ is strongly k-connected.

Fix any vertex v_0. Since any other vertex can reach v_0 in $G(V, \bar{s})$ over k internally disjoint paths, we can construct a subgraph G' of $G(V, \bar{s})$ that includes all edges needed such that every node can reach v_0 via k internally disjoint paths. Denote by $w_{\max}(v) = \max\{w_e, e = (v, v') \,\varepsilon\, G'\}$. By construction,

$$C(\bar{s}) = \sum_{v \in V} s_v^\alpha \geq \sum_{v \in V} w_{\max}(v)^\alpha$$

Next, observe that $r_v \leq \sum_{v \varepsilon V} w_{\max}(v)$, because $w_{(v,v')} \leq \sum_{v \varepsilon V} w_{\max}(v)$, for any point $v' \,\varepsilon\, V$ (since the edge lengths are euclidean distances, and satisfy the triangle inequality). Therefore, $C(\bar{r}) \leq n(\sum_{v \in V} w_{\max}(v))^\alpha$. The ratio of the cost of the Nash equilibrium to the optimal cost is therefore bounded by

$$\frac{C(\bar{r})}{C(\bar{s})} \leq \frac{n\left(\sum_{v \in V} w_{\max}(v)\right)^\alpha}{\sum_{v \in V} w_{\max}(v)^\alpha}$$

Since the denominator is fixed, this ratio is maximized when the divider is minimized. The minimum value of $\sum_{v \in V} w_{\max}(v)^\alpha$ is $n((\sum_{v \in V} w_{\max}(v))/n)^\alpha$. Therefore, the ratio $C(\bar{r})/C(\bar{s})$ is bounded by n^α. □

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, as Observation 1 shows.

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

Proof: Let n locations be placed in the plane as illustrated in Figure 1. At every location there are k vertices colocated, hence there are kn vertices in total. There are $n/8$ locations in each of the levels A to F, and the remaining $n/4$ locations are situated in the set R along the right side of those levels. The locations 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 locations in R all have interpoint distance 1. For n sufficiently big, this graph is Euclidean.

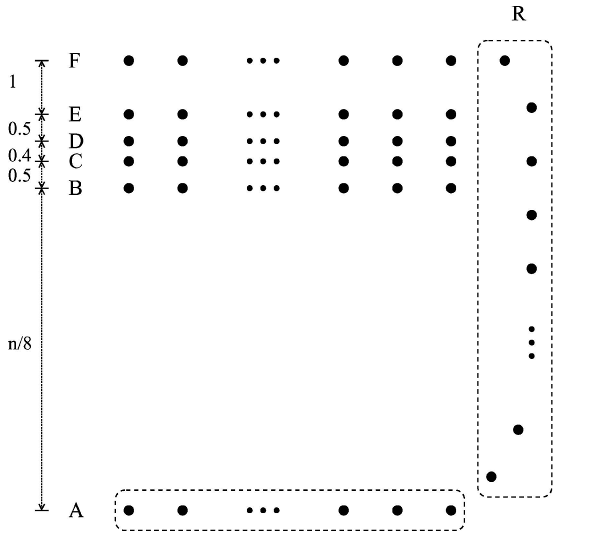

Figure 1. A graph where we can find a Nash equilibrium of cost $\Theta(n^\alpha)$ times the optimal cost.

Assume each node chooses radius $r_v = 1$. For $k = 1$, it is easy to verify that the induced graph $G(V, \bar{r})$ is strongly connected. For $k > 1$, each node can reach its neighbor locations (the locations that are within distance 1) directly, as well as via the $k-1$ vertices colocated at its own location. It can reach the vertices colocated at its own location via the k vertices of one of its neighbor locations. Therefore the transmission graph is strongly k-connected. The cost of the radius vector is $C(\bar{r}) = \sum_{v \in V} r_{v\alpha} = kn$. Now consider a second choice of radii $\bar{r}'$. We denote by r'_X the radius of any node v in the subset $X \subset V$.

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

The following proves that this choice of radii constitutes a Nash equilibrium. In order for a node to have a utility function that is bigger than $-M$, that is, for the induced graph to be strongly *k-connected*, each node has to choose a positive (nonzero) radius. If they chose a zero radius, they would reach only the $k-1$ nodes colocated at their location and therefore not have k internally disjoint paths to every other node. The nodes in B, C, 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, and the induced graph is strongly k-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 unilaterally decreased its radius, it would not reach the node in the level above through k internally disjoint paths anymore and the cycle would be broken. Assume for example, any one node v in level A chose a radius of 1 instead of $n/8$. Let us call the location of node v A1, and let us denote by B1 the location in level B that is directly above A1. The nodes at location B1 would then only be reached by the $k-1$ nodes colocated with v (nodes in C do not reach nodes in B), but not by v. Therefore, v would only have $k-1$ internally disjoint paths to the nodes at B1 and no incentive to decrease its radius to 1. For a node v at a location D1 in level D, it would get impossible to reach the nodes directly above it at location E1 in level E via k internally disjoint paths, if it decreased it radius. Node v would still reach E1 through $k-1$ internally disjoint paths, and the location C1 in level C directly below it, but the nodes in C1 would only lead back to D1 and never to E1. The same line of argument can be used for the nodes in level E.

The cost of this Nash equilibrium is $kn/8 * (n/8)^\alpha + O(kn) = O(kn^{\alpha+1})$ and we have therefore found a Nash equilibrium of cost $\Theta(n^\alpha)$ times the optimal cost, for any k.

Finding a Nash equilibrium. The following local improvement algorithm always leads to a Nash equilibrium in the STRONG k-CONNECTIVITY GAME.

Local improvement algorithm

1. Start with any choice of radii $\bar{r}^{(0)} = \bar{r}$ such that $G(V, \bar{r})$ is strongly k-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 decreases its radius to the smallest radius so that it can still reach every other vertex via k internally disjoint paths. 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 of the* STRONG k-CONNECTIVITY GAME *in polynomial time.*

Proof: Throughout the algorithm, the graph $G(V, \bar{r}^{(i)})$ stays strongly k-connected (see proof of Lemma 2). 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 $i+1, \cdots, n$. 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. $\square$

The local improvement algorithm described above is *incentive-compatible*. We call an algorithm incentive-compatible, if at every step in the algorithm, a vertex does what is in its best interest, that is, what maximizes its utility function, but we require that a vertex only acts when it is its turn to act. In our algorithm, if we take the start radius vector as given and force the vertices to follow the order given in step 2, the algorithm does exactly what each of the selfish vertices would do.[2]

There are different possibilities for finding a start vector $\bar{r}^{(0)}$ such that $G(V, \bar{r}^{(0)})$ is strongly k-connected, as needed in step 1 of the local improvement algorithm. Certainly the vector $\bar{r}$ with $r_v = \max\{w_e, e = \{v, v'\} \in H(V, E', \bar{w})\}$ yields a strongly *k-connected* transmission graph.

