Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Journal of Computer and System Sciences 69 (2004) 281–302
JOURNAL OF COMPUTER AND SYSTEM SCIENCES
http://www.elsevier.com/locate/jcss

Network QoS games: stability vs optimality tradeoff

Andrew Lomonosov,[a] Meera Sitharam,[a,*,1] and Kihong Park[b,2]

[a] CISE, University of Florida, Gainesville, FL 32611-6120, USA
[b] Department of Computer Sciences, Purdue University, West Lafayette, IN 47907-1398, USA

Received 14 February 2002; revised 5 March 2004

Abstract

We study 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. Our classes of games are based on realistic microeconomic market models of QoS provision (Proceedings of the First International Conference on Information and Computation Economics ICE'98, 1998) and have two competing characteristics—stability and optimality. Stability refers to whether the game reaches a Nash equilibrium. Optimality is a measure of how close a Nash equilibrium is to optimizing a given objective function defined on game configuration. The overall goal is to determine a minimal set of static game rules based on pricing that result in stable and efficient QoS provision. We give a new and general technique to establish stability and demonstrate a close trade-off between stability and optimality for our game classes. We also state several open problems and directions together with initial observations and conjectures.

1. Introduction

QoS provision and network resource allocation are problems relevant to Internet usage. One approach by the networking [1,6,5,12,27,26,21] research community over the past several years is to use a microeconomic model: treat the network as a market and its users and providers as players of a *noncooperative game* [9,6,5,22]. A number of related, fundamental issues have been isolated—in algorithmic mechanism design, computational aspects of game theory, and complexity of distributed computing and communication—that are of interest to theoreticians [2,16,17,19,20,25] and potentially have other applications as well.

*Corresponding author.
E-mail address: sitharam@cise.ufl.edu (M. Sitharam).
[1] Supported in part by NSF Grant EIA0096104
[2] Supported in part by NSF Grants ANI9714707 and ANI9875789

doi:10.1016/j.jcss.2004.03.001

The overall goal of noncooperative game theoretic modeling is to design a game that permits the users and providers of a network (or their agents) to behave like selfish and distributed *players* [6,18,27], realistically and fairly, with minimal intervention by any external network manager. On the other hand, despite this market anarchy, at natural *equilibrium game configurations*, this game should result in "desirable" overall QoS provision and resource allocation or assignment.

To a theoretician, one valid view of noncooperative game theoretic approaches and algorithmic mechanism design approaches to network problems is that they are simply paradigms for designing efficient algorithms [2,17,25] for distributed optimization (or approximation) on a network. Within this view, the game, i.e., the *feasible game configurations*, the players, their *utility* and reward or *pricing* functions, their *selfish moves* or dominant strategies are all free to be defined in any computationally meaningful manner.

In this paper however, we adopt the network modeler's [6,18,27] point of view that these definitions should— in addition—correspond to a meaningful, realistic and fair market architecture for the users and providers of network resources.

One difficult issue is a precise definition of a "desirable" game configuration, which takes many forms. One purely market-based point of view is that "a desirable outcome is simply any natural outcome of a fair and selfish game—further interference is undesirable." Once constraints are imposed on the rules of the game (fairness, personal freedom, and efficiency of individual moves) thereafter any equilibrium that this game naturally reaches should be accepted as desirable. The common type of equilibria studied in this context are the so-called *Nash equilibria*, defined as configurations where none of the players individually has any (selfish) reason to make a move. In a mechanism design framework, Nash equilibria often automatically correspond to a so-called "social choice" function [17] that aggregates (privately known) preferences of many people into a consistent social choice configuration. Sometimes "desirable" is defined as configurations that optimize a *communal welfare* function, optionally subject to constraints based on equitable distribution, collective efficiency etc. [3,8]; or as configurations that satisfy a prescribed set of constraints arising from measures of fairness, freedom etc; or as a combination of the two: configurations that optimize a well-defined function, subject to a set of constraints. In these cases, the game design problem is closely related to mechanism design optimization problems: I.e., to obtain a social choice function that in addition maps to desirable configurations [4]. In the game context, how to guide a selfish game towards desirable configurations, i.e., to design (realistic, fair, typically *pricing* based [17] incentive or reward) functions that alter the players' personal utility functions in such a way that their purely selfish behavior according to the altered utilities results in (Nash) equilibria that have the desirable properties.

In this paper, we consider a simple (and commonly used [6,27]) communal welfare function defined simply as the sum of the individual player's utilities (volume-adjusted and minus prices); we then design *pricing* incentive functions that result in Nash equilibria that (approximately) optimize communal welfare.

Many interesting problems lend themselves to a *static game* approach, i.e., one defines the game by specifying the set of feasible game configurations, individual player utility functions, pricing incentive functions and selfish moves or strategies, and thereafter simply studies the relevant properties of Nash equilibria. Other computational problems arise from imposing *dynamic* rules on (a discretized version of the) game such as the order or frequency of player moves. This translates to interpreting the game configurations and selfish moves as the vertices and edges of a finite *game configuration graph*, studying the lengths of particular paths in this graph, which represent *game plays*, or interpreting the graph as a Markov chain with probabilities attached to the edges or moves. In the latter randomized setting, one problem is determining the stationary probability distribution on (necessarily Nash or terminal cycle) configurations, given a natural initial distribution, and thereby determine properties of Nash configurations that hold with high probability. In both the randomized and deterministic settings, a complexity issue of interest is the time taken for *convergence* of game plays to Nash equilibria, terminal cycles, or to a stationary distribution.

Our approach in this paper is primarily static, although we touch upon simple (deterministic) dynamic aspects.

Remark 1. Further interesting issues in network resource allocation and QoS provision games—which we do not emphasize in this paper—are: game sensitivity to a small changes in total resources, disclosure of information by players and game outcome, computational complexity of the player utility functions, and the pricing function, etc. These and other issues have been listed in a comprehensive DIMACS talk [28].

One issue that is however usually ignored in the literature is *stability*: Does the game have stable configurations, i.e., Nash equilibria at all? Or are there only *terminal cycles* in the game configuration graph. i.e., a set of at least two configurations and a cyclic sequence of moves between them that the players are trapped into traversing indefinitely if they always choose their selfishly optimal move. It is usually assumed [6,5,18] that Nash equilibria always exist, and that there is a path from every game configuration to a Nash equilibrium (which ensures convergence). One rationale for this assumption relies on a version of Brouwer's fixed point theorem called Kakutani's theorem which states that if the player's selfish moves are based on maximizing utility functions that are quasiconcave, it follows that a Nash equilibrium—which is a type of fixed point—always exists.

This assumption was challenged by [22], where it was shown that for a natural class of games their realistic utility functions—based on a commonly used network model—are not quasiconcave and result in natural QoS provision games that may not have Nash equilibria.

In this paper, we show that the stability question for practically realistic classes of QoS network games of [22] gives rise to potentially fundamental new problems and techniques. Our main contributions here are described in the following section.

Remark 2. We do not practically justify our base classes of QoS network games, relate them to other commonly used classes of network games, nor provide the fundamental reason why our games cannot be assumed to have guaranteed stability. All of these issues were discussed extensively in [22] and have been generally accepted and cited [23,13–15,24,10,7].

2. Description of results

1. In Theorem 5 we give a simple but general technique to establish the stability of game classes and to establish properties of the game configuration graph such as the existence of cycles and the

existence of paths from an arbitrary configuration to a Nash configuration. In other theorems, we apply this technique to establish the stability of various realistic classes of QoS provision games based on [22].

We also use this theorem to classify all network games based on their stability. Later, this classification is illustrated by concrete examples.

