Topology Design of Communication Networks: A Game-Theoretic Perspective

Amir Nahir, Ariel Orda, Ari Freund

***Abstract*—We study the performance of non-cooperative networks in light of three major topology design considerations, namely the price of establishing a link, path delay, and path proneness to congestion, the latter being modeled through the "relaying extent" of the nodes. We analyze these considerations and the tradeoffs between them from a game theoretic perspective, where each network element attempts to optimize its individual performance.**

We show that for all considered cases but one, the existence of a *Nash equilibrium point* is guaranteed. For the latter case, we indicate, by simulations, that practical scenarios tend to admit a Nash equilibrium. In addition, we demonstrate that the *price of anarchy*, i.e., the performance penalty incurred by non-cooperative behavior, may be prohibitively large; yet, we also show that such games usually admit at least one Nash equilibrium that is system-wide optimal, i.e., their *price of stability* is 1. This finding suggests that a major improvement can be achieved by providing a central ("social") agent with the ability to impose the initial configuration on the system.

***Index Terms*—Communication Networks, Game Theory.**

I. Introduction

Topology design of computer networks focuses on finding the network configuration with the best possible performance given some optimization criteria [8]. Classical works on this subject aimed at finding optimal network topologies with respect to criteria such as fault tolerance [31], reliability [29] and delay [26].

Many network settings involve selfish agents making decisions to maximize their own profit. These non-cooperative networks give rise to game-theoretic analysis to quantify the stability of the emerging networks and the performance associated with them. Prior studies on this subject, such as [2], [23] and [19], mainly investigated the structure of the network operating points i.e., the Nash equilibria of the respective games. Such equilibria are inherently inefficient [12] and, in general, exhibit suboptimal network performance. In order to quantify this inefficiency, two conceptual measures have been proposed in the literature. The first, termed the *price of anarchy* [25], is the ratio between the *worst* Nash equilibrium and the social optimum. The second, termed the *price of stability* [4] is the ratio between the *best* Nash equilibrium and the optimum.

In this work, we perform a rigorous analysis of topology design issues of non-cooperative networks. We study the effects of several topology design considerations on the performance of non cooperative networks. These considerations are the price of establishing a (directed) link between network elements, the delay of the resulting routing paths, and the "relaying extent" of nodes along the routing paths, where the latter is a measure of the path's proneness to congestion. We analyze the effect of each of these considerations, as well as tradeoffs amongst them, on the resulting network topology and its implied performance. As mentioned, we conduct our investigation within the realm of noncooperative games; however, this requires us to consider also the more basic framework of the (system) optimization problems.

First, we consider the price of establishing a link. We consider *directional* links, and, in addition, we allow link prices to assume *different values* throughout the network. In addition, we analyze the link establishment price as a single consideration before moving to analyze the tradeoff between it and the delay consideration.

Delay is our second design consideration, and, similarly to [15], it is estimated by counting the hops of the resulting routing paths.

Finally, our third design consideration is termed the "relaying extent" of the nodes along the routing paths, and is a function of the nodal in-degree. A higher relaying extent of a node typically implies a higher number of nodes using it as a relay, resulting in higher congestion at that node. In some cases, they may affect the bandwidth allocated to a specific user when sending its data along that path, while in other cases this may increase the probability of packet loss. We capture both of these cases through two separate cost functions.

While we focus our work on network design, our findings are relevant also to topology control in wireless networks. Several studies have addressed such "topology control games" [14], [16], [22], however in the context of omnidirectional antennas, i.e., connectivity is (solely) derived by the radius covered by the power level set by the node. Our results, on the other hand, are relevant to the case of directional antennas, where a node can choose its neighbors on a per-node basis.

Our findings can be summarized as follows:

- For all considered cases but one, we establish the existence of a Nash equilibrium.
- For the remaining case, we provide a generic counter-example. Furthermore, we show, through simulations, that practical scenarios tend to admit at least one Nash equilibrium with a reasonable price of anarchy.
- We show that, typically, the considered games may perform poorly from a systemwide perspective, as exhibited by large values of the price of anarchy. Moreover, we

Conference version appeared in the Proceedings of IEEE INFOCOM 2009.

Amir Nahir is with the Department of Computer Science, Technion, Israel Institue of Technology, Haifa 32000, Israel, Email: nahira@cs.technion.ac.il

Ariel Orda is with the Department of Electrical Engineering, Technion, Israel Institue of Technology, Haifa 32000, Israel, Email: ariel@ee.technion.ac.il

Ari Freund is with Google Israel, Haifa, Israel, Email: arief@google.com

show that such poor operating points may be reached through a "natural" course of the network game.
- On the other hand, we show that, in many cases, the price of stability is 1, i.e., there is a Nash equlibrium that is systemwide optimal. Therefore, efficient performance can be often achieved by just controlling the initial configuration of the system.
- In some of the considered cases we show that, similarly to [15], Nash equilibria are obtained in very simple topologies, such as a complete graph (clique) or a ring.
- Moreover, when the relaying extent is the only design consideration, we show that the system optimum too is obtained in such a simple topology, namely a ring. This finding may suggest that a practically appealing model should address the relaying extent only in conjunction with additional design considerations. We achieve this by adding delay as a second consideration.

The rest of this paper is organized as follows. After discussing related work in the next section, we formulate the model and terminology in Section III. Sections IV-X deal with the different topology design considerations we study, as well as the tradeoffs between them. Finally, conclusions are presented in section XI.

II. Related Work

Game theoretic models have been employed in various networking contexts, such as flow control [2], [17], routing [3], [23] and bandwidth allocation [19]. These studies mainly investigated the structure of the network operating points i.e., the Nash equilibria of the respective games. Such equilibria are inherently inefficient [12] and, in general, exhibit suboptimal network performance. As a result, the question of how much worse the quality of a Nash equilibrium is with respect to a centrally enforced optimum has received considerable attention e.g., [18], [27], [28]. In order to quantify this inefficiency, two conceptual measures have been proposed in the literature. The first, termed the *price of anarchy* [25], corresponds to a worst-case analysis and is the ratio between the *worst* Nash equilibrium and the social optimum. The second, termed the *price of stability* [4], is the ratio between the *best* Nash equilibrium and the optimum, and it quantifies the degradation in performance when the solution is required to be stable (i.e., with no agent having an incentive to independently defect out of it once being there).

Bottleneck models have many practical applications in the context of network design and management. Such models can be used to reflect the remaining battery life in a wireless network [33], to minimize the usage of loaded buffers in traffic engineering [8], or to avoid congested links when routing traffic [32]. In the context of game theory, bottleneck models have also been studied in [7].

Prior game-theoretic work on topology design has focused on the tradeoff between the price of establishing a link and the delay of the implied routing paths. In [15], the authors presented a fundamental model where players set up links to construct an undirected graph, trading off link establishment costs with path lengths. They focused on the case of *undirected* networks, with *homogeneous* link prices, that is: once a link has been established between nodes v_i and v_j, either node could use that link to transmit data to its adjacent node; furthermore, the price of establishing a link was assumed identical throughout the network. In [1], the authors followed the model of [15] and improved some of the results.

In [10], the authors used a model similar to the one above, but demanded bilateral agreement of both sides for every link establishment, that is, a link is established only if both nodes choose to establish it. In [13], the authors followed up and improved some of the results of [10] and [15]).

In [9], the authors extended [15] by limiting the cost of non-connectivity to a finite value.

In [5], [6], the authors studied a similar problem, where players aspire to connect to only some of the other players, and may buy non-adjacent links.

Finally, the authors of [21] looked at a similar problem from a peer-to-peer network perspective, where the links are directional and link delay may be arbitrary.

In this study, we consider directional links, and furthermore, investigate several novel design criteria.

III. The Model

Our unified network model is based on a set of N network elements. We assume that network elements establish links among themselves during the network's setup phase, and continue using these links when transmitting throughout the remainder of the network's lifetime. We refer to the network elements as *nodes*, and denote them by $V = \{v_0, v_1, \ldots, v_i, \ldots, v_{N-1}\}$. *Edges* are used to represent links between nodes, and are denoted by a set E. Note that the existence of a link from v_i to v_j means that v_i can send data directly to v_j, but it does not imply that v_j can send data directly to v_i, therefore the network's topology is represented by a directed graph, denoted as G. We use $\eta^{in}(v_i)$ to denote v_i's in-degree, and $Adj_G(v_i)$ to denote v_i's adjacent nodes, i.e.,

$$Adj_G(v_i) = \{v_j | < v_i, v_j > \in E\}.$$

A non-cooperative game [24] is comprised of three components, namely players, strategies and costs. In the non-cooperative games we analyze, the network elements (nodes) are the players. The strategies, s_i, of each player v_i, are the sets of links that it may choose to establish, i.e., each player v_i chooses its direct neighbors, $Adj_G(v_i)$. A specific choice of the players startegies, $s = \bigcup_{i=0}^{i=N-1}\{s_i\}$, defines the *strategy profile*, which induces the network's topology G. Given the network's topology G, we denote by $l_G(v_i, v_j)$ the *routing path* from v_i to v_j. Routing paths are chosen so as to (self-) optimize some design considerations, which are captured by a cost function. More specifically, in the games we investigate, the goal of each node (player) is to achieve full connectivity with all other nodes, and, under this constraint, optimize the design considerations. We term this class of games as *connectivity games*.

As mentioned, we study the impact of three design considerations. The first is the price of establishing a link. To model it, we define a weight function $W : V \times V \to \Re$, i.e., $W(v_i, v_j)$

defines the price (to v_i) of establishing a link from node v_i to node v_j. The second design consideration is the delay of the resulting routing path, which is captured through the number of hops along the routing path. The third design consideration is the relaying extent, which is captured by the in-degree of the nodes along the path. Specifically, in-degrees quantify the relaying extent in two possible ways, namely the *product* of the in-degrees of nodes along the path, or the *bottleneck* of these values. When nodal in-degree reflects the amount of traffic expected to go through a node, a high in-degree may indicate a high probability of packet loss (due to congested buffers). In such a case, the product value of nodal in-degrees along a path represents the probability of a succesful transmission of a packet through a congested relay. Whereas the bottleneck value may be chosen when nodes share the bandwidth of the relay. In both cases, a lower value of the relaying extent implies a better routing path.

Each of the three design considerations, namely establishment price, delay and relaying extent, translates into a corresponding *nodal cost function*, which captures the performance of the node. Specifically, with each player $v_i \in V$, we associate a (non-negative) cost value $c(v_i)$, which accounts for various factors, depending on the design consideration, e.g., the node's power consumption, its distance from some other node, path interference, etc. The precise definition of the node's cost value is detailed in the following respective sections. Each player (node) strives to minimize its cost $c(v_i)$.

Player v_i's *best-response move* is a strategy which, given the strategies of all other players, yields the lowest value to $c(v_i)$. A network topology G is said to be at *Nash equilibrium* if each player considers its chosen strategy to be the best under the given choices of other players.