For the case of $k = 1$, an alternative is to start with a radius vector $\bar{r}$ such that $r_v = 0$ for all $v \in V$ and then go through all the nodes once in arbitrary order and let each of them choose the radius that maximizes its utility function, which will yield a good start vector. In other words, we start with a zero radius vector, choose an arbitrary order of the vertices and perform step 3 of the above algorithm twice.

There is a third way to compute a start vector, which has the advantage, that we can give an upper bound on the cost of the resulting Nash equilibrium. We will first consider the case $k = 1$ and then describe how the algorithm can be generalized for $k > 1$. For the construction of this start vector we follow an algorithm given in [13].

MST algorithm

1. Construct the undirected, complete graph K_V over V with edge weights $w^{\alpha}_{\{u,v\}}$ for all $u, v \,\varepsilon\, V$.
2. Find a minimum weight spanning tree T of K_V.
3. For all $v \,\varepsilon\, V$ let $r_v = \max\{w_{\{v,v'\}} \mid \{v, v'\} \varepsilon\, T\}$.

Corollary 1. *A Nash equilibrium of the* STRONG CONNECTIVITY GAME *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 [13] it is shown that $C(OPT) > C(T)$ and that

$$\begin{aligned} C\left(\bar{r}^{(0)}\right) &= \sum_{v\in V}\left(\max_{\{v'|\{v,v'\}\in T\}} w_{\{v,v'\}}\right)^{\alpha} \\ &= \sum_{v\in V}\max_{\{v'|\{v,v'\}\in T\}}\left(w_{\{v,v'\}}\right)^{\alpha} < \sum_{v\in V}\sum_{\{v'|\{v,v'\}\in T\}}\left(w_{\{v,v'\}}\right)^{\alpha} \\ &= 2 * C(T) < 2 * C(OPT). \end{aligned}$$

Clearly $C(\bar{r}) \le C(\bar{r}^{(0)})$, since throughout the local improvement algorithm, the radii only get smaller, and therefore also the cost goes down. □

In order to find a start vector that yields a strongly k-connected graph, for $k > 1$, one has to change step 2 as follows. Instead of finding a minimum weight spanning tree T of K_V, one needs to find a minimum weight k-connected subgraph G' of K_V. While this is NP-complete, there are several approximation algorithms. If β denotes the best known approximation for this problem, the cost of the resulting Nash equilibrium after using the local improvement vector with a start vector from such an algorithm will then be at most 2β times the optimal cost. For $k = 2$, $\beta = 3/2$, from the result of [10], and if the edge lengths satisfy the triangle inequality $\beta = 2 + 2(k-1)/n$, from [14]; note that the edge lengths satisfy the triangle inequality only for $\alpha = 1$, but not for $\alpha > 1$. When the edge lengths do not satisfy the triangle inequality, no constant factor approximations are known.

The* STRONG CONNECTIVITY GAME *for random points in the plane. In this section we consider the case $k = 1$ only. 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 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 $\bar{r}$ be any Nash equilibrium and let $\bar{s}$ be a radius vector that minimizes the cost $C(\bar{s})$. Then, (a) $C(\bar{s}) \ge \Omega(n/\log^{\alpha} n)$, with high probability, and (b) $C(\bar{r})/C(\bar{s}) \le n^{\alpha/2}\log^{\alpha} n$, with high probability.*

Proof: (a) 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 $\sum_{B'\varepsilon S}\sum_{i\varepsilon B'} s_i \ge \log n$. 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 $\log n$, $\sum_{B'\in S}\sum_{i\in B'} s_i \ge \log n$.

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

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

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

[2]Algorithmic incentive-compatibility is an extension of the notion of truthfulness from game theory to a distributed computation environment. Our definition is a relatively weak version as we limit the freedom of a vertex considerably by forcing it to only act when it is its turn.

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, (a) follows.

(b) 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 (a), $C(\bar{s}) = \Omega(n/\log^{\alpha} n)$ 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 6. *Let $\bar{r}^{(0)}$ be a radius vector that satisfies $r_v^{(0)} \geq max_w\{d(v,w)\}$ $\forall v \; \varepsilon \; V$. 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 5. Partition the $\sqrt{n} \times \sqrt{n}$ region, A, into n/m^2 blocks of dimensions $m \times m$ each, where m will be defined later. Observe that at the end of the local improvement algorithm, all except possibly one vertex in each block have radius $O(m)$. 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*.

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

4. Equilibria in the connectivity games

Existence of Nash equilibria. Figure 2 shows an instance of the CONNECTIVITY GAME without pure Nash equilibrium. The instance consists of three sources nodes A, B, and C and three sinks nodes A′, B′, and C′.

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_A, r_B, r_C for the three source vertices A, B, C respectively. We note immediately that $r_i \in \{1, 2\}$ for $i \in \{A, B, C\}$, as any radius $r_i < 1$ would mean that source i does not reach any other vertex in the graph and thus certainly will not reach its sink i', whereas any radius $r_i > 2$ cannot be part of a Nash equilibrium as reducing r_i to 2 would still allow source i to reach its sink i' with a better utility. The following implications hold:

$$r_A = 2 \implies r_C = 1 \tag{2}$$

$$r_B = 2 \implies r_A = 1 \tag{3}$$

$$r_C = 2 \implies r_B = 1 \tag{4}$$

$$r_A = 1 \implies r_C = 2 \tag{5}$$

$$r_B = 1 \implies r_A = 2 \tag{6}$$

$$r_C = 1 \implies r_B = 2 \tag{7}$$

Implications (2)–(4) hold because the source on the right-hand side of the implication can use the source on the left-hand side of the implication to reach its target; implications (5)– (7) hold because the source on the left-hand side would not reach its sink otherwise. Combining implications (2), (7), and (3), we obtain the following contradiction:

$$r_A = 2 \Longrightarrow r_C = 1 \Longrightarrow r_B = 2 \Longrightarrow r_A = 1.$$

Thus, no pure Nash equilibrium exists for this instance.

