

Quality of Routing Congestion Games in Wireless Sensor Networks

(Invited Paper)

Costas Busch
Computer Science
Department
Louisiana State University
286 Coates Hall
Baton Rouge, LA 70803, USA
busch@csc.lsu.edu

Rajgopal Kannan
Computer Science
Department
Louisiana State University
279 Coates Hall
Baton Rouge, LA 70803, USA
rkannan@csc.lsu.edu

Athanasios V. Vasilakos
Department of Computer
Science &
Telecommunications
Engineering
University of Western
Macedonia
Greece 45431
vasilako@ath.forthnet.gr

ABSTRACT

We consider congestion games in wireless sensor networks that offer quantitatively distinct classes of routing paths. Each routing class is characterized by a service cost. Within a routing class, the maximum link congestion is also an important metric for measuring the quality of the paths. Here, we study routing games where each player i selfishly selects a path with a respective routing class that simultaneously minimizes its maximum edge congestion C_i and service cost S_i , in other words minimizes $C_i + S_i$. We examine the quality of Nash-equilibria and prove that the price of stability is 1. The price of anarchy is bounded above by $\min(C^*, S^*) \cdot m \log n$, where m is the number of routing classes, n is the size of the graph, and C^* and S^* are the optimal coordinated congestion and service costs. Thus, under certain circumstances, the player's selfishness does not hurt the social welfare and actually the equilibria can give good approximations for the coordinated optimal social cost.

Categories and Subject Descriptors

C.2.2 [Computer-communication networks]: Network Protocols—*routing protocols*; F.2.2 [Analysis of algorithms and problem complexity]: Nonnumerical algorithms and problems—*Routing and layout*

General Terms

Algorithms, Theory

Keywords

Algorithmic game theory, congestion games, Nash equilibrium, price of stability, price of anarchy

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON'08 November 17-19, 2008, Maui, Hawaii, USA
Copyright 2008 ICST 978-963-9799-36-3 ACM ...\$5.00.

1. INTRODUCTION

Routing is a fundamental task in wireless sensor network [13]. Routing algorithms provide paths for packets that will be sent over the network. Motivated by the selfish behavior of entities in wireless sensor networks, we study routing games where each packet's path is controlled independently by a selfish player.

Wireless sensor networks are extremely energy-limited. The rate of energy consumption and hence the overall lifetime at a node depends on the rate at which information flows through it. Most of the communication energy at a node is spent on forwarding packets from various flows. In heterogeneous sensor networks, different nodes have varying capabilities for supporting traffic flows through forwarding. This aspect of heterogeneity should form an important part of the routing framework.

Game-theoretic models have been extensively used to model traffic flows in networks. Routing games based on node/link congestion reflect energy consumption and hence network lifetimes. In traditional congestion games, the optimal social objective is to route traffic so as to minimize overall congestion or maximize the minimum battery lifetime in the network whereas individual nodes are concerned with choosing routes that maximize their connection lifetimes i.e. minimizing congestion (energy consumptions) along their routes.

In this paper, we are interested in evaluating games of traffic flow in heterogeneous sensor networks. Given the increasing presence of heterogeneity in sensor networks (for example, sensor networks in SCADA systems [31]), it makes sense to consider the availability of different traffic routing classes in the network. Analogous to the idea of Quality of Service classes in wired networks. A particular routing class may be composed of nodes with a specific range of available energies. A given routing class can thus guarantee minimal connection lifetimes based on the number of nodes participating in that class. A node can pay a given service cost to participate in routes belonging to a particular traffic class. Each class guarantees a minimum routing quality. However if too many nodes select a particular class, the quality of routing will obviously deteriorate. We can assume that nodes in a particular routing form a separate subnetwork

and only route traffic flows in the class. Thus nodes in different classes do not affect energy consumptions and hence congestion in each other.

In this paper we develop routing games based on the availability of different traffic routing classes. We characterize each routing class by a fixed service cost which reflects the connection/energy guarantees provided by the class. A class with high service costs may correspond to a subnetwork of nodes with high available energies. Since these nodes can transmit at higher powers, they can guarantee lower delays and higher successful transmission probabilities. However nodes with higher service costs may be attractive for these reasons and thus have higher congestion which degrades performance. Hence the metric which reflects overall costs to participants in a routing class is the sum of a fixed service cost and variable congestion costs.

1.1 Quality of Routing Congestion Games

We consider network congestion games where there is a set of N players and each player has to select a path from a source node to a destination node in the network. We consider *atomic* congestion games where the flow of each player is unsplitable, that is, all the traffic flow of the player is sent along one path from the source to the destination. The paths selected by different players may interfere with each other when they have common edges. We use as a path interference metric the *edge congestion*, where the congestion of an edge is the number of player paths that use the edge. Each selected path has a cost which is related to the congestion of its bottleneck edge (the most congested edge along the path). The players select paths that minimize the costs, that is, they prefer paths with lower congestion. Edge congestion is an appropriate cost metric which relates to the energy required to transmit packet along that edge, that is, the higher the congestion on an edge the more the energy needed to transmit the packets. Thus, lower congestion implies increased lifetime for the network. Further congestion is related to the scheduling time for sending the packets, the higher the congestion, the faster the packet delivery schedule.

