



STABILITY, OPTIMALITY AND COMPLEXITY OF NETWORK GAMES WITH PRICING AND PLAYER DROPOUTS

ANDREW LOMONOSOV* AND MEERA SITHARAM†

Abstract. We study basic properties of a class of noncooperative games whose players are selfish, distributed users of a network and the game's broad objective is to optimize Quality of Service (QoS) provision. This class of games was previously introduced by the authors and is a generalization of well-studied network congestion games.

The overall goal is to determine a minimal set of static game rules based on pricing that result in stable and near optimal QoS provision.

We show the following. (i) Standard techniques for exhibiting stability or existence of Nash equilibria fail for these games—specifically, neither are the utility functions convex, nor does a generalized potential function exist. (ii) The problem of finding whether a specific game instance in this class has a Nash equilibrium is NP-complete.

To offset the apparent instability of these games, we show positive results. (iii) For natural subclasses of these games, although generalized potential functions do not exist, *approximate* Nash equilibria do exist and are easy to compute. (iv) These games perform well in terms of “price of stability” and “price of anarchy.” I.e., all of these approximate Nash equilibria nearly optimize a communal (or social) welfare function, and there is at least one Nash equilibrium that is optimal.

Finally, we give computer experiments illustrating the basic dynamics of these games which indicate that price thresholds could speed up convergence to Nash equilibria.

Key words. Congestion games, Selfish routing, Atomic unsplittable model, Nash Equilibria, Network pricing

1. Introduction. Recently much research has been done in applying game-theoretic concepts and general economics techniques to analysis of computer network traffic [2, 3, 5, 10, 11, 12, 16, 14, 20, 21, 24]. For a general survey see [1]. Stability in games refers to whether the game reaches a *Nash equilibrium*, a state where no player has incentive to move. Optimality is a measure of how close a Nash equilibrium is to optimizing a *social or communal welfare function*, usually the sum of the individual players' utility functions.

We consider primarily *atomic* games, where the number of players (network users) is finite. The case of *non-atomic games* where there is an infinite number of infinitesimally small players is easier to analyze. For similar reasons, *spittable* games, where network users can split their volume onto many service classes are easier to analyze and have more orderly behavior than *unsplittable games*, where each user is forced to place all their volume onto the same class.

The atomic splittable network game model has been studied [20, 12], with early results in the transportation literature. Efficiency (or optimality) of Nash equilibria in atomic splittable network games was studied in [24] and [28].

Here we consider primarily the unsplittable case that has also been studied for some time, for example [26].

Most of the research deals with *congestion games* where payoff to a player depends only on the player's strategy and on the number of players choosing the same strategy. Thanks to [26] it is known that such games always have Nash equilibrium. Two common techniques that are used to demonstrate existence of Nash Equilibria are the following. When the player utility functions are convex, Kakutani's fixed point theorem [25] directly shows existence. Also when such convexity properties are not present, *potential functions*, [18], certain functions that increase after every move, are used to show existence. These have a long history, for example, as Lyapunov stability functions classically used to describe equilibria in dynamical systems.

The [23] network games have realistic features that make them somewhat different from congestion games: in particular, players have non-convex utility functions caused by a threshold of total traffic volume in service classes that they are willing to tolerate. In addition in the [15] games, the players are allowed to refrain from participation, or to *dropout*, if their traffic quality demands are not satisfied. Hence existence of Nash equilibria or potential functions is not guaranteed for these classes of games. However, we were able to show existence of Nash equilibria for some of these classes of games by constructing *generalized potential functions*. (Generalized) potential functions have also been used by others to study versions of congestion and other games e.g., [7, 21, 22].

For the classes of games in [15, 16] we additionally showed that the Nash equilibria established via generalized potential functions are easy to compute. In general, however, while potential functions guarantee existence of Nash equilibrium, the problem of actually finding such an equilibrium remains computationally challenging.

*UGS Inc., 10824 Hope Street, Cypress, CA 90630 USA(lomonoso@ugs.com).

†CISE, University of Florida, Gainesville, FL 32611, USA(sitharam@cise.ufl.edu).