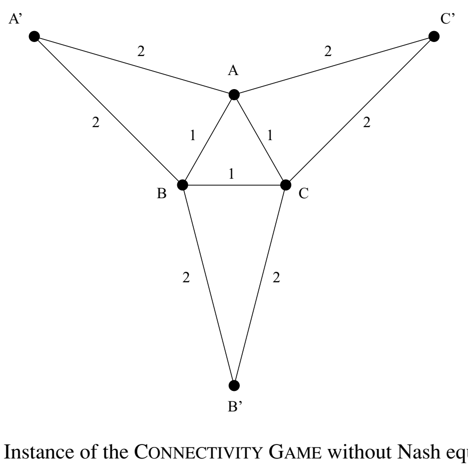

Figure 2. Instance of the CONNECTIVITY GAME without Nash equilibrium.

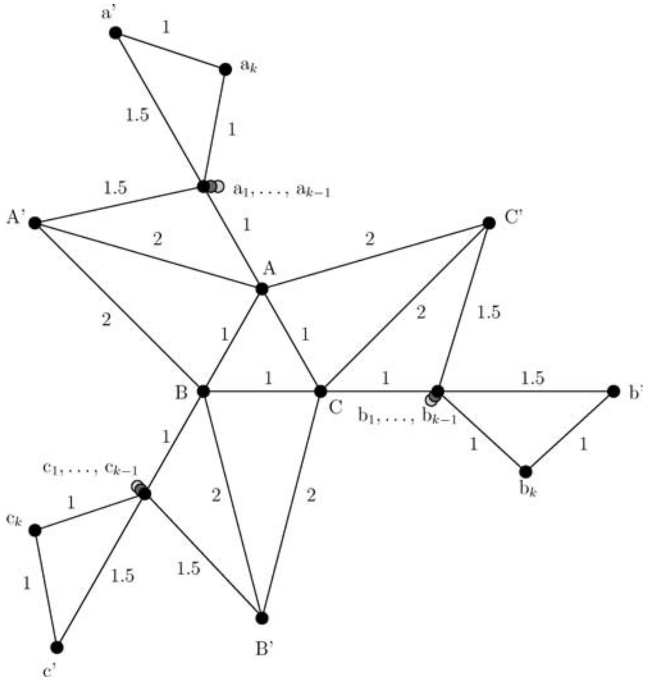

Figure 3. Instance of the k-CONNECTIVITY GAME without Nash equilibrium.

The previous result also holds for the k-CONNECTIVITY GAME for $k > 1$. Consider the graph given in figure 3. It is the same graph as the one in figure 2, but there are $k + 1$ additional *helper* nodes for each source node A, B, and C. The nodes $a_1, \ldots, a_k$ (short the a-nodes) are source nodes that have to reach the target a' via k internally disjoint paths; similarly for the b-, and c-nodes.

Observation 3. *No pure Nash equilibrium exists for the* k-CONNECTIVITY GAME *instance given in figure 3.*

Proof: Let us first determine the radii for the helper nodes. Since none of the nodes besides the a-nodes have an incentive to choose a radius large enough to reach a', the nodes $a_1, \ldots, a_{k-1}$ will choose their radius to be 1.5 in any Nash equilibrium, and a_k will choose the radius 1. With this choice of radii, all of the a-nodes reach a' via k internally disjoint paths and none of them has an incentive to decrease its radius. The same line of thought is used to determine the radii of the b-nodes and the c-nodes. Now consider the radii of nodes A, B, and C. Again, implications (2)–(7) hold. The first three implications hold because the source on the right-hand side of the implication can reach its target via the source on the left hand-side as well as over its helper nodes and therefore reaches its target via k internally disjoint paths. The remaining three implications hold, because otherwise the source on the right-hand side will reach its sink only over $k-1$ internally disjoint paths. □

Existence of approximate Nash equilibria. With respect to approximate Nash equilibria, a slight adaptation of the instance from figure 2 yields the following negative result for $k = 1$:

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 d, for an arbitrary $d > 1$. In geometric terms, this corresponds to making the three isosceles triangles longer. Each source node will now use a radius of either 1 or d. In any feasible combination of radii of the three sources as given in the proof of the previous lemma, a reduction from radius d to 1 will improve the utility of the corresponding source by a factor of $1/d^{\alpha}$. Thus, this modified instance does not have an β-approximate Nash equilibrium for any $\beta < d^{\alpha}$. Since we can choose d arbitrarily large, the corollary follows. □

It is instructive to look at the mixed Nash Equilibria for these games as well. Again, consider the instance in figure 2. In the mixed Nash equilibrium, let p denote the probability that node A chooses radius 1; so it chooses radius 2 with probability $1 - p$. By symmetry, this distribution is the same for each node. Recalling the definition of M from Section 2, it is easy to see that the utility of any node can be written as $-\,p(1-p) - Mp^2 - 2(1-p)$, for the case of $\alpha = 1$. This function is maximized for $p = \frac{1}{2(M-1)}$. As $M \to \infty$, the probability that each source chooses a radius of 2 tends to 1. Hence, it seems as if a radius vector with each source choosing radius 2 becomes a Nash equilibrium. However, for $M = \infty$, the utility functions are no longer continuous, and no mixed equilibrium exists.

Complexity of deciding whether a Nash equilibrium exists. 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.[3] 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 SATISFIABILITY to PURE NASH CONNECTIVITY WITH TRIANGLE INEQUALITY.

Definition 1. The problem MONOTONE 1-IN-3 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 SATISFIABILITY is NP-hard [11].

[3] Note that in this case we assume that the graph is a complete graph, with the edges satisfying the triangle inequality. In case an edge (u, v) is not present, a new edge of length equal to the length of the shortest path from u to v can be added.

Lemma 7. PURE NASH CONNECTIVITY WITH TRIANGLE INEQUALITY *is NP-hard.*

The proof of Lemma 7 is given in the appendix.

5. The reachability game

Existence of Nash equilibria. The REACHABILITY GAME does not necessarily have a pure Nash equilibrium, even for a 1-dimensional instance. The simple example of figure 4 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.