We introduce here the *quality of routing congestion games* (*QoR games*), where the paths are partitioned into m routing classes Q_1, \dots, Q_m , where each routing class Q_j has a service cost $S(Q_j)$. Each path in the network (and hence the path selected by a player) is assigned to exactly one of those routing classes. The paths in different routing classes do not affect each other, that is, any two paths in different routing classes do not cause congestion to each other. This non-interfering property can be achieved by selecting the paths of different classes to be edge-disjoint, or by assigning the paths of each routing class to a different transmission channel (e.g. with frequency or time-division multiplexing). Thus, only players within the same routing class can cause congestion to each other (analogous to packets on the same frequency channel causing congestion to each other in a frequency multiplexed routing scheme).

For a player i the cost function is defined as $C_i + S_i$, where S_i is the service cost of the routing class that the path of i belongs to, and C_i is the congestion of the bottleneck edge of the player's chosen path, where the congestion is measured only among the paths that belong to the same routing class as player i .

Each player selects a path with the smallest cost possible.

The player is allowed to select its path in any of the routing classes. A player is *locally optimal* with its current choice if no other path choice yields lower cost for the player. We are interested in network states where every player is locally optimal, such a state is called a *Nash Equilibrium*. In particular, we examine *pure* Nash equilibria, where a strategy is a deterministic choice of a path, in contrast to *mixed* equilibria where the strategies are probability distributions among the paths.

We first show that QoR games always have pure Nash equilibria, which can be obtained with best response dynamics, where each player greedily improves its path whenever possible. We then examine the quality of the Nash equilibria. A natural problem is to determine the effect of the players' selfishness on the welfare of the whole communication network. The welfare of the network is measured with the *social cost*, SC , which is the maximum cost experienced by any player. We can write $SC = C + S$, where C is the congestion and S is the cost of the respective routing class of the the worst cost player. Let $SC^* = C^* + S^*$ denote the *optimal social cost*, where an optimal coordinated off-line algorithm assigns paths to players with the objective to minimize the social cost without considering the local optimality of the players.

QoR games may possibly have multiple Nash equilibria. We quantify the effect of selfishness with the *price of stability* and *price of anarchy* [15, 23], which express how much larger is the social cost in a Nash equilibrium compared to the social cost in the optimal coordinated solution. Price of stability (PoS) is the ratio of the best Nash equilibrium social cost to the optimal social cost. We show that $PoS = 1$, that is, there is always a Nash equilibrium that achieves optimal social cost.

Price of anarchy (PoA) is the ratio of the worst Nash equilibrium social cost to the optimal social cost. We examine *restricted-QoR* games where the service cost of each class is at least the length of each path that the class contains. Such games are interesting when different routing classes correspond to different ranges of path lengths; for example, it could be more costly to use longer paths than shorter paths. We show that the price of anarchy in any restricted-QoR game is bounded as:

$$PoA = O(\min(C^*, S^*) \cdot m \log n),$$

where n is the number of nodes in the underlying graph. Therefore, when either of C^* or S^* is small (e.g. a constant), and the number of classes m is small, the Nash equilibrium provides a very good approximation to the coordinated routing problem. Thus, in those scenarios the effect of selfishness is small, which is an important result that shows that the local optimality and selfishness can provide good outcomes for the global social welfare.

1.2 Related Work

Routing games (on congestion) were introduced and studied in [21, 25]. The notion of price of anarchy was introduced in [15]. Since then, many routing game models have been studied which are distinguished by the topology of the network, cost functions, type of traffic (atomic or splittable), nature of strategy sets, and kind of equilibria (pure or mixed). Specifically, pure equilibria with atomic flow have been studied in [2, 3, 18, 25, 30] (our work fits into this category), and with splittable flow in [26, 27, 28, 29]. Mixed

equilibria with atomic flow have been studied in [6, 7, 8, 10, 11, 12, 14, 15, 19, 20, 23], and with splittable flow in [4, 9].

Most of the work in the literature uses a single cost metric which is related to the network congestion. The most common metric for the player cost is the sum of the congestions on all the edges of the player's path [3, 12, 27, 28, 29, 30]. Games with similar player costs have been studied in the literature [3, 4, 6, 7, 8, 9, 10, 14, 15, 20, 23, 26], and other variations have been studied in [11, 18, 19, 25]. The vast majority of the work on routing games has been performed for parallel link networks, with only a few exceptions on general network topologies [2, 3, 4, 26].