It has been shown [7] that the problem of finding Nash Equilibrium in congestion games is PLS-Complete, which intuitively means “as hard to compute as any object whose existence is guaranteed by a potential function”.

Considerable research has gone into the *price of anarchy* and *price of stability* of Nash equilibria [27]. These notions describe how far or how close Nash equilibria can be to the *System Optimum* of a game, where system optimum is a configuration (not necessarily a Nash equilibrium) that has greatest communal welfare.

We showed that for the classes of games with Nash equilibria in [15, 16], the communal welfare at these equilibria was poor, i. e., they are far from the system optimum. To rectify this, we further generalized our classes of games by introducing *pricing* incentives (not to be confused with the word “price” in the previous paragraph). The effect of pricing on congestion games has also been studied in [9, 6, 8]. Our original goal was to modify our original class of games so that the Nash equilibria would be close to system optima. However, the priced games were shown to not have Nash equilibria, in general. We instead showed that there is trade-off between game stability (existence of Nash Equilibria) and communal welfare achieved by such games. I.e., while the priced games did not always have Nash equilibria, the Nash equilibria, when they existed, were close to the system optima.

This trade-off has since been formalized by examining *approximate Nash equilibria* i. e. states where no player can improve their individual welfare by more than a certain factor, and the value of communal welfare at such approximate equilibria [4]. For example, [2] demonstrated a tradeoff between welfare and stability when costs functions are semiconvex.

In this paper, our overall goal is to analyze our classes of realistic network congestion games with respect to these stability and communal welfare measures; investigate *mechanisms* for games to optimize these measures; and to pose formal questions about the structure of game classes imposed by such measures.

More specifically, the original classes of games introduced in [15] were: the class \mathcal{Q} where players were solely motivated by their traffic quality demands and classes \mathcal{PQ} where players were also influenced by prices imposed on traffic. Stability of games in \mathcal{Q} was demonstrated by means of general potential functions, and concrete examples of instability of \mathcal{PQ} were then given.

In this paper, we establish the NP-completeness of determining existence of Nash equilibria and for computing Nash Equilibria in \mathcal{PQ} . We further study stability and communal welfare of (a modified version of) approximate Nash equilibria in \mathcal{PQ} , as compared to class \mathcal{Q} (i. e. effect of pricing on stability and social welfare in our games).

We also briefly look at game *dynamics*, i. e. number of steps that it actually takes to converge to Nash Equilibria for some of our games and conduct computer experiments to study trade-off between willingness to pay and speed of convergence.

Section 2 presents preliminary definitions, Section 3 presents previous results on the class \mathcal{Q} of games, Section 4 presents the main results of this paper concerning the class \mathcal{PQ} , and Section 5 concludes by tabulating and comparing the results of Sections 3 and 4, followed by open problems.

2. Definitions. A *game (instance)* G in the base class of QoS provision network games is specified by the *game parameters* $G = \langle n, m \in \mathbb{N}, \{\lambda_i \in \mathbb{R}^+ : 1 \leq i \leq n\}, \{b_{i,j} \in \mathbb{R}^+ : 1 \leq i \leq n; 1 \leq j \leq m\}, \{p_j : \mathbb{R}^+ \rightarrow \mathbb{R}, 1 \leq j \leq m\} \rangle$. The best way to define G is by identifying it with its finite game configuration graph (formally defined below) which consists of a set of feasible game configurations (vertices) and the valid or selfish game moves (oriented edges). The game G is played by n *users or players* each wanting to send a traffic of λ_i units through one of m network service classes and (for convenience of analysis) an overflow or Dummy Class with index 0, referred to as *DC*. Each player i additionally has a *volume threshold* $b_{i,j}$ (to be described below) for each class j . A *price* function $p_j(\cdot)$ for each service class is a nonincreasing function that maps the total (traffic) volume in the class to a unit price. (Unit price typically decreases with increasing congestion or total volume in any service class). The price for using DC is 0. A *feasible configuration* Λ of G is fully specified by an allocation $J_\Lambda : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ which describes which service class $J_\Lambda(i)$ that the user or *player* i has decided to place their chunk λ_i of traffic. This allocation J_Λ results in a *total traffic volume* $q_{\Lambda,j} = \sum_{i:1 \leq i \leq n \wedge J_\Lambda(i)=j} \lambda_i$ in each class $1 \leq j \leq m$ at the game configuration Λ . The set of feasible game configurations \bar{F} form the *vertex* set of the *game configuration graph* Ω . *Individual utility function* $U_i(\Lambda)$ is a type of step function based on i 's volume threshold being met at the configuration Λ , and on the unit price incurred by the player i in its class $j = J_\Lambda(i)$. $U_i(\Lambda)$ is:

