



Congestion games with failures[☆]

Michal Penn^a, Maria Polukarov^{b,*}, Moshe Tennenholtz^{c,a}

^a Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel

^b School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, United Kingdom

^c Microsoft Israel R&D Center, 13 Shenkar Street, Herzeliya 46725, Israel

ARTICLE INFO

Article history:

Received 3 June 2008

Received in revised form 5 December 2010

Accepted 18 January 2011

Available online 1 March 2011

Keywords:

Congestion games

Resource failures

Pure strategy Nash equilibrium

Price of anarchy

Semi-strong Nash equilibrium

Algorithms

ABSTRACT

We introduce a new class of games, *congestion games with failures* (CGFs), which allows for resource failures in congestion games. In a CGF, players share a common set of resources (service providers), where each service provider (SP) may fail with some known probability (that may be constant or depend on the congestion on the resource). For reliability reasons, a player may choose a subset of the SPs in order to try and perform his task. The cost of a player for utilizing any SP is a function of the total number of players using this SP. A main feature of this setting is that the cost for a player for successful completion of his task is the minimum of the costs of his successful attempts. We show that although CGFs do not, in general, admit a (generalized ordinal) potential function and the finite improvement property (and thus are not isomorphic to congestion games), they always possess a pure strategy Nash equilibrium. Moreover, every best reply dynamics converges to an equilibrium in any given CGF, and the SPs' congestion experienced in different equilibria is (almost) unique. Furthermore, we provide an efficient procedure for computing a pure strategy equilibrium in CGFs and show that every best equilibrium (one minimizing the sum of the players' disutilities) is semi-strong. Finally, for the subclass of symmetric CGFs we give a constructive characterization of best and worst equilibria.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Consider a system consisting of two resources, a and b , where each of them can be used in order to process a player's task, and the player chooses one of them to perform his task. However, if the resources are not fully reliable then the player may decide to assign his task to both resources. More generally, suppose that there are n players who need to process their individual tasks using one of the two identical reliable resources. The cost for each player for utilizing a particular resource is a (nondecreasing) function of the congestion experienced by this resource. Naturally, each player aims at minimizing his cost. An optimal solution would be that "half" of the players will use each of the resources. When modeled as a game, such "equal partition" is an equilibrium of the game. However, if the resources are not reliable, the players might choose *both* resources to perform their tasks. Indeed, such behavior might be obtained in equilibrium of the corresponding game. As a result of such behavior, the resources might be overloaded, and the cost to each player will be very high. More generally, each player wants to maximize the probability of successful execution of his tasks, and simultaneously, to minimize his cost; for example, a player may prefer not to use a resource if the congestion of that resource causes a long processing delay.

[☆] A preliminary version of this paper appeared in EC-2005 Penn et al. (2005) [20].

* Corresponding author. Fax: +44 0 23 8059 2865.

E-mail addresses: mpenn@ie.technion.ac.il (M. Penn), mp3@ecs.soton.ac.uk (M. Polukarov), moshet@microsoft.com, moshet@ie.technion.ac.il (M. Tennenholtz).

Hence, as we can see, reliability issues have significant implications on player behavior in congestion settings. In order to address this challenge, we introduce a model termed *congestion games with failures* (CGFs), and establish several basic results for this model. To the best of our knowledge, we were the first to incorporate the issue of failures into congestion games.

Following the preliminary conference version of this paper [20], we considered another, alternative model of Congestion Games with Load-Dependent Failures [CGLFs] [21]. Although the terms CGF and CGLF may sound similar, these models refer to very different situations and have very different motivation. In a CGF, players strive to minimize the delay caused by using a set of alternative resources, and the payoff of a player is determined by the minimum of delays of the set of selected resources (and by incompleteness costs). In a CGLF, a player receives a reward in the case of successful completion of his task, and his objective is to maximize the difference between the expected benefit from the successful task completion and the total cost of the utilized resources.

The class of congestion games was first introduced by Rosenthal [25] who proved that these games always possess a Nash equilibrium in pure strategies. Congestion games are noncooperative games in which a collection of players have to choose from a finite set of alternatives (resources). The utility of a player from using a particular resource depends only on the number of players using it, and his total utility is the sum of the utilities obtained from the resources he uses. Congestion games have been used to model traffic behavior in road and communication networks, competition among firms for production processes, migration of animals between different habitats, and received a lot of attention in the recent computer science and electronic commerce communities [15,18,24,25,27]. Rosenthal [25] studied congestion games with a finite number of players. In addition, several authors have considered *non-atomic congestion games*, where the “non-atomic” part refers to the assumption that there is a continuum of players, each controlling a negligible fraction of the overall load on the system (see, e.g., [28,29]), or, differently, to the assumption that a player may take continuous decisions, representing the amount of congestion contributed by the player to each service provider [18].

However, the above settings do not take into consideration the possibility that resources may fail to execute their assigned tasks. Typically, the resources are machines, computers, service providers, communication lines etc. These kinds of resources are obviously prone to failures because of breakage or for any other reasons. Thus, the issue of failures should not be ignored.

The notion of failures is widely used in the field of distributed systems. In a lot of situations failing components of the system may be viewed as playing against its correctly functioning parts. The issue of failures in game-theoretic setting is extensively discussed by Linial [13]. In another line of research, [23] introduced the notion of fault tolerant mechanism design which extends the standard game-theoretic framework of mechanism design to allow for uncertain executions. In the above settings failing components are self-motivated or malicious players. In our work, we initiate an investigation of failures in congestion games, where not the players, but the resources they share may fail.

In a congestion game with failures (CGF), players share a common set of resources, or service providers,¹ where each service provider (SP) may fail with some known probability. For reliability reasons, a player may choose a subset of the service providers in order to try and perform his task. Therefore, each player's set of pure strategies coincide with the power set of the set of SPs, and the total load on the system is not known in advance, but strategy-dependent. The cost for a player for successful completion of his task is the *minimum* of the costs of his successful attempts. The cost function associated with each SP is not universal but player-specific. That is, the (dis-)utility to a player depends not only on the number of players using the same SP, but also on the identity of the player in question. Congestion games with player-specific cost functions were first studied by Milchtaich [14]. This generalization was, however, accompanied by the assumption that each player chooses only one resource.

The contribution of this paper is both conceptual and technical. On the conceptual level we introduce a model to handle strategic behavior in congestion settings with unreliable resources. In addition, we provide several basic results on the existence of various forms of equilibria in such settings.

It is known that any congestion game possesses a potential function, and as a result possesses a pure strategy equilibrium [17]. Our first technical contribution is by showing that although CGFs do not, in general, admit a potential function, and thus are not isomorphic to congestion games, they always possess a pure strategy Nash equilibrium. Moreover, though, as we show, an arbitrary improvement dynamics may cycle (that is, the finite improvement property (FIP) does not hold), the convergence to an equilibrium is guaranteed from any strategy profile if the players make best response improvement steps. We also develop an efficient algorithm for computing a pure strategy equilibrium and prove that the SPs' congestion experienced in different Nash equilibria is (almost) unique. Furthermore, we consider the existence of a semi-strong equilibrium in congestion games with failures. In a strong equilibrium [2] it is required that deviations by any *subset* of the players would not be beneficial to them. This is a much more demanding concept than the Nash equilibrium, in which only unilateral deviations are considered. In a semi-strong equilibrium we require that if there exists such beneficial group deviation, then it should itself not be stable: there will exist a single player among the deviators, who can gain by deviating from the deviating strategy.² We show that although strong equilibrium does not, in general, exist in CGFs, every Nash equilibrium strategy profile that minimizes the sum of the players' disutilities is a semi-strong Nash equilibrium.

¹ In this paper, we use the terms resources and service providers interchangeably.

² Semi-strong equilibrium is a more demanding concept than coalition-proof equilibrium [5,6]. In a coalition-proof equilibrium a deviation by a subset of the players is considered stable, if there is no stable deviation from it by any subset of it.

We devote special attention to the subclass of symmetric CGFs. We use this interesting subclass in order to differentiate between different types of equilibria according to their social disutility (the sum of the players' disutilities). We give a characterization of the best and worst Nash equilibria with respect to the social disutility, present algorithms for their construction, and compare the social disutilities of the players at these points. An interesting property is that while in worst equilibrium some of the players exploit the system, the structure of the best equilibrium is such that the system exhibits fair allocation of resources to the players.³

We also consider the ratio between the social disutilities incurred by players in an equilibrium and in an optimal outcome. The worst possible ratio between equilibrium and optimum disutilities (dubbed “the price of anarchy” [19]) was proposed by Koutsoupias and Papadimitriou [11] as a measure of the inefficiency of selfish behavior in noncooperative systems, and was extensively studied for nonatomic congestion games [8,27,28,26]. Recently, Awerbuch et al. [4] and, independently, [7] provided bounds for the price of anarchy in finite congestion games with linear cost functions. We show that in congestion games with failures even the best possible ratio between equilibrium and optimum disutilities (termed the “price of stability” in [1]) depends on the parameters of the game and cannot be bounded by a constant value, even for very simple (e.g., linear) cost functions. As a result, the price of anarchy is also unbounded.

The paper is organized as follows. In Section 2 we define our model. In Section 3 we show that CGFs do not admit a potential function and the finite improvement property (FIP). In Section 4 we prove the convergence of best replies, implying the existence of a pure strategy Nash equilibrium. In addition, we provide an efficient algorithm for computing and analyzing Nash equilibrium uniqueness properties. In Section 5 we show that every best Nash equilibrium of any given CGF is semi-strong. Section 6 is devoted to symmetric CGFs where we characterize best and worst Nash equilibria in symmetric CGFs, present algorithms for their construction and provide an upper bound on the ratio between them. We also discuss the (best and worst possible) ratio between social disutility in Nash equilibrium and optimum social disutility in these games. We conclude by drawing a few directions for further research in Section 7.

2. The model

A CGF is defined as follows. Let N be a set of $n \in \mathbb{N}$ players, and let M be a set of $m \in \mathbb{N}$ service providers, each associated with a *failure probability*. We assume that the failure or success of a particular service provider is *independent* of the failure or success of other SPs. Each player has a task which can be carried out by any of the SPs. The *service cost* for player i for utilizing service provider e is a nonnegative nondecreasing function $l_e^i : \{1, \dots, n\} \rightarrow \mathbb{R}_+$ of the congestion experienced by e ; thus, the greater is the congestion experienced by the service provider, the longer it takes to complete the task execution and, as a result, the greater is the cost incurred by the player. Player i 's disutility from a successful task completion is determined by the *minimum* of the service costs of the SPs he has chosen which did not fail. Player i 's disutility from an uncompleted task is evaluated by his (nonnegative) *incompletion cost* (denoted by W_i). This is defined more precisely below.

The *success probability* of $e \in M$ is denoted by $s_e (s_e \in (0, 1))$. Similarly, $f_e = 1 - s_e$ stands for the *failure probability* of e .

Remark 1. The CGF-model can be extended to allow congestion-dependent failure/success probabilities, while leading to similar results (see Remark 8 in Section 4 for justification). We define these probabilities as constants solely in order to simplify the exposition and presentation of our results.

The set of pure strategies Σ_i for player $i \in N$ is the power set of the set of SPs: $\Sigma_i = P(M)$. Given a subset of players $S \subseteq N$, the set of strategy combinations of the members of S is denoted by $\Sigma_S = \times_{i \in S} \Sigma_i$, and the set of strategy combinations of the complement subset of players is denoted by $\Sigma_{-S} (\Sigma_{-S} = \Sigma_{N \setminus S} = \times_{i \in N \setminus S} \Sigma_i)$. The set of pure strategy profiles of all the players is denoted by $\Sigma (\Sigma = \Sigma_N)$.

Let $\sigma = (\sigma_1, \dots, \sigma_n) \in \Sigma$ be a combination of pure strategies. The (m -dimensional) *congestion vector* that corresponds to σ is $h^\sigma = (h_e^\sigma)_{e \in M}$, where $h_e^\sigma = |\{i \in N : e \in \sigma_i\}|$. The *outcome* from σ is the subset $X \subseteq M$ of the service providers that have successfully executed their assigned tasks. For any player $i \in N$, if $\sigma_i \cap X = \emptyset$ then the *disutility* of player i from a strategy profile σ and the outcome X , $\pi_i(\sigma, X)$, is equal to his incompletion cost, W_i ; otherwise, if $X \neq \emptyset$, then the i 's disutility is determined by the *minimum* among the service costs of his successful resources:

$$\pi_i(\sigma, X) = \min_{e \in \sigma_i \cap X} l_e^i(h_e^\sigma).$$

Given a strategy profile σ , let $\mathbb{X}(\sigma)$ denote the random variable representing the subset of successful SPs; $\mathbb{X}(\sigma)$ is distributed over the power set of the set of chosen service providers, $P(\cup_{i \in N} \sigma_i)$, and its distribution is determined by $(f_e)_{e \in \cup_{i \in N} \sigma_i}$. The *expected disutility* of player i from strategy profile σ , $\Pi_i(\sigma)$, is therefore:

$$\Pi_i(\sigma) = W_i \prod_{e \in \sigma_i} f_e + \sum_{A \in P(\sigma_i) \setminus \{\emptyset\}} \min_{e \in A} l_e^i(h_e^\sigma) \prod_{e \in A} s_e \prod_{e \in \sigma_i \setminus A} f_e.$$

We use the convention that $\prod_{e \in \emptyset} f_e = 1$. Hence, if agent i chooses an empty set $\sigma_i = \emptyset$ (does not assign his task to any resource), then his expected disutility, $\Pi_i(\emptyset, \sigma_{-i})$, equals his incompletion cost, W_i . The aim of each player is to minimize his own expected disutility.