The selfish behavior of the players typically leads to network-wide inefficiency. We quantify this inefficiency through the ratio between the cost of the worst possible Nash equilibrium topology and the cost of an optimal solution. In keeping with common terminology [18], [25], this ratio is called the *price of anarchy* and it quantifies the "penalty" incurred by lack of cooperation (or coordination) among the players in a noncooperative game. We also consider the *price of stability* [4], which is the ratio between the cost of the best Nash equilibrium and the cost of an optimal solution and quantifies the inefficiency of the noncooperative game in cases where the initial operating point of the system can be chosen by a social agent.

Note that, in order to quantify the price of anarchy and price of stability, the related system optimization problems need to be defined. In general, the goal of system optimization is to minimize the total cost, namely $C(G) = \sum_{v_i \in V} c(v_i)$. An exception is whenever the nodes' costs depend on a bottleneck function, as when a bottleneck relaying extent is considered; there, total cost is defined as the networkwide bottleneck, i.e., $C(G) = \max_{v_i \in V} c(v_i)$.

IV. Link Prices (Solely)

The first game we consider entails each player v_i with the total price of the links it chooses to establish, i.e., $c(v_i) = \sum_{v_j \in Adj_G(v_i)} W(v_i, v_j)$. In case player v_i fails to connect to one (or more) of the other nodes, its cost is infinite (i.e., $c(v_i) = \infty$). We term the respective game as *the connectivity game with link prices*. We show that this game has a Nash equilibrium point. In addition, we show that the game's price of stability [4] is 1. On the other hand, we show that this game has a high price of anarchy, and that an iterative best-response move by each player may lead to that state.

Theorem 1: The connectivity game with link prices has a Nash equilibrium.

Proof: By proving that this game is an exact potential game [20]. To that end, we define the following (potential) function Φ: $\Phi(G) = \sum_{v_i \in V} c(v_i)$. One needs to prove then, that when a player v_i improves its cost, the potential decreases by the exact same amount. Let G be a valid topology (i.e., all players have routing paths to all other players) induced by a strategy profile s. Assume player v_i changes its strategy from s_i to s'_i to improve its cost. It is easy to see that when starting from a valid strategy profile, player v_i must remove at least one link to improve its cost. Let G' be the topology induced by the strategy profile after player v_i has changed its strategy. For any player $v_j \neq v_i$, since G is a valid topology, it holds that v_j has a routing path to v_i. Note that when player v_i changes its strategy (i.e., adds or removes links), it cannot affect v_j's connectivity to itself. It follows that v_j is connected to v_i in G'. Since v_i improved its cost, it is clear that v_i has routing paths to all other nodes in G', and so v_j has routing paths to all nodes in G'. We conclude that for any player $v_j \neq v_i, c_G(v_j) = c_{G'}(v_j)$. Hence,

$$\Phi(G') - \Phi(G) = \sum_{v_j \in V} c_{G'}(v_j) - \sum_{v_j \in V} c_G(v_j) =$$
$$c_{G'}(v_i) + \sum_{v_j \neq v_i} c_{G'}(v_j) - c_G(v_i) - \sum_{v_j \neq v_i} c_G(v_j) =$$
$$c_{G'}(v_i) - c_G(v_i).$$

In view of [20], the theorem follows. ■

Next, we prove an important property on the correlation between the Nash equilibrium strategy profiles of the connectivity game with link prices and the local optima of the corresponding (system) optimization problem. We will later rely on this lemma to prove our results on the price of stability and price of anarchy of the game.

Lemma 1: Let G be a topology. G defines a local optimum of the connectivity problem with link prices if and only if G is at Nash equilibrium for the connectivity game with link prices.

Proof: Recall that optimum is measured with respect to the total costs of all players, i.e., the optimization objective function is $C(G) = \sum_{v_i \in V} c(v_i)$.

In one direction, let G be a local optimum for the optimization problem. It is clear that G is a valid topology. Assume, by negation, that G is not at Nash equilibrium. By definition, there exists at least one player that can unilateraly improve its cost. Let v_i be such a player. Let G' be the topology reached after v_i played its best response to G. As shown in the proof of Theorem 1, for every player $v_j \neq v_i, c_G(v_j) = c_{G'}(v_j)$. In addition, since v_i played its best response move, $c_{G'}(v_i) < c_G(v_i)$. Hence, $C(G') < C(G)$ in contradiction to G's local optimality.

The other direction holds by definition. ■

Theorem 2: The price of stability for the connectivity game with link prices is 1.

Proof: By proving that there exists an optimum point that is at Nash equilibrium.

As the global optimum is also a local one, Theorem 2 follows directly from Lemma 1.[1] ■

In order to analyze the game's price of anarchy, several definitions are required. A player's *minimal connectivity price* is the minimal price a player must pay in order to establish any link. Formally, $W^i_{min} = \min_{v_j \in V, v_j \neq v_i}\{W(v_i, v_j)\}$. A player's *maximal connectivity price* is the price a player must pay in order to establish the most costly link. Formally, $W^i_{max} = \max_{v_j \in V}\{W(v_i, v_j)\}$ (note that this is not the highest value $c(v_i)$ can take since it refers to a single link only). A game's *minimal connectivity price* is the minimal price required so that some player may establish a link to some other player. Formally, $W_{min} = \min_{v_i \in V}\{W^i_{min}\}$. A game's *maximal connectivity price* is the price of the most costly link (i.e., it's the highest value $W(\cdot,\cdot)$ takes). Formally, $W_{max} = \max_{v_i \in V}\{W^i_{max}\}$.

Theorem 3: The price of anarchy for the connectivity game with link prices is $\Theta(\frac{W_{max}}{W_{min}})$.

Proof: By establishing the corresponding lower and upper bounds, in the following lemmas.

Lemma 2: The price of anarchy for the connectivity game with link prices is $O(\frac{W_{max}}{W_{min}})$.

Proof: By showing a lower bound on the cost of the optimal topology and an upper bound on the cost of any Nash equilibrium.

First, we note that each player must establish at least one out-going link in order to have any kind of connectivity with the rest of the network. It follows that the cost of each player is at least W_{min}. Hence, the cost of the optimal topology is at least $N \cdot W_{min}$.

Next, we consider the problem of the worst Nash equilibrium. It follows from Lemma 1 that this Nash equilibrium is also the worst local optimum for the optimization problem.

The following lemma proves a general property of directed graphs. Using this property, we will establish the desired upper bound.

Lemma 3: Let $G = (V, E)$ be a directed clique of N nodes. Let $\hat{G} = (V, \hat{E})$ be a subgraph of G, such that $\hat{G}$ is strongly connected, and, in addition, for any edge $e_i \in \hat{E}$, $(V, \hat{E}\backslash\{e_i\})$ is not strongly connected. Then, $|\hat{E}| \leq 2 \cdot (N-1)$.

Proof: Assume, by negation, that there exists a strongly connected graph $\hat{G} = (V, \hat{E})$, such that for any edge $e_i \in \hat{E}$, $(V, \hat{E}\backslash\{e_i\})$ is not strongly connected, and, in addition, $|\hat{E}| > 2 \cdot (N-1)$. Let us now run Tarjan's algorithm for strongly connected components [11]. It holds, from the correctness of Tarjan's algorithm and our assumption, that the algorithm will return a single strongly connected component, which includes all of V. However, Tarjan's algorithm is comprised of two runs of the Depth-First-Search algorithm [11]. Since in each run the DFS algorithm traverses $N-1$ edges, Tarjan's algorithm traverses at most $2 \cdot (N-1)$ edges. It follows that some of the edges in $\hat{E}$ can be removed without damaging $\hat{G}$'s connectivity, in contradiction to our assumption. This completes the proof of Lemma 3. ■

Corollary 1: The cost of the worst Nash equilibrium for the connectivity game with basic link prices is $\leq 2 \cdot (N-1) \cdot W_{max}$. ■

Therefore, the cost of any Nash equilibrium is upper bounded by $2 \cdot (N-1) \cdot W_{max}$, while the optimal configuration is lower bounded by $N \cdot W_{min}$, yielding a bound on the price of anarchy of $O(\frac{W_{max}}{W_{min}})$. This completes the proof of Lemma 2. ■

Lemma 4: The price of anarchy for the connectivity game with link prices is $\Omega(\frac{W_{max}}{W_{min}})$.

Proof: By proving that any directed ring topology is at Nash equilibrium and providing an example.

Lemma 5: Every directed ring topology is at Nash equilibrium.

Proof: Assume, by negation, that there exists a directed ring topology G, such that G is not at Nash equilibrium. Since G is not at Nash equilibrium, it holds that there exists at least one player that can unilaterally reduce it cost. Let v_i be such a player, and let $v_j \in Adj_G(v_i)$. Since G is a directed ring, it holds that v_i is the only node directly connected to v_j. And so, it holds that v_i cannot remove its link with v_j, since any topology resulting from such a move will not be strongly connected. Hence, v_i's best response move can only include addition of links. Since $W(\cdot,\cdot)$ is a non-negative function, this contradicts the best response definition. This completes the proof of Lemma 5. ■

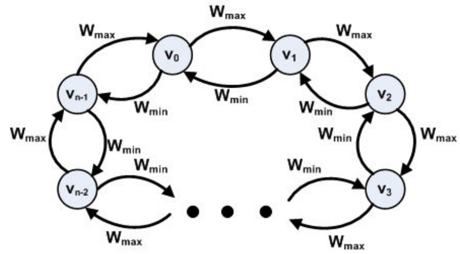

Fig. 1. A network topology with price of anarchy $\frac{W_{max}}{W_{min}}$

Consider, for example, the network setup (partly) depicted in Figure 1. In this network, the price of the link connecting v_i to v_{i-1} is W_{min}, while all other link prices are W_{max}.

In case each node v_i chooses to establish a single link to v_{i-1}, the resulting topology G_{min} is a directed ring topology, thus, following Lemma 5, it is at Nash equilibrium. In addition, we note that the cost of this topology is $C(G_{min}) = N \cdot W_{min}$ (which is optimal). Next, we consider the case in which each node v_i chooses to establish a single link to v_{i+1}, the resulting topology G_{max} is a directed ring topology, thus, following Lemma 5, it is at Nash equilibrium. The cost of this topology is $C(G_{max}) = N \cdot W_{max}$, yielding a lower bound on the price of anarchy of $\Omega(\frac{W_{max}}{W_{min}})$.

[1]Actually, the first direction of Lemma 1 proves a claim stronger than required for establishing the theorem.

This completes the proof of Lemma 4. ■

Theorem 3 follows directly from Lemmas 2 and 4. ■

A *natural order game* is a game that begins with all players having no outgoing links (i.e., $\forall v_i \in V, Adj_G(v_i) = \emptyset$), and it advances as each player, in its turn, plays its best response; the game ends when in equilibrium.

We show that, even in such a simple setting, the connectivity game with link prices still yields a price of anarchy of $\Theta(\frac{W_{max}}{W_{min}})$.