- 0 if $j = 0$ (user i is in DC)

- $-\epsilon$, for small $\epsilon > 0$ if $b_{i,j} < q_{\Lambda,j}$ (volume threshold exceeded)
- equal to $\lambda_i(1 - p_j q_{\Lambda,j})$ otherwise.

It is assumed that the price functions are always appropriately normalized so that this quantity is always *strictly positive* for all players i and their classes $J_{\Lambda}(i)$ at any configuration Λ . A typical utility function is shown on Figure 2.1. We say that user i is *satisfied* at configuration Λ if $U_i(\Lambda) \neq 0$, and not satisfied otherwise. We define




FIG. 2.1. Utility as a function of volume, volume threshold and price

a function $Sat_{\Lambda}(i) = 1$ if $U_{\Lambda}(i) \neq 0$, otherwise $Sat_{\Lambda}(i) = 0$. A *selfish move* by user i at a configuration Λ_1 is a reallocation of i 's volume λ_i from a departure class j_1 (i.e. $J_{\Lambda_1}(i) = j_1$), to a destination class j_2 resulting in a configuration Λ_2 (i.e. $J_{\Lambda_2}(i) = j_2$) that increases utility of this user, i.e. $U_i(\Lambda_1) < U_i(\Lambda_2)$. Moves to DC by a user whose volume threshold is exceeded are called *user dropouts*. Note that user dropouts qualify as selfish moves according to our definition.

Each selfish move is an ordered pair of feasible game configurations (for example $(\Lambda_1, \Lambda_2) \in F \times F$), and represents an *oriented edge* of the game configuration graph Ω . A *generalized potential function* is a function defined on configurations that increases after every player move. A *game play* for G is a sequence of valid selfish moves in G , i.e. $(\Lambda_1, \Lambda_2), (\Lambda_2, \Lambda_3), \dots, (\Lambda_{k-1}, \Lambda_k)$, or a *path* in the game configuration graph Ω . A *Nash Equilibrium* or *NE* of a game G is a configuration Λ such that there is no selfish move possible for any user i . Nash equilibria are exactly sink vertices of a game configuration graph Ω that have no outgoing edges toward other vertices. For our classes of games, the *communal welfare function* for configuration Λ is defined as $C(\Lambda) = \sum_i Sat_{\Lambda}(i)\lambda_i$. The feasible game configuration that has highest value of communal welfare function is called the *System Optimum* or *SO*. Let Λ_N be a Nash Equilibrium that has the smallest value of communal welfare function taken over all Nash Equilibria, while Λ_M be a Nash Equilibrium that has the largest value. As defined in say [27] a *price of anarchy* of a game is equal to $C(\Lambda_N)/C(\Lambda_*)$, where Λ_* is SO. A *price of stability* is equal to $C(\Lambda_M)/C(\Lambda_*)$.

Class of games that do not have pricing, i. e. $p_j(x) = 0$ for all classes j and their volumes x is denoted by \mathcal{Q} . In such games players are motivated only by their desire to satisfy their volume thresholds. Subclass $\mathcal{Q}_{\mathcal{E}} \subset \mathcal{Q}$ is a class of games with no pricing where all players have equal volume. Class of games that have only one pricing function $p(x)$ for all classes j and this function is strictly decreasing ($p(x) < p(y) \leftrightarrow x > y$) is denoted by \mathcal{PQ} . Subclass $\mathcal{PQ}_{\mathcal{E}} \subset \mathcal{PQ}$ is a class of games with single strictly decreasing price function where all players have equal volume. Here we will give a pictorial example, Figure 2.2, of some notions introduced in this section. A game configuration graph Ω and configurations Λ of a particular game G are shown. Columns represent classes, rectangles represent users, the size of a rectangle corresponds to volume of a user, volume thresholds of users are indicated on the right. In this example the game G in class \mathcal{PQ} has 2 classes, 2 users A and B that have equal volumes and the volume threshold of A is greater than that of B . Game configuration graph Ω has 4 vertices. This game G has no Nash equilibrium.