³ The issue of fairness in congestion settings is the topic of study of several papers (see e.g. [9,12]).

We note that the service cost of a particular resource to player i may be higher than his incompleteness cost, W_i . For instance, the service cost may consist of two types of costs—an *execution cost*, b_e^i , that represents a cost of processing player i 's task by service provider e , and a *fixed completion cost*, a , that models for example a payment to the network administrator for successful execution of a task, by one or more of the service providers. It is natural to assume that $b_e^i(\cdot)$ is a nonnegative nondecreasing function satisfying $b_e^i(k) \leq W_i$ for all $i \in N$, $e \in M$ and $1 \leq k \leq n$. This means that the execution of a task does not cost a player more than its failure. W.l.o.g., one can also assume that for any player i his incompleteness cost W_i is higher than the fixed completion cost a . Otherwise, the obvious dominant strategy of player i is to avoid assigning his task to any service provider. Despite the fact that $a \leq W_i$ and $b_e^i(k) \leq W_i$ for all $1 \leq k \leq n$, the service cost, $l_e^i(k) = b_e^i(k) + a$, might be higher than W_i . However, it is obvious that if $l_e^i(1) \geq W_i$ for all $e \in M$, then the dominant strategy of player i is to avoid assigning his task to any service provider, i.e. in this case player i can be actually ignored. Therefore, w.l.o.g., we assume that such cases do not take place.

Remark 2. We note that CGFs with constant failure probabilities (i.e., $f_e(k) = f_e$ for all $e \in M$ and $k = 1, \dots, n$) and a single service provider ($m = 1$) can be reduced to congestion games with player-specific payoff functions presented in [14]. If player i chooses strategy $\sigma_i = \emptyset$, then his disutility is the constant W_i (Note that W_i is player-specific); otherwise, if he submits his task to the single available service provider, e , then his expected disutility is given by $\Pi_i(\sigma) = W_i f_e + l_e^i(h_e^e)(1 - f_e)$, which is a nondecreasing, player-specific function of the congestion on e . This is equivalent to the choice between two resources, when the cost of each resource is a (possibly constant) nondecreasing, player-specific function of its congestion, which is exactly the case in [14]. Observe that $\Pi_i(\cdot)$ is nondecreasing since $l_e^i(\cdot)$ is nondecreasing and f_e is a constant value. However, if failure probabilities depend on the congestion, then $\Pi_i(\cdot)$ may be non-monotonic.

Applying Remark 2, when discussing CGFs with constant failure probabilities, we can restrict our attention to the case with 2 or more resources ($m \geq 2$).

3. The non-existence of a potential function and of a FIP

In this section we investigate the basic properties of congestion games with failures. Specifically, we observe that no CGF with symmetric failure probabilities and service cost functions admits an (exact) potential function, thus implying that the class of CGFs lies beyond the class of congestion games. Furthermore, we show that these games, in general, do not possess even a weaker – *generalized ordinal* – potential, whose existence in a game is equivalent to the so-called “finite improvement property” (FIP) that guarantees that any sequence of one-sided improving steps converges to a pure strategy Nash equilibrium from any initial combination of players’ strategies. We present an example of a CGF in which such a sequence of improvements has a cycle.

3.1. Exact potential

Monderer and Shapley [17] introduced the notion of *potential function* and defined a *potential game* to be a game which possesses a potential function. A potential function (or, an *exact potential*) is a real-valued function over the set of pure strategy profiles, with the property that the gain (or loss) of a player shifting to another strategy while keeping the other players’ strategies unchanged is equal to the corresponding increment of the potential function. That is, if Γ is a game in strategic form with a finite number of players, where the set of strategies of player i is Σ_i , the set of strategy profiles is $\Sigma = \times_i \Sigma_i$, and the payoff function of player i is $\Pi_i : \Sigma \rightarrow \mathbb{R}$, then a function $P : \Sigma \rightarrow \mathbb{R}$ is an (exact) potential for Γ if for every player i and for every $\sigma_{-i} \in \Sigma_{-i}$,

$$\Pi_i(\sigma_{-i}, x) - \Pi_i(\sigma_{-i}, y) = P(\sigma_{-i}, x) - P(\sigma_{-i}, y),$$

for any $x, y \in \Sigma_i$. The authors [17] showed that the classes of finite potential games and congestion games coincide.

In this section we show that the class of CGFs does not possess a potential function, and therefore is not isomorphic to the class of congestion games. In particular, we show that no *symmetric* CGF (in which the failure probabilities, the incompleteness costs and the service costs do not depend on the service provider or the player identity)⁴ admits a potential function. To prove this statement we employ the following technical characterization of a potential game by Monderer and Shapley [17]: if Γ is a game in strategic form with $\Pi_i : \Sigma \rightarrow \mathbb{R}$ the payoff function of player i , then Γ is a potential game if and only if for every $i, j \in N$, for every $z \in \Sigma_{-\{i,j\}}$, and for every $x_i, y_i \in \Sigma_i$ and $x_j, y_j \in \Sigma_j$,

$$\Pi_i(\alpha) - \Pi_i(\beta) + \Pi_j(\beta) - \Pi_j(\gamma) + \Pi_i(\gamma) - \Pi_i(\delta) + \Pi_j(\delta) - \Pi_j(\alpha) = 0,$$

where $\alpha = (x_i, x_j, z)$, $\beta = (y_i, x_j, z)$, $\gamma = (y_i, y_j, z)$, $\delta = (x_i, y_j, z)$.⁵

If in a CGF all service cost functions are constant (we refer to such games as “degenerate”), then the above characterization easily implies the existence of a potential function (this follows from the fact that the expected disutility of each player in this case is independent of the choices of the other players). In other cases, a potential function does not necessarily exist in CGFs. In particular, it never exists in symmetric CGFs.

⁴ See Section 6 for a detailed discussion.

⁵ $\alpha \rightarrow \beta \rightarrow \gamma \rightarrow \delta \rightarrow \alpha$ is termed “a simple closed path” of length 4.

	$\{e_2\}$	$\{e_1, e_2\}$
\emptyset	$\Pi_1(\alpha) = W$ $\Pi_2(\alpha) = fW + (1 - f)l(k)$	$\Pi_1(\delta) = W$ $\Pi_2(\delta) = f^2W + (1 - f^2)l(k)$
$\{e_1\}$	$\Pi_1(\beta) = fW + (1 - f)l(k)$ $\Pi_2(\beta) = fW + (1 - f)l(k)$	$\Pi_1(\gamma) = fW + (1 - f)l(k + 1)$ $\Pi_2(\gamma) = f^2W + (1 - f)l(k)$ $+ f(1 - f)l(k + 1)$

Fig. 1. The non-existence of a potential function in symmetric CGFs.

	\emptyset	$\{e_1\}$	$\{e_2\}$	$\{e_1, e_2\}$
\emptyset	(10, 10)	(10, 9)	(10, 10)	(10, 9)
$\{e_1\}$	(6, 10)	(7, 10)	(6, 10)	(7, 10)
$\{e_2\}$	(10, 10)	(10, 9)	(10, 10)	(10, 9)
$\{e_1, e_2\}$	(6, 10)	(7, 10)	(6, 10)	(7, 10)

Fig. 2. Players' payoffs in Γ .

Proposition 3. A non-degenerate symmetric CGFs with constant failure probabilities does not possess an exact potential function.

Proof. Let Γ be a non-degenerate symmetric CGF with f denoting the failure probability of each resource, W standing for the incomplection cost of each player, and $l(\cdot)$ representing the congestion-dependent service cost of each resource to each player. Let $k \in \{1, \dots, n - 1\}$ be an arbitrary integer satisfying $l(k) < l(k + 1)$, and consider the simple closed path of length 4 which is formed by $\alpha = (\emptyset, \{e_2\}, z)$, $\beta = (\{e_1\}, \{e_2\}, z)$, $\gamma = (\{e_1\}, \{e_1, e_2\}, z)$, $\delta = (\emptyset, \{e_1, e_2\}, z)$, where $z \in \Sigma_{-\{1,2\}}$ satisfies $h_{e_1}^z = h_{e_2}^z = k - 1$. That is, each SP is chosen by exactly $k - 1$ players, excluding players 1 and 2. For instance, let $k - 1$ players in $N \setminus \{1, 2\}$ play $\{e_1, e_2\}$ and all the others play \emptyset . The expected disutilities of the deviators (players 1 and 2) on the path $\alpha \rightarrow \beta \rightarrow \gamma \rightarrow \delta \rightarrow \alpha$ are presented in Fig. 1. Exploring Fig. 1, we get

$$\Pi_1(\alpha) - \Pi_1(\beta) + \Pi_2(\beta) - \Pi_2(\gamma) + \Pi_1(\gamma) - \Pi_1(\delta) + \Pi_2(\delta) - \Pi_2(\alpha) = (1 - f)^2(l(k + 1) - l(k)) > 0,$$

which implies the non-existence of a potential function. \square

Note that Proposition 3 is also true for many CGFs with congestion-dependent failures but not to all such games. Also the proposition applies to CGFs with symmetric service cost functions and failure probabilities but player-specific incomplection costs (the same proof, with the changes required due to player-specific incomplection costs, is valid).

Remark 4. Note that since there is no potential function in symmetric CGFs, its absence is a result of the added features of the CGF's setting (namely, the resource failures and the minimum operator in the players' objectives), and not due to the player-specific service cost functions. Moreover, some, but not many, non-degenerate CGFs with player-specific service costs possess a potential function. Such a game is demonstrated by the following example.

Example 5. Let Γ be a CGF in which two players $N = \{1, 2\}$ share a set of two resources $M = \{e_1, e_2\}$. The incomplection cost of each player is $W = 10$, the failure probability of each resource is $f = 0.5$ and the service cost function is given by $l_{e_1}^1(1) = 2, l_{e_1}^1(2) = 4; l_{e_1}^2(1) = 8, l_{e_1}^2(2) = 10$; and $l_{e_2}^1(1) = l_{e_2}^1(2) = l_{e_2}^2(1) = l_{e_2}^2(2) = 10$. Fig. 2 presents the payoff matrix of the game. A potential function of the game is presented in Fig. 3.

By exploring Figs. 2 and 3, one can verify that for any two strategy profiles differing by the choice of exactly one player, the difference in the payoff of that player between the two profiles equals the corresponding increment in the function presented in Fig. 3. Therefore, this function is a potential function.

One can easily see that any local optimum of a potential function of a finite potential game is a Nash equilibrium strategy profile as no player can improve his payoff by a unilateral deviation from this profile. Moreover, any sequence of myopic

	\emptyset	$\{e_1\}$	$\{e_2\}$	$\{e_1, e_2\}$
\emptyset	4	3	4	3
$\{e_1\}$	0	0	0	0
$\{e_2\}$	4	3	4	3
$\{e_1, e_2\}$	0	0	0	0

Fig. 3. A potential function of Γ .

improving deviations will converge to such a local optimum, regardless of where the process has started. This is called the *finite improvement property* [17].

3.2. (Generalized) ordinal potential and the FIP

Note that the convergence of a myopic improvement dynamics is guaranteed even by a weaker type of a potential function, for which it is only required that an increase in the utility of a player who unilaterally shifts to another strategy implies an increase in the potential function; thus, the potential increases along the improvement path, and the finite improvement property holds. This type of a potential is termed *generalized ordinal potential*, and it has been shown that its existence in a game is *equivalent* to the FIP [17]. Below we show that this property, in general, does not hold for congestion games with failures.