2. For these classes of games, we prove a series of results that demonstrate a close tradeoff between stability and optimality. In classes of games that are stable, i.e., where Nash equilibria are guaranteed to exist, they could be far from optimizing the communal welfare function. However, when we systematically alter such a game class to ensure that Nash equilibria are a reasonable approximation of the communal welfare optimum, then the games in the altered class are no longer guaranteed to be stable, i.e., they may not have Nash equilibria at all. In particular, we show the following:

(i) Theorem 7, and Observation 1 show that a realistic class $\mathscr{Q}$ of QoS provision games from [22] (that is formally defined later and does not use pricing functions to alter user utilities) has guaranteed stability, but Nash equilibria may be arbitrarily far from optimizing communal welfare.

(ii) Observations 2, 3 and Theorem 11 show that on expanding $\mathscr{Q}$ to a class of games $\mathscr{PQ}$ (and its natural extension $\mathscr{SPQ}$) by adding a single realistic type of pricing function to all of the individual player utilities, the new class of games is no longer guaranteed to have Nash equilibria. We additionally give examples of cases when Nash equilibria coexist with games cycles. However, when Nash equilibria do exist for games in class $\mathscr{PQ}$, these equilibria achieve optimal communal welfare, under certain conditions C on the parameters of the game. When the conditions C do not hold, arbitrarily suboptimal counterexamples exist.

(iii) These conditions C have both a practical and theoretical justification. Theorem 8 demonstrates the latter: optimization of communal welfare—over all feasible network (QoS provision) configurations—is a computational problem independent of any game-theoretic context. This optimization problem, which we call CW is NP-complete (can be seen as a general version of SUBSET SUM) and the set C arises as a natural set of conditions on the input parameters for which a greedy approach gives an optimal solution.

The greedy approach however is traditionally algorithmic, i.e., it dictates a strict *sequence or order* of steps that is *crucial* for arriving at the solution to CW. One standard interpretation of our type of result is that our games provide a more self-organizing, less externally dictated method that results in solutions to CW, under the same conditions C. In other words, simply designing the *static* rules of $\mathscr{PQ}$ games appropriately, *any* terminating sequence of valid game moves—no matter what their order—in fact terminates in a solution to the optimization problem CW. I.e., just the static property of being a Nash equilibrium configuration of a $\mathscr{PQ}$ game makes it a solution to the optimization problem CW. This could also be viewed as a type of Church-Rosser property.

(iv) Theorem 13 and associated Observation 4 show that on restricting $\mathscr{PQ}$ to a class $\mathscr{HPQ}$ of games by placing constraints on player moves that hurt other players, we guarantee stability again at the cost of deteriorated communal welfare at Nash configurations.

(v) Theorems 15 and 16 show that by modifying $\mathscr{PQ}$ (which forces the use of a single pricing function), to a class $\mathscr{DPQ}$ of games that use several carefully ordered price functions simultaneously, we guarantee both stability and optimal communal welfare at Nash configurations, under the conditions C.

3. For all of these game classes Theorem 17 introduces simple *dynamic* rules that impose a priority order in which the players take turns for moves: these guarantee fast convergence of *game-plays* to the Nash configurations. Particularly for the classes $\mathscr{P}\mathscr{Q}$ and $\mathscr{D}\mathscr{P}\mathscr{Q}$ under the conditions C, these simple dynamic rules give an equally efficient but nevertheless less dictated alternative to greedy algorithms for communal welfare optimization.

4. Finally, we state several open problems, conjectures and directions for extending our results and motivate them by initial observations.

2.1. Organization

In Section 3 we formally define our base classes of QoS Provision network games and the essential terminology appearing in italics in the Introduction. In Section 4, we demonstrate a technique for proving existence of Nash equilibria. In Section 5, we prove the main results described above. In Section 6, we state a rule that allows for a rapid convergence to Nash equilibria. In the final Section 7, we discuss open problems, conjectures, interesting directions and initial results.

3. Definitions

A *game (instance)* G in the base class of QoS provision network games is specified by the *game parameters* $G = \langle n, m \in N, \{\lambda_i \in R^+ : 1 \leqslant i \leqslant n\}, \{b_{i,j} \in R^+ : 1 \leqslant i \leqslant n; 1 \leqslant j \leqslant m\}, \{p_j : R^+ \rightarrow R, 1 \leqslant j \leqslant m\} \rangle$. 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 $p_j()$ 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* $q_{\Lambda,j} = \sum_{i:1 \leqslant i \leqslant n \wedge J_\Lambda(i)=j} \lambda_i$ in each class $1 \leqslant j \leqslant m$ at the game configuration Λ. The set of feasible game configurations 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)

–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 in Fig. 1. We say that user i is *satisfied* at configuration Λ if $U_i(\Lambda) \neq 0$, and not satisfied otherwise. We define 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. $U_i(\Lambda_1) < U_i(\Lambda_2)$. 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 *game play* for G is a sequence of valid selfish moves in G, i.e. $(\Lambda_1, \Lambda_2), (\Lambda_2, \Lambda_3), \ldots, (\Lambda_{k-1}, \Lambda_k)$, or a *path* in the game configuration graph Ω.

This concludes the *static* description of our base class of games.

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. A game is *resource plentiful* if there is a configuration Λ such that all users are satisfied. For our classes of games, the *communal welfare function* for configuration Λ is defined as $\Sigma_i Sat_\Lambda(i)\lambda_i$. The feasible game configuration that has highest value of communal welfare function is called the *System Optimum* or SO.

Dynamic augmentations of the games G (that we consider) contain the parameters of G and respect the static definition of G given above, but in addition they also include a fixed linear ordering of players which translates to a partial ordering of all the edges (selfish moves) emanating from each vertex (game configuration) in G's configuration graph. (All selfish moves corresponding to the same player are given the same priority.) A valid game play in the dynamic setting should also respect the dynamic rules.

For reasons motivated in the Introduction and detailed in Section 5 we alter the base class of QoS games by adding or removing appropriate pricing function(s) to the individual user utilities. More specifically, $\mathscr{Q}$ denotes the class of games where $\forall j, x,\ p_j(x) = 0$ (or any fixed positive constant). The class $\mathscr{DPQ}$ of games has strictly decreasing price functions: $\forall j$, $p_j(x_1) < p_j(x_2) \Leftrightarrow x_1 > x_2$; and in addition, they are strictly differentiated between classes j, i.e., $\forall j,\ p_j(\infty) > p_{j+1}(0)$. The class $\mathscr{PQ}$ of games satisfies both: $\forall j,\ p_j(x_1) < p_j(x_2) \Leftrightarrow x_1 > x_2$, and in addition, the price functions are the same for all classes, i.e., $\forall x,\ p_1(x) = p_2(x) = \cdots = p_m(x)$. The class $\mathscr{SPQ}$ is a modification of $\mathscr{PQ}$ that allows price function to be constant on predefined intervals around points corresponding to volume thresholds. Finally, $\mathscr{HPQ}$ is the subclass of $\mathscr{PQ}$ where selfish moves are restricted to those that do not exceed volume threshold of another player, i.e., do not cause any other player to become dissatisfied.

Here we will give a pictorial example, Fig. 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

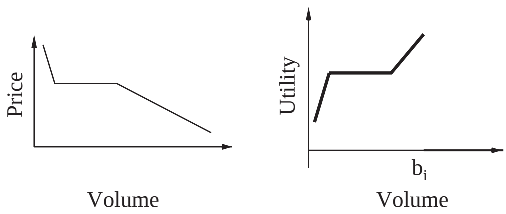

Fig. 1. Utility as a function of volume, volume threshold and price.