Throughout this paper we assume wlog that every player i has the same volume threshold $b_i = b_{i,1} = b_{i,2} = \dots b_{i,m}$ in every class $j = 1 \dots m$. We also assume that players are sorted in the increasing order of their thresholds, i.e. $b_1 \leq b_2 \leq \dots \leq b_n$. (The former assumption could be easily generalized for all results in this paper, the latter assumption is realistic and commonly made [23]).

In proofs when describing a game configuration Λ , we will specify values of game parameters n and m , provide a list of users in the form User(Volume, Volume Threshold) (for example $A(5,12)$ means that User A has volume 5 and volume threshold 12), as well as specify where these users are, i.e. $\{J_{\Lambda}(i)\}$.




FIG. 2.2. Game configuration graph and individual configurations

3. Previously known properties of \mathcal{Q} . We list relevant properties of the class \mathcal{Q} of games established in [15] concerning existence, optimality and complexity of computing Nash equilibria.

THEOREM 3.1. *Every game in \mathcal{Q} has a generalized potential function and therefore every such game has a Nash Equilibrium.*

THEOREM 3.2. *For any $\epsilon > 0$ there is a game in \mathcal{Q} that has price of anarchy and price of stability equal to ϵ .*

THEOREM 3.3. *A Nash Equilibrium that is also a System Optimum of a game in \mathcal{Q}_ϵ can be found in time linear in the game parameters.*

THEOREM 3.4. *Any Nash Equilibrium of any game $G \in \mathcal{Q}_\epsilon$ has communal welfare of at least a half of that of G 's System Optimum.*

THEOREM 3.5. *For any initial configuration of every game in \mathcal{Q}_ϵ there is a sequence of selfish moves by players that will terminate at Nash Equilibrium after $O(n^2)$ steps. This sequence can be determined by considering players in decreasing order of their volume thresholds and letting them make their selfish choices.*

4. New results. In this section we consider stability of games in class \mathcal{PQ} and various properties of their Nash equilibria. Results will be compared to those of \mathcal{Q} in Table 5.

We begin by establishing the following simple result about the prices of anarchy and stability of general games in the class \mathcal{PQ} , showing that they are not particularly well behaved.

THEOREM 4.1. *For any $\epsilon > 0$ there is a game in \mathcal{PQ} that has a unique Nash equilibrium, whose communal welfare is ϵ , while the system optimum of this game has communal welfare equal to 1. This implies that prices of anarchy and stability of such a game are equal to ϵ .*

Proof. Consider a game with one non-DC class, and two players, $A(\epsilon, 1+\epsilon)$ and $B(1, 1)$. The only equilibrium this game has is when player A is in class 1 and player B is in DC, as opposed to the system optimum when their positions are reversed. \square

4.1. Approximate Nash equilibria. As we have noted in the Introduction and Figure 2.2, Nash equilibria do not necessarily exist in games \mathcal{PQ} that involve pricing. One approach to examining such games involves α -approximate Nash equilibria, defined in for example [4]. A configuration is said to be α -approximate Nash equilibrium if no player can move and decrease her cost by more than an α multiplicative factor.

Note that since pricing functions of \mathcal{PQ} are arbitrary decreasing linear functions, we will instead use a more appropriate notion of δ -approximate Nash equilibrium instead, where δ is an additive factor.

Let \mathcal{PQ}_ϵ be the subset of \mathcal{PQ} where all players have volume $\epsilon = \delta$. In such a game a configuration where all players are satisfied and all classes have equal total volume would be a ϵ -approximate Nash equilibrium, since no player would have an incentive to move.