Theorem 4: The price of anarchy for the natural order connectivity game with link prices is $\Theta(\frac{W_{max}}{W_{min}})$.

Proof: By showing both upper and lower bounds. Note that the upper bound of $O(\frac{W_{max}}{W_{min}})$ proven in Lemma 2 still holds. We prove the lower bound by using the example defined in Lemma 4. First, we show that if players perform best response moves in ascending order of index (i.e., player v_0 plays first, followed by v_1, v_2, and so on), then the optimal Nash equilibrium is achieved. When player v_0 first plays, all other players have no outgoing links established, and so v_0 is forced to directly connect with all other nodes at a cost of $W_{min}+(N-2)\cdot W_{max}$. When player v_1 enters the game, it can either transmit through v_0, or transmit directly to all players. Since transmission through v_0 yields v_1 a cost of W_{min}, it will strictly prefer this startegy. When player v_2 enters the game, its best response is to transmit through v_1 at a cost of W_{min}. In a similar fashion, each player $v_i, 3 \leq i \leq N-1$ will establish a single link to v_{i-1} at a cost of W_{min}. Once all players have played once, only v_0 can improve its cost by removing all links except for a single link to v_{N-1} (which has the price of W_{min}). It can be easily seen that, after v_0's second move, the resulting topology is the same as G_{min} defined in Lemma 4, and thus at Nash equilibrium. The total cost of this topology is $N \cdot W_{min}$.

Next, we present an entry order that yields a Nash equilibrium with a cost of $N \cdot W_{max}$. In this case players enter in decending order of index (i.e., v_{N-1} plays first, followed by v_{N-2}, v_{N-3}, and so on).

When player v_{N-1} first plays, all other players have no outgoing links established, and so v_{N-1} is forced to directly connect with all other nodes at a cost of $W_{min} + (N-2)\cdot W_{max}$. When player v_{N-2} enters the game, its best response move is to transmit through v_{N-1} at a cost of W_{max}. In a similar fashion, each player $v_i, 0 \leq i \leq N-3$ will establish a single link to v_{i+1} at a cost of W_{max} (note that in this case, other options with the same cost exist. making other choices will result in a Nash equilibrium with an equally bad cost). Once all players have played once, only v_{N-1} can improve its cost by removing all links but the one connecting it to v_0. It can be easily seen, that after v_{N-1}'s second move, the resulting topology is the same as G_{max} defined in Lemma 4, and thus at Nash equilibrium. The total cost of this topology is $N \cdot W_{max}$, and the lower bound on the price of anarchy follows. This completes the proof of Theorem 4. ■

V. Link Prices and Delay

In this section we discuss a game in which the cost function of each player v_i entails it both with the total price of the links it chooses to establish, as well as its distance from its destinations. This is perhaps the most elementary model that exhibits the link-delay tradeoff: on the one hand, each node aspires to reduce its expenses on link establishment; on the other hand, transmitting to many other nodes implies less hops between the node and its destinations, thus reducing the hops component of the cost.[2] Formally, the cost of each node v_i is defined as

$$c(v_i) = \alpha \cdot \sum_{v_j \in Adj_G(v_i)} W(v_i, v_j) + \sum_{v_j \in V} d_G(v_i, v_j),$$

where α is a parameter quantifying the relative importance of the hops component in comparison with the link establishment price component, and $d_G(v_i, v_j)$ denotes the length of the shortest path connecting v_i to v_j in G. As before, in case node v_i fails to connect to one (or more) of the nodes in the network, its cost is infinite (i.e., $c(v_i) = \infty$). We term the respective game as *the connectivity game with price-delay costs*.

We show that, while this game resembles the game analyzed in [15] (there, for homogeneous prices and undirected links), it may fail to admit a Nash equilibrium. Furthermore, we provide a generic example that does not admit any Nash equilibrium, regardless of the value of α. However, this example is rather particular; indeed, through simulations, we indicate that more typical scenarios do admit a Nash equilibrium. Our simulation results also indicate a reasonable (lower) bound on the price of anarchy. Finally, we show that a special case of the game, in which all link prices are identical, admits similar results to those described in [15]; we denote this special case as the *directed network creation game*.

Theorem 5: For any α, there exist instances of the connectivity game with price-delay costs that do not admit any Nash equilibrium point.

Proof: First, we prove two properties of the connectivity game with price-delay costs.

Lemma 6: Let $< V, W >$ define a network. Any Nash equilibrium includes all links with price $< \frac{1}{\alpha}$.

Proof: Assume, by negation, that there exists a Nash equilibrium topology G and a pair of node $v_i, v_j \in V$, such that $W(v_i, v_j) < \frac{1}{\alpha}$ and $v_j \notin Adj_G(v_i)$. It holds that the length of the path connecting v_i to v_j, $d_G(v_i, v_j)$, is greater than 1. We construct a new configuration G', which is identical to G, with the addition that $Adj_{G'}(v_i) = Adj_G(v_i) \cup \{v_j\}$ (i.e., v_i establishes an additional link to v_j). It holds that v_i's price of link establishment has increased by $\alpha \cdot W(v_i, v_j) < 1$, while its distance to v_j has decreased by at least one hop. It follows that $C_{G'}(v_i) < C_G(v_i)$, which contradicts our assumption that G is a Nash equilibrium topology. This completes the proof of Lemma 6. ■

Lemma 7: Let $< V, W >$ define a network. Let G be a topology defined in the following manner: $\forall v_i, v_j \in V, v_j \in Adj_G(v_i) \Leftrightarrow W(v_i, v_j) < \frac{1}{\alpha}$ (G includes all links with price

[2]Note that there is no interest in dealing with delay as the only design consideration, as it boils down to a trivial case where a clique, i.e., complete graph, is the optimal solution as well as the unique Nash equilibrium.

$\frac{1}{\alpha}$, and only such links). If G is strongly connected, then no Nash equilibrium includes any link with a price $> \frac{N^2}{2\cdot\alpha}$.

Proof: Assume, by negation, that there exists a Nash equilibrium topology G' and a pair of nodes $v_i, v_j \in V$, such that $W(v_i, v_j) > \frac{N^2}{2\cdot\alpha}$ and $v_j \in Adj_{G'}(v_i)$. Assume now that v_i removes its link to v_j. Note that following Lemma 6 the network remains strongly connected. It holds that v_i's price of link establishment has decreased by $\alpha \cdot W(v_i, v_j) > \frac{N^2}{2}$, while the value of it distances component of the cost is always upper bound by $\sum_{i=1}^{N-1}(i) = \frac{N\cdot(N-1)}{2}$. It follows that by removing its link to v_j, v_i has reduced its own cost, which contradicts our assumption that G' is a Nash equilibrium topology. This completes the proof of Lemma 7. ■

We continue by providing a specific example of a network which does not admit any Nash equilibrium for the connectivity game with price-delay costs. Consider the following network:

- $V = \{v_0, v_1, v_2, v_3, v_4\}$;
- W :
  - $\forall v_i \in V, W(v_i, v_{i+1}) = \frac{1}{2\cdot\alpha}$;
  - $\forall v_i \in V, W(v_i, v_{i+2}) = \frac{2.5}{\alpha}$;
  - for every other $v_i, v_j \in V, W(v_i, v_j) = \frac{13}{\alpha}$.

Assume, by negation, that the network described above has a Nash equilibrium topology G. Following Lemmas 6 and 7, it holds that G includes all links with price $< \frac{1}{\alpha}$, and does not include any of the links with price $> \frac{12.5}{\alpha}$. Hence, each node $v_i \in V$ has only two possible strategies:

- establish a single link to v_{i+1}, i.e., $Adj_G(v_i) = \{v_{i+1}\}$, we denote this strategy as s_G^1;
- establish two links, one to v_{i+1} and another to v_{i+2}, i.e., $Adj_G(v_i) = \{v_{i+1}, v_{i+2}\}$, we denote this strategy as s_G^2.

Lemma 8: For any player v_i, if v_{i+1} uses strategy s_G^1, v_i strictly prefers using strategy s_G^2.

Proof: Assume, w.l.o.g., that player v_1 uses the strategy s_G^1. First, we evaluate v_0's cost of using strategy s_G^1. Since v_1 establishes a single link to v_2, it is clear that the distance between v_0 and players v_3, v_4, is at least 3. It follows that v_0's cost of using strategy s_G^1 is at least $\alpha\cdot\frac{1}{2\cdot\alpha}+1+2+3+3 = 9.5$. By switching to strategy s_G^2, v_0 reduces its distance to three destinations (v_2, v_3 and v_4) by one hop each, at a cost of $\alpha\cdot\frac{2.5}{\alpha}$, and Lemma 8 follows. ■

Lemma 9: For any player v_i, if v_{i+1} uses the s_G^2 strategy, v_i strictly prefers using s_G^1.

Proof: Assume, w.l.o.g., that player v_1 uses the strategy s_G^2. It is clear that v_0's distance from v_3 is exactly 2. Player v_0's cost of using strategy s_G^2 is $\alpha \cdot \frac{3.5}{\alpha} + 1 + 1 + 2 + d(v_0, v_4)$, where $d(v_0, v_4)$ depends on player v_2's strategy, and may be either 2 or 3. By switching to strategy s_G^1, player v_0 increases its distance from v_2 and v_4 by one hop, but saves $\alpha \cdot \frac{2.5}{\alpha}$ in power, and Lemma 9 follows. ■

We consider two cases, based on v_4's strategy. In the first case, assume v_4 uses the strategy s_G^1 It follows from Lemma 8 that player v_3 must use strategy s_G^2. It then follows from Lemma 9 that player v_2 must use strategy s_G^1. Applying the same reasoning again, yields the conclusion that v_0 must use strategy s_G^1, hence v_4 must use strategy s_G^2, which contradicts our assumption. The second case in which v_4 uses the strategy s_G^2 yields a similar contradiction. This completes the proof of Theorem 5. ■

While the example above does not admit any Nash equilibrium, it is obtained within a rather particular case. We therefore turn to simulations in order to explore the existence of Nash equilibria in more general cases. In our simulation environment, five nodes are randomly placed on a grid. Next, the nodes' initial strategies are constructed. Initial strategies are constructed, following one out of three policies: *all empty*, that is, the game begins with no links; *all connected*, where the game begins with all possible links established; and *random connections* where the probability of each link to exist when the game begins is determined randomly, *i.i.d*, with 0.5 probability per link.

We consider three different weight functions, based on the Eucleadan distance between the two node, that is, if node v_i is placed at $< x_i, y_i >$, and node v_j is placed at $< x_j, y_j >$, then

$$d(v_i, v_j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}.$$

Specifically, we consider the following weight functions:

$$W(v_i, v_j) = d(v_i, v_j)^2,$$

$$W(v_i, v_j) = d(v_i, v_j),$$

and

$$W(v_i, v_j) = log_{10}(d(v_i, v_j)).$$

The players then take turns, by random order, and make best-response moves. We term a sequence in which each player makes a single best response move as an *iteration*. In case an iteration completes and all player strategies remain unchanged, we conclude that this strategy profile constitutes a Nash equilibrium point. On average, very few iterations (under 10) were required to reach a Nash equilibrium point. In case 100 iterations are played, and no Nash equilibrium is reached, we conclude that (apparently) there is no Nash equilibrium.