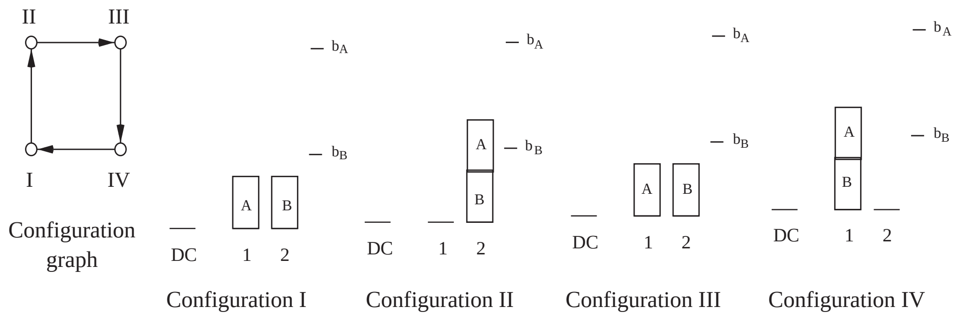

Fig. 2. Game configuration graph and individual configurations.

thresholds of users are indicated on the right. In this example, the game G in class $\mathscr{P}\mathscr{Q}$ 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. We will use this game in Observation 2.

Remark 3. Throughout this paper we assume wlog that every player i has the same volume threshold $b_i = b_{i,1} = b_{i,2} = \cdots b_{i,m}$ in every class $j = 1 \cdots m$. We also assume that players are sorted in the increasing order of their thresholds, i.e. $b_1 \leqslant b_2 \leqslant \cdots \leqslant b_n$. (The former assumption could be easily generalized for all results in this paper, the latter assumption is realistic and commonly made [22].)

Remark 4. 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. $\{J_\Lambda(i)\}$.

4. General technique for establishing stability of network games

First we give a simple, general result that however yields a clean technique for establishing stability in game configuration graphs.

Theorem 5. *The following statements are equivalent*:

(i) *There is a function defined on configuration graph Ω that increases after every selfish move* (*a so-called stability function*).

(ii) *In configuration graph Ω there is no oriented cycle C of selfish moves $\Lambda_1, \Lambda_2, \ldots, \Lambda_k$* (*i.e. such that there is an oriented edge from Λ_1 to Λ_2, from Λ_2 to $\Lambda_3, \ldots,$ from Λ_k* to Λ_1).

(iii) *Every maximal oriented* (*simple*) *path starting from any initial vertex of Ω terminates at a vertex corresponding to a Nash configuration.*

Proof. (iii) $\Rightarrow$ (i) Let $f(\Lambda)$ be equal to a $2^n - d$, where d is the maximum *oriented distance* (number of edges in the longest oriented path) from Λ to a Nash configuration. Because of (iii), f is well-defined. Let $e = (\Lambda_1, \Lambda_2)$ be an oriented edge of Ω, then $f(\Lambda_2) - f(\Lambda_1) \geqslant 1$ since for every oriented path P from Λ_2 to a Nash configuration Λ there is a longer oriented path $((\Lambda_1, \Lambda_2), P)$ from Λ_1 to Λ. Thus f is a stability function.

(i) $\Rightarrow$ (ii) Suppose that there is an oriented cycle C. Then f will continually increase over C, which contradicts the fact that C is a cycle and f is a function.

(ii) $\Rightarrow$ (iii) Let P be a maximal, simple, oriented path. Due to finiteness, maximality, and the fact that there are no cycles, the P must terminate at a vertex with no outgoing edges, i.e., at a Nash configuration. $\square$

Remark 6. Formally, a cycle mentioned in the Theorem 5 can be defined as a sequence of selfish moves that begins and ends at the same configuration Λ. This cycle explicitly specifies which player makes the first move, which makes the second move etc. There are two different types of cycles. One is where all possible sequences of selfish moves originating at any cycle configuration Λ will revisit Λ eventually. Such cycles are called terminal cycles. Another type of cycles is where there is some configuration Λ and some sequence of selfish moves that would never visit Λ again. Such cycles are called nonterminal. Note that according to the Theorem 5 a game cannot lack both Nash equilibria and selfish cycles. In Section 5.2, we will give examples of terminal and nonterminal cycles, as well as of all 3 other possible Nash/cycle combinations: (1) games that have Nash equilibria and do not have any selfish cycles, (2) games that have both Nash equilibria and selfish cycles and (3) games that have no Nash equilibria and have selfish cycles.

5. Stability vs optimality

5.1. Class $\mathcal{Q}$ of games with no pricing

First we consider the class of games $\mathcal{Q}$ where there is no pricing, i.e. $p_j(x) = 0$, for all classes j and their volumes x, and users are only motivated by their desire to satisfy their volume thresholds.

A *selfish move* by user i in a game $G \in \mathcal{Q}$ is a reallocation of i's volume from a departure Class j_1 to destination Class $j_2 \neq 0$, provided that the volume threshold of i was exceeded in Class j_1 prior to the move and it is not exceeded in Class j_2 after the move.

A corollary of the following result is that all games in $\mathcal{Q}$ always have a Nash equilibrium.

Theorem 7. *For a game in $\mathcal{Q}$, any maximal sequence of selfish moves starting at an arbitrary initial feasible configuration will terminate at a Nash configuration.*

Proof. We will give two independent proofs of this theorem, one by constructing the stability function of item (i) of Theorem 5, second by proving nonexistence of cycle of item (ii) of Theorem 5.

By construction of stability function: Suppose that players $1, \ldots, n$ have thresholds $b_1 \leqslant b_2 \leqslant \cdots \leqslant b_n$. Recall that $Sat_\Lambda(i) = 1$, if in Configuration Λ, Player i is satisfied; otherwise $Sat_\Lambda(i) = 0$. Define $f(\Lambda) = \sum_i 2^i Sat_\Lambda(i)$. Note that after every selfish move, the function $f()$ is increasing (since $Sat_\Lambda(i)$ of a moving Player i changes from 0 to 1 and if $Sat_\Lambda(i')$ changes for any other i', then $i' < i$, and $\sum_{i'<i} 2^{i'} < 2^i$). Therefore $f()$ satisfies the criteria for being a stability function.

Proof of nonexistence of a cycle: Assume that such a cycle C exists. Let M be the maximum value of b_i over all *active* players i defined as those players whose moves correspond to edges in C. Let k be the corresponding player, i.e. $b_k = M$. Suppose that a move by Player k changes Configuration Λ_1 into Configuration Λ_2 in C. By definition of selfish move, Player k must be unsatisfied at Λ_1 and satisfied at Λ_2. Therefore, there must be a selfish move by some Player t that makes k unsatisfied. However by definition of k, $b_t \leqslant b_k$, hence t cannot make a move that will dissatisfy k while satisfying t, contradiction. Thus such a cycle C cannot exist. □

Next we consider the complexity of finding a System Optimum configuration Λ of a game G. This problem is in general NP-Complete. It can, however, be solved greedily for two special subclasses of games defined by conditions we refer to collectively as C.

Theorem 8. (1) *The problem of finding a System Optimum of a network game is NP-complete.*

(2) *A System Optimum of a network game when all players have the same volume λ can be found in time linear in the game parameters.*

(3) *A System Optimum of a network game when all players have superincreasing volumes λ_i (i.e., if $b_1 \leqslant \cdots \leqslant b_n$ then $\lambda_i > \sum_{j<i} \lambda_j, \forall i$ and also $b_i \geqslant 2\lambda_i, \forall i$) can be found in time linear in the game parameters.*

Proof. 1. By reduction from MAXIMUM SUBSET SUM problem (i.e. given set $S = \{s_1, \ldots, s_n\}$ and target t, find $A \subseteq S$ such that $\sum_{i \in A} s_i \leqslant t$ and this sum is maximum). Reduction to System Optimum of network game can be done as follows. Suppose that players $1, \ldots, n$ all have same threshold $b_1 = b_2 = \cdots b_n = t$, individual volumes $\lambda_i = s_i$ and there is one non-DC class. Then System Optimum configuration corresponds to subset A described above.

2. The greedy algorithm solves this problem (let $b_1 \leqslant \cdots \leqslant b_n$; place Player n in Class 1, place Player $n-1$ in Class 1 if $b_{n-1} \geqslant 2\lambda$, otherwise place Player $n-1$ in Class 2; place Player $n-2$ in Class 1 if $b_{n-2} \geqslant 3\lambda$ etc.).