When ϵ goes to zero and number of players goes to infinity, the class \mathcal{PQ}_ϵ will be denoted as \mathcal{PQ}_∞ . This class of games has similar behavior to the class of games where players are allowed to split their volume between several classes.

THEOREM 4.2. *A Nash equilibrium (δ -approximate Nash equilibrium) that is also system optimum can be constructed for any game in \mathcal{PQ}_∞ (\mathcal{PQ}_ϵ) in time of $O(n)$.*

Proof. A greedy algorithm solves this problem. Here is the algorithm for \mathcal{PQ}_ϵ . Let $b_1 \leq \dots \leq b_n$; place player n in class 1, place player $n-1$ in class 1 if $b_{n-1} \geq 2\epsilon$, otherwise place player $n-1$ in class 2; place player

$n - 2$ in class 1 if $b_{n-2} \geq 3\epsilon$ etc. The resulting configuration is a system optimum and a δ -approximate Nash equilibrium. \square

Note that while the preceding theorem guarantees existence of an approximate Nash equilibrium for games \mathcal{PQ}_ϵ , it does not promise that *every* sequence of selfish moves will arrive at an approximate Nash equilibrium. Consider the following observation, which also disproves existence of general potential functions for all games in \mathcal{PQ}_ϵ . This is also true for games in \mathcal{PQ}_∞ .

THEOREM 4.3. *There is a game in \mathcal{PQ}_ϵ where there is a cycle of selfish moves.*

Proof. Let $\delta = 1$. Consider a game with 2 non-DC classes and 12 players:

$$A_1(1, 9), A_2(1, 9), A_3(1, 9), B_1(1, 6), B_2(1, 6), B_3(1, 6), C_1(1, 3), \dots, C_6(1, 3).$$

Initial configuration Λ : players C_4, C_5 and C_6 are in class 2, all other players are in class 1. First players B_1, B_2 and B_3 move to class 2, after that players C_1, C_2, C_3 move to DC, then players A_1, A_2 and A_3 move to class 2 and finally players C_1, C_2, C_3 move from DC to class 1. The resulting configuration is essentially isomorphic to Λ , hence a cycle has occurred. \square

Now we will examine properties of corresponding Nash equilibria.

THEOREM 4.4. *Price of anarchy of games in \mathcal{PQ}_∞ is equal to $1/2$. Price of stability of such games is equal 1.*

If price of anarchy and price of stability were redefined over ϵ -approximate Nash equilibria instead of regular Nash equilibria, then it would hold that price of anarchy of games in \mathcal{PQ}_ϵ is equal to $1/2$ and price of stability of such games is equal 1.

Proof. Price of stability follows from the fact that Nash equilibria constructed in Theorem 4.2 are system optima.

Price of anarchy can be demonstrated by following argument for games in \mathcal{PQ}_ϵ , and the proof for \mathcal{PQ}_∞ is similar. Let Λ be a Nash equilibrium when all players have the same volume ϵ . Consider the unsatisfied player i that has the largest volume threshold b_i . (If there are no unsatisfied players then such a Nash equilibrium is a system optimum). Total traffic volume q_j in every class j is strictly greater than $b_i - \epsilon$, hence communal welfare of Λ is greater than or equal to $m(b_i - \epsilon)$ but communal welfare of system optimum cannot be more than $2(m(b_i - \epsilon))$. \square

4.2. Finding a Nash equilibrium. It was shown in [16] that the problem of finding system optimum of a game in class \mathcal{Q} is NP-Complete. It was also shown that the problem of finding a Nash equilibrium in \mathcal{Q} can be solved in $O(n^2)$ time. Similarly the problem of finding a system optimum of a game in class \mathcal{PQ} is NP-Complete. Now we will examine the problem of finding a Nash equilibrium (or determining that it does not exist) for games in \mathcal{PQ} .