We employed the above simulation environment in order to evaluate the impact of the weight function, the grid size and the α value on the existence of Nash equilibria points. We used grids of three different sizes: $10x10$, $100x100$ and $1000x1000$, and 16 different α values, ranging from 0.25 to 1000. For each grid size and α value, we simulated 20 different random placements. We tested each random placement with 50 different play sequences and all strategy initialization policies. All in all, $432,000$ different simulation runs were executed. Our results are conclusive: in each out of the 2880 tested cases (3 weight functions, 3 grid sizes, 16 α values, 20 placements), at least one Nash equilibrium point was found. In fact, the vast majority (over 94%) of simulation runs (i.e., obtained for different play sequences) completed with a Nash equilibrium point.

In addition, for each node placement, we computed the optimal configuration (by converting the problem into a Binary Integer Problem, and using MATLAB to solve it). We used the optimal value, along with the value of the worst Nash equilibrium identified in the simulations, to compute a lower bound on the price of anarchy.

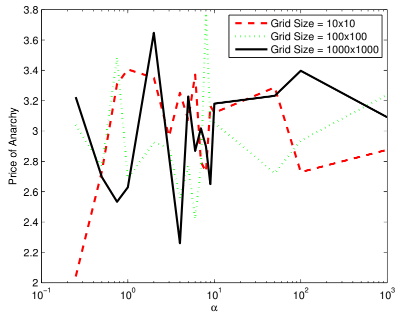

Fig. 2. The worst price of anarchy computed based on simulation results, $W(v_i, v_j) = d(v_i, v_j)^2$

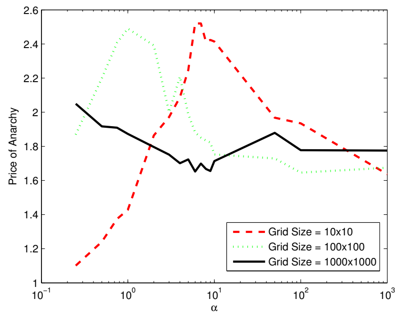

Fig. 3. The worst price of anarchy computed based on simulation results, $W(v_i, v_j) = d(v_i, v_j)$

Figures 2, 3 and 4 show the results of our experiment for the three weight functions. As can be seen in the figures, the lower bound on the price of anarchy is rather low (the worst value is 3.77).

A. The Directed Network Creation Game

Here, we address the directed network creation game, which is the *directed* version of the game first introduced in [15] (and later extended in [1], [13], all considering the undirected case). The directed network creation game is, in fact, a special case of the connectivity game with price-delay costs, in which the prices of all links are equal. For simplicity, we assume the for

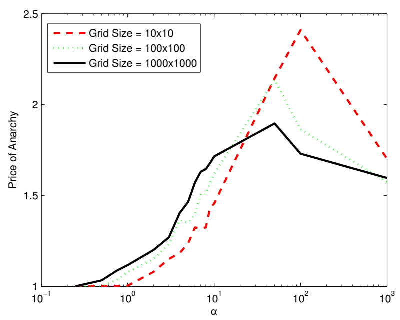

Fig. 4. The worst price of anarchy computed based on simulation results, $W(v_i, v_j) = log_{10}d(v_i, v_j)$

every $v_i, v_j \in V, W(v_i, v_j) = 1$. Similar results hold for any other link price value, $\bar{w}$, by substituting $\bar{\alpha} \equiv \frac{\alpha}{\bar{w}}$.

Before analyzing the game's price of anarchy, we note that a Nash equilibrium is guaranteed to exist. More precisely, and similarly to the undirected case [15], for any $\alpha < 1$, the directed clique is at Nash equilibrium, whereas for any $\alpha \geq 1$, the directed star is at Nash equilibrium.

Theorem 6: For any $\alpha < N^2$ the price of anarchy for the directed network creation game is $O(\sqrt{\alpha})$.

Proof: The proof goes along similar lines to that of Theorem 1 in [15]. First, we note that, since in any strongly connected graph there are at least N edges and the distance between any two nodes is at least 1, the cost of the optimal configuration is at least $\alpha \cdot N + N^2$. In addition, we note that the cost of a directed star is less than $2 \cdot (\alpha \cdot N + N^2)$. Hence, as in the undirected version, the price of anarchy is $\Theta\left(\frac{\alpha \cdot |E| + \sum_{v_j \in V} d_G(v_i, v_j)}{\alpha \cdot N + N^2}\right)$. Notice that, for every $v_i, v_j \in V$, it holds that $d_G(v_i, v_j) < 2\sqrt{\alpha}$, since otherwise v_i would connect to a node that is mid-way on the path to v_j, and since the path to v_j is at least $2\sqrt{\alpha}$ hops long, it will thus shorten its distance to at least $\sqrt{\alpha}$ nodes (all the nodes from the mid-way node to v_j) by $\sqrt{\alpha}$ (by substituting the $\sqrt{\alpha}$ long path from v_i to the mid-way node by a single link), at the cost of adding a single link - α. Therefore, it suffices to prove that $|E| = O\left(\frac{N^2}{\alpha}\right)$.

Consider the edges emanating from node v_i: $e_1, e_2, \ldots,$ etc. For any edge we will count vertices v_j for which $< v_i, v_j >$ is not in the graph. In other words, we will associate several non-edges with each edge. Ideally, we want the ratio between the number of edges and the number of non-edges to be $1 : \sqrt{\alpha}$.

Let $T(v_j)$ be the set of all nodes $v_k \in V$ such that the shortest path connecting v_i to v_k, $l_G(v_i, v_k)$, goes through v_j, i.e.,

$$T(v_j) = \{v_k \in V | l_G(v_i, v_k) \text{ goes through } < v_i, v_j >\}.$$

We ensure that $T(v_j)$ are disjoint by considering a canonical shortest path for each node. Let G' be the graph without G, i.e., $G' = (V, E \setminus \{e\})$. The alternative shortest path from v_i to $v_j \in T(v_j)$ in G' is either infinity of $< 4\sqrt{\alpha}$. We consider these two cases separately.

If v_i is connected to $T(v_j)$ in G', it holds that the distance from v_i to each node $v_k \in T(v_j)$ is at most $4\sqrt{\alpha}$. To justify the addition of e, the total improvement from the addition of e must be greater then α, i.e., $\sum_{v_k \in T(v_j)}(d_{G'}(v_i, v_k) - d_G(v_i, v_k)) \geq \alpha$, which implies that $|T(v_j)| = \Omega(\sqrt{\alpha})$.

If v_i is not connected to $T(v_j)$ in G', it holds that G' can be partitioned into two parts: all the nodes reachable from v_i in G' and $T(v_j)$. It holds that these two sets are disjoint. The lack of connectivity implies that none of the edges in the set reachable from v_i has an edge connecting it to a node in $T(v_j)$. Hence, we can count $(N - |T(v_j)|) \cdot |T(v_j)|$ non-edges in G. The minimal value this expression takes is $\frac{N^2}{4}$ when $|T(v_j)| = \frac{N}{2}$. Since $\alpha < N^2$, it follows that $|T(v_j)| = \Omega(\sqrt{\alpha})$.

This completes the proof of Theorem 6. ■

VI. Relaying Extent - Product Form

In the game presented in this section, players strive to achieve full connectivity while minimizing the use of routing paths that go through “hot spots”, which are characterized as paths with a high ”relaying extent”. The latter is captured by the in-degree of the nodes along the path. As explained in Section III, in-degrees quantify the relaying extent in two possible forms, namely the *product* of the in-degrees of nodes along the path, or the *bottleneck* of these values. The reader is referred to Section III for the motivation of these two forms.

In this section we focus on the product case, namely, each player chooses routes such that the product of nodal in-degrees along the path is minimal. We term this cost structure as *PRE (Product of Relaying Extent) costs*, and formally define them as follows:

$$c(v_i) = \sum_{v_k \in V} \prod_{v_j \in l_{v_i,v_k}} \eta_G^{in}(v_j)$$

As before, in case node v_i fails to achieve connectivity with one (or more) other nodes, its cost is infinite, i.e., $c(v_i) = \infty$. We term the resulting game as *the connectivity game with PRE costs*.

We show that both a directed clique, as well as a ring graph, are always at Nash equilibrium; but while the ring is also optimal, the clique configuration has a high cost, and therefore the game has a high price of anarchy.

Theorem 7: The clique configuration, i.e., $\forall v_i \in V$, $Adj_G(v_i) = V$, is a Nash equilibrium configuration for the connectivity game with PRE costs, for which $\forall v_i \in V, c(v_i) = (N-1)^2$.

Proof: Assume, by negation, that the clique configuration is not a Nash equilibrium. Then, there exists some node v_i, such that v_i can unilaterally reduce its cost. v_i can do this only by removing links. Assume, then, that v_i can reduce its cost by removing its link to v_j. It holds that v_i must remain connected to v_j. Let v_m be the first node on v_i's path to v_j. $\eta_{v_m}^{in} = N-1$. Hence, v_i's new cost is

$$c(v_i) = \sum_{v_k \in V} \prod_{v_l \in l_{v_i,v_k}} \eta_G^{in}(v_l) =$$

$$\sum_{v_k \in V \setminus \{v_j\}} \prod_{v_l \in l_{v_i,v_k}} \eta_G^{in}(v_l) + \sum \prod_{v_l \in l_{v_i,v_j}} \eta_G^{in}(v_l) =$$

$$\sum_{v_k \in V \setminus \{v_j\}} (N-1) + (N-1) \cdot \prod_{v_l \in l_{v_m,v_j}} \eta_G^{in}(v_l) =$$

$$(N-2) \cdot (N-1) + (N-1) \cdot \prod_{v_l \in l_{v_m,v_j}} \eta_G^{in}(v_l) >$$

$$(N-1)^2,$$

which contradicts our assumption. ■

Theorem 8: Any ring configuration, e.g., $\forall v_i \in V, Adj_G(v_i) = \{v_{i+1}\}$, is a Nash equilibrium configuration for the connectivity game with PRE costs, for which $\forall v_i \in V, c(v_i) = N-1$.

Proof: It holds that, in the ring configuration, $\forall v_i \in V, \eta_{v_i}^{in} = 1$, hence the cost of each node v_i is $c(v_i) = \sum_{v_k \in V} \prod_{v_j \in l_{v_i,v_k}} \eta_G^{in}(v_j) = \sum_{v_k \in V} \prod_{v_j \in l_{v_i,v_k}} 1 = \sum_{v_k \in V} 1 = N-1$. Since 1 is the minimal cost of any path, it holds that no node can reduce its cost, and the claim follows. ■

Theorem 9: Any ring configuration, e.g., $\forall v_i \in V, Adj_G(v_i) = \{v_{i+1}\}$, is an optimal configuration for the connectivity problem with PRE costs, for which $\forall v_i \in V, c(v_i) = N-1$.