3. The greedy algorithm solves this problem, similarly to 2. □

Remark 9. When we consider the subset of a class of games $\mathscr{X}$ that satisfies conditions of item 2 (resp. 3) of Theorem 8 above we will denote this subclass of games by $\mathscr{X}_{\mathscr{E}}$ (resp. $\mathscr{X}_{\mathscr{S}}$).

We know that for games in $\mathscr{Q}$, Nash equilibria always exist. The next result states how far these Nash equilibria could be from System Optimum of their games.

Observation 1. *For any λ, there is a game $G \in \mathscr{Q}$ with a Nash equilibrium Λ whose communal welfare is $O(\frac{OPT}{\lambda})$ where OPT is the communal welfare of G's System Optimum, and λ is one of G's player volumes.*

Proof. Consider a configuration that has two classes plus the dummy class DC, users $A_1(1, 2\lambda), A_2(1, 2\lambda), B(\lambda, \lambda), \lambda \geqslant 2$. Then there is a Nash equilibria Λ when A_1 is in Class 1, A_2 is in Class 2, B is in DC. The communal welfare of Λ is 2. On the other hand, the System Optimum has A_1 and A_2 in Class 1, B in Class 2 and the communal welfare of System Optimum is $\lambda + 2$. $\square$

Theorem 10. (1) *Any Nash equilibrium of any game* $G \in \mathcal{Q}_{\mathcal{E}}$, *has communal welfare of at least a half of that of* G's *System Optimum.*

(2) *Any Nash equilibrium of any game* $G \in \mathcal{Q}_{\mathcal{G}}$, *has communal welfare of at least* $(1 - 1/2^m)OPT$ *where* m *is the number of classes and* OPT *is communal welfare of* G's *System Optimum.*

Proof. Case of games in class Q_E: Let Λ be a Nash equilibrium when all players have the same volume λ. 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 $b_i - \lambda_i$ hence communal welfare of Λ is greater than or equal to $m(b_i - \lambda_i)$ but communal welfare of System Optimum cannot be more than $2(m(b_i - \lambda_i))$.

Case of games in class $\mathcal{Q}_{\mathcal{G}}$: At Nash equilibria say the players $n, \cdots n-m$ are satisfied. Due to conditions C, $\lambda_n + \cdots + \lambda_{n-m} \geqslant (\lambda_n + \cdots \lambda_1)(1 - 1/2^m)$ and communal welfare of System Optimum is at most $\lambda_n + \cdots \lambda_1$. $\square$

Classes of games with pricing functions

As we have shown so far, games without pricing may result in Nash equilibria that are arbitrarily far from System Optima. Now we will examine effects that (nondegenerate) pricing has on existence and optimality of Nash equilibria.

Recall that a pricing function is a pricing per unit volume (or pricing for short) function $p_j() : R^+ \to R$. This is a nonincreasing function, i.e. $p_j(x) \geqslant p_j(y) \Leftrightarrow x \leqslant y$. As defined for games with no pricing, a selfish move by user i (in games with pricing) is a reallocation of i's volume from a departure class j_1 to destination class $j_2 \neq 0$, that increases utility of this user. The difference is that now utility of user i depends both on satisfaction of volume threshold of i and prices in j_1 and j_2.

Our motivation for introducing pricing is to increase communal welfare of the resulting Nash equilibrium. Consider for example Fig. 3. It shows an original Nash equilibrium in a game G without pricing and a new Nash equilibrium of a game G' that has a pricing function (H denotes highly demanding users, M moderately demanding and L low demanding). In this case, the new NE clearly has greater value of community welfare function (i.e. volume of all rectangles in non-DC classes) than the original Nash equilibrium. In the remainder of this section we will examine various pricing schema and the effects they have on stability and optimality of the corresponding network games.

5.2. *Class* $\mathcal{PQ}$ *of games with strictly decreasing pricing function*

Here we examine the class of games $\mathcal{PQ}$ when there is only one pricing function $p(x)$ for all classes j and this pricing function is strictly decreasing, i.e. $p(x) < p(y) \Leftrightarrow x > y$.

We first show that even under the conditions C, the existence of NE is not guaranteed.

Observation 2. *There are games in* $\mathcal{PQ}_{\mathcal{E}}$ *and in* $\mathcal{PQ}_{\mathcal{G}}$ *that do not have Nash equilibria.*

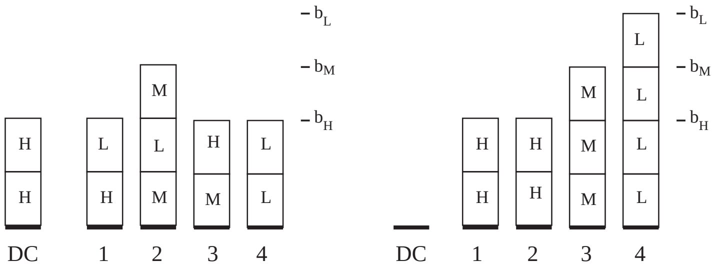

Fig. 3. Nash equilibrium without pricing and one without.

Proof. Case of games in class $\mathcal{PQ}_{\mathcal{E}}$: Consider the game depicted in Fig. 2. This game is in the class $\mathcal{PQ}_{\mathcal{E}}$ since all players have equal volumes. Consider Configuration 1. Current utility of both players is positive. However, Player A can improve his utility by moving into Class 2, since price $p(2\lambda)$, when volumes $(=\lambda)$ of A and B are combined, is lower than the price $p(\lambda)$ when A is alone, due to the strictly decreasing nature of the pricing function $p()$. This move by Player A to Class 2 results in a Configuration 2. Utility of Player B is now equal to 0, since his volume threshold is exceeded. Player B can improve his utility by moving into Class 1, Player A will follow him, and so on. This sequence of selfish moves will never terminate. Since every configuration in this game has a selfish move leading from it, this game has no Nash equilibrium.

Case of games in class $\mathcal{PQ}_{\mathcal{S}}$: Consider the game consisting of two classes, and two users A(100,300) and B(10,30). User A will always want to move to the class where B is, and B will always want to move away from A, thus creating a cycle. □

The next result states however that when Nash equilibria do exist under conditions C they optimize communal welfare.

Theorem 11. *If a game G in $\mathcal{PQ}_{\mathcal{E}}$ or in $\mathcal{PQ}_{\mathcal{S}}$ has a Nash equilibrium Λ, then Λ is a System Optimum of G.*

Proof. Case of games in class $\mathcal{PQ}_{\mathcal{E}}$: Consider a System Optimum Λ_1. If Player i is in DC then all the players k such that $b_k < b_i$ are also in DC (otherwise i could have moved into the class where k is increasing communal welfare, contradiction). Consider a Nash equilibria Λ_2. If Player i is in DC then all the players k such that $b_k < b_i$ are also in DC (otherwise i would move into the class where k is). Hence any Nash equilibrium is a System Optimum.

Case of games in class $\mathcal{PQ}_{\mathcal{S}}$: We can transform any System Optimum into any Nash equilibrium by means of either exchanging players between DC and non-DC classes or between non-DC classes. Since all players have equal volumes the resulting Nash equilibria will have the same communal welfare as the original System Optimum. □

Class of games $\mathcal{SPQ}$

So far we have shown that introduction of strictly decreasing price function tends to cause instability by creating cycles and destroying Nash Equilibria. Intuitively cycles are created by higher threshold players "chasing" lower threshold players, as in Fig. 4.

It was shown in Section 5.1 that if the price is a constant function then Nash equilibrium always exists. Fig. 4 seems to indicate that if price function were constant in a small neighborhood around volume thresholds and strictly decreasing elsewhere, as shown in Fig. 5, then Nash equilibrium would always exist. Unfortunately, this is not the case, as results below indicate.