Observation 4. *The* REACHABILITY GAME *instance in figure 4 with* $\alpha = 1$ *does not have a pure Nash equilibrium.*

The proof is given in the appendix.

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 8. *For an instance* $\mathcal{P}$ *of random points in the plane, the* REACHABILITY GAME *with a fixed* 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.

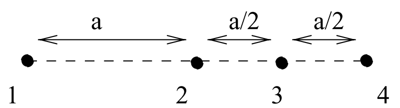

Figure 4. Instance of REACHABILITY GAME with no pure Nash equilibrium: there are $a/2-1$ vertices located together at points 1 and 3, three vertices at point 4 and a single vertex at point 2. The distances between the points are as shown. The value a is any number larger than 4.

6. Extensions for wireless networks with directional antennas

In this section, we study the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME for wireless networks with directional antennas. We denote by $N'(v)$ the set of neighbors of node v in the input graph H. Recall that we require the input graph H to have a direct link from each source to its link, thus in the STRONG CONNECTIVITY GAME it is a complete graph. In a directional antenna network, a strategy of a node v is to choose a set of nodes $N(v) \subset N'(v)$ and set up a directional antenna to each node in this set. The strategy vector $\bar{N}$ induces a graph $G(V, E)$ such that there is a directed edge $(u, v) \in E$ if $v \in N(u)$. $G(V, \bar{N})$ then denotes the graph induced by $\bar{N}$. The cost a node v incurs are $C(v) = \sum_{v' \in N(v)} w_{\{v,v'\}}$ for strategy $N(v)$. We can then define the cost of a strategy vector as $C(\bar{N}) = \sum_v \sum_{v' \in N(v)} w_{\{v,v'\}}$. The utility function remains as defined for the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME.

6.1. *The* DIRECTIONAL ANTENNA STRONG CONNECTIVITY GAME

Existence of Nash equilibria. In the DIRECTIONAL ANTENNA STRONG CONNECTIVITY GAME every instance has a Nash equilibrium, which follows from the following lemma.

Lemma 9. *Let* $G(V, E)$ *be a directed cycle in* $H(V, E', \bar{w})$ *that visits each node in* V *exactly once. For every edge* $(u, v) \in E$*, let* $v \in N(u)$*. Then the neighbor vector* $\bar{N}$ *constitutes a Nash equilibrium of the* DIRECTIONAL ANTENNA STRONG CONNECTIVITY GAME.

Proof: Every vertex reaches every other vertex, therefore there is no vertex with an incentive to reach an additional node, or a node that is further away. Moreover, no vertex can unilaterally decrease its cost, because using a shorter edge would not allow it to reach all other nodes. □

Finding a Nash equilibrium. Clearly, every permutation of the nodes yields a cycle and hence a Nash equilibrium for this game. However, these are not the only equilibria that exist and finding cycles does not necessarily happen in an incentive-compatible manner. The following algorithm is incentive-compatible and finds a Nash equilibrium for every instance of the DIRECTIONAL ANTENNA STRONG CONNECTIVITY GAME. The resulting Nash equilibrium is not necessarily a cycle.

Local improvement algorithm for DIRECTIONAL ANTENNA STRONG CONNECTIVITY GAME

1. Start with $\bar{N}^{(0)} = \bar{N}$ such that every node has a direct link to every other node. More formally, for each node v, let $N(v) = \{u \in V, u \neq v\}$.
2. Order the vertices arbitrarily as $v_1, \ldots, v_n$ and consider the vertices in this order.

3. In step $i = 1, \ldots, n-1$, the neighbor vector is $\bar{N}^{(i-1)}$ initially. Vertex v_i keeps the smallest link such that it still reaches all other vertices in $G(V, \bar{N}^{(i-1)})$ and stops paying for all others. $\bar{N}^{(i)}$ denotes the neighbor vector after v_i updates its neighbors.
4. In step n, vertex v_n stops paying for all links except the ones pointing to nodes that would otherwise have indegree 0.

Lemma 10. *The above algorithm leads to a Nash equilibrium of the* Directional Antenna Strong Connectivity Game.

Proof: We will first prove that the transmission graph remains strongly connected throughout the algorithm, which implies that no vertex ever has an incentive to pay a higher cost. In a second step we will prove that no vertex can decrease its cost further while still reaching all other nodes.

To prove that the transmission graph remains strongly connected throughout the algorithm, let us first consider steps $i = 1, \ldots, n-1$. In each of these steps, vertex v_i chooses the shortest link such that it can still reach all other nodes. At least one such link always exists, namely the link to vertex v_n, which still has a link to all nodes. Therefore the graph remains strongly connected until step $n-1$. In step n, vertex v_n keeps all the links to nodes that have indegree 0 otherwise. This allows vertex v_n to reach every node. If there is a node v that v_n does not reach directly, v has an incoming edge from some node v'. If v_n does not reach v' directly either, we can keep tracing back incoming edges until we arrive at a node u to which v_n has a direct link. It is not possible that we encounter node v again before we get to such a node u, since that would imply that v does not reach v_n.

Now let us prove that at the end of the algorithm, no node has an incentive to further decrease its cost. We will prove that in each step, a vertex decreases its cost as much as possible. As we only delete edges in the process of the algorithm and never establish new ones, a node will not develop an incentive to further decrease its cost after its turn.

For steps $1, \ldots, n-1$, we have to prove that the smallest one link that allows a vertex to reach all other nodes cannot be replaced by a set of links with even smaller cost. The proof will be by contradiction. Let e be the smallest link that vertex v can use to reach all other nodes, and let $f, f', \ldots$ be a set of links, that also allow v to reach all other nodes and for which it further holds that $w_f \ldots w_{f'} + \ldots < w_e$. At least one of the endvertices of $f, f', \ldots$ must be able to reach the endvertex of e, which we call v' (see figure 5). Let us say, without loss of generality, that this is the endvertex of f. But we know that using e allows v to reach all other nodes, therefore v' must be able to reach all nodes, which implies that edge f is sufficient to replace e. Since $w_f < w_e$, this contradicts our assumptions that e is the shortest single edge that allows v to reach all other nodes.