THEOREM 4.5. *Problem of finding Nash equilibrium for games in \mathcal{PQ} is NP-Complete.*

Proof. Consider the following version of MAXIMUM SUBSET SUM problem—given set $S = \{s_1, \dots, s_n\}$ and targets t_1, t_2 , find $A \subseteq S$ such that $t_1 \leq \sum_{i \in A} s_i \leq t_2$. This problem can be reduced to problem of finding a Nash equilibrium as follows. There are $n + 1$ players and two non-DC classes. Players $1, \dots, n$ all have same threshold $b_1 = b_2 = \dots = b_n = t_2$, individual volumes $\lambda_i = s_i$. Player $n + 1$ has volume $\lambda_{n+1} = t_2$ and threshold $b_{n+1} = t_1 + t_2$. Then this game will have a Nash equilibrium if and only if the original MAXIMUM SUBSET SUM problem had a feasible solution. \square

4.3. Price thresholds. In [16] it was shown that games in class \mathcal{Q} will terminate in $O(n^2)$ steps, given certain assumptions on order of player moves. Here we will describe a computer experiment that examined speed of convergence of games where there was no such ordering of player moves.

This experiment involved a following natural assumption about players behavior. In practice, there could be a limit on how much a user is willing to pay, and this concept can be easily added to our games, resulting in the new classes of games. This concept has a desirable effect on the dynamics of the game, as explained below. Formally, for players i we define *price thresholds* (in addition to the old volume thresholds) t_i that have the following property. If the price in a class exceeds player i 's price threshold, then player i is not satisfied. We assume that $b_i \leq b_j$ if and only if $t_i \geq t_j$, i.e users who demand better quality of service (smaller traffic volume in their class) are willing to pay more.

We conjecture that in addition to being realistic, such price thresholds also tend to improve the speed of convergence to Nash equilibria. This is because of players spending less time looping in non-terminal cycles.

To test this conjecture we ran a computer program simulating a game in class \mathcal{PQ} . Later we added pricing thresholds to the game which has considerably improved time lapsed before convergence to Nash equilibria. Game parameters were chosen such that Nash equilibrium would always exist. Parameters of the game were M = number of classes, M/T = number of *types* of users that have the same volume and volume threshold, K = number of users of the same type that can fit in one class without exceeding their volume threshold. Volumes were in increments of one, i.e there are $T * K$ users that have volume 1 and volume threshold K , $T * K$ users that have volume 2 and threshold $2K$, \dots , $T * K$ users that have volume M/T and threshold $M * K/T$. Thus there are a total of $M * K$ users. For example let $K = 10, M = 20, T = 5$. This means that there are 20 classes, 4 types of users and at most 10 users of any one type can fit into one class. Users are

$$A_1(1, 10), \dots, A_{50}(1, 10), B_1(2, 20), \dots, B_{50}(2, 20), C_1(3, 30), \dots, C_{50}(3, 30), \\ D_1(4, 40), \dots, D_{50}(4, 40).$$

Initially all users are in the dummy class (DC). A game proceeds by picking one of the $M * K$ users at random and this user moves either to the largest class where his threshold would not be exceeded or to the DC. Even if this move exceeds the volume threshold of some other users in the destination class of the moving user, these unsatisfied users cannot move until it is their turn to move and turns are determined at random. Eventually a Nash equilibrium was always reached, where all users of the first type were in T classes, all users of the second type were in the second set of T classes etc. Results are shown in table 4.1. ‘‘Moves1’’ denotes the total number of user moves until Nash equilibrium was reached.

Later a simulation of pricing thresholds was added to the experiment. Effectively it would prohibit a user i that has volume threshold b_i to move into any class j such that $q_j + \lambda_i < b_i - \Delta$ where Δ is some constant. The reason for this is that class j is too expensive for the i^{th} user.

TABLE 4.1

K	M	T	Moves1	Δ	Moves2
5	20	1	161,000	5	7,000
10	20	1	17,077,000	10	9,000
20	20	2	1,354,000	20	25,000
50	20	1	56,000	50	35,000
100	20	1	49,000	100	46,000
100	20	10	3,000	100	5,000
1000	20	10	35,000	1000	49000
5	40	1	2,360,000	5	190,000
5	50	1	8,391,000	5	940000

When $\Delta = \infty$ this is equivalent to the old experiment without pricing thresholds. In general introduction of small Δ significantly improved number of moves that was needed to reach the Nash equilibrium. See ‘‘Moves2’’ in the table 4.1.