First we will formally define the pricing function shown in Fig. 5. Let players $1, \dots, n$ have volume thresholds $b_1 \leqslant b_2 \leqslant \cdots \leqslant b_n$ and volumes $\lambda_1, \dots, \lambda_n$. Define a *stopping price function* $p(x)$ as a function that is flat on intervals $(b_i - \lambda, b_i + \lambda), \forall i$, where $\lambda = \max_j \lambda_j$, and $p(x)$ is strictly decreasing between these intervals. $\mathscr{SPQ}$ denotes the class of games that have stopping price functions.

Observation 3. *There is a game in* $\mathscr{SPQ}$ *where there is a cycle of selfish moves.*

Proof. 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 1, 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. Current configuration is essentially isomorphic to Λ, hence a cycle has occurred. □

Remark 12. We mentioned in Section 4 that there are 3 different possibilities for Nash/cycle existence. By now we have seen examples of all such possibilities. Any game in class $\mathscr{Q}$ has Nash

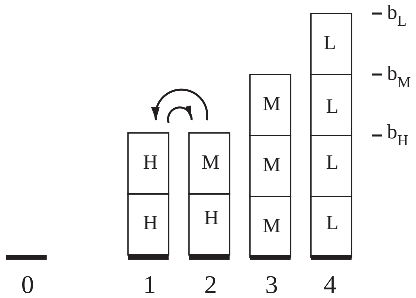

Fig. 4. Game without Nash equilibrium.

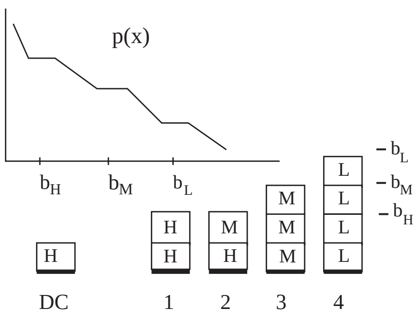

Fig. 5. Stopping price function.

equilibria and has no cycles (terminal or nonterminal). In Observation 2 we have seen an example of a game with a (terminal) cycle and no Nash equilibria. In Observation 3 we gave an example of a game where there is both a (nonterminal) cycle of selfish moves and a Nash equilibrium (all A and B players in one class, three C players in another class, remaining C players in DC).

5.3. Class $\mathscr{HPQ}$ of games with nonhurting moves

So far we have shown that the introduction of pricing tends to cause instability by creating cycles and destroying Nash equilibria. Now we will impose natural restrictions on types of selfish player moves allowed. We will show that the class of games $\mathscr{HPQ}$ with such restrictions will be free of instabilities. We will also examine optimality of games in $\mathscr{HPQ}$.

A (selfish) *nonhurting move* by Player i is a reallocation of i's volume from a departure Class j_1 to destination Class j_2, changing the Configuration Λ_1 to Configuration Λ_2 such that $U_i(\Lambda_1) < U_i(\Lambda_2)$ and there is no player k such that $Sat_{\Lambda_1}(k) \neq 0$ and $Sat_{\Lambda_2}(k) = 0$. In other words, player i improves his utility without violating volume thresholds of any other players.

A corollary of the following result is that all games in $\mathscr{HPQ}$ always have a Nash equilibrium.

Theorem 13. *For any game $G \in \mathscr{HPQ}$ any maximal sequence of selfish nonhurting moves starting at an arbitrary feasible configuration Λ will terminate at a Nash equilibrium.*

Proof. We will give two proofs of this theorem, one by constructing a stability function of item (i) of Theorem 5 and secondly proving nonexistence of cycle of item (ii) of Theorem 5.

Construction of stability function: Let $f(\Lambda) = \sum_{j=1}^{m} q_j^2$ (DC contributes zero). Note that whenever a selfish move by some Player i changes Configuration Λ_1 into a Configuration Λ_2 then $f(\Lambda_2) > f(\Lambda_1)$. This is because when Player i moves from class j_1 to class j_2 the following holds:

$$(q_{j_1} - \lambda_i)^2 + (q_{j_2} + \lambda_i)^2 > q_{j_1}^2 + q_{j_2}^2,$$

provided $q_{j_1} < q_{j_2} + \lambda_i$; and when Player i moves from DC to Class j_2, it holds that $(q_{j_2} + \lambda_i)^2 > q_{j_2}^2$. Finally, no player ever moves to DC.

Nonexistence of cycle: Assume that such a cycle C exists. Let M be the minimum value of $p(q_j)$, where j is taken over all destination and departure classes j_2 and j_1 of *active* players i defined as those players whose moves correspond to edges in C. Since M is the minimum, no active player will move away from the Class O with price M. The price M of the Class O also will not be reduced by any active players moving into O since the function $p()$ is decreasing. Since active players that are stuck at O will not be able to move out, there cannot be such a cycle C, contradiction. Note that none of the players in DC can participate in such a cycle C either since no selfish move ever causes a player to return to DC. □

How far can a Nash equilibrium of a $\mathscr{HPQ}$ game be away from a System Optimum of this game? The following theorem states that it can be arbitrarily far, even when restricted by conditions C.

Observation 4. *For any n, M, there is a game $G \in \mathscr{HPQ}_{\mathscr{E}}$ (resp. in $\mathscr{HPQ}_{\mathscr{S}}$) that has a NE Λ such that communal welfare of Λ is $O(\frac{OPT}{n})$ (resp. $O(\frac{OPT}{M})$) where OPT is communal welfare of SO of G, n is number of players of G, and M is the ratio of two of G's player volumes.*

Proof. Case of games in class $\mathscr{HPQ}_{\mathscr{E}}$: Consider a game that has 1 class plus DC and users $A_1(\lambda, n\lambda), A_2(\lambda, n\lambda), \ldots, A_n(\lambda, n\lambda), B(\lambda, \lambda)$. Then there is a Nash equilibrium Λ in which B is in Class 1 and all other users are in DC. The communal welfare of Λ is λ. On the other hand, a System Optimum has $A_1, \ldots, A_n$ in Class 1, B in DC, communal welfare of System Optimum is $n\lambda$.

Case of games in class $\mathscr{HPQ}_{\mathscr{S}}$: Consider a game that has 1 class plus DC, with users $A(M, M), B(1, 2), M \geqslant 1$. Then there is a Nash equilibrium Λ when B is in Class 1 and A is in DC, and communal welfare of Λ is 1. On the other hand, a System Optimum has A in Class 1, B in DC, and communal welfare of System Optimum is M. □

5.4. Class of games $\mathscr{DPQ}$ with (different) separating price functions

So far we have considered classes of games when there was one price function in effect for all classes. Examples of the price functions we have seen would either not induce Nash equilibria or induce suboptimal Nash equilibria. However if we were allowed to introduce special different price functions for different classes then we can show that games in this class $\mathscr{DPQ}$ always terminate at a Nash equilibrium and under conditions C, these Nash equilibria are also System Optimal.

Definition 14. A set of strictly decreasing functions $p_1(), \ldots, p_m()$ are separating price functions if $p_m(0) > p_m(\infty) > p_{m-1}(0) > p_{m-1}(\infty) > p_{m-2}(0) > p_{m-2}(\infty) > \cdots > p_1(0) > p_1(\infty)$. The class of games with such pricing functions is denoted by $\mathscr{DPQ}$. See Fig. 6.

A corollary of the following result is that all games in $\mathscr{DPQ}$ always have a Nash equilibrium.

Theorem 15. *For any game $G \in \mathscr{DPQ}$ any maximal sequence of selfish moves starting at an arbitrary initial feasible configuration will terminate at a Nash equilibrium.*

Proof. We will give two proofs of this theorem, one by constructing stability function of item (i) of Theorem 5, second by proving nonexistence of cycle of item (ii) of Theorem 5.