Proof: Trivial. In a ring configuration, the in-degree of each node is 1. Therefore, the cost of each path is 1, which is the lowest cost possible. Hence, the cost of each node is $c(v_i) = \sum_{v_j \in V} 1 = N-1$. ■

Corollary 2: The optimal configuration cost for the connectivity problem with PRE costs is $N \cdot (N-1)$, and it is achieved by any ring configuration. ■

Corollary 3: The price of stability for the connectivity game with PRE costs is 1. ■

Theorem 10: The price of anarchy for the connectivity game with PRE costs is $\Theta(N)$.

Proof: By establishing the corresponding lower and upper bounds, in the following lemmas.

Lemma 10: The price of anarchy for the connectivity game with PRE costs is $O(N)$.

Proof: By establishing, in the next lemma, an upper bound on the cost of each node in any Nash equilibrium.

Lemma 11: Let s be a strategy profile, such that there exists a node v_i for which $c_s(v_i) > N \cdot (N-1)$. s is not a Nash equilibrium strategy profile.

Proof: Assume, by negation, that there exists a strategy profile s and a node v_i, such that $c_s(v_i) > N \cdot (N-1)$ and s is a Nash equilibrium strategy profile. It follows that there exists a node v_j, such that $\prod_{v_k \in l_{v_i,v_j}} \eta_G^{in}(v_k) > (N-1)$, and v_j is not used by v_i as a relay in any other path. Since the in-degree of each node is at most $N-1$, it holds that v_i is not directly connected to v_j in G. Assume now that v_i establishes a direct link to v_j. It holds that the cost v_i pays for the path connecting it to v_j has decreased, while the cost of any other path has not changed, which contradicts our assumption that s is a Nash equilibrium strategy profile.

This establishes Lemma 11. ■

Following Lemma 11, we derive an upper bound of $N^2 \cdot (N-1)$ on the cost of any Nash equilibrium configuration. Since the value of the optimal configuration is exactly $N \cdot (N-1)$, Lemma 10 follows. ■

The correctness of the lower bound follows directly from Theorem 7 and Theorem 9, thus Theorem 10 follows. ■

VII. Relaying Extent - Bottleneck Form

Following the previous section, in this section we focus on the *bottleneck* case namely, the cost entailed by each player reflects the player's “greatest relaying extent point”, i.e., the node along all of the player's routing paths that has the highest relaying extent value, which in our case is represented by the node's in-degree. We term this cost structure as *BRE (Bottleneck Relaying Extent) costs*, and formally define them as follows:

$$c(v_i) = \max_{v_k \in V} \max_{v_j \in l_{v_i,v_k}} \eta_G^{in}(v_k).$$

The resulting game is termed as *the connectivity game with BRE costs*. Recall that paths are chosen to optimize the considered design criteria, i.e., in this section, l_{v_i,v_j} denotes the path connecting v_i to v_j which has the lowest BRE cost value.

Note that, since each node is required to connect to all other nodes, it holds (in a strongly connected configuration) that $c(v_i) = \max_{v_j \in V, v_j \neq v_i} \eta_G^{in}(v_j)$. In this model, the social welfare, $C(G)$, is the bottleneck of the entire configuration; formally, $C(G) = \max_{v_i \in V} c(v_i)$ which, in strongly connected configurations, is equivalent to $\max_{v_i \in V} \eta_G^{in}(v_i)$.

We shall show that the directed ring is an optimal solution to the connectivity problem with BRE costs, and that any graph in which all in-degrees are equal is induced by a strategy profile that is at Nash equilibrium. With these findings, we shall establish that the game has a low price of stability and a high price of anarchy.

Theorem 11: The directed ring configuration, e.g., $\forall v_i \in V$, $Adj_G(v_i) = \{v_{i+1}\}$, is an optimal configuration for the connectivity problem with BRE costs, for which $\forall v_i \in V, c(v_i) = 1$.

Proof: Trivial. In a directed ring configuration, the in-degree of each node is 1. A ring configuration is strongly connected, and therefore the cost of each node is 1, which is the lowest cost possible. ■

Theorem 12: Let s be a strategy profile for the connectivity game with BRE costs, and let G be the directed graph induced by s. If $\forall v_i, v_j \in V \; \eta_G^{in}(v_i) = \eta_G^{in}(v_j)$, then s is at Nash equilibrium.

Proof: Let s be a strategy profile such that for each player v_i, $\eta_G^{in}(v_i) = \hat{\eta}$ for some $1 \leq \hat{\eta} \leq N-1$. Assume, by negation, that s is not at Nash equilibrium. Then, there exists some player v_i, such that v_i can unilaterally reduce its cost. v_i can do this in either one out of three ways: removing existing links; adding new links; and replacing links (i.e., removing some links, and adding others). Recall that v_i's cost is $\max_{v_j \in V, v_j \neq v_i} \eta_G^{in}(v_j) = \hat{\eta}$. It holds that if v_i adds any new links, connecting it to, say, v_j, it increases v_j's in-degree to $\hat{\eta}+1$, increasing v_i's cost to $\hat{\eta}+1$ thus contradicting the best response move definition. It remains to be shown that v_i cannot improve its cost by removing links. Note that v_i must keep at least one of its out-going links in order to be connected. Assume that v_i keeps its link connecting it to, say, v_j. It holds that v_j's in-degree remains unchanged, and therefore, v_i's cost does not change, in contradiction to our assumption. ■

Corollary 4: The directed ring configuration, e.g., $\forall v_i \in V$, $Adj_G(v_i) = \{v_{i+1}\}$, is a Nash equilibrium configuration for the connectivity game with BRE costs, for which $\forall v_i \in V, c(v_i) = 1$. ■

Corollary 5: The price of stability for the connectivity game with BRE costs is 1. ■

Theorem 13: The directed clique configutaion, i.e., $\forall v_i \in V$, $Adj_G(v_i) = V$, is a Nash equilibrium configuration for the connectivity game with BRE costs, for which $\forall v_i \in V, c(v_i) = N-1$.

Corollary 6: The price of anarchy for the connectivity game with BRE costs is $N-1$. In addition, since $N-1$ is the maximal in-degree possible, this result is tight. ■

The above results suggest that, when the relaying extent of the bottleneck type is the only design consideration, the network nodes should not be left to play their game without any intervention of a social agent. Indeed, without such intervention, they may end up in an equilibrium where each node connects to each other node through a direct link. On the other hand, by allowing a social agent to be able to impose just the initial configuration of the system, it can guarantee a stable working point that is systemwide optimal.

VIII. Relaying Extent (Product Form) and Delay

In this section we discuss a combination of two design considerations, namely relaying extent and delay. Specifically, we consider the product form of relaying extent costs, i.e., PRE costs, defined in Section VI, however paths longer than d hops are not allowed. Accordingly, the cost is defined as Section VI, namely:

$$c(v_i) = \sum_{v_k \in V} \prod_{v_j \in l_{v_i,v_k}} \eta_G^{in}(v_j),$$

however, in case node v_i fails to connect to one (or more) of the nodes in the network, or if one (or more) of its routing paths are more than d hops long, its cost is infinite (i.e., $c(v_i) = \infty$). We term this cost structure as *H-PRE (Hop-constrained PRE) costs*, while the resulting game is termed *the connectivity game with H-PRE costs*.

We show that the optimal solution to the connectivity problem with H-PRE costs is $\Theta(\frac{N^3}{d})$. In addition, we show that the corresponding game has a unique Nash equilibrium and that the price of anarchy is $\Theta(d)$.

Theorem 14: The optimal value for the connectivity problem with H-PRE costs is $\Omega(\frac{N^3}{d})$.

Proof: First, recall the cost of a strongly connected configuration, namely:

$$c(G) = \sum_{v_i \in V} \sum_{v_j \in V} \prod_{v_k \in l_{v_i,v_j}} \eta_G^{in}(v_k).$$

This cost function can also be written as

$$\sum_{v_j \in V} \sum_{v_i \in V} \prod_{v_k \in l_{v_i,v_j}} \eta_G^{in}(v_k).$$

We proceed by showing that, for any $v_j \in V$,

$$\sum_{v_i \in V} \prod_{v_k \in l_{v_i,v_j}} \eta_G^{in}(v_k) \geq \frac{(N-1)^2}{d}.$$

Lemma 12: Let $G_d^*(V, v_j) = (V, E)$ be an inversed graph in which every node has a directed path connecting it to v_j in at most d hops. Let $C(G_d^*(V, v_j)) = \sum_{v_i \in V} \prod_{v_k \in l_{v_i,v_j}} \eta_G^{in}(v_k)$. If $G_d^*(V, v_j)$ is optimal w.r.t $C(G_d^*(V, v_j))$, then $G_d^*(V, v_j)$ is a tree.

Proof: We partition the nodes to sets according to their distance (on their chosen path) to v_j. Let U_k denote the nodes at distance k from v_j. We show that for every k, $1 \leq k \leq d$ and every node $v_k \in U_k$, v_k has exactly one out-going edge to a node in U_{k-1}.

First, we show that for every k, $1 \leq k \leq d$ and every node $v_k \in U_k$, v_k has exactly one out-going edge to nodes in sets U_m, $m < k$. Assume, by negation, that $v_k \in U_k$ has more than one out-going edge to nodes in sets U_m, $m < k$. It holds that all but one of these edges are not on l_{v_k,v_j} (i.e., they are not used by v_k to route data to v_j, but are there for other nodes). Let $e = (v_k, v_r)$ be such an edge, and let U_e be the set of nodes which are directly connected to v_k and their paths include e. Let v_n be a leaf (i.e., it is not on the path of any other node) such that $e \in l_{v_n,v_j}$. We construct a new graph $G_d(V, v_j)) = (V, E')$, where $E' = E \bigcup\{(v_n, v_r)\} \bigcup\{(v_w, v_n)|v_w \in U_e\} \setminus \{e\} \setminus \{(v_w, v_k)|v_w \in U_e\}$, i.e., v_n *replaces* v_k on all paths going through e. It is easy to see that $C(G_d(V, v_j)) < C(G^*_d(V, v_j))$, which contradicts $G^*_d(V, v_j)$'s optimality.

Next, following the above, it is easy to see that the only out-going edge from node $v_k \in U_k$ must connect it to a node in U_{k-1}.