5. Conclusions, Directions. Here we summarize known results about Nash Equilibria for various subclasses of \mathcal{Q} and \mathcal{PQ} .

	NE/GenPotential always exists	Price of anarchy	Price of stability	Complexity of finding NE
\mathcal{Q}	Yes/Yes	ϵ	ϵ	$O(n^2)$
\mathcal{Q}_ϵ	Yes/Yes	1/2	1	$O(n)$
\mathcal{PQ}	No/No	ϵ	ϵ	NP-Complete
\mathcal{PQ}_ϵ	Yes/No	1/2	1	$O(n)$
\mathcal{PQ}_∞	Yes/No	1/2	1	$O(n)$

Existence of Nash Equilibria for \mathcal{Q} (and \mathcal{Q}_ϵ , since $\mathcal{Q}_\epsilon \subset \mathcal{Q}$) is shown in Theorem 3.1. Example of nonexistence of Nash Equilibria in \mathcal{PQ} is demonstrated in Figure 2.2. For \mathcal{PQ}_ϵ entry ‘‘Yes’’ refers to δ -approximate Nash Equilibria, not regular Nash Equilibria. This (and \mathcal{PQ}_∞ case) is shown in Theorem 4.2. The nonexistence

of generalized potential functions for these classes is shown in Theorem 4.3. Prices of anarchy and stability of \mathcal{Q} are shown in Theorem 3.2, of $\mathcal{Q}_{\mathcal{E}}$ in Theorem 3.4, of \mathcal{PQ} in Theorem 4.1 (assuming that Nash Equilibrium exists), of $\mathcal{PQ}_{\mathcal{E}}$ and \mathcal{PQ}_{∞} in Theorem 4.4. Complexity of finding a Nash Equilibrium in games of class \mathcal{Q} is shown in Theorem 3.5, case of $\mathcal{Q}_{\mathcal{E}}$ is Theorem 3.3, for games in \mathcal{PQ} this problem is NP-Complete (Theorem 4.5), for games in $\mathcal{PQ}_{\mathcal{E}}$ and \mathcal{PQ}_{∞} result follows from Theorem 4.2.

5.1. Open questions. The class \mathcal{PQ} contains both games that have Nash equilibria and those who do not.

What is the structure of games in class \mathcal{PQ} where Nash equilibria or approximate Nash equilibria (additive or multiplicative) are guaranteed to exist but they are hard to compute? For example, are there PLS-complete games in the class \mathcal{PQ} ? For the subclasses such as $\mathcal{PQ}_{\mathcal{E}}$ Nash equilibria existence is easy to determine, and (approximate) Nash equilibria are easy to compute. Formally state and prove the conjecture of Section 4.3 concerning the usage of price thresholds and speed of convergence to Nash equilibria.