Construction of stability function: Let $f(\Lambda) = \sum_i (m - J_\Lambda(i))2^{2^i}$ where the summation is taken over all satisfied players (i.e. those not in DC and whose volume thresholds are not exceeded). Due to the structure of pricing functions, all selfish moves are either by currently satisfied players A to a lower indexed class j_1, or by currently unsatisfied players B to a different class j_2. In the former case, a gain in $f()$ caused by decrease in $J_\Lambda(A)$ is greater than the loss in $f()$ caused by all players in j_1 who become unsatisfied (their indexes are less than that of A). Therefore the function $f()$ increases. Similarly, in the later case, a gain in $f()$ caused by adding a summation term for B is greater than the loss in $f()$ caused by all players in j_1 who become unsatisfied. Thus the function $f()$ increases after every selfish step.

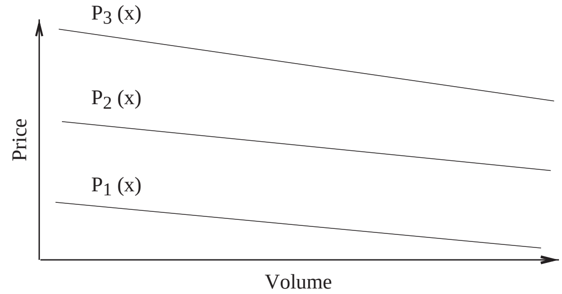

Fig. 6. Separating price functions.

Nonexistence of a cycle: Suppose that there is a cycle of selfish moves C. Let i be the highest threshold player that participates in this cycle. Let j be the smallest numbered class that i moves into during C. Then since price at Class j is less than price at any other class (regardless of total volume values) Player i will never leave class j, thus cycle C cannot exist, contradiction. □

The next result states that in general, Nash equilibria of games in class $\mathscr{DPQ}$ can be arbitrarily far from corresponding System Optimum.

Observation 5. *For any λ, there is a game $G \in \mathscr{DPQ}$ with a Nash equilibrium Λ whose communal welfare is $O(\frac{OPT}{\lambda})$, where OPT is communal welfare of G's System Optimum, and λ is one of G's player volumes.*

Proof. Similar to the proof of Observation 1. □

Under conditions C, however, Nash equilibria have the largest possible value of communal welfare.

Theorem 16. *For any game in $\mathscr{DPQ}_{\mathscr{E}}$ or $\mathscr{DPQ}_{\mathscr{S}}$ every Nash equilibrium is a System Optimum.*

Proof. Similar to the proof of Theorem 11. □

6. Dynamics

Here we briefly examine speed of convergence to the Nash configurations for various game classes. First we introduce simple rules that impose a priority order in which users move.

A dynamic game rule that orders user moves at any configuration proportionally to user's thresholds (i.e. if $b_1 \leqslant b_2 \leqslant \cdots \leqslant b_n$ then user n has a right to move before everybody else does, then $n-1, n-2$ etc.) is called *increasing-threshold-order rule*.

Theorem 17. *For any game in $\mathscr{Q}$, $\mathscr{DPQ}$* (resp in $\mathscr{DPQ}_{\mathscr{S}}$, $\mathscr{PQ}_{\mathscr{S}}$, $\mathscr{Q}_{\mathscr{S}}$) *and for any initial configuration Λ, every maximal increasing-threshold-order sequence of selfish moves will terminate at a Nash* *equilibrium after* $O(n^2)$ *steps* (*resp after* $O(n)$ *steps*), *where n is the number of players*. (*Note this results holds for games in* $\mathcal{PQ}_{\mathcal{S}}$ *provided these games actually have Nash equilibria*.)

Proof. Case of games in class $\mathcal{Q}$: Note that once Player n has moved, it will not move again. Suppose by induction that there were $O((n-1)^2)$ moves before and after Player n has moved, hence total time is $O(n^2)$.

Cases of games in classes $\mathcal{Q}_{\mathcal{S}}, \mathcal{PQ}_{\mathcal{S}}$ and $\mathcal{DPQ}_{\mathcal{S}}$: Note that Player n can always move first (to Class 1 in case of $\mathcal{DPQ}_{\mathcal{S}}$ and $\mathcal{Q}_{\mathcal{S}}$, to class where Player $n-1$ resides in case of $\mathcal{PQ}_{\mathcal{S}}$). Player n will not move after that (unless in case of $\mathcal{PQ}_{\mathcal{S}}$ volume threshold of Player $n-1$ was exceeded, so Player $n-1$ would move to another class and Player n will follow him, creating a cycle. But we are only considering cases where Nash equilibrium exists). Therefore, every player will move at most a constant number of times, hence the total time is $O(n)$.

Case of games in class $\mathcal{DPQ}$: Note that Player n can move at most n times. After Player n has stopped moving, Player $n-1$ can move at most n times etc. Therefore the total time is $O(n^2)$. □

7. Directions, conjectures, initial results

7.1. Advantages of pricing

We now argue that games with pricing in general have greater communal welfare at Nash equilibria than similar games without pricing. Unfortunately, counterexamples such as Observations 2 and 4 indicate the existence of unstable games and games with arbitrarily suboptimal communal welfare at Nash equilibria for some game classes. Even Theorem 16 for (approximate) optimal communal welfare at Nash equilibria of certain game classes relies on the conditions C; and counterexamples such as Observation 5 indicate that these conditions are necessary. In addition, we have the following Observation 6 that apparently questions the efficacy of using of pricing to increase communal welfare at Nash equilibria.

In this section, we first describe Observation 6 and then conjecture that when averaged over all games in certain classes, pricing tends to improve communal welfare at Nash equilibria. In order to compare games with and without pricing, we need to introduce appropriate definitions.

Let G_1 be a game in class $\mathcal{Q}$ of games without pricing. Let Λ_1 be a Nash equilibrium of G_1. Let G_2 be a game in $\mathcal{PQ}$ that has the same game parameters as G_1 plus a pricing function $p()$. Let Λ_2 be a configuration in G_2 that corresponds to Λ_1 in G_1 (i.e. configurations Λ_1 and Λ_2 have identical assignment of users to classes). Note that Λ_2 may or may not be a Nash equilibrium in G_2. A Nash equilibrium Λ_3 in game G_2 is said to be *induced* by Nash equilibrium Λ_1 and pricing function $p()$ if there is a game play in G_2 that leads from Λ_2 to Λ_3. Similarly game G_2 is said to be *induced* by G_1 and $p()$.

Observation 6. *For any strictly decreasing price function* $p()$, *there is a Nash equilibrium* Λ_1 *of game* $G_1 \in \mathcal{Q}$ *such that a Nash equilibrium* Λ_3 (*of a game* $G_2 \in \mathcal{PQ}$) *induced by* Λ_1 *and* $p()$ *has strictly smaller communal welfare than* Λ_1.

Proof. Consider a configuration that has two classes plus DC, and users $A_1(4,9), A_2(4,9), B(1,13), C(10,11), D(6,14)$. A NE Λ_1 without pricing has Users A_1, A_2, and B in Class 1, User C in Class 2, User D in DC, and communal welfare equal to 19. A NE induced by Λ_1 and any strictly decreasing price function $p()$ has User D in Class 1, Users B and C in Class 2, Users A_1 and A_2 in DC, and communal welfare $=$ 17. See Fig. 7. □

One way to offset Observation 6 is by establishing an approximate upper bound on the possible deterioration of communal welfare caused by the introduction of pricing. We have the following weak conjecture and believe that significantly stronger statements should be provable.

Conjecture 18. *Let Λ_1 be a Nash equilibrium of a game without pricing. Let $p()$ be any strictly decreasing cost function. Let Λ_3 be a Nash equilibrium induced by Λ_2 and $p()$ (assuming that Λ_3 exists, which is not always guaranteed in $\mathscr{P}\mathscr{Q}$). Let $\lambda = \max_i \lambda_i$. Then communal welfare of Λ_1 minus communal welfare of $\Lambda_3 \leqslant \sum_j x_j$ where $x_j = q_j - \lfloor \frac{q_j}{\lambda} \rfloor \lambda$, and q_j denotes the total volume in Class j in configuration Λ_1.*

Another way of offsetting Observation 6 is to use probabilistic analysis to compare communal welfare of all Nash equilibria of the original game to communal welfare of all Nash equilibria of the induced game. Furthermore such analysis should be applied to entire classes of games (for example $\mathscr{Q}$ vs $\mathscr{P}\mathscr{Q}$) instead of specific individual games.