Bottleneck congestion games have been studied in [1]. Our work is closer to [2] where they also study bottleneck congestion routing problems. We extended some results presented in [2] to apply to different routing classes and service costs, instead of using the single criterium of congestion. Specifically, the techniques that we use to prove existence of Nash equilibria with best response dynamics, and also to prove upper bounds on the price of anarchy, were originally introduced in [2]. Here, we modified and extended appropriately these techniques in a non-trivial way to apply to the setting of different quality of routing classes.

Outline of Paper

We proceed as follows. In Section 2 we give basic definitions. We study the stability of QoR games in Section 3. In Section 4 we estimate the price of anarchy. We finish with a discussion in Section 5.

2. DEFINITIONS

Here we describe the *quality of routing congestion games* (QoR games). An instance of a QoR game is a tuple $\mathcal{R} = (\mathcal{N}, G, \mathcal{P}, \mathcal{Q})$, where the constituents of the tuple are defined as follows:

- $\mathcal{N} = \{1, 2, \dots, N\}$ is the set of players,
- $G = (V, E)$ is a graph with nodes V and edges E ,
- \mathcal{P} is a set of available paths in G that the players can choose. We can write $\mathcal{P} = \bigcup_{i \in \mathcal{N}} \mathcal{P}_i$, where \mathcal{P}_i is a collection of available paths in G for player i . Each path in \mathcal{P}_i is from a source node $u_i \in V$ to destination node $v_i \in V$.
- $\mathcal{Q} = \{Q_1, \dots, Q_m\}$ is a set of *routing classes*. Each routing class Q_j imposes a fixed *service cost* $S(Q_j)$ to each path that uses it, where $1 \leq S(Q_i) < S(Q_j)$, for any $1 \leq i < j \leq m$. The paths in \mathcal{P} are partitioned so that each path $p \in \mathcal{P}$ belongs to one routing class, denoted $Q(p)$.

Each path in \mathcal{P}_i is a *pure strategy* available to player i . A *pure strategy profile* $\mathbf{p} = [p_1, p_2, \dots, p_N]$ is a collection of pure strategies (paths), one for each player, where $p_i \in \mathcal{P}_i$. We refer to a pure strategy profile as a *routing*. We use the standard notation \mathbf{p}_{-i} to refer to the collection of paths $\{p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_N\}$, and $(p_i; \mathbf{p}_{-i})$ as an alternative notation for \mathbf{p} which emphasizes the dependence on p_i .

Consider now a routing \mathbf{p} . For any $p_i \in \mathbf{p}$, we will use the notation $S_i(\mathbf{p}) = S(p_i) = S(Q(p_i))$. We will denote by $Q(\mathbf{p})$ the maximum class that any path in \mathbf{p} belongs to, and $S(\mathbf{p}) = S(Q(\mathbf{p}))$. For an edge e and routing class Q_j we define the *congestion of edge e at class Q_j* , denoted $C_{(e, Q_j)}(\mathbf{p})$,

as the number of paths in \mathbf{p} that belong to class Q_j and use edge e . We denote with $C_e(\mathbf{p})$ the largest congestion of edge e at any class, namely $C_e(\mathbf{p}) = \max_{j \in \{1, \dots, m\}} C_{(e, Q_j)}(\mathbf{p})$. For any path $p_i \in \mathbf{p}$ we define the *path congestion* $C_{p_i}(\mathbf{p})$ (alternatively, player i 's congestion), to be the maximum congestion at class $Q(p_i)$ of any of the edges of path p_i ; namely, $C_{p_i}(\mathbf{p}) = \max_{e \in p_i} C_{(e, Q(p_i))}(\mathbf{p})$. Note that only paths of the same routing class can cause congestion to each other. We will use the notation $C_i(\mathbf{p})$ to denote the congestion of player i 's path in \mathbf{p} , that is, $C_i(\mathbf{p}) = C_{p_i}(\mathbf{p})$. Let $C_{Q_j}(\mathbf{p})$ denote the *congestion of routing \mathbf{p} at routing class Q_j* , which is the maximum congestion experienced by any player in that class; namely, $C_{Q_j}(\mathbf{p}) = \max_{p_i \in Q_j} C_{p_i}(\mathbf{p})$. We denote by $C(\mathbf{p})$ the maximum congestion experienced by any player in routing \mathbf{p} , namely, $C(\mathbf{p}) = \max_{i \in \mathcal{N}} C_i(\mathbf{p})$. Whenever the context is clear, in the notations we will drop the dependence on \mathbf{p} , for example, we will write C instead of $C(\mathbf{p})$.

The *player cost* pc_i is defined as the sum of the congestion of the chosen path and the service cost of the routing class that the chosen path belongs; namely,

$$pc_i(\mathbf{p}) = C_i(\mathbf{p}) + S_i(\mathbf{p}).$$

The *social cost* of the game is defined as the largest player cost in any routing class, namely,

$$SC(\mathbf{p}) = \max_{j \in \{1, \dots, m\}} (C_{Q_j}(\mathbf{p}) + S(Q_j)).$$

Player i is *locally optimal* in routing \mathbf{p} if $pc_i(\mathbf{p}) \leq pc_i(p'_i; \mathbf{p}_{-i})$ for all paths $p'_i \in \mathcal{P}_i$. A routing \mathbf{p} is in a Nash Equilibrium (we say \mathbf{p} is a *Nash-routing*) if every player is locally optimal. A routing \mathbf{p}^* is an optimal pure strategy profile if it has minimum attainable social cost: for any other pure strategy profile \mathbf{p} , $SC(\mathbf{p}^*) \leq SC(\mathbf{p})$.