Theorem 6. *There exist CGFs with no finite improvement property.*

Proof. Let Γ be a CGF with 2 players $N = \{1, 2\}$ sharing 4 resources $M = \{e_1, e_2, e_3, e_4\}$, the failure probability of each of which is given by $f < 1$. The incompleteness costs of both players are $W_1 = W_2 = 2$. The service costs are given by $l_{e_1}^1(1) = l_{e_4}^2(1) = 1, l_{e_1}^1(2) = l_{e_4}^2(2) = 3$, and the rest take value 2 (i.e., for any $k = 1, 2$ we have $l_e^i(k) = 2$ for $e = e_2, e_3, e_4$ when $i = 1$, and for $e = e_1, e_2, e_3$ when $i = 2$). Consider the cycle of better replies which is formed by $\alpha = (\{e_1, e_2, e_3, e_4\}, \{e_1, e_3\}) \rightarrow \beta = (\{e_2, e_3\}, \{e_1, e_3\}) \rightarrow \gamma = (\{e_2, e_3\}, \{e_2, e_4\}) \rightarrow \delta = (\{e_1, e_2, e_3, e_4\}, \{e_2, e_4\}) \rightarrow \alpha$ (see Fig. 4).

The expected disutilities of the deviators satisfy the following:

$$\begin{aligned} \Pi_1(\alpha) &= 2 + f^3(1 - f) > 2 = \Pi_1(\beta) \\ \Pi_2(\beta) &= 2 > 1 + f = \Pi_2(\gamma) \\ \Pi_1(\gamma) &= 2 > 1 + f = \Pi_1(\delta) \\ \Pi_2(\delta) &= 2 + f(1 - f) > 2 = \Pi_2(\alpha), \end{aligned}$$

implying that $\alpha \rightarrow \beta \rightarrow \gamma \rightarrow \delta \rightarrow \alpha$ is indeed a better reply cycle.⁶ □

Thus, CGFs have no FIP and have no generalized ordinal potential function.

4. Pure strategy Nash equilibria and best reply dynamics

We have shown in Section 3 that congestion games with failures have no FIP and no generalized ordinal potential function. Note, however, that this fact, in general, does not contradict the existence of an equilibrium in pure strategies or the convergence of particular one-sided better reply dynamics (e.g., *best responses*). In this section, we prove that any sequence of best responses always converges in CGFs, thus implying the existence of a pure strategy Nash equilibrium in these games. Moreover, we provide an efficient algorithm for its construction, and show that given a CGF, its different equilibria correspond to (almost) the same congestion vector.

4.1. Existence and convergence

We start by proving an intuitive and useful property of CGFs (see Proposition 7) which is the basis of our algorithm for constructing a pure strategy Nash equilibrium.

⁶ Though the first inequality is not satisfied by every f , it does hold for all $f \leq \frac{3}{4}$; the rest are true for any $f < 1$.

e_1	e_2	e_3	e_4
1	1	1	1
2		2	
2	1	1	
		2	
	1	1	2
	2		
1	1	1	1
	2		2
1	1	1	1
2		2	

Fig. 4. Example for the CGF without FIP.

Proposition 7. Let $\sigma \in \Sigma$ be a pure strategy profile of a given CGF, and let h^σ be its corresponding congestion vector. The strategy profile σ is a Nash equilibrium if and only if the following conditions are satisfied for all $i \in N$:

- (i) $l_e^i(h_e^\sigma) \leq W_i \quad \forall e \in \sigma_i$;
 - (ii) $l_e^i(h_e^\sigma + 1) \geq W_i \quad \forall e \notin \sigma_i$.
- (1)

Note that if $h_e^\sigma = n$ then $\{i \in N : e \notin \sigma_i\} = \emptyset$, and (1)(ii) is satisfied vacuously.

Proof. (\Rightarrow) Suppose σ is a Nash equilibrium pure strategy profile then we prove that σ satisfies inequalities (1).

- (i) Suppose there is a player $i \in N$ such that the subset $L_i = \{e \in \sigma_i : l_e^i(h_e^\sigma) > W_i\}$ of his chosen resources is not empty. Let $a \in L_i$ be such a resource with maximum service cost, that is, $a \in \arg \max_{e \in L_i} l_e^i(h_e^\sigma)$ (note that since $l_e^i(h_e^\sigma) \leq W_i$ for all $e \in \sigma_i \setminus L_i$, we have that $l_a^i(h_a^\sigma) \geq l_e^i(h_e^\sigma)$ for all $e \in \sigma_i$). We show below that $\Pi_i(\sigma) > \Pi_i(\sigma_{-i}, \sigma_i \setminus \{a\})$, i.e. player i can improve his expected disutility by removing resource a from his strategy, which contradicts $\sigma \in NE$. To prove this we show, on a *sample path* basis, that the above deviation is profitable for i . That is, for each realization of the resource failures, player i 's disutility obeying his modified strategy, is at most as in his original strategy, and there exists a realization for which his disutility is strictly smaller. For any given realization, let X denote the outcome from the strategy profile σ —the subset of SPs that have succeeded to execute their tasks, and let $\pi_i(\sigma, X)$ be player i 's disutility from σ and X . If all SPs have failed to execute their tasks (the subset of successful resources is empty, $X = \emptyset$) then $\pi_i(\sigma, \emptyset) = \pi_i((\sigma_{-i}, \sigma_i \setminus a), \emptyset) = W_i$. If $X = \{a\}$, then player i benefits from removing a from his set of utilized SPs: $\pi_i((\sigma_{-i}, \sigma_i \setminus \{a\}), X) = W_i < l_a^i(h_a^\sigma) = \pi_i(\sigma, X)$. If $X \neq \emptyset, \{a\}$, then $\pi_i((\sigma_{-i}, \sigma_i \setminus \{a\}), X) = \min_{e \in (\sigma_i \setminus \{a\}) \cap X} l_e^i(h_e^\sigma) = \min_{e \in \sigma_i \cap X} l_e^i(h_e^\sigma) = \pi_i(\sigma, X)$. Therefore, $\Pi_i(\sigma_{-i}, \sigma_i \setminus \{a\}) < \Pi_i(\sigma)$, contradicting σ being a Nash equilibrium strategy profile.

- (ii) Suppose there are player i and service provider $a \in M \setminus \sigma_i$ such that $l_a^i(h_a^\sigma + 1) < W_i$. We show below that $\Pi_i(\sigma) > \Pi_i(\sigma_{-i}, \sigma_i \cup \{a\})$, which contradicts $\sigma \in NE$.

We denote $(\sigma_{-i}, \sigma_i \cup \{a\})$ by σ' . If all SPs have failed to execute their assigned tasks, then the i 's disutility from both σ and σ' is equal to his incompleteness cost, W_i . If all the resources, excluding resource a , have failed ($X = \{a\}$), then the i 's disutility from σ' is less than the one from σ : $\pi(\sigma', X) = l_a^i(h_a^{\sigma'}) = l_a^i(h_a^\sigma + 1) < W_i = \pi(\sigma, X)$. If $X \neq \emptyset, \{a\}$, then $\pi(\sigma', X) = \min_{e \in (\sigma_i \cup \{a\}) \cap X} l_e^i(h_e^{\sigma'}) \leq \min_{e \in \sigma_i \cap X} l_e^i(h_e^\sigma) = \pi(\sigma, X)$ (note that $h_e^{\sigma'} = h_e^\sigma$ for all $e \in M \setminus \{a\}$). Therefore, $\Pi_i(\sigma') < \Pi_i(\sigma)$, contradicting σ being a Nash equilibrium strategy profile.

(\Leftarrow) Let σ be a strategy profile that satisfies conditions (1). We prove that σ is a Nash equilibrium, that is, $\Pi_i(\sigma) \leq \Pi_i(\sigma_{-i}, \sigma'_i)$ for each player $i \in N$ and for all $\sigma'_i \in \Sigma_i$.

Let $i \in N$ be any player and let $\sigma'_i \in \Sigma_i$ be any strategy of player i . We denote (σ_{-i}, σ'_i) by σ' and note that

$$h_e^{\sigma'} = \begin{cases} h_e^\sigma & e \in \sigma_i \cap \sigma'_i; \\ h_e^\sigma + 1 & e \in \sigma'_i \setminus \sigma_i. \end{cases}$$

We show that for any realization of the resource failures, player i 's payoff from obeying σ is no worse than his payoff from σ' , implying σ is an equilibrium. For a given realization, let X represent the (possibly empty) subset of successful service providers and note that $\pi_i(\sigma, X) \leq W_i$ (this is since if $\sigma_i \cap X = \emptyset$ then $\pi_i(\sigma, X) = W_i$; otherwise, $\pi_i(\sigma, X) = \min_{e \in \sigma_i \cap X} l_e^i(h_e^\sigma) \leq_{(1)(i)} W_i$). If $\sigma_i \cap \sigma'_i \cap X = \emptyset$ then $\pi_i(\sigma', X) \geq W_i \geq \pi_i(\sigma, X)$ (if $\sigma'_i \cap X = \emptyset$ then $\pi_i(\sigma', X) = W_i$; otherwise, $\pi_i(\sigma', X) = \min_{e \in \sigma'_i \cap X} l_e^i(h_e^{\sigma'}) \geq_{(1)(ii)} W_i$). Otherwise, if $\sigma_i \cap \sigma'_i \cap X \neq \emptyset$, then $\pi_i(\sigma', X) = \min_{e \in \sigma'_i \cap X} l_e^i(h_e^{\sigma'}) =_{(1)} \min_{e \in \sigma_i \cap \sigma'_i \cap X} l_e^i(h_e^\sigma) \geq \min_{e \in \sigma_i \cap X} l_e^i(h_e^\sigma) = \pi_i(\sigma, X)$ (the equality holds since $l_e^i(h_e^\sigma) \leq_{(1)(i)} W_i \leq_{(1)(ii)} l_{e'}^i(h_{e'}^\sigma + 1) = l_{e'}^i(h_{e'}^{\sigma'})$ for any $e \in \sigma_i \cap \sigma'_i$ and $e' \in \sigma'_i \setminus \sigma_i$; the inequality is due to $\sigma_i \cap \sigma'_i \cap X \subseteq \sigma_i \cap X$). Therefore, $\Pi_i(\sigma) \leq \Pi_i(\sigma')$ for all $i \in N$ and $\sigma'_i \in \Sigma_i$, as claimed. That is, σ is a Nash equilibrium strategy profile. \square

Remark 8. Proposition 7 reflects the idea that a player's decision regarding any SP, that is, to sign up or not to an SP, in an equilibrium outcome is independent of his decisions regarding other SPs, and is a function of his incompleteness cost and the service cost of that SP. In addition, the proposition implies that the equilibrium outcomes are independent of the values of resource failure probabilities. Thus, Proposition 7 also holds for CGF-models with congestion-dependent failure probabilities.

Based on Proposition 7, at any strategy profile of a given CGF, a best response of a player to his opponents' chosen strategies is uniquely defined by his incompleteness cost and the congestion on the resources caused by the rest of the players. We now show that any sequence of best replies always converges to a Nash equilibrium profile, regardless of the initial point at which the improvement process has started. This property is termed the *finite best response property* (FBRP) [14].

Theorem 9. Every congestion game with failures has the FBRP.

Proof. On the contrary, assume there exists a game with a cyclic path of best responses. Fix such a game and such a cycle, and an arbitrary service provider, e , which is involved in this cycle. Consider changes made to e along the best response path; since changes made to other service providers have no effect on e , without loss of generality, we can focus on only those steps at which a player have added or dropped resource e . Note that since the game is finite with n players, the cycle must include both adds and drops, so fix any add with resource e and let it be the starting point, denoted σ^0 .

Consider the function $P^{BR}(\sigma) = L \cdot W^{\max}(\sigma) + n^{\max}(\sigma)$, where L is a large number so that $\min_{i,j \in N: W_i \neq W_j} |W_i - W_j| \cdot L > n$; W^{\max} is the maximal incompleteness cost among all users of e : $W^{\max}(\sigma) = \max_{i \in N: e \in \sigma_i} W_i$; and n^{\max} is the number of e 's users with the maximal incompleteness cost: $n^{\max}(\sigma) = |\{i \in N | e \in \sigma_i, W_i = W^{\max}(\sigma)\}|$.

First, observe that the function P^{BR} does not decrease along the best response path. Recall that at step σ^0 some player adds resource e . By Proposition 7, the incompleteness cost of the deviator has to be at least as great as the service cost of e at σ^0 . By definition, the same holds for any player i with $W_i = W^{\max}(\sigma^0)$. Hence, no such player will drop resource e (meaning that W^{\max} , n^{\max} , and hence P^{BR} , will not decrease), unless its service cost exceeds $W^{\max}(\sigma^0)$. But for this to happen, a player with a higher incompleteness cost must add resource e , implying that W^{\max} strictly increases. Note that the terms in the function P^{BR} are scaled in a way that any possible change in W^{\max} is greater by the absolute value than the total number of players, n . So, any time W^{\max} grows, the function P^{BR} grows too, although n^{\max} could decrease. Next, we show that either P^{BR} strictly increases along the improvement path, or there is another (very similar) function that does. This will complete the proof.