Here we will introduce some straightforward probability notions dealing with Markov chains, that could allow us to compare Nash equilibria in entire classes of original and induced games.

The *transition probability* of a game configuration graph $\Omega = (V, E)$ is an assignment of weights $W : E \to R$ that has the following properties. Weight of an edge $w(\Lambda_1, \Lambda_2)$ is equal to the probability of a move from configuration Λ_1 to Λ_2. Thus weights of all the edges leaving any node should add up to one. Nodes corresponding to Nash configurations have one outgoing looping edge of weight one. The transition probability is *uniform* if weight of an edge $w(\Lambda_1, \Lambda_2)$ is a reciprocal of the number of edges leaving Λ_1. The *probability distribution P* over the configuration graph $\Omega = (V, E)$ is a probability assignment $P : V \to R^+$ such that $\sum_{v \in V} P(v) = 1$. The *uniform probability distribution* assigns $1/|V|$ to every vertex v. A sequence of probability distributions $P_0, P_1, \ldots, P_i, \ldots$ is *induced* by an *initial probability distribution* P_0 and transition probability W if

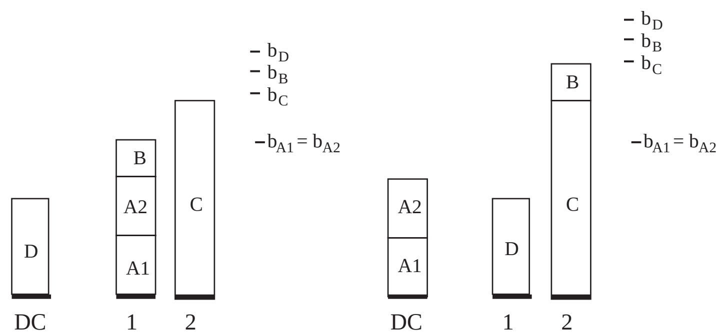

Fig. 7. Nash equilibria without and with pricing.

for all configurations Λ

$$P_i(\Lambda) = \sum_{\pi} P_0(\Lambda_{\pi(1)}) w(\Lambda_{\pi(1)}, \Lambda_{\pi(2)}) \dots w(\Lambda_{\pi(i-1)}, \Lambda_{\pi(i)}),$$

where π is a game play to Λ of length i in Ω, i.e. $\Lambda_{\pi(j)}$ represents the jth vertex in this path and $\Lambda_{\pi(i)} = \Lambda$. Wlog, we assume that initial probability distribution and transition probabilities are uniform. The limit of induced probability distributions $\lim_{i \to \infty} P_i(\Lambda)$ is called the *stationary distribution* and is denoted $P_\infty(\Lambda)$. If this limit does not exist then $P_\infty(\Lambda) = 0$. The *stationary communal welfare* $E_G(W)$ of a game G is defined as $\sum_{\Lambda \in G} P_\infty(\Lambda) W(\Lambda)$ where $W(\Lambda)$ is the communal welfare of a configuration Λ. The *expected value of communal welfare* $E_{\mathscr{A}}(W)$ for a class of games $\mathscr{A}$ is a

$$E_{\mathscr{A}}(W) = \sum_{G \in A} Prob(G) * E_G(W),$$

where Prob(G) is the probability attached to a particular game G in $\mathscr{A}$ (we generally assume that G is picked uniformly from the space of game parameters m, n, b_i, λ_i).

The definitions above would allow us to compare various classes of network games.

Conjecture 19. $E(\mathscr{DPQ}) > E(\mathscr{SPQ}) > E(\mathscr{HPQ}) > E(\mathscr{Q})$.

Move-correlated welfare functions

A different way of comparing various classes of games to class $\mathscr{Q}$ is by using *move-correlated welfare functions* defined on game configurations. Intuitively a function $g()$ is an increasing (resp. decreasing) move-correlated welfare function (for games G in class $\mathscr{A}$) if $g()$ increases (resp. decreases) on average after a selfish move in G and $g()$ is positively correlated with communal welfare. The existence of such a function for a class of games would indicate that the Nash equilibria of such games tend to have high communal welfare values.

One possible candidate for such a decreasing move-correlated welfare function for games in class $\mathscr{PQ}$ is the *volume function* defined as

$$v(\Lambda) = \sum_{j} \sum_{i: J_\Lambda(i) = j} (b_i - q_j).$$

Intuitively this function assigns small values to those configurations where most of the players in any given class j have thresholds close to the total volume in j. Hence at Nash equilibria the volume function is (inversely) correlated with communal welfare.

The volume function decreases after every selfish move for games in $\mathscr{HPQ}_{\mathscr{E}}$, since users always move from classes with smaller total volumes to the larger ones. Unfortunately for games in $\mathscr{PQ}$, this is no longer the case, since users might move to the smaller class if their volume threshold was exceeded in a larger class (which was explicitly disallowed in $\mathscr{HPQ}_{\mathscr{E}}$). However, we conjecture that it is possible to amortize such moves from larger to smaller classes by the moves that caused those thresholds to be exceeded in the first place. Since there are selfish moves that do not exceed any player's volume threshold and since the function $v()$ would decrease for such moves, the overall expected change in $v()$ would be negative, motivating the following claim.

Conjecture 20. *The volume function is a decreasing move-correlated welfare function for games in class $\mathscr{P}\mathscr{Q}$.*

Remark 21. The technique described above can be extended for studying interesting functions other than communal welfare on Nash equilibria. For example we conjecture that for games in $\mathscr{P}\mathscr{Q}$, the volume function is highly correlated with the function that measures the number of occupied classes, i.e. the number of classes used by at least one user.

7.2. Provider participation and price thresholds

So far, we considered network users as the only players. These users move according to their individual preferences and fixed price functions set by a benevolent network manager. Now we would like to extend this model, so it would include network providers as players as well. The role of network provider is to determine the price functions that will be used by network users. During the network game a move by a provider replaces current price function by a new price function. There are two types of providers: *selfish providers* that try to choose price functions that maximize the total amount paid by all network users and *benevolent providers* that try to choose price functions that will result in a Nash equilibria that have high value of communal welfare.

In this section, we only study benevolent providers since simple counterexamples [11] show that games with selfish providers generally are either highly unstable or have highly suboptimal Nash equilibria. However if the provider is benevolent then we show that in fact convergence to Nash with optimal communal welfare can be ensured for our classes of games.

Now we introduce appropriate definitions. A *move by a provider* replaces the current price function by a new price function (wlog in this section we only consider the case when there is only one price function in effect for all classes). The class of games with benevolent provider is denoted by $\mathscr{PRBE}$.

Theorem 22. *Any game in $\mathscr{PRBE}_{\mathscr{E}}$ or $\mathscr{PRBE}_{\mathscr{S}}$ (i.e. when provider is benevolent and users have equal or superincreasing volumes) terminates at a Nash equilibrium which is also System Optimum.*

Proof. Consider the following sequence of moves: (wlog we assume that moves are done under increasing order threshold rule defined in Section 6) first the provider introduces a strictly decreasing price function. Than all users move until there is a Class X that contains users $n, \dots, k$ such that $\sum_{i=k}^{n} \lambda_i \leqslant b_k, \sum_{i=k-1}^{n} \lambda_i > b_{k-1}$. In other words, all users in class X are satisfied and of the remaining users, the ones with the highest volume threshold cannot move into X. At this point, provider adjusts the price function so as to prevent users $n, \dots, k$ from ever leaving X again. This is done by making $p()$ constant $(= p(\sum_{i=k}^{n} \lambda_i))$ on $[\sum_{i=k}^{n} \lambda_i, \infty]$. The process is repeated for the remaining users $1, \dots, k-1$, until Nash equilibria is reached. Proof that this Nash equilibrium is System Optimum is similar to the proof of Theorem 11.