We quantify the quality of the Nash-routings by the *price of anarchy* (PoA) (sometimes referred to as the coordination ratio) and the *price of stability* (PoS). Let \mathbf{P} denote the set of distinct Nash-routings, and let SC^* denote the social cost of an optimal routing \mathbf{p}^* . Then,

$$PoS = \inf_{\mathbf{p} \in \mathbf{P}} \frac{SC(\mathbf{p})}{SC^*}, \quad PoA = \sup_{\mathbf{p} \in \mathbf{P}} \frac{SC(\mathbf{p})}{SC^*}.$$

3. PRICE OF STABILITY

Here, we show that QoR games have Nash-routings and we also estimate the price of stability. The existence of Nash routing relies on finding an appropriate potential function and a total order of the routings. We prove that given an arbitrary initial state any greedy move of a player can only give a new routing with smaller order. Thus, repeated greedy moves converge either to the smallest order routing or to a routing where no player can improve further. In either case, a Nash-routing is reached.

Here we define the potential function that we will use. Let $\mathcal{R} = (\mathcal{N}, G, \mathcal{P}, \mathcal{Q})$ be a QoR game. Let $r = N + S(Q_m)$. For any routing \mathbf{p} we define the *routing vector* $M(\mathbf{p}) = [m_1(\mathbf{p}), \dots, m_r(\mathbf{p})]$, where $m_i(\mathbf{p})$ is the number of players in \mathbf{p} whose cost is i . Note that if $SC(\mathbf{p}) = k$ then $m_k \neq 0$ and $m'_k = 0$ for all $k' > k$.

We give total order on the routings based on their lexicographic ordering. Let \mathbf{p} and \mathbf{p}' be two routings, with $M(\mathbf{p}) = [m_1, \dots, m_r]$, and $M(\mathbf{p}') = [m'_1, \dots, m'_r]$. We say that $M(\mathbf{p}) = M(\mathbf{p}')$ if $m_i = m'_i$ for all $1 \leq i \leq r$, and

$M(\mathbf{p}) < M(\mathbf{p}')$ if there is a j , $1 \leq j \leq r$, such that $m_k = m'_k$ for all $k > j$, and $m_j < m'_j$. We order the p and p' according to the order of their respective vectors, that is $p \leq p'$ if and only if $M(p) \leq M(p')$. Note that for any two p and p' it either holds that $p = p'$ or $p < p'$, that is, the routings are totally ordered.

Consider an arbitrary routing \mathbf{p} . If \mathbf{p} is not a Nash-routing, there is at least one user i which is not locally optimal. Then a *greedy move* is available to player i in which the player can obtain lower cost by changing the path from p_i to some other path p'_i with lower cost. In other words, the greedy move takes the original routing $\mathbf{p} = (p_i; \mathbf{p}_{-i})$ to a routing $\mathbf{p}' = (p'_i; \mathbf{p}_{-i})$ with improved player cost $pc_i(\mathbf{p}') < pc_i(\mathbf{p})$. We show now that any greedy move gives a smaller order routing:

LEMMA 3.1. *If a greedy move by player i takes a routing \mathbf{p} to a new routing \mathbf{p}' , then $\mathbf{p}' < \mathbf{p}$.*

PROOF. Let p_i and p'_i denote the paths of i in routings \mathbf{p} and \mathbf{p}' , respectively. Let $pc_i(\mathbf{p}) = C_i(\mathbf{p}) + S(p_i) = z_1$ and $pc_i(\mathbf{p}') = C_i(\mathbf{p}') + S(p'_i) = z_2$. Since player i decreases its cost in \mathbf{p}' , $z_2 < z_1$. Consider now the vectors of the routings $M(\mathbf{p}) = [m_1, \dots, m_r]$ and $M(\mathbf{p}') = [m'_1, \dots, m'_r]$. We will show that $M(\mathbf{p}) < M(\mathbf{p}')$.

Let $S(p_i) = k_1$ and $S(p'_i) = k_2$. Let Y denote the set of players whose cost increases in \mathbf{p}' with respect to their cost in \mathbf{p} . Next, we show that for any $j \in Y$ it holds that $pc_j(\mathbf{p}') \leq pc_j(\mathbf{p})$, which will help us to prove the desired result.