If P^{BR} is strictly increasing, we are done. So assume otherwise. Recall that the function can only stay unchanged if none of the $n^{\max}(\sigma^0)$ users of e with the maximal incompleteness cost at σ^0 made any change along the improvement path. Therefore, we can "ignore" those players and either remove them from the game and modify the service cost functions of the rest of the players accordingly, or, alternatively, consider a function $P^{BR}_{(2)}(\sigma) = L \cdot W^{(2)\max}(\sigma) + n^{(2)\max}(\sigma)$, where $W^{(2)\max}$ and $n^{(2)\max}$ are defined in a similar way, but for the 2nd highest incompleteness cost, and repeat the argument with respect to $P^{BR}_{(2)}$. Since the game is finite, there exists $P^{BR}_{(k)}$ for $k = 1, \dots, n$ (defined similarly for the k th highest incompleteness cost) that strictly increases along the cycle, in contradiction to the contrary assumption. \square

Theorem 9 implies that every CGF possesses a pure strategy Nash equilibrium, and it can be achieved by a sequence of best replies. However, it gives no guarantees on how long such an improvement dynamics may take. In what follows, we strengthen the result of Theorem 9 by showing that fixing the initial strategy profile and the order of players in which they apply their best responses, one can easily obtain an equilibrium in time, polynomial in the number of players and service providers.

4.2. Efficient construction

We develop the NE-Algorithm for constructing Nash equilibria in CGFs. The algorithm is initialized with an empty strategy set for each player. At each iteration, the algorithm selects (arbitrarily) a service provider. Then it sorts the players in a non-increasing order of their maximum congestion for which the service cost of a player for using this service provider is less than his incompleteness cost. If the ordering number of the player is at most the above value of the congestion, then the algorithm adds the selected service provider to the player's strategy set. In such a way, at the end of the algorithm, for each i , the service cost of player i for using any of the SPs in σ_i is lower than W_i , while utilizing any other SP costs him more. Then, by Proposition 7, in such situations no player wishes to deviate (unilaterally) from his strategy. The NE-Algorithm is presented below.

4.2.1. NE-algorithm

Initialization:	For all $1 \leq i \leq n$, set $\sigma_i := \emptyset$;
Main	For all $e \in M$:
Step:	(1) Sort the players in a non-increasing order of \mathbf{k}_e^i , where $\mathbf{k}_e^i = \begin{cases} 0 & \text{if } \{k : W_i > l_e^i(k), k = 1, \dots, n\} = \emptyset \\ \max\{k : W_i > l_e^i(k), k = 1, \dots, n\} & \text{otherwise;} \end{cases}$ Let $\varphi_e : N \rightarrow \{1, \dots, n\}$ $i \mapsto i_e = \varphi_e(i)$ be the corresponding permutation function; (2) For $i_e = 1$ to n : if $i_e \leq \mathbf{k}_e^i$, then $\sigma_i := \sigma_i \cup \{e\}$.

Theorem 10. *The NE-algorithm finds a pure strategy Nash equilibrium in a given CGF with time complexity of $O(mn \log n)$.*

Proof (Validity). Let us first briefly describe the intuition behind the proof. For any resource e and for any player i we show that if the algorithm assigns player i 's task to resource e then the total congestion experienced by e in the outcome of the algorithm is at most \mathbf{k}_e^i . By the definition of \mathbf{k}_e^i , player i 's service cost for using resource e is less than the W_i , as required. Otherwise, if the algorithm does not assign player i 's task to resource e , then we show that the total congestion on e in the algorithm's outcome, increased by 1, is greater than \mathbf{k}_e^i , implying that the service cost of e is at least W_i , as required. The formal proof is presented below.

Let $\sigma = (\sigma_i)_{i \in N}$ be the combination of pure strategies constructed by the NE-Algorithm, and let $h^\sigma = (h_e^\sigma)_{e \in M}$ be the corresponding congestion vector. Choose any service provider e and any player i . If $e \in \sigma_i$ then $h_e^\sigma \geq 1$ and $\mathbf{k}_e^i \geq \mathbf{k}_e^{\varphi_e^{-1}(x)}$, where

$$x = \max_{j \in N} \{\varphi_e(j) : e \in \sigma_j\} = h_e^\sigma.$$

Since $e \in \sigma_x$, then $\mathbf{k}_e^{\varphi_e^{-1}(x)} \geq x$ and $\mathbf{k}_e^i \geq \mathbf{k}_e^{\varphi_e^{-1}(x)} \geq x = h_e^\sigma$. By the definition of \mathbf{k}_e^i , for all $1 \leq k \leq \mathbf{k}_e^i$, we have $W_i > l_e^i(k)$. Thus, since $1 \leq h_e^\sigma \leq \mathbf{k}_e^i$, we get $W_i > l_e^i(h_e^\sigma)$.

If $e \notin \sigma_i$ then $\mathbf{k}_e^i \leq \mathbf{k}_e^{\varphi_e^{-1}(y)}$, where

$$y = \min_{j \in N} \{\varphi_e(j) : e \notin \sigma_j\} = \max_{j \in N} \{\varphi_e(j) : e \in \sigma_j\} + 1 = h_e^\sigma + 1.$$

Since $e \notin \sigma_y$, then $\mathbf{k}_e^{\varphi_e^{-1}(y)} < y$, and $\mathbf{k}_e^i \leq \mathbf{k}_e^{\varphi_e^{-1}(y)} < y = h_e^\sigma + 1$. By the definition of \mathbf{k}_e^i , for all $k > \mathbf{k}_e^i$, we have $W_i \leq l_e^i(k)$. Thus, since $h_e^\sigma + 1 > \mathbf{k}_e^i$, we get $W_i \leq l_e^i(h_e^\sigma + 1)$.

Hence, the congestion vector h^σ satisfies the conditions in (1). Therefore, by Proposition 7, σ is a Nash equilibrium strategy profile.

Complexity The number of iterations of the NE-algorithm is m , where each iteration takes $O(n \log n)$ operations. Hence, the time complexity of the NE-algorithm is $O(mn \log n)$. \square

Therefore, a pure strategy Nash equilibrium in a given CGF is easy to find. It turns out, that equilibrium points in CGFs have some additional nice properties, as is shown in the sequel.

4.3. (Almost) Uniqueness

We now discuss some uniqueness properties of Nash equilibria in CGFs. More precisely, we consider the uniqueness of congestion experienced by any service provider in any equilibrium. The uniqueness of equilibrium may fail due to indifference in the choice between strategies. Hence we confine our attention to games with strictly increasing service cost functions. We show that in such CGFs the difference between the congestion experienced by any SP in any two different Nash equilibria is bounded by 1. Furthermore, in games with $l_e^i(\cdot) \neq W_i$, for all e and i , any SP has the same congestion in all equilibrium profiles. In particular, all generic CGFs have this uniqueness property. Let $NE \subseteq \Sigma$ be a set of Nash equilibrium pure strategy profiles. Then,

Proposition 11. *Given a CGF, if for all $e \in M$ and $i \in N$, $l_e^i(\cdot)$ is a strictly increasing monotone function, then for any pair of equilibrium profiles $\sigma^1, \sigma^2 \in NE$, the inequality $|h_e^{\sigma^1} - h_e^{\sigma^2}| \leq 1$ holds for all $e \in M$.*

Proof. Let $\sigma^1, \sigma^2 \in NE$ be Nash equilibrium strategy profiles, and assume that $h_e^{\sigma^1} > h_e^{\sigma^2} + 1$ for some $e \in M$. Then, there is a player i such that $e \in \sigma_i^1$, but $e \notin \sigma_i^2$. By Proposition 7, for player i we have $l_e^i(h_e^{\sigma^1}) \leq W_i$ and $l_e^i(h_e^{\sigma^2} + 1) \geq W_i$. Therefore, $l_e^i(h_e^{\sigma^1}) \leq l_e^i(h_e^{\sigma^2} + 1)$. Now, $h_e^{\sigma^1} > h_e^{\sigma^2} + 1$ coupled with the monotonicity of $l_e^i(\cdot)$, leads to $l_e^i(h_e^{\sigma^1}) > l_e^i(h_e^{\sigma^2} + 1)$, in contradiction to $l_e^i(h_e^{\sigma^1}) \leq l_e^i(h_e^{\sigma^2} + 1)$. \square

It can be easily seen that if in addition to the requirements of Proposition 11, the cost function $l_e^i(\cdot)$ satisfies $l_e^i(k) \neq W_i$ for $1 \leq k \leq n$, then all Nash equilibria of a given CGF correspond to the same congestion vector, i.e. the congestion of any SP is fixed for all equilibrium points.

5. Semi-strong Nash equilibria in CGFs

In the previous section we proved that CGFs, while not admitting a potential function, always possess a Nash equilibrium in pure strategies. An equilibrium strategy profile has the property that, once it has been agreed upon, no player has an incentive to deviate unilaterally. This interpretation has been criticized on various grounds (see, e.g., [3]). In particular, if the coalition of the whole can communicate to agree on a particular strategy choice, then smaller coalitions may also be able to communicate and coordinate their actions. According to this line of argument, one should not be satisfied with strategy profiles that are immune only to individual deviations, but instead should insist also on immunity to deviations of coalitions. These considerations have led to new solution concepts of *strong Nash equilibrium* [2] and *coalition-proof Nash equilibrium* [5,6]. Milgrom and Roberts [16] focused on coalition deviations that are robust against further individual deviations. They called the corresponding Nash equilibria *strongly coalition proof*. The same concept has been proposed by Kaplan [10], who called the equilibria *semi-strong*.

Strong Nash equilibrium is a strategy profile for which there is no profitable deviation available to any coalition of players. This requirement is too strong since the profile must be resistant to deviations which are not themselves resistant to further deviations. The notion of coalition-proof Nash equilibrium, instead, requires only that a strategy profile be immune to profitable deviations which are *self-enforcing*.⁷ Strongly coalition-proof (or, semi-strong) Nash equilibrium is a strategy profile for which there are no profitable coalition deviations which are robust against further individual deviations (we define these equilibria more precisely below). Given a semi-strong Nash equilibrium, one expects no coalition deviations, since any such deviation is not an equilibrium among the members of coalition.⁸

In this section, we show the existence of semi-strong Nash equilibria in the class of CGFs. Moreover, we show that every best Nash equilibrium of a given CGF is semi-strong.

Definition 12. Let $\Gamma = (N, \Sigma, (\Pi_i(\cdot))_{i \in N})$ be a game in strategic form, where N denotes the set of players, Σ denotes the set of strategy profiles, and Π_i denotes the disutility function of player $i \in N$. A profile σ is a **semi-strong Nash equilibrium** if

- (i) σ is a Nash equilibrium of Γ , and
- (ii) for every $S \subseteq N, S \neq \emptyset$, and every Nash equilibrium ρ_S of $\Gamma^{S,\sigma}$, where $\Gamma^{S,\sigma} = (S, (\Sigma_i)_{i \in S}, (\Pi_i(\sigma_{-S}, \cdot))_{i \in S})$ is the reduced game of Γ w.r.t. S and σ , there exists $i \in S$ such that $\Pi_i(\sigma) \leq \Pi_i(\rho_S, \sigma_{-S})$.

We restrict our attention to CGFs with service cost functions satisfying $l_e^i(k) \neq W_i$, for all $e \in M, i \in N, k \in \{1, \dots, n\}$. That is, we assume that a player is never indifferent between adding or ignoring a particular service provider. Let σ be a Nash equilibrium strategy profile that minimizes the sum of the players' disutilities, then σ is termed a *best Nash Equilibrium*. (See Section 6.2 for a detailed discussion on best Nash-equilibria.) The following theorem states that any best Nash Equilibrium is semi-strong.

Theorem 13. Let G be a CGF satisfying $l_e^i(k) \neq W_i$, for all $e \in M, i \in N, k \in \{1, \dots, n\}$, and let σ be a best Nash equilibrium. Then, σ is a semi-strong Nash equilibrium of G .

Proof. Assume on the contrary that there exists $S \subseteq N, S \neq \emptyset$, and a Nash equilibrium ρ_S of $\Gamma^{S,\sigma}$ such that $\Pi_i(\rho_S, \sigma_{-S}) < \Pi_i(\sigma)$, for all $i \in S$. Since ρ_S is a Nash equilibrium of $\Gamma^{S,\sigma}$, then, by Proposition 7, for all $i \in S$,

$$(i) l_e^i(h_e^{(\rho_S, \sigma_{-S})}) < W_i \quad \forall e \in \rho_i; \tag{2}$$

$$(ii) l_e^i(h_e^{(\rho_S, \sigma_{-S})} + 1) > W_i \quad \forall e \in M \setminus \rho_i. \tag{3}$$

If there exists a service provider $e \in M$ such that $h_e^{(\rho_S, \sigma_{-S})} < h_e^\sigma$, then there is a player $i \in S$ with $e \in \sigma_i \setminus \rho_i$. Hence,

$$h_e^{(\rho_S, \sigma_{-S})} < h_e^\sigma \Rightarrow h_e^{(\rho_S, \sigma_{-S})} + 1 \leq h_e^\sigma \Rightarrow l_e^i(h_e^{(\rho_S, \sigma_{-S})} + 1) \leq l_e^i(h_e^\sigma). \tag{4}$$

⁷ A deviation is self-enforcing if there is no further self-enforcing and profitable deviation available to a proper sub-coalition of players.