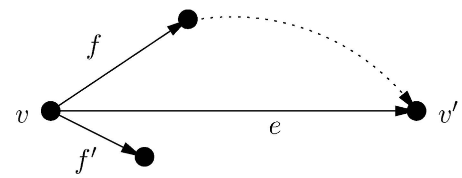

Figure 5. If edge e is the shortest single edge that allows vertex v to reach all other nodes, than no set of edges f, f', f'' of shorter length exists, that could replace e.

Finally, it is easy to see that vertex v_n cannot decrease its cost, since the only way to reach a node with indegree 0 is to establish a link to this node.

Quality of Nash equilibria

Lemma 11. *Let $\bar{N}$ be a neighbor vector that constitutes a Nash equilibrium in the* Directional Antenna Strong Connectivity Game *and let $\bar{M}$ be an optimal neighbor assignment. Then $C(\bar{N})/C(\bar{M}) \leq 2n$, if the edge lengths satisfy the triangle inequality.*

Proof: Let $w_{\max} = \max\{w_e, e \in E'\}$. Let us first prove the following claim: $C(\bar{N}) \leq (2n-2)\, w_{\max}$. Fix an arbitrary vertex v_0. Since v_0 has a path to every other vertex in $G(V, \bar{N})$, we can construct a rooted outtree $T \subset G(V, \bar{N})$, rooted at v_0: v_0 has a directed path of shortest hop length to all other nodes. There are $n-1$ edges in T, therefore $C(v_0) \leq C(T) \leq (n-1)\, w_{\max}$. All the remaining nodes would reach every other node if they established an edge to the root. Therefore, the costs of any node $v \neq v_0$ will be at most $C(v) = w\{v, v_0\} \leq w_{\max}$. Summing up the individual costs proves the claim. Since the triangle inequality holds in the input graph $H(V, E', \bar{w})$, $C(\bar{M}) \geq w_{\max}$. □

We have not found an example to show that the above is tight. However, the following instance has a ratio of $n/2$: Form two groups of vertices, each consisting of $n/2$ of the nodes. The two groups are far apart from each other, but the distances of the nodes within one group are very small. In this example, the social optimum has cost $2w_{\max}$. The worst Nash equilibrium is a cycle that alternatively visits nodes from each of the groups. It has cost $nw_{\max}$.

6.2. *The* Directional Antenna Connectivity Game

Existence of Nash equilibria

Lemma 12. *There is an instance of the* Directional Antenna Connectivity Game *for which no pure Nash equilibrium exists.*

The proof is given in the appendix.

Quality of Nash equilibria. An immediate conclusion from Lemma 12 is that the social optimum need not necessarily be a Nash equilibrium in the Directional Antenna Connec-

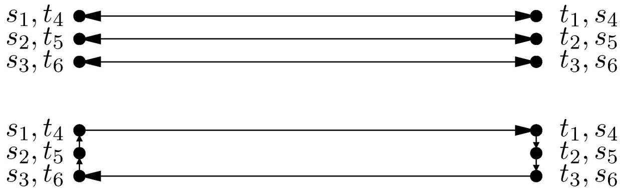

Figure 6. The upper figure shows the worst case Nash equilibrium for the given instance of the DIRECTIONAL ANTENNA CONNECTIVITY GAME. The lower figure shows the social optimum. The cost of the Nash is $O(j)$ times higher than the cost of the social optimum.

TIVITY GAME. However, in all instances where Nash equilibria exist, we can compare their cost to the cost of the social optimum, and this ratio does not depend on the interpoint distances, but only on the number of source-sink pairs.

Lemma 13. *Any Nash equilibrium for the* DIRECTIONAL ANTENNA CONNECTIVITY GAME *has cost at most j times the optimal cost, where j is the number of source-sink pairs* (s_i, t_i)*, if the edge lengths satisfy the triangle inequality.*

Proof: Let us denote by $v(s_i)$ and $v(t_i)$ the vertices s_i and t_i are located on, respectively, and let $w'_{\max} = \max \{w_{\{v(si),v(ti)\}} \mid 1 \le i \le j\}$. Let $\bar{M}$ be an optimal neighbor vector and $\bar{N}$ be a neighbor vector that constitutes a Nash equilibrium respectively. Then $C(\bar{M}) \ge w'_{\max}$ by triangle inequality. Further, $C(\bar{N}) \le jw'_{\max}$, since a Nash equilibrium can cost at most as much as when every source has a direct link to its target. Hence $C(\bar{N}) \le jC(\bar{M})$. □

The above bound is tight, as one can see with in the example given in figure 6.

7. Extensions for wireline networks

In order to study the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME for wireline networks, we adapt our games as follows. We again denote by $N'(v)$ the set of neighbors of node v in the input graph H and require H to have a direct link between each source and its sink. A strategy of a node v is to choose a set of nodes $N(v) \subset N'(v)$ and establish (or pay) a wire to each node in this set. The strategy vector $\bar{N}$ induces a graph $G(V, E)$ such that there is a bidirectional (or undirected) edge $\{u, v\} \in E$ if either $u \in N(v)$ or $v \in N(u)$. $G(V, \bar{N})$ then denotes the graph induced by $\bar{N}$. The cost a node v incurs are $C(v) = \sum_{v' \varepsilon N(v)} w_{\{v,v'\}}$ for strategy $N(v)$. We can then define the cost of a strategy vector as $C(\bar{N}) = \sum_v \sum_{v' \in N(v)} w_{\{v,v'\}}$. The utility function remains as defined for the STRONG CONNECTIVITY GAME and the CONNECTIVITY GAME.

7.1. The WIRELINE STRONG CONNECTIVITY GAME

In the WIRELINE STRONG CONNECTIVITY GAME, every Nash equilibrium is a spanning tree, but not every spanning tree

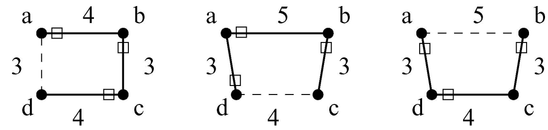

Figure 7. Three examples for the WIRELINE STRONG CONNECTIVITY GAME. A box by a node on an edge indicates, which node will pay for the wire.

is a Nash equilibrium. The first example in figure 7 shows a spanning tree that is not a Nash equilibrium: node *a* would increase its utility function, if it replaced its wire to node *b* by a wire to node *d*.

Existence of Nash equilibria. Every instance of the WIRELINE STRONG CONNECTIVITY GAME has a Nash equilibrium. This follows from the following lemma.