Let p_j be the path of player $j \in Y$, and let also $S(p_j) = k_3$ (note that j does not switch paths and routing classes between \mathbf{p} and \mathbf{p}'). If $k_3 \neq k_1$ and $k_3 \neq k_2$ then the cost of j remains unaffected between \mathbf{p} and \mathbf{p}' . Thus, either $k_3 = k_1$ or $k_3 = k_2$. If $k_3 = k_1$, then the cost of j can only decrease from \mathbf{p} to \mathbf{p}' , since path p_i no longer affects path p_j . Therefore, it has to be that $k_3 = k_2$. Suppose, for the sake of contradiction, that $pc_j(\mathbf{p}') > pc_j(\mathbf{p})$. Then, $C_j(\mathbf{p}') + k_3 > C_i(\mathbf{p}') + k_2$. Since $k_3 = k_2$, it has to be that $C_j(\mathbf{p}') > C_i(\mathbf{p}')$. Since $j \in Y$, $C_j(\mathbf{p}) < C_j(\mathbf{p}')$. The increase in congestion of p_j in \mathbf{p}' can only be caused by p'_i due to a common edge e with p_j that has congestion $C_{(e,k_3)}(\mathbf{p}') = C_j(\mathbf{p}')$. Therefore, $C_i(\mathbf{p}') \geq C_j(\mathbf{p}')$, a contradiction. Therefore, $pc_j(\mathbf{p}') \leq pc_j(\mathbf{p})$.

Consequently, in vector $M(\mathbf{p}')$ all the entries in positions $z_2 + 1, \dots, r$ do not increase with respect to $M(\mathbf{p})$. Further, since player i switches paths from cost z_1 to z_2 with $z_1 > z_2$, we obtain $m_{z_1} > m'_{z_1}$. Thus, $M(\mathbf{p}) < M(\mathbf{p}')$, as needed. \square

Since there are only a finite number of routings, Lemma 3.1 implies that starting from arbitrary initial state, every best response dynamic converges in a finite time to a Nash-routing, where every player is locally optimal.

THEOREM 3.2 (EXISTENCE OF NASH-ROUTINGS). *In any QoR game, starting from an arbitrary initial routing, best response dynamics converge to a Nash-routing.*

Since the routings are totally ordered, there is a routing \mathbf{p}_{\min} which is the minimum, that is, for all routings \mathbf{p} , $\mathbf{p}_{\min} \leq \mathbf{p}$. The minimum routing is also a Nash-routing, since no greedy move can improve from it. As we show next, the minimum routing \mathbf{p}_{\min} achieves also optimal social cost, which implies that the price of stability is 1.

LEMMA 3.3. *The minimum routing \mathbf{p}_{\min} achieves optimal social cost, that is, $SC(\mathbf{p}_{\min}) \leq SC(\mathbf{p})$ for any routing $\mathbf{p} \neq \mathbf{p}_{\min}$.*

PROOF. Suppose for contradiction that there exists a routing $\mathbf{p} \geq \mathbf{p}_{\min}$ with $SC(\mathbf{p}) < SC(\mathbf{p}_{\min})$. Let $SC(\mathbf{p}) = k_1$, and $SC(\mathbf{p}_{\min}) = k_2$. Clearly, $k_1 < k_2$. Therefore, in the vector $M(\mathbf{p}) = [m_1, \dots, m_r]$ it holds that $m_{k_1} \neq 0$ and $m_k = 0$ for $k > k_1$. Similarly, in the vector $M(\mathbf{p}_{\min}) = [\hat{m}_1, \dots, \hat{m}_r]$ it holds that $\hat{m}_{k_2} \neq 0$ and $\hat{m}_k = 0$ for $k > k_2$. Therefore, $M(\mathbf{p}_{\min}) > M(\mathbf{p})$, contradicting the fact that $\mathbf{p} \geq \mathbf{p}_{\min}$. \square

COROLLARY 3.4 (PRICE OF STABILITY). *For any QoR game, the price of stability is $PoS = 1$.*

4. PRICE OF ANARCHY

From Theorem 3.2, every QoR game has at least one Nash-routing. Here, we bound the price of anarchy. Consider a QoR game $\mathcal{R} = (\mathcal{N}, G, \mathcal{P}, \mathcal{Q})$, where G has n nodes, and \mathcal{Q} has m routing classes. We will consider *restricted-QoR* games where the routing classes have the following restriction for the path lengths of the packets that they contain:

For any path $p \in Q(p)$ it holds that $|p| \leq S(Q(p))$, where $|p|$ denotes the length of p .