⁸ Obviously, every strong equilibrium is semi-strong, and every semi-strong equilibrium is coalition-proof.

Since σ is a Nash equilibrium of Γ and $e \in \sigma_i$, then by Proposition 7,

$$l_e^i(h_e^\sigma) < W_i \Rightarrow (4) l_e^i(h_e^{(\rho_S, \sigma_{-S})} + 1) < W_i,$$

in contradiction to $e \notin \rho_i$ and (3).

If there exists a service provider $e \in M$ such that $h_e^{(\rho_S, \sigma_{-S})} > h_e^\sigma$, then there is a player $i \in S$ with $e \in \rho_i \setminus \sigma_i$. Then,

$$h_e^{(\rho_S, \sigma_{-S})} > h_e^\sigma \Rightarrow h_e^{(\rho_S, \sigma_{-S})} \geq h_e^\sigma + 1 \Rightarrow l_e^i(h_e^{(\rho_S, \sigma_{-S})}) \geq l_e^i(h_e^\sigma + 1). \tag{5}$$

Since σ is a Nash equilibrium of Γ and $e \notin \sigma_i$, then by Proposition 7,

$$l_e^i(h_e^\sigma + 1) > W_i \Rightarrow (5) l_e^i(h_e^{(\rho_S, \sigma_{-S})}) > W_i,$$

in contradiction to $e \in \rho_i$ and (2).

Otherwise, $h_e^{(\rho_S, \sigma_{-S})} = h_e^\sigma$, for all $e \in M$. Then, $\Pi_i(\rho_S, \sigma_{-S}) = \Pi_i(\sigma)$ for all $i \notin S$ and $\Pi_i(\rho_S, \sigma_{-S}) < \Pi_i(\sigma)$ for all $i \in S$, which implies $\sum_{j \in N} \Pi_j(\rho_S, \sigma_{-S}) < \sum_{j \in N} \Pi_j(\sigma)$. That is, the sum of the players' disutilities from (ρ_S, σ_{-S}) is strictly less than the one from σ . In addition, since ρ_S is a Nash equilibrium of $\Gamma^{S, \sigma}$, then inequalities (2) and (3) hold for all $i \in S$; since σ is a Nash equilibrium of Γ and $h_e^{(\rho_S, \sigma_{-S})} = h_e^\sigma$ for all $e \in M$, then Proposition 7 implies that (2) and (3) hold for all $i \notin S$. Thus, (2) and (3) hold for all $i \in N$, implying that (ρ_S, σ_{-S}) is a Nash equilibrium of Γ (by Proposition 7), in contradiction to σ being a best Nash equilibrium of Γ . \square

6. Symmetric CGFs

In this section, we consider the properties of Nash equilibria in symmetric CGFs. In a symmetric CGF, the parameters of the game do not depend on service provider or player identity, i.e. for all $i \in N$ and $e \in M$ we have $W_i = W, f_e = f$, and $l_e^i(k) = l(k)$, for all $k \in \{1, \dots, n\}$. Recall that if $l(1) \geq W$ then the dominant strategy of each player is to avoid assigning a task. Therefore, it is assumed that $l(1) < W$.

First, we show that in symmetric CGFs all SPs have almost the same congestion at Nash equilibrium. Furthermore, we characterize and compare best and worst equilibria in symmetric CGFs, and provide efficient simple algorithms for their construction. We also evaluate the inefficiency of Nash equilibrium in these games by considering the (best and worst possible) ratio between Nash equilibrium and optimum outcomes.

6.1. (Almost) Even congestion

We start by showing that as a result of the symmetry among service providers in symmetric CGFs, the SPs are (almost) evenly congested in any equilibrium. Let $\mathbf{k} = \max\{k : l(k) < W, k = 1, \dots, n\}$ and $\mathbf{k}^* = \max\{k : l(k) \leq W, k = 1, \dots, n\}$. Then,

Proposition 14. *Given a symmetric CGF, let $\sigma \in NE$ be a Nash equilibrium strategy profile with its corresponding congestion vector h^σ . Then, $\mathbf{k} \leq h_e^\sigma \leq \mathbf{k}^*$ for all $e \in M$.*

Proof. Assume on the contrary that $h_a^\sigma < \mathbf{k}$ for some $a \in M$. Then, $h_a^\sigma + 1 \leq \mathbf{k}$, implying $l(h_a^\sigma + 1) < W$ (by the monotonicity of $l(\cdot)$ and the definition of \mathbf{k}). Since $h_a^\sigma < \mathbf{k} \leq n$ then $a \notin \sigma_i$ for some $i \in N$. Then, Proposition 7 implies $l(h_a^\sigma + 1) \geq W$, a contradiction.

Assume now that $h_b^\sigma > \mathbf{k}^*$ for some $b \in M$. Then, $l(h_b^\sigma) > W$ (by the monotonicity of $l(\cdot)$ and the definition of \mathbf{k}^*). Since $h_b^\sigma > \mathbf{k}^* \geq 1$ then $b \in \sigma_i$ for some $i \in N$. Then, Proposition 7 implies $l(h_b^\sigma) \leq W$, a contradiction. \square

Note that if $l(\cdot)$ strictly increases with the number of users then $\mathbf{k}^* - \mathbf{k} \leq 1$. Then, Corollary 15 follows directly from Proposition 14.

Corollary 15. *Given a symmetric CGF, if $l(k)$ is a strictly increasing monotone function on the interval $1 \leq k \leq \mathbf{k}^*$, then at any Nash equilibrium $\sigma \in NE$, the difference between the congestions of different SPs is bounded by 1, i.e. for all $\sigma \in NE$ and for all $a, b \in M$, the inequality $|h_a^\sigma - h_b^\sigma| \leq 1$ holds.*

As we can see, symmetric CGFs have the property that all resources suffer almost the same congestion, regardless of the particular equilibrium the players may converge to. Hence, the load on the system is expected to be fairly distributed among the resources.

6.2. Best and worst equilibria

Given a profile σ , define the (expected) social disutility $\Pi(\sigma)$ as the sum of the players' expected disutilities in this strategy profile: $\Pi(\sigma) = \sum_{i \in N} \Pi_i(\sigma)$. Social disutility is a standard mean to measure the cost suffered by the society as a whole.

A strategy profile that minimizes the social disutility over the set of strategy profiles is called a *social optimum*. A *best (worst) equilibrium* is a strategy profile that minimizes (maximizes) the social disutility over the set of equilibrium strategies. Let $BNE \subseteq NE(WNE \subseteq NE)$ denote the subset of all best (worst) pure strategy Nash equilibria. The social disutility in a best equilibrium describes the best result that can be obtained in a system with noncooperative selfish players. This notion is of particular interest when there exists a mediator who can suggest rational players a particular behavior. Such behavior should be stable against unilateral deviations, but the mediator can choose it to be the best behavior among the stable points. The worst equilibrium is of particular interest when no such mediator exists, and one is interested in the worst rational behavior the system may converge to.

In this section we characterize, construct and compare best and worst Nash equilibria in symmetric CGFs. Specifically, we show that the values of the social disutility in best and worst equilibrium points are very close. The ratios between these values and the optimum social disutility are discussed in the following subsection. We start by providing two lemmas that will be needed in the sequel.

Lemma 16. *Given a symmetric CGF, assume there exists $\sigma \in \Sigma \setminus \{(\emptyset, \dots, \emptyset)\}$ with $\max_{e \in M} l(h_e^\sigma) \geq W$. Let $a \in \arg \max_{e \in M} l(h_e^\sigma)$. Then, for any $i \in N$, $\hat{\sigma} = (\sigma_{-i}, \sigma_i \setminus \{a\})$ satisfies $\Pi(\hat{\sigma}) \leq \Pi(\sigma)$.*

Proof. Let $i \in N$. If $a \notin \sigma_i$, then clearly $\Pi(\hat{\sigma}) \leq \Pi(\sigma)$. Otherwise $a \in \sigma_i$ and the strategy profile $\hat{\sigma}$ is obtained from σ by removing resource a from the strategy set of player i . As a result, the congestion on resource a is reduced by 1, and the other resources' congestion is unchanged. Then, by the monotonicity of $l(\cdot)$, the expected disutility of any player $k \neq i$ cannot increase. It remains to show that player i 's expected disutility, Π_i , does not increase, either.

For any realization of resource failures, as before, let X be the subset of all successful SPs. If $a \notin X$ then $\hat{\sigma}_i \cap X = (\sigma_i \setminus \{a\}) \cap X = \sigma_i \cap X$, and hence $\pi_i(\hat{\sigma}, X) = \pi_i(\sigma, X)$. If $\sigma_i \cap X = \{a\}$ then $\hat{\sigma}_i \cap X = \emptyset$ and $\pi_i(\hat{\sigma}, X) = W \leq l(h_a^\sigma) = \pi_i(\sigma, X)$. Otherwise, if $a \subsetneq \sigma_i \cap X$ then $\hat{\sigma}_i \cap X \neq \emptyset$ and $\pi_i(\hat{\sigma}, X) = \min_{e \in \hat{\sigma}_i \cap X} l(h_e^\sigma) = \min_{e \in (\sigma_i \setminus \{a\}) \cap X} l(h_e^\sigma) = \min_{e \in \arg \max_{e \in M} l(h_e^\sigma)} l(h_e^\sigma) = \pi_i(\sigma, X)$. Therefore, the expected disutility of i satisfies $\Pi(\hat{\sigma}) \leq \Pi(\sigma)$, on a sample path basis. \square

Lemma 16 reflects the simple idea that keeping resources with high service costs is not beneficial for the society. The following lemma, Lemma 17, captures the idea that strategy profiles which are better for the society are also fair⁹ ones.

Lemma 17. *Given a symmetric CGF, let $\sigma \in \Sigma$ be a combination of strategies satisfying $l(h_e^\sigma) \leq W$ for all $e \in M$. If there are two players $i, j \in N$ such that $|\sigma_i| > |\sigma_j| + 1$, then the strategy profile $\hat{\sigma} = (\sigma_{-\{i,j\}}, \sigma_i \setminus \{b\}, \sigma_j \cup \{b\})$, where $b \in \arg \max_{e \in \sigma_i \setminus \sigma_j} h_e^\sigma$, satisfies $\Pi(\hat{\sigma}) \leq \Pi(\sigma)$.*

The proof of Lemma 17 is technical and relatively lengthy, therefore we have chosen to present it in the Appendix.

Now we are ready to study best and worst equilibria in symmetric games. Recall that $\mathbf{k} = \max\{k : l(k) < W, k = 1, \dots, n\}$, and let $\Sigma^* \subseteq \Sigma$ be the subset of all pure strategy profiles satisfying the following conditions: for all $\sigma \in \Sigma^*$,

$$\begin{aligned} h_e^\sigma &= \mathbf{k} \quad \forall e \in M; \\ ||\sigma_i| - |\sigma_j|| &\leq 1 \quad \forall i, j \in N. \end{aligned} \tag{6}$$

Assume Σ^* is not empty and let $\sigma \in \Sigma^*$. Since the congestion on each service provider in σ equals \mathbf{k} , then $\sum_{i \in N} |\sigma_i| = m\mathbf{k}$. Also, since in σ all the players use (almost) the same number of resources, σ has the following structure: x players choose $\lfloor \frac{m\mathbf{k}}{n} \rfloor$ service providers and y players choose $\lfloor \frac{m\mathbf{k}}{n} \rfloor + 1$ service providers, where x and y satisfy the following equations:

$$\begin{cases} x \lfloor \frac{m\mathbf{k}}{n} \rfloor + y \left(\lfloor \frac{m\mathbf{k}}{n} \rfloor + 1 \right) = m\mathbf{k} \\ x + y = n. \end{cases}$$

Therefore, the values of x and y are

$$\begin{aligned} x &= n \left(\left\lfloor \frac{m\mathbf{k}}{n} \right\rfloor + 1 \right) - m\mathbf{k}; \\ y &= m\mathbf{k} - n \left\lfloor \frac{m\mathbf{k}}{n} \right\rfloor. \end{aligned} \tag{7}$$

Note that if n divides $m\mathbf{k}$, then $x = n, y = 0$.

Using the above characterization of profiles in Σ^* , one can apply the following simple procedure to greedily allocate the resources to the players to get such a profile, proving Σ^* is not empty.