Lemma 14. *Let $G(V, E)$ be a minimum weight spanning tree of the input graph $H(V, E', \bar{w})$. For every edge $(u, v) \varepsilon E$, let either $u \varepsilon N(v)$ or $v \varepsilon N(u)$ (but never both). Then the wire vector $\bar{N}$ constitutes a Nash equilibrium of the* WIRELINE STRONG CONNECTIVITY GAME.

Proof: The graph $G(V, E)$ is connected, therefore no vertex has an incentive to pay for more or longer edges. Suppose node v has an incentive to cease paying for a wire e. Since $G(V, E)$ is a tree, this implies that v must want to pay for another wire e' instead (otherwise v would not reach every other node anymore) which is shorter than e and still allows v to reach all other nodes. But if such an edge e' exists, $G(V, E)$ is not a minimum weight spanning tree. □

Remark 1. Not every Nash equilibrium must yield a transmission graph that is a minimum weight spanning tree. The second example in figure 7 shows a Nash equilibrium which does not yield a minimum weight spanning tree. The third example in figure 7 shows, for the same instance, a Nash equilibrium which yields a minimum weight spanning tree.

Finding a Nash equilibrium. We can use any algorithm that finds a minimum weight spanning tree in order to find a Nash equilibrium of the WIRELINE STRONG CONNECTIVITY GAME. However, these algorithms are usually not incentive-compatible.

The following is an incentive-compatible algorithm that finds a Nash equilibrium for this game. We start with a neighbor vector, where every node pays a wire to every other node. This implies a fully connected transmission graph, where every wire is paid for twice, once by each of its endnodes. Then one after another the vertices get a chance to improve their utility function, which means they cease paying for wires that are either already paid for, or that they do not need in order to reach all other nodes. This algorithm always ends in a Nash equilibrium. In particular, it always ends in the Nash equilibrium, in which the node that gets to update its wires last, pays for a link to *each* of the other nodes, and none of the other nodes pay for any links. Therefore, the node that

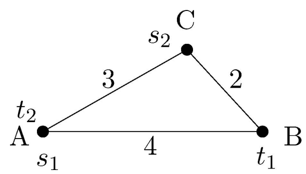

Figure 8. An instance of the WIRELINE CONNECTIVITY GAME without Nash equilibrium.

gets to change wires last, has the burden of the total resulting cost.

Quality of Nash equilibria

Lemma 15. *Let $\bar{N}$ be a neighbor vector that constitutes a Nash equilibrium in the* WIRELINE STRONG CONNECTIVITY GAME *and let $\bar{M}$ be an optimal neighbor assignment. Further let $w_{\min} = \min\{w_e, e \in E'\}$ and $w_{\max} = \max\{w_e, e \in E'\}$. If the edge lengths satisfy the triangle inequality,*

$$C(\bar{N})/C(\bar{M}) \leq \min\{w_{\max}/w_{\min}, n-1\}.$$

Proof: If $\bar{N}$ is a Nash equilibrium, then $G(V, \bar{N})$ must be a tree. Thus $C(\bar{N}) \leq (n-1)\, w_{\max}$. For $\bar{M}$ to be optimal, $G(V, \bar{M})$ must have at least $n-1$ edges, otherwise not every node would be able to reach every other node. This implies $C(\bar{M}) \geq (n-1)\, w_{\min}$, which proves $C(\bar{N})/C(\bar{M}) \leq w_{\max}/w_{\min}$. Since the triangle inequality holds for the input graph, we must also have that $C(\bar{M}) = \sum_{e \in G(V,M)} w_e \geq w_{\max}$, which proves $C(\bar{N})/C(\bar{M}) \leq n-1$.

7.2. The WIRELINE CONNECTIVITY GAME

Existence of Nash equilibria. Figure 8 shows an instance without Nash equilibrium for the WIRELINE CONNECTIVITY GAME.

Observation 5. *No pure Nash equilibrium exists for the instance of the* WIRELINE CONNECTIVITY GAME *given in figure 8.*

Proof: Obviously, no neighbor vector $\bar{N}$ that implies a transmission graph $G(\bar{N}, V)$ with zero, one or three links can be a Nash equilibrium. It is easy to check, that none of the six possibilities of neighbor vectors that lead to a transmission graph with two edges, constitute a Nash equilibrium either.□

Quality of Nash equilibria. *Comparing the cost of a Nash equilibrium in case it exists to the cost of a social optimum, yields the same result as we have found in Lemma 13 for the* DIRECTIONAL ANTENNA CONNECTIVITY GAME.

Lemma 16. *Any Nash equilibrium for the* WIRELINE CONNECTIVITY GAME *has cost at most j times the optimal cost, where j is the number of source-sink pairs (s_i, t_i).*

That this bound is also tight can be shown with an example similar to the one in figure 6.

8. Conclusions and open problems

We studied 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. Intermediate nodes in the games we study have to be paid by the source for forwarding each message. Design mechanisms or pricing schemes that include this price.
2. Develop a topology control protocol that is incentive-compatible.
3. Our work leaves a multitude of obvious problems open, among them: Is it *NP*-hard to decide whether an instance of the REACHABILITY GAME has a Nash equilibrium? Can the concept of *k*-connectivity be extended to the directional antenna and the wireline game? Are there instances where no approximate Nash equilibria exist in the DIRECTIONAL ANTENNA and WIRELINE CONNECTIVITY GAME?

Appendix

Proof of Lemma 7: Given a MONOTONE 1-IN-3 SATISFIABILITY instance I consisting of variables $x_1, \ldots, x_n$ 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 9 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 9) 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.[4] This completes the description of the PURE

[4]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.

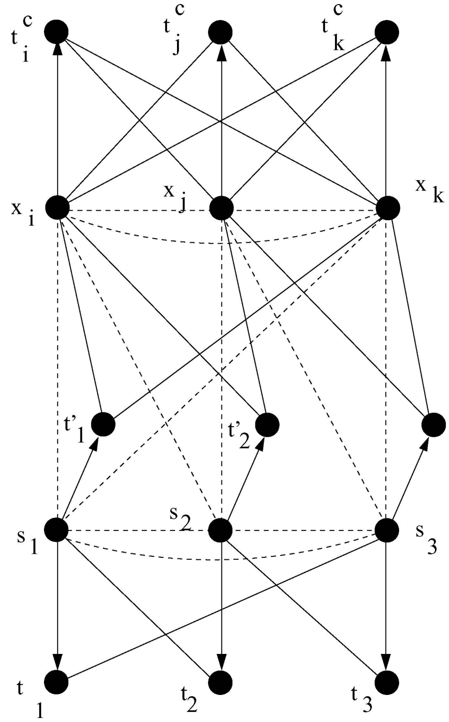

Figure 9. Clause gadget: source-sink relationships are indicated by arrows; dashed lines denote edges with weight 1, solid lines denote edges with weight 2.