So, assume that \mathcal{R} is a restricted-QoR game. Consider also a Nash-routing \mathbf{p} . Let $C = C(\mathbf{p})$ and $S = S(\mathbf{p})$. Let \mathbf{p}^* be the optimum (coordinated) routing with minimum social cost. Denote $C^* = C(\mathbf{p}^*)$ and $S^* = S(\mathbf{p}^*)$. Note that each payer $i \in \mathcal{N}$ has a path in $p_i \in \mathbf{p}$ (with congestion C_i and service cost S_i) and a corresponding ‘‘optimal’’ path $p_i^* \in \mathbf{p}^*$ (with congestion C_i^* and service cost S_i^*). We obtain:

LEMMA 4.1. *In Nash-routing \mathbf{p} , for any player i with $C_i \geq C - x$, where $x \geq 0$, it holds that $S_i \leq S_i^* + x + 1$.*

PROOF. Suppose for the sake of contradiction that there is a player i with $C_i \geq C - x$ and $S_i > S_i^* + x + 1$. Then, $pc_i = C_i + S_i > C - x + S_i^* + x + 1 = C + S_i^* + 1$. If user i was to switch to path p_i^* its cost would be $pc'_i \leq C + 1 + S_i^*$, since p_i^* has congestion at most C before player i switches its path, and the congestion of path p_i^* increases to at most $C + 1$ after player i switches to it. Therefore, $pc'_i \leq C + 1 + S_i^* < pc_i$. Thus, in \mathbf{p} player i would not be optimal, which is a contradiction, since \mathbf{p} is a Nash-routing. Therefore, $S_i \leq S_i^* + x + 1$, as needed. \square

For each edge $e \in G$ let $\Pi_e(\mathbf{p})$ denote the set of players whose paths in routing \mathbf{p} use edge e . For any edge e , we define $F(e, i)$ to be the set that contains all edges $e' \in p_i^*$ with $C_{(e', Q(p_i^*))}(\mathbf{p}) \geq C_i - S^*$. Let $F(e) = \cup_{i \in \Pi_e(\mathbf{p})} F(e, i)$, and for any set of edges X , $F(X) = \cup_{e \in X} F(e)$. We have:

LEMMA 4.2. *In Nash-routing \mathbf{p} , for every player i and edge e it holds that $|F(e, i)| \geq 1$.*

PROOF. Suppose that $|F(e, i)| = 0$. Then for every edge $e' \in p_i^*$ it holds that $C_{(e', Q(p_i^*))}(\mathbf{p}) < C_i - S^*$. Therefore, $C_{p_i^*}(\mathbf{p}) < C_i - S^*$. If player i was to choose path p_i^* its cost would be: $pc'_i \leq C_{p_i^*}(\mathbf{p}) + 1 + S_i^* < C_i - S^* + 1 + S^* = C_i + 1 \leq C_i + S_i = pc_i(\mathbf{p})$. Thus, $pc'_i < pc_i(\mathbf{p})$ which implies that player i is not locally optimal in Nash-routing \mathbf{p} , a contradiction. \square

We continue with the following result:

LEMMA 4.3. *Let Z be the set that contains all edges e with congestion $C_e(\mathbf{p}) \geq C - 2S^* \cdot \lg n$. If $C > 2S^* \cdot \lg n$, then there is a set of edges $X \subseteq Z$ with $|F(X)| \leq 2|X|$.*

PROOF. We recursively define sets of edges $E_0, \dots, E_{2 \lg n}$, such that $E_i = E_{i-1} \cup F(E_{i-1})$, and set E_0 contains all the edges e with congestion $C_e(\mathbf{p}) = C$. We show now that there is a j , $0 \leq j \leq 2 \lg n$, such that $|F(E_j)| \leq 2|E_j|$. Suppose for contradiction that such a j does not exist. Thus, for all j , $0 \leq j \leq 2 \lg n$, it holds that $|F(E_j)| > 2|E_j|$. In this case, it is straightforward to show that $|E_k| > 2|E_{k-1}|$, for any $1 \leq k \leq 2 \lg n$. Since $|E_0| \geq 1$, it holds that $|E_{2 \lg n}| > 2^{2 \lg n} = n^2$. However, this is a contradiction, since the number of edges in G cannot exceed n^2 . Thus, there is a j , $0 \leq j \leq 2 \lg n$, with $|F(E_j)| \leq 2|E_j|$; we set $X = E_j$.

It only remains to show that $X \subseteq Z$. It suffices to show that for any E_k and $e \in E_k$, $C_e(\mathbf{p}) \geq C - kS^*$, where $0 \leq k \leq 2 \lg n$. We prove this claim by induction on k . For $k = 0$ we have that every edge in e has $C_e(\mathbf{p}) = C = C - 0 \cdot S^*$, thus the claim trivially holds. For the induction hypothesis, suppose that the claim holds for any $k = t < 2 \lg n$. In the induction step we will prove that the claim holds also for $k = t + 1$. We have that $E_{t+1} = E_t \cup F(E_t)$. By induction hypothesis, for any $e \in E_t$, $C_e(\mathbf{p}) \geq C - tS^*$. Therefore, for any $e \in E_t$, $C_e(\mathbf{p}) \geq C - (t + 1)S^*$. Since $C > 2S^* \cdot \lg n$ and $t + 1 \leq 2 \lg n$, $C_e(\mathbf{p}) \geq 1$. Thus, for any $e \in E_t$ there is at least one player $i \in \Pi_e(\mathbf{p})$ whose path p_i uses edge e , and therefore $C_i \geq C - tS^*$. By definition, every edge $e' \in F(e, i)$ has the property that $C_{(e', Q(p_i))}(\mathbf{p}) \geq C_i - S^* \geq C - (t + 1)S^*$. Thus, $C_{e'}(\mathbf{p}) \geq C - (t + 1)S^*$. From the definition of $F(E_t)$, it follows that for any edge $e' \in F(E_t)$, $C_{e'}(\mathbf{p}) \geq C - (t + 1)S^*$. By considering the union $E_{t+1} = E_t \cup F(E_t)$, we obtain that for any $e \in E_{t+1}$, $C_e(\mathbf{p}) \geq C - (t + 1)S^*$, as needed. \square