⁹ By "fair" we mean that the resources are evenly distributed among the players.

6.2.1. BNE-procedure

Number the players (resp., resources) by $1, \dots, n$ (resp., $1, \dots, m$), and let x and y be given by (7). Assign the first $\lfloor \frac{mk}{n} \rfloor$ resources to the first player. Then continue with the second player and assign him the next $\lfloor \frac{mk}{n} \rfloor$ resources with the lowest congestion, where the assignment is according to the resources' numbers from the lowest to the highest. Continue until you have assigned the allocated resources to the first x players, $\lfloor \frac{mk}{n} \rfloor$ resources each. Proceed assigning the resources to the remaining y players by assigning each player $\lfloor \frac{mk}{n} \rfloor + 1$ resources.

Thus, the set Σ^* is not empty and any solution obtained by the BNE Procedure belongs to Σ^* . The following proposition, Proposition 18, implies that every profile in Σ^* is a best Nash equilibrium; that is, this simple procedure finds a best equilibrium strategy profile in a given symmetric CGF.

Proposition 18. *Let BNE be the subset of all best pure strategy Nash equilibria of a given symmetric CGF, then $\Sigma^* \subseteq BNE$.*

Proof. First, we show that every $\sigma \in \Sigma^*$ is a Nash equilibrium strategy profile. For all $e \in M$, $h_e^\sigma = \mathbf{k}$, which implies $l(h_e^\sigma) = l(\mathbf{k}) < W$ (by the definition of \mathbf{k}). That is, for all $i \in N$ and $e \in \sigma_i$ we have $l(h_e^\sigma) < W$. If $\mathbf{k} = n$ then $e \in \sigma_i$ for all $i \in N$. Otherwise, by the definition of \mathbf{k} , $l(h_e^\sigma + 1) = l(\mathbf{k} + 1) \geq W$. That is, for all $i \in N$ and $e \notin \sigma_i$ we have $l(h_e^\sigma + 1) \geq W$. Hence, by Proposition 7, σ is a Nash equilibrium strategy profile. That is, $\Sigma^* \subseteq NE$.

It remains to show that $\Sigma^* \subseteq BNE$. We begin by showing that all profiles in Σ^* have the same social disutility. Let $\sigma \in \Sigma^*$ and recall that x players in σ choose $\lfloor \frac{mk}{n} \rfloor$ service providers and y players choose $\lfloor \frac{mk}{n} \rfloor + 1$ service providers, where x and y are given by (7). Thus, the social disutility of σ equals

$$x \cdot (Wf^{\lfloor \frac{mk}{n} \rfloor} + l(\mathbf{k})(1 - f^{\lfloor \frac{mk}{n} \rfloor})) + y \cdot (Wf^{\lfloor \frac{mk}{n} \rfloor + 1} + l(\mathbf{k})(1 - f^{\lfloor \frac{mk}{n} \rfloor + 1})).$$

This implies that $\Pi(\sigma_1) = \Pi(\sigma_2)$ for all $\sigma_1, \sigma_2 \in \Sigma^*$. Therefore, to prove that $\Sigma^* \subseteq BNE$, it suffices to show that $\Sigma^* \cap BNE \neq \emptyset$.

Let $BNE_1 \subseteq BNE$ be the (nonempty) subset of all best Nash equilibria with minimum total congestion on the resources. That is,

$$BNE_1 = \arg \min_{\sigma \in BNE} \sum_{e \in M} h_e^\sigma,$$

and let $BNE_2 \subseteq BNE_1$ be the (nonempty) subset of all best equilibria with minimum total congestion and minimum sum of the differences between the cardinalities of the players' strategies:

$$BNE_2 = \arg \min_{\sigma \in BNE_1} \sum_{\{i,j\}:i,j \in N} ||\sigma_i| - |\sigma_j||.$$

We show below that $BNE_2 \subseteq \Sigma^*$, implying $\Sigma^* \cap BNE \neq \emptyset$.

Let $\sigma \in BNE_2$. If $\sigma \in \Sigma^*$ we are done. Otherwise, there exists a resource $e \in M$ with $h_e^\sigma \neq \mathbf{k}$ (implying by Proposition 14 that $h_e^\sigma > \mathbf{k}$), or there exist $i, j \in N$ such that $|\sigma_i| > |\sigma_j| + 1$. Assume the former case and let $a \in \arg \max_{e \in M} h_e^\sigma$. Note that $l(h_a^\sigma) = W$ (this is since $h_a^\sigma > \mathbf{k}$ implies $l(h_a^\sigma) \geq W$ and $\exists i \in N$ with $a \in \sigma_i$, yielding $l(h_a^\sigma) \leq W$ by Proposition 7 and $\sigma \in NE$). Let $i \in N$ with $a \in \sigma_i$, and consider $\sigma' = (\sigma_{-i}, \sigma_i \setminus \{a\})$. Now, $l(h_a^\sigma) = W$ implies $l(h_a^{\sigma'}) \leq W$ and $l(h_a^{\sigma'} + 1) \geq W$. For any other resource $e \in M \setminus \{a\}$, $h_e^{\sigma'} = h_e^\sigma$ implies $l(h_e^{\sigma'}) \leq W$ and $l(h_e^{\sigma'} + 1) \geq W$ (or $h_e^{\sigma'} = n$). Then, by Proposition 7, σ' is a Nash equilibrium strategy profile. By Lemma 16, $\Pi(\sigma') \leq \Pi(\sigma)$, implying $\sigma' \in BNE$. Thus, since $\sum_{e \in M} h_e^{\sigma'} < \sum_{e \in M} h_e^\sigma$, we get a contradiction to $\sigma \in BNE_1$. Therefore, σ satisfies $h_e^\sigma = \mathbf{k}$ for all $e \in M$. Now, since $\sigma \notin \Sigma^*$, there exist $i, j \in N$ such that $|\sigma_i| > |\sigma_j| + 1$. Let $\sigma'' = (\sigma_{-\{i,j\}}, \sigma_i \setminus \{b\}, \sigma_j \cup \{b\})$, where $b \in \sigma_i \setminus \sigma_j$. We note that for all $e \in M$, $h_e^{\sigma''} = h_e^\sigma = \mathbf{k}$ yields $l(h_e^{\sigma''}) < W$ and $l(h_e^{\sigma''} + 1) \geq W$ (or $h_e^{\sigma''} = n$). Then, by Proposition 7, σ'' is a Nash equilibrium strategy profile. By Lemma 17, $\Pi(\sigma'') \leq \Pi(\sigma)$, implying $\sigma'' \in BNE$. In addition, since $h_e^{\sigma''} = \mathbf{k}$ for all $e \in M$, then, by Proposition 14, $\sigma'' \in \arg \min_{\sigma \in BNE} \sum_{e \in M} h_e^\sigma = BNE_1$. Thus, $\sum_{\{i,j\}:i,j \in N} ||\sigma_i''| - |\sigma_j''|| < \sum_{\{i,j\}:i,j \in N} ||\sigma_i| - |\sigma_j||$ contradicts $\sigma \in BNE_2$. This completes the proof. \square

By Proposition 18, the BNE Procedure constructs a best equilibrium strategy profile in which the players are distributed “evenly” over the service providers. Hence, it is possible to suggest to the players rational behaviors, which are fair, and also benefit the society.

Next, in the following proposition, Proposition 19, we identify some worst equilibria in symmetric CGFs. These equilibrium points have very simple form and can be easily constructed, also in a greedily fashion, but in a different way to the BNE Procedure. As opposed to the fairness property of a best Nash equilibrium obtained by the BNE Procedure, the suggested worst Nash equilibrium suffers from the largest possible imparity among the players. The proof of the following proposition follows similar lines to the ones used in the proof of Proposition 18 above and thus appears in the Appendix. Recall that $\mathbf{k}^* = \max\{k : l(k) \leq W, k = 1, \dots, n\}$. Note that if $l(k) \neq W$ for all k , then $\mathbf{k}^* = \mathbf{k}$.

Proposition 19. *Given a symmetric CGF, let $\Sigma^{**} \subseteq \Sigma$ be the subset of all pure strategy profiles in which exactly \mathbf{k}^* players play M , $n - \mathbf{k}^*$ players play \emptyset and $h_e^\sigma = \mathbf{k}^*$ for all $e \in M$. Then, $\Sigma^{**} \subseteq WNE$, where WNE denotes the subset of all worst Nash equilibria.*

Next we compare the best and the worst Nash equilibria. Let us denote the social disutility of a best Nash equilibrium strategy profile by Π_B , and the worst one by Π_W . Then,

$$\begin{aligned} \Pi_W &= \mathbf{k}^*(Wf^m + l(\mathbf{k}^*)(1 - f^m)) + (n - \mathbf{k}^*)W \\ &= \mathbf{k}^*(1 - f^m)(l(\mathbf{k}^*) - W) + nW; \\ \Pi_B &= x \left(Wf^{\lfloor \frac{mk}{n} \rfloor} + l(\mathbf{k}) \left(1 - f^{\lfloor \frac{mk}{n} \rfloor} \right) \right) + y \left(Wf^{\lfloor \frac{mk}{n} \rfloor + 1} + l(\mathbf{k}) \left(1 - f^{\lfloor \frac{mk}{n} \rfloor + 1} \right) \right) \\ &= f^{\lfloor \frac{mk}{n} \rfloor} (x + fy)(W - l(\mathbf{k})) + nl(\mathbf{k}), \end{aligned}$$

where x and y are given by (7).

Therefore, the ratio between the social disutilities in a worst and a best equilibrium is

$$\frac{\Pi_W}{\Pi_B} = \frac{\mathbf{k}^*(1 - f^m)(l(\mathbf{k}^*) - W) + nW}{f^{\lfloor \frac{mk}{n} \rfloor} (x + fy)(W - l(\mathbf{k})) + nl(\mathbf{k})}.$$

Since $l(\mathbf{k}) < W$ and $l(\mathbf{k}^*) \leq W$, we have that

$$\frac{\Pi_W}{\Pi_B} < \frac{nW}{nl(\mathbf{k})} = \frac{W}{l(\mathbf{k})}.$$

This implies that the values of the social disutilities in different Nash equilibrium points lie in a very narrow range. In the context of social performance of Nash equilibria, one has to ask how far these values are from the social optimum.

6.3. Nash equilibria and social optimum

In this section we discuss the social performance of Nash equilibria in congestion games with failures. Previous results of Awerbuch et al. [4] and Christodoulou and Koutsoupias [7] showed that the price of anarchy (the worst possible ratio between social disutilities at an equilibrium and an optimum outcome) of pure equilibria in congestion games with nonnegative linear cost functions is $\frac{5}{2}$. We show that in CGFs with such cost functions even the best possible ratio between pure equilibrium and optimum social disutilities (“the price of stability”) depends on the parameters of the game and cannot be bounded by a constant value. As a result, the price of anarchy in CGFs does not have a constant upper bound, implying that it is unbounded in general. Consider the following example.

Example 20. Suppose we have the set N of $n \geq 2$ players sharing the set M of $m \geq 2$ service providers. Each service provider $e \in M$ has the failure probability f , and the service cost of each SP for each player is $l(k) = k$, where a is a fixed completion cost and $k \in \{1, \dots, n\}$. The incompleteness cost of each player is $W = n$.