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 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-in-3" 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 one 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 radius. 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,

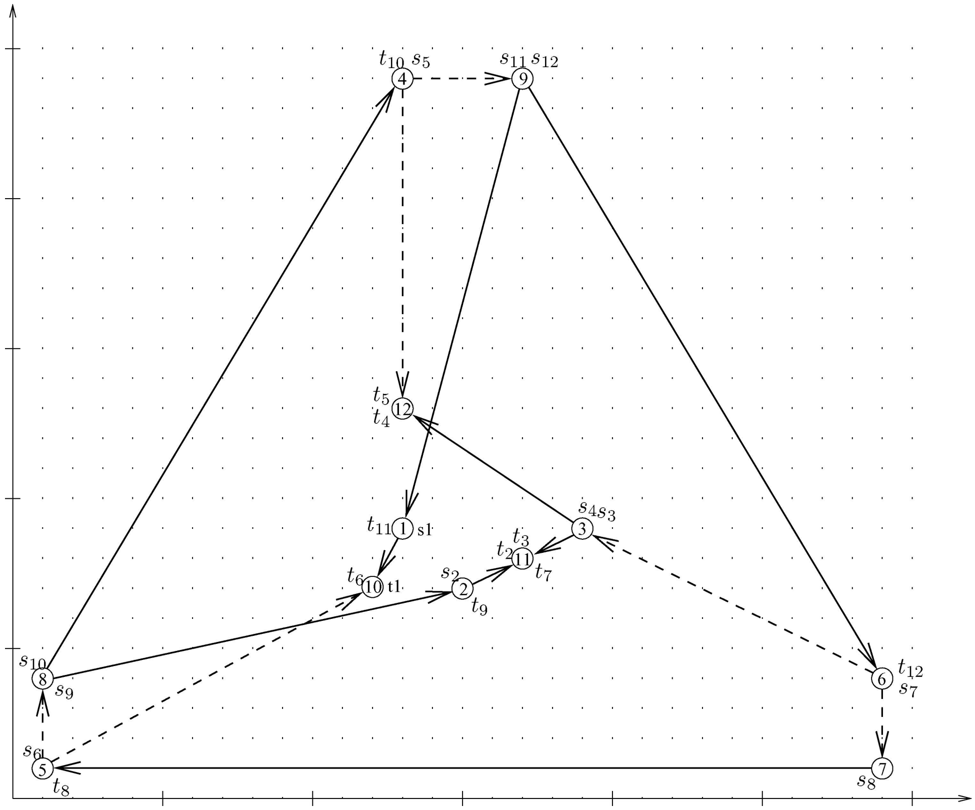

Figure 10. An instance of the DIRECTIONAL ANTENNA CONNECTIVITY GAME without Nash equilibrium.

we can assign a “1-in-3” satisfying truth assignment to the variables of I immediately. This completes our proof. □

Proof of Observation 4: 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 $U_1(a, \sigma_{-1}) \leq 2$, 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.

Proof of Lemma 12: Assume the instance in figure 10 does have a Nash equilibrium. Note that any node that needs to connect to another node, can do this either by establishing a direct link to its target, or by using any node that is closer than its target as an intermediate node. Some nodes need to connect to more than one target. In that case, the node will determine the set of links that is shortest in sum and allows it to reach all its targets.

Note that vertices 10, 11, and 12 do not need to connect to any node. Thus in any Nash equilibrium, those vertices have no outgoing links.

Vertex 1 needs to connect to vertex 10. Since vertex 10 is also its closest neighbor, vertex 1 will establish a direct link to vertex 10 in any Nash equilibrium. Equivalently, vertex 2 will establish a direct link to vertex 11, which is its only target as well as its closest neighbor.

Vertex 3 needs to reach vertex 11 as well as vertex 12. In order to reach vertex 12, vertex 3 could use 11, 2, 1, or 10 as intermediate nodes. We have already shown though, that vertices 10 and 11 never have an outgoing edge in a Nash equilibrium, and further that the only outgoing edges of vertices 1 and 2 go to 10 and 11 respectively. Thus, none of this nodes can be used as an intermediate node to reach vertex 12 and therefore vertex 3 must establish a direct link to vertex 12. Since vertex 12 does not have any outgoing links in a Nash equilibrium, the established link from vertex 3 to vertex 12 will not help vertex 3 to reach its second target, vertex 11. Hence it will establish a second direct link to vertex 11, which is its closest neighbor.

Let us call the vertices 4–9 the outer vertices, and the remaining ones the inner vertices. We have established all the links for the inner vertices. All those vertices will establish those links in any Nash equilibrium, independent of the links of the outer vertices. Note that none of the inner vertices established a link to an outer vertex.

Vertices 4, 5, and 6 need to reach vertices 12, 10, and 11 respectively. Vertex 4 and 5 will either reach their target directly, or via their closest neighbor 9 and 8 respectively, since those are the only nodes that are closer to them than their target nodes. Vertex 6 can reach its target either via the vertex 3 which has a link to vertex 11, or via vertex 7, which is its closest neighbor. Thus each of the vertices 4, 5, and 6 have two possibilities to reach its targets. We cannot rule out one of this two possibilities at this point.

Now let us consider vertex 7, which needs to reach vertex 5. Its closest neighbor is vertex 6, but vertex 6 either has a link to an inner vertex, which never reaches any outer vertex, or a link back to 7 itself. Thus vertex 6 cannot be used as an intermediate link. None of the inner vertices can be used as an intermediate link, so the next closest vertex is vertex 5 itself and 7 establishes a direct link to vertex 5.