Finally, we show that for every k, $1 \le k \le d$ and every node $v_k \in U_k$, v_k has no out-going edges to nodes in sets U_m, $m \ge k$. Assume, by negation, that $v_k \in U_k$ has at least one out-going edge to a node in U_m, $m \ge k$. Let $e = (v_k, v_m)$ be such an edge. Since v_k is at distance k from the root, and v_m is at distance m from the root, where $m \ge k$, it holds that there must be a node v_n on l_{v_m,v_j}, such that $v_n \in U_n$, and v_n has an out-going edge to a node in U_i, $i < n - 1$, which contradicts the properties established above, and lemma 12 follows. ■

Lemma 13: Let $G_d(V, v_j)$ be an inversed tree of depth $\le d$ rooted at v_j spanning V. For any $N \ge d \ge 1$, (where $N = |V|$, and $N \ge 1$) it holds that

$$\sum_{v_i \in V} \prod_{v_k \in l_{v_i,v_j}} \eta^{in}_G(v_k) \ge \frac{(N-1)^2}{d}.$$

Proof: Note that for $N = 1$ the claim holds trivialy (since the lower bound is 0 and the cost is always non-negative). Throughout the rest of the proof we assume $N > 1$ and prove by induction on d. For $d = 1$ it holds that $G_1(V, v_j)$ is an inversed star (i.e., each node $v_i \in V \setminus \{v_j\}$ has a link connecting it to v_j), hence, v_j's in-degree is $N - 1$ and

$$\sum_{v_i \in V} \prod_{v_k \in l_{v_i,v_j}} \eta^{in}_G(v_k) = (N-1)^2 = \frac{(N-1)^2}{d}.$$

Assume now that the claim holds for $d - 1$, and consider $G_d(V, v_j)$. In case $G_d(V, v_j)$ is of depth $< d$ the claim holds, otherwise, let U denote v_j's siblings in $G_d(V, v_j)$, i.e., $U = \{v_i | v_j \in Adj_G(v_i)\}$. For each $v_i \in U$ we define $T(v_i) = \{v_k | v_i \in l_{v_k,v_j}\}$. Note that the different T_i's are mutually disjoint. By the inductive hypothesis, the cost of $G_{d-1}(T(v_i), v_i)$ is at least $\frac{(|T(v_i)|-1)^2}{d-1}$. Consider now $G_d(V, v_j)$. It holds that the cost of $G_d(V, v_j)$ is at least

$$|U| \cdot \sum_{v_i \in U} \frac{(|T(v_i)| - 1)^2}{d-1} + |U|^2.$$

Since $\sum_{v_i \in U} |T(v_i)| = N - 1 - |U|$, it holds that $\sum_{v_i \in U}(|T(v_i)| - 1)^2$ is minimal when all $T(v_i)$'s are of the same size. Hence, the cost of $G_d(V, v_j)$ is at least

$$|U| \cdot \sum_{v_i \in U} \frac{(N-1-|U|)^2}{d-1} + |U|^2.$$

This expression is minimized for $|U| = \frac{N-1}{d}$, yielding a lower bound of $\frac{(N-1)^2}{d}$ as required by the lemma. ■

Let G^* be an optimal solution to the connectivity problem with H-PRE costs. It holds that for each node $v_i \in V$, there exists a graph $G_d(V, v_i)$ which is a subgraph of G^*. Following Lemma 13 it holds that the cost of each such graph is at least $\frac{(N-1)^2}{d}$, and thus the total cost of G^* is lower bounded by $\frac{N \cdot (N-1)^2}{d}$. This concludes the proof of Theorem 14. ■

Theorem 15: The optimal value for the connectivity problem with H-PRE costs is $O(\frac{N^3}{d})$.

Proof: By presenting a directed graph whose cost is at most $\frac{3 \cdot N^3}{d}$. First, we partition $V \setminus \{v_0\}$ into sets of $\frac{d}{2} - 1$ nodes, and denote these sets by $V_1, V_2, \ldots V_{\frac{2\cdot(N-1)}{d-2}}$. For each such set, we construct a directed ring that includes all the nodes in the set and v_0. It holds that, for each node $v_i \in V \setminus \{v_0\}$, v_i's in-degree is exactly 1. Since v_0 has a single in-coming link for each ring, its in-degree is $\frac{2\cdot(N-1)}{d-2}$. First, we show that the maximal distance between any two nodes is d. Let $v_i, v_j \in V$ denote two nodes. In case $v_i, v_j \in V_l$ for some l, or one of the nodes is v_0, it holds that both nodes are on the same ring, hence the maximal distance between them is $\frac{d}{2}$. Otherwise, it holds that v_0 is reachable from v_i within $\frac{d}{2}$ hops, and v_j is reachable from v_0 within $\frac{d}{2}$ hops. Finally, we evaluate the cost of the constructed graph. Note that each path goes through v_0 at most once, hence the cost of each path is $\le \frac{2\cdot(N-1)}{d-2}$. Thus, the total cost of the graph is $\le N^2 \cdot \frac{2\cdot(N-1)}{d-2}$. ■

Theorem 16: For any $d < N - 1$, the connectivity game with H-PRE costs has a unique Nash equilibrium configuration, which is the directed clique graph.

Proof: By showing, through the next lemma, that no other directed graph can be at Nash equilibrium.

Lemma 14: Let s be a strategy profile, and let G be the directed graph induced by s. Assume that there exists a pair of nodes $v_i, v_j \in V$ for which $\prod_{v_k \in l_{v_i,v_j}} \eta^{in}_G(v_k) > 2$ and, in addition, either (i) $\prod_{v_k \in l_{v_i,v_j} \setminus \{v_j\}} \eta^{in}_G(v_k) > 2$ or (ii) $\prod_{v_k \in l_{v_i,v_j} \setminus \{v_j\}} \eta^{in}_G(v_k) = 2$ and $\eta^{in}_G(v_j) \ge 2$ (i.e., either the product of in-degrees on the path connecting v_i to v_j, excluding v_j, is more than 2, or it is exactly 2 and v_j's in-degree is at least 2). Then, s is not at Nash equilibrium for the connectivity game with H-PRE costs.

Proof: Let s be a strategy profile as described in the claim above, and let G be the directed graph induced by it. We assume that v_j is a leaf in v_i's routing tree (i.e., v_i does not use v_j as a relay in any other path); otherwise, there exists another node v_k that conforms to the same conditions and is a leaf in v_i's routing tree. We show that v_i can unilaterally reduce its cost by adding a link connecting it to v_j. Let s' denote a strategy profile identical to s with the addition of a link connecting v_i to v_j, and let G' be the directed graph induced by s'.

$$c_{s'}(v_i) = \sum_{v_k \in V} \prod_{v_m \in l_{v_i,v_k}} \eta^{in}_{G'}(v_m) =$$

$$c_s(v_i) - \prod_{v_m \in l_{v_i,v_j}} \eta_G^{in}(v_m) + \left(\eta_G^{in}(v_j) + 1\right).$$

It holds that v_i reduces its cost only if

$$\prod_{v_m \in l_{v_i,v_j}} \eta_G^{in}(v_m) - \left(\eta_G^{in}(v_j) + 1\right) > 0.$$

Hence,

$$\prod_{v_m \in l_{v_i,v_j}} \eta_G^{in}(v_m) - \left(\eta_G^{in}(v_j) + 1\right) =$$

$$\prod_{v_m \in l_{v_i,v_j} \setminus \{v_j\}} \eta_G^{in}(v_m) \cdot \eta_G^{in}(v_j) - \left(\eta_G^{in}(v_j) + 1\right)$$

which is greater than 0 when

$$\prod_{v_m \in l_{v_i,v_j} \setminus \{v_j\}} \eta_G^{in}(v_m) > \frac{\eta_G^{in}(v_j) + 1}{\eta_G^{in}(v_j)}.$$

By assumption, either $\prod_{v_k \in l_{v_i,v_j} \setminus \{v_j\}} \eta_G^{in}(v_k) > 2$ in which case the inequality holds since

$$\frac{\eta_G^{in}(v_j) + 1}{\eta_G^{in}(v_j)} \leq 2,$$

or $\prod_{v_k \in l_{v_i,v_j} \setminus \{v_j\}} \eta_G^{in}(v_k) = 2$ in which case the inequality holds since

$$\frac{\eta_G^{in}(v_j) + 1}{\eta_G^{in}(v_j)} < 2$$

(following our assumption that $\eta_G^{in}(v_j) \geq 2$). This concludes the proof of Lemma 14. ■

Next, we show that, in any directed well-connected graph G such that, for each pair of nodes v_i, v_j, there exists a path connecting v_i to v_j that is d hops or less, where $d < |V| - 1$, there exists a pair of nodes that adhere to the conditions of Lemma 14.

Lemma 15: Let G be a directed well-connected graph such that, for each pair of nodes v_i, v_j, there exists a path connecting v_i to v_j that is d hops or less, where $d < |V| - 1$. Then, one of the following holds:

- there is at least one relay node v_k such that $\eta_G^{in}(v_k) > 2$; or
- there are at least two relay nodes v_k, v_m such that $\eta_G^{in}(v_k) \geq 2$ and $\eta_G^{in}(v_m) \geq 2$.

Proof: Assume, by negation, that the claim is wrong. It holds that there exists a directed graph G, such that G is well connected and for each pair of nodes, v_i, v_j, there exists a path in G connecting v_i to v_j which is d hops or less, $d < |V| - 1$. In addition, either every node in G has an in-degree of 1, or it holds that G includes a single node v_0 such that $\eta_G^{in}(v_0) = 2$ whereas for any other node v_i, $\eta_G^{in}(v_i) = 1$. In the former case, it is clear that G is a *ring*, and therefore there exist pairs of nodes v_i, v_j such that the shortest path connecting v_i to v_j is $|V| - 1$ hops long, contradicting our assumption.

In the latter case, it holds that G is a "figure-8" graph, such as the one depicted in Figure 5. Note that the different parts of the graph may be of various size, i.e., it may be that $v_b \equiv v_z$, and/or $v_0 \equiv v_a$. Consider the shortest path connecting v_b to v_y.

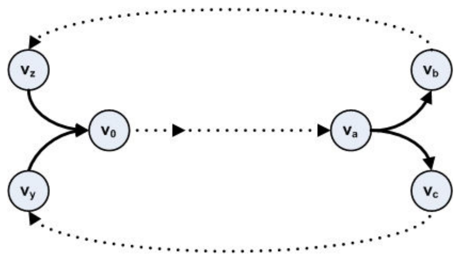

Fig. 5. A "figure-8" graph

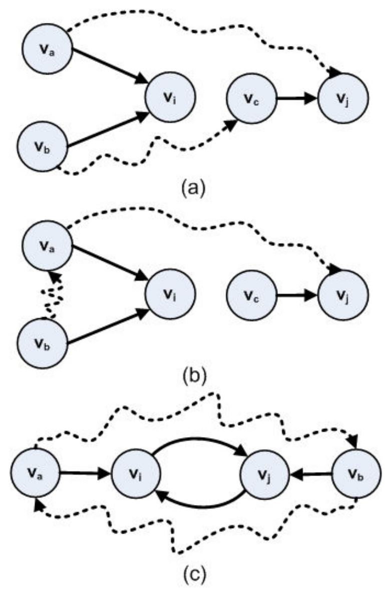

Fig. 6. well connected graphs including relays with in-degrees eqal to 2.