LEMMA 4.4. *In Nash-routing \mathbf{p} it holds that $C \leq 10C^* \cdot S^* \cdot \lg n$.*

PROOF. If $C \leq 2S^* \cdot \lg n$, then the claim holds immediately. So, assume that $C > 2S^* \cdot \lg n$. From Lemma 4.3, there is a set of edges X with $|F(X)| \leq 2|X|$, where for each edge $e \in X$ it holds that $C_e(\mathbf{p}) \geq C - 2S^* \cdot \lg n$. Let $\Pi = \bigcup_{e \in X} \Pi_e(\mathbf{p})$, that is, Π is the set of players that their paths in \mathbf{p} use edges in X . Let $M = \sum_{e \in X} C_e(\mathbf{p})$ denote the total utilization of the edges in X . We have that $M \geq |X|(C - 2S^* \cdot \lg n)$.

For any player $i \in \Pi$ we have $C_i(\mathbf{p}) \geq C - 2S^* \cdot \lg n$. From Lemma 4.1, we obtain $S_i \leq S^* + 2S^* \cdot \lg n + 1 \leq 4S^* \cdot \lg n$. Let L denote the maximum path length of any player in Π . Since \mathcal{R} is a restricted-QoR game, for every player $i \in \Pi$ we have that $|p_i| \leq S_i$, which implies that $L \leq 4S^* \cdot \lg n$.

The parameter M can also be bounded as $M \leq L|\Pi|$. Consequently, $|X|(C - 2S^* \cdot \lg n) \leq L|\Pi|$, which gives:

$$C \leq \frac{L|\Pi|}{|X|} + 2S^* \cdot \lg n \leq \frac{4S^*|\Pi| \lg n}{|X|} + 2S^* \cdot \lg n \quad (1)$$

Since from Lemma 4.2, $|F(e, i)| \geq 1$, in the optimal routing \mathbf{p}^* the path of each user in Π has to use at least one edge in $F(X)$. Thus, in the optimal routing \mathbf{p}^* , the edges in $F(X)$ are used at least $|\Pi|$ times. Thus, there is some edge $e \in F(X)$ which in \mathbf{p}^* is used by at least $|\Pi|/|F(X)|$ paths. Since there are m service classes, the congestion of e in one of those service classes is at least $|\Pi|/(|F(X)| \cdot m)$. Therefore,

$C^* \geq |\Pi|/(|F(X)| \cdot m)$. Since $|F(X)| \leq 2|X|$, we obtain:

$$|\Pi| \leq 2C^* \cdot |X| \cdot m. \quad (2)$$

By Combining Equations 1 and 2, we get: $C \leq 8C^* \cdot S^* \cdot m \lg n + 2S^* \lg n \leq 10C^* \cdot S^* \cdot m \lg n$. \square

Finally, we obtain the main result:

THEOREM 4.5 (PRICE OF ANARCHY). *For any restricted-QoR game, $PoA = O(\min(C^*, S^*) \cdot m \lg n)$.*

PROOF. Suppose that \mathbf{p} is the worst Nash-routing with maximum social cost. We have that $PoA \leq SC(\mathbf{p})/SC(\mathbf{p}^*) \leq (C + S)/(C^* + S^*)$. We examine two cases:

- $C > S/4$: In this case $S = O(C)$. From Lemma 4.4, $C \leq 10C^* \cdot S^* \cdot m \lg n$. Therefore: $PoA = O(C/(C^* + S^*)) = O((C^* \cdot S^* \cdot m \lg n)/(C^* + S^*)) = O(\min(C^*, S^*) \cdot m \lg n)$.
- $C \leq S/4$: In this case $C = O(S)$. Let $p_i \in \mathbf{p}$ be the path with maximum cost in \mathbf{p} . Clearly, $pc_i(\mathbf{p}) = C_i + S_i \geq S$. Further, $0 \leq C_i(\mathbf{p}) \leq C \leq S/4$. Thus, $S_i \geq S - C_i \geq S - S/4 = 3S/4$. Since $C_i \geq 0 = C - C$, Lemma 4.1 gives $S_i \leq S_i^* + C + 1 \leq S^* + S/4 + 1$. Therefore, $3S/4 < S^* + S/4 + 1$. Thus, $S < 2(S^* + 1) \leq 3S^*$. In other words, $S = O(S^*)$. We obtain: $PoA = O(S/(C^* + S^*)) = O(S^*/(C^* + S^*)) = O(1)$.