Recall that there is a best Nash equilibrium, σ , in which x players choose $\lfloor \frac{mk}{n} \rfloor$ service providers and y players choose $\lfloor \frac{mk}{n} \rfloor + 1$ service providers, where $\mathbf{k} = \max\{k : l(k) < W, k = 1, \dots, n\}$ and x and y are given by (7). In our example, $\mathbf{k} = n - 1$, which implies $\lfloor \frac{mk}{n} \rfloor = m - \lceil \frac{m}{n} \rceil$. The disutility of player $i \in N$ at this point is given by

$$\begin{aligned} \Pi_i^x(\sigma) &= f^{m - \lceil \frac{m}{n} \rceil} W + \left(1 - f^{m - \lceil \frac{m}{n} \rceil} \right) l(n - 1) \\ &= f^{m - \lceil \frac{m}{n} \rceil} n + \left(1 - f^{m - \lceil \frac{m}{n} \rceil} \right) (n - 1) \\ &= n - \left(1 - f^{m - \lceil \frac{m}{n} \rceil} \right), \\ \text{or } \Pi_i^y(\sigma) &= f^{m - \lceil \frac{m}{n} \rceil + 1} W + \left(1 - f^{m - \lceil \frac{m}{n} \rceil + 1} \right) l(n - 1) \\ &= f^{m - \lceil \frac{m}{n} \rceil + 1} n + \left(1 - f^{m - \lceil \frac{m}{n} \rceil + 1} \right) (n - 1) \\ &= n - \left(1 - f^{m - \lceil \frac{m}{n} \rceil + 1} \right), \end{aligned}$$

with $\Pi_i^x(\sigma) \geq \Pi_i^y(\sigma)$. Then, the social disutility of σ ,

$$\Pi(\sigma) = \sum_{i \in N} \Pi_i(\sigma) \geq n \left(n - \left(1 - f^{m - \lceil \frac{m}{n} \rceil} \right) \right).$$

Consider the combination of strategies $\hat{\sigma}$ that corresponds to the following players’ behavior: each player chooses only one SP and the players divide up the SPs in a uniform way, i.e. each SP is chosen by $\frac{n}{m}$ players (assume m divides n). The disutility of player $i \in N$ at this point is

$$\Pi_i(\hat{\sigma}) = fW + (1 - f)l\left(\frac{n}{m}\right) = fn + (1 - f)\frac{n}{m},$$

and the social disutility is

$$\Pi(\hat{\sigma}) = \sum_{i \in n} \Pi_i(\hat{\sigma}) = n \left(fn + (1-f) \frac{n}{m} \right).$$

Then, the ratio between the outcomes of a best Nash equilibrium and a social optimum is

$$\begin{aligned} \frac{\Pi(\sigma)}{\Pi(OPT)} &\geq \frac{\Pi(\sigma)}{\Pi(\hat{\sigma})} \geq \frac{n \left(n - (1-f)^{m-\lceil \frac{m}{n} \rceil} \right)}{n \left(fn + (1-f) \frac{n}{m} \right)} = \frac{n - \left(1 - f^{m-\lceil \frac{m}{n} \rceil} \right)}{fn + (1-f) \frac{n}{m}} \\ &= \frac{m \left(n - \left(1 - f^{m-\lceil \frac{m}{n} \rceil} \right) \right)}{fmn + (1-f)n} \xrightarrow{f \rightarrow 0} \frac{m(n-1)}{n} \xrightarrow{n \rightarrow \infty} m. \end{aligned}$$

This implies that the ratio between the disutilities in a best Nash equilibrium and the social optimum (and therefore, the price of anarchy—the ratio between the social disutilities in a worst equilibrium and the social optimum) in CGFs, unlike in congestion games, is not bounded above by a constant, but is game-dependent. Therefore, the price of stability and the price of anarchy are unbounded in general.

It is of interest to know, whether it is possible to find an upper bound in terms of the game parameters. To address this problem one has to deal with different types of service cost functions (linear, polynomial etc.), and they may lead to different bounds on the price of stability/anarchy. In addition, there is a need to find a method for evaluating an optimal social disutility, which appears to be a non-trivial problem even in the symmetric games. We see evaluating the price of stability/anarchy in CGFs as one of the most interesting and challenging directions to continue our study.

7. Summary and future work

The study of congestion in systems is central to many disciplines. This congestion may be a result of the actions taken by self-motivated participants, and therefore congestion settings deserve extensive study. Surprisingly, although the notion of machine failures is widely discussed in the OR and CS literature, the relationships between congestion settings with self-motivated participants and machine failures have hardly been studied. In order to address this need we introduced in this paper the notion of Congestion Games with Failures [CGFs]. As it turns out, this new setting leads to interesting observations about the interplay between the need to deal with failures and the emergence of congestion in non-cooperative systems. Indeed, the classical idea of using several resources in order to overcome the possibility of failure, may result in a highly congested system, hurting all players in the system.

Our results show that although CGFs do not possess a potential function, and not even a generalized ordinal potential function, they always have a Nash equilibrium in pure strategies. Moreover, although an arbitrary improvement dynamics may cycle, any sequence of the players' best responses converges to an equilibrium profile. We also propose a procedure that guarantees convergence in polynomial time.

We further explore the properties of pure strategy equilibria in CGFs. In particular, for symmetric games we characterize best and worst equilibria and show that best equilibria possess fairness properties. The disutilities of best and worst equilibria are shown to be quite close to one another, while the relation between their disutility and the optimal social disutility might be unbounded, in difference to the results known for standard congestion games.

Although the price of stability and the price of anarchy in CGFs are unbounded in general, it is a challenging question whether it is possible to find an upper bound in terms of the instance parameters. In addition, the inefficiency of Nash equilibria motivates the study of methods for improving the social outcome obtained by selfish players. In this context, one may consider the use of taxation in order to improve the social utility. Another interesting direction is to characterize the semi-strong Nash-equilibria.

The model of CGFs can be extended in various ways. For instance, a natural generalization is to introduce fixed costs (or, taxes) that the players would be required to pay for all service providers they use. As it turns out, incorporating taxes significantly complicates the model, and basic results such as Proposition 7, hold no more. However, for the special case with symmetric taxes and failure probabilities we were able to prove the existence of a pure strategy Nash equilibrium and develop a polynomial-time algorithm for its construction. These results are presented in a separate paper on Taxed Congestion Games with Failures [TCGFs] [22]. We intend to continue our study of the TCGF-model: in particular, it is a challenge to generalize the model to allow for resource-dependent taxes and failure probabilities, and explore the relations between the taxation scheme and the social welfare.

Acknowledgements

We thank the anonymous reviewers for their useful suggestions. Part of this research was supported by the German Israeli Foundation (GIF) under contract 877/05. The second authors work was done when the author was a Ph.D. student at the Technion, Haifa 32000, Israel.

Appendix

A.1. Proof of Lemma 17

Since $|\sigma_i| > |\sigma_j| + 1$, then $\sigma_i \setminus \sigma_j$ is not empty, and let $b \in \sigma_i \setminus \sigma_j$ be such a resource with maximal service cost: $b \in \arg \max_{e \in \sigma_i \setminus \sigma_j} l(h_e^\sigma)$. We want to show that if we take this resource from player i and give it to player j , the social disutility of the players can only improve. Obviously, for all $e \in M$, $h_e^{\hat{\sigma}} = h_e^\sigma$ (we will denote it by h_e). Then, since $\hat{\sigma}_k = \sigma_k$ for all $k \neq i, j$, we have $\Pi_k(\hat{\sigma}) = \Pi_k(\sigma)$, for all $k \neq i, j$. Therefore, it suffices to show that $\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma}) \leq \Pi_i(\sigma) + \Pi_j(\sigma)$. The sample path technique that we used in the proofs of Proposition 7 and Lemma 16 is not valid here, and we operate directly with expected disutilities.

$$\begin{aligned} [\Pi_i(\sigma) + \Pi_j(\sigma)] - [\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma})] &= Wf^{|\sigma_i|} + \sum_{A' \in P(\sigma_i) \setminus \{\emptyset\}} \min_{e \in A'} l(h_e) s^{|A'|} f^{|\sigma_i \setminus A'|} \\ &+ Wf^{|\sigma_j|} + \sum_{B \in P(\sigma_j) \setminus \{\emptyset\}} \min_{e \in B} l(h_e) s^{|B|} f^{|\sigma_j \setminus B|} \\ &- Wf^{|\sigma_i| - 1} - \sum_{A \in P(\sigma_i \setminus \{b\}) \setminus \{\emptyset\}} \min_{e \in A} l(h_e) s^{|A|} f^{|\sigma_i \setminus A| - 1} \\ &- Wf^{|\sigma_j| + 1} - \sum_{B' \in P(\sigma_j \cup \{b\}) \setminus \{\emptyset\}} \min_{e \in B'} l(h_e) s^{|B'|} f^{|\sigma_j \setminus B'| + 1}. \end{aligned}$$

For any set X , we henceforth denote by $\bar{P}(X)$ the set of all *nonempty* subsets of X : $\bar{P}(X) = P(X) \setminus \{\emptyset\}$, and observe that the following equality holds for every pair of sets X, Y :

$$\bar{P}(X) = \bar{P}(X \cap Y) \cup \bar{P}(X \setminus Y) \cup \{\Omega \cup \Psi \mid \Omega \in \bar{P}(X \cap Y), \Psi \in \bar{P}(X \setminus Y)\}. \tag{8}$$

By (8), and since $b \in \sigma_i \setminus \sigma_j$,

$$\begin{aligned} \bar{P}(\sigma_i) &= \bar{P}(\sigma_i \setminus \{b\}) \cup \{b\} \cup \{A \cup \{b\} \mid A \in \bar{P}(\sigma_i \setminus \{b\})\}; \\ \bar{P}(\sigma_j \cup \{b\}) &= \bar{P}(\sigma_j) \cup \{b\} \cup \{B \cup \{b\} \mid B \in \bar{P}(\sigma_j)\}. \end{aligned} \tag{9}$$

Then, by (9),

$$\begin{aligned} [\Pi_i(\sigma) + \Pi_j(\sigma)] - [\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma})] &= Wf^{|\sigma_i| - 1} (f - 1) + Wf^{|\sigma_j|} (1 - f) + l(h_b) s^f f^{|\sigma_i| - 1} - l(h_b) s^f f^{|\sigma_j|} \\ &+ \sum_{A \in \bar{P}(\sigma_i \setminus \{b\})} \min_{e \in A} l(h_e) s^{|A|} f^{|\sigma_i \setminus A|} - \sum_{A \in \bar{P}(\sigma_i \setminus \{b\})} \min_{e \in A} l(h_e) s^{|A|} f^{|\sigma_i \setminus A| - 1} \\ &+ \sum_{A \in \bar{P}(\sigma_i \setminus \{b\})} \min_{e \in A \cup \{b\}} l(h_e) s^{|A| + 1} f^{|\sigma_i \setminus A| - 1} - \sum_{B \in \bar{P}(\sigma_j)} \min_{e \in B} l(h_e) s^{|B|} f^{|\sigma_j \setminus B| + 1} \\ &+ \sum_{B \in \bar{P}(\sigma_j)} \min_{e \in B} l(h_e) s^{|B|} f^{|\sigma_j \setminus B|} - \sum_{B \in \bar{P}(\sigma_j)} \min_{e \in B \cup \{b\}} l(h_e) s^{|B| + 1} f^{|\sigma_j \setminus B|}. \end{aligned}$$

Simplifying this expression, we get

$$\begin{aligned} [\Pi_i(\sigma) + \Pi_j(\sigma)] - [\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma})] &= s \left[(f^{|\sigma_j|} - f^{|\sigma_i| - 1})(W - l(h_b)) + \sum_{A \in \bar{P}(\sigma_i \setminus \{b\})} s^{|A|} f^{|\sigma_i \setminus A| - 1} \right. \\ &\left. \times \left(\min_{e \in A \cup \{b\}} l(h_e) - \min_{e \in A} l(h_e) \right) + \sum_{B \in \bar{P}(\sigma_j)} s^{|B|} f^{|\sigma_j \setminus B|} \left(\min_{e \in B} l(h_e) - \min_{e \in B \cup \{b\}} l(h_e) \right) \right]. \end{aligned}$$

By (8), and since $b \in \sigma_i \setminus \sigma_j$,

$$\begin{aligned} \bar{P}(\sigma_i \setminus \{b\}) &= \bar{P}(\sigma_i \cap \sigma_j) \cup \bar{P}((\sigma_i \setminus \{b\}) \setminus \sigma_j) \\ &\quad \cup \{\Omega \cup A' \mid \Omega \in \bar{P}(\sigma_i \cap \sigma_j), A' \in \bar{P}((\sigma_i \setminus \{b\}) \setminus \sigma_j)\}; \\ \bar{P}(\sigma_j) &= \bar{P}(\sigma_i \cap \sigma_j) \cup \bar{P}(\sigma_j \setminus (\sigma_i \setminus \{b\})) \\ &\quad \cup \{\Omega \cup B' \mid \Omega \in \bar{P}(\sigma_i \cap \sigma_j), B' \in \bar{P}(\sigma_j \setminus (\sigma_i \setminus \{b\}))\}. \end{aligned} \tag{10}$$