It is clear that this path is $|V| - 1$ hops long, in contradiction to our assumption. This concludes the proof of Lemma 15. ■

Note, however, that Lemma 15 is not enough to complete the proof of Theorem 16, since in the case of relay nodes with an in-degree of at least 2 we must prove the existence of a path in G that uses at least two such nodes (one of which may be the end node). Since the case of relay nodes with an in-degree strictly larger than 2 has been covered by Lemma 14, we need to address only the case of relay nodes whose in-degree is precisely 2.

Lemma 16: Let G be directed well-connected graph such that for each pair of nodes, v_i, v_j, there exists a path connecting v_i to v_j which is d hops or less, $d < |V| - 1$, and, in addition, G includes at least two nodes whose in-degree is 2. Then, there is a path in G that uses at least two relay nodes whose in-degree is 2 (one of which may be the end node).

Proof: Assume, by negation, that the claim is wrong. Let v_i, v_j be two nodes whose in-degree is 2. We consider three cases. The first and second cases are that either $v_i \notin Adj_G(v_j)$ or $v_j \notin Adj_G(v_i)$, as depicted in Figure 6(a) and Figure 6(b). In the first case, the two nodes that have links to v_i (v_a, v_b in the figure) must have paths to v_j that go through nodes with an in-degree of 1. Hence, when we consider the path connecting v_i to v_j we need to create an additional relay along one of these paths, contradicting our assumption. In the second case, the path connecting v_b to v_j goes through v_a; this example yields a contradiction in a similar manner (from the need to connect v_c to v_i without going through v_j). Finally, the third case is one in which $v_i \in Adj_G(v_j)$ and $v_j \in Adj_G(v_i)$, depicted in Figure 6(c). Using the same reasoning as above (from the need to connect v_i to v_b and v_j to v_a), a similar contradiction is obtained. ■

Theorem 16 follows from Lemmas 14, 15 and 16. ■

Corollary 7: The price of anarchy for the connectivity game with H-PRE costs is $\Theta(d)$. ■

IX. Relaying Extent (Bottleneck) And Delay

In this section we analyze the combination of relaying extent in its bottleneck form (i.e., BRE costs) with delay. As in Section VIII, we restrict players from using routing paths longer than d hops by entailing them with an infinite cost in case such paths are used. Formally, the cost in this model is defined as:

$$c(v_i) = \max_{v_k \in V} \max_{v_j \in l_{v_i,v_k}} \eta_G^{in}(v_k),$$

however, in case node v_i fails to connect to one (or more) of the nodes in the network, or if one (or more) of its routing paths are more than d hops long, its cost is infinite (i.e., $c(v_i) = \infty$). We term this cost structure as *H-BRE (Hop-constrained BRE) costs*, while the resulting game is termed *the connectivity game with H-BRE costs*.

As in Section VII, the social welfare in this game, $C(G)$, is the bottleneck of the entire configuration; formally, $C(G) = \max_{v_i \in V} c(v_i)$.

In the following, we obtain the following results. First, we establish upper and lower bounds on the solution of the systemwide (i.e., social) optimization problem (Theorems 17-19). In particular, we show that, if the restriction d on the number of hops is loose, namely $d = \theta(N)$, then the value of the solution is small, namely the bottleneck in-degree is $\theta(1)$ (Corollary 8). With these bounds at hand, we proceed to characterize the properties of the game. First, we characterize sufficient conditions for a topology to be at Nash equilibrium (Theorem 20). This allows us to show that a Nash equilibrium exists, and, moreover, there is one that is systemwide optimal (Theorem 21). On the other hand, we show that the clique configuration is the worst possible Nash equilibrium (Corollary 9), which implies a very large price of anarchy (Theorems 22-23).

Theorem 17: The optimal value for the connectivity problem with H-BRE costs is upper bounded by $\frac{2 \cdot N}{d}$.

Proof: By showing a configuration with a maximal in-degree of $\frac{2\cdot(N-1)}{d-1}$. First, we partition $V \setminus \{v_0\}$ into sets of $d-1$ nodes, denote these groups by $V_1, V_2, \ldots V_{\frac{N-1}{d-1}}$. Next, for each set, we construct a doubly-linked ring including all the nodes in the set and v_0. It holds that for each node $v_i \in V \setminus \{v_0\}$, v_i's in-degree is exactly 2. Since v_0 has two in-coming links on each ring, its in-degree is $\frac{2\cdot(N-1)}{d-1}$. Finally, we show that the maximal distance between any two nodes is d. Let $v_i, v_j \in V$ denote two nodes. We show that the maximal distance from v_i to v_j is d. In case $v_i, v_j \in V_l$ for some l, or one of the nodes is v_0, it holds that both nodes are on the same ring, hence the maximal distance between them is $\frac{d}{2}$. Otherwise, it holds that v_0 is reachable from v_i within $\frac{d}{2}$ hops, and v_j is reachable from v_0 within $\frac{d}{2}$ hops, and the claim follows. ■

Theorem 18: For any d larger than $log(N)$, the optimal value for the connectivity problem with H-BRE costs is upper bounded by $\min\{\frac{2\cdot N}{d}, log(N)\}$.

Proof: First, we note that the upper bound of $\frac{2\cdot N}{d}$ proven in Theorem 17 holds for any value of d. Next, we note that if we arrange the nodes in a chord-skip-list structure [30], where each node v_i is connected to all nodes v_j such that $j = (i + 2^k)\%N$ (where $\%$ denotes the modulo operator) for some $0 \le k \le log(N)$, the distance between any two nodes will be at most $log(N)$, and the in-degree of each node will be $log(N)$. ■

Theorem 19: The optimal value for the connectivity problem with H-BRE costs is lower bounded by $N^{\frac{1}{d}} \equiv 2^{\frac{1}{d}\cdot log(N)}$.

Proof: Let $\eta*$ denote the optimal in-degree value. It holds that at most $\eta*$ nodes can reach v_0 within a single hop. In general, at most $\sum_{i=0}^{l}(\eta*)^i$ nodes can reach v_0 within l hops. Therefore, since all nodes must be able to reach v_0 within d hops, it holds that $N \le \sum_{i=0}^{d}(\eta*)^i \le \eta*^d$, which means that $\eta* \ge N^{\frac{1}{d}}$, and the claim follows. ■

Corollary 8: In case $d \in \theta(N)$, the bounds defined in Theorems 17 and 19 are asymptotically tight, i.e., the optimal solution has $\theta(1)$ in-degree. ■

Theorem 20: Let s be a strategy profile for the connectivity problem with H-BRE costs, and let G be the directed graph induced by s. If $\forall v_i, v_j \in V$ $\eta_G^{in}(v_i) = \eta_G^{in}(v_j)$ and $d_G(v_i, v_j) \le d$, then s is at Nash equilibrium.

Proof: Let s be a strategy profile such that for each player v_i, $\eta_G^{in}(v_i) = \hat{\eta}$ for some $1 \le \hat{\eta} \le N-1$, that is, all in-degrees in the induced graph are the same (a ring, a clique, and a chord-skip-list are examples of such graphs), and, $\forall v_i, v_j \in V, d_G(v_i, v_j) \le d$. Assume, by negation, that s is not at Nash equilibrium. Then, there exists some player v_i, such that v_i can unilaterally reduce its cost. v_i can do this in either one out of three ways: removing existing links; adding new links; and replacing links (i.e., removing some links, and adding others). Recall that v_i's cost is $\max_{v_j \in V, v_j \neq v_i} \eta_G^{in}(v_j) = \hat{\eta}$. It holds that if v_i adds any new links, connecting it to, say, v_j, it increases v_j's in-degree to $\hat{\eta}+1$, increasing v_i's cost to $\hat{\eta}+1$ thus contradicting the best response move definition. It remains to be shown that v_i cannot improve its cost by removing links. Note that v_i must keep at least one of its out-going links in order to be connected. Assume that v_i keeps its link connecting it to, say, v_j. It holds that v_j's in-degree remains unchanged, and therefore, v_i's cost does not change, in contradiction to our assumption. ■

Theorem 21: The connectivity game with H-BRE costs admits a Nash equilibrium; moreover, its price of stability is 1.

Proof: By showing that there exists a strategy profile s which is at Nash equilibrium, and that the directed graph induced by it, G, is optimal. Let G' denote an optimal solution for the connectivity problem with H-BRE costs. Let s' be a strategy profile such that G' is induced by s'. Let η^* denote the maximal in-degree in G', i.e.,

$$\eta^* = \max_{v_i \in V} \eta_{G'}^{in}(v_i).$$

It holds that $C(G') = \eta^*$. In addition, it holds that $\forall v_i, v_j \in V, d_{G'}(v_i, v_j) \leq d$. We construct a new strategy profile, s in the following manner: starting with s', we add links connecting arbitrary nodes to nodes that have in-degree smaller than η^*. We continue to add links until all nodes have an in-degree of η^*. Let G be the graph induced by s. It holds that in s, $\forall v_i, v_j \in V, \eta_G^{in}(v_i) = \eta_G^{in}(v_j) = \eta^*$ and $d_G(v_i, v_j) \leq d$, and, in addition, $C(G) = \max_{v_i \in V} \eta_G^{in}(v_i) = \eta^*$. And so, it holds that s satisfies the conditions of Theorem 20, while its social welfare is optimal, and the claim follows. ■

Corollary 9: The clique configuration, i.e., $\forall v_i \in V, Adj(v_i) = V$, is a Nash equilibrium configuration for the connectivity problem for which $\forall v_i \in V, c(v_i) = N - 1$. Since any node's in-degree in upper bounded by $N - 1$, it holds that the clique configuration is the worst Nash equilibrium configuration. ■

Theorem 22: The price of anarchy for the connectivity game with H-BRE costs is $O(N^{1-\frac{1}{d}})$.

Proof: The correctness of the theorem follows directly from Theorem 19 and Corollary 9. ■

Theorem 23: The price of anarchy for the connectivity game with H-BRE costs is $\Omega(d)$.

Proof: The correctness of the theorem follows directly from Theorem 17 and Corollary 9. ■

Corollary 10: If d is larger than $log(N)$, then the price of anarchy for the connectivity game with H-BRE costs is $\Omega(\max\{d, \frac{N}{log(N)}\})$. ■

X. Relaying Extent and Link Costs

In this section we discuss a combination of two design considerations, namely relaying extent and link establishment costs. That is, each player v_i is entailed with both the total price of the links it chooses to establish, as well as the relaying extent along the paths it uses. Formally, the cost of each node v_i is defined as

$$c(v_i) = \alpha \cdot \sum_{v_j \in Adj_G(v_i)} W(v_i, v_j) + \sum_{v_j \in V} R_G(l_{v_i, v_j}),$$

where α is a parameter quantifying the relative importance of the relaying extent component in comparison with the link establishment price component, and $R(l_{v_i, v_j})$ denotes the relaying extent along the shortest path connecting v_i to v_j in G, which may reflect either the product or the bottleneck forms of the relaying extent (RE). As before, in case node v_i fails to connect to one (or more) of the nodes in the network, its cost is infinite (i.e., $c(v_i) = \infty$). We term the respective game as *the connectivity game with price-RE costs*.