By combining the two above cases we obtain the desirable result. \square

5. DISCUSSION

Our work has further consequences to the problem of routing and scheduling in general networks. Given a set of paths in a network, a well known lower bound for delivering the packets along the paths is $\Omega(C + D)$, where C is the maximum congestion in any edge, and D is the maximum path length. Leighton et al. [16] have shown the existence of packet scheduling algorithms that deliver packets along paths in time close to optimal $O(C + D)$ [5, 16, 17, 22, 24]. Our restricted-QoR games relate to this problem. We can assign the paths into routing classes according to their lengths. We can have $m = O(\log n)$ routing classes, where routing class Q_i can hold paths with lengths in range $[2^i, 2^{i+1})$, and the service cost of each class is 2^{i+1} . Clearly, the corresponding games are restricted. The Nash-equilibria for these games give poly-logarithmic approximations to the $C + S$ problem (when C^* or S^* is small), which in turn give poly-logarithmic approximations to the $C + D$ routing problem, when either C^* or D^* is small.

6. REFERENCES

- [1] R. Banner and A. Orda. Bottleneck routing games in communication networks. In *Proc. INFOCOM*, 2006.
- [2] C. Busch and M. Magdon-Ismail. Atomic routing games on maximum congestion. In *Proceedings of the 2nd International Conference on Algorithmic Aspects in Information and Management (AAIM)*, pages 79–91, Hong Kong, China, June 2006.
- [3] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proc. STOC*, 2005.

- [4] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Computational complexity, fairness, and the price of anarchy of the maximum latency problem. In *Proc. 10th Conf. on Integer Programming and Combinatorial Optimization (IPCO)*, 2004.
- [5] R. Cypher, F. Meyer auf der Heide, C. Scheideler, and B. Vöcking. Universal algorithms for store-and-forward and wormhole routing. In *In Proc. of the 28th ACM Symp. on Theory of Computing*, pages 356–365, 1996.
- [6] A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. In *Proc. STOC*, 2002.
- [7] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. SODA*, 2002.
- [8] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proc. ICALP*, 2002.
- [9] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. In *Proc. ICALP*, 2004.
- [10] M. Garing, T. Lücking, M. Mavronicolas, and B. Monien. Computing nash equilibria for scheduling on restricted parallel links. In *Proc. STOC*, 2004.
- [11] M. Garing, T. Lücking, M. Mavronicolas, and B. Monien. The price of anarchy for polynomial social cost. In *Proc. MFCS*, 2004.
- [12] M. Garing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. In *Proc. ICALP*, 2004.
- [13] M. Hollick, I. Martinovic, T. Krop, and I. Rimac. A survey on dependable routing in sensor networks, ad hoc networks, and cellular networks. In *Proc. 30th EUROMICRO Conference*, pages 495–502, 2004.
- [14] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate equilibria and ball fusion. In *Proc. SIROCCO*, 2002.
- [15] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. STACS*, 1999.
- [16] F. T. Leighton, B. M. Maggs, and S. B. Rao. Packet routing and job-scheduling in $O(\text{congestion} + \text{dilation})$ steps. *Combinatorica*, 14:167–186, 1994.
- [17] T. Leighton, B. Maggs, and A. W. Richa. Fast algorithms for finding $O(\text{congestion} + \text{dilation})$ packet routing schedules. *Combinatorica*, 19:375–401, 1999.
- [18] L. Libman and A. Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.
- [19] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *Proc. STACS*, 2004.
- [20] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. STOC*, 2001.
- [21] D. Monderer and L. S. Shapely. Potential games. *Games and Economic Behavior*, 1996.
- [22] R. Ostrovsky and Y. Rabani. Universal $O(\text{congestion} + \text{dilation} + \log^{1+\epsilon} N)$ local control packet switching algorithms. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing*, pages 644–653, New York, May 1997.
- [23] C. H. Papadimitriou. Algorithms, games and the internet. In *Proc. STOC*, pages 749–753, 2001.
- [24] Y. Rabani and É. Tardos. Distributed packet switching in arbitrary networks. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, pages 366–375, Philadelphia, Pennsylvania, 22–24 May 1996.
- [25] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 1973.
- [26] T. Roughgarden. The maximum latency of selfish routing. In *Proc. SODA*, 2004.
- [27] T. Roughgarden. Selfish routing with atomic players. In *Proc. SODA*, 2005.
- [28] T. Roughgarden and Éva Tardos. How bad is selfish routing. *Journal of the ACM*, 49(2):236–259, March 2002.
- [29] T. Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.
- [30] S. Suri, C. D. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. In *Proc. SPAA*, 2004.
- [31] Wikipedia. SCADA. <http://en.wikipedia.org/wiki/SCADA>.