Remark 23. In practice there might 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. We also conjecture that this concept has effect on the dynamics of the game, as is mentioned 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 $b_i \leqslant b_j$ iff $t_i \geqslant t_j$, 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 due to the game plays in games with price thresholds spending less time looping in non-terminal cycles. We have performed a set of computer experiments that support this conjecture, see [11].

7.3. New ways of proving existence of Nash equilibria

Our stability function technique of Theorems 7, 13 etc. is only useful in establishing existence of Nash equilibria in the situations where no cycle is present in the game configuration graph. As was shown in Section 5.2, there are game classes that are generally not cycle-free. Hence, we would like to extend the stability function technique to be able to show Nash existence in games that may have Nash equilibria coexisting with cycles. One way to do this is by using the concept of local stability *local stability functions*.

Intuitively a function $g()$ is a local stability function for a subgraph A of a game configuration graph Ω, if it increases after a selfish move within A, i.e. $g(\Lambda_1) < g(\Lambda_2)$ whenever $\Lambda_1, \Lambda_2 \in A$ and there is an oriented edge in A from Λ_1 to Λ_2. The existence of a local stability function for a subgraph A (and a graph Ω) would imply existence of a Nash equilibrium in this subgraph A, provided that A is closed under selfish moves, i.e. any edge from any vertex in A points to another vertex in A (and not in $\Omega \backslash A$).

One example of such a local stability function for games in class $\mathscr{P}\mathscr{Q}$ is the *satisfied volume function* defined as

$$sv(\Lambda) = \sum_j (q_j^s)^2 + q_0,$$

where q_j^s is the total volume of all satisfied players in Class j, i.e.

$$q_j^s = \sum_{i | J_\Lambda(i) = j \& Sat_\Lambda(i) \neq 0} \lambda_i.$$

Recall that q_0 denotes total volume in DC. Consider a subgraph A that consists of a Nash equilibrium Λ_1 and all configurations $\Lambda_2, \ldots, \Lambda_p$ that have only one outgoing edge each, and this edge points toward Λ_1. It is easy to see, that on such a subgraph A, the satisfied volume function increases after every selfish move in A (due to the reasons similar to the ones used in the proof of Theorem 13).

Note that the subgraph A above is small and more significantly defined in terms of Nash equilibrium Λ_1, and hence A is useless for proving the general existence of Nash equilibria. This presents a problem when analyzing a class of games using local stability functions, since natural local stability functions may correspond to unnatural subgraphs and vice versa.

A different approach to proving existence of Nash equilibria would be to relax the condition that forces stability functions to increase after every selfish move. If say the average value of a function were shown to be increasing for any sequence of c consecutive selfish moves, for some fixed c, then this would guarantee the existence of Nash equilibria. By constructing such functions, it would be possible to identify and completely classify subclasses of $\mathscr{P}Q$ that have both cycles of selfish moves and Nash equilibria.

References

[1] R. Cocchi, D. Estrin, S. Shenker, Pricing in computer networks: motivation, IEEE/ACM Trans. Network. 1 (6) (1993) 614–627.
[2] J. Feigenbaum, A. Krishnamurthy, R. Sami, S. Shenker, Approximation and collusion in multicast cost sharing. in: Proceedings of the Third Conference on Electronic Commerce, ACM Press, New York, 2001, pp. 253–255.
[3] J.M. Kleinberg, Y. Rabani, E. Tardos, Fairness in routing and load balancing, in: IEEE Symposium on Foundations of Computer Science, 1999, pp. 568–578.
[4] Y. Korilis, A. Lazar, Why is flow control hard: optimality etc, CTR Technical Report CU/CTR/TR 332-93-11, 1992.
[5] Y.A. Korilis, A.A. Lazar, On the existence of equilibria in noncooperative optimal flow control, J. ACM 42 (3) (1995) 584–613.
[6] Y.A. Korilis, A.A. Lazar, A. Orda, Architecting noncooperative networks, IEEE J. Selected Areas Comm. 13 (7) (1995) 1241–1251.
[7] E. Koutsoupias, C.H. Papadimitriou, Worst-case equilibria, in: Symposium on Theoretical Aspects of Computer Science, 1999, pp. 404–413.
[8] A. Kumar, J.M. Kleinberg, Fairness measures for resource allocation, in: IEEE Symposium on Foundations of Computer Science, 2000, pp. 75–85.
[9] L. Libman, A. Orda, Atomic resource sharing in noncooperative networks, in: INFOCOM, 1997, pp. 1006–1013.
[10] L. Libman, A. Orda, Optimal transport strategies for best-effort traffic over priced connections, Technical Report, Technion, 2000.
[11] A. Lomonosov, M. Sitharam, K. Park, The effect of pricing on the structure of a noncooperative networking architecture, Technical Report, University of Florida, http://www.cise.ufl.edu/˜sitharam/pub.html, 2002.
[12] S.H. Low, P. Varaiya, An algorithm for optimal service provisioning using resource pricing, in: Proceedings of INFOCOM, 1994, pp. 368–373.
[13] G. Malewicz, A. Shvartsman, An auction-based flexible pricing scheme for renegotiated qos connections and its evaluation, in: Seventh International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), 1999, pp. 244–251.
[14] P. Marbach, The role of pricing in differentiated services networks, Technical Report, CSRG-421, Department of Computer Science, University of Toronto.
[15] P. Marbach, Pricing differentiated services networks: bursty traffic, in: INFOCOM, 2001.
[16] N. Nisan, Algorithms for selfish agents, in: Symposium on Theoretical Aspects of Computer Science, 1999, pp. 1–15.
[17] N. Nisan, A. Ronen, Algorithmic mechanism design, in: Proceedings of the 31st ACM Symposium on Theory of Computing, 1999, pp. 129–140.
[18] A. Orda, R. Rom, N. Shimkin, Competitive routing in multiuser communication networks, IEEE/ ACM Trans. Network. 1 (5) (1993) 510–521.
[19] C. Papadimitriou, Computational aspects of organization theory, Proceedings of the 1996 European Symposium on Algorithms, Lecture Notes in Computer Science, Springer, Berlin, 1996.
[20] C.H. Papadimitriou, algorithms, games, and the internet, in: Proceedings of STOC, 2001.
[21] K. Park, Self-organized multi-class QoS provision for ABR traffic in ATM networks, in: Proceedings of the 15th IEEE International Phoenix Conference on Computers and Communications, 1996, pp. 446–453.
[22] K. Park, M. Sitharam, S. Chen, Quality of service provision in noncooperative networks: heterogeneous preferences, in: Proceedings of the First International Conference on Information and Computation Economics ICE'98, 1998.

[23] L. Rasmusson, Agents, self-interest and electronic markets, Knowledge Eng. Rev. 14 (2) (1999) 143–150.
[24] A. Ronen, Solving optimization problems among selfish agents, Ph.D. Thesis, Hebrew University of Jerusalem, 2000.
[25] T. Roughgarden, E. Tardos, How bad is selfish routing? in: IEEE Symposium on Foundations of Computer Science, 2000, pp. 93–102.
[26] J. Sairamesh, D. Ferguson, Y. Yemini, Optimal allocation and quality of service provisioning in high-speed packet networks, in: Proceedings of Conference on Computer Communications, Boston, Massachusetts, 1995, pp. 1111–1119.
[27] S. Shenker, Making greed work in networks: a game-theoretic analysis of switch service disciplines, in: SIGCOMM Symposium on Communications Architectures and Protocols, 1994, pp. 47–57.
[28] S. Shenker, Mechanism design and the internet, in: Presentation in DIMACS Workshop on Computational Issues in Game Theory and Mechanism Design, 2001.