REFERENCES

- [1] E. ALTMAN, T. BOULOGNE, R. EL-AZOUZI, T. JIMENEZ AND L. WYNTER, *A survey on networking games in telecommunications*, Comput. Oper. Res., pages 286–311, 2006.
- [2] E. ANSHELEVICH, A. DASGUPTA, J. KLEINBERG, E. TARDOS, T. WEXLER AND T. ROUGHGARDEN, *The Price of Stability for Network Design with Fair Cost Allocation*, FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04), pages 295–304, 2004.
- [3] B. AWERBUCH, Y. AZAR AND A. EPSTEIN, *Large The price of routing unsplittable flow*, STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing, pages 57–66, 2005.
- [4] H. CHEN AND T. ROUGHGARDEN, *Network design with weighted players*, ACM Symposium on parallel algorithms and architecture, pages 29–38, 2006.
- [5] R. COCCHI AND D. ESTRIN AND S. SHENKER, *Pricing in Computer Networks: Motivation*, IEEE/ACM Transactions on Networking, Vol. 1, No. 6, December 1993, pages 614–627.
- [6] R. COLE, Y. DODIS AND T. ROUGHGARDEN, *Pricing network edges for heterogeneous selfish users*, STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, pages 521–530, 2003.
- [7] A. FABRIKANT, C. PAPADIMITRIOU AND K. TALWAR, *The complexity of pure Nash equilibria*, STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, pages 604–612, 2004.
- [8] L. FLEISCHER, K. JAIN AND M. MAHDIAN, *Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games*, FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04), pages 277–285, 2004.
- [9] L. FLEISCHER, *Linear tolls suffice: new bounds and algorithms for tolls in single source networks*, Theor. Comput. Sci., (348), pages 217–225, 2005.
- [10] D. FOTAKIS, S. C. KONTOGIANNIS, E. KOUTSOPIAS, M. MAVRONICOLAS AND P. G. SPIRAKIS, *The Structure and Complexity of Nash Equilibria for a Selfish Routing Game*, ICALP '02: Proceedings of the 29th International Colloquium on Automata, Languages and Programming, pages 123–134, 2002.
- [11] D. FOTAKIS, S. KONTOGIANNIS, P. SPIRAKIS, *Selfish unsplittable flows*, Theor. Comput. Sci., (348), pages 226–239, 2005.
- [12] Y. KORILIS, A. LAZAR, AND A. ORDA, *Architecting noncooperative networks*, IEEE Journal of Selected Areas in Communications, (13) pages 1241–1251, 1995.
- [13] E. KOUTSOPIAS AND C. H. PAPADIMITRIOU, *Worst-case equilibria*, Symposium on Theoretical Aspects of Computer Science, pages 404–413, 1999.
- [14] ———, *Optimal transport strategies for best-effort traffic over priced connections*, Technical report, Technion, 2000.
- [15] A. LOMONOSOV, M. SITHARAM AND K. PARK, *Stability vs optimality tradeoff in game theoretic mechanisms for QoS provision*, SAC '03: Proceedings of the 2003 ACM symposium on Applied computing, pages 28–32, 2003.
- [16] ———, *Network QoS games: stability vs optimality tradeoff*, J. Comput. Syst. Sci., (69), pages 281–302, 2004.
- [17] I. MILCHTAICH, *Congestion games with player-specific payoff functions*, Games and Economic Behavior (13), pages 111–124, 1996.
- [18] D. MONDERER AND L. SHAPLEY, *Potential Games*, Games and Economic Behavior (14), pages 124–143, 1996.
- [19] N. NISAN AND A. RONEN, *Algorithmic mechanism design*, Proc. 31st ACM Symp. on Theory of Computing, pages 129–140, 1999.
- [20] A. ORDA, R. ROM, AND N. SHIMKIN, *Competitive routing in multiuser communication networks*, IEEE/ACM Transactions on Networking, (1), pages 510–521, 1993.
- [21] ———, *Computing correlated equilibria in multi-player games*, STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing, pages 49–56, 2005.
- [22] C. PAPADIMITRIOU AND T. ROUGHGARDEN, *Computing equilibria in multi-player games*, SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, pages 82–91, 2005.
- [23] K. PARK, M. SITHARAM, AND S. CHEN, *Quality of service provision in noncooperative networks: Heterogeneous preferences*, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, 1998.
- [24] K. PARK, M. SITHARAM, AND S. CHEN, *Quality of service provision in noncooperative networks with diverse user requirements*, Decision Support Systems, Special Issue on Information and Computation Economics, vol. 28, pages 101–122, 2000.
- [25] J. B. ROSEN, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica, pages 520–534, 1965.

- [26] R. W. ROSENTHAL, *A Class of games possessing pure strategy nash equilibria*, International Journal of Game Theory, pages 65–67, 1973.
- [27] T. ROUGHGARDEN, *Selfish Routing and the Price of Anarchy*, The MIT Press, 2005.
- [28] T. ROUGHGARDEN AND E. TARDOS, *How bad is selfish routing?* IEEE Symposium on Foundations of Computer Science, pages 93–102, 2000.
- [29] ———, *Mechanism design and the Internet*, Presentation in DIMACS Workshop on Computational Issues in Game Theory and Mechanism Design, 2001.

Edited by: Marcin Paprzycki, Niranjan Suri

Received: October 1, 2006

Accepted: December 10, 2006