Using the same arguments, vertex 8 has to establish a direct link to vertex 4, and vertex 9 has to establish a direct link to vertex 6. But vertices 8 and 9 also need to connect to vertices 2 and 1 respectively. In order for vertex 8 to reach vertex 2, it could either use the direct link to vertex 4, or establish a link to one of the nodes 5, 10, or 1. Node 10 does not have any outgoing links though, and node 1 only leads back to 10. Node 5 leads to 10 as well, or back to 8. Therefore node 8 could either use the existing link, or establish a direct link to 2. In order for vertex 9 to reach vertex 1, it could also use the existing link to node 6 or establish a link to one of the nodes 4, 12, 3, or 1. Again we can rule out nodes 4, 12, and 3 as intermediate links. Now that we have listed all the possible links for all the nodes, we find that in any Nash equilibrium, every node but node 8 never has an incentive to establish a link to node 2. Therefore node 8 cannot use its existing link to node 4 in order to reach 2. Similarly, no node but node 9 ever has an incentive to establish a link to node 1. Therefore node 8 and node 9 will have to establish a direct link to node 2 and 1 respectively. The two links of vertices 8 and 9 cannot be replaced by any single link to an intermediate node, because no other vertex will ever have a link to any of these nodes.

Finally we have established the link for all nodes except node 4, 5, and 6. For those nodes the following implications hold.

4 has link to 9 $\Rightarrow$ 5 has link to 8.
5 has link to 8 $\Rightarrow$ 6 has link to 7.
6 has link to 7 $\Rightarrow$ 4 has link to 12.
4 has link to 12 $\Rightarrow$ 5 has link to 10.
5 has link to 10 $\Rightarrow$ 6 has link to 3.
6 has link to 3 $\Rightarrow$ 4 has link to 9.

These implications lead to a contradiction, and therefore, this instance does not have a Nash equilibrium for the DIRECTIONAL ANTENNA CONNECTIVITY GAME.

Acknowledgments

We are very grateful to Madhav Marathe for his constant encouragement and valuable suggestions. We also thank the anonymous referees for their comments.

References

[1] L. Anderegg and S. Eidenbenz, Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents, in: *Proceedings of the Ninth Annual International Conference on Mobile Computing and Networking (MOBICOM)* (San Diego, CA, 2003) pp. 245–259.

[2] E. Ansheleich, A. Dasgupta, E. Tardos and T. Wexler, Near-optimal network design with selfish agents, in: *Proceedings of the 35th Symposium on Theory of Computing* (2003) pp. 511–520.

[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, ACM/Kluwer Mobile Networks and Applications (MONET) 8(5) (2003) 579–592.

[5] H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations, Annals of Mathematical Statistics 23 (1952) 493–509.

[6] 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 73B(2) (1969) 91–118.

[7] S. Doshi, S. Bhandare, and T. Brown, An On-demand minimum energy routing protocol for a wireless ad hoc network, in: *ACM Mobile Computing and Communications Review* 6(3) (2002).

[8] S. Eidenbenz, V.S. Anil Kumar and S. Zust. Equilibria in topology control games for ad hoc networks, in: *Proceedings in DIALM-POMC Mobicom*, (San Diego, CA, 2003) 2–11.

[9] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou and S. Shenker, On a network creation game, in: *Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing (PODC)* (2003) pp. 347–351.

[10] G.N. Frederickson and J. Jájá, On the relationship between the biconnectivity augmentation and traveling salesman problems, *Theoretical Comput. Sci.* 19 (1982) 189–201.

[11] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of* NP-*completeness*. (Freeman, San Francisco, 1979).

[12] M. T. Hajiaghayi, N. Immorlica and V. S. Mirrokni. Power optimization in fault-tolerant topology control algorithms for wireless multi-hop

networks, in: *Proceedings of the Ninth Annual International Conference on Mobile Computing and Networking (MOBICOM)* (San Diego, CA, 2003) pp. 300–312.
[13] L.M. Kirousis, E. Kranakis, D. Krizanc and A. Pelc, Power consumption in packet radio networks, *Theoretical Computer Science* 243 (2000) 289–305.
[14] S. Khuller and B. Raghavachari, Improved approximation algorithms for uniform connectivity problems, J. Algorithms 21 (1996) 434–450.
[15] A. Mas-Collel, W. Whinston and J. Green, *Microeconomic Theory*. (Oxford University Press, 1994).
[16] C. Perkins, Ad hoc networking: An introduction, in: C. Perkins, ed. *Ad Hoc Networking* (2001) pp. 314–318.
[17] R. Rajaraman, Topology control and routing in ad hoc networks: A survey. SIGACT News 33 (2002) 60–73.
[18] T. Roughgarden, *Selfish Routing*. Ph.D. Thesis, Cornell University (2002).
[19] T. Roughgarden, Designing networks for selfish users is hard, in: *Proceedings of the 42nd Symposium on the Foundations of Computer Science* (2001) pp 472–481.
[20] R. Roychoudhury, X. Yang, R. Ramanathan and N. Vaidya, Using directional antennas for medium access control in ad hoc networks, in: *Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking (MOBICOM)* (2002) pp. 59–70.

Stephan Eidenbenz is a technical staff member in Discrete Simulation Sciences (CCS-5) at Los Alamos National Laboraotry. He received his Ph.D. in Computer Science from the Swiss Federal Institute of Technology, Zurich, Switzerland in 2000. Stephan's research covers areas in approximability, algorithms, computational geometry, computational biology, large-scale discrete simulation, selfish networking, efficient networking, protocol design and optimization.
E-mail: eidenben@lanl.gov

V. S. Anil Kumar is currently an Assistant Professor in the Dept. of Computer Science and a Senior Research Associate at Virginia Bioinformatics Institute, Virginia Tech. Prior to this, he was a technical staff member in Los Alamos National Laboratory. He received a Ph.D. in Computer Science from the Indian Institute of Science in 1999. His research interests include approximation algorithms, mobile computing, combinatorial optimization and simulation of large socio-technical systems.
E-mail: akumar@vbi.vt.edu

Sibylle Zust received her Masters degree in mathematics from ETH Zurich in Switzerland in 2002. She wrote her diploma thesis at the University of Copenhagen in Denmark. Sibylle Zust spent two and a half years (2002–2005) as a graduate research assistant at the Los Alamos National Laboratory in New Mexico, USA, where she worked on algorithmic aspects of game theory and scheduling problems. She now works for an insurance company in Zurich, Switzerland.
E-mail: zust@lanl.gov