Lemma 17: In the connectivity game with price-RE costs, every directed ring topology is at Nash equilibrium.

Proof: The proof is identical to that of Lemma 5. ■

In the proof of the next Theorem, we make use of the notations defined in Section IV, namely, W^i_{min} denotes player i's minimal connectivity cost, W^i_{max} denotes player i's maximal connectivity cost, and W_{min} and W_{max} denote the game's minimal and maximal connectivity prices, respectively.

Lemma 18: The price of anarchy for the connectivity game with price-RE costs is $\Omega(\frac{W_{max}+N}{W_{min}+N})$.

Proof: The proof follows the same lines as that of Lemma 4. Specifically, consider the network setup (partly) depicted in Figure 1. In this network, the price of the link connecting v_i to v_{i-1} is W_{min}, while all other link prices are W_{max}.

In case each node v_i chooses to establish a single link to v_{i-1}, the resulting topology G_{min} is a directed ring topology, thus, following Lemma 17, it is at Nash equilibrium. In addition, we note that the cost of this topology is $C(G_{min}) = N \cdot (W_{min} + N)$ (which is optimal). Next, we consider the case in which each node v_i chooses to establish a single link to v_{i+1}, the resulting topology G_{max} is a directed ring topology, thus, following Lemma 17, it is at Nash equilibrium. The cost of this topology is $C(G_{max}) = N \cdot (W_{max} + N)$, yielding a lower bound on the price of anarchy of $\Omega(\frac{W_{max}+N}{W_{min}+N})$.

This completes the proof of Lemma 18. ■

Finally, we show that, in the case of homogeneous link costs, i.e., $(v_i, v_j) = w$ for every pair of players v_i and v_j, and for some non-negative cost w, the price of stability is 1.

Theorem 24: The price of stability for the connectivity game with homogeneous price-RE costs is 1.

Proof: We consider (any) directed ring. The cost for every player v_i is

$$c(v_i)\alpha \cdot w + (N - 1).$$

We note that this cost is minimum, as each player must establish at least one out-going edge, and, in addition, the bottleneck on the path to every other player is at least 1.

Therefore, in the case of homogeneous link costs, every directed ring is optimal. By Lemma 17, it is also at Nash equilibrium, and the theorem follows. ■

XI. Conclusion

We considered three major topology design considerations, namely the price of establishing a link, path delay and the relaying extent, the latter being a measure of the potential congestion along a path. To our knowledge, this is the first work that studies the game-theoretic aspects of these problems. We established that, while all but one of the considered games are guaranteed to have a Nash equilibrium point, their *price of anarchy* is potentially large. We also indicated that the *price of stability* is typically 1; hence, often optimal (networkwide) performance is achievable by just being able to impose the initial configuration on the players.

Several issues merit further research. First, we note that, even when the price of stability is 1, computing the actual initial configuration may be computationally hard. Therefore, the computational tractability of the basic system optimization variants of the problems considered in this paper is an intriguing question. In addition, the question of designing mechanisms to induce cooperation among the network nodes without enforcing an initial configuration is an interesting one. Moreover, as one of the considered games fails to admit a

Nash equilibrium, it is important to investigate how the system behavior can be characterized in such cases; another related question is whether a Nash equilibrium (and its implied stability) could be obtained in such cases through the intervention of a central ("social") agent.

While the present study provides important insight for the above considerations, several extensions of the model assumptions are of interest. Examples of such assumptions are the rather simple cost structures and the complete topological information available to each user. Finally, while this study focused on stationary networks, future research should investigate the impact of nodal mobility in wireless settings, that is, a case where nodes may move around and thus the cost of links would change dynamically.

Acknowledgments

We would like to thank Danny Raz for fruitful discussions that paved the way for some of the results.

References

[1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty, "On Nash equilibria for a network creation game," in *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. New York, NY, USA: ACM Press, 2006, pp. 89–98.

[2] E. Altman, "Flow control using the theory of zero sum markov games," *Decision and Control, 1992., Proceedings of the 31st IEEE Conference on*, pp. 1632–1637 vol.2, 1992.

[3] E. Altman, T. Basar, T. Jiménez, and N. Shimkin, "Competitive routing in networks with polynomial cost," in *INFOCOM*, 2000, pp. 1586–1593.

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

[5] ——, "The price of stability for network design with fair cost allocation," in *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, ser. FOCS '04. Washington, DC, USA: IEEE Computer Society, 2004, pp. 295–304. [Online]. Available: http://dx.doi.org/10.1109/FOCS.2004.68

[6] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler, "Near-optimal network design with selfish agents," in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, ser. STOC '03. New York, NY, USA: ACM, 2003, pp. 511–520. [Online]. Available: http://doi.acm.org/10.1145/780542.780617

[7] R. Banner and A. Orda, "Bottleneck routing games in communication networks," *IEEE J.Sel. A. Commun.*, vol. 25, no. 6, pp. 1173–1179, Aug. 2007. [Online]. Available: http://dx.doi.org/10.1109/JSAC.2007.070811

[8] D. Bertsekas and R. Gallager, *Data networks (2nd ed.)*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1992.

[9] U. Brandes, M. Hoefer, and B. Nick, "Network creation games with disconnected equilibria," in *Proceedings of the 4th International Workshop on Internet and Network Economics*, ser. WINE '08. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 394–401. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-92185-1_45

[10] J. Corbo and D. Parkes, "The price of selfish behavior in bilateral network formation," in *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, ser. PODC '05. New York, NY, USA: ACM, 2005, pp. 99–107. [Online]. Available: http://doi.acm.org/10.1145/1073814.1073833

[11] T. Corman, C. Leiserson, and R. Rivest, *Introduction to Algorithms*. Cambridge MA. The MIT Press, 1990.

[12] G. Debreu, "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Science*, vol. 38, pp. 886–893, Oct. 1952.

[13] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam, "The price of anarchy in network creation games," in *PODC '07: Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*. New York, NY, USA: ACM, 2007, pp. 292–298.

[14] S. Eidenbenz, V. Anil Kumar, and S. Zust, "Equilibria in topology control games for ad hoc networks," *Mobile Networks and Applications*, vol. 11, pp. 143–159, 2006.

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

[16] R. S. Komali, A. B. MacKenzie, and R. P. Gilles, "Effect of selfish node behavior on efficient topology design." *IEEE Trans. Mob. Comput.*

[17] Y. A. Korilis and A. A. Lazar, "On the existence of equilibria in noncooperative optimal flow control," *Journal of the ACM*, vol. 42, no. 3, pp. 584–613, 1995.

[18] E. Koutsoupias and C. H. Papadimitriou, "Worst-case equilibria," *Lecture Notes in Computer Science*, no. 1563, pp. 404–413, 1999.

[19] A. A. Lazar, A. Orda, and D. E. Pendarakis, "Virtual path bandwidth allocation in multiuser networks," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 861–871, 1997.

[20] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, no. 14, pp. 142–143, 1996.

[21] T. Moscibroda, S. Schmid, and R. Wattenhofer, "On the topologies formed by selfish peers," in *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. New York, NY, USA: ACM Press, 2006, pp. 133–142.

[22] A. Nahir and A. Orda, "The energy-delay tradeoff in wireless networks - system-wide optimization and game-theoretic perspectives," Department of Electrical Engineering, Technion, Haifa, Israel, Tech. Rep., 2007. [Online]. Available: http://www.ee.technion.ac.il/Sites/People/ArielOrda/Info/Other/NO07.pdf

[23] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Transactions on Networking*, vol. 1, no. 5, pp. 510–521, oct 1993.

[24] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge MA. The MIT Press, 1994.

[25] C. Papadimitriou, "Algorithms, games, and the internet," in *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 2001, pp. 749–753.

[26] S. Pierre and G. Legault, "A genetic algorithm for designing distributed computer network topologies," *Systems, Man, and Cybernetics, Part B, IEEE Transactions on*, vol. 28, no. 2, pp. 249–258, Apr 1998.

[27] T. Roughgarden, "Designing networks for selfish users is hard," in *FOCS '01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 2001, p. 472.

[28] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.

[29] F.-M. Shao, X. Shen, and P.-H. Ho, "Reliability optimization of distributed access networks with constrained total cost," *Reliability, IEEE Transactions on*, vol. 54, no. 3, pp. 421–430, Sept. 2005.

[30] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," *SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 4, pp. 149–160, 2001.

[31] E. Szlachcic, "Fault tolerant topological design for computer networks," in *DEPCOS-RELCOMEX '06: Proceedings of the International Conference on Dependability of Computer Systems*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 150–159.

[32] Y. Wang and Z. Wang, "Explicit routing algorithms for internet traffic engineering," in *ICCN*, 1999.

[33] G. Zussman and A. Segall, "Energy efficient routing in ad hoc disaster recovery networks," in *INFOCOM*, 2003.

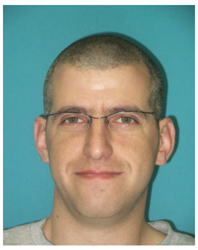

Amir Nahir received his B.Sc. (summa cum laude) in computer science from the Technion, Israel Institute of Technology in 2005, and is currently pursuing his doctoral degree there. He has been a research staff member at the IBM Research Labs in Haifa since 2006. His research interests include distributed systems and Cloud Computing, communication networks, with an emphasis on the application of game-theoretical methods to networking problems, as well as functional verification of hardware designs.

Ariel Orda (S'84-M'92-SM'97-F'06) received the B.Sc. (summa cum laude), M.Sc., and D.Sc. degrees in Electrical Engineering from the Technion in 1983, 1985, and 1991, respectively.

Since 1994, he has been with the Department of Electrical Engineering at the Technion, where he is the Herman and Gertrude Gross Professor of Communications. His research interests include network routing, the application of game theory to computer networking, survivability, QoS provisioning, wireless networks and network pricing.

He received several awards for research, teaching and service, most recently the 2009 Henry Taub Prize for Academic Excellence and the 2011 TCCC Outstanding Service Award. He served as Program co-chair of IEEE Infocom'2002, Program Chair of WiOpt'10 and General Chair of NETGCOOP'2012, and was an Editor of the IEEE/ACM Transactions on Networking and of the Journal of Computer Networks.

Ari Freund received the PhD degree in computer science from the Technion—Israel Institute of Technology in 2002. Following that, he was a postdoctoral fellow at the Caesarea Edmond Benjamin de Rothschild Foundation Institute for Interdisciplinary Applications of Computer Science. In 2003 he joined IBM Haifa Research Lab. He has been working at Google since 2010.