Then, by (10),

$$\begin{aligned}
 [\Pi_i(\sigma) + \Pi_j(\sigma)] - [\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma})] &= s \left[(f^{|\sigma_j|} - f^{|\sigma_i|-1})(W - l(h_b)) + \sum_{\Omega \in \bar{P}(\sigma_i \cap \sigma_j)} (\min_{e \in \Omega} l(h_e) - \min_{e \in \Omega \cup \{b\}} l(h_e)) s^{|\Omega|} \right. \\
 &\quad \times (f^{|\sigma_j \setminus \Omega|} - f^{|\sigma_i \setminus \Omega|-1}) + \sum_{B' \in \bar{P}(\sigma_j \setminus (\sigma_i \setminus \{b\}))} (\min_{e \in B'} l(h_e) - \min_{e \in B' \cup \{b\}} l(h_e)) \\
 &\quad \times s^{|B'|} f^{|\sigma_j \setminus B'|} + \sum_{B' \in \bar{P}(\sigma_j \setminus (\sigma_i \setminus \{b\}))} \sum_{\Omega \in \bar{P}(\sigma_i \cap \sigma_j)} (\min_{e \in B' \cup \Omega} l(h_e) - \min_{e \in B' \cup \Omega \cup \{b\}} l(h_e)) \\
 &\quad \times s^{|B' \cup \Omega|} f^{|\sigma_j \setminus (B' \cup \Omega)|} + \sum_{A' \in \bar{P}((\sigma_i \setminus \{b\}) \setminus \sigma_j)} \left(\min_{e \in A' \cup \{b\}} l(h_e) - \min_{e \in A'} l(h_e) \right) \\
 &\quad \times s^{|A'|} f^{|\sigma_i \setminus A'|-1} + \sum_{A' \in \bar{P}((\sigma_i \setminus \{b\}) \setminus \sigma_j)} \sum_{\Omega \in \bar{P}(\sigma_i \cap \sigma_j)} \\
 &\quad \left. \times \left(\min_{e \in A' \cup \Omega \cup \{b\}} l(h_e) - \min_{e \in A' \cup \Omega} l(h_e) \right) s^{|A' \cup \Omega|} f^{|\sigma_i \setminus (A' \cup \Omega)|-1} \right].
 \end{aligned}$$

Let $X \subseteq M$ represent any (even empty) set of resources. Then,

- $f(X), s(X) \geq 0$ (as probabilities);
- $f(\sigma_j \setminus X) - f((\sigma_i \setminus \{b\}) \setminus X) > 0$ (follows from $|\sigma_i| > |\sigma_j| + 1$);
- $W \geq l(h_b)$ (given);
- $l_X(h_X) \geq l_{X \cup \{b\}}(h_{X \cup \{b\}})$ (since $X \subseteq X \cup \{b\}$);
- from the choice of b , for all $e \in \sigma_i \setminus \sigma_j, l(h_e) \leq l(h_b)$
 \Rightarrow for all $Y \in \bar{P}(\sigma_i \setminus \sigma_j)$ (and, in particular, for $Y \in \bar{P}((\sigma_i \setminus \{b\}) \setminus \sigma_j)$), $l_{Y \cup X \cup \{b\}}(h_{Y \cup X \cup \{b\}}) = l_{Y \cup X}(h_{Y \cup X})$.

Therefore, we finally get $[\Pi_i(\sigma) + \Pi_j(\sigma)] - [\Pi_i(\hat{\sigma}) + \Pi_j(\hat{\sigma})] \geq 0$, as required. \square

A.2. Proof of Proposition 19

Clearly Σ^{**} is not empty, and all strategy profiles in Σ^{**} have the same social disutility $\mathbf{k}^*(Wf^m + l(\mathbf{k}^*)(1 - f^m))$. We turn now to show that every $\sigma \in \Sigma^{**}$ is a Nash equilibrium strategy profile. For all $e \in M, h_e^\sigma = \mathbf{k}^*$, implying $l(h_e^\sigma) = l(\mathbf{k}^*) \leq W$ (by the definition of \mathbf{k}^*). Hence, for all $i \in N$ and $e \in \sigma_i$ we have $l(h_e^\sigma) \leq W$. If $\mathbf{k}^* = n$ then $e \in \sigma_i$ for all $i \in N$. Otherwise, by the definition of \mathbf{k}^* and the symmetry, $l(h_e^\sigma + 1) = l(\mathbf{k}^* + 1) > W$ for each $e \in M$. That is, for all $i \in N$ and $e \notin \sigma_i$ we have $l(h_e^\sigma + 1) > W$. Then, by Proposition 7, σ is a Nash equilibrium strategy profile. That is, $\Sigma^{**} \subseteq NE$.

It remains to prove that $\Sigma^{**} \subseteq WNE$. Since $\Pi(\sigma^1) = \Pi(\sigma^2)$ for all $\sigma^1, \sigma^2 \in \Sigma^{**}$, it suffices to show that $\Sigma^{**} \cap WNE \neq \emptyset$. Let $WNE_1 \subseteq WNE$ be the (nonempty) subset of all worst Nash equilibria with maximum total congestion on the resources. That is,

$$WNE_1 = \arg \max_{\sigma \in WNE} \sum_{e \in M} h_e^\sigma,$$

and let $WNE_2 \subseteq WNE_1$ be the (nonempty) subset of all worst equilibria with maximum total congestion and maximum sum of the differences between the cardinalities of the players' strategies:

$$WNE_2 = \arg \max_{\sigma \in WNE_1} \sum_{\{i,j\}: i,j \in N} ||\sigma_i| - |\sigma_j||.$$

We show below that $WNE_2 \subseteq \Sigma^{**}$, implying $\Sigma^{**} \cap WNE \neq \emptyset$. Let $\sigma \in WNE_2$. If $\sigma \in \Sigma^{**}$ we are done. Otherwise, if there exists a resource $a \in M$ such that $h_a^\sigma \neq \mathbf{k}^*$ then by Proposition 14 $h_a^\sigma < \mathbf{k}^* \leq n$. Let $i \in N$ be a player with $a \notin \sigma_i$, and consider $\sigma' = (\sigma_{-i}, \sigma_i \cup \{a\})$. Since σ is a Nash equilibrium strategy profile, by Proposition 7, $l(h_a^\sigma + 1) \geq W$, implying $l(h_a^{\sigma'} + 1) \geq W$. In addition, $h_a^\sigma < \mathbf{k}^*$ yields $l(h_a^{\sigma'}) \leq l(\mathbf{k}^*) \leq W$. For any other resource $e \in M \setminus \{a\}, h_e^{\sigma'} = h_e^\sigma$ implies $l(h_e^{\sigma'}) \leq W$ and $l(h_e^{\sigma'} + 1) \geq W$ (or $h_e^{\sigma'} = n$). Thus, by Proposition 7, σ' is a Nash equilibrium strategy profile. We note that $l(h_a^\sigma + 1) = l(h_a^{\sigma'}) \geq W$, coupled with the definition of \mathbf{k}^* , imply that $l(h_a^{\sigma'}) \geq l(\mathbf{k}^*) \geq l(h_e^\sigma) = l(h_e^{\sigma'})$ for all $e \in M \setminus \{a\}$. Then, by Lemma 16, $\Pi(\sigma') \geq \Pi(\sigma)$, implying $\sigma' \in WNE$. Thus, since $\sum_{e \in M} h_e^{\sigma'} > \sum_{e \in M} h_e^\sigma$, we get a contradiction to $\sigma \in WNE_1$. Therefore, σ satisfies $h_e^\sigma = \mathbf{k}^*$ for all $e \in M$. Then, since $\sigma \notin \Sigma^{**}$, there exist $i, j \in N$ such that $\sigma_i, \sigma_j \neq \emptyset, M$ (w.l.o.g., assume that $|\sigma_i| \leq |\sigma_j|$). Let $\sigma'' = (\sigma_{-\{i,j\}}, \sigma_i \setminus \{b\}, \sigma_j \cup \{b\})$, where $b \in \sigma_i \setminus \sigma_j$. We note that for all $e \in M, h_e^{\sigma''} = h_e^\sigma$ yields $l(h_e^{\sigma''}) \leq W$ and $l(h_e^{\sigma''} + 1) \geq W$ (or $h_e^{\sigma''} = n$), implying that σ'' is a Nash equilibrium strategy profile. By Lemma 17, $\Pi(\sigma'') \geq \Pi(\sigma)$, implying $\sigma'' \in WNE$. In addition, since $h_e^{\sigma''} = \mathbf{k}^*$ for all $e \in M$, then, by Proposition 14, $\sigma'' \in \arg \max_{\sigma \in WNE} \sum_{e \in M} h_e^\sigma = WNE_1$. Thus, $\sum_{\{i,j\}: i,j \in N} ||\sigma_i''| - |\sigma_j''|| > \sum_{\{i,j\}: i,j \in N} ||\sigma_i| - |\sigma_j||$ contradicts $\sigma \in WNE_2$. This completes the proof. \square

References

- [1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden, The price of stability for network design with fair cost allocation, in: *The Proceedings of the 45th Symposium on Foundations of Computer Science, FOCS-04, 2004*, pp. 295–304.
- [2] R. Aumann, Acceptable points in general cooperative n -person games, in: *Contributions to the Theory of Games*, vol. IV, Princeton University Press, Princeton, NJ, 1959, pp. 287–324.
- [3] R. Aumann, Nash equilibria are not self-enforcing, in: J. Gabrszewicz, J. Richard, L. Wolsey (Eds.), *Economics Decision-Making: Games, Econometrics and Optimization: Essays in Honor of Jacques Dreze*, Elsevier Science, 1990, pp. 201–206.
- [4] B. Awerbuch, Y. Azar, A. Epstein, Large the price of routing unsplitable flow, in: *Proceedings of the 37th Annual ACM Symposium on Theory and Computing, STOC-05, 2005*, pp. 57–66.
- [5] D. Bernheim, B. Peleg, M. Whinston, Coalition-proof Nash equilibria. I. Concepts, *International Journal of Game Theory* 42 (1987) 1–12.
- [6] D. Bernheim, M. Whinston, Coalition-proof Nash equilibria. II. Applications, *International Journal of Game Theory* 42 (1987) 13–29.
- [7] G. Christodoulou, E. Koutsoupias, The price of anarchy of finite congestion games, in: *STOC-05, 2005*, pp. 67–73.
- [8] J. Correa, A. Schulz, N. Stier-Moses, Fast, fair, and efficient flows in networks, *Operations Research* 55 (2) (2007) 215–225.
- [9] O. Jahn, R. Möhring, A. Schulz, N. Stier-Moses, System-optimal routing of traffic flows with user constraints in networks with congestion, *Operations Research* 53 (2005) 600–616.
- [10] G. Kaplan, Sophisticated outcomes and coalitional stability, Master's Thesis, Tel-Aviv University, 1992.
- [11] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, in: *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, 1999*, pp. 404–413.
- [12] A. Kumar, J. Kleinberg, Fairness measures for resource allocation, *SIAM Journal on Computing* 36 (3) (2006) 657–680.
- [13] N. Linial, Game-theoretic aspects of computing, in: R. Aumann, S. Hart (Eds.), *Handbook of Game Theory with Economic Applications*, vol. 2, Elsevier Science B.V., 1994, pp. 1339–1395 (Chapter 38).
- [14] I. Milchtaich, Congestion games with player-specific payoff functions, *Games and Economic Behavior* 13 (1996) 111–124.
- [15] I. Milchtaich, Congestion models of competition, *American Naturalist* 147 (5) (1996) 760–783.
- [16] P. Milgrom, J. Roberts, Coalition-proofness and correlation with arbitrary communication possibilities, *Games and Economic Behavior* 17 (1996) 113–128.
- [17] D. Monderer, L. Shapley, Potential games, *Games and Economic Behavior* 14 (1996) 124–143.
- [18] A. Orda, R. Rom, N. Shimkin, Competitive routing in multi-user communication networks, *IEEE/ACM Transactions on Networking* 1 (1993) 510–521.
- [19] C. Papadimitriou, Algorithms, games, and the Internet, in: *Proceedings of the 33th Annual ACM Symposium on the Theory of Computing, STOC-01, 2001*, pp. 749–753.
- [20] M. Penn, M. Polukarov, M. Tennenholtz, Congestion games with failures, in: *Proceedings of the 6th ACM Conference on Electronic Commerce, EC-05, 2005*, pp. 259–268.
- [21] M. Penn, M. Polukarov, M. Tennenholtz, Congestion games with load-dependent failures: identical resources, *Games and Economic Behavior* 67 (1) (2009) 156–173.
- [22] M. Penn, M. Polukarov, M. Tennenholtz, Taxed congestion games with failures, *Annals of Mathematics and Artificial Intelligence* 56 (2) (2009) 133–151.
- [23] R. Porter, A. Ronen, Y. Shoham, M. Tennenholtz, Mechanism design with execution uncertainty, in: *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence, UAI-02, Morgan Kaufmann, San Francisco, CA, 2002*, pp. 414–442.
- [24] T. Quint, M. Shubik, A model of migration, No. 1088 in *Cowles Foundation Discussion Papers*, Cowles Foundation, Yale University, 1994.
- [25] R. Rosenthal, A class of games possessing pure-strategy Nash equilibria, *International Journal of Game Theory* 2 (1973) 65–67.
- [26] T. Roughgarden, *Selfish Routing and the Price of Anarchy*, The MIT Press, 2005.
- [27] T. Roughgarden, E. Tardos, How bad is selfish routing, *Journal of the ACM* 49 (2) (2002) 236–259.
- [28] T. Roughgarden, E. Tardos, Bounding the inefficiency of equilibria in nonatomic congestion games, *Games and Economic Behavior* 47 (2004) 389–403.
- [29] W. Sandholm, Potential games with continuous player sets, *Journal of Economic Theory* 97 (2001) 81–108.